


DB2 Version 9.5
for Linux, UNIX, and Windows



Data Servers, Databases, and Database Objects Guide
Updated April, 2009

DB2 Version 9.5
for Linux, UNIX, and Windows



Data Servers, Databases, and Database Objects Guide
Updated April, 2009

Note

Before using this information and the product it supports, read the general information under Appendix B, "Notices," on page 617.

Edition Notice

This document contains proprietary information of IBM. It is provided under a license agreement and is protected by copyright law. The information contained in this publication does not include any product warranties, and any statements provided in this manual should not be interpreted as such.

You can order IBM publications online or through your local IBM representative.

- To order publications online, go to the IBM Publications Center at www.ibm.com/shop/publications/order
- To find your local IBM representative, go to the IBM Directory of Worldwide Contacts at www.ibm.com/planetwide

To order DB2 publications from DB2 Marketing and Sales in the United States or Canada, call 1-800-IBM-4YOU (426-4968).

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

© Copyright International Business Machines Corporation 1993, 2009.

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

Contents

About this book ix

Part 1. Data servers 1

Chapter 1. DB2 data servers 3

Management of data server capacity 3
Enabling large page support in a 64-bit environment (AIX) 4

Chapter 2. Multiple DB2 copies 7

Default IBM database client interface copy 7
Setting the DAS when running multiple DB2 copies 10
Setting the default instance when using multiple DB2 copies (Windows) 11
Multiple instances of the database manager 12
Multiple instances (Windows) 13
Updating DB2 copies (Windows) 13
Running multiple instances concurrently (Windows) 14
Working with instances on the same or different DB2 copies 15

Chapter 3. Autonomic computing 17

Automatic features 17
Automatic maintenance 19
 Maintenance windows. 19
Self-tuning memory 20
 Memory allocation in DB2 21
 Self tuning memory operational details and limitations. 24
 Operational details, limitations, and interaction between memory parameters 25
 Enabling self tuning memory 27
 Disabling self tuning memory 27
 Determining which memory consumers are enabled for self tuning. 28
 Self tuning memory in partitioned database environments. 29
 Using self-tuning memory in partitioned database environments 31
Configuring memory and memory heaps 32
 Agent and process model configuration 35
 Agent, process model, and memory configuration 36
Automatic storage 38
 Automatic storage table spaces 38
 Automatic storage databases. 44
 Automatic storage restrictions 47
Automatic (compression) dictionary creation (ADC) 47
 Data row compression. 49
Automatic statistics collection 50
 Enabling automatic statistics collection 54
Configuration Advisor. 55
 Tuning configuration parameters using the Configuration Advisor. 55

Generating database configuration recommendations 56
 Example: Requesting configuration recommendations using the Configuration Advisor 56
Utility throttling. 58
 Asynchronous index cleanup 58
 Asynchronous index cleanup for MDC tables . . . 60

Chapter 4. Instances 63

Designing instances. 64
 Default instance 65
 Instance directory 66
 Multiple instances (Linux, UNIX) 66
 Multiple instances (Windows) 67
Creating instances 68
Modifying instances 69
 Updating the instance configuration (Linux, UNIX) 69
 Updating the instance configuration (Windows) 70
Working with instances 71
 Auto-starting instances 71
 Starting instances (Linux, UNIX) 71
 Starting instances (Windows) 72
 Attaching to and detaching from instances . . . 72
 Working with instances on the same or different DB2 copies 73
 Stopping instances (Linux, UNIX) 73
 Stopping instances (Windows) 74
Dropping instances. 75

Chapter 5. Lightweight Directory Access Protocol (LDAP) 77

Security considerations in an LDAP environment. . 77
LDAP object classes and attributes used by DB2 . . 78
Extending the LDAP directory schema with DB2 object classes and attributes 88
Supported LDAP client and server configurations 88
 LDAP support and DB2 Connect 89
 Extending the directory schema for IBM Tivoli Directory Server. 90
 Netscape LDAP directory support and attribute definitions. 91
 Extending the directory schema for Sun One Directory Server. 93
 Windows Active Directory 95
 Enabling LDAP support after DB2 installation is complete 98
Registering LDAP entries. 98
 Registration of DB2 servers after installation . . 98
 Catalog a node alias for ATTACH 100
 Registration of databases in the LDAP directory 100
Deregistering LDAP entries. 100
 Deregistering the DB2 server 100

Deregistering the database from the LDAP directory	101
Configuring LDAP users	101
Creating an LDAP user	101
Configuring the LDAP user for DB2 applications	102
Setting DB2 registry variables at the user level in the LDAP environment	102
Disabling LDAP support	102
Updating the protocol information for the DB2 server	102
Rerouting LDAP clients to another server	103
Attaching to a remote server in the LDAP environment.	104
Refreshing LDAP entries in local database and node directories	104
Searching the LDAP servers	105

Part 2. Databases. 107

Chapter 6. Databases 109

Designing databases	109
Database directories and files	110
Space requirements for database objects	118
Space requirements for log files	118
Lightweight Directory Access Protocol (LDAP) directory service	119
Creating databases	120
Automatic storage databases	121
Cataloging databases	128
Binding utilities to the database	129
Creating database aliases	130
Connecting to distributed relational databases	131
Remote unit of work for distributed relational databases.	131
Application-directed distributed unit of work	134
Application process connection states	135
Connection states	136
Options that govern unit of work semantics	137
Data representation considerations	138
Viewing the local or system database directory files	138
Dropping databases	138
Dropping aliases	139

Chapter 7. Database partitions 141

Chapter 8. Buffer pools 143

Designing buffer pools	143
Buffer pool memory protection (AIX running on POWER6)	145
Creating buffer pools	146
Modifying buffer pools	147
Dropping buffer pools	148

Chapter 9. Table spaces 151

Designing table spaces	152
Types of table spaces	154
Comparison of SMS and DMS table spaces	167

Considerations when choosing table spaces for your tables	170
Automatic re-sizing of table spaces	171
Automatic prefetchsize adjustment after adding or dropping containers	175
Table spaces without file system caching	176
Table space extent sizes	182
Table space page sizes	183
Table space disk I/O	183
Defining initial table spaces	185
Attaching DMS direct disk access devices	186
Configuring and setting up DMS direct disk access (Linux)	187
Creating table spaces	189
Altering table spaces	193
Altering SMS table spaces	193
Altering DMS table spaces	193
Altering automatic storage table spaces.	206
Renaming a table space	206
Switching table spaces from offline to online	207
Optimizing table space performance when data is on RAID devices	207
Dropping table spaces	208

Chapter 10. Schemas 211

Designing schemas	212
Grouping objects by schema	214
Schema name restrictions and recommendations	215
Creating schemas	216
Copying schemas	216
Example of schema copy using the ADMIN_COPY_SCHEMA procedure	218
Examples of schema copy using the db2move utility	218
Restarting a failed copy schema operation.	219
Dropping schemas.	222

Part 3. Database objects 223

Chapter 11. Tables 225

Types of tables	225
Designing tables	227
Table design concepts	227
Space requirements for tables	236
Space requirements for user table data	238
Space compression for tables	240
Optimistic locking.	244
Table partitioning and data organization schemes	253
Creating tables	253
Declaring global temporary tables	254
Creating tables like existing tables	255
Creating tables for staging data	255
Modifying tables	256
Altering materialized query table properties	257
Refreshing the data in a materialized query table	257
Changing column properties	257
Renaming tables	260
Recovering inoperative summary tables	260
Viewing table definitions	261

Table or view aliases	261
Dropping tables	261
Dropping materialized query or staging tables	262
Scenarios and examples of tables	262
Scenarios: Optimistic locking and time-based detection	263

Chapter 12. Constraints 267

Types of constraints	267
NOT NULL constraints	268
Unique constraints	268
Primary key constraints	269
(Table) Check constraints	269
Foreign key (referential) constraints	269
Informational constraints	274
Designing constraints.	274
Designing unique constraints	274
Designing primary key constraints	275
Designing check constraints	275
Designing foreign key (referential) constraints	277
Designing informational constraints	282
Creating and modifying constraints	284
Viewing constraint definitions for a table	285
Dropping constraints	286

Chapter 13. Indexes 289

Types of indexes	290
Designing indexes.	292
Tools for designing indexes.	294
Space requirements for indexes	294
Creating indexes	297
Modifying indexes	298
Renaming indexes.	298
Rebuilding indexes	299
Dropping indexes	300

Chapter 14. Triggers 301

Types of triggers	302
BEFORE triggers	303
AFTER triggers.	303
INSTEAD OF triggers	304
Designing triggers.	305
Specifying what makes a trigger fire (triggering statement or event)	307
Specifying when a trigger fires (BEFORE, AFTER, and INSTEAD OF clauses)	308
Defining conditions for when trigger-action will fire (WHEN clause)	310
Supported SQL PL statements in triggers	311
Accessing old and new column values in triggers using transition variables	312
Referencing old and new table result sets using transition tables	313
Creating triggers	314
Modifying and dropping triggers.	316
Examples of triggers and trigger use	316
Examples of interaction between triggers and referential constraints.	316
Examples of defining actions using triggers	318

Example of defining business rules using triggers	319
Example of preventing operations on tables using triggers	320

Chapter 15. Sequences 321

Designing sequences	321
Managing sequence behavior	322
Application performance and sequences	323
Sequences compared to identity columns	324
Creating sequences	325
Generating sequential values	326
Determining when to use identity columns or sequences	326
Modifying sequences	327
Viewing sequence definitions	328
Dropping sequences	328
Examples of how to code sequences.	329
Sequence reference	330

Chapter 16. Views 335

Designing views	336
System catalog views.	336
Views with the check option	337
Deletable views	339
Insertable views	340
Updatable views	340
Read-only views	341
Creating views	341
Creating views that use user-defined functions (UDFs)	342
Modifying typed views	343
Recovering inoperative views	343
Dropping views	344

Part 4. Reference 345

Chapter 17. Conforming to naming rules 347

Naming rules	347
DB2 object naming rules.	348
Delimited identifiers and object names	349
User, user ID and group naming rules	350
Naming rules in an NLS environment	350
Naming rules in a Unicode environment	351

Chapter 18. SQL and XML limits 353

Chapter 19. Registry and environment variables 363

Environment variables and the profile registry	363
Declaring, showing, changing, resetting, and deleting registry and environment variables	365
Setting environment variables on Windows	367
Setting environment variables on Linux and UNIX operating systems.	369
Setting the current instance environment variables	370
Aggregate registry variables	371

DB2 registry and environment variables	372
General registry variables	377
System environment variables	385
Communications variables	394
Command-line variables	397
Partitioned database environment variables	398
Query compiler variables	399
Performance variables	404
Miscellaneous variables	421

Chapter 20. Configuration parameters 439

Configuring the DB2 database manager with configuration parameters	440
Configuration parameters summary	443
Configuration parameters that affect the number of agents	455
Configuration parameters that affect query optimization.	456
Restrictions and behavior when configuring max_coordagents and max_connections	458
Database Manager configuration parameters	460
agent_stack_sz - Agent stack size	460
agentpri - Priority of agents	461
aslheapsz - Application support layer heap size	463
audit_buf_sz - Audit buffer size	464
authentication - Authentication type.	465
catalog_noauth - Cataloging allowed without authority	466
clnt_krb_plugin - Client Kerberos plug-in	467
clnt_pw_plugin - Client userid-password plug-in	467
cluster_mgr - Cluster manager name	468
comm_bandwidth - Communications bandwidth	468
conn_elapse - Connection elapse time	469
cpuspeed - CPU speed	469
dft_account_str - Default charge-back account	470
dft_monswitches - Default database system monitor switches	471
dftdbpath - Default database path	472
diaglevel - Diagnostic error capture level	473
diagpath - Diagnostic data directory path	474
dir_cache - Directory cache support	475
discover - Discovery mode	476
discover_inst - Discover server instance	477
fcm_num_buffers - Number of FCM buffers	477
fcm_num_channels - Number of FCM channels	478
fed_noauth - Bypass federated authentication	479
federated - Federated database system support	479
federated_async - Maximum asynchronous TQs per query configuration parameter	480
fenced_pool - Maximum number of fenced processes	480
group_plugin - Group plug-in	482
health_mon - Health monitoring	482
indexrec - Index re-creation time	483
instance_memory - Instance memory	485
intra_parallel - Enable intra-partition parallelism	487
java_heap_sz - Maximum Java interpreter heap size.	487
jdk_path - Software Developer's Kit for Java installation path	488

keepfenced - Keep fenced process	489
local_gssplugin - GSS API plug-in used for local instance level authorization.	490
max_connections - Maximum number of client connections	490
max_connretries - Node connection retries.	491
max_coordagents - Maximum number of coordinating agents	491
max_querydegree - Maximum query degree of parallelism	492
max_time_diff - Maximum time difference among nodes	493
maxagents - Maximum number of agents	493
maxcagents - Maximum number of concurrent agents	494
mon_heap_sz - Database system monitor heap size.	495
nodetype - Machine node type	496
notifylevel - Notify level.	496
num_initagents - Initial number of agents in pool	497
num_initfenced - Initial number of fenced processes	498
num_poolagents - Agent pool size	498
numdb - Maximum number of concurrently active databases including host and System i databases.	499
query_heap_sz - Query heap size.	500
release - Configuration file release level	501
resync_interval - Transaction resync interval	501
rqrioblk - Client I/O block size	502
sheapthres - Sort heap threshold	503
spm_log_file_sz - Sync point manager log file size.	504
spm_log_path - Sync point manager log file path	505
spm_max_resync - Sync point manager resync agent limit	506
spm_name - Sync point manager name.	506
srvcon_auth - Authentication type for incoming connections at the server	506
srvcon_gssplugin_list - List of GSS API plug-ins for incoming connections at the server	507
srvcon_pw_plugin - Userid-password plug-in for incoming connections at the server	507
srv_plugin_mode - Server plug-in mode	508
start_stop_time - Start and stop timeout	508
svcname - TCP/IP service name.	509
sysadm_group - System administration authority group name	510
sysctrl_group - System control authority group name	510
sysmaint_group - System maintenance authority group name	511
sysmon_group - System monitor authority group name	512
tm_database - Transaction manager database name	512
tp_mon_name - Transaction processor monitor name	513
trust_allclnts - Trust all clients.	514

trust_clntauth - Trusted clients authentication	515	hadr_remote_inst - HADR instance name of the remote server	546
util_impact_lim - Instance impact policy	516	hadr_remote_svc - HADR remote service name	547
Database configuration parameters	517	hadr_syncmode - HADR synchronization mode for log write in peer state	547
alt_collate - Alternate collating sequence	517	hadr_timeout - HADR timeout value	548
app_ctl_heap_sz - Application control heap size	517	indexrec - Index re-creation time	548
appgroup_mem_sz - Maximum size of application group memory set	518	jdk_64_path - 64-Bit Software Developer's Kit for Java installation path DAS	550
appl_memory - Application Memory configuration parameter	519	locklist - Maximum storage for lock list	551
applheapsz - Application heap size	520	locktimeout - Lock timeout	553
archretrydelay - Archive retry delay on error	521	log_retain_status - Log retain status indicator	554
auto_del_rec_obj - Automated deletion of recovery objects configuration parameter	521	logarchmeth1 - Primary log archive method	555
auto_maint - Automatic maintenance	522	logarchmeth2 - Secondary log archive method	556
autorestart - Auto restart enable	524	logarchopt1 - Primary log archive options	557
avg_appls - Average number of active applications	524	logarchopt2 - Secondary log archive options	557
backup_pending - Backup pending indicator	525	logbufsz - Log buffer size	558
blk_log_dsk_ful - Block on log disk full	525	logfilsiz - Size of log files	558
catalogcache_sz - Catalog cache size	526	loghead - First active log file	560
chnngpgs_thresh - Changed pages threshold	527	logindexbuild - Log index pages created	560
codepage - Code page for the database	528	logpath - Location of log files	560
codeset - Codeset for the database	528	logprimary - Number of primary log files	560
collate_info - Collating information	528	logretain - Log retain enable	562
country/region - Database territory code	529	logsecond - Number of secondary log files	562
database_consistent - Database is consistent	529	max_log - Maximum log per transaction	564
database_level - Database release level	530	maxappls - Maximum number of active applications	564
database_memory - Database shared memory size	530	maxfilop - Maximum database files open per application	565
db_mem_thresh - Database memory threshold	532	maxlocks - Maximum percent of lock list before escalation	566
dbheap - Database heap	533	min_dec_div_3 - Decimal division scale to 3	568
decflt_rounding - Decimal floating point rounding configuration parameter	534	mincommit - Number of commits to group	569
dft_degree - Default degree	536	mirrorlogpath - Mirror log path	570
dft_extent_sz - Default extent size of table spaces	536	multipage_alloc - Multipage file allocation enabled	571
dft_loadrec_ses - Default number of load recovery sessions	537	newlogpath - Change the database log path	571
dft_mttb_types - Default maintained table types for optimization	537	num_db_backups - Number of database backups	573
dft_prefetch_sz - Default prefetch size	538	num_freqvalues - Number of frequent values retained	573
dft_queryopt - Default query optimization class	539	num_iocleaners - Number of asynchronous page cleaners	574
dft_refresh_age - Default refresh age	539	num_ioservers - Number of I/O servers	576
dft_sqlmathwarn - Continue upon arithmetic exceptions	540	num_log_span - Number log span	576
discover_db - Discover database	541	num_quantiles - Number of quantiles for columns	577
dlchktime - Time interval for checking deadlock	541	numarchretry - Number of retries on error	578
dyn_query_mgmt - Dynamic SQL and XQuery query management	542	numsegs - Default number of SMS containers	578
enable_xmlchar - Enable conversion to XML configuration parameter	543	overflowlogpath - Overflow log path	579
failarchpath - Failover log archive path	543	pagesize - Database default page size	580
groupheap_ratio - Percent of memory for application group heap	544	pckcachesz - Package cache size	580
hadr_db_role - HADR database role	544	priv_mem_thresh - Private memory threshold	582
hadr_local_host - HADR local host name	544	rec_his_retentn - Recovery history retention period	583
hadr_local_svc - HADR local service name	545	restore_pending - Restore pending	583
hadr_peer_window - HADR peer window configuration parameter	545	restrict_access - Database has restricted access configuration parameter	583
hadr_remote_host - HADR remote host name	546	rollfwd_pending - Roll forward pending indicator	584
		self_tuning_mem - Self-tuning memory	584

seqdetect - Sequential detection flag	585
sheapthres_shr - Sort heap threshold for shared sorts	586
softmax - Recovery range and soft checkpoint interval	587
sortheap - Sort heap size	589
stat_heap_sz - Statistics heap size.	590
stmtheap - Statement heap size	591
territory - Database territory	591
trackmod - Track modified pages enable	592
tsm_mgmtclass - Tivoli Storage Manager management class	592
tsm_nodename - Tivoli Storage Manager node name	592
tsm_owner - Tivoli Storage Manager owner name	593
tsm_password - Tivoli Storage Manager password.	593
user_exit_status - User exit status indicator	594
userexit - User exit enable	594
util_heap_sz - Utility heap size	594
vendoropt - Vendor options	595
wlm_collect_int - Workload management collection interval configuration parameter	595
DB2 Administration Server (DAS) configuration parameters	596
authentication - Authentication type DAS	596
contact_host - Location of contact list	597
das_codepage - DAS code page	597
das_territory - DAS territory	598
dasadm_group - DAS administration authority group name	598
db2system - Name of the DB2 server system	599
discover - DAS discovery mode	599

exec_exp_task - Execute expired tasks	600
jdk_path - Software Developer's Kit for Java installation path DAS.	600
sched_enable - Scheduler mode	601
sched_userid - Scheduler user ID.	601
smtp_server - SMTP server	601
toolscat_db - Tools catalog database	602
toolscat_inst - Tools catalog database instance	602
toolscat_schema - Tools catalog database schema	602

Part 5. Appendixes 605

Appendix A. Overview of the DB2 technical information 607

DB2 technical library in hardcopy or PDF format	607
Ordering printed DB2 books	610
Displaying SQL state help from the command line processor	611
Accessing different versions of the DB2 Information Center.	611
Displaying topics in your preferred language in the DB2 Information Center	611
Updating the DB2 Information Center installed on your computer or intranet server	612
DB2 tutorials	614
DB2 troubleshooting information	614
Terms and Conditions	614

Appendix B. Notices 617

Index 621

About this book

The Data Servers, Databases, and Database Objects Guide provides information necessary to use and administer the DB2[®] relational database management system (RDBMS) products. It contains information about database planning and design, and implementation and management of database objects. This book also contains reference information for database configuration and tuning.

Who should use this book

This book is intended primarily for database and system administrators who need to design, implement and maintain a database to be accessed by local or remote clients. It can also be used by programmers and other users who require an understanding of the administration and operation of the DB2 relational database management system.

How this book is structured

This book is structured in four parts, as follows:

Part 1. Data Servers

This section briefly describes DB2 data servers, including management of their capacity and large page support in 64-bit environments on AIX[®]. In addition, it also provides information on running multiple DB2 copies on a single computer, information on the automatic features that assist you in managing your database system, information on designing, creating, and working with instances, and optional information on configuring Lightweight Directory Access Protocol (LDAP) servers.

Part 2. Databases

This section describes the design, creation, and maintenance of databases, buffer pools, table spaces, and schemas. Detailed information about database partitions is found in the new *Partitioning and Clustering Guide*.

Part 3. Database objects

This section describes the design, creation, and maintenance of the following database objects: tables, constraints, indexes, triggers, sequences and views.

Part 4. Reference

This section contains reference information for configuring and tuning your database system with environment and registry variables, and configuration parameters. It also lists the various naming rules and SQL and XML limits.

Part 1. Data servers

Chapter 1. DB2 data servers

Data servers provide software services for the secure and efficient management of structured information. DB2 is a hybrid relational and XML data server.

A data server refers to a machine where the DB2 database engine is installed. The DB2 engine is a full-function, robust database management system that includes optimized SQL support based on actual database usage and tools to help manage the data.

IBM offers a number data server products, including data server clients that can access all the various data servers. For a complete list of DB2 data server products, features available, and detailed descriptions and specifications, see: <http://www-306.ibm.com/software/data/db2/9/>.

Management of data server capacity

If data server capacity does not meet your present or future needs, you can expand its capacity by adding disk space and creating additional containers, or by adding memory. If these simple strategies do not add the capacity you need, also consider adding processors or physical partitions. When you scale your system by changing the environment, you should be aware of the impact that such a change can have on your database procedures such as loading data, or backing up and restoring databases.

Adding processors

If a single-partition database configuration with a single processor is used to its maximum capacity, you might either add processors or add database partitions. The advantage of adding processors is greater processing power. In an SMP system, processors share memory and storage system resources. All of the processors are in one system, so there are no additional overhead considerations such as communication between systems and coordination of tasks between systems. Utilities such as load, backup, and restore can take advantage of the additional processors.

Note: Some operating systems, such as the Solaris operating system, can dynamically turn processors on- and off-line.

If you add processors, review and modify some database configuration parameters that determine the number of processors used. The following database configuration parameters determine the number of processors used and might need to be updated:

- Default degree (dft_degree)
- Maximum degree of parallelism (max_querydegree)
- Enable intra-partition parallelism (intra_parallel)

You should also evaluate parameters that determine how applications perform parallel processing.

In an environment where TCP/IP is used for communication, review the value for the DB2TCPCONNMGRS registry variable.

Adding physical partitions

If your database manager is currently in a partitioned database environment, you can increase both data-storage space and processing power by adding separate single-processor or multiple-processor physical partitions. The memory and storage system resources on each database partition are not shared with the other database partitions. Although adding database partitions might result in communication and task-coordination issues, this choice provides the advantage of balancing data and user access across more than one system. The database manager supports this environment.

You can add database partitions either while the database manager system is running or while it is stopped. If you add database partitions while the system is running, however, you must stop and restart the system before databases migrate to the new database partition.

When you add a new database partition, you cannot drop or create a database that takes advantage of the new database partition until the procedure is complete, and the new server is successfully integrated into the system.

Enabling large page support in a 64-bit environment (AIX)

In addition to the traditional page size of 4 KB, the POWER4™ processors (and higher) in the IBM® eServer™ pSeries® systems also support a 16 MB page size. Applications such as IBM DB2 Version 9.1 for AIX 64-bit Edition, that require intensive memory access and that use large amounts of virtual memory can gain performance improvements by using large pages.

Note: Enabling large pages prevents the self-tuning memory manager from automatically tuning overall database memory consumption, so should only be considered for well-defined workloads that have relatively static database memory requirements.

1. For detailed instructions on how to run the `vmo` command, refer to your AIX manuals.
2. You should be extremely cautious when configuring your system for pinning memory and supporting large pages. Pinning too much memory results in heavy paging activities for the memory pages that are not pinned. Allocating too much physical memory to large pages will degrade system performance if there is insufficient memory to support the 4 KB pages.
3. Setting the `DB2_LARGE_PAGE_MEM` registry variable also implies that the memory is pinned.

You must have root authority to work with the AIX operating system commands.

1. Configure your AIX server for large page support by issuing the `vmo` command with the following flags: :

```
vmo -r -o lgpg_size=LargePageSize -o lgpg_regions=LargePages
```

where *LargePageSize* specifies the size in bytes of the hardware-supported large pages, and *LargePages* specifies the number of large pages to reserve. For example, if you need to allocate 25 GB for large page support, run the command as follows:

```
vmo -r -o lgpg_size=16777216 -o lgpg_regions=1600
```

2. Run the `bosboot` command so that the `vmo` command that you previously run will take effect following the next system boot.
3. After the server comes up, enable it for pinned memory:

- Issue the vmo command with the following flags:
`vmo -o v_pinshm=1`
- Use the db2set command to set the **DB2_LARGE_PAGE_MEM** registry variable to DB, then start DB2:
`db2set DB2_LARGE_PAGE_MEM=DB`
`db2start`

Chapter 2. Multiple DB2 copies

With Version 9 and later, you can install and run multiple DB2 copies on the same computer. A DB2 copy refers to one or more installations of DB2 database products in a particular location on the same computer. Each DB2 Version 9 copy can be at the same or different code levels.

The benefits of doing this include:

- The ability to run applications that require different DB2 versions on the same computer at the same time
- The ability to run independent copies of DB2 products for different functions
- The ability to test on the same computer before moving the production database to the latter version of the DB2 product
- For independent software vendors, the ability to embed a DB2 server product into your product and hide the DB2 database from your users. For COM+ applications, use and distribute the IBM Data Server Driver for ODBC and CLI with your application instead of the Data Server Runtime Client as only one Data Server Runtime Client can be used for COM+ applications at a time. The IBM Data Server Driver for ODBC and CLI does not have this restriction.

Default IBM database client interface copy

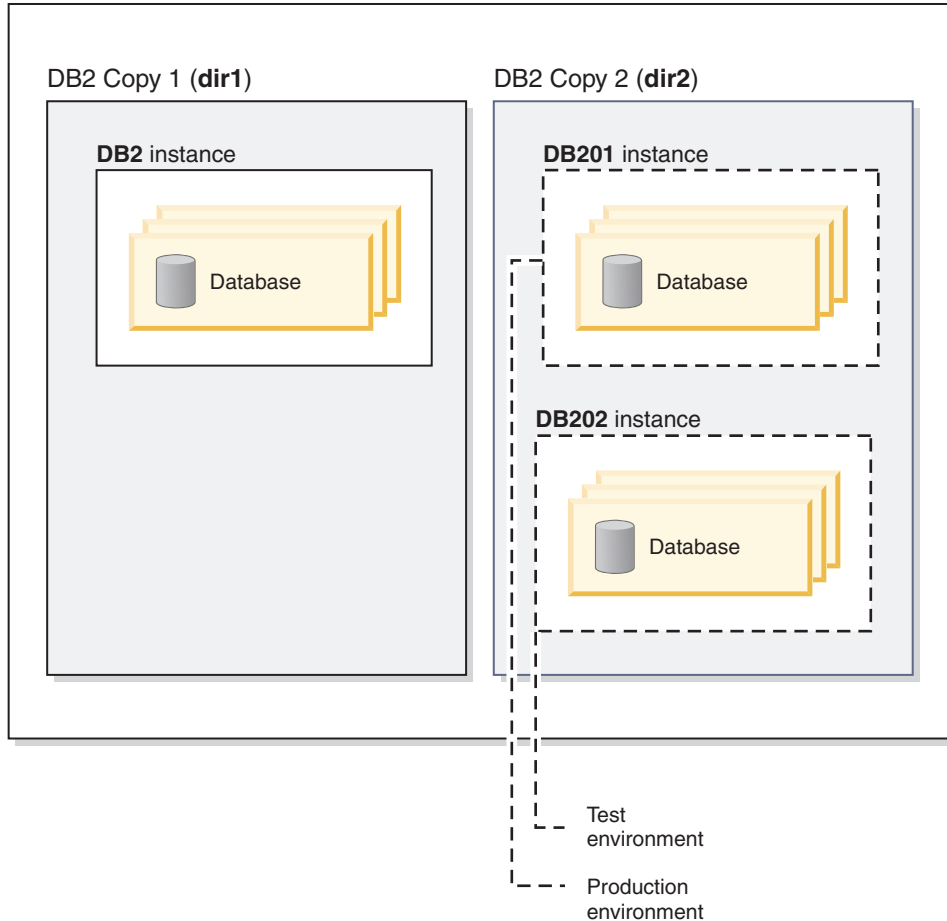
You can have multiple DB2 copies on a single computer, as well as a default IBM database client interface copy, which is the means by which a client application has the ODBC driver, CLI driver, and .NET data provider code needed to interface with the database by default.

In Version 9.1 (and later), the code for the IBM database client interface copy is included with the DB2 copy. With Version 9.5 (and later) there is a new product you can choose to install that has the needed code to allow a client application to interface with a database. This product is IBM Data Server Driver Package (DSDRIVER). With Version 9.5 (and later) you can install DSDRIVER on an IBM data server driver copy separate from the installation of a DB2 copy.

Following Version 9.1, you can have multiple DB2 copies installed on your computer; following Version 9.5, you can have multiple IBM database client interface copies and multiple DB2 copies installed on your computer. During the time of installation of a new DB2 copy or new IBM data server driver copy you would have had the opportunity to change the default DB2 copy and the default IBM database client interface copy.

The following diagram shows multiple DB2 copies installed on a DB2 server, which can be any combination of the DB2 database products:

DB2 server



Version 8 and Version 9 (or later) copies can coexist on the same computer, however Version 8 must be the default DB2 and IBM database client interface copy. You cannot change from the Version 8 copy to the Version 9 (or later) copy as the default DB2 copy or default IBM database client interface copy during installation, nor can you later run the switch default copy command, `db2swtch`, unless you first migrate or uninstall Version 8 copy. If you run the `db2swtch` command when Version 8 exists on the system, you will receive an error message indicating that you cannot change the default copy because Version 8 is found on the system.

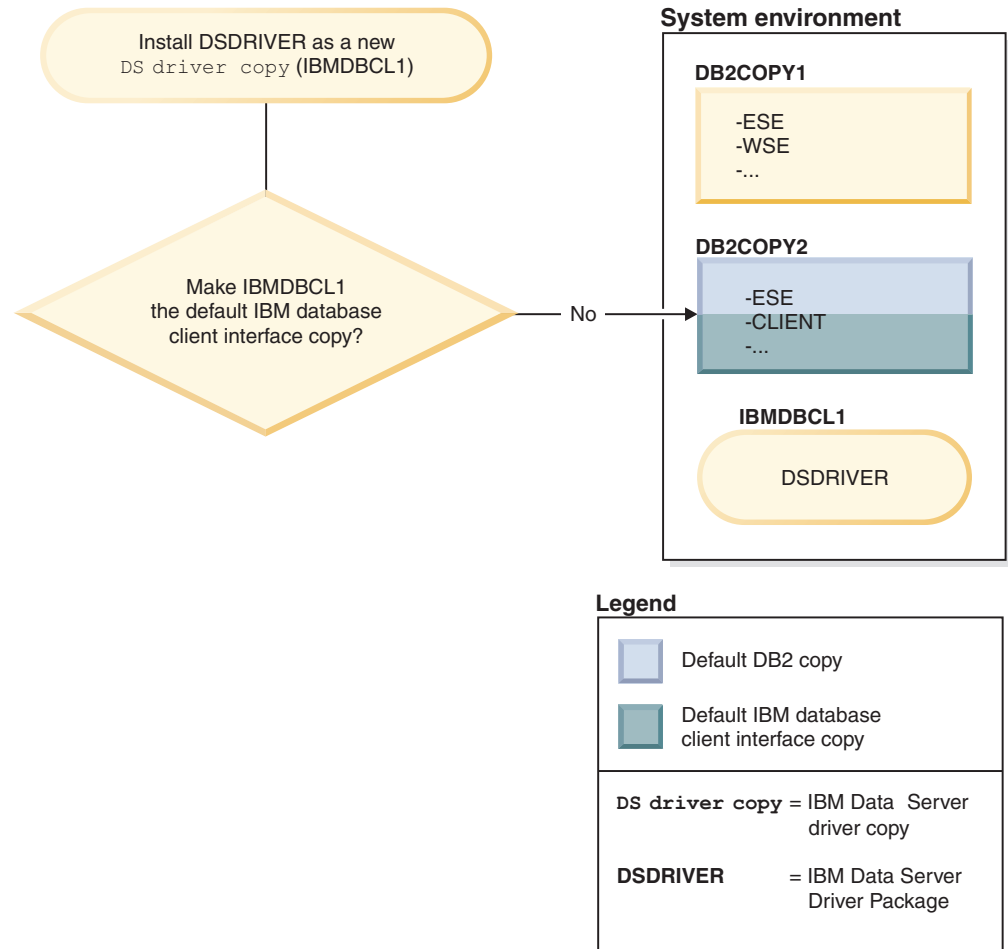
Sometime after installing multiple DB2 copies or multiple IBM data server driver copies, you may want to change either the default DB2 copy or the default IBM[®] database client interface copy. If you have Version 8 installed, you need to uninstall the product or migrate it to at least Version 9 before you can change the default DB2 copy or to at least Version 9.5 before you can change the default IBM database client interface copy.

Client applications can always choose to go directly to a data server driver location which is the directory where the `DSDRIVER` is installed.

When you uninstall either the DB2 copy or the IBM data server driver copy that had been the default IBM database client interface copy, the defaults are managed for you. Chosen default copies are removed and new defaults are selected for you. When you uninstall the default DB2 copy which is not the last DB2 copy on the system, you will be asked to switch the default to another DB2 copy first.

Choosing a default when installing a new IBM database client interface copy

Following Version 9.5, consider the scenario where you have installed two DB2 copies (DB2COPY1 and DB2COPY2). DB2COPY2 is the default DB2 copy and the default IBM database client interface copy.

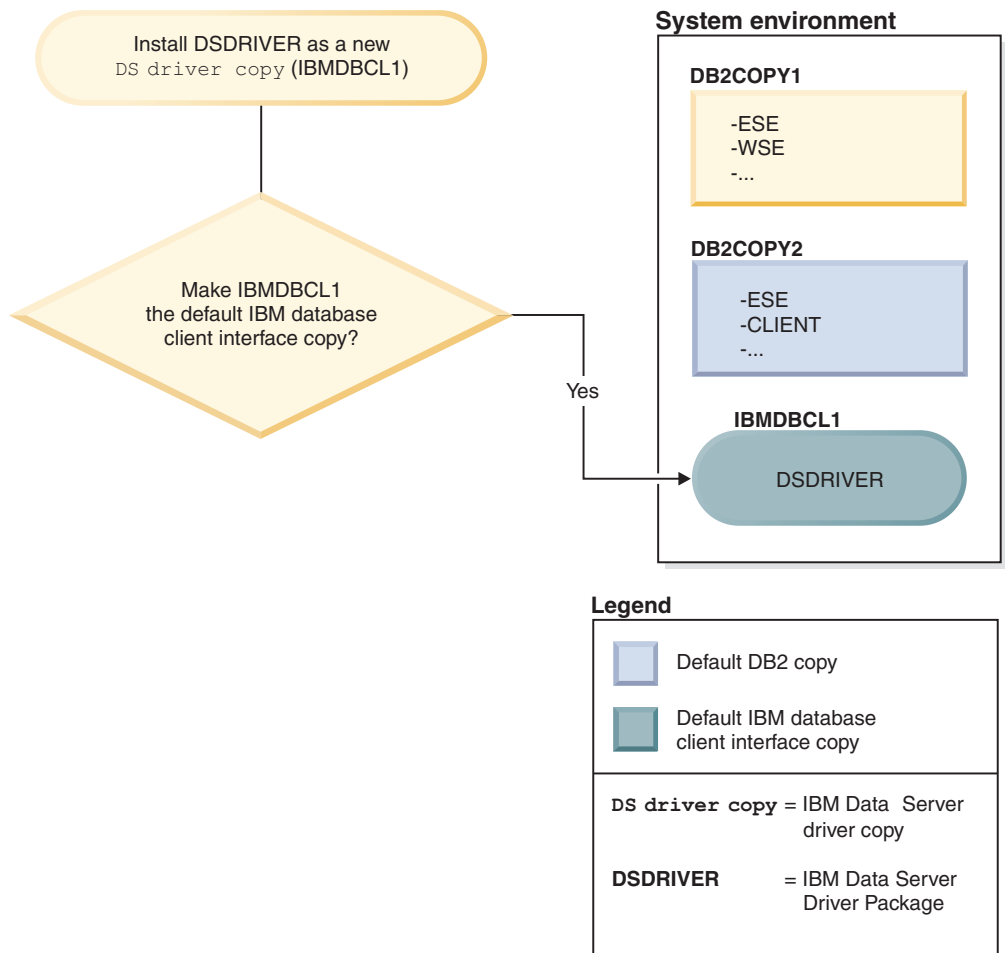


Install IBM Data Server Driver Package (DSDRIVER) on a new IBM data server driver copy.

During the install of the new IBM data server driver copy (IBMDBCL1) you are asked if you want to make the new IBM data server driver copy the default IBM database client interface copy.

If you respond “No”, then DB2COPY2 remains the default IBM database client interface copy. (And it continues to be the default DB2 copy.)

However, consider the same scenario but you respond “Yes” when asked if you want to make the new IBM data server driver copy the default IBM database client interface copy.



In this case, IBMDBCL1 becomes the default IBM database client interface copy. (DB2COPY2 remains the default DB2 copy.)

Setting the DAS when running multiple DB2 copies

Starting with Version 9, you can have multiple DB2 copies running on the same computer. This affects how the DB2 Administration Server (DAS) operates. The DAS is a unique component within the database manager that is limited to having only one version active, despite how many DB2 copies are installed on the same computer. For this reason the following restrictions and functional requirements apply.

On the server, there can be only one DAS version and it administers instances as follows:

- If the DAS is on Version 9.1 or Version 9.5, then it can administrator Version 8 and Version 9.1 or Version 9.5 instances.
- If the DAS is on Version 8, then it can administer only Version 8 instances. You can migrate your Version 8 DAS, or drop it and create a new Version 9.5 DAS to administer the Version 8 and Version 9.1 instances. This is required only if you want to use the Control Center to administer the instances.

Only one DAS can be created on a given computer at any given time despite the number of DB2 copies that are installed on the same computer. This DAS will be

used by all the DB2 copies that are on the same computer. In Version 9 or later, the DAS can belong to any DB2 copy that is currently installed.

To move the DAS between one Version 9.5 copy to another Version 9.5 copy, use the `dasupdt` command. To move the DAS between a Version 9.1 copy to a Version 9.5 copy, you cannot use `dasupdt`, you must migrate from Version 9.1 to Version 9.5.

On Windows operating systems, you can also use the `dasupdt` command when you need to move the DAS to a new Default DB2 copy in the same version.

Note:

- The `dasupdt` command can only be used to move the DAS between various DB2 copies of the same release (that is, between different Fix Packs). It cannot be used to setup DAS.
- For migration from Version 8 to Version 9.1 to Version 9.5 DAS, use the `dasmigr` command.
- If DAS is not set up, then a regular DAS setup procedure should be followed to set it up on one of the DB2 copies.

Setting the default instance when using multiple DB2 copies (Windows)

Starting with Version 9.1, the `DB2INSTANCE` environment is set according to the DB2 copy that your environment is currently set up to use. If you do not set it explicitly to an instance in the current copy, it defaults to the default instance that is specified with the `DB2INSTDEF` profile registry variable.

`DB2INSTDEF` is the default instance variable that is specific to the current DB2 copy in use. Every DB2 copy has its own `DB2INSTDEF` profile registry variable. Instance names must be unique on the system; when an instance is created, the database manager scans through existing copies to ensure its uniqueness.

Use the following guidelines to set the default instance when using multiple DB2 copies:

- If `DB2INSTANCE` is not set for a particular DB2 copy, then the value of `DB2INSTDEF` is used for that DB2 copy. This means:
 - If `DB2INSTANCE=ABC` and `DB2INSTDEF=XYZ`, `ABC` is the value that is used
 - If `DB2INSTANCE` is *not set* and `DB2INSTDEF=XYZ`, `XYZ` is used
 - If `DB2INSTANCE` is *not set* and `DB2INSTDEF` is *not set*, then any application or command that depends on a valid `DB2INSTANCE` will not work.
- You can use either the `db2envar.bat` command or the `db2SelectDB2Copy` API to switch DB2 copies. Setting all the environment variables appropriately (for example, `PATH`, `INCLUDE`, `LIB` and `DB2INSTANCE`) will also work, but you need to ensure that they are set properly.

Note: Using the `db2envar.bat` command is not quite the same as setting the environment variables. The `db2envar.bat` command determines which DB2 copy it belongs to, and then adds the path of this DB2 copy to the front of the `PATH` environment variable.

When there are multiple DB2 copies on the same machine, the PATH environment variable can only point to one of them: the DEFAULT COPY. For example, if DB2COPY1 is under c:\sqlib\bin and is the default copy; and DB2COPY2 is under d:\sqlib\bin. If you want to use DB2COPY2 in a regular command window, you would run d:\sqlib\bin\db2envar.bat in that command window. This adjusts the PATH (and some other environment variables) for this command window so that it will pick up binaries from d:\sqlib\bin.

- DB2INSTANCE is only valid for instances under the DB2 copy that you are using. However, if you switch copies by running the db2envar.bat command, DB2INSTANCE will be updated to the value of DB2INSTDEF for the DB2 copy that you switched to initially.
- DB2INSTANCE is the current DB2 instance that will be used by applications that are executing in that DB2 copy. When you switch between copies, by default, DB2INSTANCE is changed to the value of DB2INSTDEF for that copy. DB2INSTDEF is less meaningful on a one copy system because all the instances are in the current copy; however, it is still applicable as being the default instance, if another instance is not set.
- All global profile registry variables are specific to a DB2 copy, unless you specify them using SET VARIABLE=<variable_name>.

Multiple instances of the database manager

Multiple instances of the database manager might be created on a single server. This means that you can create several instances of the same product on a physical computer, and have them running concurrently. This provides flexibility in setting up environments.

Note: The same instance name cannot be used in two different DB2 copies.

You might want to have multiple instances to create the following environments:

- Separate your development environment from your production environment.
- Separately tune each environment for the specific applications it will service.
- Protect sensitive information from administrators. For example, you might want to have your payroll database protected on its own instance so that owners of other instances will not be able to see payroll data.

Note:

- (On UNIX[®] operating systems only:) To prevent environmental conflicts between two or more instances, you should ensure that each instance home directory is on a local file system.
- (On Windows[®] platforms only:) Instances are cataloged as either local or remote in the node directory. Your default instance is defined by the DB2INSTANCE environment variable. You can ATTACH to other instances to perform maintenance and utility tasks that can only be done at an instance level, such as creating a database, forcing off applications, monitoring a database, or updating the database manager configuration. When you attempt to attach to an instance that is not in your default instance, the node directory is used to determine how to communicate with that instance.
- (On any platform:) DB2 database program files are physically stored at one location and each instance points back to the copy to which that instance belongs so that the program files are not duplicated for each instance that is created. Several related databases can be located within a single instance.

Multiple instances (Windows)

It is possible to run multiple instances of the database manager on the same computer. Each instance of the database manager maintains its own databases and has its own database manager configuration parameters.

Note: The instances can also belong to different DB2 copies on a computer that can be at different levels of the database manager.

An instance of the database manager consists of the following:

- A Windows service that represents the instance. The name of the service is same as the instance name. The display name of the service (from the Services panel) is the instance name, prefixed with the “DB2 - ” string. For example, for an instance named “DB2”, there exists a Windows service called “DB2” with a display name of “DB2 - <DB2 Copy Name> - DB2”.

Note: A Windows service is not created for client instances.

- An instance directory. This directory contains the database manager configuration files, the system database directory, the node directory, the Database Connection Services (DCS) directory, all the diagnostic log and dump files that are associated with the instance. The instance directory is by default a sub-directory inside the SQLLIB directory and has the same name as the instance name. For example, the instance directory for instance “DB2” is C:\SQLLIB\DB2, where C:\SQLLIB is where the database manager is installed. You can use the registry variable DB2INSTPROF to change the default location of the instance directory. If the DB2INSTPROF registry variable is set to another location, then the instance directory is created under the directory pointed to by DB2INSTPROF. For example, if DB2INSTPROF=D:\DB2PROFS, then the instance directory will be D:\DB2PROFS\DB2.
 - Set DB2INSTPROF to c:\DB2PROFS using the db2set.exe -g command
 - Run DB2ICRT.exe command to create the instance.
- When you create an instance on Windows operating systems, the default locations for user data files, such as instance directories and the db2cli.ini file, are the following directories:
 - Documents and Settings\All Users\Application Data\IBM\DB2\copy name on the Windows XP and Windows 2003 operating systems
 - ProgramData\IBM\DB2\copy name on the Windows Vista operating system

Updating DB2 copies (Windows)

When updating your DB2 product, you will be required to specify whether you want to update an existing DB2 copy, or whether to install a new one. To update a DB2 copy, you must select the Work with Existing option.

You will not be able to update more than one DB2 copy at the same time. In order to update other DB2 copies that may be installed on the same computer, you need to rerun the installation. The installation provides the option to migrate Version 8 - or Version 9.1 copy (in the same path) or to install a new Version 9.1 or Version 9.5 copy without modifying the Version 8 installation.

- If you select to migrate, your Version 8 installation will be removed.
- If you select to install a new DB2 copy, you can later choose to migrate your instances using the db2ckmig and db2imigr commands.

You can use the `db2iupdt` command to move a DB2 instance between different Version 9.1 or Version 9.5 DB2 copies, and the `db2imigr` command to move a Version 8 instance to Version 9.1 or Version 9.5.

Note:

- Coexistence of Version 7 and Version 9.1 or Version 9.5 is not supported.
- Coexistence of a 32-bit DB2 data server and a 64-bit DB2 data server on the same Windows X64 computer is not supported.
It is not possible to migrate from a 32-bit X64 DB2 installation at Version 8 to a 64-bit installation at Version 9.1 or Version 9.5. Instead, you need to migrate to Version 9.1 or Version 9.5 32-bit to use the X64 DB2 data server installation to move to 64-bit. The 32-bit version will be removed. If you have more than one 32-bit DB2 copy installed, you will need to move all of your instances to one DB2 copy and remove these copies from the computer.
- If you have multiple Version 9.1 or Version 9.5 copies, the installation options are install a new copy or work with an existing DB2 copy, which you can upgrade or add new features. The migrate action is available if you also have a Version 8 copy in addition to the Version 9.1 or Version 9.5 copies.
- If Version 8 or Version 9.1 is installed, your installation options are to migrate the existing Version 8 or Version 9.1 to Version 9.5 copy or install a new DB2 copy.
- If Version 7 or earlier is installed, the installation displays a message to indicate that migration to Version 9.1 or Version 9.5 is not supported. You can only install a new DB2 copy after uninstalling Version 7. In other words, Version 7 and Version 9.1 or Version 9.5 cannot coexist.
- To move an instance from one Version 9.1 or Version 9.5 copy to another, you can use the `db2iupdt` command.
- If you use the `db2imigr` command to migrate your instances from Version 8, you will need to reconfigure any ODBC data sources.

Running multiple instances concurrently (Windows)

You can run multiple instances concurrently in the same DB2 copy, or in different DB2 copies.

To run multiple instances concurrently in the same DB2 copy, using the command line:

1. Set the `DB2INSTANCE` variable to the name of the other instance that you want to start by entering:

```
set db2instance=<another_instName>
```

2. Start the instance by entering the `db2start` command.

To run multiple instances concurrently in different DB2 copies, use either of the following methods:

- Using the DB2 Command window from the Start → Programs → IBM DB2 → <DB2 Copy Name> → Command Line Tools → DB2 Command Window: the Command window is already set up with the correct environment variables for the particular DB2 copy chosen.
- Using `db2envar.bat` from a Command window:
 1. Open a Command window.
 2. Run the `db2envar.bat` file using the fully qualified path for the DB2 copy that you want the application to use:

```
<DB2 Copy install dir>\bin\db2envvar.bat
```

After you switch to a particular DB2 copy, use the method specified in the section above, "To run multiple instances concurrently in the same DB2 copy", to start the instances.

Working with instances on the same or different DB2 copies

You can run multiple instances concurrently, in the same DB2 copy or in different DB2 copies.

To work with instances in the same DB2 copy, you need to:

1. Create or migrate all instances to the same DB2 copy.
2. Set the DB2INSTANCE environment variable to the name of the instance you are working with before issuing commands against that instance.

To prevent one instance from accessing the database of another instance, the database files for an instance are created under a directory that has the same name as the instance name. For example, when creating a database on drive C: for instance "DB2", the database files are created inside a directory called C:\DB2. Similarly, when creating a database on drive C: for instance TEST, the database files are created inside a directory called C:\TEST. By default, its value is the drive letter where DB2 product is installed. For more information, see the *dftdbpath* database manager configuration parameter.

To work with an instance in a system with multiple DB2 copies, use either of the following methods:

- Using the Command window from the Start → Programs → IBM DB2 → <DB2 Copy Name> → Command Line Tools → Command Window: the Command window is already set up with the correct environment variables for the particular DB2 copy chosen.
- Using db2envvar.bat from a Command window:
 1. Open a Command window.
 2. Run the db2envvar.bat file using the fully qualified path for the DB2 copy that you want the application to use:

```
<DB2 Copy install dir>\bin\db2envvar.bat
```

Chapter 3. Autonomic computing

The DB2 autonomic computing environment is self-configuring, self-healing, self-optimizing, and self-protecting. By sensing and responding to situations that occur, autonomic computing shifts the burden of managing a computing environment from database administrators to technology.

The following automatic features can assist you in managing your database system:

- Self-tuning memory
- Automatic storage
- Automatic (compression) dictionary creation (ADC)
- Automatic database backups
- Automatic statistics collection
- Configuration Advisor
- Health monitor
- Utility throttling

Automatic features

Automatic features assist you in managing your database system. They allow your system to perform self-diagnosis and to anticipate problems before they happen by analyzing real-time data against historical problem data. You can configure some of the automatic tools to make changes to your system without intervention to avoid service disruptions.

When you create a database, some of the following automatic features are enabled by default, but others you need to enable manually:

Self-tuning memory (single-partition databases only)

The self-tuning memory feature simplifies the task of memory configuration. This feature responds to significant changes in workload by automatically and iteratively adjusting the values of several memory configuration parameters and the sizes of the buffer pools, thus optimizing performance. The memory tuner dynamically distributes available memory resources among several memory consumers, including the sort function, the package cache, the lock list, and buffer pools. You can disable self-tuning memory after creating a database by setting the database configuration parameter `self_tuning_mem` to OFF.

Automatic storage

The automatic storage feature simplifies storage management for table spaces. When you create a database, you specify the storage paths where the database manager will place your table space data. Then, the database manager manages the container and space allocation for the table spaces as you create and populate them.

Automatic (compression) dictionary creation (ADC)

Compression dictionaries are automatically created during data population operations on tables for which you defined the COMPRESS attribute to YES if a compression dictionary does not already exist within the physical table or partition and after a table reaches approximately 1 MB in size as the result of data being added (through insert or load processing, for

example), the dictionary is created and is inserted into the table. Provided that the table COMPRESS attribute remains enabled, all data moved into the table after creation of the compression dictionary is subject to compression.

Automatic database backups

A database can become unusable due to a wide variety of hardware or software failures. Ensuring that you have a recent, full backup of your database is an integral part of planning and implementing a disaster recovery strategy for your system. Use automatic database backups as part of your disaster recovery strategy to enable the database manager to back up your database both properly and regularly.

Automatic statistics collection

Automatic statistics collection helps improve database performance by ensuring that you have up-to-date table statistics. The database manager determines which statistics are required by your workload and which statistics need to be updated. Statistics can be collected either asynchronously (in the background) or synchronously, by gathering runtime statistics when SQL statements are compiled. The DB2 optimizer can then choose an access plan based on accurate statistics. You can disable automatic statistics collection after creating a database by setting the database configuration parameter `auto_runstats` to OFF. Real-time statistics gathering can be enabled only when automatic statistics collection is enabled. Real-time statistics gathering is controlled by the `auto_stmt_stats` configuration parameter.

Configuration Advisor

When you create a database, this tool is automatically run to determine and set the database configuration parameters and the size of the default buffer pool (IBMDEFAULTBP). The values are selected based on system resources and the intended use of the system. This initial automatic tuning means that your database performs better than an equivalent database that you could create with the default values. It also means that you will spend less time tuning your system after creating the database. You can run the Configuration Advisor at any time (even after your databases are populated) to have the tool recommend and optionally apply a set of configuration parameters to optimize performance based on the current system characteristics.

Health monitor

The health monitor is a server-side tool that proactively monitors situations or changes in your database environment that could result in a performance degradation or a potential outage. A range of health information is presented without any form of active monitoring on your part. If a health risk is encountered, the database manager informs you and advises you on how to proceed. The health monitor gathers information about the system by using the snapshot monitor and does not impose a performance penalty. Further, it does not turn on any snapshot monitor switches to gather information.

Utility throttling

This feature regulates the performance impact of maintenance utilities so that they can run concurrently during production periods. Although the *impact policy* for throttled utilities is defined by default, you must set the *impact priority* if you want to run a throttled utility. The throttling system ensures that the throttled utilities run as frequently as possible without

violating the impact policy. Currently, you can throttle statistics collection, backup operations, rebalancing operations, and asynchronous index cleanup.

Automatic maintenance

The database manager provides automatic maintenance capabilities for performing database backups, keeping statistics current, and reorganizing tables and indexes as necessary. Performing maintenance activities on your databases is essential in ensuring that they are optimized for performance and recoverability.

Maintenance of your database includes some or all of the following activities:

- **Backups.** When you back up a database, the database manager takes a copy of the data in the database and stores it on a different medium in case of failure or damage to the original. Automatic database backups help to ensure that your database is backed up properly and regularly so that you don't have to worry about when to back up or know the syntax of the BACKUP command.
- **Data defragmentation (table or index reorganization).** This maintenance activity can increase the efficiency with which the database manager accesses your tables. Automatic reorganization manages an offline table and index reorganization so that you don't need to worry about when and how to reorganize your data.
- **Data access optimization (statistics collection).** The database manager updates the system catalog statistics on the data in a table, the data in indexes, or the data in both a table and its indexes. The optimizer uses these statistics to determine which path to use to access the data. Automatic statistics collection attempts to improve the performance of the database by maintaining up-to-date table statistics. The goal is to allow the optimizer to choose an access plan based on accurate statistics.
- **Statistics profiling.** Automatic statistics profiling advises when and how to collect table statistics by detecting outdated, missing, or incorrect statistics, and by generating statistical profiles based on query feedback.

It can be time-consuming to determine whether and when to run maintenance activities, but automatic maintenance removes the burden from you. You can manage the enablement of the automatic maintenance features simply and flexibly by using the automatic maintenance database configuration parameters. Using the Configure Automatic Maintenance wizard, you can specify your maintenance objectives. The database manager uses these objectives to determine whether the maintenance activities need to be done and runs only the required ones during the next available maintenance window (a time period that you define).

Maintenance windows

A maintenance window is a time period that you define for the running of automatic maintenance activities, which are backups, statistics collection, statistics profiling, and reorganizations. An offline window might be the time period when access to a database is unavailable. An online window might be the time period when users are permitted to connect to a database.

A maintenance window is different from a task schedule. During a maintenance window, each automatic maintenance activity is not necessarily run. Instead, the database manager evaluates the system to determine the need for each

maintenance activity to be run. If the maintenance requirements are not met, the maintenance activity is run. If the database is already well maintained, the maintenance activity is not run.

Think about when you want the automatic maintenance activities to be run. Automatic maintenance activities consume resources on your system and might affect the performance of your database when the activities are run. Some of these activities also restrict access to tables, indexes, and databases. Therefore, you must provide appropriate windows when the database manager can run maintenance activities. You specify these periods as offline and online maintenance time windows using the Automatic Maintenance wizard from the Control Center or Health Center.

Offline maintenance activities

Offline maintenance activities (offline database backups and table and index reorganizations) are maintenance activities that can occur only in the offline maintenance window. The extent to which user access is affected depends on which maintenance activity is running:

- During an offline backup, no applications can connect to the database. Any currently connected applications are forced off.
- During an offline table or index reorganization (data defragmentation), applications can access but not update the data in tables.

Offline maintenance activities run to completion even if they go beyond the window specified. Over time, the internal scheduling mechanism learns how to best estimate job completion times. If the offline maintenance window is too small for a particular database backup or reorganization activity, the scheduler will not start the job the next time and relies on the health monitor to provide notification of the need to increase the offline maintenance window.

Online maintenance activities

Online maintenance activities (automatic statistics collection and profiling, online index reorganizations, and online database backups) are maintenance activities that can occur only in the online maintenance window. When online maintenance activities run, any currently connected applications are allowed to remain connected, and new connections can be established. To minimize the impact on the system, online database backups and automatic statistics collection and profiling are throttled by the adaptive utility throttling mechanism.

Online maintenance activities run to completion even if they go beyond the window specified.

Self-tuning memory

Starting in DB2 Version 9, a new memory-tuning feature simplifies the task of memory configuration by automatically setting values for several memory configuration parameters. When enabled, the memory tuner dynamically distributes available memory resources among the following memory consumers: buffer pools, package cache, locking memory and sort memory.

The tuner works within the memory limits defined by the **database_memory** configuration parameter. The value of **database_memory** itself can be automatically tuned as well. When self-tuning is enabled for **database_memory** (when you set it to AUTOMATIC), the tuner determines the overall memory requirements for the database and increases or decreases the amount of memory allocated for database

shared memory depending on the current database requirements. For example, if the current database requirements are high and there is sufficient free memory on the system, more memory will be consumed by database shared memory. If the database memory requirements decrease or if the amount of free memory on the system becomes too low, some database shared memory is released.

If you do not enable the **database_memory** parameter for self-tuning (you do not set it to AUTOMATIC), the entire database uses the amount of memory that you specify for the parameter, distributing it across the database memory consumers as required. You can specify the amount of memory used by the database in two ways: by setting **database_memory** to a numeric value or by setting it to COMPUTED. In the second case, the total amount of memory is computed based on the sum of the initial values of the database memory heaps at database startup time.

In addition to tuning database shared memory by using the **database_memory** configuration parameter, you can enable other memory consumers for self-tuning as follows:

- For buffer pools, use the ALTER BUFFERPOOL and CREATE BUFFERPOOL statements.
- For the package cache, use the **pckcachesz** configuration parameter.
- For locking memory, use the **locklist** and **maxlocks** configuration parameters.
- For sort memory, use the **sheapthres_shr** and **sortheap** configuration parameters.

Memory allocation in DB2

Memory allocation and de-allocation occurs at various times in DB2. Memory may be allocated to a particular memory area when a specified event occurs, such as when an application connects, or it may be re-allocated based on a change in a configuration parameter setting.

The figure below shows the different areas of memory that the database manager allocates for various uses and the configuration parameters that allow you to control the size of this memory. Note that in an Enterprise Server Edition environment that comprises multiple logical database partitions, each database partition has its own Database Manager Shared Memory set.

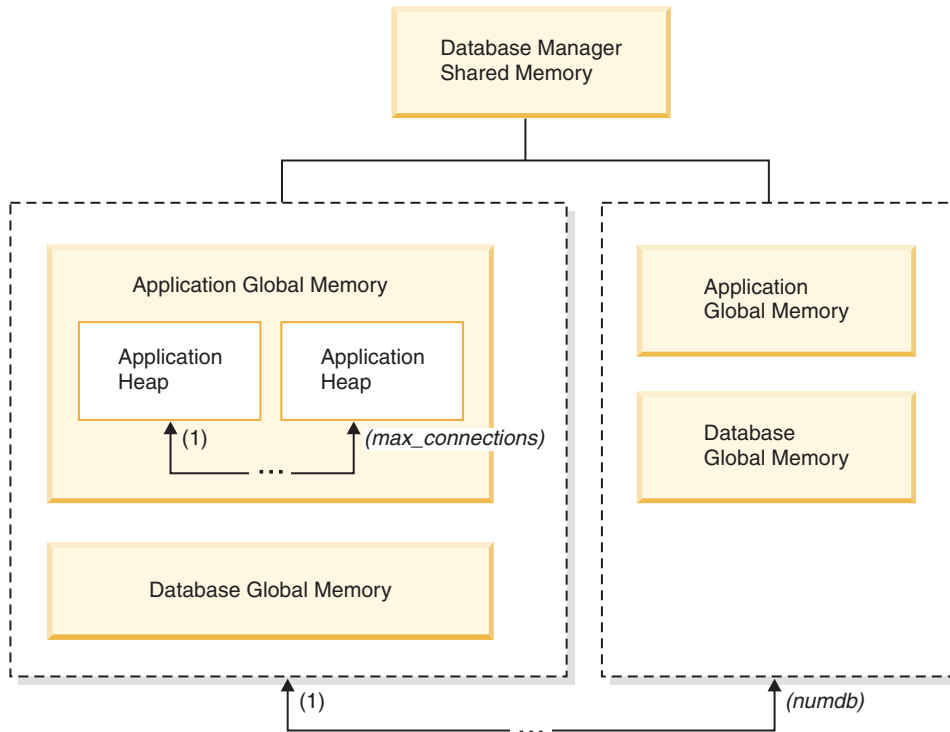


Figure 1. Types of memory used by the Database Manager

Memory is allocated for each instance of the database manager when the following events occur:

- **When the database manager is started (db2start):** Database manager global shared memory (also known as instance shared memory) is allocated and remains allocated until the database manager is stopped (db2stop). This area contains information that the database manager uses to manage activity across all database connections. DB2 automatically controls the database manager global shared memory size.
- **When a database is activated or connected to for the first time:** Database global memory is allocated. Database global memory is used across all applications that connect to the database. The size of the database global memory is specified by the **database_memory** configuration parameter. By default, this parameter is set to automatic, allowing DB2 to calculate the initial amount of memory allocated for the database and to automatically configure the database memory size during runtime based on the needs of the database. You can set **database_memory** to allocate more memory than is needed initially so that the additional memory can be dynamically distributed later.

The following memory areas can be dynamically adjusted, for example, to decrease memory allocated to one area and increase memory in another area.

- Buffer pools (using the ALTER BUFFERPOOL DDL statement)
- Database heap (including log buffers)
- Utility heap
- Package cache
- Catalog cache
- Lock list (This memory area can only be increased dynamically, and not decreased.)

In an environment in which the database manager intra-partition parallelism configuration parameter (`intra_parallel`) is enabled, in an environment in which the connection concentrator is enabled, or in an environment where the database partitioning feature (DPF) is enabled, the shared sort heap is also allocated as part of the database global memory. Also, if the **`sheapthres`** database manager configuration parameter is set to 0 (default), all sorts will use database global memory.

- **When an application connects to a database:** An application heap is allocated. Each application has its own application heap. If desired, you can limit the amount of memory any one application can allocate using the **`applheapsz`** configuration parameter, or limit overall application memory consumption using the **`appl_memory`** configuration parameter.

The database manager configuration parameter **`max_connections`** sets an upper limit on the number of applications that can attach to the instance or connect to any databases that exist in the instance. Since each application that attaches to a database involves the allocation of some memory, allowing a larger number of concurrent applications will potentially use more memory.

- **When an agent is created:** Agent private memory is allocated for an agent when the agent is assigned as the result of a connect request or a new SQL request in a parallel environment. Agent private memory is allocated for the agent and contains memory that is used only by this specific agent, such as the private sort heap.

The figure also specifies the following configuration parameter settings, which limit the amount of memory that is allocated for each type of memory area. Note that in a partitioned database environment, this memory is allocated on each database partition.

- **`numdb`**

This parameter specifies the maximum number of concurrent active databases that different applications can use. Because each database has its own global memory area, the amount of memory that might be allocated increases if you increase the value of this parameter.

- **`maxappls`**

This parameter specifies the maximum number of applications that can simultaneously connect to a single database. It affects the amount of memory that might be allocated for agent private memory and application global memory for that database. Note that this parameter can be set differently for every database.

Two other parameters that need to be considered are **`max_coordagents`** and **`max_connections`** both of which apply at the instance level (per node on a DPF instance).

- **`max_connections`**

This parameter limits the number of connections or instance attachments that can access the DB2 server at any one time (per node on a DPF instance).

- **`max_coordagents`**

This parameter limits the number of database manager coordinating agents that can exist simultaneously across all active databases in an instance (per node on a DPF instance). Together with **`maxappls`** and **`max_connections`**, this parameter limits the amount of memory allocated for agent private memory and application global memory.

The memory tracker, invoked by the `db2mtrk` command, allows you to view the current allocation of memory within the instance, including the following types of information for each memory pool:

- Current size
- Maximum size (hard limit)
- Largest size (high water mark)

Self tuning memory operational details and limitations

Determining tuning requirements

In order to ensure a fair and relevant comparison between memory consumers, a new common metric has been developed. Each tuned memory consumer calculates the predicted benefit from additional memory, and reports this to the self tuning memory process. Self tuning memory uses these figures as the basis for memory tuning, taking memory from consumers with the least need and reallocating it to those memory areas that will benefit the most.

Frequency of memory tuning

When enabled, self tuning memory will periodically check the variability of database workload. If the workload is not constant (that is, if the queries being run do not exhibit similar memory characteristics), the memory tuner will reallocate memory less frequently - up to 10 minutes between tuning cycles - to achieve more stable trend prediction. For workloads with more constant memory profiles, the memory tuner will tune memory more frequently - as little as 30 seconds between tuning cycles - in order to converge more quickly.

Tracking the progress of self tuning memory

Your current memory configuration can be obtained using the `GET DATABASE CONFIGURATION` command, or using a snapshot. Changes made by self tuning are recorded in the memory tuning log files in the `stmmlog` directory. The memory tuning log files contain summaries of the resource demands for each memory consumer at each tuning interval. These intervals can be determined based on the timestamps in the log entries.

Expected time to converge on best configuration

Leaving this feature enabled should result in quick tuning of parameters to optimize memory usage. A system can be tuned from an initial configuration in as little as one hour. In most cases, tuning will usually be complete in at most 10 hours. This worst case occurs when queries run against the database exhibit markedly different memory characteristics.

Limitations of self tuning memory

In cases where low amounts of memory are available (for example, because the value of `database_memory` is set very low, or because multiple databases, instances or other applications are running on the server) performance benefits of self tuning memory will be limited.

Because self tuning memory bases tuning decisions on database workload, workloads with changing memory characteristics limit the ability of self tuning memory to tune effectively. If your workload's memory characteristics are constantly changing, self tuning memory will tune memory less frequently, and

will repeatedly tune towards shifting target conditions. In this case, self tuning memory will not achieve absolute convergence, but will instead try to maintain a memory configuration that is tuned to the current workload.

Operational details, limitations, and interaction between memory parameters

While you can enable self-tuning memory and use the default AUTOMATIC setting on most memory-related configuration parameters, it might be useful to know the operational details, limitations, and interactions between the different memory parameters, especially the interactions between **instance_memory**, **database_memory**, and **appl_memory** parameters, in order to have more control over their setting, and also to understand why “out of memory” errors are still possible under certain conditions.

Purpose

Basically, the DB2 database manager uses two types of memory:

- *Cache-based memory*, controlled and distributed to the various performance heaps by the self-tuning memory manager (STMM). The **database_memory** configuration parameter can be used to limit the maximum amount of cache-based memory that can be used, or can be set to AUTOMATIC to let the self-tuning memory manager (STMM) manage the overall amount of cache-based memory.
- *Functional memory*, used by application programs. The **appl_memory** configuration parameter is used to control the maximum amount of application memory, or functional memory, that is allocated by DB2 database agents to service application requests. By default, its value is set to AUTOMATIC, meaning that application memory requests are allowed if the total amount of memory allocated by the database partition is within the **instance_memory** limits.

Process

In previous releases, various operating system and DB2 tools were available to see different parts of the memory, such as shared memory, private memory, buffer pool memory, locklists, sortheaps, and so forth, but it was almost impossible to see the total memory used by the DB2 database manager. When one of the heaps reached the memory limit, a statement in an application would fail with an “out of memory” error message. The DBA could increase the memory for that heap, and rerun the application, only to get an “out of memory” error on another statement for another heap. Now, individual hard upper limits on functional memory heaps can be removed by using the default AUTOMATIC configuration parameter setting.

If desired (for instance, to avoid scenarios where a poorly-behaving database application requires extremely large amounts of memory), a limit on overall application memory can be applied at the database level using the **appl_memory** configuration parameter. Individual heap limits can also be applied if desired by changing the appropriate database configuration parameter for that heap from the AUTOMATIC setting to a fixed value. If all the functional memory heaps are left at the default AUTOMATIC setting, and **appl_memory** is also left at the default AUTOMATIC setting, then the only limit on application memory consumption is the **instance_memory** setting. If **instance_memory** is also set to AUTOMATIC, then DB2 will automatically determine an upper limit on memory consumption. DBAs can easily see the total amount of **instance_memory** consumed, as well as the

current **instance_memory** limit, using the `admin_get_dbp_mem_usage` table function.

Interaction between the `self_tuning_mem`, `instance_memory`, `database_memory`, and `appl_memory` configuration parameters

When self-tuning memory is fully enabled (**self_tuning_mem** set to ON, and all the memory parameters set to AUTOMATIC), the self-tuning memory manager checks the free memory available on the system and automatically determines how much memory should be dedicated for cache-based heaps for optimal performance. All the cache-based heaps contribute to the overall **database_memory** size. In addition to the cache-based memory requirements, some memory is also required to ensure the operation and integrity of the DB2 database manager. The difference between **instance_memory** and these two memory consumers is what is left available for application memory (**appl_memory**) use. Functional memory for application programs is then allocated as needed, as long as it falls within the **instance_memory** limit - there are no additional restrictions to how much memory a single application can allocate.

The self-tuning memory manager will also periodically query how much free system memory is remaining, and how much free **instance_memory** is remaining. The self-tuning memory manager weighs application requirements more heavily than performance criteria (to prevent application failures), so will sacrifice performance by lowering cache-based heaps in order to ensure enough free system memory and **instance_memory** is available for application memory requests. As applications complete, the used memory is freed, ready to be re-used either by other applications, or to be reclaimed for **database_memory** use by the self-tuning memory manager. If performance of the database system becomes unacceptable during periods of heavy application activity, it may be useful to either apply controls on how many applications are allowed into the database manager (for instance, using either the connection concentrator or the new Workload Manager feature of DB2 9.5), or to consider adding additional memory resources to the system.

Limitations (cases where “out of memory” errors are still possible)

In some cases, you might still get “out of memory” errors if the self-tuning memory manager does not have enough time to react to sudden spikes in memory usage, for example, when an application suddenly requires a very large amount of memory, or if there is a sudden spike in your database workload (i.e. many new applications connecting to your database at the same time). In this case, or in cases where a DBA knows that most applications use a set amount of memory, it might be better to use a hard-coded value for **appl_memory** instead of the AUTOMATIC setting. If **appl_memory** is set to a hard value, for instance 2GB, then DB2 will not allow total application memory consumption to exceed that amount. Each application is then allowed to consume as much memory as required as long as total application memory consumption is less than the **appl_memory** limit. If either the **appl_memory** limit or the **instance_memory** limit is reached, the application request causing the database manager to hit the limit will fail, returning a suitable SQL code (the actual error code returned depends on exactly where in the operation of the application the “out of memory” failure was encountered). When an “out of memory” error is encountered, the DBA can view the `db2diag.log` to determine how much memory was being used when the error occurred, which can help determine whether any memory parameters need to be adjusted.

Enabling self tuning memory

Self tuning memory simplifies the task of memory configuration by automatically setting values for memory configuration parameters and sizing buffer pools. When enabled, the memory tuner dynamically distributes available memory resources between several memory consumers including sort, package cache and lock list areas and buffer pools.

1. Enable self tuning for the database by setting **self_tuning_mem** to ON. You can set **self_tuning_mem** to ON using the UPDATE DATABASE CONFIGURATION command, the sqlfupd API, or through the Change Database Configuration Parameter window in the Control Center.
2. To enable self tuning of memory areas controlled by memory configuration parameters, set the relevant configuration parameters to AUTOMATIC using the UPDATE DATABASE CONFIGURATION command, the sqlfupd API, or through the Change Database Configuration Parameter window in the Control Center.
3. To enable self tuning of buffer pools, set the buffer pool size to AUTOMATIC. You can do this using the ALTER BUFFER POOL statement for existing buffer pools or the CREATE BUFFER POOL statement for new buffer pools. If the size of a buffer pool is set to AUTOMATIC in a partitioned database environment, that buffer pool should not have any entries defined in SYSIBM.SYSBUFFERPOOLNODES.

Note:

1. Because self tuning redistributes memory between different memory areas, there must be at least two memory areas enabled for self tuning to occur, for example the lock memory area and the database shared memory area. The only exception to this is the memory controlled by the **sortheap** configuration parameter. When **sortheap** alone is set to AUTOMATIC, self tuning of **sortheap** is enabled.
2. In order to enable the **locklist** configuration parameter for self tuning, **maxlocks** must also be enabled for self tuning, therefore **maxlocks** is set to AUTOMATIC when **locklist** is set to AUTOMATIC. In order to enable the **sheapthres_shr** configuration parameter for self tuning, **sortheap** must also be enabled for self tuning, therefore **sortheap** is set to AUTOMATIC when **sheapthres_shr** is set to AUTOMATIC.
3. Automatic tuning of **sheapthres_shr** or **sortheap** is allowed only when the database manager configuration parameter **sheapthres** is set to 0.
4. Self tuning memory runs only on the HADR primary server. When self tuning memory is activated on an HADR system, it will never run on the secondary server and it will run on the primary server only if the configuration is set properly. If a command is run that switches the HADR database roles, self tuning memory operations will also switch so that they run on the new primary server. After the primary database is started or the standby database is converted to a primary database by takeover, the STMM EDU may not start until the first client connection comes in.

Disabling self tuning memory

Self tuning can be disabled for the entire database by setting **self_tuning_mem** to OFF. When **self_tuning_mem** is set to OFF, the memory configuration parameters and buffer pools that are set to AUTOMATIC remain AUTOMATIC and the memory areas remain at their current size.

You can set **self_tuning_mem** to OFF using the UPDATE DATABASE CONFIGURATION command, the sqlfupd API, or through the Change Database Configuration Parameter window in the Control Center.

Self tuning can also be effectively deactivated for the entire database if only a single memory consumer is enabled for self tuning. This is because memory cannot be redistributed when only one memory area is enabled.

For example, to disable self tuning of the **sortheap** configuration parameter, you could enter the following:

```
UPDATE DATABASE CONFIGURATION USING SORTHEAP MANUAL
```

To disable self tuning of the **sortheap** configuration parameter and change the current value of **sortheap** to 2000 at the same time, enter the following:

```
UPDATE DATABASE CONFIGURATION USING SORTHEAP 2000
```

In some cases, one memory configuration parameter can only be enabled for self tuning if another related memory configuration parameter is also enabled. For example, self tuning of the **maxlocks** configuration parameter is only permitted when the **locklist** configuration parameter is also enabled. Similarly, self tuning of the **sheapthres_shr** configuration parameter can only be enabled if self tuning of the **sortheap** configuration parameter is also enabled. This means that disabling self tuning of the **locklist** or **sortheap** parameters disables self tuning of the **maxlocks** or **sheapthres_shr** parameters, respectively.

Self tuning can be disabled for a buffer pool by setting the buffer pool to a specific size. For example, the following statement will disable self tuning for bufferpool1:

```
ALTER BUFFERPOOL bufferpool1 SIZE 1000
```

Determining which memory consumers are enabled for self tuning

To view the self tuning settings for memory consumers controlled by configuration parameters, use one of the following methods.

- To view the self tuning settings for configuration parameters from the command line, use the GET DATABASE CONFIGURATION command specifying the SHOW DETAIL parameter.

The memory consumers that can be enabled for self tuning are grouped together in the output as follows:

Description	Parameter	Current Value	Delayed Value
Self tuning memory	(SELF_TUNING_MEM)	= ON (Active)	ON
Size of database shared memory (4KB)	(DATABASE_MEMORY)	= AUTOMATIC(37200)	AUTOMATIC(37200)
Max storage for lock list (4KB)	(LOCKLIST)	= AUTOMATIC(7456)	AUTOMATIC(7456)
Percent. of lock lists per application	(MAXLOCKS)	= AUTOMATIC(98)	AUTOMATIC(98)
Package cache size (4KB)	(PCKCACHESZ)	= AUTOMATIC(5600)	AUTOMATIC(5600)
Sort heap thres for shared sorts (4KB)	(SHEAPTHRES_SHR)	= AUTOMATIC(5000)	AUTOMATIC(5000)
Sort list heap (4KB)	(SORTHEAP)	= AUTOMATIC(256)	AUTOMATIC(256)

- You can also use the db2CfgGet API to determine whether or not tuning is enabled. The following values are returned:

```
SQLF_OFF          0
SQLF_ON_ACTIVE    2
SQLF_ON_INACTIVE  3
```

SQLF_ON_ACTIVE describes a situation where self tuning is enabled and active while SQLF_ON_INACTIVE indicates that self tuning is enabled but is currently inactive.

- You can also view the configuration settings in the **Database Configuration** window in the Control Center.

To view the self tuning settings for buffer pools, use one of the following methods.

- To retrieve the list of buffer pools that are enabled for self tuning from the command line, enter:

```
db2 "select BPNAME, NPAGES from sysibm.sysbufferpools"
```

When self tuning is enabled for a buffer pool, the NPAGES field in the sysibm.sysbufferpools table for that particular buffer pool will be set to -2. When self tuning is disabled, the NPAGES field will be set to the buffer pool's current size.

- To determine the current size of buffer pools that have been enabled for self tuning, use the snapshot monitor as follows and examine the current size of the buffer pool (the value of the bp_cur_buffsz monitor element):

```
db2 get snapshot for bufferpools on db_name
```

- To view the self tuning settings of your buffer pools using the Control Center, right-click on a buffer pool and view the attributes of the buffer pools in the object details pane.

It is important to note that responsiveness of the memory tuner is limited by the time required to resize a memory consumer. For example, reducing the size of a buffer pool can be a lengthy process and therefore the performance benefits of trading buffer pool memory for sort area memory may not be immediately realized.

Self tuning memory in partitioned database environments

When using the self tuning memory feature in partitioned database environments, there are a few factors that determine whether the feature will tune the system appropriately.

When self tuning memory is enabled in partitioned databases, a single database partition is designated as the tuning partition, and all memory tuning decisions are based on the memory and workload characteristics of that database partition. Once tuning decisions are made on the tuning partition, the memory adjustments are distributed to all other database partitions to ensure that all database partitions maintain similar configurations.

The single tuning partition model necessitates that the feature be used only on database partitions that have similar memory requirements. The following are guidelines to use when determining whether to enable self tuning memory on your partitioned database.

Cases where self tuning is recommended in partitioned databases

When all database partitions have similar memory requirements and are running on similar hardware, self tuning memory can be enabled without any modifications. These types of environments share the following characteristics:

- All database partitions on identical hardware, including an even distribution of multiple logical nodes to multiple physical nodes
- Perfect or near-perfect distribution of data

- Workload running on the database partitions is distributed evenly across database partitions. This means that no one database partition has elevated memory requirements for one or more heaps.

In such an environment, it is desirable to have all database partitions configured equally, and self tuning memory will properly configure the system.

Cases where self tuning is recommended in partitioned databases with care

In cases where most of the database partitions in an environment have similar memory requirements and are running on similar hardware, it is possible to use self tuning memory as long as some care is taken with the initial configuration. These systems might have a set of database partitions for data, and a much smaller set of coordinator partitions and a catalog partitions. In such environments, it might be beneficial to configure the coordinator partitions and catalog partitions differently than the database partitions that contain your data.

In this environment, it is still possible to benefit from the self tuning memory feature with some minor setup. Since the database partitions containing the data comprise the bulk of the database partitions, self tuning should be enabled on all of these database partitions and one of these database partitions should be specified as the tuning partition. Additionally, since the catalog and coordinator partitions might be configured differently, self tuning memory should be disabled on these partitions. To disable self tuning on the catalog and coordinator partitions, update the *self_tuning_mem* database configuration parameter on these partitions to OFF.

Cases where self tuning is not recommended in partitioned databases

In environments where the memory requirements of each database partition are different or when different database partitions are running on dramatically different hardware, it is advisable to disable the self tuning memory feature. This can be done by setting the *self_tuning_mem* database configuration parameter to OFF on all partitions.

Comparing memory requirements of different database partitions

The best way to determine if the memory requirements of different database partitions are sufficiently similar is to consult the snapshot monitor. If the following snapshot elements are similar on all partitions (differing by no more than 20%), then the partitions can be considered similar.

Collect the following data by issuing the command `get snapshot for database on <dbname>`.

```
Total Shared Sort heap allocated           = 0
Shared Sort heap high water mark          = 0
Post threshold sorts (shared memory)      = 0
Sort overflows                             = 0

Package cache lookups                     = 13
Package cache inserts                     = 1
Package cache overflows                   = 0
Package cache high water mark (Bytes)     = 655360

Number of hash joins                      = 0
Number of hash loops                      = 0
```

```

Number of hash join overflows           = 0
Number of small hash join overflows     = 0
Post threshold hash joins (shared memory) = 0

Locks held currently                    = 0
Lock waits                              = 0
Time database waited on locks (ms)     = 0
Lock list memory in use (Bytes)         = 4968
Lock escalations                        = 0
Exclusive lock escalations              = 0

```

Collect the following data by issuing the command `get snapshot for bufferpools on <dbname>`

```

Buffer pool data logical reads          = 0
Buffer pool data physical reads        = 0
Buffer pool index logical reads        = 0
Buffer pool index physical reads       = 0
Total buffer pool read time (milliseconds) = 0
Total buffer pool write time (milliseconds) = 0

```

Using self-tuning memory in partitioned database environments

When self-tuning is enabled in partitioned database environments, there is a single database partition, known as the *tuning partition*, that monitors the memory configuration and propagates any configuration changes to all other database partitions to maintain a consistent configuration across all the participating database partitions.

The tuning partition is selected based on a number of characteristics, such as the number of database partitions in the partition group and the number of buffer pools defined.

- To determine which database partition is currently specified as the tuning partition, use the following `ADMIN_CMD`:
`CALL SYSPROC.ADMIN_CMD('get stmm tuning dbpartitionnum')`
- To change the tuning partition, use the following `ADMIN_CMD`:
`CALL SYSPROC.ADMIN_CMD('update stmm tuning dbpartitionnum <db_partition_num>')`

When you issue this command, the tuning partition will be updated asynchronously or at the next database startup.

- To have the memory tuner automatically re-select the tuning partition, enter `-1` for the `<db_partition_num>` value.

Starting the memory tuner on DPF systems

The memory tuner will only be started in a DPF environment if the database is activated by an explicit `ACTIVATE DATABASE` command because self-tuning requires all partitions to be active before it can properly tune the memory on a multi-partition system.

Disabling self-tuning for a given database partition

- To disable self-tuning for a subset of database partitions, set the `self_tuning_mem` configuration parameter to `OFF` for the database partitions you want to leave untuned.
- To disable self-tuning for a subset of the memory consumers controlled by configuration parameters on a particular database partition, set the value of the

relevant configuration parameter or buffer pool size to MANUAL or a specific value on that database partition. However, it is recommended that self-tuning configuration parameter values be consistent across all running partitions.

- To disable tuning for a particular buffer pool on a database partition, issue an ALTER BUFFER POOL command specifying a size value and a value for the PARTITIONNUM parameter for the partition where self-tuning is to be disabled.

An ALTER BUFFERPOOL statement that specifies the size on a particular database partition using the PARTITIONNUM clause will create an exception entry for the given buffer pool in the SYSCAT.SYSBUFFERPOOLNODES catalog, or update the exception entry if one already exists. When an exception entry exists for a buffer pool in this catalog, that buffer pool will not participate in self-tuning when the default buffer pool size is set to AUTOMATIC. To remove an exception entry so that a buffer pool can be re-enabled for self-tuning:

1. Disable tuning for this buffer pool by issuing an ALTER BUFFERPOOL statement setting the buffer pool size to a specific value.
2. Issue another ALTER BUFFERPOOL statement with the PARTITIONNUM clause specified to set the size of the buffer pool on this database partition to the default buffer pool size.
3. Enable self-tuning by issuing another ALTER BUFFERPOOL statement setting the size to AUTOMATIC.

Enabling self-tuning memory in non-uniform environments

Ideally, your data should be distributed evenly across all of your database partitions and the workload run on each partition should have similar memory requirements. If the data distribution is skewed so that one or more of your database partitions contain significantly more or less data than other database partitions, these anomalous database partitions should not be enabled for self-tuning. The same is true if the memory requirements are skewed across the database partitions, which can happen, for example, if resource-intensive sorts are only performed on one partition, or if some database partitions are associated with different hardware and more available memory than others. Self-tuning can still be enabled on some database partitions in this type of environment. To take advantage of self-tuning memory in environments with skew, identify a set of database partitions that have similar data and memory requirements and enable them for self-tuning. Memory configuration in the remaining partitions should be configured manually.

Configuring memory and memory heaps

With the simplified memory configuration feature, you can configure memory and memory heaps required by the DB2 data server by using the default AUTOMATIC setting for most memory-related configuration parameters, thereby, requiring much less tuning.

The simplified memory configuration feature provides the following benefits:

- You can use a single parameter, **instance_memory**, to specify all of the memory that the database manager is allowed to allocate from its private and shared memory heaps. Also, you can use the **appl_memory** configuration parameter to control the maximum amount of application memory that is allocated by DB2 database agents to service application requests.
- You do not need to manually tune parameters used solely for functional memory.

- You can query how much total memory is currently being consumed by the private and shared memory heaps of the database manager by using the Memory Visualizer. You can also use the db2mtrk command to monitor heap usage and the ADMIN_GET_DBP_MEM_USAGE() table function to query overall memory consumption.
- The default DB2 configuration requires much less tuning, a benefit for new instances that you create.

The following table lists the memory configuration parameters whose values default to the AUTOMATIC setting. These parameters can also be configured dynamically, if necessary. Note that the meaning of the AUTOMATIC setting differs with each parameter, as described in the rightmost column.

Table 1. Memory configuration parameters whose values default to AUTOMATIC

Configuration parameter name	Description	Meaning of the AUTOMATIC setting
appl_memory	Controls the maximum amount of application memory that is allocated by DB2 database agents to service application requests.	The AUTOMATIC setting allows all application memory requests as long as the total amount of memory allocated by the database partition is within the instance_memory limits.
applheapsz	Prior to Version 9.5, referred to the amount of application memory that each database agent working for an application could consume. In Version 9.5, this parameter refers to the total amount of application memory that can be consumed by the entire application. For DPF, Concentrator, or SMP configurations, this means that you might need to increase the applheapsz value used in previous releases unless you use the AUTOMATIC setting.	The AUTOMATIC setting allows the application heap size to increase as needed until either the appl_memory or the instance_memory limit is reached.
database_memory (Prior to Version 9.5, the default setting of AUTOMATIC applied only to Windows and AIX platforms. As of Version 9.5, AUTOMATIC is the default setting for all DB2 server products.)	Specifies the amount of shared memory that is reserved for the database shared memory region.	When enabled, the memory tuner determines the overall memory requirements for the database and increases or decreases the amount of memory allocated for database shared memory depending on the current database requirements.
dbheap	Determines the maximum memory used by the database heap.	The AUTOMATIC setting allows the database heap to increase as needed until either the database_memory or the instance_memory limit is reached.

Table 1. Memory configuration parameters whose values default to AUTOMATIC (continued)

Configuration parameter name	Description	Meaning of the AUTOMATIC setting
instance_memory	Specifies the maximum amount of memory that can be allocated for a database partition.	The AUTOMATIC setting allows the overall memory consumed by the entire database manager instance to increase to a limit of 75 - 95% of the physical RAM on the machine. This limit is calculated during db2start processing.
mon_heap_sz	Determines the amount of the memory, in pages, to allocate for database system monitor data.	The AUTOMATIC setting allows the monitor heap to increase as needed until the instance_memory limit is reached.
stat_heap_sz	Indicates the maximum size of the heap used in collecting statistics using the RUNSTATS command.	The AUTOMATIC setting allows the statistics heap size to increase as needed until either the appl_memory or the instance_memory limit is reached.
stmtheap	Specifies the size of the statement heap, which is used as a work space for the SQL or XQuery compiler to compile an SQL or XQuery statement.	The AUTOMATIC setting allows the statement heap to increase as needed until either the appl_memory or the instance_memory limit is reached.

Note: The DBMCFG and DBCFG administrative views retrieve database manager configuration parameter information for the currently connected database for all database partitions. For the **mon_heap_sz**, **stmtheap**, and **stat_heap_sz** configuration parameters, the DEFERRED_VALUE column on this view does not persist across database activations. That is, when you issue the get dbm cfg show detail or get db cfg show detail command, the output from the query shows updated (in memory) values.

The following table shows whether configuration parameters are set to the default AUTOMATIC value during instance migration or creation and during database migration or creation.

Table 2. Configuration parameters set to AUTOMATIC during instance and database migration and creation

Configuration parameters	Set to AUTOMATIC upon instance migration or instance creation	Set to AUTOMATIC upon database migration	Set to AUTOMATIC upon database creation
applheapsz ¹		X	X
dbheap		X	X
instance_memory	X		
mon_heap_sz ¹	X		
stat_heap_sz ¹		X	X

Table 2. Configuration parameters set to AUTOMATIC during instance and database migration and creation (continued)

Configuration parameters	Set to AUTOMATIC upon instance migration or instance creation	Set to AUTOMATIC upon database migration	Set to AUTOMATIC upon database creation
stmtheap ¹			X

As part of the move to simplified memory configuration, the following elements have been deprecated:

- Configuration parameters **appgroup_mem_sz**, **groupheap_ratio**, and **app_ctl_heap_sz**. These configuration parameters are replaced with the new **appl_memory** configuration parameter.
- The **-p** parameter of the db2mtrk memory tracker command. This option, which lists private agent memory heaps, is replaced with the **-a** parameter, which lists all application memory consumption.

The Memory Visualizer displays the maximum application memory consumption by a database using the new **appl_memory** configuration parameter, and the maximum memory consumption by an instance using the updated **instance_memory** configuration parameter. The Memory Visualizer also displays the values for all of the configuration parameters that allow the AUTOMATIC setting. Values for the deprecated configuration parameters are not displayed in the Memory Visualizer for Version 9.5 databases, but they are displayed for earlier versions of the databases.

Attempts to update the **instance_memory** parameter to values larger than those specified in this list will fail with a SQL5130N return code:

- 4 GB (1 048 576 * 4 KB pages) for DB2 Express Edition and DB2 Express-C
- 16 GB (4 194 304 * 4 KB pages) for DB2 Workgroup Server Edition

When fast communications manager (FCM) shared memory is allocated, each local database partition's share of the overall FCM shared memory size for the system is accounted for in the **instance_memory** limit of that database partition. Due to the nature of FCM memory (failure to allocate FCM buffers can bring down the instance), FCM memory requests never fail due to the **instance_memory** limit. However, they can fail if memory cannot be allocated from the operating system. If an FCM memory request causes a database partition to exceed its **instance_memory** limit, other memory requests will fail until the memory usage of the partition returns to a level below the **instance_memory** limit.

Agent and process model configuration

Version 9.5 provides a less complex and more flexible mechanism for configuring process model-related parameters. This simplified configuration eliminates the need for regular adjustments to these parameters and reduces the time and effort required to configure them. It also eliminates the need to shut down and restart DB2 instances to have the new values take effect.

To allow for dynamic and automatic agent and memory configuration, slightly more memory resources are required when an instance is activated.

Agent, process model, and memory configuration

DB2 data servers exploit multithreaded architecture on both 32-bit and 64-bit platforms to provide you with a number of benefits, such as enhanced usability, better sharing of resources, memory footprint reduction, and consistent threading architecture across all operating systems.

Configuring databases across multiple partitions

The database manager provides a single view of all database configuration elements across multiple partitions. This means that you can update or reset a database configuration across all database partitions without invoking the `db2_all` command against each database partition.

You can update a database configuration across partitions by issuing only one SQL statement or only one administration command from any partition on which the database resides. By default, the method of updating or resetting a database configuration is *on all database partitions*.

For backward compatibility of command scripts and applications, you have three options:

- Use the `db2set` command to set the `DB2_UPDDBCFG_SINGLE_DBPARTITION` registry variable to `TRUE`, as follows:

```
DB2_UPDDBCFG_SINGLE_DBPARTITION=TRUE
```

Note: Setting the registry variable does not apply to `UPDATE DATABASE CONFIGURATION` or `RESET DATABASE CONFIGURATION` requests that you make using the `ADMIN_CMD` procedure.

- Use the `DBPARTITIONNUM` parameter with either the `UPDATE DATABASE CONFIGURATION` or the `RESET DATABASE CONFIGURATION` command or with the `ADMIN_CMD` procedure. For example, to update the database configurations on all database partitions, call the `ADMIN_CMD` procedure as follows:

```
CALL SYSPROC.ADMIN_CMD  
('UPDATE DB CFG USING sortheap 1000')
```

To update a single database partition, call the `ADMIN_CMD` procedure as follows:

```
CALL SYSPROC.ADMIN_CMD  
('UPDATE DB CFG DBPARTITIONNUM 10 USING sortheap 1000')
```

- Use the `DBPARTITIONNUM` parameter with the `db2CfgSet` API. The flags in the `db2Cfg` structure indicate whether the value for the database configuration is to be applied to a single database partition. If you set a flag, you must also provide the `DBPARTITIONNUM` value, for example:

```
#define db2CfgSingleDbpartition      256
```

If you do not set the `db2CfgSingleDbpartition` value, the value for the database configuration applies to all database partitions unless you set the `DB2_UPDDBCFG_SINGLE_DBPARTITION` registry variable to `TRUE` or you set `versionNumber` to anything that is less than the version number for Version 9.5, for the `db2CfgSet` API that sets the database manager or database configuration parameters.

When migrating your databases to Version 9.5, existing database configuration parameters, as a general rule, retain their values after migration. However, new parameters are added using their default values and some existing parameters are set to their new Version 9.5 default values. Refer to the "DB2 server behavior

changes" topic in the *Migration Guide* for details about the changes to existing database configuration parameters. Any subsequent update or reset database configuration requests for the migrated databases will use the Version 9.5 method of updating or resetting configuration requests.

For existing update or reset command scripts, the same rules mentioned previously apply: you can use the pre-Version 9.5 method, you can modify your scripts to include the **DBPARTITIONNUM** option of the **UPDATE DATABASE CONFIGURATION** or **RESET DATABASE CONFIGURATION** command, or you can set the **DB2_UPDDBCFG_SINGLE_DBPARTITION** registry variable.

For existing applications that call the **db2CfgSet** API, you must use the Version 9.5 method. If you want the pre-Version 9.5 method, you can set the **DB2_UPDDBCFG_SINGLE_DBPARTITION** registry variable or modify your applications to call the API with the Version 9.5 version number, including the new **db2CfgSingleDbpartition** flag and the new **dbpartitionnum** field to update or reset database configurations for a specific database partition.

Note: If you find that database configuration values are inconsistent, you can update or reset each database partition individually.

Shared file handle table

The threaded database manager maintains a single shared file handle table for each database and all agents working on each database so that I/O requests made on the same file do not require the file to be reopened and closed.

Prior to Version 9.5, the file handle table was maintained separately by each DB2 agent, and the size of the per-agent file handle table was controlled by the **maxfilop** configuration parameter. Starting in Version 9.5, the database manager maintains a single shared file handle table for the entire database, such that the same file handle can be shared among all agents working on the same database file. As a result, the **maxfilop** configuration parameter is used to control the size of the shared file handle table.

Because of this change, the **maxfilop** configuration parameter has a new default value and new minimum and maximum values. During database migration, the **maxfilop** configuration parameter is automatically set to the new default values.

Running vendor library functions in fenced-mode processes

The database manager supports vendor library functions in fenced-mode processes that perform such tasks as data compression, TSM backups, and log data archiving.

Prior to Version 9.5, vendor library functions, vendor utilities, or routines were run inside agent processes. Since Version 9.5, because the DB2 database manager itself is a multithreaded application, vendor library functions that are no longer threadsafe and cause memory or stack corruption or, worse, data corruption in DB2 databases. For these reasons, a new fenced-mode process is created for each invocation of a vendor utility, and vendor library functions or routines run inside this fenced-mode process. This does not result in significant performance degradation.

Note: The fenced-mode feature is not available for Windows platforms.

Automatic storage

Automatic storage simplifies storage management for table spaces. When you create a database, you specify the storage paths where the database manager will place your table space data. Then, the database manager will manage the container and space allocation for the table spaces as you create and populate them.

Automatic storage table spaces

When you create a table space in a database that is not enabled for automatic storage, you must specify the `MANAGED BY SYSTEM` or `MANAGED BY DATABASE` clause. Using these clauses results in the creation of a system-managed space (SMS) table space or database-managed space (DMS) table space, respectively. You must provide an explicit list of containers in both cases.

If a database is enabled for automatic storage, other choices exist: you can specify the `MANAGED BY AUTOMATIC STORAGE` clause or omit the `MANAGED BY` clause (which implies the use of automatic storage). You do not need to provide container definitions in this case because the database manager assigns the containers automatically.

Following are some examples of statements that create automatic storage table spaces:

```
CREATE TABLESPACE TS1
CREATE TABLESPACE TS2 MANAGED BY AUTOMATIC STORAGE
CREATE TEMPORARY TABLESPACE TEMPTS
CREATE USER TEMPORARY TABLESPACE USRTMP MANAGED BY AUTOMATIC STORAGE
CREATE LONG TABLESPACE LONGTS
```

Although the automatic storage table space type appears to be a different table space type, it is really just an extension of the existing SMS and DMS types. If you create a table space as a regular or large table space, it is created as a DMS table space with file containers. If you create a table space as a user or system temporary table space, it is created as an SMS table space with directory containers.

Note: This behavior might change in future versions of the database manager.

The names associated with these containers have the following format:

```
storage path/instance name/NODE####/database name/T#####/C#####.EXT
```

where:

storage path

Is a storage path associated with the database

instance name

Is the instance under which the database was created

database name

Is the name of the database

NODE####

Is the database partition number (for example, NODE0000)

T#####

Is the table space ID (for example, T0000003)

C#####

Is the container ID (for example, C0000012)

- EXT** Is an extension based on the type of data being stored:
- CAT** System catalog table space
- TMP** System temporary table space
- UTM** User temporary table space
- USR** User or regular table space
- LRG** Large table space

Differences between regular and large automatic storage table spaces and DMS table spaces

Regular and large automatic storage table spaces are created as DMS table spaces, and all of the rules and behaviors associated with DMS table spaces apply. However, there are differences with respect to how storage is managed, as shown in the following table:

Table 3. Differences between managing non-automatic storage and automatic storage table spaces

Non-automatic storage	Automatic storage
You must explicitly provide a list of containers when creating the table space.	You cannot provide a list of containers when creating the table space; instead, the database manager automatically assigns and allocates containers.
Automatic resizing of table spaces is off (AUTORESIZE is set to NO) by default.	Automatic resizing of table spaces is on (AUTORESIZE is set to YES) by default.
You cannot use the INITIALSIZE clause to specify the initial size for the table space.	You can use the INITIALSIZE clause to specify the initial size for the table space.
You can perform container operations using the ALTER TABLESPACE statement (specifying ADD, DROP, BEGIN NEW STRIPE SET, and so on).	You cannot perform container operations because the database manager manages space.
You can use a redirected restore operation to redefine the containers associated with the table space.	You cannot use a redirected restore operation to redefine the containers associated with the table space because the database manager manages space.

As mentioned in the previous table, when you create a regular or large automatic storage table space, you can specify an initial size as part of the CREATE TABLESPACE statement, as shown in the following example:

```
CREATE TABLESPACE TS1 INITIALSIZE 100 M
```

If you do not specify an initial size, the database manager uses a default value of 32 megabytes.

To create a table space with a given size, the database manager creates file containers within the storage paths. If there is an uneven distribution of space among the paths, containers might be created with different sizes. As a result, it is important that all of the storage paths have a similar amount of free space on them.

If you enable automatic resizing for a table space, as space is used within it, the database manager automatically extends existing containers and adds new ones

(using stripe sets). Whether containers are extended or added, no rebalancing takes place.

Automatic re-sizing of table spaces

Enabling automatic storage table spaces for automatic resizing allows the database manager to handle the full file system condition automatically by adding a new stripe set of containers.

Two table space types can exist within a database system: system-managed space (SMS) and database-managed space (DMS). The containers associated with SMS table spaces are file system directories, and the files within these directories grow as the objects in the table space grow. The files grow until a file system limit is reached for one of the containers or until the table space size limit of the database is reached (see SQL and XML limits).

DMS table spaces are made up of file containers or raw device containers, and their sizes are set when the containers are assigned to the table space. The table space is considered to be full when all of the space within the containers has been used. However, unlike for SMS table spaces, you can add or extend containers using the ALTER TABLESPACE statement, allowing more storage space to be given to the table space. DMS table spaces also have a feature called *auto-resize*: as space is consumed in a DMS table space that can be automatically resized, the database system might extend the table space by one or more file containers. SMS table spaces have similar capabilities for growing automatically, but the term auto-resize is used exclusively for DMS.

Automatic resizing of table spaces has the following implications:

- Table spaces that are enabled for automatic resizing have metadata associated with them that is not recognized by Version 8.2.1 or earlier releases. Any attempt to use a database with table spaces enabled for automatic resizing with these versions results in a failure (most likely, returning an SQL0980C or SQL0902C error). An error might be sent if you try to connect to a database or try to restore a database. If you enabled table spaces for automatic resizing, disabling the auto-resize feature for these table spaces removes the metadata, allowing the database to be used with Version 8.2.1 or earlier releases.
- If you disable the auto-resize feature, the values that are associated with INCREASESIZE and MAXSIZE are lost if you subsequently enable this feature.
- You cannot this feature for table spaces that use raw device containers, and you cannot add raw device containers to a table space that can be automatically resized. Attempting these operations results in errors (SQL0109N). If you need to add raw device containers, you must disable the feature first.
- A redirected restore operation cannot change the container definitions to include a raw device container. Attempting this kind of operation results in an error (SQL0109N).
- Because the maximum size limits how the database manager automatically increases a table space, the maximum size also limits how you can increase a table space. That is, when you perform an operation that adds space to a table space, the resulting size must be less than or equal to the maximum size. You can add space by using the ADD, EXTEND, RESIZE, or BEGIN NEW STRIPE SET clause of the ALTER TABLESPACE statement.

Enabling and disabling the auto-resize feature

By default, the auto-resize feature is not enabled for a DMS table space. The following statement creates a DMS table space without enabling auto-resize:

```
CREATE TABLESPACE DMS1 MANAGED BY DATABASE
  USING (FILE '/db2files/DMS1' 10 M)
```

To enable the auto-resize feature, specify the `AUTORESIZE YES` clause for the `CREATE TABLESPACE` statement:

```
CREATE TABLESPACE DMS1 MANAGED BY DATABASE
  USING (FILE '/db2files/DMS1' 10 M) AUTORESIZE YES
```

You can also enable or disable the auto-resize feature after creating a DMS table space by using `ALTER TABLESPACE` statement with the `AUTORESIZE` clause:

```
ALTER TABLESPACE DMS1 AUTORESIZE YES
ALTER TABLESPACE DMS1 AUTORESIZE NO
```

Two other attributes, `MAXSIZE` and `INCREASESIZE`, are associated with auto-resize table spaces:

Maximum size (`MAXSIZE`)

The `MAXSIZE` clause of the `CREATE TABLESPACE` statement defines the maximum size for the table space. For example, the following statement creates a table space that can grow to 100 megabytes (per database partition if the database has multiple database partitions):

```
CREATE TABLESPACE DMS1 MANAGED BY DATABASE
  USING (FILE '/db2files/DMS1' 10 M)
  AUTORESIZE YES MAXSIZE 100 M
```

The `MAXSIZE NONE` clause specifies that there is no maximum limit for the table space. The table space can grow until a file system limit or table space limit is reached (see `SQL` and `XML` limits). If you do not specify the `MAXSIZE` clause, there is no maximum limit when the auto-resize feature is enabled.

Use the `ALTER TABLESPACE` statement to change the value of `MAXSIZE` for a table space that has auto-resize already enabled, as shown in the following examples:

```
ALTER TABLESPACE DMS1 MAXSIZE 1 G
ALTER TABLESPACE DMS1 MAXSIZE NONE
```

If you specify a maximum size, the actual value that the database manager enforces might be slightly smaller than the value specified because the database manager attempts to keep container growth consistent.

Increase size (`INCREASESIZE`)

The `INCREASESIZE` clause of the `CREATE TABLESPACE` statement defines the amount of space used to increase the table space when there are no free extents within the table space but a request for one or more extents was made. You can specify the value as an explicit size or as a percentage, as shown in the following examples:

```
CREATE TABLESPACE DMS1 MANAGED BY DATABASE
  USING (FILE '/db2files/DMS1' 10 M)
  AUTORESIZE YES INCREASESIZE 5 M

CREATE TABLESPACE DMS1 MANAGED BY DATABASE
  USING (FILE '/db2files/DMS1' 10 M)
  AUTORESIZE YES INCREASESIZE 50 PERCENT
```


A percentage value means that the amount by which to increase, as specified by the INCREASESIZE value, is calculated every time that the table space needs to grow and that growth is based on a percentage of the table space size at that time. For example, if the table space is 20 MB in size and the INCREASESIZE value is 50% , the table space grows by 10 MB the first time (to a size of 30 MB) and by 15 MB the next time.

If you do not specify the INCREASESIZE clause when you enable the auto-resize feature, the database manager determines an appropriate value to use, which might change over the life of the table space. As with AUTORESIZE and MAXSIZE, you can change the value of INCREASESIZE using the ALTER TABLESPACE statement.

If you specify a size increase, the actual value that the database manager will use might be slightly different than the value that you provide. This adjustment in the value used is done to keep growth consistent across the containers in the table space.

How table spaces are extended

For table spaces that can be automatically resized, the database manager attempts to increase the size of the table space when all of the existing space has been used and a request for more space is made. The database manager determines which of the containers can be extended in the table space so that a rebalance does not occur. The database manager extends only those containers that exist within the last range of the table space map (the map describes the storage layout for the table space) and extends them by an equal amount.

For example, consider the following statement:

```
CREATE TABLESPACE TS1 MANAGED BY DATABASE
  USING (FILE 'C:\TS1CONT' 1000, FILE 'D:\TS1CONT' 1000,
        FILE 'E:\TS1CONT' 2000, FILE 'F:\TS1CONT' 2000)
  EXTENTSIZE 4
  AUTORESIZE YES
```

Keeping in mind that the database manager uses a small portion (one extent) of each container for metadata, following is the table space map that is created for the table space based on the CREATE TABLESPACE statement. (The table space map is part of the output from a table space snapshot.)

Table space map:

Range Number	Stripe Set	Stripe Offset	Max Extent	Max Page	Start Stripe	End Stripe	Adj.	Containers
[0]	[0]	0	995	3983	0	248	0	4 (0,1,2,3)
[1]	[0]	0	1495	5983	249	498	0	2 (2,3)

The table space map shows that the containers with an identifier of 2 or 3 (E:\TS1CONT and F:\TS1CONT) are the only containers in the last range of the map. Therefore, when the database manager automatically extends the containers in this table space, it extends only those two containers.

Note: If you create a table space with all of the containers having the same size, there is only one range in the map. In such a case, the database manager extends each of the containers. To prevent restricting extensions to only a subset of the containers, create a table space with containers of equal size.

As discussed previously, you can specify a limit on the maximum size of the table space, or you can specify a value of NONE, which does not limit growth. If you specify NONE or no limit, the upper limit is defined by the file system limit or by the table space limit; the database manager does not attempt to increase the table space size past the upper limit. However, before that limit is reached, an attempt to increase a container might fail due to a full file system. In this case, the database manager does not increase the table space size any further and returns an out-of-space condition to the application. There are two ways to resolve this situation:

- Increase the amount of space available on the file system that is full.
- Perform container operations on the table space such that the container in question is no longer in the last range of the table space map. The easiest way to do this is to add a new stripe set to the table space with a new set of containers, and the best practice is to ensure that the containers are all the same size. You can add new stripe sets by using the ALTER TABLESPACE statement with the BEGIN NEW STRIPE SET clause. By adding a new stripe set, a new range is added to the table space map. With a new range, the containers that the database manager automatically attempts to extend are within this new stripe set, and the older containers remain unchanged.

Note: When a user-initiated container operation is pending or a subsequent rebalance is in progress, the auto-resize feature is disabled until the operation is committed or the rebalance is complete.

For example, for DMS table spaces, suppose that a table space has three containers that are the same size and that each resides on its own file system. As work is done on the table space, the database manager automatically extends these three containers. Eventually, one of the file systems becomes full, and the corresponding container can no longer grow. If more free space cannot be made available on the file system, you must perform container operations on the table space such that the container in question is no longer in the last range of the table space map. In this case, you could add a new stripe set specifying two containers (one on each of the file systems that still has space), or you could specify more or fewer containers (again, making sure that each container being added is the same size and that there is sufficient room for growth on each of the file systems being used). When the database manager attempts to increase the size of the table space, it now attempts to extend the containers in this new stripe set instead of attempting to extend the older containers.

Monitoring

Information about automatic resizing for DMS table spaces is displayed as part of the table space monitor snapshot output. The increase size and maximum size values are included in the output, as shown in the following sample:

```

Auto-resize enabled           = Yes or No
Current tablespace size (bytes) = ###
Maximum tablespace size (bytes) = ### or NONE
Increase size (bytes)        = ###
Increase size (percent)      = ###
Time of last successful resize = DD/MM/YYYY HH:MM:SS.SSSSSS
Last resize attempt failed    = Yes or No

```

Automatic storage databases

The database manager creates all databases as “automatic storage” databases, by default. To create a database that is not an “automatic storage” database, specify **AUTOMATIC STORAGE NO** when issuing the **CREATE DATABASE** command.

Databases that are enabled for automatic storage have a set of one or more storage paths associated with them. A table space can be defined as *managed by automatic storage* and its containers assigned and allocated by the database manager based on those storage paths.

You can enable a database for automatic storage only when you create it; similarly, you cannot disable automatic storage for a database that was originally designed to use it.

All databases are created as automatic storage databases by default. To create a database that is not an automatic storage database, specify **AUTOMATIC STORAGE NO** when issuing the **CREATE DATABASE** command.

Examples of disabling automatic storage:

```
CREATE DATABASE ASNODB1 AUTOMATIC STORAGE NO
CREATE DATABASE ASNODB2 AUTOMATIC STORAGE NO ON X:
```

Examples of automatic storage being enabled either explicitly or implicitly:

```
CREATE DATABASE DB1
CREATE DATABASE DB2 AUTOMATIC STORAGE YES ON X:
CREATE DATABASE DB3 ON /data/path1, /data/path2
CREATE DATABASE DB4 ON D:\StoragePath DBPATH ON C:
```

Based on the syntax used, the database manager extracts the following two pieces of information that pertain to storage locations:

- The database path (where the database manager stores various control files for the database):
 - If you specify **DBPATH ON**, this indicates the database path.
 - If you do not specify **DBPATH ON**, the first path listed in **ON** indicates the database path (and the storage path).
 - If you specify neither **DBPATH ON** nor **ON**, the **dftdbpath** database manager configuration parameter is used to determine the database path.
- The storage paths (where the database manager creates automatic storage table space containers):
 - If you specify **ON**, all of the listed paths are storage paths.
 - If you do not specify **ON**, there is a single storage path that is set to the value of the **dftdbpath** database manager configuration parameter.

For the examples shown previously, the following table summarizes the database and storage paths used:

Table 4. Automatic storage database and storage paths

CREATE DATABASE command	Database path	Storage paths
CREATE DATABASE DB1 AUTOMATIC STORAGE YES	Value of the dftdbpath configuration parameter	Value of the dftdbpath configuration parameter
CREATE DATABASE DB2 AUTOMATIC STORAGE YES ON X:	X:	X:

Table 4. Automatic storage database and storage paths (continued)

CREATE DATABASE command	Database path	Storage paths
CREATE DATABASE DB3 ON /data/path1, /data/path2	/data/path1	/data/path1, /data/path2
CREATE DATABASE DB4 ON D:\StoragePath DBPATH ON C:	C:	D:\StoragePath

The storage paths provided must exist and be accessible. In a partitioned database environment, the same storage paths are used on each database partition. You cannot specify a unique set of storage paths for a particular database partition unless you use database partition expressions as part of the storage path names. Doing this allows the database partition number to be reflected in the storage paths such that the resulting path names are different on each database partition.

Use the argument \$N (that is, \$N preceded by a blank) to indicate a database partition expression. You can use a database partition expression anywhere in the storage path, and you can specify multiple database partition expressions. Terminate the database partition expression with a space character; whatever follows the space is appended to the storage path after the database partition expression is evaluated. If there is no space character in the storage path after the database partition expression, it is assumed that the rest of the string is part of the expression. The following table lists the only valid forms of the \$N argument. Operators are evaluated from left to right, and % represents the modulus operator. The database partition number in the examples is 10.

Table 5. Database partition expressions

Syntax	Example	Value
[blank]\$N	" \$N"	10
[blank]\$N+[number]	" \$N+100"	110
[blank]\$N%[number]	" \$N%5"	0
[blank]\$N+[number]%[number]	" \$N+1%5"	1
[blank]\$N%[number]+[number]	" \$N%4+2"	4

The following is an example of using database partition expressions:

```
CREATE DATABASE TESTDB ON "/path1ForNode $N",
"/path2ForNode $N" DBPATH ON "/dbpathForNodes"
```

The following is an example of a database partition expression embedded in the middle of a path:

```
CREATE DATABASE TESTDB ON "/path1ForNode $N",
"/path2ForNode $N suffix" DBPATH ON "/dbpathForNodes"
```

Note: Database partition expressions are not valid in database paths, whether you specify them explicitly in **DBPATH ON** or implicitly by using a database partition expression in the first storage path.

When free space is calculated for a storage path for a given database partition, the database manager checks for the existence of the following directories or mount points within the storage path and uses the first one that it finds:

```
storage path/instance name/NODE####/database name
storage path/instance name/NODE####
storage path/instance name
storage path
```

where:

storage path

Is a storage path associated with the database

instance name

Is the instance under which the database resides

NODE####

Is the database partition number (for example, NODE0000 or NODE0001)

database name

Is the name of the database

File systems can be mounted at a point beneath the storage path, and the database manager recognizes that the actual amount of free space available for table space containers might not be the same amount that is associated with the storage path directory itself.

Consider the example where two logical database partitions exist on one physical computer and there is a single storage path: /db2data. Each database partition can use this storage path, but you might want to isolate the data from each database partition by creating for each one a separate file system. The file system is mounted at /db2data/*instance*/NODE####. When creating containers on the storage path and determining free space, the database manager does not retrieve free space information for /db2data but instead retrieves it for the corresponding /db2data/*instance*/NODE#### directory.

There are three default table spaces created whenever you create a database. If you do not provide explicit table space definitions as part of the CREATE DATABASE command, the table spaces are created as automatic storage table spaces.

After you create a database, you can add new storage paths to it by using the ADD STORAGE clause of the ALTER DATABASE statement, as shown in the following example:

```
ALTER DATABASE ADD STORAGE ON '/data/path3', '/data/path4'
```

Adding automatic storage paths to databases enabled for automatic storage

Using the ALTER DATABASE statement, you can add an automatic storage path to a database that is enabled for automatic storage. You can enable a database for automatic storage only when you create it.

When you add a storage path for a multi-partition database environment, the storage path must exist on each database partition. If the specified path does not exist on every database partition, the statement is rolled back.

To add a storage path to an existing database, issue the following ALTER DATABASE statement:

```
ALTER DATABASE PATH pathname
```

Automatic storage restrictions

When deciding whether to create a database using automatic storage, there are some restrictions to consider.

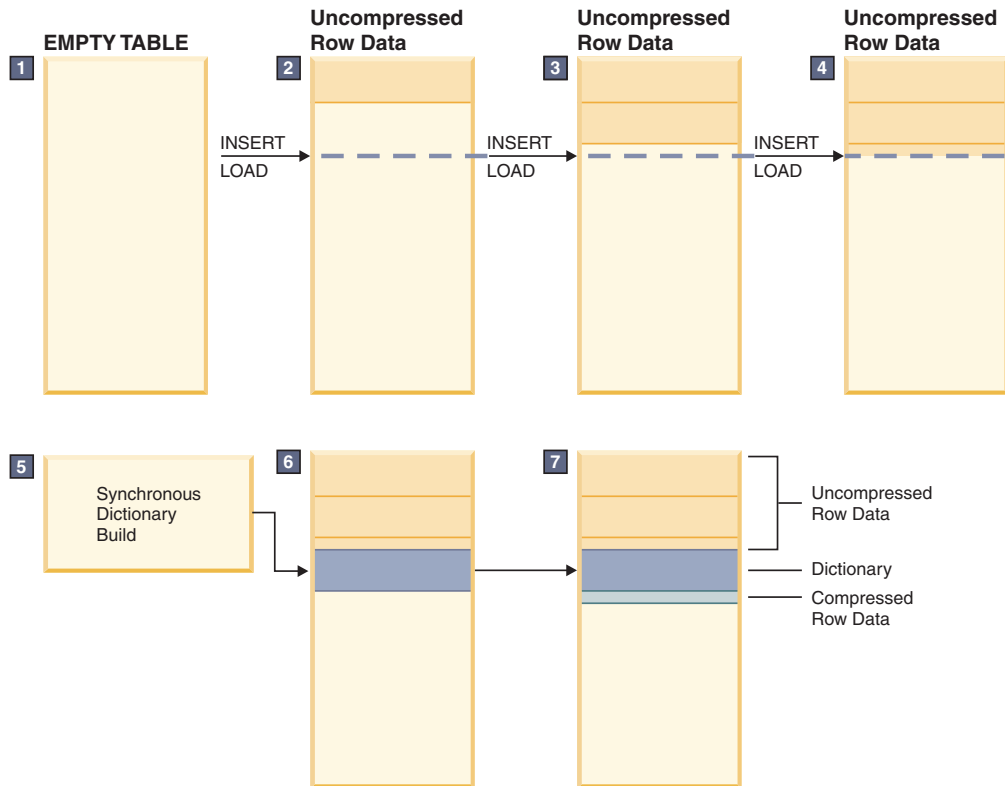
- You cannot disable or enable automatic storage for a database after you created it.
- Storage paths must be absolute path names. They can be paths or drive letters on the Windows operating system. The database path must be a drive letter. The maximum path length is 175 characters.
- For partitioned databases, you must use the same set of storage paths on each database partition (unless you use database partition expressions).
- Database partition expressions are not valid in database paths, whether you specify them explicitly by using the **DBPATH ON** option of the **CREATE DATABASE** command, or implicitly by using a database partition expression in the first storage path.

Automatic (compression) dictionary creation (ADC)

A compression dictionary is used to compress data moved into a table to free up space so that more data can be added in the table. A compression dictionary is automatically created and inserted or appended to a table during a data population operation (such as a load or an insert operation) if certain conditions are met.

Automatic (compression) dictionary creation (ADC) occurs for a table if you defined the **COMPRESS** attribute on the table, if a compression dictionary does not already exist within that physical table or partition, and if the table has sufficient data. Data subsequently moved into the table is compressed using the compression dictionary (if the table **COMPRESS** attribute remains enabled).

The following diagram shows the process by which the compression dictionary is automatically created:



1. A compression dictionary is not created because the table is empty.
2. Data is inserted into the table using insert or load operations and remains uncompressed.
3. As more data is inserted or loaded into the table, it remains uncompressed.
4. After a threshold is reached, dictionary creation is triggered automatically if the COMPRESS attribute is set to YES.
5. The dictionary is created.
6. The dictionary is appended to the table.
7. From this point forward, the data is compressed.

The following table shows the compression dictionary creation differences by release:

Table 6. Compression dictionary creation differences by release

Commands and attributes	Version 9.1	Version 9.5
LOAD REPLACE command with the RESETDICTIONARY option	Not applicable.	When you set the table COMPRESS attribute to YES, any existing compression dictionary is removed and a new one is generated if at least one row of data is loaded or inserted into the table.
CREATE or ALTER TABLE statement with the COMPRESS attribute set to YES	Dictionary creation was not automatic. To compress table data, you had to explicitly create a compression dictionary using a table reorganization process.	Setting the table COMPRESS attribute to YES makes the table eligible for ADC if at least one row of data is loaded or inserted into the table.

Table 6. Compression dictionary creation differences by release (continued)

Commands and attributes	Version 9.1	Version 9.5
INSERT, LOAD INSERT, IMPORT INSERT, or REDISTRIBUTE command	Not applicable.	When you set the table COMPRESS attribute to YES, a table that does not already have a compression dictionary is subject to ADC, if the table has enough data (meaning that the threshold is passed). Note: The REDISTRIBUTE command triggers ADC only on newly added database partitions.
REORG TABLE command with the KEEPDICTIONARY option	If you set the table COMPRESS attribute to YES and a compression dictionary did not yet exist in the table, an attempt was made to build, insert, or append a compression dictionary into the table independent of the volume of data in the table.	A dictionary is inserted into a table only if the table size is equal to the ADC table size threshold and if enough data exists in the table when it passes the threshold.

Data row compression

The purpose of data row compression is to achieve disk storage space savings, and it can also lead to disk I/O savings. Also, more data can be cached in the buffer pool, thereby increasing bufferpool hit ratios. Data row compression uses a static dictionary-based compression algorithm to compress data by row.

Compressing data at the row level allows repeating patterns that span multiple column values within a row to be replaced with shorter symbol strings.

Note: There is an associated cost in the form of extra CPU cycles needed to compress and decompress data. The storage savings and performance impact of data row compression are tied to the characteristics of the data within the database, the layout and tuning of the database, and the application workload. Only the data on a data page or in log records is compressed.

To compress table data, a compression dictionary must exist for the table, you must set the COMPRESS attribute of the CREATE TABLE or ALTER TABLE statement to YES, and there needs to be sufficient data in the table. If these compression conditions exist for the table, then when you issue an INSERT statement or a LOAD INSERT, IMPORT INSERT, or REDISTRIBUTE command, data added to the table is compressed.

In Version 9.5, data row compression is automatically enabled if a table has the COMPRESS attribute set to YES and after the data compression dictionary has been created. If you created or altered a table with a COMPRESS attribute set to YES, no manual operation or database request is required on your part: that is, you do not need to perform an explicit classic (offline) table reorganization to create the data compression dictionary.

Note: If you set the COMPRESS attribute to YES and a compression dictionary exists, compression applies to any insert row operation, including an insert

through an import or a load operation. Compression is enabled for an entire table; however, each row is compressed individually. Therefore, a table could contain both compressed and non-compressed rows at the same time.

To explicitly build a compression dictionary (and subsequently compress a table), perform a classic (offline) table reorganization. All of the data rows that exist in a table participate in the building of the compression dictionary. The dictionary is stored with the table data rows in the data object portions of the table.

To decompress a table, set the table COMPRESS attribute to NO, and then perform a classic (offline) table reorganization.

Restrictions

- Data row compression is not applicable to index, long, LOB, and XML data objects.
- Row compression is not compatible with table data replication support.
- You can generate row compression statistics by using the RUNSTATS command. These are stored in the system catalog table SYSCAT.TABLES. A compression estimation option, which estimates the effectiveness of row compression for a table, is available with the INSPECT utility. The query optimizer includes decompression cost in its costing model.
- Depending upon update activity and the positioning of update changes within a data row, there might be an increase in log space consumption.
- If a row is increasing in size, the new version of the row might not fit on the current data page. In that case, the new image of the row is stored on an overflow page. To minimize the creation of such pointer-overflow records, you can add more free space within a data page. For example, if 5% free space was used without compression, allocate 10% free space with compression. This recommendation is especially important for data that is heavily updated.

Automatic statistics collection

The DB2 optimizer uses catalog statistics to determine the most efficient access plan for any given query. Having out-of-date or incomplete statistics for a table or an index might lead the optimizer to select a plan that is not optimal, slowing down query execution. However, deciding which statistics to collect for a given workload is complex, and keeping these statistics up-to-date is time-consuming.

With automatic statistics collection, part of the DB2 automated table maintenance feature, you can let the DB2 database manager determine whether database statistics need to be updated. Automatic statistics collection can occur at statement compilation time using the real-time statistics (RTS) feature or can be collected by running the RUNSTATS utility in the background. Background statistics collection can be enabled while real-time statistics collection is disabled. Background statistics collection must be enabled in order for real-time statistics collection to be enabled. Background automatic statistics collection is enabled by default when you create a new database. Real-time automated statistics are enabled by the dynamic configuration parameter `auto_stmt_stats`.

Understanding asynchronous and real-time statistics collection

Automatic statistics can be collected synchronously or asynchronously by running the RUNSTATS utility. *Asynchronous* collection occurs in the background. When the real-time statistics feature is enabled, statistics can also be collected *synchronously*, at statement compilation time. When real-time statistics collection is enabled,

statistics can also be fabricated using meta-data maintained by the index and data manager. *Fabrication* means deriving or creating statistics, rather than collecting them as part of normal RUNSTATS activity. For example, the number of rows in a table can be derived from knowing the number of pages in the table, the page size, and the average row width. In some cases, statistics are not actually derived, but are maintained by the index and data manager and can be stored directly in the catalog. For example, the index manager maintains a count of the number of leaf pages and levels in each index.

The query optimizer determines how the statistics should be collected, based on the needs of the query and the amount of table update activity. Table update activity is measured by the number of update, insert, or delete operations.

Real-time statistics collection is determined by the needs of an SQL statement before it is optimized. This provides more timely statistics collection and more accurate statistics. Accurate statistics can result in better query execution plans and improved performance. When real-time statistics collection is not enabled, asynchronous statistics collection occurs at two-hour intervals. This might not be frequent enough to provide accurate statistics for some applications.

When real-time statistics collection is enabled, asynchronous statistics collection checking still occurs at two-hour intervals. Real-time statistics collection also initiates asynchronous collection requests when:

- Table activity is not high enough to require synchronous collection, but is high enough to require asynchronous collection
- Synchronous statistics collection used sampling because the table was large
- Synchronous statistics were fabricated
- Synchronous statistics collection failed because the collection time was exceeded

At most, two asynchronous requests can be processed at the same time, but only for different tables. One request will have been initiated by real-time statistics, and the other will have been initiated by asynchronous statistics collection checking.

The performance impact of automatic statistics collection is minimized in several ways:

- Asynchronous statistics collection is performed using a throttled RUNSTATS utility. Throttling controls the amount of resources consumed by the RUNSTATS utility based on current database activity: as database activity increases, the RUNSTATS utility runs more slowly, reducing its resource demands.
- Synchronous statistics collection is limited to five seconds per query. This value can be controlled by the RTS optimization guideline. If synchronous collection exceeds the time limit, an asynchronous collection request is submitted.
- Synchronous statistics collection does not store the statistics in the system catalog. Instead, the statistics are stored in a statistics cache and are later stored in the system catalog by an asynchronous operation. This avoids the overhead and possible lock contention involved with updating the system catalog. Statistics in the statistics cache are available for subsequent SQL compilation requests.
- Only one synchronous statistics collection operation will occur per table. Other agents requiring synchronous statistics collection will fabricate statistics, if possible, and continue with statement compilation. This behavior is also enforced in a partitioned database environment, where agents on different database partitions might require synchronous statistics.

- By default, statistics collected for synchronous and asynchronous operations are basic table statistics with distribution information and detailed index statistics using sampling. (The RUNSTATS command is issued specifying the WITH DISTRIBUTION and SAMPLED DETAILED INDEXES ALL options.) You can customize the type of statistics collected by enabling statistics profiling, which uses information about previous database activity to determine which statistics are required by the database workload. You can also customize the type of statistics collected for a particular table by creating your own statistics profile for that table.
- Only tables with missing statistics or high levels of activity (as measured by the number of update, insert, or delete operations) are considered for statistics collection. Even if a table meets the statistics collection criteria, synchronous statistics are not collected unless query optimization requires them. In some cases, the query optimizer can choose an access plan without statistics.
- For asynchronous statistics collection checking, large tables (with more than 4000 pages) are sampled to determine whether high table activity has changed the statistics. Statistics for such large tables are collected only if warranted.
- For asynchronous statistics collection, the RUNSTATS utility is automatically scheduled to run during the optimal maintenance window specified in your maintenance policy definition. This policy also specifies the set of tables that are within the scope of automatic statistics collection, further minimizing unnecessary resource consumption.
- Synchronous statistics collection and fabrication do not follow the optimal maintenance window specified in your maintenance policy definition, because synchronous requests must occur immediately and have limited collection time. Synchronous statistics collection and fabrication follow the policy that specifies the set of tables that are within the scope of automatic statistics collection.
- While automatic statistics collection is being performed, the affected tables are still available for regular database activity (update, insert, or delete operations).
- For asynchronous statistics collection, the SYSPROC.NNSTAT stored procedure is run using the catalog-based collection method to refresh nickname statistics automatically. Real-time statistics (synchronous or fabricated) are not collected for nicknames.

Real-time synchronous statistics collection is performed for regular tables, materialized query tables (MQTs), and declared global temporary tables (DGTTs). Asynchronous statistics are not collected for DGTTs. This means that real-time statistics processing will not initiate asynchronous requests for DGTTs.

Automatic statistics collection (synchronous or asynchronous) does not occur for:

- Statistical views
- Tables that are marked VOLATILE (tables with the VOLATILE field set in the SYSCAT.TABLES catalog view)
- Tables that have had their statistics manually updated, by issuing UPDATE statements directly against SYSSTAT catalog views

When you modify the statistics manually for tables, the database manager assumes that you are now responsible for maintaining their statistics. To allow the database manager to reconsider and maintain statistics for a table that has had its statistics manually updated, collect statistics using the RUNSTATS command or specify to collect statistics when using the LOAD command. Tables created prior to Version 9.5 that had their statistics updated manually prior to migration are not affected and their statistics are automatically maintained by the database manager until you manually update their statistics.

Statistics fabrication does not occur for:

- Statistical views
- Tables that have had their statistics manually updated, by issuing UPDATE statements directly against SYSSTAT catalog views. Note that if real-time statistics collection is not enabled, some statistics fabrication will still occur for tables that have had their statistics manually updated.

In a partitioned database environment, statistics are collected on a single database partition and then extrapolated. The database manager always collects statistics (both synchronous and asynchronous) on the first database partition of the database partition group.

No real-time statistics collection activity will occur until at least five minutes after database activation.

When real-time statistics are enabled, you should schedule a defined maintenance window; the maintenance window is undefined by default. If there is no defined maintenance window, only synchronous statistics collection will occur. In this case, the collected statistics are only in-memory, and are typically collected using sampling (except in the case of small tables).

Real-time statistics processing occurs for both static and dynamic SQL.

A table that has been truncated using the IMPORT command is automatically recognized as having stale statistics.

Automatic statistics collection, both synchronous and asynchronous, invalidates cached dynamic statements that reference tables for which statistics have been collected. This is done so that cached dynamic statements can be re-optimized with the latest statistics.

Real-time statistics and explain processing

There is no real-time processing for a query which is just explained (not executed) using the explain facility. The following table summarizes the behavior for the different values of the CURRENT EXPLAIN MODE special register.

Table 7. Real-time statistics collection as a function of the value of the CURRENT EXPLAIN MODE special register

CURRENT EXPLAIN MODE value	Real-time statistics collection considered
YES	Yes
EXPLAIN	No
NO	Yes
REOPT	Yes
RECOMMEND INDEXES	No
EVALUATE INDEXES	No

Automatic statistics collection and the statistics cache

A statistics cache was introduced in DB2 Version 9.5 to make synchronously-collected statistics available to all queries. This cache is part of the catalog cache. In a partitioned database environment, this cache resides only on the catalog database partition. The catalog cache can store multiple entries for the same SYSTABLES

object, which increases the size of the catalog cache on all database partitions. Consider increasing the value of the `catalogcache_sz` database configuration parameter when real-time statistics collection is enabled.

Starting with DB2 Version 9, you can use the Configuration Advisor to determine the initial configuration for new databases. The Configuration Advisor recommends that the `auto_stmt_stats` database configuration parameter be set to ON.

Automatic statistics collection and statistical profiles

Synchronous and asynchronous statistics are collected according to a statistical profile that is in effect for a table, with the following exceptions:

- To minimize the overhead of synchronous statistics collection, the database manager might collect statistics using sampling. In this case, the sampling rate and method might be different from those specified in the statistical profile.
- Synchronous statistics collection might choose to fabricate statistics, but it might not be possible to fabricate all statistics specified in the statistical profile. For example, column statistics such as COLCARD, HIGH2KEY, and LOW2KEY cannot be fabricated unless the column is leading in some index.

If synchronous statistics collection cannot collect all statistics specified in the statistical profile, an asynchronous collection request is submitted.

Although real-time statistics collection is designed to minimize statistics collection overhead, try it in a test environment first to ensure that there is no negative performance impact. This might be the case in some online transaction processing (OLTP) scenarios, especially if there is an upper boundary for how long a query can run.

Enabling automatic statistics collection

Having accurate and complete database statistics is critical to efficient data access and optimal workload performance. Use the automatic statistics collection feature of the automated table maintenance functionality to update and maintain relevant database statistics. You can optionally enhance this functionality in environments where a single database partition operates on a single processor (serial environment) by collecting query data and generating statistics profiles that help DB2 automatically collect the exact set of statistics that is required by your workload. This option is not available in MPP environments, certain federated environments, or environments in which intra-partition parallelism is enabled.

To enable automatic statistics collection:

1. Configure your database instance by using the Configure Automatic Maintenance wizard or the command line:
 - To use the Configure Automatic Maintenance wizard:
 - a. Open the wizard either from the Control Center by right-clicking a database object or from the Health Center by right-clicking a database instance.
 - b. Select **Configure Automatic Maintenance** from the pop-up window. Within this wizard, you can enable automatic statistics collection, specify the tables from which you want to automatically collect statistics, and specify a maintenance window for the execution of the RUNSTATS utility.
 - To use the command line, set each of the following configuration parameters to ON:

- `AUTO_MAINT`
 - `AUTO_TBL_MAINT`
 - `AUTO_RUNSTATS`
2. Optional: To enable the automatic statistics profile generation, set the following two configuration parameters to ON:
 - `AUTO_STATS_PROF`
 - `AUTO_PROF_UPD`
 3. Optional: To enable real-time statistics gathering, set the `AUTO_STMT_STATS` configuration parameter to ON. If this configuration parameter is set to ON, table statistics are automatically compiled at statement compilation time, whenever they are needed to optimize a query.

Configuration Advisor

You can use the Configuration Advisor to obtain recommendations for the initial values of the buffer pool size, database configuration parameters, and database manager configuration parameters.

To use the Configuration Advisor, specify the `AUTOCONFIGURE` command for an existing database, or specify `AUTOCONFIGURE` as an option of the `CREATE DATABASE` command. To configure your database, you must have `SYSADM`, `SYSCTRL`, or `SYSMAINT` authority.

You can display the recommended values or apply them by using the `APPLY` option of the `CREATE DATABASE` command. The recommendations are based on input that you provide and system information that the advisor gathers.

The values suggested by the Configuration Advisor are relevant for only one database per instance. If you want to use this advisor on more than one database, each database must belong to a separate instance.

Tuning configuration parameters using the Configuration Advisor

The Configuration Advisor helps you to tune performance and to balance memory requirements for a single database per instance by suggesting which configuration parameters to modify and suggesting values for them. The Configuration Advisor is automatically run when you create a database.

To disable this feature or to explicitly enable it, use the `db2set` command before creating a database, as follows:

```
db2set DB2_ENABLE_AUTOCONFIG_DEFAULT=NO
db2set DB2_ENABLE_AUTOCONFIG_DEFAULT=YES
```

To define values for several of the configuration parameters and to determine the scope of the application of those parameters, use the `AUTOCONFIGURE` command, specifying one of the following options:

- `NONE`, meaning that none of the values are applied
- `DB ONLY`, meaning that only database configuration and buffer pool values are applied
- `DB AND DBM`, meaning that all parameters and their values are applied

Note: Even if you automatically enabled the Configuration Advisor when you ran the `CREATE DATABASE` command, you can still specify `AUTOCONFIGURE`

command options. If you did not enable the Configuration Advisor when you ran the CREATE DATABASE command, you can run the Configuration Advisor manually afterwards.

Generating database configuration recommendations

The Configuration Advisor is automatically run when you create a database. You can also run the Configuration Advisor by specifying the AUTOCONFIGURE command in the command line processor (CLP) or by calling the db2AutoConfig API.

To request configuration recommendations using the CLP, enter the following command:

```
AUTOCONFIGURE
  USING input_keyword param_value
  APPLY value
```

Following is an example of an AUTOCONFIGURE command that requests configuration recommendations based on input about how the database is used but specifies that the recommendations not be applied:

```
DB2 AUTOCONFIGURE USING
  MEM_PERCENT 60
  WORKLOAD_TYPE MIXED
  NUM_STMTS 500
  ADMIN_PRIORITY BOTH
  IS_POPULATED YES
  NUM_LOCAL_APPS 0
  NUM_REMOTE_APPS 20
  ISOLATION RR
  BP_RESIZEABLE YES
APPLY NONE
```

Example: Requesting configuration recommendations using the Configuration Advisor

This scenario demonstrates to run the Configuration Advisor from the command line to generate recommendations and shows the output that the Configuration Advisor produces.

To run the Configuration Advisor:

1. Connect to the PERSONL database by specifying the following command from the command line:

```
DB2 CONNECT TO PERSONL
```

2. Issue the AUTOCONFIGURE command from the CLP, specifying how the database is used. As shown in the following example, set a value of NONE for the **APPLY** option to indicate that you want to view the configuration recommendations but not apply them:

```
DB2 AUTOCONFIGURE USING
  MEM_PERCENT 60
  WORKLOAD_TYPE MIXED
  NUM_STMTS 500
  ADMIN_PRIORITY BOTH
  IS_POPULATED YES
  NUM_LOCAL_APPS 0
  NUM_REMOTE_APPS 20
  ISOLATION RR
  BP_RESIZEABLE YES
APPLY NONE
```

If you are unsure about the value of a parameter for the command, you can omit it, and the default will be used. You can pass up to 10 parameters without values: MEM_PERCENT, WORKLOAD_TYPE, and so on, as shown in the previous example.

The recommendations generated by the AUTOCONFIGURE command are displayed on the screen in table format:

Table 8. Configuration Advisor sample output: Part 1

Former and Applied Values for Database Manager Configuration			
Description	Parameter	Current Value	Recommended Value
Application support layer heap size (4KB)	(ASLHEAPSZ) = 15		15
No. of int. communication buffers(4KB)	(FCM_NUM_BUFFERS) = AUTOMATIC		AUTOMATIC
Enable intra-partition parallelism	(INTRA_PARALLEL) = NO		NO
Maximum query degree of parallelism	(MAX_QUERYDEGREE) = ANY		1
Agent pool size	(NUM_POOLAGENTS) = 100(calculated)		200
Initial number of agents in pool	(NUM_INITAGENTS) = 0		0
Max requester I/O block size (bytes)	(RQRIOBLK) = 32767		32767
Sort heap threshold (4KB)	(SHEAPTHRES) = 0		0

Table 9. Configuration Advisor sample output (continued)

Former and Applied Values for Database Configuration			
Description	Parameter	Current Value	Recommended Value
Default application heap (4KB)	(APPLHEAPSZ) = 256		256
Catalog cache size (4KB)	(CATALOGCACHE_SZ) = (MAXAPPLS*4)		260
Changed pages threshold	(CHNGPGS_THRESH) = 60		80
Database heap (4KB)	(DBHEAP) = 1200		2791
Degree of parallelism	(DFT_DEGREE) = 1		1
Default tablespace extent size (pages)	(DFT_EXTENT_SZ) = 32		32
Default prefetch size (pages)	(DFT_PREFETCH_SZ) = AUTOMATIC		AUTOMATIC
Default query optimization class	(DFT_QUERYOPT) = 5		5
Max storage for lock list (4KB)	(LOCKLIST) = 100		AUTOMATIC
Log buffer size (4KB)	(LOGBUFSZ) = 8		99
Log file size (4KB)	(LOGFILSIZ) = 1000		1024
Number of primary log files	(LOGPRIMARY) = 3		8
Number of secondary log files	(LOGSECOND) = 2		3
Max number of active applications	(MAXAPPLS) = AUTOMATIC		AUTOMATIC
Percent. of lock lists per application	(MAXLOCKS) = 10		AUTOMATIC
Group commit count	(MINCOMMIT) = 1		1
Number of asynchronous page cleaners	(NUM_IOCLEANERS) = 1		1
Number of I/O servers	(NUM_IOSERVERS) = 3		4
Package cache size (4KB)	(PCKCACHESZ) = (MAXAPPLS*8)		1533
Percent log file reclaimed before soft chckpt	(SOFTMAX) = 100		320
Sort list heap (4KB)	(SORTHEAP) = 256		AUTOMATIC
statement heap (4KB)	(STMTHEAP) = 4096	4096	
Statistics heap size (4KB)	(STAT_HEAP_SZ) = 4384		4384
Utilities heap size (4KB)	(UTIL_HEAP_SZ) = 5000		113661
Self tuning memory	(SELF_TUNING_MEM) = ON		ON
Automatic runstats	(AUTO_RUNSTATS) = ON		ON
Sort heap thres for shared sorts (4KB)	(SHEAPTHRES_SHR) = 5000		AUTOMATIC

Table 10. Configuration Advisor sample output (continued)

Former and Applied Values for Bufferpool(s)			
Description	Parameter	Current Value	Recommended Value
IBMDEFAULTBP	Bufferpool size =	-2	340985

DB210203I AUTOCONFIGURE completed successfully. Database manager or database configuration values may have been changed. The instance must be restarted before any changes come into effect. You may also want to rebind your packages after the new configuration parameters take effect so that the new values will be used.

If you agree with all of the recommendations, either reissue the AUTOCONFIGURE command but specify that you want the recommended values to be applied by using the APPLY option, or update individual configuration parameters using the UPDATE DATABASE MANAGER CONFIGURATION command and the UPDATE DATABASE CONFIGURATION command.

Utility throttling

Utility throttling regulates the performance impact of maintenance utilities so that they can run concurrently during production periods. Although the impact policy, a setting that allows utilities to run in throttled mode, is defined by default, you must set the impact priority, a setting that each cleaner has indicating its throttling priority, when you run a utility if you want to throttle it.

The throttling system ensures that the throttled utilities are run as frequently as possible without violating the impact policy. You can throttle statistics collection, backup operations, rebalancing operations, and asynchronous index cleanups.

You define the impact policy by setting the **util_impact_lim** configuration parameter.

Cleaners are integrated with the utility throttling facility. By default, each (index) cleaner has a utility impact priority of 50 (acceptable values are between 1 and 100, with 0 indicating no throttling). You can change the priority by using the SET UTIL_IMPACT_PRIORITY command or the db2UtilityControl API.

Asynchronous index cleanup

Asynchronous index cleanup (AIC) is the deferred cleanup of indexes following operations that invalidate index entries. Depending on the type of index, the entries can be row identifiers (RIDs) or block identifiers (BIDs). Either way, these entries are removed by the index cleaners, which operate asynchronously in the background.

AIC accelerates the detach of a data partition from a partitioned table. If the partitioned table contains one or more non-partitioned indexes, AIC is initiated. In this case, AIC removes all non-partitioned index entries that refer to the detached data partition and any pseudo-deleted entries. After all of the indexes have been cleaned, the identifier associated with the detached data partition is removed from the system catalog.

Note: If the partitioned table has dependent materialized query tables (MQTs) defined, AIC is not initiated until after a SET INTEGRITY operation is performed.

While AIC is in progress, normal table access is maintained. Queries accessing the indexes ignore any non-valid entries that have not yet been cleaned.

In most cases, one cleaner is started for each non-partitioned index associated with the partitioned table. An internal task distribution daemon is responsible for distributing the AIC tasks to the appropriate database partitions and assigning database agents.

Both the distribution daemon and cleaner agents are internal system applications. They appear in the LIST APPLICATION output with the application names **db2taskd** and **db2aic**, respectively. To prevent accidental disruption, system applications cannot be forced. The distribution daemon remains online as long as the database is active. The cleaners remain active until the cleaning is complete. If the database is deactivated while cleaning is in progress, AIC resumes when you reactivate the database.

Performance

AIC incurs minimal performance impact.

An instantaneous row lock test is required to determine whether a pseudo-deleted entry is committed. However, because the lock is never acquired, concurrency is unaffected.

Each cleaner acquires a minimal table space lock (IX) and table lock (IS). These locks are released if the cleaner determines that other applications are waiting for the locks. If this occurs, the cleaner suspends processing for 5 minutes.

Cleaners are integrated with the utility throttling facility. By default, each cleaner has a utility impact priority of 50. You can change the priority by using the SET UTIL_IMPACT_PRIORITY command or the db2UtilityControl API.

Monitoring

You can monitor AIC with the LIST UTILITIES command. Each index cleaner appears in the monitor as a separate utility.

The following example demonstrates AIC activity in the WSDB database at the current database partition using the Command Line Processor (CLP) interface:

```
$ db2 list utilities show detail
```

```
ID = 2
Type = ASYNCHRONOUS INDEX CLEANUP
Database Name = WSDB
Partition Number = 0
Description = Table: USER1.SALES, Index: USER1.I2
Start Time = 12/15/2005 11:15:01.967939
State = Executing
Invocation Type = Automatic
Throttling:
  Priority = 50
Progress Monitoring:
  Total Work = 5 pages
  Completed Work = 0 pages
  Start Time = 12/15/2005 11:15:01.979033

ID = 1
Type = ASYNCHRONOUS INDEX CLEANUP
Database Name = WSDB
```

```

Partition Number           = 0
Description                 = Table: USER1.SALES, Index: USER1.I1
Start Time                 = 12/15/2005 11:15:01.978554
State                      = Executing
Invocation Type           = Automatic
Throttling:
  Priority                  = 50
Progress Monitoring:
  Total Work               = 5 pages
  Completed Work          = 0 pages
  Start Time              = 12/15/2005 11:15:01.980524

```

In this case, there are two cleaners operating on the USERS1.SALES table. One cleaner is processing index I1; the other is processing index I2. The progress monitoring section shows the estimated total number of index pages that need cleaning and the current number of clean index pages.

The State field indicates the current state of the cleaner. Normally, the state is Executing. The cleaner might be in the Waiting state if the cleaner is waiting to be assigned to an available database agent or if the cleaner is temporarily suspended due to lock contention.

Note: Different tasks on different database partitions can have the same utility ID because each database partition assigns IDs to tasks on that database partition only.

Asynchronous index cleanup for MDC tables

You can enhance the performance of a rollout deletion, an efficient method of deleting qualifying blocks of data from multidimensional clustering (MDC) tables, by using asynchronous index cleanup (AIC). AIC is the deferred cleanup of indexes following operations that invalidate index entries.

During a standard rollout delete, indexes are cleaned up synchronously with the delete. For tables that contain many record ID (RID) indexes, a significant portion of the delete time is spent removing index keys that reference the table rows being deleted. You can speed up the rollout by specifying that these indexes are to be cleaned up after the delete has been committed.

To take advantage of AIC for MDC tables, you need to explicitly enable the *deferred index cleanup rollout* mechanism. There are two methods of specifying a deferred rollout: setting the **DB2_MDC_ROLLOUT** registry variable to DEFER and issuing the SET CURRENT MDC ROLLOUT MODE statement. During a deferred index cleanup rollout, blocks are marked as rolled out without an update to the RID indexes until after the transaction commits. Block identifier (BID) indexes are still cleaned up during the delete because they don't require row-level processing.

Rollout AIC is invoked when a rollout delete is committed or, if the database was shut down, when the table is first accessed following a restart of the database. While AIC is in progress, any queries against the indexes work, including those accessing the index being cleaned up.

There one coordinating cleaner per MDC table. The index cleanup for multiple rollouts is consolidated in the cleaner. The cleaner spawns a cleanup agent for each RID index, and the cleanup agents update the RID indexes in parallel. Cleaners are also integrated with the utility throttling facility. By default, each cleaner has a utility impact priority of 50 (acceptable values are between 1 and 100, with 0 indicating no throttling). You can change the priority by using the SET UTIL_IMPACT_PRIORITY command or the db2UtilityControl API.

Monitoring

Because the rolled-out blocks on an MDC table are not reusable until after the cleanup is complete, it is useful to monitor the progress of a deferred index cleanup rollout. Use the LIST UTILITIES monitor command to display a utility monitor entry for each index being cleaned up. You can also query the number of blocks in the table currently being cleaned up through deferred index cleanup rollout (BLOCKS_PENDING_CLEANUP) by using the SYSPROC.ADMIN_GET_TAB_INFO_V95 table function. To query the number of MDC table blocks pending cleanup at the database level, use the GET SNAPSHOT command.

In the following sample output for the LIST UTILITIES, progress is indicated by the number of pages in each index that have been cleaned up. Each phase listed in the output represents one of the RID indexes being cleaned for the table.

```
db2 LIST UTILITIES SHOW DETAILS output.
ID = 2
Type = MDC ROLLOUT INDEX CLEANUP
Database Name = WSDB
Partition Number = 0
Description = TABLE.<schema_name>.<table_name>
Start Time = 06/12/2006 08:56:33.390158
State = Executing
Invocation Type = Automatic
Throttling:
  Priority = 50
Progress Monitoring:
  Estimated Percentage Complete = 83
  Phase Number = 1
    Description = <schema_name>.<index_name>
    Total Work = 13 pages
    Completed Work = 13 pages
    Start Time = 06/12/2006 08:56:33.391566
  Phase Number = 2
    Description = <schema_name>.<index_name>
    Total Work = 13 pages
    Completed Work = 13 pages
    Start Time = 06/12/2006 08:56:33.391577
  Phase Number = 3
    Description = <schema_name>.<index_name>
    Total Work = 9 pages
    Completed Work = 3 pages
    Start Time = 06/12/2006 08:56:33.391587
```

Chapter 4. Instances

An *instance* is a logical database manager environment where you catalog databases and set configuration parameters. Depending on your needs, you can create more than one instance on the same physical server providing a unique database server environment for each instance.

Note: For non-root installations on Linux[®] and UNIX operating systems, a single instance is created during the installation of your DB2 product. Additional instances cannot be created.

You can use multiple instances to do the following:

- Use one instance for a development environment and another instance for a production environment.
- Tune an instance for a particular environment.
- Restrict access to sensitive information.
- Control the assignment of SYSADM, SYSCTRL, and SYSMAINT authority for each instance.
- Optimize the database manager configuration for each instance.
- Limit the impact of an instance failure. In the event of an instance failure, only one instance is affected. Other instances can continue to function normally.

Multiple instances will require:

- Additional system resources (virtual memory and disk space) for each instance.
- More administration because of the additional instances to manage.

The instance directory stores all information that pertains to a database instance. You cannot change the location of the instance directory once it is created. The directory contains:

- The database manager configuration file
- The system database directory
- The node directory
- The node configuration file (db2nodes.cfg)
- Any other files that contain debugging information, such as the exception or register dump or the call stack for the DB2 database processes.

Terminology:

Bit-width

The number of bits used to address virtual memory: 32-bit and 64-bit are the most common. This term might be used to refer to the bit-width of an instance, application code, external routine code. 32-bit application means the same things as 32-bit width application.

32-bit DB2 instance

A DB2 instance that contains all 32-bit binaries including 32-bit shared libraries and executables.

64-bit DB2 instance

A DB2 instance that contains 64-bit shared libraries and executables, and

also all 32-bit client application libraries (included for both client and server), and 32-bit external routine support (included only on a server instance).

Designing instances

DB2 databases are created within DB2 instances on the database server. The creation of multiple instances on the same physical server provides a unique database server environment for each instance.

For example, you can maintain a test environment and a production environment on the same machine, or you can create an instance for each application and then fine-tune each instance specifically for the application it will service, or, to protect sensitive data, you can have your payroll database stored in its own instance so that owners of other instances (on the same server) cannot see payroll data.

The installation process creates a default DB2 instance, which is defined by the DB2INSTANCE environment variable. This is the instance that is used for most operations. However, instances can be created (or dropped) after installation.

When determining and designing the instances for your environment, note that each instance controls access to one or more databases. Every database within an instance is assigned a unique name, has its own set of system catalog tables (which are used to keep track of objects that are created within the database), and has its own configuration file. Each database also has its own set of grantable authorities and privileges that govern how users interact with the data and database objects stored in it. Figure 2 shows the hierarchical relationship among systems, instances, and databases.




Figure 2. Hierarchical relationship among DB2 systems, instances, and databases

You also need to be aware of another particular type of instance called the *DB2 administration server* (DAS). The DAS is a special DB2 administration control point used to assist with the administration tasks only on other DB2 servers. A DAS must be running if you want to use the Client Configuration Assistant to discover the remote databases or the graphical tools that come with the DB2 product, for example, the Control Center or the Task Center. There is only one DAS in a DB2 database server, even when there are multiple instances.

Once your instances are created, you can attach to any other instance available (including instances on other systems). Once attached, you can perform maintenance and utility tasks that can only be done at the instance level, for example, create a database, force applications off a database, monitor database activity, or change the contents of the database manager configuration file that is associated with that particular instance.

Default instance

As part of your DB2 installation procedure, you create an initial instance of the database manager called DB2, if there is no other instance called “DB2”. If you have DB2 Version 8 installed, and you upgrade to Version 9.1 or Version 9.5, the default instance is “DB2_01”.

On Linux and UNIX, the initial instance can be called anything you want within the naming rules guidelines. The instance name is used to set up the directory structure.

To support the immediate use of this instance, the following are set during installation:

- The environment variable DB2INSTANCE is set to “DB2”.
- The registry variable DB2INSTDEF is set to “DB2”.

These settings establish “DB2” as the default instance. You can change the instance that is used by default, but first you have to create an additional instance.

Before using the database manager, the database environment for each user must be updated so that it can access an instance and run the DB2 database programs. This applies to all users (including administrative users).

On Linux and UNIX operating systems, sample script files are provided to help you set the database environment. The files are: db2profile for Bourne or Korn shell, and db2cshrc for C shell. These scripts are located in the sqllib subdirectory under the home directory of the instance owner. The instance owner or any user belonging to the instance’s SYSADM group can customize the script for all users of an instance. Use sqllib/userprofile and sqllib/usercshrc to customize a script for each user.

The blank files sqllib/userprofile and sqllib/usercshrc are created during instance creation to allow you to add your own instance environment settings. The db2profile and db2cshrc files are overwritten during an instance update in a DB2 FixPak installation. If you do not want the new environment settings in the db2profile or db2cshrc scripts, you can override them using the corresponding *user script*, which is called at the end of the db2profile or db2cshrc script. During an instance migration (using the db2imigr command), the user scripts are copied over so that your environment modifications will still be in use.

The sample script contains statements to:

- Update a user's PATH by adding the following directories to the existing search path: the bin, adm, and misc subdirectories under the sqllib subdirectory of the instance owner's home directory.
- Set the DB2INSTANCE environment variable to the instance name.

Instance directory

The instance directory stores all information that pertains to a database instance. The location of the instance directory cannot be changed once it is created.

The instance directory contains:

- The database manager configuration file
- The system database directory
- The node directory
- The node configuration file (db2nodes.cfg)
- Other files that contain debugging information, such as the exception or register dump or the call stack for the DB2 processes.

On Linux and UNIX operating systems, the instance directory is located in the INSTHOME/sqllib directory, where INSTHOME is the home directory of the instance owner. The default instance can be called anything you want within the naming rules guidelines.

On Windows operating systems, the instance directory is located under the /sqllib directory where the DB2 database product was installed. The instance name is the same as the name of the service, so it should not conflict. No instance name should be the same as another service name. You must have the correct authorization to create a service.

In a partitioned database environment, the instance directory is shared between all database partition servers belonging to the instance. Therefore, the instance directory must be created on a network share drive that all computers in the instance can access.

db2nodes.cfg

The db2nodes.cfg file is used to define the database partition servers that participate in a DB2 instance. The db2nodes.cfg file is also used to specify the IP address or host name of a high-speed interconnect, if you want to use a high-speed interconnect for database partition server communication.

Multiple instances (Linux, UNIX)

It is possible to have more than one instance on a Linux or UNIX operating system if the DB2 product was installed with root privileges. Although each instance runs simultaneously, each is independent. Therefore, you can only work within one instance of the database manager at a time.

Note: To prevent environmental conflicts between two or more instances, you should ensure that each instance has its own home directory. Errors will be returned when the home directory is shared. Each home directory can be in the same or a different file system.

The instance owner and the group that is the System Administration (SYSADM) group are associated with every instance. The instance owner and the SYSADM

group are assigned during the process of creating the instance. One user ID or username can be used for only one instance, and that user ID or username is also referred to as the *instance owner*.

Each instance owner must have a unique home directory. All of the configuration files necessary to run the instance are created in the home directory of the instance owner's user ID or username. If it becomes necessary to remove the instance owner's user ID or username from the system, you could potentially lose files associated with the instance and lose access to data stored in this instance. For this reason, you should dedicate an instance owner user ID or username to be used exclusively to run the database manager.

The primary group of the instance owner is also important. This primary group automatically becomes the system administration group for the instance and gains SYSADM authority over the instance. Other user IDs or usernames that are members of the primary group of the instance owner also gain this level of authority. For this reason, you might want to assign the instance owner's user ID or username to a primary group that is reserved for the administration of instances. (Also, ensure that you assign a primary group to the instance owner user ID or username; otherwise, the system-default primary group is used.)

If you already have a group that you want to make the system administration group for the instance, you can simply assign this group as the primary group when you create the instance owner user ID or username. To give other users administration authority on the instance, add them to the group that is assigned as the system administration group.

To separate SYSADM authority between instances, ensure that each instance owner user ID or username uses a different primary group. However, if you choose to have a common SYSADM authority over multiple instances, you can use the same primary group for multiple instances.

Multiple instances (Windows)

It is possible to run multiple instances of the database manager on the same computer. Each instance of the database manager maintains its own databases and has its own database manager configuration parameters.

Note: The instances can also belong to different DB2 copies on a computer that can be at different levels of the database manager.

An instance of the database manager consists of the following:

- A Windows service that represents the instance. The name of the service is same as the instance name. The display name of the service (from the Services panel) is the instance name, prefixed with the "DB2 - " string. For example, for an instance named "DB2", there exists a Windows service called "DB2" with a display name of "DB2 - <DB2 Copy Name> - DB2".

Note: A Windows service is not created for client instances.

- An instance directory. This directory contains the database manager configuration files, the system database directory, the node directory, the Database Connection Services (DCS) directory, all the diagnostic log and dump files that are associated with the instance. The instance directory is by default a sub-directory inside the SQLLIB directory and has the same name as the instance name. For example, the instance directory for instance "DB2" is C:\SQLLIB\DB2, where C:\SQLLIB is where the database manager is installed. You

can use the registry variable DB2INSTPROF to change the default location of the instance directory. If the DB2INSTPROF registry variable is set to another location, then the instance directory is created under the directory pointed to by DB2INSTPROF. For example, if DB2INSTPROF=D:\DB2PROFS, then the instance directory will be D:\DB2PROFS\DB2.

- Set DB2INSTPROF to c:\DB2PROFS using the db2set.exe -g command
- Run DB2ICRT.exe command to create the instance.
- When you create an instance on Windows operating systems, the default locations for user data files, such as instance directories and the db2cli.ini file, are the following directories:
 - Documents and Settings\All Users\Application Data\IBM\DB2\copy name on the Windows XP and Windows 2003 operating systems
 - ProgramData\IBM\DB2\copy name on the Windows Vista operating system

Creating instances

Although an instance is created as part of the installation of the database manager, your business needs might require you to create additional instances.

Prerequisites

If you belong to the Administrative group on Windows, or you have root authority on Linux or UNIX platforms, you can add additional instances. The computer where you add the instance becomes the instance-owning computer (node zero). Ensure that you add instances on a computer where a DB2 administration server resides. Instance IDs should not be root or have password expired.

Restrictions

- On Linux and UNIX operating systems, additional instances cannot be created for non-root installations.
- If existing user IDs are used to create DB2 instances, make sure that the user IDs:
 - Are not locked
 - Do not have expired passwords

To add an instance using the command line, enter:

```
db2icrt <instance_name>
```

When creating instance on an AIX server, you must provide the fenced user id, for example:

```
DB2DIR/instance/db2icrt -u db2fenc1 db2inst1
```

When using the db2icrt command to add another DB2 instance, you should provide the login name of the instance owner and optionally specify the authentication type of the instance. The authentication type applies to all databases created under that instance. The authentication type is a statement of where the authenticating of users will take place.

You can change the location of the instance directory from DB2PATH using the DB2INSTPROF environment variable. You require write-access for the instance directory. If you want the directories created in a path other than DB2PATH, you have to set DB2INSTPROF before entering the db2icrt command.

For DB2 Enterprise Server Edition (ESE), you also need to declare that you are adding a new instance that is a partitioned database system. In addition, when working with a ESE instance having more than one database partition, and working with Fast Communication Manager (FCM), you can have multiple connections between database partitions by defining more TCP/IP ports when creating the instance.

For example, for Windows operating systems, use the `db2icrt` command with the `-r <port range>` parameter. The port range is shown as follows, where the `base_port` is the first port that can be used by FCM, and the `end_port` is the last port in a range of port numbers that can be used by FCM:

```
-r:<base_port,end_port>
```

Modifying instances

Instances are designed to be as independent as possible from the effects of subsequent installation and removal of products. On Linux and UNIX, you can update instances after the installation or removal of executables or components. On Windows, you run the `db2iupdt` command.

In most cases, existing instances automatically inherit or lose access to the function of the product being installed or removed. However, if certain executables or components are installed or removed, existing instances do not automatically inherit the new system configuration parameters or gain access to all the additional function. The instance must be updated.

If the database manager is updated by installing a Program Temporary Fix (PTF) or a patch, all the existing database instances should be updated using the `db2iupdt` command (root installations) or the `db2nrupdt` command (non-root installations).

You should ensure you understand the instances and database partition servers you have in an instance before attempting to change or delete an instance.

Updating the instance configuration (Linux, UNIX)

This topic applies to root instances only. To update non-root instances, run the `db2nrupdt` command.

Running the `db2iupdt` command updates the specified instance by performing the following:

- Replaces the files in the `sqllib` subdirectory under the instance owner's home directory.
- If the node type has changed, then a new database manager configuration file is created. This is done by merging relevant values from the existing database manager configuration file with the default database manager configuration file for the new node type. If a new database manager configuration file is created, the old file is backed up to the `backup` subdirectory of the `sqllib` subdirectory under the instance owner's home directory.

The `db2iupdt` command is found in `/usr/opt/db2_09_05/instance/` directory on AIX. The `db2iupdt` command is found in `/opt/IBM/db2/V9.5/instance/` directory on HP-UX, Solaris, or Linux.

To update an instance using the command line, enter:

```
db2iupdt InstName
```

The InstName is the login name of the instance owner.

There are other optional parameters associated with this command:

-h or -?

Displays a help menu for this command.

-d

Sets the debug mode for use during problem determination.

-a AuthType

Specifies the authentication type for the instance. Valid authentication types are SERVER, SERVER_ENCRYPT, or CLIENT. If not specified, the default is SERVER, if a DB2 server is installed. Otherwise, it is set to CLIENT. The authentication type of the instance applies to all databases owned by the instance.

-e

Allows you to update each instance that exists. Usedb2ilist to list the existing instances.

-u Fenced ID

Names the user under which the fenced user-defined functions (UDFs) and stored procedures will execute. This is not required if you install the Data Server Client or the DB2 Software Developer's Kit. For other DB2 products, this is a required parameter. Note: Fenced ID might not be "root" or "bin".

-k

This parameter preserves the current instance type. If you do not specify this parameter, the current instance is upgraded to the highest instance type available in the following order:

- Partitioned database server with local and remote clients
- Database Server with local and remote clients
- Client

Examples:

- If you installed DB2 Workgroup Server Edition or DB2 Enterprise Server Edition after the instance was created, enter the following command to update that instance:

```
db2iupdt -u db2fenc1 db2inst1
```
- If you installed the DB2[®] Connect[™] Enterprise Edition after creating the instance, you can use the instance name as the Fenced ID also:

```
db2iupdt -u db2inst1 db2inst1
```
- To update client instances, invoke the following command:

```
db2iupdt db2inst1
```

Updating the instance configuration (Windows)

To update the instance configuration on Windows, use the db2iupdt command.

Running the db2iupdt command updates the specified instance by performing the following:

- Replaces the files in the sql1ib subdirectory under the instance owner's home directory.
- If the node type is changed, then a new database manager configuration file is created. This is done by merging relevant values from the existing database manager configuration file with the default database manager configuration file for the new node type. If a new database manager configuration file is created, the old file is backed up to the backup subdirectory of the sql1ib subdirectory under the instance owner's home directory.

The db2iupdt command is found in \sql11b\bin directory.

The command is used as shown:

```
db2iupdt InstName
```

The InstName is the login name of the instance owner.

There are other optional parameters associated with this command:

/h: hostname

Overrides the default TCP/IP host name if there are one or more TCP/IP host names for the current computer.

/p: instance profile path

Specifies the new instance profile path for the updated instance.

/r: baseport,endport

Specifies the range of TCP/IP ports used by the partitioned database instance when running with multiple database partitions.

/u: username,password

Specifies the account name and password for the DB2 service.

Working with instances

When working with instances, you can start or stop instances, and attach to or detach from instances.

Each instance is managed by users who belong to the SYSADM_GROUP defined in the *instance configuration file*, also known as the *database manager configuration file*. Creating user IDs and user groups is different for each operating environment.

Auto-starting instances

On Windows operating systems, the database instance that is created during install is set as auto-started by default. An instance created using db2icrt is set as a manual start. To change the start type, you need to go to the Services panel and change the property of the DB2 service there.

On UNIX operating systems, to enable an instance to auto-start after each system restart, enter the following command:

```
db2iauto -on <instance name>
```

where <instance name> is the login name of the instance. On UNIX operating systems, to prevent an instance from auto-starting after each system restart, enter the following command:

```
db2iauto -off <instance name>
```

where <instance name> is the login name of the instance.

Starting instances (Linux, UNIX)

You might need to start or stop a DB2 database during normal business operations, for example, you must start an instance before you can perform some of the following tasks: connect to a database on the instance, precompile an application, bind a package to a database, or access host databases.

Before you start an instance on your Linux or UNIX system:

1. Log in with a user ID or name that has SYSADM, SYSCTRL, or SYSMANT authority on the instance; or log in as the instance owner.
2. Run the startup script as follows, where INSTHOME is the home directory of the instance you want to use:

```
. INSTHOME/sql1lib/db2profile      (for Bourne or Korn shell)
source INSTHOME/sql1lib/db2cshrc   (for C shell)
```

To start the instance using the command line, enter:

```
db2start
```

Note: When you run commands to start or stop an instance's database manager, the DB2 database manager applies the command to the current instance.

Starting instances (Windows)

You might need to start or stop a DB2 instance during normal business operations, for example, you must start an instance before you can perform some of the following tasks: connect to a database on the instance, precompile an application, bind a package to a database, or access a host database.

In order to successfully launch the DB2 database instance as a service from db2start, the user account must have the correct privilege as defined by the Windows operating system to start a Windows service. The user account can be a member of the Administrators, Server Operators, or Power Users group. When extended security is enabled, only members of the DB2ADMNS and Administrators groups can start the database by default.

To start an instance using the command line, enter:

```
db2start
```

Note: When you run commands to start or stop an instance's database manager, the DB2 database manager applies the command to the current instance.

The db2start command will launch the DB2 database instance as a Windows service. The DB2 database instance on Windows can still be run as a process by specifying the "/D" switch when invoking db2start. The DB2 database instance can also be started as a service using the Control Panel or the NET START command.

When running in a partitioned database environment, each database partition server is started as a Windows service. You can not use the "/D" switch to start a DB2 instance as a process in a partitioned database environment.

Attaching to and detaching from instances

On all platforms, to attach to another instance of the database manager, which might be remote, use the ATTACH command. To detach from an instance, use the DETACH command.

More than one instance must already exist.

To attach to an instance using the command line, enter:

```
db2 attach to <instance name>
```

For example, to attach to an instance called testdb2 that was previously cataloged in the node directory:

```
db2 attach to testdb2
```


After performing maintenance activities for the testdb2 instance, for example, to detach from an instance using the command line, enter:

```
db2 detach
```

Attaching to and detaching from client applications:

- To attach to an instance from a client application, call the sqleatin API,
- To detach from an instance from a client application, call the sqledtin API.

Working with instances on the same or different DB2 copies

You can run multiple instances concurrently, in the same DB2 copy or in different DB2 copies.

To work with instances in the same DB2 copy, you need to:

1. Create or migrate all instances to the same DB2 copy.
2. Set the DB2INSTANCE environment variable to the name of the instance you are working with before issuing commands against that instance.

To prevent one instance from accessing the database of another instance, the database files for an instance are created under a directory that has the same name as the instance name. For example, when creating a database on drive C: for instance "DB2", the database files are created inside a directory called C:\DB2. Similarly, when creating a database on drive C: for instance TEST, the database files are created inside a directory called C:\TEST. By default, its value is the drive letter where DB2 product is installed. For more information, see the *dftdbpath* database manager configuration parameter.

To work with an instance in a system with multiple DB2 copies, use either of the following methods:

- Using the Command window from the Start → Programs → IBM DB2 → <DB2 Copy Name> → Command Line Tools → Command Window: the Command window is already set up with the correct environment variables for the particular DB2 copy chosen.
- Using db2envar.bat from a Command window:
 1. Open a Command window.
 2. Run the db2envar.bat file using the fully qualified path for the DB2 copy that you want the application to use:

```
<DB2 Copy install dir>\bin\db2envar.bat
```

Stopping instances (Linux, UNIX)

You might need to stop the current instance of the database manager.

Before you begin

To stop an instance on your Linux or UNIX system, you must do the following:

1. Log in or attach to an instance with a user ID or name that has SYSADM, SYSCTRL, or SYSMANT authority on the instance; or, log in as the instance owner.
2. Display all applications and users that are connected to the specific database that you want to stop. To ensure that no vital or critical applications are running, list applications. You need SYSADM, SYSCTRL, or SYSMANT authority for this.

3. Force all applications and users off the database. You require SYSADM or SYSCTRL authority to force users.

About this task

The db2stop command can only be run at the server. No database connections are allowed when running this command; however, if there are any instance attachments, they are forced off before the instance is stopped.

Note: If command line processor sessions are attached to an instance, you must run the terminate command to end each session before running the db2stop command. The db2stop command stops the instance defined by the DB2INSTANCE environment variable.

Note: When you run commands to start or stop an instance's database manager, the DB2 database manager applies the command to the current instance. For more information, see Setting the current instance environment variables.

Procedure

To stop the instance using the command line, enter: db2stop

You can use the db2stop command to stop, or drop, individual database partitions within a partitioned database environment. When working in a partitioned database environment and you are attempting to drop a logical partition using

```
db2stop drop nodenum <0>
```

You must ensure that no users are attempting to access the database. If they are, you will receive an error message SQL6030N.

Stopping instances (Windows)

You might need to stop the current instance of the database manager.

Before you begin

To stop an instance on your system, you must do the following:

1. The user account stopping the DB2 database service must have the correct privilege as defined by the Windows operating system. The user account can be a member of the Administrators, Server Operators, or Power Users group.
2. Display all applications and users that are connected to the specific database that you want to stop. To ensure that no vital or critical applications are running, list applications. You need SYSADM, SYSCTRL, or SYSMANT authority for this.
3. Force all applications and users off the database. You require SYSADM or SYSCTRL authority to force users.

About this task

The db2stop command can only be run at the server. No database connections are allowed when running this command; however, if there are any instance attachments, they are forced off before the DB2 database service is stopped.

Note: If command line processor sessions are attached to an instance, you must run the terminate command to end each session before running the db2stop command. The db2stop command stops the instance defined by the DB2INSTANCE environment variable.

Recall that when you are using the database manager in a partitioned database environment, each database partition server is started as a service. Each service must be stopped.

Note: When you run commands to start or stop an instance's database manager, the database manager applies the command to the current instance. For more information, see *Setting the current instance environment variables*.

Procedure

To stop an instance on your system, use one of the following methods:

- Stop using the `db2stop` command.
- Stop using the `NET STOP` command.
- Stop the instance from within an application.

Dropping instances

This topic applies to root instances only on all platforms. To drop non-root instances, you must uninstall your DB2 product.

To remove an instance, from the command line, enter:

```
db2idrop <instance_name>
```

The preparation and details to removing an instance using the command line are:

1. Stop all applications that are currently using the instance.
2. Stop the Command Line Processor by running the terminate commands in each Command window.
3. Stop the instance by running the `db2stop` command.
4. Back up the instance directory indicated by the `DB2INSTPROF` registry variable.

On Linux and UNIX operating systems, consider backing up the files in the `INSTHOME/sqllib` directory (where `INSTHOME` is the home directory of the instance owner). For example, you might want to save the database manager configuration file, `db2system`, the `db2nodes.cfg` file, user-defined functions (UDFs), or fenced stored procedure applications.

5. (On Linux and UNIX operating systems only) Log off as the instance owner.
6. (On Linux and UNIX operating systems only) Log in as a user with root authority.
7. Issue the `db2idrop` command:

```
db2idrop InstName
```

where `InstName` is the name of the instance being dropped.

This command removes the instance entry from the list of instances and removes the instance directory.

8. (On Linux and UNIX operating systems only) Optionally, as a user with root authority, remove the instance owner's user ID and group (if used only for that instance). Do not remove these if you are planning to re-create the instance.

This step is optional since the instance owner and the instance owner group might be used for other purposes.

The `db2idrop` command removes the instance entry from the list of instances and removes the `sqllib` subdirectory under the instance owner's home directory.

Note: On Linux and UNIX operating systems, when attempting to drop an instance using the `db2idrop` command, a message is generated saying that the `sqllib` subdirectory cannot be removed, and in the `adm` subdirectory several files with the `.nfs` extension are being generated. The `adm` subdirectory is an NFS-mounted system and the files are controlled on the server. You must delete the `*.nfs` files from the file server from where the directory is being mounted. Then you can remove the `sqllib` subdirectory.

Chapter 5. Lightweight Directory Access Protocol (LDAP)

Lightweight Directory Access Protocol (LDAP) is an industry standard access method to directory services. A directory service is a repository of resource information about multiple systems and services within a distributed environment; and it provides client and server access to these resources.

Each database server instance publishes its existence to an LDAP server and provides database information to the LDAP directory when the databases are created. When a client connects to a database, the catalog information for the server can be retrieved from the LDAP directory. Each client is no longer required to store catalog information locally on each machine. Client applications search the LDAP directory for information required to connect to the database.

A caching mechanism exists so that the client only needs to search the LDAP directory server once. After the information is retrieved from the LDAP directory server, it is stored or cached on the local machine based on the values of the *dir_cache* database manager configuration parameter and the DB2LDAPCACHE registry variable. The *dir_cache* database manager configuration parameter is used to store database, node, and DCS directory files in a memory cache. The directory cache is used by an application until the application closes. The DB2LDAPCACHE registry variable is used to store database, node, and DCS directory files in a local disk cache.

- If DB2LDAPCACHE=NO and *dir_cache*=NO, then always read the information from LDAP.
- If DB2LDAPCACHE=NO and *dir_cache*=YES, then read the information from LDAP once and insert it into the DB2 cache.
- If DB2LDAPCACHE=YES or is not set, then read the information from LDAP once and cache it into the local database, node, and DCS directories.

Note: The DB2LDAPCACHE registry variable is only applicable to the database and node directories.

Security considerations in an LDAP environment

Before accessing information in the LDAP directory, an application or user is authenticated by the LDAP server. The authentication process is called *binding* to the LDAP server. It is important to apply access control on the information stored in the LDAP directory to prevent anonymous users from adding, deleting, or modifying the information.

Access control is inherited by default and can be applied at the container level. When a new object is created, it inherits the same security attribute as the parent object. An administration tool available for the LDAP server can be used to define access control for the container object.

By default, access control is defined as follows:

- For database and node entries in LDAP, everyone (or any anonymous user) has read access. Only the Directory Administrator and the owner or creator of the object has read/write access.

- For user profiles, the profile owner and the Directory Administrator have read/write access. One user cannot access the profile of another user if that user does not have Directory Administrator authority.

Note: The authorization check is always performed by the LDAP server and not by DB2. The LDAP authorization check is not related to DB2 authorization. An account or authorization ID that has SYSADM authority may not have access to the LDAP directory.

When running the LDAP commands or APIs, if the bind Distinguished Name (bindDN) and password are not specified, DB2 binds to the LDAP server using the default credentials which may not have sufficient authority to perform the requested commands and an error will be returned.

You can explicitly specify the user's bindDN and password using the USER and PASSWORD clauses for the DB2 commands or APIs.

LDAP object classes and attributes used by DB2

The following tables describe the object classes that are used by the DB2 database manager:

Table 11. cimManagedElement

Class	cimManagedElement
Active Directory LDAP Display Name	Not applicable
Active Directory Common Name (cn)	Not applicable
Description	Provides a base class of many of the system management object classes in the IBM Schema
SubClassOf	top
Required Attribute(s)	
Optional Attribute(s)	description
Type	abstract
OID (Object Identifier)	1.3.18.0.2.6.132
GUID (Global Unique Identifier)	b3afd63f-5c5b-11d3-b818-002035559151

Table 12. cimSetting

Class	cimSetting
Active Directory LDAP Display Name	Not applicable
Active Directory Common Name (cn)	Not applicable
Description	Provides a base class for configuration and settings in the IBM Schema
SubClassOf	cimManagedElement
Required Attribute(s)	
Optional Attribute(s)	settingID
Type	abstract
OID (object identifier)	1.3.18.0.2.6.131
GUID (Global Unique Identifier)	b3afd64d-5c5b-11d3-b818-002035559151

Table 13. eProperty

Class	eProperty
Active Directory LDAP Display Name	ibm-eProperty
Active Directory Common Name (cn)	ibm-eProperty
Description	Used to specify any application specific settings for user preference properties
SubClassOf	cimSetting
Required Attribute(s)	
Optional Attribute(s)	propertyType cisPropertyType cisProperty cesPropertyType cesProperty binPropertyType binProperty
Type	structural
OID (object identifier)	1.3.18.0.2.6.90
GUID (Global Unique Identifier)	b3afd69c-5c5b-11d3-b818-002035559151

Table 14. DB2Node

Class	DB2Node
Active Directory LDAP Display Name	ibm-db2Node
Active Directory Common Name (cn)	ibm-db2Node
Description	Represents a DB2 Server
SubClassOf	eSap / ServiceConnectionPoint
Required Attribute(s)	db2nodeName
Optional Attribute(s)	db2nodeAlias db2instanceName db2Type host / dNSHostName (see Note 2) protocolInformation/ServiceBindingInformation
Type	structural
OID (object identifier)	1.3.18.0.2.6.116
GUID (Global Unique Identifier)	b3afd65a-5c5b-11d3-b818-002035559151

Table 14. DB2Node (continued)

Class	DB2Node
Special Notes [®]	<ol style="list-style-type: none"> 1. The <i>DB2Node</i> class is derived from <i>eSap</i> object class under IBM Tivoli[®] Directory Server and from <i>ServiceConnectionPoint</i> object class under Microsoft[®] Active Directory. 2. The <i>host</i> is used under the IBM Tivoli Directory Server environment. The <i>dNSHostName</i> attribute is used under Microsoft Active Directory. 3. The <i>protocolInformation</i> is only used under the IBM Tivoli Directory Server environment. For Microsoft Active Directory, the attribute <i>ServiceBindingInformation</i>, inherited from the <i>ServiceConnectionPoint</i> class, is used to contain the protocol information.

The *protocolInformation* (in IBM Tivoli Directory Server) or *ServiceBindingInformation* (in Microsoft Active Directory) attribute in the *DB2Node* object contains the communication protocol information to bind the DB2 database server. It consists of tokens that describe the network protocol supported. Each token is separated by a semicolon. There is no space between the tokens. An asterisk (*) may be used to specify an optional parameter.

The tokens for TCP/IP are:

- "TCPIP"
- Server hostname or IP address
- Service name (svcname) or port number (e.g. 50000)
- (Optional) security ("NONE" or "SOCKS")

The tokens for Named Pipe are:

- "NPIPE"
- Computer name of the server
- Instance name of the server

Table 15. DB2Database

Class	DB2Database
Active Directory LDAP Display Name	ibm-db2Database
Active Directory Common Name (cn)	ibm-db2Database
Description	Represents a DB2 database
SubClassOf	top
Required Attribute(s)	db2databaseName db2nodePtr

Table 15. DB2Database (continued)

Class	DB2Database
Optional Attribute(s)	db2databaseAlias db2additionalParameter db2ARLibrary db2authenticationLocation db2gwPtr db2databaseRelease DCEPrincipalName db2altgwPtr db2altnodePtr
Type	structural
OID (object identifier)	1.3.18.0.2.6.117
GUID (Global Unique Identifier)	b3afd659-5c5b-11d3-b818-002035559151

Table 16. db2additionalParameters

Attribute	db2additionalParameters
Active Directory LDAP Display Name	ibm-db2AdditionalParameters
Active Directory Common Name (cn)	ibm-db2AdditionalParameters
Description	Contains any additional parameters used when connecting to the host database server
Syntax	Case Ignore String
Maximum Length	1024
Multi-Valued	Single-valued
OID (object identifier)	1.3.18.0.2.4.426
GUID (Global Unique Identifier)	b3afd315-5c5b-11d3-b818-002035559151

Table 17. db2authenticationLocation

Attribute	db2authenticationLocation
Active Directory LDAP Display Name	ibm-db2AuthenticationLocation
Active Directory Common Name (cn)	ibm-db2AuthenticationLocation
Description	Specifies where authentication takes place
Syntax	Case Ignore String
Maximum Length	64
Multi-Valued	Single-valued
OID (object identifier)	1.3.18.0.2.4.425
GUID (Global Unique Identifier)	b3afd317-5c5b-11d3-b818-002035559151
Notes	Valid values are: CLIENT, SERVER, DCS, DCE, KERBEROS, SVRENCRYPT, or DCSENCRIPT

Table 18. db2ARLibrary

Attribute	db2ARLibrary
Active Directory LDAP Display Name	ibm-db2ARLibrary
Active Directory Common Name (cn)	ibm-db2ARLibrary
Description	Name of the Application Requestor library
Syntax	Case Ignore String
Maximum Length	256
Multi-Valued	Single-valued
OID (object identifier)	1.3.18.0.2.4.427
GUID (Global Unique Identifier)	b3afd316-5c5b-11d3-b818-002035559151

Table 19. db2databaseAlias

Attribute	db2databaseAlias
Active Directory LDAP Display Name	ibm-db2DatabaseAlias
Active Directory Common Name (cn)	ibm-db2DatabaseAlias
Description	Database alias name(s)
Syntax	Case Ignore String
Maximum Length	1024
Multi-Valued	Multi-valued
OID (object identifier)	1.3.18.0.2.4.422
GUID (Global Unique Identifier)	b3afd318-5c5b-11d3-b818-002035559151

Table 20. db2databaseName

Attribute	db2databaseName
Active Directory LDAP Display Name	ibm-db2DatabaseName
Active Directory Common Name (cn)	ibm-db2DatabaseName
Description	Database name
Syntax	Case Ignore String
Maximum Length	1024
Multi-Valued	Single-valued
OID (object identifier)	1.3.18.0.2.4.421
GUID (Global Unique Identifier)	b3afd319-5c5b-11d3-b818-002035559151

Table 21. db2databaseRelease

Attribute	db2databaseRelease
Active Directory LDAP Display Name	ibm-db2DatabaseRelease
Active Directory Common Name (cn)	ibm-db2DatabaseRelease
Description	Database release number
Syntax	Case Ignore String
Maximum Length	64
Multi-Valued	Single-valued
OID (object identifier)	1.3.18.0.2.4.429

Table 21. db2databaseRelease (continued)

Attribute	db2databaseRelease
GUID (Global Unique Identifier)	b3afd31a-5c5b-11d3-b818-002035559151

Table 22. db2nodeAlias

Attribute	db2nodeAlias
Active Directory LDAP Display Name	ibm-db2NodeAlias
Active Directory Common Name (cn)	ibm-db2NodeAlias
Description	Node alias name(s)
Syntax	Case Ignore String
Maximum Length	1024
Multi-Valued	Multi-valued
OID (object identifier)	1.3.18.0.2.4.420
GUID (Global Unique Identifier)	b3afd31d-5c5b-11d3-b818-002035559151

Table 23. db2nodeName

Attribute	db2nodeName
Active Directory LDAP Display Name	ibm-db2NodeName
Active Directory Common Name (cn)	ibm-db2NodeName
Description	Node name
Syntax	Case Ignore String
Maximum Length	64
Multi-Valued	Single-valued
OID (object identifier)	1.3.18.0.2.4.419
GUID (Global Unique Identifier)	b3afd31e-5c5b-11d3-b818-002035559151

Table 24. db2nodePtr

Attribute	db2nodePtr
Active Directory LDAP Display Name	ibm-db2NodePtr
Active Directory Common Name (cn)	ibm-db2NodePtr
Description	Pointer to the Node (DB2Node) object that represents the database server which owns the database
Syntax	Distinguished Name
Maximum Length	1000
Multi-Valued	Single-valued
OID (object identifier)	1.3.18.0.2.4.423
GUID (Global Unique Identifier)	b3afd31f-5c5b-11d3-b818-002035559151
Special Notes	This relationship allows the client to retrieve protocol communication information to connect to the database

Table 25. db2altnodePtr

Attribute	db2altnodePtr
Active Directory LDAP Display Name	ibm-db2AltNodePtr

Table 25. db2altnodePtr (continued)

Attribute	db2altnodePtr
Active Directory Common Name (cn)	ibm-db2AltNodePtr
Description	Pointer to the Node (DB2Node) object that represents the alternate database server
Syntax	Distinguished Name
Maximum Length	1000
Multi-Valued	Multi-valued
OID (object identifier)	1.3.18.0.2.4.3093
GUID (Global Unique Identifier)	5694e266-2059-4e32-971e-0778909e0e72

Table 26. db2gwPtr

Attribute	db2gwPtr
Active Directory LDAP Display Name	ibm-db2GwPtr
Active Directory Common Name (cn)	ibm-db2GwPtr
Description	Pointer to the Node object that represents the gateway server and from which the database can be accessed
Syntax	Distinguished Name
Maximum Length	1000
Multi-Valued	Single-valued
OID (object identifier)	1.3.18.0.2.4.424
GUID (Global Unique Identifier)	b3afd31b-5c5b-11d3-b818-002035559151

Table 27. db2altgwPtr

Attribute	db2altgwPtr
Active Directory LDAP Display Name	ibm-db2AltGwPtr
Active Directory Common Name (cn)	ibm-db2AltGwPtr
Description	Pointer to the Node object that represents the alternate gateway server
Syntax	Distinguished Name
Maximum Length	1000
Multi-Valued	Multi-valued
OID (object identifier)	1.3.18.0.2.4.3092
GUID (Global Unique Identifier)	70ab425d-65cc-4d7f-91d8-084888b3a6db

Table 28. db2instanceName

Attribute	db2instanceName
Active Directory LDAP Display Name	ibm-db2InstanceName
Active Directory Common Name (cn)	ibm-db2InstanceName
Description	The name of the database server instance
Syntax	Case Ignore String
Maximum Length	256
Multi-Valued	Single-valued

Table 28. *db2instanceName* (continued)

Attribute	db2instanceName
OID (object identifier)	1.3.18.0.2.4.428
GUID (Global Unique Identifier)	b3afd31c-5c5b-11d3-b818-002035559151

Table 29. *db2Type*

Attribute	db2Type
Active Directory LDAP Display Name	ibm-db2Type
Active Directory Common Name (cn)	ibm-db2Type
Description	Type of the database server
Syntax	Case Ignore String
Maximum Length	64
Multi-Valued	Single-valued
OID (object identifier)	1.3.18.0.2.4.418
GUID (Global Unique Identifier)	b3afd320-5c5b-11d3-b818-002035559151
Notes	Valid types for database server are: SERVER, MPP, and DCS

Table 30. *DCEPrincipalName*

Attribute	DCEPrincipalName
Active Directory LDAP Display Name	ibm-DCEPrincipalName
Active Directory Common Name (cn)	ibm-DCEPrincipalName
Description	DCE principal name
Syntax	Case Ignore String
Maximum Length	2048
Multi-Valued	Single-valued
OID (object identifier)	1.3.18.0.2.4.443
GUID (Global Unique Identifier)	b3afd32d-5c5b-11d3-b818-002035559151

Table 31. *cesProperty*

Attribute	cesProperty
Active Directory LDAP Display Name	ibm-cesProperty
Active Directory Common Name (cn)	ibm-cesProperty
Description	Values of this attribute may be used to provide application-specific preference configuration parameters. For example, a value may contain XML-formatted data. All values of this attribute must be homogeneous in the cesPropertyType attribute value.
Syntax	Case Exact String
Maximum Length	32700
Multi-Valued	Multi-valued
OID (object identifier)	1.3.18.0.2.4.307
GUID (Global Unique Identifier)	b3afd2d5-5c5b-11d3-b818-002035559151

Table 32. cesPropertyType

Attribute	cesPropertyType
Active Directory LDAP Display Name	ibm-cesPropertyType
Active Directory Common Name (cn)	ibm-cesPropertyType
Description	Values of this attribute may be used to describe the syntax, semantics, or other characteristics of all of the values of the cesProperty attribute. For example, a value of "XML" might be used to indicate that all the values of the cesProperty attribute are encoded as XML syntax.
Syntax	Case Ignore String
Maximum Length	128
Multi-Valued	Multi-valued
OID (object identifier)	1.3.18.0.2.4.308
GUID (Global Unique Identifier)	b3afd2d6-5c5b-11d3-b818-002035559151

Table 33. cisProperty

Attribute	cisProperty
Active Directory LDAP Display Name	ibm-cisProperty
Active Directory Common Name (cn)	ibm-cisProperty
Description	Values of this attribute may be used to provide application-specific preference configuration parameters. For example, a value may contain an INI file. All values of this attribute must be homogeneous in their cisPropertyType attribute value.
Syntax	Case Ignore String
Maximum Length	32700
Multi-Valued	Multi-valued
OID (object identifier)	1.3.18.0.2.4.309
GUID (Global Unique Identifier)	b3afd2e0-5c5b-11d3-b818-002035559151

Table 34. cisPropertyType

Attribute	cisPropertyType
Active Directory LDAP Display Name	ibm-cisPropertyType
Active Directory Common Name (cn)	ibm-cisPropertyType
Description	Values of this attribute may be used to describe the syntax, semantics, or other characteristics of all of the values of the cisProperty attribute. For example, a value of "INI File" might be used to indicate that all the values of the cisProperty attribute are INI files.
Syntax	Case Ignore String
Maximum Length	128
Multi-Valued	Multi-valued
OID (object identifier)	1.3.18.0.2.4.310
GUID (Global Unique Identifier)	b3afd2e1-5c5b-11d3-b818-002035559151

Table 35. binProperty

Attribute	binProperty
Active Directory LDAP Display Name	ibm-binProperty
Active Directory Common Name (cn)	ibm-binProperty
Description	Values of this attribute may be used to provide application-specific preference configuration parameters. For example, a value may contain a set of binary-encoded Lotus® 123 properties. All values of this attribute must be homogeneous in their binPropertyType attribute values.
Syntax	binary
Maximum Length	250000
Multi-Valued	Multi-valued
OID (object identifier)	1.3.18.0.2.4.305
GUID (Global Unique Identifier)	b3afd2ba-5c5b-11d3-b818-002035559151

Table 36. binPropertyType

Attribute	binPropertyType
Active Directory LDAP Display Name	ibm-binPropertyType
Active Directory Common Name (cn)	ibm-binPropertyType
Description	Values of this attribute may be used to describe the syntax, semantics, or other characteristics of all of the values of the binProperty attribute. For example, a value of “Lotus 123” might be used to indicate that all the values of the binProperty attribute are binary-encoded Lotus 123 properties.
Syntax	Case Ignore String
Maximum Length	128
Multi-Valued	Multi-valued
OID (object identifier)	1.3.18.0.2.4.306
GUID (Global Unique Identifier)	b3afd2bb-5c5b-11d3-b818-002035559151

Table 37. PropertyType

Attribute	PropertyType
Active Directory LDAP Display Name	ibm-propertyType
Active Directory Common Name (cn)	ibm-propertyType
Description	Values of this attribute describe the semantic characteristics of the eProperty object
Syntax	Case Ignore String
Maximum Length	128
Multi-Valued	Multi-valued
OID (object identifier)	1.3.18.0.2.4.320
GUID (Global Unique Identifier)	b3afd4ed-5c5b-11d3-b818-002035559151

Table 38. settingID

Attribute	settingID
Active Directory LDAP Display Name	Not applicable
Active Directory Common Name (cn)	Not applicable
Description	A naming attribute that may be used to identify the cimSetting derived object entries such as eProperty
Syntax	Case Ignore String
Maximum Length	256
Multi-Valued	Single-valued
OID (object identifier)	1.3.18.0.2.4.325
GUID (Global Unique Identifier)	b3afd596-5c5b-11d3-b818-002035559151

Extending the LDAP directory schema with DB2 object classes and attributes

The LDAP Directory Schema defines object classes and attributes for the information stored in the LDAP directory entries. An object class consists of a set of mandatory and optional attributes. Every entry in the LDAP directory has an object class associated with it.

Before the DB2 database manager can store information in LDAP, the Directory Schema for the LDAP server must include the object classes and attributes that the DB2 database system uses. The process of adding new object classes and attributes to the base schema is called *schema extension*.

Supported LDAP client and server configurations

The following table summarizes the supported LDAP client and server configurations.

Table 39. Supported LDAP client and server configurations

LDAP clients	LDAP servers		
	IBM Tivoli Directory server ¹	Microsoft Active Directory server ²	Sun One LDAP server
IBM LDAP Client ³	Supported	Supported	Supported
Microsoft LDAP/ADSI Client ⁴	Supported	Supported	Supported

- ¹ IBM Tivoli Directory Server is an LDAP Version 3 server and is available for Windows, AIX, Solaris, Linux, and HP-UX. It is shipped as part of the base operating system on AIX and System i™, and with OS/390® Security Server.
- ² Microsoft Active Directory server is an LDAP Version 3 server and is available as part of the Windows 2000 Server and Windows Server 2003 family of operating systems.
- ³ The DB2 database system supports IBM LDAP client on AIX, Solaris, HP-UX 11.11, Windows, and Linux.
- ⁴ The Microsoft LDAP Client is included with the Windows operating system.

Note: When running on Windows operating systems, the DB2 database manager supports using either the IBM LDAP client or the Microsoft LDAP client. To

explicitly select the IBM LDAP client, use the db2set command to set the **DB2LDAP_CLIENT_PROVIDER** registry variable to IBM. The Microsoft LDAP Client is included with the Windows operating system.

LDAP support and DB2 Connect

If LDAP support is available at the DB2 Connect gateway, and the database is not found at the gateway database directory, then the DB2 database manager will look up the database location in LDAP and will attempt to keep the found information.

Registering host databases in LDAP

When you register host databases in LDAP, there are two possible configurations: direct connection to the host databases or, connection to the host database through a gateway.

For direct connection to the host databases, you register the host server in LDAP, then catalog the host database in LDAP specifying the node name of the host server. For connection to the host database through a gateway, you register the gateway server in LDAP, then catalog the host database in LDAP specifying the node name of the gateway server.

If LDAP support is available at the DB2 Connect gateway, and the database is not found at the gateway database directory, the DB2 database system looks up LDAP and attempts to keep the found information.

The following example shows both cases, consider the following: Suppose there is a host database called NIAGARA_FALLS. It can accept incoming connections using TCP/IP. If the client cannot connect directly to the host because it does not have DB2 Connect, then it connects using a gateway called goto@niagara.

The following steps need to be done:

1. Register the host database server in LDAP for TCP/IP connectivity. The TCP/IP hostname of the server is "myhost" and the port number is "446". The **NODETYPE** clause is set to DCS to indicate that this is a host database server.

```
db2 register ldap as nftcpip tcpip hostname myhost svcname 446
remote mvssys instance mvsinst nodetype dcs
```

2. Register a DB2 Connect gateway server in LDAP for TCP/IP connectivity. The TCP/IP hostname for the gateway server is "niagara" and the port number is "50000".

```
db2 register ldap as whasf tcpip hostname niagara svcname 50000
remote niagara instance goto nodetype server
```

3. Catalog the host database in LDAP using TCP/IP connectivity. The host database name is "NIAGARA_FALLS", the database alias name is "nftcpip". The **GWNODE** clause is used to specify the nodename of the DB2 Connect gateway server.

```
db2 catalog ldap database NIAGARA_FALLS as nftcpip at node nftcpip
gwnode whasf authentication server
```

After completing the registration and cataloging shown above, if you want to connect to the host using TCPIP, you connect to nftcpip. If you do not have DB2 Connect on your client workstation, the connection goes through the gateway using TCPIP. From the gateway, it connects to the host using TCP/IP.

In general, you can manually configure host database information in LDAP so that each client does not need to manually catalog the database and node locally on each machine. The process follows:

1. Register the host database server in LDAP. You must specify the remote computer name, instance name, and the node type for the host database server in the REGISTER command using the **REMOTE**, **INSTANCE**, and **NODETYPE** clauses respectively. The **REMOTE** clause can be set to either the host name or the LU name of the host server machine. The **INSTANCE** clause can be set to any character string that has eight characters or less. (For example, the instance name can be set to "DB2".) The **NODETYPE** clause must be set to DCS to indicate that this is a host database server.
2. Register the host database in LDAP using the CATALOG LDAP DATABASE command. Any additional DRDA[®] parameters can be specified by using the **PARMS** clause. The database authentication type should be set to SERVER.

Extending the directory schema for IBM Tivoli Directory Server

If you are using IBM Tivoli Directory Server, all the object classes and attributes that are required by the DB2 database before Version 8.2 are included in the base schema.

Run the following command to extend the base schema with new DB2 database attributes introduced in Version 8.2, and later:

```
ldapmodify -c -h machine_name:389 -D dn -w password -f altgwnode.ldif
```

The following is the content of the altgwnode.ldif file:

```

dn: cn=schema
changetype: modify
add: attributetypes
attributetypes: (
  1.3.18.0.2.4.3092
  NAME 'db2altgwPtr'
  DESC 'DN pointer to DB2 alternate gateway (node) object'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.12)
-

```

```

add: ibmattributetypes
ibmattributetypes: (
  1.3.18.0.2.4.3092
  DBNAME ('db2altgwPtr' 'db2altgwPtr')
  ACCESS-CLASS NORMAL
  LENGTH 1000)

```

```

dn: cn=schema
changetype: modify
add: attributetypes
attributetypes: (
  1.3.18.0.2.4.3093
  NAME 'db2altnodePtr'
  DESC 'DN pointer to DB2 alternate node object'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.12)
-

```

```

add: ibmattributetypes
ibmattributetypes: (
  1.3.18.0.2.4.3093
  DBNAME ('db2altnodePtr' 'db2altnodePtr')
  ACCESS-CLASS NORMAL
  LENGTH 1000)

```

```

dn: cn=schema
changetype: modify
replace: objectclasses
objectclasses: (
  1.3.18.0.2.6.117
  NAME 'DB2Database'
  DESC 'DB2 database'
  SUP cimSetting
  MUST ( db2databaseName $ db2nodePtr )
  MAY ( db2additionalParameters $ db2altgwPtr $ db2altnodePtr
        $ db2ARLibrary $ db2authenticationLocation $ db2databaseAlias
        $ db2databaseRelease $ db2gwPtr $ DCEPrincipalName ) )

```

The altgwnode.ldif and altgwnode.readmefiles can be found at URL:
<ftp://ftp.software.ibm.com/ps/products/db2/tools/ldap>

After adding the DB2 schema definition, the Directory Server must be restarted for all changes to be active.

Netscape LDAP directory support and attribute definitions

The supported level for Netscape LDAP Server is Version 4.12, or later.

Within Netscape LDAP Server Version 4.12, or later, the Netscape Directory Server allows applications to extend the schema by adding attribute and object class definitions to the following two files, slapd.user_oc.conf and slapd.user_at.conf. These two files are located in the <Netscape_install path>\slapd-<machine_name>\config directory.

Note: If you are using Sun One Directory Server 5.0, refer to the topic about extending the directory schema for the Sun One Directory Server.

The DB2 attributes must be added to the `slapd.user_at.conf` as follows:

```
#####
#
# IBM DB2 Database
# Attribute Definitions
#
# bin -> binary
# ces -> case exact string
# cis -> case insensitive string
# dn -> distinguished name
#
#####

attribute binProperty                1.3.18.0.2.4.305    bin
attribute binPropertyType            1.3.18.0.2.4.306    cis
attribute cesProperty                1.3.18.0.2.4.307    ces
attribute cesPropertyType            1.3.18.0.2.4.308    cis
attribute cisProperty                1.3.18.0.2.4.309    cis
attribute cisPropertyType            1.3.18.0.2.4.310    cis
attribute propertyType                1.3.18.0.2.4.320    cis
attribute systemName                 1.3.18.0.2.4.329    cis
attribute db2nodeName                1.3.18.0.2.4.419    cis
attribute db2nodeAlias                1.3.18.0.2.4.420    cis
attribute db2instanceName            1.3.18.0.2.4.428    cis
attribute db2Type                     1.3.18.0.2.4.418    cis
attribute db2databaseName            1.3.18.0.2.4.421    cis
attribute db2databaseAlias            1.3.18.0.2.4.422    cis
attribute db2nodePtr                  1.3.18.0.2.4.423    dn
attribute db2gwPtr                    1.3.18.0.2.4.424    dn
attribute db2additionalParameters    1.3.18.0.2.4.426    cis
attribute db2ARLibrary                1.3.18.0.2.4.427    cis
attribute db2authenticationLocation  1.3.18.0.2.4.425    cis
attribute db2databaseRelease         1.3.18.0.2.4.429    cis
attribute DCEPrincipalName           1.3.18.0.2.4.443    cis
```

The DB2 object classes must be added to the `slapd.user_oc.conf` file as follows:

```
#####
#
# IBM DB2 Database
# Object Class Definitions
#
#####

objectclass eProperty
    oid 1.3.18.0.2.6.90
    requires
        objectClass
    allows
        cn,
        propertyType,
        binProperty,
        binPropertyType,
        cesProperty,
        cesPropertyType,
        cisProperty,
        cisPropertyType

objectclass eApplicationSystem
    oid 1.3.18.0.2.6.84
    requires
        objectClass,
        systemName
```

```

objectclass DB2Node
  oid 1.3.18.0.2.6.116
  requires
    objectClass,
    db2nodeName
  allows
    db2nodeAlias,
    host,
    db2instanceName,
    db2Type,
    description,
    protocolInformation

objectclass DB2Database
  oid 1.3.18.0.2.6.117
  requires
    objectClass,
    db2databaseName,
    db2nodePtr
  allows
    db2databaseAlias,
    description,
    db2gwPtr,
    db2additionalParameters,
    db2authenticationLocation,
    DCEPrincipalName,
    db2databaseRelease,
    db2ARLibrary

```

After adding the DB2 schema definition, the Directory Server must be restarted for all changes to be active.

Extending the directory schema for Sun One Directory Server

The Sun One Directory Server is also known as the Netscape or iPlanet directory server.

To have the Sun One Directory Server work in your environment, add the 60ibmdb2.ldif file to the following directory:

On Windows, if you have iPlanet installed in C:\iPlanet\Servers, add the above file to .\sldap-<machine_name>\config\schema.

On UNIX, if you have iPlanet installed in /usr/iplanet/servers, add the above file to ./slapd-<machine_name>/config/schema.

The following is the contents of the file:

```

#####
# IBM DB2 Database
#####
dn: cn=schema
#####
# Attribute Definitions (Before V8.2)
#####
attributetypes: ( 1.3.18.0.2.4.305 NAME 'binProperty'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.5 X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.306 NAME 'binPropertyType'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.307 NAME 'cesProperty'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.308 NAME 'cesPropertyType'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 X-ORIGIN 'IBM DB2' )

```

```

attributetypes: ( 1.3.18.0.2.4.309 NAME 'cisProperty'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.310 NAME 'cisPropertyType'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.320 NAME 'propertyType'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.329 NAME 'systemName'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 SINGLE-VALUE X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.419 NAME 'db2nodeName'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 SINGLE-VALUE X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.420 NAME 'db2nodeAlias'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.428 NAME 'db2instanceName'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 SINGLE-VALUE X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.418 NAME 'db2Type'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 SINGLE-VALUE X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.421 NAME 'db2databaseName'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 SINGLE-VALUE X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.422 NAME 'db2databaseAlias'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.426 NAME 'db2additionalParameters'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 SINGLE-VALUE X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.427 NAME 'db2ARLibrary'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 SINGLE-VALUE X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.425 NAME 'db2authenticationLocation'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 SINGLE-VALUE X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.429 NAME 'db2databaseRelease'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 SINGLE-VALUE X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.443 NAME 'DCEPrincipalName'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 SINGLE-VALUE X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.423 NAME 'db2nodePtr'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.12 SINGLE-VALUE X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.424 NAME 'db2gwPtr'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.12 SINGLE-VALUE X-ORIGIN 'IBM DB2' )
#####
# Attribute Definitions (V8.2 and later)
#####
attributetypes: ( 1.3.18.0.2.4.3092 NAME 'db2altgwPtr'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.12 X-ORIGIN 'IBM DB2' )
attributetypes: ( 1.3.18.0.2.4.3093 NAME 'db2altnodePtr'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.12 X-ORIGIN 'IBM DB2' )
#####
# Object Class Definitions
# DB2Database for V8.2 has the above two new optional attributes.
#####
objectClasses: ( 1.3.18.0.2.6.90 NAME 'eProperty'
  SUP top STRUCTURAL MAY ( cn $ propertyType $ binProperty
    $ binPropertyType $ cesProperty $ cesPropertyType $ cisProperty
    $ cisPropertyType ) X-ORIGIN 'IBM DB2' )
objectClasses: ( 1.3.18.0.2.6.84 NAME 'eApplicationSystem'
  SUP top STRUCTURAL MUST systemName
  X-ORIGIN 'IBM DB2' )
objectClasses: ( 1.3.18.0.2.6.116 NAME 'DB2Node'
  SUP top STRUCTURAL MUST db2nodeName MAY ( db2instanceName $ db2nodeAlias
    $ db2Type $ description $ host $ protocolInformation )
  X-ORIGIN 'IBM DB2' )
objectClasses: ( 1.3.18.0.2.6.117 NAME 'DB2Database'
  SUP top STRUCTURAL MUST (db2databaseName $ db2nodePtr ) MAY
    ( db2additionalParameters $ db2altgwPtr $ db2altnodePtr $ db2ARLibrary
    $ db2authenticationLocation $ db2databaseAlias $ db2databaseRelease
    $ db2gwPtr $ DCEPrincipalName $ description )
  X-ORIGIN 'IBM DB2' )

```

The 60ibmdb2.ldif and 60ibmdb2.readmefiles can be found at URL:
<ftp://ftp.software.ibm.com/ps/products/db2/tools/ldap>

After adding the DB2 schema definition, the Directory Server must be restarted for all changes to be active.

Windows Active Directory

The DB2 database servers are published in the Active Directory as the `ibm_db2Node` objects. The `ibm_db2Node` object class is a subclass of the `ServiceConnectionPoint` (SCP) object class.

Each `ibm_db2Node` object contains protocol configuration information to allow client applications to connect to the DB2 database server. When a new database is created, the database is published in the Active Directory as the `ibm_db2Database` object under the `ibm_db2Node` object.

When connecting to a remote database, a DB2 client queries the Active Directory, through the LDAP interface, for the `ibm_db2Database` object. The protocol communication to connect to the database server (binding information) is obtained from the `ibm_db2Node` object, which the `ibm_db2Database` object is created under.

Property pages for the `ibm_db2Node` and `ibm_db2Database` objects can be viewed or modified using the *Active Directory Users and Computer* Management Console (MMC) at a domain controller. To set up the property page, run the `regsvr32` command to register the property pages for the DB2 objects as follows:

```
regsvr32 %DB2PATH%\bin\db2ads.dll
```

You can view the objects by using the *Active Directory Users and Computer* Management Console (MMC) at a domain controller. To get to this administration tool, follow Start—> Program—> Administration Tools—> Active Directory Users and Computer.

Note: You must select *Users, Groups, and Computers as containers* from the View menu to display the DB2 database objects under the computer objects.

Note: If the DB2 database system is not installed on the domain controller, you can still view the property pages of DB2 database objects by copying the `db2ads.dll` file from `%DB2PATH%\bin` and the resource DLL `db2adsr.dll` from `%DB2PATH%\msg\locale-name` to a local directory on the domain controller. (The directory where you place these two copied files must be one of those found in the `PATH` environment variable.) Then, you run the `regsvr32` command from the local directory to register the DLL.

Configuring the DB2 database manager to use Active Directory

In order to access Microsoft Active Directory, ensure that the following conditions are met:

1. The machine that runs DB2 database must belong to a Windows 2000 or Windows Server 2003 domain.
2. The Microsoft LDAP client is installed. The Microsoft LDAP client is part of the Windows 2000, Windows XP, and Windows Server 2003 operating systems.
3. Enable LDAP support. For more information, see “Extending the Active Directory Schema for LDAP directory services (Windows)” in *Quick Beginnings for DB2 Servers*.
4. Log on to a domain user account when running the DB2 database system to read information from the Active Directory.

Security considerations for Active Directory

The DB2 database and node objects are created under the computer object of the machine where the DB2 server is installed in the Active Directory. To register a database server or to catalog a database in the Active Directory, you need to have sufficient access to create or update the objects under the computer object.

By default, objects under the computer object are readable by any authenticated users and can be updated by administrators (users that belong to the Administrators, Domain Administrators, and Enterprise Administrators groups). To grant access for a specific user or a group, use the *Active Directory Users and Computer* Management Console (MMC) as follows:

1. Start the *Active Directory Users and Computer* administration tool
(Start—> Program—> Administration Tools—> Active Directory Users and Computer)
2. Under *View*, select *Advanced Features*
3. Select the *Computers* container
4. Right click on the computer object that represents the server machine where DB2 is installed and select *Properties*
5. Select the *Security* tab, then add the required access to the specified user or group

The DB2 registry variables and CLI settings at the user level are maintained in the DB2 property object under the user object. To set the DB2 registry variables or CLI settings at the user level, a user needs to have sufficient access to create objects under the User object.

By default, only administrators have access to create objects under the User object. To grant access to a user to set the DB2 registry variables or CLI settings at the user level, use the *Active Directory Users and Computer* Management Console (MMC) as follows:

1. Start the *Active Directory Users and Computer* administration tool
(Start—> Program—> Administration Tools—> Active Directory Users and Computer)
2. Select the user object under the Users container
3. Right click on the user object and select *Properties*
4. Select the *Security* tab
5. Add the user name to the list by using the Add button
6. Grant “Write”, and “Create All Child Objects” access
7. Using the Advanced setting, set permissions to apply onto “This object and all child objects”
8. Select the check box “Allow inheritable permissions from parent to propagate to this object”

DB2 objects in the Active Directory

The DB2 database manager creates objects in the Active Directory at two locations:

1. The DB2 database and node objects are created under the computer object of the machine where the DB2 server is installed. For the DB2 server machine that does not belong to the Windows domain, the DB2 database and node objects are created under the “System” container.
2. The DB2 registry variables and CLI settings at the user level are stored in the DB2 property objects under the User object. These objects contain information that is specific to that user.

Extending the directory schema for Active Directory

Before the DB2 database manager can store information in the Active Directory, the directory schema needs to be extended to include the new DB2 database object classes and attributes. The process of adding new object classes and attributes to the directory schema is called *schema extension*.

You must extend the schema for Active Directory by running the DB2 Schema Installation program, db2schex. You should run this command before installing DB2 products and creating databases, otherwise you have to manually register the node and catalog the databases.

The db2schex program is included on the product CD-ROM in the following location: x:\db2\windows\utilities\ where x: is the DVD drive letter.

To update the schema, you must be a member of the Schema Administrators group or have been delegated the rights to update the schema. Run the following command on any machine that is part of the Windows domain:

```
runas /user:MyDomain\Administrator x:\db2\Windows\utilities\db2schex.exe
```

where x: represents the DVD drive letter.

If you have run the db2schex command in an earlier version of the DB2 database management system, when you run this same command again on DB2 UDB Version 8.2, or later, the following two optional attributes are added to the ibm-db2Database class:

```
ibm-db2AltGwPtr  
ibm-db2NodePtr
```

If you have not run the db2schex command on an earlier version of the DB2 database management system on Windows, when you run this same command on DB2 Version 9.5, or later, all the classes and attributes for DB2 database system LDAP support are added.

There are other optional clauses associated with this command. For more information, refer to the “db2schex - Active Directory schema extension command” topic.

Examples:

- To install the DB2 database schema:

```
db2schex
```

- To install the DB2 database schema and specify a bind DN and password:

```
db2schex -b "cn=A Name,dc=toronto1,dc=ibm,dc=com"  
-w password
```

Or,

```
db2schex -b Administrator -w password
```

- To uninstall the DB2 database schema:

```
db2schex -u
```

- To uninstall the DB2 database schema and ignore errors:

```
db2schex -u -k
```

Enabling LDAP support after DB2 installation is complete

Before you can use LDAP, you must enable it after the DB2 product installation is complete.

To enable LDAP support:

1. On any machine that is part of a Windows domain, perform the following steps:
 - a. If you did not do so before installing the DB2 product, you must extend the directory schema if you want to use Microsoft Active Directory. For more information, see the “Extending the directory schema for Active Directory” topic.
 - b. Install the LDAP support binary files by running the DB2 Setup program and selecting the LDAP Directory Exploitation support from Custom install. The Setup program sets automatically the DB2 registry variable **DB2_ENABLE_LDAP** to YES which is a required setting to enable LDAP support.
 - c. Optional: To use the IBM LDAP client instead of the Microsoft LDAP client, set the **DB2LDAP_CLIENT_PROVIDER** registry variable to IBM.
2. On each LDAP client, perform the following steps:
 - a. Specify the TCP/IP host name and optionally the port number of the LDAP server by running the following command: `db2set DB2LDAPHOST=base_domain_name[:port_number]` where *base_domain_name* is the TCP/IP hostname, and [*port_number*] is the port number. If you do not specify a port number, the default LDAP port number 389 is used.
DB2 objects are located in the LDAP base distinguished name (baseDN). You can configure the baseDN on each machine by running the following command:

```
db2set DB2LDAP_BASEDN=baseDN
```

where *baseDN* is the name of the LDAP suffix that is defined at the LDAP server.
 - b. Optional: To use LDAP to store DB2 user-specific information, enter the distinguished name (DN) and password of the LDAP user.
3. If you extended the directory schema after installing the DB2 product, perform the following steps:
 - a. Register the current instance of the DB2 server in LDAP by running the following command:

```
db2 register ldap as node-name protocol tcpip
```
 - b. Register specific databases in LDAP by running the following command:

```
db2 catalog ldap database dbname as alias_dbname
```

You can now register the LDAP entries.

Registering LDAP entries

Registration of DB2 servers after installation

Each DB2 server instance must be registered in LDAP to publish the protocol configuration information that is used by the client applications to connect to the DB2 server instance.

When registering an instance of the database server, you need to specify a *node name*. The node name is used by client applications when they connect or attach to the server. You can catalog another alias name for the LDAP node by using the CATALOG LDAP NODE command.

Note: If you are working in a Windows domain environment, then during installation the DB2 server instance is automatically registered in the Active Directory with the following information:

```
nodename: TCP/IP hostname
protocol type: TCP/IP
```

If the TCP/IP hostname is longer than 8 characters, it will be truncated to 8 characters.

The REGISTER command appears as follows:

```
db2 register db2 server in ldap
as <ldap_node_name>
protocol tcpip
```

The protocol clause specifies the communication protocol to use when connecting to this database server.

When creating an instance for DB2 Enterprise Server Edition that includes multiple physical machines, the REGISTER command must be invoked once for each machine. Use the rah command to issue the REGISTER command on all machines.

Note: The same ldap_node_name cannot be used for each machine since the name must be unique in LDAP. You will want to substitute the hostname of each machine for the ldap_node_name in the REGISTER command. For example:

```
rah ">DB2 REGISTER DB2 SERVER IN LDAP AS <> PROTOCOL TCP/IP"
```

The "<>" is substituted by the hostname on each machine where the rah command is run. In the rare occurrence where there are multiple DB2 Enterprise Server Edition instances, the combination of the instance and host index may be used as the node name in the rah command.

The REGISTER command can be issued for a remote DB2 server. To do so, you must specify the remote computer name, instance name, and the protocol configuration parameters when registering a remote server. The command can be used as follows:

```
db2 register db2 server in ldap
as <ldap_node_name>
protocol tcpip
hostname <host_name>
svcname <tcpip_service_name>
remote <remote_computer_name>
instance <instance_name>
```

The following convention is used for the computer name:

- If TCP/IP is configured, the computer name must be the same as the TCP/IP hostname.

When running in a high availability or failover environment, and using TCP/IP as the communication protocol, the cluster IP address must be used. Using the cluster IP address allows the client to connect to the server on either machine without having to catalog a separate TCP/IP node for each machine. The cluster IP address is specified using the hostname clause, shown as follows:

```
db2 register db2 server in ldap
as <ldap_node_name>
protocol tcpip
hostname n.nn.nn.nn
```

where n.nn.nn.nn is the cluster IP address.

To register the DB2 server in LDAP from a client application, call the db2LdapRegister API.

Catalog a node alias for ATTACH

A node name for the DB2 server must be specified when registering the server in LDAP. Applications use the node name to attach to the database server.

If you require a different node name, such as when the node name is hard-coded in an application, use the CATALOG LDAP NODE command to make the change, for example:

```
db2 catalog ldap node <ldap_node_name>
as <new_alias_name>
```

To uncatalog a LDAP node, use the UNCATALOG LDAP NODE command, for example:

```
db2 uncatalog ldap node <ldap_node_name>
```

Registration of databases in the LDAP directory

During the creation of a database within an instance, the database is automatically registered in LDAP. Registration allows remote client connection to the database without having to catalog the database and node on the client machine. When a client attempts to connect to a database, if the database does not exist in the database directory on the local machine then the LDAP directory is searched.

If the name already exists in the LDAP directory, the database is still created on the local machine but a warning message is returned stating the naming conflict in the LDAP directory. For this reason you can manually catalog a database in the LDAP directory. The user can register databases on a remote server in LDAP by using the CATALOG LDAP DATABASE command. When registering a remote database, you specify the name of the LDAP node that represents the remote database server. You must register the remote database server in LDAP using the REGISTER DB2 SERVER IN LDAP command before registering the database. To register a database manually in LDAP, use the CATALOG LDAP DATABASE command:

```
db2 catalog ldap database <dbname>
at node <node_name>
with "My LDAP database"
```

To register a database in LDAP from a client application, call the db2LdapCatalogDatabase API.

Deregistering LDAP entries

Deregistering the DB2 server

Deregistration of an instance from LDAP also removes all the node, or alias, objects and the database objects referring to the instance.

Deregistration of the DB2 server on either a local or a remote machine requires the LDAP node name be specified for the server:

```
db2 deregister db2 server in ldap
node <node_name>
```

To deregister the DB2 server from LDAP from a client application, call the `db2LdapDeregister` API.

When the DB2 server is deregistered, any LDAP node entry and LDAP database entries referring to the same instance of the DB2 server are also uncataloged.

Deregistering the database from the LDAP directory

The database is automatically deregistered from LDAP when the database is dropped, or the owning instance is deregistered from LDAP.

You can manually deregister the database from LDAP using the following command:

```
db2 uncatalog ldap database <dbname>
```

To deregister a database from LDAP from a client application, call the `db2LdapUncatalogDatabase` API.

Configuring LDAP users

Creating an LDAP user

The DB2 database system supports setting DB2 registry variables and CLI configuration at the user level. (This is not available on the Linux and UNIX platforms.) User level support provides user-specific settings in a multi-user environment. An example is Windows Terminal Server where each logged on user can customize his or her own environment without interfering with the system environment or another user's environment.

When using the IBM Tivoli directory, you must define an LDAP user before you can store user-level information in LDAP. You can create an LDAP user by creating an LDIF file to contain all attributes for the user object, then run the LDIF import utility to import the object into the LDAP directory. The LDIF utility for the IBM Tivoli Directory Server is `LDIF2DB`.

LDIF file containing the attributes for a person object appears similar to the following:

```
File name: newuser.ldif

dn: cn=Mary Burnnet, ou=DB2 Development, ou=Toronto, o=ibm, c=ca
objectclass: ePerson
cn: Mary Burnnet
sn: Burnnet
uid: mburnnet
userPassword: password
telephonenumber: 1-416-123-4567
facsimiletelephonenumber: 1-416-123-4568
title: Software Developer
```

Following is an example of the LDIF command to import an LDIF file using the IBM LDIF import utility:

```
LDIF2DB -i newuser.ldif
```

Note:

1. You must run the LDIF2DB command from the LDAP server machine.
2. You must grant the required access (ACL) to the LDAP user object so that the LDAP user can add, delete, read, and write to his own object. To grant ACL for the user object, use the LDAP Directory Server Web Administration tool.

Configuring the LDAP user for DB2 applications

When you use the Microsoft LDAP client, the LDAP user is the same as the operating system user account. However, when you use the IBM LDAP client, before you use the DB2 database manager, you must configure the LDAP user distinguished name (DN) and password for the current logged on user.

To configure the LDAP user distinguished name (DN) and password, use the db2ldcfg utility:

```
db2ldcfg -u <userDN> -w <password> --> set the user's DN and password
-r --> clear the user's DN and password
```

For example:

```
db2ldcfg -u "cn=Mary Burnnet,ou=DB2 Development,ou=Toronto,o=ibm,c=ca"
-w password
```

Setting DB2 registry variables at the user level in the LDAP environment

Under the LDAP environment, the DB2 profile registry variables can be set at the user level which allows a user to customize their own DB2 environment.

To set the DB2 profile registry variables at the user level, use the -ul option:

```
db2set -ul <variable>=<value>
```

Note: This is not supported on AIX or Solaris operating systems.

DB2 has a caching mechanism. The DB2 profile registry variables at the user level are cached on the local machine. If the -ul parameter is specified, DB2 always reads from the cache for the DB2 registry variables. The cache is refreshed when:

- You update or reset a DB2 registry variable at the user level.
- The command to refresh the LDAP profile variables at the user level is:

```
db2set -ur
```

Disabling LDAP support

To disable LDAP support, use the following procedure:

1. For each instance of the DB2 server, deregister the DB2 server from LDAP:

```
db2 deregister db2 server in ldap node <nodename>
```
2. Set the DB2 profile registry variable DB2_ENABLE_LDAP to "NO".

Updating the protocol information for the DB2 server

The DB2 server information in LDAP must be kept current. For example, changes to the protocol configuration parameters or the server network address require an update to LDAP.

To update the DB2 server in LDAP on the local machine, use the following command:

```
db2 update ldap ...
```

Examples of protocol configuration parameters that can be updated include a TCP/IP hostname and service name or port number parameters.

To update a remote DB2 server protocol configuration parameters use the UPDATE LDAP command with a node clause:

```
db2 update ldap
node <node_name>
hostname <host_name>
svcname <tcpip_service_name>
```

Rerouting LDAP clients to another server

Just as with the ability to reroute clients on a system failure, the same ability is also available to you when working with LDAP.

The DB2_ENABLE_LDAP registry variable must be set to “Yes”.

Within an LDAP environment, all database and node directory information is maintained at an LDAP server. The client retrieves information from the LDAP directory. This information is updated in its local database and node directories if the DB2LDAPCACHE registry variable is set to “Yes”.

Use the UPDATE ALTERNATE SERVER FOR LDAP DATABASE command to define the alternate server for a database that represents the DB2 database in LDAP. Alternatively, you can call the db2LdapUpdateAlternateServerForDB API from a client application to update the alternate server for the database in LDAP.

Once established, this alternate server information is returned to the client upon connection.

Note: It is strongly recommended to keep the alternate server information stored in the LDAP server synchronized with the alternate server information stored at the database server instance. Issuing the UPDATE ALTERNATE SERVER FOR DATABASE command (notice that it is not “FOR LDAP DATABASE”) at the database server instance will help ensure synchronization between the database server instance and the LDAP server.

When you enter UPDATE ALTERNATE SERVER FOR DATABASE command at the server instance, and if LDAP support is enabled (DB2_ENABLE_LDAP=Yes) on the server, and if the LDAP user ID and password is cached (db2ldcfg was previously run), then the alternate server for the database is automatically, or implicitly, updated on the LDAP server. This works as if you entered UPDATE ALTERNATE SERVER FOR LDAP DATABASE explicitly.

If the UPDATE ALTERNATE SERVER FOR LDAP DATABASE command is issued from an instance other than the database server instance, ensure the alternate server information is also identically configured at the database server instance using the UPDATE ALTERNATE SERVER FOR DATABASE command. After the client initially connects to the database server instance, the alternate server information returned from the database server instance will take precedence over

what is configured in the LDAP server. If the database server instance has no alternate server information configured, client reroute will be considered disabled after the initial connect.

Attaching to a remote server in the LDAP environment

In the LDAP environment, you can attach to a remote database server using the LDAP node name on the ATTACH command: `db2 attach to <ldap_node_name>`.

When a client application attaches to a node or connects to a database for the first time, since the node is not in the local node directory, the database manager searches the LDAP directory for the target node entry. If the entry is found in the LDAP directory, the protocol information of the remote server is retrieved. If you connect to the database and if the entry is found in the LDAP directory, then the database information is also retrieved. Using this information, the database manager automatically catalogs a database entry and a node entry on the local machine. The next time the client application attaches to the same node or database, the information in the local database directory is used without having to search the LDAP directory.

In more detail: A caching mechanism exists so that the client only searches the LDAP server once. After the information is retrieved, it is stored or cached on the local machine based on the values of the *dir_cache* database manager configuration parameter and the DB2LDAPCACHE registry variable.

- If DB2LDAPCACHE=NO and dir_cache=NO, then always read the information from LDAP.
- If DB2LDAPCACHE=NO and dir_cache=YES, then read the information from LDAP once and insert it into the DB2(R) cache.
- If DB2LDAPCACHE=YES or is not set, then read the information from LDAP server once and cache it into the local database, node, and DCS directories.

Note: The caching of LDAP information is not applicable to user-level CLI or DB2 profile registry variables.

Refreshing LDAP entries in local database and node directories

The DB2 database system provides a caching mechanism to reduce the number of times a client searches the LDAP server.

After the information is retrieved, it is stored or cached on the local machine based on the values of the *dir_cache* database manager configuration parameter and the DB2LDAPCACHE registry variable.

- If DB2LDAPCACHE=NO and dir_cache=NO, then always read the information from LDAP.
- If DB2LDAPCACHE=NO and dir_cache=YES, then read the information from LDAP once and insert it into the DB2 cache.
- If DB2LDAPCACHE=YES or is not set, then read the information from LDAP server once and cache it into the local database, node, and DCS directories.

Note: The caching of LDAP information is not applicable to user-level CLI or DB2 profile registry variables. Since information in LDAP is subject to change, it may be necessary to refresh the LDAP entries cached in the local database and node directories. There are a few ways to do this.

To refresh all the local database and node entries that were retrieved from LDAP, use the following command:

```
db2 refresh ldap immediate
```

Similarly, the following command can be used to both refresh existing local database and node entries and add new entries from LDAP:

```
db2 refresh ldap immediate all
```

Specifying the IMMEDIATE ALL option will add all the NODE and DB entries contained with the LDAP server into the local directories.

Alternatively, to force DB2 to refresh the database entries that refer to LDAP resources on the next database connection or instance attachment, use the following command:

```
db2 refresh ldap database directory
```

Likewise, to force the DB2 database manager to refresh the nodes entries that refer to LDAP resources on the next database connection or instance attachment, use the following command:

```
db2 refresh ldap node directory
```

As part of the refresh, all the LDAP entries that are saved in the local database and node directories are removed. The next time that the application accesses the database or node, it will read the information directly from LDAP and generate a new entry in the local database or node directory.

To ensure the refresh is done in a timely way, you may want to:

- Schedule a refresh that is run periodically.
- Run the REFRESH command during system bootup.
- Use an available administration package to invoke the REFRESH command on all client machines.
- Set DB2LDAPCACHE="NO" to avoid LDAP information being cached in the database, node, and DCS directories.

Searching the LDAP servers

The DB2 database system searches the current LDAP server but in an environment where there are multiple LDAP servers, you can define the scope of the search.

For example, if the information is not found in the current LDAP server, you can specify automatic search of all other LDAP servers, or, alternatively, you can restrict the search scope to only the current LDAP server, or to the local DB2 database catalog.

When you set the search scope, this sets the default search scope for the entire enterprise. The search scope is controlled through the DB2 database profile registry variable, DB2LDAP_SEARCH_SCOPE. To set the search scope value, use the -gl option, which means global in LDAP, on the db2set command:

```
db2set -gl db2ldap_search_scope=<value>
```

Possible values include: local, domain, or global. If it is not set, the default value is domain which limits the search scope to the directory on the current LDAP server.

For example, you may want to initially set the search scope to “global” after a new database is created. This allows any DB2 client configured to use LDAP to search all the LDAP servers to find the database. Once the entry has been recorded on each machine after the first connect or attach for each client, if you have caching enabled, the search scope can be changed to “local”. Once changed to “local”, each client will not scan any LDAP servers.

Note: The DB2 database profile registry variables DB2LDAP_KEEP_CONNECTION and DB2LDAP_SEARCH_SCOPE are the only registry variables that can be set at the global level in LDAP.

Part 2. Databases

Chapter 6. Databases

A DB2 database is a *relational database*. The *database* stores all data in tables that are related to one another. Relationships are established between tables such that data is shared and duplication is minimized.

A *relational database* is a database that is treated as a set of tables and manipulated in accordance with the relational model of data. It contains a set of objects used to store, manage, and access data. Examples of such objects are tables, views, indexes, functions, triggers, and packages. Objects can be either defined by the system (system-defined objects) or defined by the user (user-defined objects).

A *distributed relational database* consists of a set of tables and other objects that are spread across different but interconnected computer systems. Each computer system has a relational database manager to manage the tables in its environment. The database managers communicate and cooperate with each other in a way that allows a given database manager to execute SQL statements on another computer system.

A *partitioned relational database* is a relational database whose data is managed across multiple database partitions. This separation of data across database partitions is transparent to users of most SQL statements. However, some data definition language (DDL) statements take database partition information into consideration (for example, CREATE DATABASE PARTITION GROUP). DDL is the subset of SQL statements used to describe data relationships in a database.

A *federated database* is a relational database whose data is stored in multiple data sources (such as separate relational databases). The data appears as if it were all in a single large database and can be accessed through traditional SQL queries. Changes to the data can be explicitly directed to the appropriate data source.

Designing databases

When designing a database, you are modeling a real business system that contains a set of entities and their characteristics, or *attributes*, and the rules or relationships between those entities.

The first step is to describe the system that you want to represent. For example, if you were creating a database for publishing system, the system would contain several types of entities, such as books, authors, editors, and publishers. For each of these entities, there are certain pieces of information, or attributes, that you need to record:

- *Books*: titles, ISBN, date published, location, publisher,
- *Authors*: name, address, phone and fax numbers, e-mail address,
- *Editors*: name, address, phone and fax numbers, e-mail address,
- *Publishers*: name, address, phone and fax numbers, e-mail address,

You will need the database to represent not only these types of entities and their attributes, but you also need a way to relate these entities to each other. For example, you need to represent the relationship between books and their authors, the relationship between books/authors and editors, and the relationship between books/authors and publishers.

There are three types of relationships between the entities in a database:

One-to-one relationships

In this type of relationship, each instance of an entity relates to only one instance of another entity. Currently, no one-to-one relationships exist in the scenario described above.

One-to-many relationships

In this type of relationship, each instance of an entity relates to one or more instances of another entity. For example, an author could have written multiple books, but certain books have only one author. This is the most common type of relationship modeled in relational databases.

Many-to-many relationships

In this type of relationship, many instances of a given entity relate to one or more instances of another entity. For example, co-authors could write a number of books.

Because databases consist of tables, you need to construct a set of tables that will best hold this data, with each cell in the table holding a single view. There are many possible ways to perform this task. As the database designer, your job is to come up with the best set of tables possible.

For example, you could create a single table, with many rows and columns, to hold all of the information. However, using this method, some information would be repeated. Secondly, data entry and data maintenance would be time-consuming and error prone. In contrast to this single-table design, a *relational database* allows you to have multiple simple tables, reducing redundancy and avoiding the difficulties posed by a large and unmanageable table. In a relational database, tables should contain information about a single type of entity.

Also, the integrity of the data in a relational database must be maintained as multiple users access and change the data. Whenever data is shared, there is a need to ensure the accuracy of the values within database tables.

You can:

- Use isolation levels to determine how data is locked or isolated from other processes while the data is being accessed.
- Protect data and define relationships between data by defining constraints to enforce business rules.
- Create triggers that can do complex, cross-table data validation.
- Implement a recovery strategy to protect data so that it can be restored to a consistent state.

Database design is a much more complex task than is indicated here, and there are many items that need to be considered, such as space requirements, keys, indexes, constraints, security and authorization, and so forth. You can find some of this information in the DB2 Information Center, and in the many DB2 retail books that are available on this subject.

Database directories and files

When you create a database, information about the database including default information is stored in a directory hierarchy.

The hierarchical directory structure is created for you at a location that is determined by the information you provide in the CREATE DATABASE command.

If you do not specify the location of the directory path or drive when you create the database, the default location is used. It is suggested that you explicitly state where you would like the database created.

In the directory you specify as the database path in the CREATE DATABASE command, a subdirectory that uses the name of the instance is created. This subdirectory ensures that databases created in different instances under the same directory do not use the same path. Below the instance-name subdirectory, a subdirectory named NODE0000 is created. This subdirectory differentiates database partitions in a logically partitioned database environment. Below the node-name directory, a subdirectory named SQL00001 is created. This name of this subdirectory uses the database token and represents the database being created. SQL00001 contains objects associated with the first database created, and subsequent databases are given higher numbers: SQL00002, and so on. These subdirectories differentiate databases created in this instance on the directory that you specified in the CREATE DATABASE command.

The directory structure appears as follows:*your_database_path/your_instance/NODE0000/SQL00001/*.

The database directory contains the following files that are created as part of the CREATE DATABASE command.

- The files SQLBP.1 and SQLBP.2 contain buffer pool information. These files are duplicates of each other for backup purposes.
- The files SQLSPCS.1 and SQLSPCS.2 contain table space information. These files are duplicates of each other for backup purposes.
- The files SQLSGF.1 and SQLSGF.2 contain storage path information associated with the database's automatic storage. These files are duplicates of each other for backup purposes.
- The SQLDBCONF file contains database configuration information. Do not edit this file.

Note: The SQLDBCON file was used in previous releases and contains similar information that can be used if SQLDBCONF is corrupted.

To change configuration parameters, use the UPDATE DATABASE CONFIGURATION and RESET DATABASE CONFIGURATION commands.

- The DB2RHIST.ASC history file and its backup DB2RHIST.BAK contain history information about backups, restores, loading of tables, reorganization of tables, altering of a table space, and other changes to a database.

The DB2TSCHG.HIS file contains a history of table space changes at a log-file level. For each log file, DB2TSCHG.HIS contains information that helps to identify which table spaces are affected by the log file. Table space recovery uses information from this file to determine which log files to process during table space recovery. You can examine the contents of both history files in a text editor.

- The log control files, SQLOGCTL.LFH.1, its mirror copy SQLOGCTL.LFH.2, and SQLOGMIR.LFH, contain information about the active logs.

Recovery processing uses information from these files to determine how far back in the logs to begin recovery. The SQLOGDIR subdirectory contains the actual log files.

Note: You should ensure the log subdirectory is mapped to different disks than those used for your data. A disk problem could then be restricted to your data or the logs but not both. This can provide a substantial performance benefit

because the log files and database containers do not compete for movement of the same disk heads. To change the location of the log subdirectory, change the **newlogpath** database configuration parameter.

- The SQLINSLK file helps to ensure that a database is used by only one instance of the database manager.

At the same time a database is created, a detailed deadlocks event monitor is also created. The detailed deadlocks event monitor files are stored in the database directory of the catalog node. When the event monitor reaches its maximum number of files to output, it will deactivate and a message is written to the notification log. This prevents the event monitor from consuming too much disk space. Removing output files that are no longer needed will allow the event monitor to activate again on the next database activation.

Additional information for SMS database directories in non-automatic storage databases

In non-automatic storage databases, the SQLT* subdirectories contain the default System Managed Space (SMS) table spaces required for an operational database. Three default table spaces are created:

- SQLT0000.0 subdirectory contains the catalog table space with the system catalog tables.
- SQLT0001.0 subdirectory contains the default temporary table space.
- SQLT0002.0 subdirectory contains the default user data table space.

Each subdirectory or container has a file created in it called SQLTAG.NAM. This file marks the subdirectory as being in use so that subsequent table space creation does not attempt to use these subdirectories.

In addition, a file called SQL*.DAT stores information about each table that the subdirectory or container contains. The asterisk (*) is replaced by a unique set of digits that identifies each table. For each SQL*.DAT file there might be one or more of the following files, depending on the table type, the reorganization status of the table, or whether indexes, LOB, or LONG fields exist for the table:

- SQL*.BKM (contains block allocation information if it is an MDC table)
- SQL*.LF (contains LONG VARCHAR or LONG VARGRAPHIC data)
- SQL*.LB (contains BLOB, CLOB, or DBCLOB data)
- SQL*.XDA (contains XML data)
- SQL*.LBA (contains allocation and free space information about SQL*.LB files)
- SQL*.INX (contains index table data)
- SQL*.IN1 (contains index table data)
- SQL*.DTR (contains temporary data for a reorganization of an SQL*.DAT file)
- SQL*.LFR (contains temporary data for a reorganization of an SQL*.LF file)
- SQL*.RLB (contains temporary data for a reorganization of an SQL*.LB file)
- SQL*.RBA (contains temporary data for a reorganization of an SQL*.LBA file)

Database configuration file

A *database configuration file* is created for each database. This file is called SQLDBCON prior to Version 8.2, and SQLDBCONF in Version 8.2 and later. The creation of this file is done for you.

This file contains values for various *configuration parameters* that affect the use of the database, such as:

- Parameters specified or used when creating the database (for example, database code page, collating sequence, DB2 database release level)
- Parameters indicating the current state of the database (for example, backup pending flag, database consistency flag, roll-forward pending flag)
- Parameters defining the amount of system resources that the operation of the database might use (for example, buffer pool size, database logging, sort memory size).

Note: If you edit the `db2system`, `SQLDBCON` (prior to Version 8.2), or `SQLDBCONF` (Version 8.2 and later) file using a method other than those provided by the DB2 database manager, you might make the database unusable. Therefore, do not change these files using methods other than those documented and supported by the database manager.

Performance Tip: Many of the configuration parameters come with default values, but might need to be updated to achieve optimal performance for your database. By default, the Configuration Advisor is invoked as part of the `CREATE DATABASE` command so that the initial values for some parameters are already configured for your environment.

For multi-partition databases: When you have a database that is distributed across more than one database partition, the configuration file should be the same on all database partitions. Consistency is required since the query compiler compiles distributed SQL statements based on information in the local node configuration file and creates an access plan to satisfy the needs of the SQL statement. Maintaining different configuration files on database partitions could lead to different access plans, depending on which database partition the statement is prepared.

Node directory

The database manager creates the *node directory* when the first database partition is cataloged.

To catalog a database partition, use the `CATALOG NODE` command. To list the contents of the local node directory, use the `LIST NODE DIRECTORY` command.

The node directory is created and maintained on each database client. The directory contains an entry for each remote workstation having one or more databases that the client can access. The DB2 client uses the communication end point information in the node directory whenever a database connection or instance attachment is requested.

The entries in the directory also contain information on the type of communication protocol to be used to communicate from the client to the remote database partition. Cataloging a local database partition creates an alias for an instance that resides on the same computer.

Local database directory

A *local database directory* file exists in each path (or “drive” for Windows operating systems) in which a database has been defined. This directory contains one entry for each database accessible from that location.

Each entry contains:

- The database name provided with the CREATE DATABASE command
- The database alias name (which is the same as the database name, if an alias name is not specified)
- A comment describing the database, as provided with the CREATE DATABASE command
- The name of the root directory for the database
- Other system information.

System database directory

A *system database directory* file exists for each instance of the database manager, and contains one entry for each database that has been cataloged for this instance.

Databases are implicitly cataloged when the CREATE DATABASE command is issued and can also be explicitly cataloged with the CATALOG DATABASE command.

For each database created, an entry is added to the directory containing the following information:

- The database name provided with the CREATE DATABASE command
- The database alias name (which is the same as the database name, if an alias name is not specified)
- The database comment provided with the CREATE DATABASE command
- The location of the local database directory
- An indicator that the database is *indirect*, which means that it resides on the current database manager instance
- Other system information.

On UNIX platforms and in a partitioned database environment, you must ensure that all database partitions always access the same system database directory file, `sqlbdir`, in the `sqlbdir` subdirectory of the home directory for the instance. Unpredictable errors can occur if either the system database directory or the system intention file `sqlbins` in the same `sqlbdir` subdirectory are symbolic links to another file that is on a shared file system.

Creating node configuration files

If your database is to operate in a partitioned database environment, you must create a node configuration file called `db2nodes.cfg`.

This file must be located in the `sqllib` subdirectory of the home directory for the instance before you can start the database manager with parallel capabilities across multiple database partitions. The file contains configuration information for all database partitions in an instance, and is shared by all database partitions for that instance.

Windows considerations

If you are using DB2 Enterprise Server Edition on Windows, the node configuration file is created for you when you create the instance. You should not attempt to create or modify the node configuration file manually. You can use the `db2nrcrt` command to add a database partition server to an instance. You can use the `db2ndrop` command to drop a database partition server from an instance. You can use the `db2nchg` command to modify a database partition server configuration

including moving the database partition server from one computer to another; changing the TCP/IP host name; or, selecting a different logical port or network name.

Note: You should not create files or directories under the `sqllib` subdirectory other than those created by the database manager to prevent the loss of data if an instance is deleted. There are two exceptions. If your system supports stored procedures, put the stored procedure applications in the `function` subdirectory under the `sqllib` subdirectory. The other exception is when user-defined functions (UDFs) have been created. UDF executables are allowed in the same directory.

The file contains one line for each database partition that belongs to an instance. Each line has the following format:

```
dbpartitionnum hostname [logical-port [netname]]
```

Tokens are delimited by blanks. The variables are:

dbpartitionnum

The database partition number, which can be from 0 to 999, uniquely defines a database partition. Database partition numbers must be in ascending sequence. You can have gaps in the sequence.

Once a database partition number is assigned, it cannot be changed. (Otherwise the information in the distribution map, which specifies how data is distributed, would be compromised.)

If you drop a database partition, its database partition number can be used again for any new database partition that you add.

The database partition number is used to generate a database partition name in the database directory. It has the format:

```
NODE nnnn
```

The *nnnn* is the database partition number, which is left-padded with zeros. This database partition number is also used by the `CREATE DATABASE` and `DROP DATABASE` commands.

hostname

The hostname of the IP address for inter-partition communications. Use the fully-qualified name for the hostname. The `/etc/hosts` file also should use the fully-qualified name. If the fully-qualified name is not used in the `db2nodes.cfg` file and in the `/etc/hosts` file, you might receive error message `SQL30082N RC=3`.

(There is an exception when `netname` is specified. In this situation, `netname` is used for most communications, with `hostname` only being used for `db2start`, `db2stop`, and `db2_all`.)

logical-port

This parameter is optional, and specifies the logical port number for the database partition. This number is used with the database manager instance name to identify a TCP/IP service name entry in the `etc/services` file.

The combination of the IP address and the logical port is used as a well-known address, and must be unique among all applications to support communications connections between database partitions.

For each hostname, one *logical-port* must be either 0 (zero) or blank (which defaults to 0). The database partition associated with this *logical-port* is the

default node on the host to which clients connect. You can override this with the **DB2NODE** environment variable in db2profile script, or with the `sqlsetc()` API.

netname

This parameter is optional, and is used to support a host that has more than one active TCP/IP interface, each with its own hostname.

The following example shows a possible node configuration file for an RS/6000® SP™ system on which SP2EN1 has multiple TCP/IP interfaces, two logical partitions, and uses SP2SW1 as the DB2 database interface. It also shows the database partition numbers starting at 1 (rather than at 0), and a gap in the *dbpartitionnum* sequence:

Table 40. Database partition number example table.

<i>dbpartitionnum</i>	<i>hostname</i>	<i>logical-port</i>	<i>netname</i>
1	SP2EN1.mach1.xxx.com	0	SP2SW1
2	SP2EN1.mach1.xxx.com	1	SP2SW1
4	SP2EN2.mach1.xxx.com	0	
5	SP2EN3.mach1.xxx.com		

You can update the db2nodes.cfg file using an editor of your choice. (The exception is: an editor should not be used on Windows.) You must be careful, however, to protect the integrity of the information in the file, as database partitioning requires that the node configuration file is locked when you issue db2start and unlocked after db2stop ends the database manager. The db2start command can update the file, if necessary, when the file is locked. For example, you can issue db2start with the **RESTART** option or the **ADD DBPARTITIONNUM** option.

Note: If the db2stop command is not successful and does not unlock the node configuration file, issue db2stop **FORCE** to unlock it.

Changing node and database configuration files

To update the database configuration file, run the AUTOCONFIGURE command with the appropriate options.

The Configuration Advisor helps you to tune performance and to balance memory requirements for a single database per instance by suggesting which configuration parameters to modify and providing suggested values for them.

If you plan to change any database partition groups (adding or deleting database partitions, or moving existing database partitions), the node configuration file must be updated. If you plan to change the database, you should review the values for the configuration parameters. You can adjust some values periodically as part of the ongoing changes made to the database that are based on how it is used.

Note: If you modify any parameters, the values are not updated until:

- For database parameters, the first new connection to the database after all applications are disconnected
- For database manager parameters, the next time that you stop and start the instance

In most cases, the values recommended by the Configuration Advisor will provide better performance than the default values because they are based on information about your workload and your own particular server. However, the values are designed to improve the performance of, though not necessarily optimize, your database system. Think of the values as a starting point on which you can make further adjustments to obtain optimized performance.

In Version 9.1, the Configuration Advisor is automatically invoked when you create a database. To disable this feature, or to explicitly enable it, use the `db2set` command before creating the database. Examples:

```
db2set DB2_ENABLE_AUTOCONFIG_DEFAULT=NO
db2set DB2_ENABLE_AUTOCONFIG_DEFAULT=YES
```

See “Automatic features” on page 17 for other features that are enabled by default.

To use the Configuration Advisor from the command line, use the `AUTOCONFIGURE` command.

To update individual parameters in the database manager configuration using the command line, enter:

```
UPDATE DBM CFG USING <config_keyword>=<value>
```

To update individual parameters in the database configuration using the command line, enter:

```
UPDATE DB CFG FOR <database_alias>
USING <config_keyword>=<value>
```

You can update one or more `<config_keyword>=<value>` combinations in a single command. Most changes to the database manager configuration file become effective only after they are loaded into memory. For a server configuration parameter, this occurs during the running of the `START DATABASE MANAGER` command. For a client configuration parameter, this occurs when the application is restarted.

To view or print the current database manager configuration parameters, use the `GET DATABASE MANAGER CONFIGURATION` command.

To access the Configuration Advisor from a client application, call the `db2AutoConfig` API. To update individual parameters in the database manager configuration or database configuration file from a client application, call the `db2CfgSet` API.

Database recovery log

A *database recovery log* keeps a record of all changes made to a database, including the addition of new tables or updates to existing ones.

This log is made up of a number of *log extents*, each contained in a separate file called a *log file*.

The database recovery log can be used to ensure that a failure (for example, a system power outage or application error) does not leave the database in an inconsistent state. In case of a failure, the changes already made but not committed are rolled back, and all committed transactions, which might not have been physically written to disk, are redone. These actions ensure the integrity of the database.

Space requirements for database objects

Estimating the size of database objects is an imprecise undertaking. Overhead caused by disk fragmentation, free space, and the use of variable length columns makes size estimation difficult, because there is such a wide range of possibilities for column types and row lengths.

After initially estimating your database size, create a test database and populate it with representative data. Then use the db2look utility to generate data definition statements for the database.

When estimating the size of a database, the contribution of the following must be considered:

- System catalog tables
- User table data
- Long field (LF) data
- Large object (LOB) data
- Index space
- Log file space
- Temporary work space

Also consider the overhead and space requirements for the following:

- The local database directory file
- The system database directory file
- The file management overhead required by the operating system, including:
 - File block size
 - Directory control space

Space requirements for log files

You will require 56 KB of space for log control files.

You will also need at least enough space for your active log configuration, which you can calculate as

$$(\text{logprimary} + \text{logsecond}) * (\text{logfilsiz} + 2) * 4096$$

where:

- *logprimary* is the number of primary log files, defined in the database configuration file
- *logsecond* is the number of secondary log files, defined in the database configuration file; in this calculation, *logsecond* cannot be set to -1. (When *logsecond* is set to -1, you are requesting an infinite active log space.)
- *logfilsiz* is the number of pages in each log file, defined in the database configuration file
- 2 is the number of header pages required for each log file
- 4096 is the number of bytes in one page.

If the database is enabled for circular logging, the result of this formula is all the space that will be allocated for logging; that is, more space will not be allocated, and you will not receive insufficient disk space errors for any of your log files.

If the database is enabled for roll-forward recovery, special log space requirements should be taken into consideration:

- With the *logarchmeth1* configuration parameter set to *logretain*, the log files will be archived in the log path directory. The online disk space will eventually fill up, unless you move the log files to a different location.
- With the *logarchmeth1* configuration parameter set to *userexit*, *DISK*, or *VENDOR*, a user exit program moves the archived log files to a different location. Extra log space is still required to allow for:
 - Online archived logs that are waiting to be moved by the user exit program
 - New log files being formatted for future use

If the database is enabled for infinite logging (that is, you set *logsecond* to *-1*), the *logarchmeth1* configuration parameter must be set to a value other than *OFF* or *LOGRETAIN* to enable archive logging. The database manager will keep at least the number of active log files specified by *logprimary* in the log path, so you should not use the value of *-1* for *logsecond* in the above formula. Ensure that you provide extra disk space to allow for the delay caused by archiving log files.

If you are mirroring the log path, you will need to double the estimated log file space requirements.

Lightweight Directory Access Protocol (LDAP) directory service

A directory service is a repository of resource information about multiple systems and services within a distributed environment; and it provides client and server access to these resources.

Clients and servers would use the directory service to find out how to access other resources. Information about these other resources in the distributed environment must be entered into the directory service repository.

Lightweight Directory Access Protocol (LDAP) is an industry standard access method to directory services. Each database server instance will publish its existence to an LDAP server and provide database information to the LDAP directory when the databases are created. When a client connects to a database, the catalog information for the server can be retrieved from the LDAP directory. Each client is no longer required to store catalog information locally on each computer. Client applications search the LDAP directory for information required to connect to the database.

Note: Publishing of the database server instance to the LDAP server is not an automatic process, but must be done manually by the administrator.

As an administrator of a DB2 system, you can establish and maintain a directory service. The Configuration Assistant can assist in the maintenance of this directory service. The directory service is made available to the DB2 database manager through Lightweight Directory Access Protocol (LDAP) directory services. To use LDAP directory services, there must first exist an LDAP server that is supported by the DB2 database manager so that directory information can be stored there.

Note: When running in a Windows domain environment, an LDAP server is already available because it is integrated with the Windows Active Directory. As a result, every computer running Windows can use LDAP.

An LDAP directory is helpful in an enterprise environment where it is difficult to update local directory catalogs on each client computer because of the large

number of clients. In this situation, you should consider storing your directory entries in an LDAP server so that maintaining catalog entries is done in one place: on the LDAP server.

Creating databases

You can create a database using the CREATE DATABASE command. To create a database from a client application, call the sqlecrea API.

You should have spent sufficient time designing the contents, layout, potential growth, and use of your database before you create it.

The following database privileges are automatically granted to PUBLIC: CREATETAB, BINDADD, CONNECT, IMPLICIT_SCHEMA, and SELECT on the system catalog views. However, if the RESTRICTIVE option is present, no privileges are automatically granted to PUBLIC. For more information on the RESTRICTIVE option, see the CREATE DATABASE command.

When you create a database, each of the following tasks are done for you:

- Setting up of all the system catalog tables that are needed by the database
- Allocation of the database recovery log
- Creation of the database configuration file and the default values are set
- Binding of the database utilities to the database

To create a database using the command line processor, enter:

```
CREATE DATABASE <database name>
```

For example, the following command creates a database called PERSON1, in the default location, with the associated comment "Personnel DB for BSchiefer Co".

```
CREATE DATABASE person1  
WITH "Personnel DB for BSchiefer Co"
```

Configuration Advisor

The Configuration Advisor helps you to tune performance and to balance memory requirements for a single database per instance by suggesting which configuration parameters to modify and providing suggested values for them. The Configuration Advisor is automatically invoked when you create a database. To disable this feature, or to explicitly enable it, use the db2set command before creating the database. Examples:

```
db2set DB2_ENABLE_AUTOCONFIG_DEFAULT=NO  
db2set DB2_ENABLE_AUTOCONFIG_DEFAULT=YES
```

See "Automatic features" on page 17 for other features that are enabled by default.

Event Monitor

At the same time a database is created, a detailed deadlocks event monitor is also created. As with any monitor, there is some overhead associated with this event monitor. If you do not want the detailed deadlocks Event Monitor, then the Event Monitor can be dropped using the command:

```
DROP EVENT MONITOR db2detaildeadlock
```

To limit the amount of disk space that this event monitor consumes, the event monitor deactivates, and a message is written to the administration notification log, once it has reached its maximum number of output files.

Removing output files that are no longer needed allows the event monitor to activate again on the next database activation.

Remote databases

You have the ability to create a database in a different, possibly remote, instance. To create a database at another (remote) database partition server, you must first attach to that server. A database connection is temporarily established by the following command during processing:

```
CREATE DATABASE <database name> AT DBPARTITIONNUM <options>
```

In this type of environment you can perform instance-level administration against an instance other than your default instance, including remote instances. For instructions on how to do this, see the `db2iupdt` (update instance) command.

Database code pages

By default, databases are created in the UTF-8 (Unicode) code set.

To override the default code page for the database, it is necessary to specify the desired code set and territory when creating the database. See the `CREATE DATABASE` command or the `sqlcrea` API for information on setting the code set and territory.

Automatic storage databases

The database manager creates all databases as “automatic storage” databases, by default. To create a database that is not an “automatic storage” database, specify `AUTOMATIC STORAGE NO` when issuing the `CREATE DATABASE` command.

Databases that are enabled for automatic storage have a set of one or more storage paths associated with them. A table space can be defined as *managed by automatic storage* and its containers assigned and allocated by the database manager based on those storage paths.

You can enable a database for automatic storage only when you create it; similarly, you cannot disable automatic storage for a database that was originally designed to use it.

All databases are created as automatic storage databases by default. To create a database that is not an automatic storage database, specify **AUTOMATIC STORAGE NO** when issuing the `CREATE DATABASE` command.

Examples of disabling automatic storage:

```
CREATE DATABASE ASNOB1 AUTOMATIC STORAGE NO
CREATE DATABASE ASNOB2 AUTOMATIC STORAGE NO ON X:
```

Examples of automatic storage being enabled either explicitly or implicitly:

```
CREATE DATABASE DB1
CREATE DATABASE DB2 AUTOMATIC STORAGE YES ON X:
CREATE DATABASE DB3 ON /data/path1, /data/path2
CREATE DATABASE DB4 ON D:\StoragePath DBPATH ON C:
```

Based on the syntax used, the database manager extracts the following two pieces of information that pertain to storage locations:

- The database path (where the database manager stores various control files for the database):
 - If you specify **DBPATH ON**, this indicates the database path.

- If you do not specify **DBPATH ON**, the first path listed in **ON** indicates the database path (and the storage path).
- If you specify neither **DBPATH ON** nor **ON**, the **dftdbpath** database manager configuration parameter is used to determine the database path.
- The storage paths (where the database manager creates automatic storage table space containers):
 - If you specify **ON**, all of the listed paths are storage paths.
 - If you do not specify **ON**, there is a single storage path that is set to the value of the **dftdbpath** database manager configuration parameter.

For the examples shown previously, the following table summarizes the database and storage paths used:

Table 41. Automatic storage database and storage paths

CREATE DATABASE command	Database path	Storage paths
CREATE DATABASE DB1 AUTOMATIC STORAGE YES	Value of the dftdbpath configuration parameter	Value of the dftdbpath configuration parameter
CREATE DATABASE DB2 AUTOMATIC STORAGE YES ON X:	X:	X:
CREATE DATABASE DB3 ON /data/path1, /data/path2	/data/path1	/data/path1, /data/path2
CREATE DATABASE DB4 ON D:\StoragePath DBPATH ON C:	C:	D:\StoragePath

The storage paths provided must exist and be accessible. In a partitioned database environment, the same storage paths are used on each database partition. You cannot specify a unique set of storage paths for a particular database partition unless you use database partition expressions as part of the storage path names. Doing this allows the database partition number to be reflected in the storage paths such that the resulting path names are different on each database partition.

Use the argument **\$N** (that is, **\$N** preceded by a blank) to indicate a database partition expression. You can use a database partition expression anywhere in the storage path, and you can specify multiple database partition expressions. Terminate the database partition expression with a space character; whatever follows the space is appended to the storage path after the database partition expression is evaluated. If there is no space character in the storage path after the database partition expression, it is assumed that the rest of the string is part of the expression. The following table lists the only valid forms of the **\$N** argument. Operators are evaluated from left to right, and **%** represents the modulus operator. The database partition number in the examples is 10.

Table 42. Database partition expressions

Syntax	Example	Value
[blank]\$N	" \$N"	10
[blank]\$N+[number]	" \$N+100"	110
[blank]\$N%[number]	" \$N%5"	0
[blank]\$N+[number]%[number]	" \$N+1%5"	1
[blank]\$N%[number]+[number]	" \$N%4+2"	4

The following is an example of using database partition expressions:

```
CREATE DATABASE TESTDB ON "/path1ForNode $N",
"/path2ForNode $N" DBPATH ON "/dbpathForNodes"
```

The following is an example of a database partition expression embedded in the middle of a path:

```
CREATE DATABASE TESTDB ON "/path1ForNode $N",
"/path2ForNode $N suffix" DBPATH ON "/dbpathForNodes"
```

Note: Database partition expressions are not valid in database paths, whether you specify them explicitly in **DBPATH ON** or implicitly by using a database partition expression in the first storage path.

When free space is calculated for a storage path for a given database partition, the database manager checks for the existence of the following directories or mount points within the storage path and uses the first one that it finds:

```
storage path/instance name/NODE####/database name
storage path/instance name/NODE####
storage path/instance name
storage path
```

where:

storage path

Is a storage path associated with the database

instance name

Is the instance under which the database resides

NODE####

Is the database partition number (for example, NODE0000 or NODE0001)

database name

Is the name of the database

File systems can be mounted at a point beneath the storage path, and the database manager recognizes that the actual amount of free space available for table space containers might not be the same amount that is associated with the storage path directory itself.

Consider the example where two logical database partitions exist on one physical computer and there is a single storage path: /db2data. Each database partition can use this storage path, but you might want to isolate the data from each database partition by creating for each one a separate file system. The file system is mounted at /db2data/*instance*/NODE####. When creating containers on the storage path and determining free space, the database manager does not retrieve free space information for /db2data but instead retrieves it for the corresponding /db2data/*instance*/NODE#### directory.

There are three default table spaces created whenever you create a database. If you do not provide explicit table space definitions as part of the CREATE DATABASE command, the table spaces are created as automatic storage table spaces.

After you create a database, you can add new storage paths to it by using the ADD STORAGE clause of the ALTER DATABASE statement, as shown in the following example:

```
ALTER DATABASE ADD STORAGE ON '/data/path3', '/data/path4'
```

Automatic storage restrictions

When deciding whether to create a database using automatic storage, there are some restrictions to consider.

- You cannot disable or enable automatic storage for a database after you created it.
- Storage paths must be absolute path names. They can be paths or drive letters on the Windows operating system. The database path must be a drive letter. The maximum path length is 175 characters.
- For partitioned databases, you must use the same set of storage paths on each database partition (unless you use database partition expressions).
- Database partition expressions are not valid in database paths, whether you specify them explicitly by using the **DBPATH ON** option of the **CREATE DATABASE** command, or implicitly by using a database partition expression in the first storage path.

Adding automatic storage paths to databases enabled for automatic storage

Using the **ALTER DATABASE** statement, you can add an automatic storage path to a database that is enabled for automatic storage. You can enable a database for automatic storage only when you create it.

When you add a storage path for a multi-partition database environment, the storage path must exist on each database partition. If the specified path does not exist on every database partition, the statement is rolled back.

To add a storage path to an existing database, issue the following **ALTER DATABASE** statement:

```
ALTER DATABASE PATH pathname
```

Monitoring storage paths

A database snapshot includes the list of storage paths associated with the database.

If the number of automatic storage paths is 0, automatic storage is not enabled for the database:

```
Number of automatic storage paths = ##  
Automatic storage path           = <1st path>  
Automatic storage path           = <2nd path>  
...
```

If the bufferpool monitor switch is on, the following elements are also set:

```
File system ID                   = 12345  
File system free space (bytes)    = 20000000000  
File system used space (bytes)    = 4000000000000  
File system total space (bytes)   = 4002000000000
```

This data is set on a per path basis: on a single database partition system per path, and per each database partition on a multi-database partitioned environment.

In addition, the following information is set within a table space snapshot. The information indicates whether or not the table space was created as an automatic storage table space:

```
Using automatic storage          = Yes or No
```

Restore database implications

The RESTORE DATABASE command is used to restore a database from a backup image.

During a restore operation it is possible to choose the location of the database path and its also possible to redefine the storage paths that are associated with the database. The database path and the storage paths are set by using a combination of TO, ON, and DBPATH ON with the RESTORE DATABASE command.

For example, here are some valid RESTORE commands for databases enabled for automatic storage:

```
RESTORE DATABASE TEST1
RESTORE DATABASE TEST2 TO X:
RESTORE DATABASE TEST3 DBPATH ON D:
RESTORE DATABASE TEST3 ON /path1, /path2, /path3
RESTORE DATABASE TEST4 ON E:\newpath1, F:\newpath2 DBPATH ON D:
```

Like the CREATE DATABASE command, the database manager extracts the following two pieces of information that pertain to storage locations:

- The **database path** (which is where the database manager stores various control files for the database)
 - If TO or DBPATH ON is specified, this indicates the database path.
 - If ON is used but DBPATH ON is not specified with it, the first path listed with ON is used as the database path (in addition to it being a storage path).
 - If none of TO, ON, or DBPATH ON are specified, the *dftdbpath* database manager configuration parameter determines the database path.

Note: If a database with the same name exists on disk, the database path is ignored, and the database is placed into the same location as the existing database.

- The **storage paths** (where the database manager creates automatic storage table space containers)
 - If ON is specified, all of the paths listed are considered storage paths, and these paths are used instead of the ones stored within the backup image.
 - If ON is not specified, no change is made to the storage paths (the storage paths stored within the backup image are maintained).

To make this concept clearer, the same five RESTORE command examples presented above are shown in the following table with their corresponding storage paths:

Table 43. Restore implications regarding database and storage paths

RESTORE DATABASE command	No database on disk exists with same name		Database exists on disk with same name	
	Database path	Storage paths	Database path	Storage paths
RESTORE DATABASE TEST1	<dftdbpath>	Uses storage paths defined in the backup image	Uses database path of existing database	Uses storage paths defined in the backup image
RESTORE DATABASE TEST2 TO X:	X:	Uses storage paths defined in the backup image	Uses database path of existing database	Uses storage paths defined in the backup image

Table 43. Restore implications regarding database and storage paths (continued)

RESTORE DATABASE command	No database on disk exists with same name		Database exists on disk with same name	
	Database path	Storage paths	Database path	Storage paths
RESTORE DATABASE TEST3 DBPATH ON /db2/databases	/db2/databases	Uses storage paths defined in the backup image	Uses database path of existing database	Uses storage paths defined in the backup image
RESTORE DATABASE TEST4 ON /path1, /path2, /path3	/path1	/path1, /path2, /path3	Uses database path of existing database	/path1, /path2, /path3
RESTORE DATABASE TEST5 ON E:\newpath1, F:\newpath2 DBPATH ON D:	D:	E:\newpath1, F:\newpath2	Uses database path of existing database	E:\newpath1, F:\newpath2

For those cases where storage paths have been redefined as part of the restore operation, the table spaces that are defined to use automatic storage are automatically redirected to the new paths. However, you cannot explicitly redirect containers associated with automatic storage table spaces using the SET TABLESPACE CONTAINERS command; this action is not permitted.

Use the -s option of the db2ckbkp command to show whether or not automatic storage is enabled for a database within a backup image. The storage paths associated with the database are displayed if automatic storage is enabled.

For multi-partition automatic storage enabled databases, the RESTORE DATABASE command has a few extra implications:

1. The database must use the same set of storage paths on all database partitions.
2. Issuing a RESTORE command with new storage paths can only be done on the catalog database partition, which will set the state of the database to RESTORE_PENDING on all non-catalog database partitions.

Table 44. Restore implications for multi-partition databases

RESTORE DATABASE command	Issued on database partition #	No database on disk exists with same name		Database exists on disk with same name (includes skeleton databases)	
		Result on other database partitions	Storage paths	Result on other database partitions	Storage paths
RESTORE DATABASE TEST1	Catalog database partition	A skeleton database is created using the storage paths from the backup image on the catalog database partition. All other database partitions are placed in a RESTORE_PENDING state.	Uses storage paths defined in the backup image	Nothing. Storage paths have not changed so nothing happens to other database partitions	Uses storage paths defined in the backup image
	Non-catalog database partition	SQL2542N or SQL2551N is returned. If no database exists, the catalog database partition must be restored first.	N/A	Nothing. Storage paths have not changed so nothing happens to other database partitions	Uses storage paths defined in the backup image

Table 44. Restore implications for multi-partition databases (continued)

RESTORE DATABASE command	Issued on database partition #	No database on disk exists with same name		Database exists on disk with same name (includes skeleton databases)	
		Result on other database partitions	Storage paths	Result on other database partitions	Storage paths
RESTORE DATABASE TEST2 ON /path1, /path2, /path3	Catalog database partition	A skeleton database is created using the storage paths specified in the RESTORE command. All other database partitions are place in a RESTORE_PENDING state.	/path1, /path2, /path3		/path1, /path2, /path3
	Non-catalog database partition	SQL1174N is returned. If no database exists, the catalog database partition must be restored first. Storage paths cannot be specified on the RESTORE of a non-catalog database partition.	N/A	SQL1172N is returned. New storage paths cannot be specified on the RESTORE of a non-catalog database partition.	N/A

Cataloging databases

When you create a new database, it is automatically cataloged in the system database directory file. You might also use the CATALOG DATABASE command to explicitly catalog a database in the system database directory file.

The CATALOG DATABASE command allows you to catalog a database with a different alias name, or to catalog a database entry that was previously deleted using the UNCATALOG DATABASE command.

Although databases are cataloged automatically when a database is created, you still might have a need to catalog the database. When you do so, the database must exist.

By default directory files, including the database directory, are cached in memory using the Directory Cache Support (*dir_cache*) configuration parameter. When directory caching is enabled, a change made to a directory (for example, using a CATALOG DATABASE or UNCATALOG DATABASE command) by another

application might not become effective until your application is restarted. To refresh the directory cache used by a command line processor session, issue the TERMINATE command.

In a partitioned database, a cache for directory files is created on each database partition.

In addition to the application level cache, a database manager level cache is also used for internal, database manager look-up. To refresh this “shared” cache, issue the db2stop and db2start commands.

To catalog a database with a different alias name using the command line processor, use the CATALOG DATABASE command. For example, the following command line processor command catalogs the PERSON1 database as HUMANRES:

```
CATALOG DATABASE person1 AS humanres
WITH "Human Resources Database"
```

Here, the system database directory entry will have HUMANRES as the database alias, which is different from the database name (PERSON1).

To catalog a database in the system database directory from a client application, call the sqlecadb API.

To catalog a database on an instance other than the default using the command line processor, use the CATALOG DATABASE command. In the following example, connections to database B are to INSTNC_C. The instance instnc_c must already be cataloged as a local node before attempting this command.

```
CATALOG DATABASE b as b_on_ic AT NODE instnc_c
```

Note: The CATALOG DATABASE command is also used on client nodes to catalog databases that reside on database server computers.

Binding utilities to the database

When a database is created, the database manager attempts to bind the utilities in db2ubind.lst and in db2cli.lst to the database. These files are stored in the bnd subdirectory of your sqllib directory.

Binding a utility creates a *package*, which is an object that includes all the information needed to process specific SQL and XQuery statements from a single source file.

Note: If you want to use these utilities from a client, you must bind them explicitly. You must be in the directory where these files reside to create the packages in the sample database. The bind files are found in the bnd subdirectory of the sqllib directory. You must also bind the db2schema.bnd file when you create or upgrade the database from a client. Refer to “DB2 CLI bind files and package names” for details.

To bind or rebind the utilities to a database, from the command line, invoke the following commands, where sample is the name of the database:

```
connect to sample
bind @db2ubind.lst
```

Creating database aliases

An *alias* is an indirect method of referencing a table, nickname, or view, so that an SQL or XQuery statement can be independent of the qualified name of that table or view.

Only the alias definition must be changed if the table or view name changes. An alias can be created on another alias. An alias can be used in a view or trigger definition and in any SQL or XQuery statement, except for table check-constraint definitions, in which an existing table or view name can be referenced.

An alias can be defined for a table, view, or alias that does not exist at the time of definition. However, it must exist when the SQL or XQuery statement containing the alias is compiled.

An alias name can be used wherever an existing table name can be used, and can refer to another alias if no circular or repetitive references are made along the chain of aliases.

The alias name cannot be the same as an existing table, view, or alias, and can only refer to a table within the same database. The name of a table or view used in a CREATE TABLE or CREATE VIEW statement cannot be the same as an alias name in the same schema.

You do not require special authority to create an alias, unless the alias is in a schema other than the one owned by your current authorization ID, in which case DBADM authority is required.

When an alias, or the object to which an alias refers, is dropped, all packages dependent on the alias are marked invalid and all views and triggers dependent on the alias are marked inoperative.

To create an alias using the command line, enter:

```
CREATE ALIAS <alias_name> FOR <table_name>
```

The alias is replaced at statement compilation time by the table or view name. If the alias or alias chain cannot be resolved to a table or view name, an error results. For example, if WORKERS is an alias for EMPLOYEE, then at compilation time:

```
SELECT * FROM WORKERS
```

becomes in effect

```
SELECT * FROM EMPLOYEE
```

The following SQL statement creates an alias WORKERS for the EMPLOYEE table:

```
CREATE ALIAS WORKERS FOR EMPLOYEE
```

Note: DB2 for OS/390 or z/Series employs two distinct concepts of aliases: ALIAS and SYNONYM. These two concepts differ from DB2 database as follows:

- ALIASes in DB2 for OS/390 or z/Series:
 - Require their creator to have special authority or privilege
 - Cannot reference other aliases
- SYNONYMs in DB2 for OS/390 or z/Series:
 - Can only be used by their creator
 - Are always unqualified

- Are dropped when a referenced table is dropped
- Do not share namespace with tables or views

Connecting to distributed relational databases

Distributed relational databases are built on formal requester-server protocols and functions.

An *application requester* supports the application end of a connection. It transforms a database request from the application into communication protocols suitable for use in the distributed database network. These requests are received and processed by a *database server* at the other end of the connection. Working together, the application requester and the database server handle communication and location considerations, so that the application can operate as if it were accessing a local database.

An application process must connect to a database manager's application server before SQL statements that reference tables or views can be executed. The CONNECT statement establishes a connection between an application process and its server.

There are two types of CONNECT statements:

- CONNECT (Type 1) supports the single database per unit of work (Remote Unit of Work) semantics.
- CONNECT (Type 2) supports the multiple databases per unit of work (Application-Directed Distributed Unit of Work) semantics.

The DB2 call level interface (CLI) and embedded SQL support a connection mode called *concurrent transactions*, which allows multiple connections, each of which is an independent transaction. An application can have multiple concurrent connections to the same database.

The application server can be local to or remote from the environment in which the process is initiated. An application server is present, even if the environment is not using distributed relational databases. This environment includes a local directory that describes the application servers that can be identified in a CONNECT statement.

The application server runs the bound form of a static SQL statement that references tables or views. The bound statement is taken from a package that the database manager has previously created through a bind operation.

For the most part, an application connected to an application server can use statements and clauses that are supported by the application server's database manager. This is true even if an application is running through the application requester of a database manager that does *not* support some of those statements and clauses.

Remote unit of work for distributed relational databases

The *remote unit of work facility* provides for the remote preparation and execution of SQL statements.

An application process at computer system "A" can connect to an application server at computer system "B" and, within one or more units of work, execute any

number of static or dynamic SQL statements that reference objects at “B”. After ending a unit of work at B, the application process can connect to an application server at computer system C, and so on.

Most SQL statements can be remotely prepared and executed, with the following restrictions:

- All objects referenced in a single SQL statement must be managed by the same application server.
- All of the SQL statements in a unit of work must be executed by the same application server.

At any given time, an application process is in one of four possible *connection states*:

- Connectable and connected

An application process is connected to an application server, and `CONNECT` statements can be executed.

If implicit connect is available:

- The application process enters this state when a `CONNECT TO` statement or a `CONNECT` without operands statement is successfully executed from the connectable and unconnected state.
- The application process may enter this state from the implicitly connectable state if any SQL statement other than `CONNECT RESET`, `DISCONNECT`, `SET CONNECTION`, or `RELEASE` is issued.

Whether or not implicit connect is available, this state is entered when:

- A `CONNECT TO` statement is successfully executed from the connectable and unconnected state.
- A `COMMIT` or `ROLLBACK` statement is successfully issued, or a forced rollback occurs from the unconnectable and connected state.

- Unconnectable and connected

An application process is connected to an application server, but a `CONNECT TO` statement cannot be successfully executed to change application servers. The application process enters this state from the connectable and connected state when it executes any SQL statement other than the following: `CONNECT TO`, `CONNECT` with no operand, `CONNECT RESET`, `DISCONNECT`, `SET CONNECTION`, `RELEASE`, `COMMIT`, or `ROLLBACK`.

- Connectable and unconnected

An application process is not connected to an application server. `CONNECT TO` is the only SQL statement that can be executed; otherwise, an error (SQLSTATE 08003) is raised.

Whether or not implicit connect is available, the application process enters this state if an error occurs when a `CONNECT TO` statement is issued, or an error occurs within a unit of work, causing the loss of a connection and a rollback. An error that occurs because the application process is not in the connectable state, or because the server name is not listed in the local directory, does not cause a transition to this state.

If implicit connect is not available:

- The application process is initially in this state
- The `CONNECT RESET` and `DISCONNECT` statements cause a transition to this state.

- Implicitly connectable (if implicit connect is available).

If implicit connect is available, this is the initial state of an application process. The CONNECT RESET statement causes a transition to this state. Issuing a COMMIT or ROLLBACK statement in the unconnectable and connected state, followed by a DISCONNECT statement in the connectable and connected state, also results in this state.

Availability of implicit connect is determined by installation options, environment variables, and authentication settings.

It is not an error to execute consecutive CONNECT statements, because CONNECT itself does not remove the application process from the connectable state. It is, however, an error to execute consecutive CONNECT RESET statements. It is also an error to execute any SQL statement other than CONNECT TO, CONNECT RESET, CONNECT with no operand, SET CONNECTION, RELEASE, COMMIT, or ROLLBACK, and then to execute a CONNECT TO statement. To avoid this error, a CONNECT RESET, DISCONNECT (preceded by a COMMIT or ROLLBACK statement), COMMIT, or ROLLBACK statement should be executed before the CONNECT TO statement.

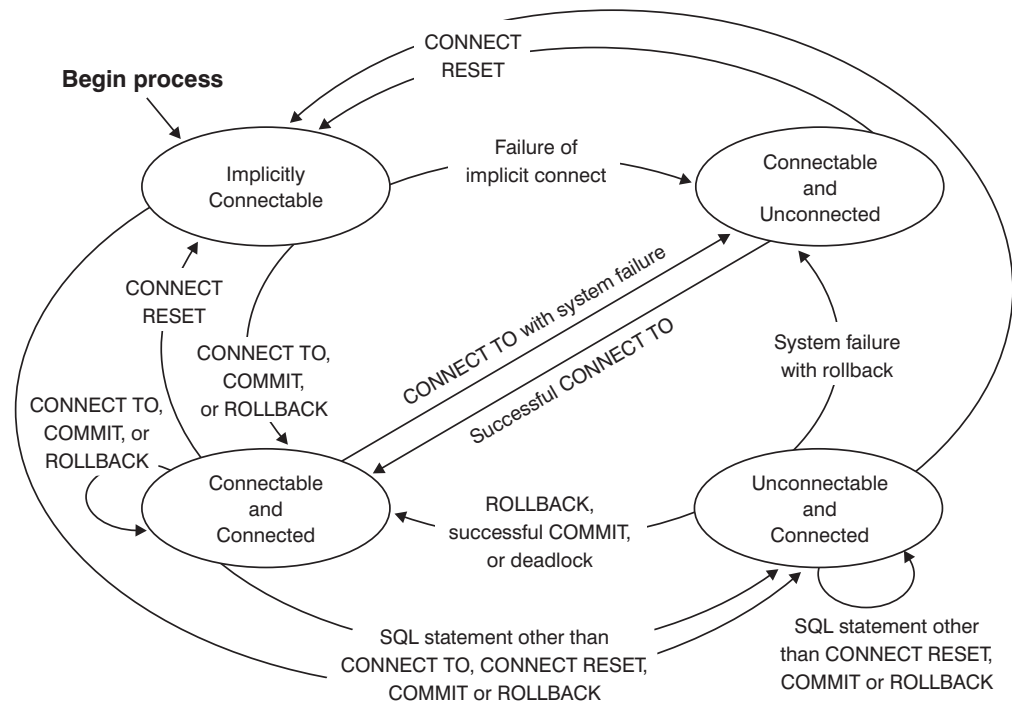


Figure 3. Connection State Transitions If Implicit Connect Is Available




Figure 4. Connection State Transitions If Implicit Connect Is Not Available

Application-directed distributed unit of work

The *application-directed distributed unit of work facility* provides for the remote preparation and execution of SQL statements.

An application process at computer system "A" can connect to an application server at computer system "B" by issuing a `CONNECT` or a `SET CONNECTION` statement. The application process can then execute any number of static and dynamic SQL statements that reference objects at "B" before ending the unit of work. All objects referenced in a single SQL statement must be managed by the same application server. However, unlike the remote unit of work facility, any number of application servers can participate in the same unit of work. A commit or a rollback operation ends the unit of work.

An application-directed distributed unit of work uses a type 2 connection. A *type 2* connection connects an application process to the identified application server, and establishes the rules for application-directed distributed units of work.

A type 2 application process:

- Is always connectable
- Is either in the connected state or in the unconnected state
- Has zero or more connections.

Each connection of an application process is uniquely identified by the database alias of the application server for the connection.

An individual connection always has one of the following connection states:

- current and held
- current and release-pending
- dormant and held
- dormant and release-pending

A type 2 application process is initially in the unconnected state, and does not have any connections. A connection is initially in the current and held state.




Figure 5. Application-Directed Distributed Unit of Work Connection State Transitions

Application process connection states

There are certain rules that apply to the execution of a CONNECT statement.

The following rules apply to the execution of a CONNECT statement:

- A context cannot have more than one connection to the same application server at the same time.
- When an application process executes a SET CONNECTION statement, the specified location name must be an existing connection in the set of connections for the application process.
- When an application process executes a CONNECT statement, and the SQLRULES(STD) option is in effect, the specified server name must *not* be an existing connection in the set of connections for the application process. For a description of the SQLRULES option, see “Options that govern unit of work semantics” on page 137.

If an application process has a current connection, the application process is in the *connected* state. The CURRENT SERVER special register contains the name of the application server for the current connection. The application process can execute SQL statements that refer to objects managed by that application server.

An application process that is in the unconnected state enters the connected state when it successfully executes a CONNECT or a SET CONNECTION statement. If there is no connection, but SQL statements are issued, an implicit connect is made, provided the DB2DBDFT environment variable has been set with the name of a default database.

If an application process does not have a current connection, the application process is in the *unconnected* state. The only SQL statements that can be executed are CONNECT, DISCONNECT ALL, DISCONNECT (specifying a database), SET CONNECTION, RELEASE, COMMIT, ROLLBACK, and local SET statements.

An application process in the *connected state* enters the *unconnected state* when its current connection intentionally ends, or when an SQL statement fails, causing a rollback operation at the application server and loss of the connection. Connections end intentionally following the successful execution of a DISCONNECT statement, or a COMMIT statement when the connection is in release-pending state. (If the DISCONNECT precompiler option is set to AUTOMATIC, all connections end. If it is set to CONDITIONAL, all connections that do not have open WITH HOLD cursors end.)

Connection states

There are two types of connection states: “held and release-pending states” and “current and dormant states”.

If an application process executes a CONNECT statement, and the server name is known to the application requester but is not in the set of existing connections for the application process: (i) the current connection is placed into the *dormant connection state*, the server name is added to the set of connections, and the new connection is placed into both the *current connection state* and the *held connection state*.

If the server name is already in the set of existing connections for the application process, and the application is precompiled with the SQLRULES(STD) option, an error (SQLSTATE 08002) is raised.

Held and release-pending states. The RELEASE statement controls whether a connection is in the held or the release-pending state. The *release-pending* state means that a disconnect is to occur at the next successful commit operation. (A rollback has no effect on connections.) The *held* state means that a disconnect is *not* to occur at the next commit operation.

All connections are initially in the held state and can be moved to the release-pending state using the RELEASE statement. Once in the release-pending state, a connection cannot be moved back to the held state. A connection remains in release-pending state across unit of work boundaries if a ROLLBACK statement is issued, or if an unsuccessful commit operation results in a rollback operation.

Even if a connection is not explicitly marked for release, it may still be disconnected by a commit operation if the commit operation satisfies the conditions of the DISCONNECT precompiler option.

Current and dormant states. Regardless of whether a connection is in the held state or the release-pending state, it can also be in the current state or the dormant state. A connection in the *current* state is the connection being used to execute SQL statements while in this state. A connection in the *dormant* state is a connection that is not current.

The only SQL statements that can flow on a dormant connection are COMMIT, ROLLBACK, DISCONNECT, or RELEASE. The SET CONNECTION and CONNECT statements change the connection state of the specified server to current, and any existing connections are placed or remain in dormant state. At any point in time, only one connection can be in current state. If a dormant connection becomes current in the same unit of work, the state of all locks, cursors, and prepared statements is the same as the state they were in the last time that the connection was current.

When a connection ends

When a connection ends, all resources that were acquired by the application process through the connection, and all resources that were used to create and maintain the connection are de-allocated. For example, if the application process executes a RELEASE statement, any open cursors are closed when the connection ends during the next commit operation.

A connection can also end because of a communications failure. If this connection is in current state, the application process is placed in unconnected state.

All connections for an application process end when the process ends.

Options that govern unit of work semantics

The semantics of type 2 connection management are determined by a set of precompiler options. These options are summarized below with default values indicated by bold and underlined text.

- CONNECT (1 | 2). Specifies whether CONNECT statements are to be processed as type 1 or type 2.
- SQLRULES (DB2 | STD). Specifies whether type 2 CONNECTs are to be processed according to the DB2 rules, which allow CONNECT to switch to a dormant connection, or the SQL92 Standard rules, which do not allow this.
- DISCONNECT (EXPLICIT | CONDITIONAL | AUTOMATIC). Specifies what database connections are to be disconnected when a commit operation occurs:
 - Those that have been explicitly marked for release by the SQL RELEASE statement (EXPLICIT)
 - Those that have no open WITH HOLD cursors, and those that are marked for release (CONDITIONAL)
 - All connections (AUTOMATIC).
- SYNCPOINT (ONEPHASE | TWOPHASE | NONE). Specifies how COMMITs or ROLLBACKs are to be coordinated among multiple database connections. This option is ignored, and is included for backwards compatibility only.
 - Updates can only occur against one database in the unit of work, and all other databases are read-only (ONEPHASE). Any update attempts to other databases raise an error (SQLSTATE 25000).
 - A transaction manager (TM) is used at run time to coordinate two-phase COMMITs among those databases that support this protocol (TWOPHASE).

- Does not use a TM to perform two-phase COMMITs, and does not enforce single updater, multiple reader (NONE). When a COMMIT or a ROLLBACK statement is executed, individual COMMITs or ROLLBACKs are posted to all databases. If one or more ROLLBACKs fail, an error (SQLSTATE 58005) is raised. If one or more COMMITs fail, another error (SQLSTATE 40003) is raised.

To override any of the above options at run time, use the SET CLIENT command or the sqlesetc application programming interface (API). Their current settings can be obtained using the QUERY CLIENT command or the sqleqryc API. Note that these are not SQL statements; they are APIs defined in the various host languages and in the command line processor (CLP).

Data representation considerations

Different systems represent data in different ways. When data is moved from one system to another, data conversion must sometimes be performed.

Products supporting DRDA automatically perform any necessary conversions at the receiving system.

To perform conversions of numeric data, the system needs to know the data type and how it is represented by the sending system. Additional information is needed to convert character strings. String conversion depends on both the code page of the data and the operation that is to be performed on that data. Character conversions are performed in accordance with the IBM Character Data Representation Architecture (CDRA). For more information about character conversion, see the *Character Data Representation Architecture: Reference & Registry* (SC09-2190-00) manual.

Viewing the local or system database directory files

Use the LIST DATABASE DIRECTORY command to view the information associated with the databases that you have on your system.

Before viewing either the local or system database directory files, you must have previously created an instance and a database.

To see the contents of the local database directory file, issue the following command, where <location> specifies the location of the database:

```
LIST DATABASE DIRECTORY ON <location>
```

To see the contents of the system database directory file, issue the LIST DATABASE DIRECTORY command *without* specifying the location of the database directory file.

Dropping databases

Dropping a database can have far-reaching effects, because this action deletes all its objects, containers, and associated files. The dropped database is removed (uncataloged) from the database directories.

To drop a database using the command line, enter:

```
DROP DATABASE <name>
```

The following command deletes the database SAMPLE:

```
DROP DATABASE SAMPLE
```

Note: If you drop the SAMPLE database and find that you need it again, you can re-create it.

To drop a database from a client application, call the sqledrpd API. To drop a database at a specified database partition server, call the sqledpan API.

Dropping aliases

When you drop an alias, its description is deleted from the catalog, any packages and cached dynamic queries that reference the alias are invalidated, and all views and triggers dependent on the alias are marked inoperative.

To drop aliases, from the command line, issue the DROP statement:

```
DROP ALIAS EMPLOYEE-ALIAS
```

Chapter 7. Database partitions

A *database partition* is a part of a database that consists of its own data, indexes, configuration files, and transaction logs. A database partition is sometimes called a node or a database node. A partitioned database environment is a database installation that supports the distribution of data across database partitions.

For complete details about database partitions, see the *Partitioning and Clustering Guide*.

Chapter 8. Buffer pools

A *buffer pool* is an area of main memory that has been allocated by the database manager for the purpose of caching table and index data as it is read from disk. Every DB2 database must have a buffer pool.

Each new database has a default bufferpool defined, called IBMDEFAULTBP. Additional buffer pools can be created, dropped, and modified, using the CREATE BUFFERPOOL, DROP BUFFERPOOL, and ALTER BUFFERPOOL statements. The SYSCAT.BUFFERPOOLS catalog view accesses the information for the buffer pools defined in the database.

How buffer pools are used

When a row of data in a table is first accessed, the database manager places the page that contains that data into a buffer pool. Pages stay in the buffer pool until the database is shut down or until the space occupied by the page is required by another page.

Pages in the buffer pool can be either in-use or not, and they can be dirty or clean:

- In-use pages are currently being read or updated. To maintain data consistency, the database manager only allows one agent to be updating a given page in a bufferpool at one time. If a page is being updated, it is being accessed exclusively by one agent. If it is being read, it may be read by multiple agents simultaneously.
- "Dirty" pages contain data that has been changed but has not yet been written to disk.
- After a changed page is written to disk, it is clean and might remain in the buffer pool.

A large part of tuning a database involves setting the configuration parameters that control the movement of data into the buffer pool and the writing of data from the buffer out to disk. If not needed by a recent agent, the page space can be used for new page requests from new applications. Database manager performance is degraded by extra disk I/O.

You can use the snapshot monitor to calculate the buffer pool hit ratio, which can help you tune your buffer pools.

Designing buffer pools

The sizes of all buffer pools can have a major impact on the performance of your database.

Before you create a new buffer pool, resolve the following items:

- What buffer pool name do you want to use?
- Whether the buffer pool is to be created immediately or following the next time that the database is deactivated and reactivated?
- Whether the buffer pool should exist for all database partitions, or for a subset of the database partitions?

- What page size you want for the buffer pool? See “Buffer pool page sizes” below.
- Whether the buffer pool will be a fixed size, or whether the database manager will automatically adjust the size of the buffer pool in response to your workload? It is suggested that you allow the database manager to tune your buffer pool automatically by leaving the SIZE parameter unspecified during buffer pool creation. For details, see the SIZE parameter of the “CREATE BUFFERPOOL statement” and “Buffer pool memory considerations.”
- Whether you want to reserve a portion of the buffer pool for block based I/O? For details, see: “Block-based buffer pools for improved sequential prefetching”.

Relationship between table spaces and buffer pools

When designing buffer pools, you need to understand the relationship between table spaces and buffer pools. Each table space is associated with a specific buffer pool. IBMDEFAULTBP is the default buffer pool. The database manager also allocates these system buffer pools: IBMSYSTEMBP4K, IBMSYSTEMBP8K, IBMSYSTEMBP16K, and IBMSYSTEMBP32K (formerly known as the “hidden buffer pools”). To associate another buffer pool with a table space, the buffer pool must exist and the two must have the same page size. The association is defined when the table space is created (using the CREATE TABLESPACE statement), but it can be changed at a later time (using the ALTER TABLESPACE statement).

Having more than one buffer pool allows you to configure the memory used by the database to improve overall performance. For example, if you have a table space with one or more large (larger than available memory) tables that are accessed randomly by users, the size of the buffer pool can be limited, because caching the data pages might not be beneficial. The table space for an online transaction application might be associated with a larger buffer pool, so that the data pages used by the application can be cached longer, resulting in faster response times. Care must be taken in configuring new buffer pools.

Buffer pool page sizes

The page size for the default buffer pool is set when you use the CREATE DATABASE command. This default represents the default page size for all future CREATE BUFFERPOOL and CREATE TABLESPACE statements. If you do not specify the page size when creating the database, the default page size is 4 KB.

Note: If you have determined that a page size of 8 KB, 16 KB, or 32 KB is required by your database, you must have at least one bufferpool of the matching page size defined and associated with table space in your database.

However, you might need a buffer pool that has different characteristics than the system buffer pool. You can create new buffer pools for the database manager to use. You may have to restart the database for table space and bufferpool changes to take effect. The page sizes that you specify for your table spaces should determine the page sizes that you choose for your buffer pools. The choice of page size used for a buffer pool is important because you cannot alter the page size after you create a buffer pool.

Buffer pool memory considerations

Memory requirements

When designing buffer pools, you should also consider the memory requirements based on the amount of installed memory on your machine

and the memory required by other applications running concurrently with the database manager on the same machine. Operating system data swapping occurs when there is insufficient memory to hold all the data being accessed. This occurs when some data is written or swapped to temporary disk storage to make room for other data. When the data on temporary disk storage is needed, it is swapped back into main memory.

Buffer pool memory protection

With Version 9.5, data pages in buffer pool memory are protected using storage keys, which are available only if explicitly enabled by the `DB2_MEMORY_PROTECT` registry variable, and only on AIX (5.3 TL06 5.4), running on POWER6™.

Buffer pool memory protection works on a per-agent level; any particular agent will only have access to buffer pool pages when that agent needs access. Memory protection works by identifying at which times the DB2 engine threads should have access to the buffer pool memory and at which times they should not have access. For details, see: “Buffer pool memory protection (AIX running on POWER6).”

Address Windowing Extensions (AWE) and Extended Storage (ESTORE)

Note: AWE and ESTORE features have been discontinued, including the ESTORE-related keywords, monitor elements, and data structures. To allocate more memory, you need to upgrade to a 64-bit hardware operating system, and associated DB2 products. You should also modify applications and scripts to remove references to this discontinued functionality.

Buffer pool memory protection (AIX running on POWER6)

The database manager uses the buffer pool to apply additions, modifications and deletions to much of the database data. On AIX 5.3 TL06+ running on POWER6, you can use *storage keys* to protect the buffer pool memory.

Storage keys is a new feature in IBM® Power6 processors and the AIX® operating system that allows the protection of ranges of memory using hardware keys at a kernel thread level. Storage key protection reduces buffer pool memory corruption problems and limits errors that might halt the database. Attempts to illegally access the buffer pool by programming means cause an error condition that the database manager can detect and deal with.

Note: Buffer pool memory protection works on a per-agent level; any particular agent will only have access to buffer pool pages when that agent needs access.

The database manager protects buffer pools by restricting access to buffer pool memory. When an agent requires access to the buffer pools to perform its work, it will temporarily be granted access to the buffer pool memory. When the agent no longer requires access to the buffer pools, access will be revoked. This ensures that agents are only allowed to modify buffer pool contents when absolutely needed, reducing the likelihood of buffer pool corruptions. Any illegal access to buffer pool memory will result in a segmentation error. Tools to diagnose these errors are provided, such as the `db2diag`, `db2fodc`, `db2pdcfg`, and `db2support` commands.

For complete enablement of this buffer pool memory protection feature, in order to increase the resilience of the database engine, you should enable both the `DB2_MEMORY_PROTECT` and the `DB2_THREAD_SUSPENSION` registry variables:

DB2_MEMORY_PROTECT registry variable

This registry variable enables and disables the buffer pool memory protection feature. When DB2_MEMORY_PROTECT is enabled (set to YES), and a DB2 engine thread tries to illegally access buffer pool memory, that engine thread traps. The default is NO.

DB2_THREAD_SUSPENSION registry variable

This registry variable enables and disables the DB2 thread suspension feature. It allows you to control whether a DB2 instance sustains a trap by suspending a faulty engine thread (that is, a thread which has tried to illegally access memory protected with storage keys).

Note: The buffer pool memory protection feature depends on the implementation of AIX Storage Protect Keys and it may not work with the pinned shared memory. If DB2_MEMORY_PROTECT is specified with DB2_PINNED_BP or DB2_LARGE_PAGE_MEM setting, AIX Storage Protect Keys may not be enabled. For more information about AIX Storage Protect Keys, refer to the following link: http://publib.boulder.ibm.com/infocenter/systems/scope/aix/index.jsp?topic=/com.ibm.aix.genprogc/doc/genprogc/storage_protect_keys.htm

DB2_THREAD_SUSPENSION can be enabled only if DB2_MEMORY_PROTECT is enabled. If DB2_THREAD_SUSPENSION is enabled:

- Whatever failure in the thread, whether or not caused by an attempt to access memory that is protected using storage keys and to which the thread has no access, if it happens at a point in which the thread has no access to this memory protected using storage keys, then the database manager guarantees that the protected memory has not been corrupted, and consequently, lets the DB2 engine continue running.
- When running a User Defined Function without SQL in unfenced mode, if a buffer pool memory protection violation is detected, the database manager will return the error to the caller of the UDF and the database will continue running without being affected.

Creating buffer pools

Use the CREATE BUFFERPOOL statement to define a new buffer pool to be used by the database manager.

Example of a basic CREATE BUFFERPOOL statement is:

```
CREATE BUFFERPOOL <buffer pool name>  
    PAGESIZE 4096
```

The buffer pool may become active immediately if there is sufficient memory available. By default new buffer pools are created using the IMMEDIATE keyword, and on most platforms, the database manager will be able to acquire more memory. The expected return is successful memory allocation. Only in cases where the database manager is unable to allocate the extra memory will it return a warning condition stating that the buffer pool could not be started, and this is done on the subsequent database startup. For immediate requests, you do not need to restart the database. When this statement is committed, the buffer pool is reflected in the system catalog tables, but the buffer pool does not become active until the next time the database is started. For more information about this statement, including other options, see the "CREATE BUFFERPOOL statement".

If you issue a `CREATE BUFFERPOOL DEFERRED`, the buffer pool is not immediately activated; instead, it is created at the next database startup. Until the database is restarted, any new table spaces will use an existing buffer pool, even if that table space is created to explicitly use the deferred buffer pool.

There needs to be enough real memory on the machine for the total of all the buffer pools that you have created. The operating system also needs some memory to operate.

To create a buffer pool using the command line, do the following:

1. Get the list of buffer pool names that already exist in the database by issuing the following SQL statement:

```
SELECT BPNAME FROM SYSCAT.BUFFERPOOLS
```

2. Choose a buffer pool name that is not currently found in the result list.
3. Determine the characteristics of the buffer pool you are going to create.
4. Ensure that you have the correct authorization ID to run the `CREATE BUFFERPOOL` statement.
5. Issue the `CREATE BUFFERPOOL` statement.

On partitioned databases, you can also define the buffer pool to be created differently, including different sizes, on each database partition. The default `ALL DBPARTITIONNUMS` clause indicates that this buffer pool will be created on all database partitions in the database.

In the following example, the optional `DATABASE PARTITION GROUP` clause identifies the database partition group or groups to which the buffer pool definition applies:

```
CREATE BUFFERPOOL <buffer pool name>  
  PAGESIZE 4096  
  DATABASE PARTITION GROUP <db partition group name>
```

If this parameter is specified, the buffer pool will only be created on database partitions in these database partition groups. Each database partition group must currently exist in the database. If the `DATABASE PARTITION GROUP` clause is not specified, this buffer pool will be created on all database partitions (and on any database partitions that are subsequently added to the database).

For more information, see the “`CREATE BUFFERPOOL` statement”.

Modifying buffer pools

There are a number of reasons why you might want to modify a buffer pool, for example, to enable self-tuning memory. To do this, you use the `ALTER BUFFERPOOL` statement.

The authorization ID of the statement must have `SYSCTRL` or `SYSADM` authority.

When working with buffer pools, you might need to do one of the following tasks :

- Enable self tuning for a buffer pool, allowing the database manager to adjust the size of the buffer pool in response to your workload.
- Modify the block area of the buffer pool for block-based I/O.
- Add this buffer pool definition to a new database partition group.
- Modify the size of the buffer pool on some or all database partitions.

To alter a buffer pool using the command line, do the following:

1. To get the list of the buffer pool names that already exist in the database, issue the following statement:

```
SELECT BPNAME FROM SYSCAT.BUFFERPOOLS
```

2. Choose the buffer pool name from the result list.
3. Determine what changes need to be made.
4. Ensure that you have the correct authorization ID to run the ALTER BUFFERPOOL statement.

Note: Two key parameters are IMMEDIATE and DEFERRED. With IMMEDIATE, the buffer pool size is changed without having to wait until the next database activation for it to take effect. If there is insufficient database shared memory to allocate new space, the statement is run as DEFERRED.

With DEFERRED, the changes to the buffer pool will not be applied until the database is reactivated. Reserved memory space is not needed; the database manager allocates the required memory from the system at activation time.

5. Use the ALTER BUFFERPOOL statement to alter a single attribute of the buffer pool object. For example:

```
ALTER BUFFERPOOL buffer pool name SIZE number of pages
```

- The *buffer pool name* is a one-part name that identifies a buffer pool described in the system catalogs.
- The *number of pages* is the new number of pages to be allocated to this specific buffer pool. You can also use a value of -1, which indicates that the size of the buffer pool should be the value found in the **buffpage** database configuration parameter.

The statement can also have the DBPARTITIONNUM <db partition number> clause that specifies the database partition on which the size of the buffer pool is modified. If this clause is not specified, the size of the buffer pool is modified on all database partitions except those that have an exception entry in SYSCAT.BUFFERPOOLDBPARTITIONS. For details on using this clause for database partitions, see the ALTER BUFFERPOOL statement.

Changes to the buffer pool as a result of this statement are reflected in the system catalog tables when the statement is committed. However, no changes to the actual buffer pool take effect until the next time the database is started, except for successful ALTER BUFFERPOOL requests specified with the default IMMEDIATE keyword.

There must be enough real memory on the machine for the total of all the buffer pools that you have created. There also needs to be sufficient real memory for the rest of the database manager and for your applications.

Dropping buffer pools

When dropping buffer pools, ensure that no table spaces have been assigned to those buffer pools. You cannot drop the IBMDEFAULTBP buffer pool.

Disk storage may not be released until the next connection to the database. Storage memory is not actually released from a dropped buffer pool until the database is stopped. Buffer pool memory is released immediately, to be used by the database manager.

You can use the `DROP BUFFERPOOL` statement to drop buffer pools, as follows:

```
DROP BUFFERPOOL <buffer pool name>
```

Chapter 9. Table spaces

A *table space* is a storage structure containing tables, indexes, large objects, and long data. Table spaces reside in database partition groups. They allow you to assign the location of database and table data directly onto containers. (A container can be a directory name, a device name, or a file name.) This can provide improved performance and more flexible configuration.

Since table spaces reside in database partition groups, the table space selected to hold a table defines how the data for that table is distributed across the database partitions in a database partition group. A single table space can span several containers. It is possible for multiple containers (from one or more table spaces) to be created on the same physical disk (or drive). If you are using automatic storage table spaces, this is handled by the database manager. If you are *not* using automatic storage table spaces, for improved performance, each container should use a different disk.

Figure 6 illustrates the relationship between tables and table spaces within a database, and the containers associated with that database.

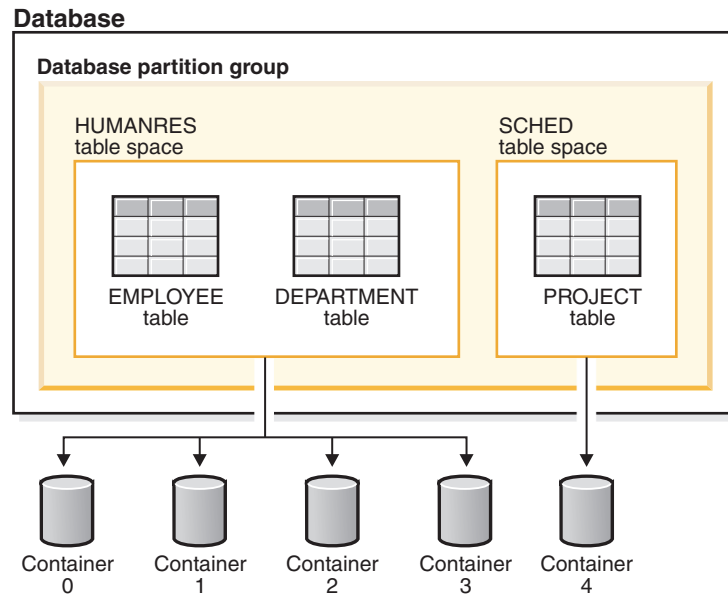


Figure 6. Table spaces and tables in a database

The EMPLOYEE and DEPARTMENT tables are in the HUMANRES table space, which spans containers 0, 1, 2 and 3. The PROJECT table is in the SCHED table space in container 4. This example shows each container existing on a separate disk.

The database manager attempts to balance the data load across containers. As a result, all containers are used to store data. The number of pages that the database manager writes to a container before using a different container is called the *extent size*. The database manager does not always start storing table data in the first container.

Figure 7 shows the HUMANRES table space with an extent size of two 4 KB pages, and four containers, each with a small number of allocated extents. The DEPARTMENT and EMPLOYEE tables both have seven pages, and span all four containers.

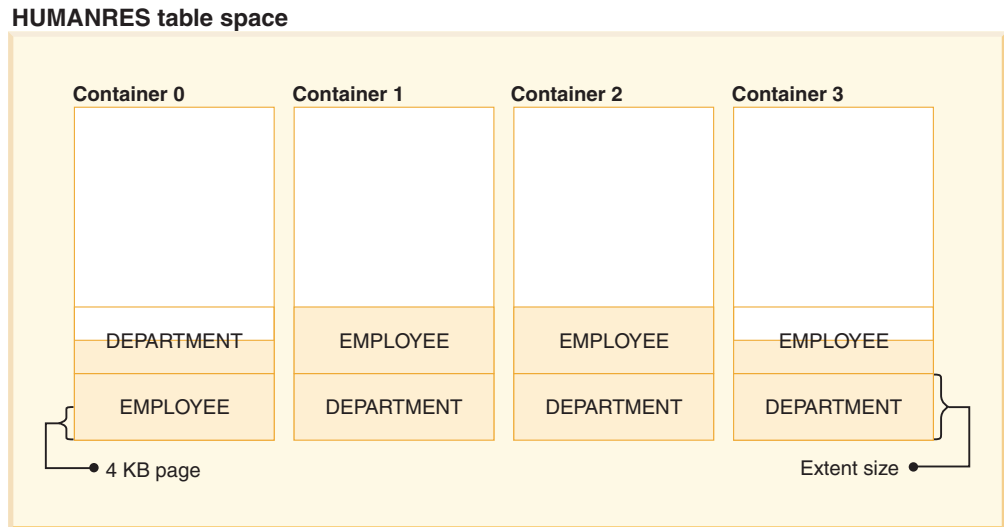


Figure 7. Containers and extents in a table space

Designing table spaces

Table spaces are used to specify where data in a database is physically stored on a system and to provide a layer of indirection between the database and the container objects in which the actual data resides.

There are many reasons for creating table spaces, including recoverability, and the ability to isolate objects in different bufferpools. With automatic storage, you no longer need to be concerned about the physical disk location of table spaces, nor about the physical location of containers. The database manager automatically assigns or creates containers for table spaces.

A database must contain at least three table spaces:

- One *catalog table space*, which contains all of the system catalog tables for the database. This table space is called SYSCATSPACE, and it cannot be dropped. IBMCATGROUP is the default database partition group for this table space.
- One or more *user table spaces*, which contain all user defined tables. By default, one table space, USERSPACE1, is created. IBMDEFAULTGROUP is the default database partition group for this table space.

You should specify a table space name when you create a table, or the results may not be what you intend.

A table's page size is determined either by row size, or the number of columns. The maximum allowable length for a row is dependent upon the page size of the table space in which the table is created. Possible values for page size are 4 KB, 8 KB, 16 KB, and 32 KB. Before Version 9.1, the default page size was 4 KB. In Version 9.1 and following, the default page size may be one of the other supported values. The default page size is declared when creating a new database. Once the default page size has been declared, you are still free to create a table space with one page size for the table, and a different table space with a different page size for long or LOB data. (Recall that SMS does not

support tables that span table spaces, but that DMS does.) If the number of columns or the row size exceeds the limits for a table space's page size, an error is returned (SQLSTATE 42997).

- One or more *temporary table spaces*, which contain temporary tables. Temporary table spaces can be *system temporary table spaces* or *user temporary table spaces*.

System temporary table spaces hold temporary data required by the database manager while performing operations such as sorts or joins. These types of operations require extra space to process the results set. A database must have at least one system temporary table space; by default, one system temporary table space called TEMPSPACE1 is created at database creation. IBMTEMPGROUP is the default database partition group for this table space.

User temporary table spaces hold temporary data from tables created with a DECLARE GLOBAL TEMPORARY TABLE statement. To allow the definition of declared temporary tables, at least one user temporary table space should be created with the appropriate USE privileges. USE privileges are granted using the GRANT statement. A user temporary table spaces is *not* created by default at database creation.

If a database uses more than one temporary table space and a new temporary object is needed, the optimizer will choose an appropriate page size for this object. That object will then be allocated to the temporary table space with the corresponding page size. If there is more than one temporary table space with that page size, then the table space will be chosen in a round-robin fashion. In most circumstances, it is not recommended to have more than one temporary table space of any one page size.

If queries are running against tables in table spaces that are defined with page sizes larger than the default, some of them may fail. This will occur if there are no temporary table spaces defined with a larger page size. You may need to create a temporary table space with a larger page size (if the default was 4 KB, then you would need to create a temporary table space with a page size of 8 KB, 16 KB, or 32 KB). Any DML (Data Manipulation Language) statement could fail unless there exists a temporary table space with the same page size as the largest page size in the user table space.

You should define a single SMS temporary table space with a page size equal to the page size used in the majority of your user table spaces. This should be adequate for typical environments and workloads.

Table spaces and database partition groups

In a partitioned database environment, each table space is associated with a specific database partition group. This allows the characteristics of the table space to be applied to each database partition in the database partition group.

The database partition group must exist (it is defined with the CREATE DATABASE PARTITION GROUP statement), and the association between the table space and the database partition group is defined when the table space is created using the CREATE TABLESPACE statement.

You cannot change the association between table space and database partition group using the ALTER TABLESPACE statement. You can only change the table space specification for individual database partitions within the database partition group. In a single-partition environment, each table space is associated with the default database partition group. The default database partition group, when defining a table space, is IBMDEFAULTGROUP, unless a system temporary table space is being defined; then IBMTEMPGROUP is used.

Types of table spaces

A database must contain at least three types table spaces: one catalog table space, one or more user table spaces, and one or more temporary table spaces.

There are two types of table space, both of which can be used in a single database:

- System managed space, in which the operating system's file manager controls the storage space.
- database managed space, in which the database manager controls the storage space.

In a partitioned database environment, the catalog partition will contain all three default table spaces, and the other database partitions will each contain only TEMPSPACE1 and USERSPACE1.

Automatic storage table spaces can also be created, which will use SMS or DMS as the underlying table space type. The actual type (SMS or DMS) is chosen by the database manager depending on the type of data that will reside in it (SMS for temporary table spaces, DMS otherwise).

System managed space

In an SMS (System Managed Space) table space, the operating system's file system manager allocates and manages the space where the table is stored.

The storage model typically consists of many files, representing table objects, stored in the file system space. You decide the location of the files, the database manager controls their names, and the file system is responsible for managing them. By controlling the amount of data written to each file, the database manager distributes the data evenly across the table space containers.

Each table has at least one SMS physical file associated with it.

The data in the table spaces is striped by extent across all the containers in the system. An *extent* is a group of consecutive pages defined to the database. The file extension denotes the type of the data stored in the file. To distribute the data evenly across all containers in the table space, the starting extents for tables are placed in round-robin fashion across all containers. Such distribution of extents is particularly important if the database contains many small tables. DB2 striping is recommended when writing data into multiple containers. If you choose to implement disk striping along with DB2 striping, the extent size of the table space and the strip size of the disk should be identical.

In an SMS table space, space for tables is allocated on demand. The amount of space that is allocated is dependent on the setting of the *multipage_alloc* database configuration parameter. If this configuration parameter is set to YES, then a full extent (typically made up of two or more pages) will be allocated when space is required. Otherwise, space will be allocated one page at a time.

Multipage file allocation is enabled by default. The value of the *multipage_alloc* database configuration parameter will indicate if multipage file allocation is enabled.

Note: Multipage file allocation is not applicable to temporary table spaces.

Multi-page file allocation only affects the data and index portions of a table. This means that the .LF, .LB, and .LBA files are not extended one extent at a time.

When all space in a single container in an SMS table space is allocated to tables, the table space is considered full, even if space remains in other containers. You can add containers to an SMS table space only on a database partition that does not yet have any containers.

Note: SMS table spaces can take advantage of file-system prefetching and caching.

SMS table spaces are defined using the `MANAGED BY SYSTEM` option on the `CREATE DATABASE` command, or on the `CREATE TABLESPACE` statement. You must consider two key factors when you design your SMS table spaces:

- Containers for the table space.

You must specify the number of containers that you want to use for your table space. It is very important to identify all the containers you want to use, because you cannot add or delete containers after an SMS table space is created. In a partitioned database environment, when a new database partition is added to the database partition group for an SMS table space, the `ALTER TABLESPACE` statement can be used to add containers to the new database partition.

Each container used for an SMS table space identifies an absolute or relative directory name. Each of these directories can be located on a different file system (or physical disk). The maximum size of the table space can be estimated by:

$$\text{number of containers} * (\text{maximum file system size supported by the operating system})$$

This formula assumes that there is a distinct file system mapped to each container, and that each file system has the maximum amount of space available. In practice, this may not be the case, and the maximum table space size may be much smaller. There are also SQL limits on the size of database objects, which may affect the maximum size of a table space.

Note: Care must be taken when defining the containers. If there are existing files or directories on the containers, an error (SQL0298N) is returned.

- Extent size for the table space.

The extent size can be specified only when the table space is created. Because it cannot be changed later, it is important to select an appropriate value for the extent size.

If you do not specify the extent size when creating a table space, the database manager will create the table space using the default extent size, defined by the `dft_extent_sz` database configuration parameter. This configuration parameter is initially set based on information provided when the database is created. If the `dft_extent_sz` parameter is not specified on the `CREATE DATABASE` command, the default extent size will be set to 32.

To choose appropriate values for the number of containers and the extent size for the table space, you must understand:

- The limitation that your operating system imposes on the size of a logical file system.

For example, some operating systems have a 2 GB limit. Therefore, if you want a 64 GB table object, you will need at least 32 containers on this type of system.

When you create the table space, you can specify containers that reside on different file systems and, as a result, increase the amount of data that can be stored in the database.

- How the database manager manages the data files and containers associated with a table space.

The first table data file (SQL00002.DAT) is created in the first container specified for the table space, and this file is allowed to grow to the extent size. After it reaches this size, the database manager writes data to SQL00002.DAT in the next container. This process continues until all of the containers contain SQL00002.DAT files, at which time the database manager returns to the first container. This process (known as *striping*) continues through the container directories until a container becomes full (SQL0289N), or no more space can be allocated from the operating system (disk full error). Striping applies to the block map files (SQLnnnnn.BKM), to index objects, as well as other objects used to store table data.

Note: The SMS table space is full as soon as any one of its containers is full. Thus, it is important to have the same amount of space available to each container.

To help distribute data across the containers more evenly, the database manager determines which container to use first by taking the table identifier (SQL00002.DAT in the above example) and factoring into account the number of containers. Containers are numbered sequentially, starting at 0.

Database managed space

In a DMS (database managed space) table space, the database manager controls the storage space.

The storage model consists of a limited number of devices or files whose space is managed by the database manager. The database administrator decides which devices and files to use, and manages the space on those devices and files. The table space is essentially an implementation of a special purpose file system designed to best meet the needs of the database manager.

DMS table spaces are different from SMS table spaces in that space for DMS table spaces is allocated when the table space is created. For SMS table spaces, space is allocated as needed. A DMS table space containing user defined tables and data can be defined as a *regular* or *large* table space that stores any table data or index data.

When designing your DMS table spaces and containers, you should consider the following:

- The database manager uses striping to ensure an even distribution of data across all containers. This writes the data evenly across all containers in the table space, placing the extents for tables in round-robin fashion across all containers. DB2 striping is recommended when writing data into multiple containers. If you choose to implement disk striping along with DB2 striping, the extent size of the table space and the strip size of the disk should be identical.
- The maximum size of a regular table space is 512 GB for 32 KB pages. The maximum size of a large table space is 16 TB. See SQL and XML limits for the maximum size of regular table spaces for other page sizes.
- Unlike SMS table spaces, the containers that make up a DMS table space do not need to be the same size; however, this is not normally recommended, because it results in uneven striping across the containers, and sub-optimal performance. If any container is full, DMS table spaces use available free space from other containers.
- Because space is pre-allocated, it must be available before the table space can be created. When using device containers, the device must also exist with enough space for the definition of the container. Each device can have only one container defined on it. To avoid wasted space, the size of the device and the

size of the container should be equivalent. If, for example, the device is allocated with 5 000 pages, and the device container is defined to allocate 3 000 pages, 2 000 pages on the device will not be usable.

- By default, one extent in every container is reserved for overhead. Only full extents are used, so for optimal space management, you can use the following formula to determine an appropriate size to use when allocating a container:

$$\text{extent_size} * (n + 1)$$

where *extent_size* is the size of each extent in the table space, and *n* is the number of extents that you want to store in the container.

- The minimum size of a DMS table space is five extents. Attempting to create a table space smaller than five extents will result in an error (SQL1422N).
 - Three extents in the table space are reserved for overhead.
 - At least two extents are required to store any user table data. (These extents are required for the regular data for one table, and not for any index, long field or large object data, which require their own extents.)
- Device containers must use logical volumes with a “character special interface,” not physical volumes.
- You can use files instead of devices with DMS table spaces. The default table space attribute - NO FILE SYSTEM CACHING in Viper 2 allows files to perform close to devices with the advantage of not requiring to set up devices. For more information, see “Table spaces without file system caching” on page 176.
- If your workload involves LOBs or LONG VARCHAR data, you may derive performance benefits from file system caching.

Note: LOBs and LONG VARCHARs are not buffered by the database manager’s buffer pool.

- Some operating systems allow you to have physical devices greater than 2 GB in size. You should consider dividing the physical device into multiple logical devices, so that no container is larger than the size allowed by the operating system.

Note: Like SMS table spaces, DMS file containers can take advantage of file system prefetching and caching. However, DMS table spaces that use raw device containers cannot.

There are two container options when working with DMS table spaces: raw devices and files. When working with file containers, the database manager allocates the entire container at table space creation time. A result of this initial allocation of the entire table space is that the physical allocation is typically, but not guaranteed to be, contiguous even though the file system is doing the allocation. When working with raw device containers, the database manager takes control of the entire device and always ensures the pages in an extent are contiguous.

When working with DMS table spaces, you should consider associating each container with a different disk. This allows for a larger table space capacity and the ability to take advantage of parallel I/O operations.

The CREATE TABLESPACE statement creates a new table space within a database, assigns containers to the table space, and records the table space definition and attributes in the catalog. When you create a table space, the extent size is defined as a number of contiguous pages. The extent is the unit of space allocation within a table space. Only one table or object, such as an index, can use the pages in any

single extent. All objects created in the table space are allocated extents in a logical table space address map. Extent allocation is managed through space map pages.

The first extent in the logical table space address map is a header for the table space containing internal control information. The second extent is the first extent of SMP for the table space. SMP extents are spread at regular intervals throughout the table space. Each SMP extent is a bit map of the extents from the current SMP extent to the next SMP extent. The bit map is used to track which of the intermediate extents are in use.

The next extent following the SMP is the object table for the table space. The object table is an internal table that tracks which user objects exist in the table space and where their first extent map page (EMP) extent is located. Each object has its own EMPs which provide a map to each page of the object that is stored in the logical table space address map. Figure 8 shows how extents are allocated in a logical table space address map.

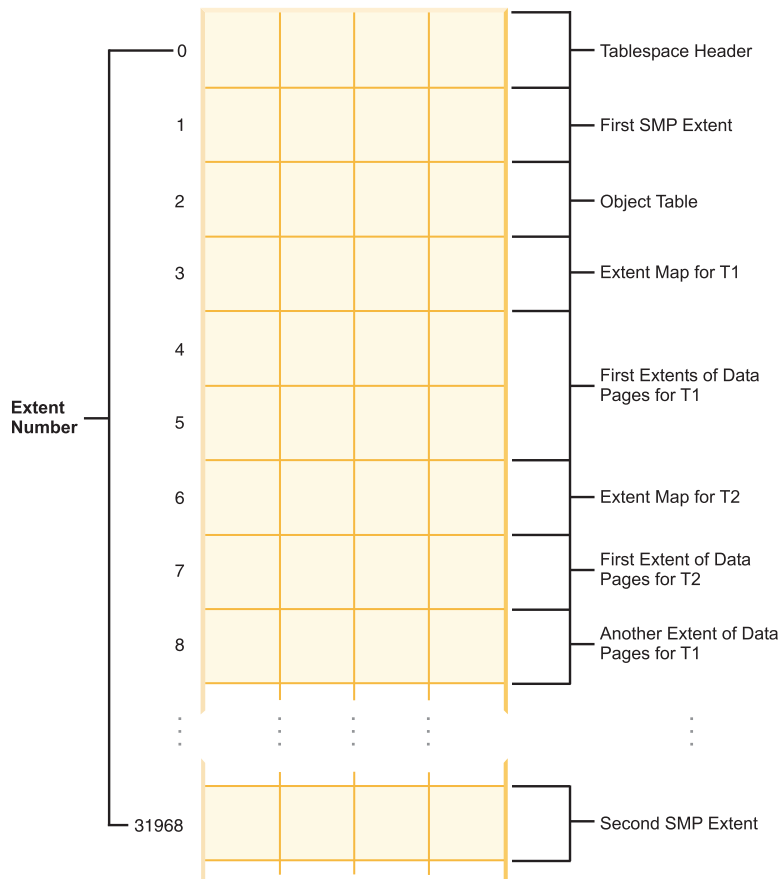


Figure 8. Logical table space address map

DMS table space maps:

A table space map is the database manager’s internal representation of a DMS table space that describes the logical to physical conversion of page locations in a table space. This topic describes why a table space map is useful, and where the information in a table space map comes from.

In a partitioned database, pages in a DMS table space are logically numbered from 0 to (N-1), where N is the number of usable pages in the table space.

The pages in a DMS table space are grouped into extents, based on the extent size, and from a table space management perspective, all object allocation is done on an extent basis. That is, a table might use only half of the pages in an extent but the whole extent is considered to be in use and owned by that object. By default, one extent is used to hold the container tag, and the pages in this extent cannot be used to hold data. However, if the DB2_USE_PAGE_CONTAINER_TAG registry variable is turned on, only one page is used for the container tag.

Figure 9 shows the logical address map for a DMS table space.

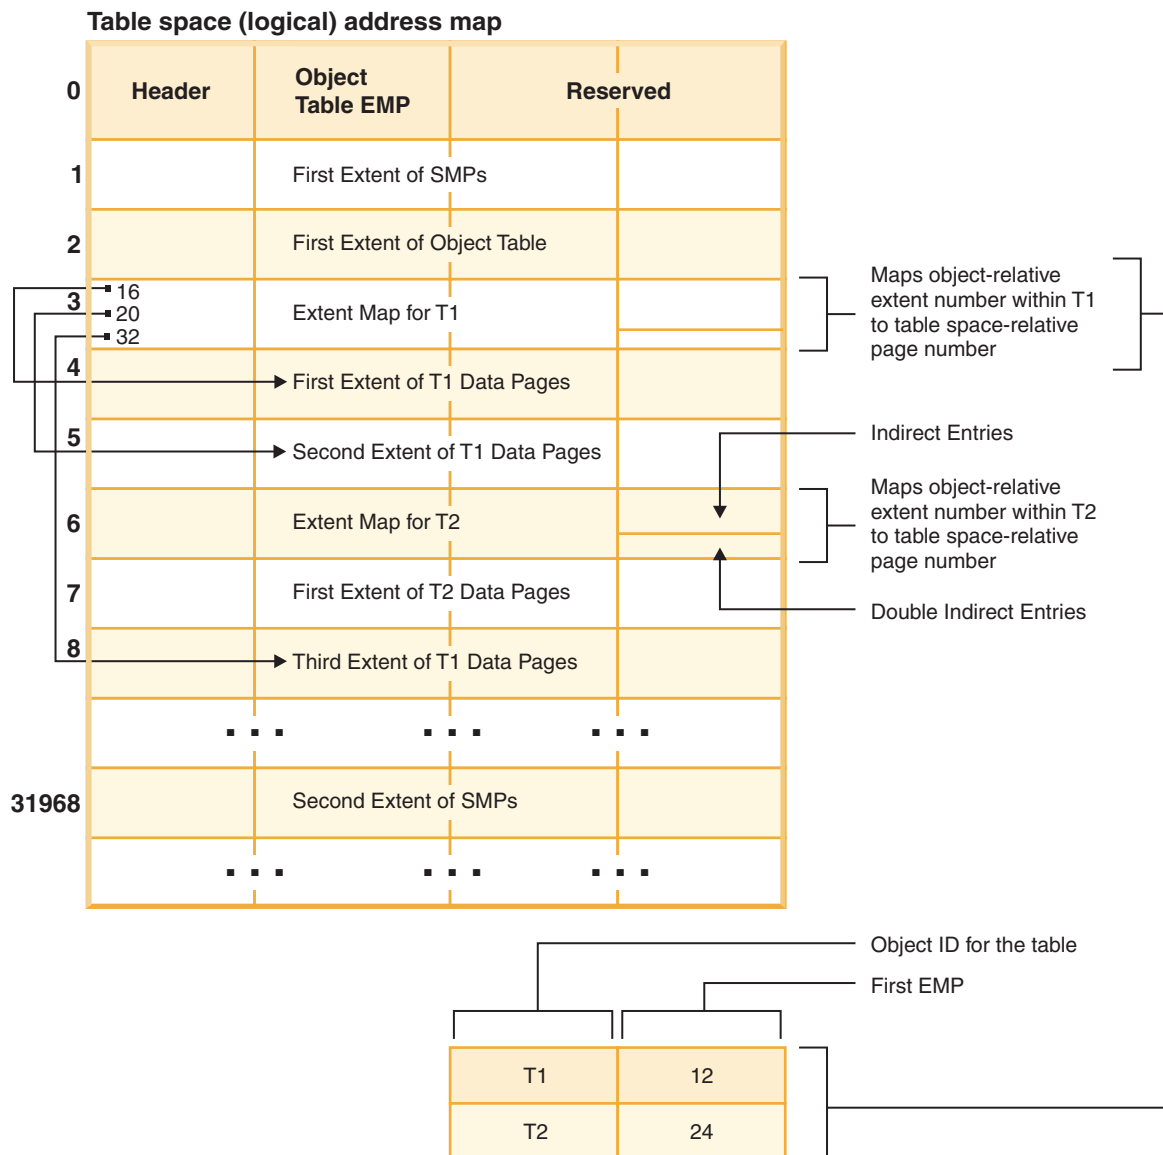


Figure 9. DMS table spaces

Within the table space address map there are two types of map pages: extent map pages (EMP) and space map pages.

The object table is an internal relational table that maps an object identifier to the location of the first EMP extent in the table. This EMP extent, directly or indirectly, maps out all extents in the object. Each EMP contains an array of entries. Each entry maps an object-relative extent number to a table space-relative page number where the object extent is located. Direct EMP entries directly map object-relative addresses to table space-relative addresses. The last EMP page in the first EMP extent contains indirect entries. Indirect EMP entries map to EMP pages which then map to object pages. The last 16 entries in the last EMP page in the first EMP extent contain double-indirect entries.

The extents from the logical table-space address map are striped in round-robin order across the containers associated with the table space.

Because space in containers is allocated by extent, pages that do not make up a full extent will not be used. For example, if you have a 205-page container with an extent size of 10, one extent will be used for the tag, 19 extents will be available for data, and the five remaining pages are wasted.

If a DMS table space contains a single container, the conversion from logical page number to physical location on disk is a straightforward process where pages 0, 1, 2, are located in that same order on disk.

It is also a fairly straightforward process when there is more than one container and each of the containers is the same size. The first extent in the table space, containing pages 0 to (extent size - 1), is located in the first container, the second extent will be located in the second container, and so on. After the last container, the process repeats starting back at the first container. This cyclical process keeps the data balanced.

For table spaces containing containers of different sizes, a simple approach that proceeds through each container in turn cannot be used as it will not take advantage of the extra space in the larger containers. This is where the table space map comes in – it dictates how extents are positioned within the table space, ensuring that all of the extents in the physical containers are available for use.

Note: In the following examples, the container sizes do not take the size of the container tag into account. The container sizes are very small, and are just used for the purpose of illustration, they are not recommended container sizes. The examples show containers of different sizes within a table space, but you are advised to use containers of the same size.

Example 1:

There are 3 containers in a table space, each container contains 80 usable pages, and the extent size for the table space is 20. Each container therefore has 4 extents (80 / 20) for a total of 12 extents. These extents are located on disk as shown in Figure 10 on page 161.

Table space

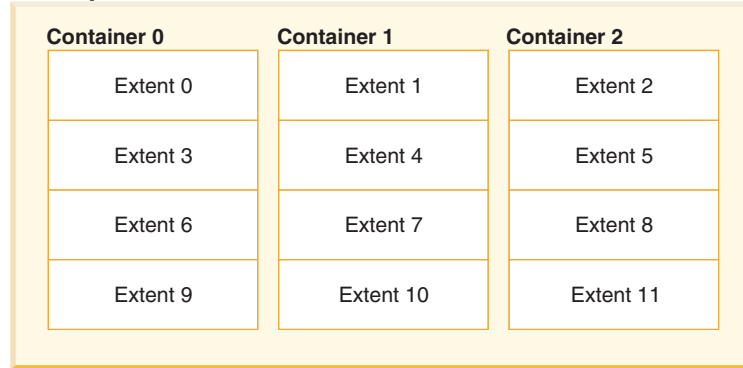


Figure 10. Table space with three containers and 12 extents

To see a table space map, take a table space snapshot using the snapshot monitor. In Example 1, where the three containers are of equal size, the table space map looks like this:

```

Range  Stripe Stripe Max   Max Start End Adj. Containers
Number Set  Offset Extent  Page Stripe Stripe
[0]    [0]   0   11  239   0   3  0  3 (0, 1, 2)

```

A *range* is the piece of the map in which a contiguous range of stripes all contain the same set of containers. In Example 1, all of the stripes (0 to 3) contain the same set of 3 containers (0, 1, and 2) and therefore this is considered a single range.

The headings in the table space map are Range Number, Stripe Set, Stripe Offset, Maximum extent number addressed by the range, Maximum page number addressed by the range, Start Stripe, End Stripe, Range adjustment, and Container list. These will be explained in more detail for Example 2.

This table space can also be diagrammed as shown in Figure 11 on page 162, in which each vertical line corresponds to a container, each horizontal line is called a *stripe*, and each cell number corresponds to an extent.

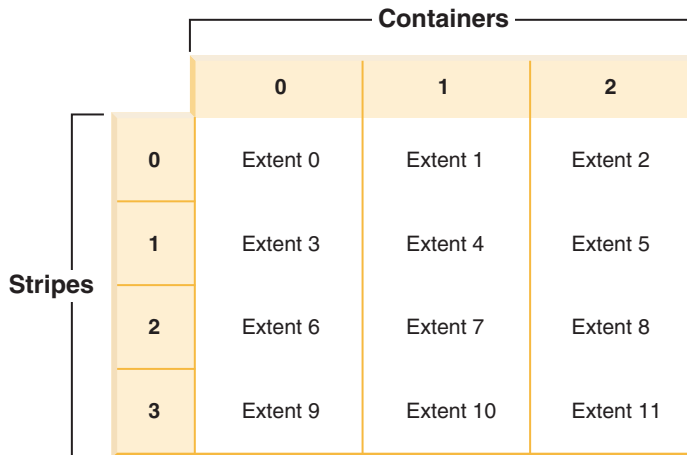


Figure 11. Table space with three containers and 12 extents, with stripes highlighted

Example 2:

There are two containers in the table space: the first is 100 pages in size, the second is 50 pages in size, and the extent size is 25. This means that the first container has four extents and the second container has two extents. The table space can be diagrammed as shown in Figure 12.

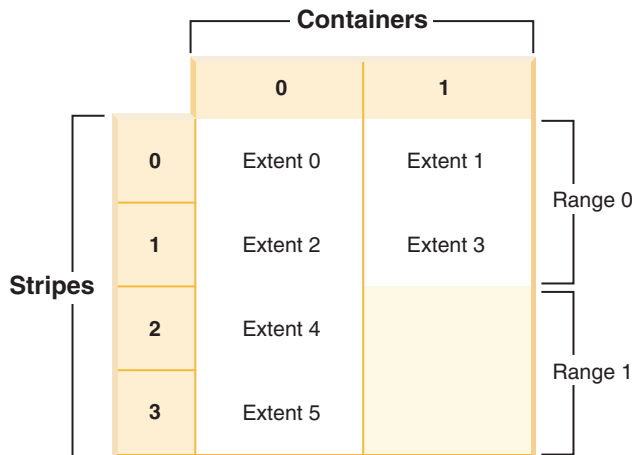


Figure 12. Table space with two containers, with ranges highlighted

Stripes 0 and 1 contain both of the containers (0 and 1) but stripes 2 and 3 only contain the first container (0). Each of these sets of stripes is a range. The table space map, as shown in a table space snapshot, looks like this:

Range Number	Stripe Set	Stripe Offset	Stripe Max Extent	Max Page	Start Stripe	End Stripe	Adj. Stripe	Containers
[0]	[0]	0	3	99	0	1	0 2	(0, 1)
[1]	[0]	0	5	149	2	3	0 1	(0)

There are four extents in the first range, and therefore the maximum extent number addressed in this range (Max Extent) is 3. Each extent has 25 pages and therefore there are 100 pages in the first range. Since page numbering also starts at 0, the maximum page number addressed in this range (Max Page) is 99. The first

stripe (Start Stripe) in this range is 0 and the last stripe (End Stripe) in the range is stripe 1. There are two containers in this range and those are 0 and 1. The stripe offset is the first stripe in the stripe set, which in this case is 0 because there is only one stripe set. The range adjustment (Adj.) is an offset used when data is being rebalanced in a table space. (A rebalance may occur when space is added or dropped from a table space.) When a rebalance is not taking place, this is always 0.

There are two extents in the second range and because the maximum extent number addressed in the previous range is 3, the maximum extent number addressed in this range is 5. There are 50 pages (2 extents * 25 pages) in the second range and because the maximum page number addressed in the previous range is 99, the maximum page number addressed in this range is 149. This range starts at stripe 2 and ends at stripe 3.

Automatic storage table spaces

When you create a table space in a database that is not enabled for automatic storage, you must specify the `MANAGED BY SYSTEM` or `MANAGED BY DATABASE` clause. Using these clauses results in the creation of a system-managed space (SMS) table space or database-managed space (DMS) table space, respectively. You must provide an explicit list of containers in both cases.

If a database is enabled for automatic storage, other choices exist: you can specify the `MANAGED BY AUTOMATIC STORAGE` clause or omit the `MANAGED BY` clause (which implies the use of automatic storage). You do not need to provide container definitions in this case because the database manager assigns the containers automatically.

Following are some examples of statements that create automatic storage table spaces:

```
CREATE TABLESPACE TS1
CREATE TABLESPACE TS2 MANAGED BY AUTOMATIC STORAGE
CREATE TEMPORARY TABLESPACE TEMPTS
CREATE USER TEMPORARY TABLESPACE USRTMP MANAGED BY AUTOMATIC STORAGE
CREATE LONG TABLESPACE LONGTS
```

Although the automatic storage table space type appears to be a different table space type, it is really just an extension of the existing SMS and DMS types. If you create a table space as a regular or large table space, it is created as a DMS table space with file containers. If you create a table space as a user or system temporary table space, it is created as an SMS table space with directory containers.

Note: This behavior might change in future versions of the database manager.

The names associated with these containers have the following format:

```
storage path/instance name/NODE####/database name/T#####/C#####.EXT
```

where:

storage path

Is a storage path associated with the database

instance name

Is the instance under which the database was created

database name

Is the name of the database

NODE####

Is the database partition number (for example, NODE0000)

T#####

Is the table space ID (for example, T0000003)

C#####

Is the container ID (for example, C0000012)

EXT Is an extension based on the type of data being stored:

CAT System catalog table space

TMP System temporary table space

UTM User temporary table space

USR User or regular table space

LRG Large table space

Differences between regular and large automatic storage table spaces and DMS table spaces

Regular and large automatic storage table spaces are created as DMS table spaces, and all of the rules and behaviors associated with DMS table spaces apply. However, there are differences with respect to how storage is managed, as shown in the following table:

Table 45. Differences between managing non-automatic storage and automatic storage table spaces

Non-automatic storage	Automatic storage
You must explicitly provide a list of containers when creating the table space.	You cannot provide a list of containers when creating the table space; instead, the database manager automatically assigns and allocates containers.
Automatic resizing of table spaces is off (AUTORESIZE is set to NO) by default.	Automatic resizing of table spaces is on (AUTORESIZE is set to YES) by default.
You cannot use the INITIALSIZE clause to specify the initial size for the table space.	You can use the INITIALSIZE clause to specify the initial size for the table space.
You can perform container operations using the ALTER TABLESPACE statement (specifying ADD, DROP, BEGIN NEW STRIPE SET, and so on).	You cannot perform container operations because the database manager manages space.
You can use a redirected restore operation to redefine the containers associated with the table space.	You cannot use a redirected restore operation to redefine the containers associated with the table space because the database manager manages space.

As mentioned in the previous table, when you create a regular or large automatic storage table space, you can specify an initial size as part of the CREATE TABLESPACE statement, as shown in the following example:

```
CREATE TABLESPACE TS1 INITIALSIZE 100 M
```

If you do not specify an initial size, the database manager uses a default value of 32 megabytes.

To create a table space with a given size, the database manager creates file containers within the storage paths. If there is an uneven distribution of space among the paths, containers might be created with different sizes. As a result, it is important that all of the storage paths have a similar amount of free space on them.

If you enable automatic resizing for a table space, as space is used within it, the database manager automatically extends existing containers and adds new ones (using stripe sets). Whether containers are extended or added, no rebalancing takes place.

Temporary table spaces

System temporary table spaces hold temporary data required by the database manager while performing operations such as sorts or joins.

These types of operations require extra space to process the results set. A database must have at least one system temporary table space with the same page size as the catalog table space; by default, one system temporary table space called `TEMPSPACE1` is created at database creation time. `IBMTEMPGROUP` is the default database partition group for this table space.

User temporary table spaces hold temporary data from tables created with a `DECLARE GLOBAL TEMPORARY TABLE` statement. To allow the definition of declared temporary tables, at least one user temporary table space should be created with the appropriate `USE` privileges. `USE` privileges are granted using the `GRANT` statement. A user temporary table space is *not* created by default at database creation time.

It is recommended that you define a single temporary table space with a page size equal to the page size used in the majority of your user table spaces. This should be suitable for typical environments and workloads. However, it can be advantageous to experiment with different temporary table space configurations and workloads. The following points should be considered:

- Temporary tables are in most cases accessed in batches and sequentially. That is, a batch of rows are inserted, or a batch of sequential rows are fetched. Therefore, a larger page size typically results in better performance, because fewer logical and physical page requests are required to read a given amount of data.
- When reorganizing a table using a temporary table space, the page size of the temporary table space must match that of the table. For this reason, you should ensure that there are temporary table spaces defined for each different page size used by existing tables that you may reorganize using a temporary table space. You can also reorganize without a temporary table space by reorganizing the table directly in the same table space. Of course, this type of reorganization requires that there be extra space in the table space(s) of the table for the reorganization process.
- When using SMS system temporary table spaces, you may want to consider using the registry variable `DB2_SMS_TRUNC_TMPTABLE_THRESH`. When dropped, files created for the system temporary tables are truncated to a size of 0. The `DB2_SMS_TRUNC_TMPTABLE_THRESH` can be used to avoid visiting the file systems and potentially leave the files at a non-zero size to avoid the performance cost of repeated extensions and truncations of the files.
- In general, when temporary table spaces of different page sizes exist, the optimizer will choose the temporary table space whose buffer pool can hold the most number of rows (in most cases that means the largest buffer pool). In such cases, it is often wise to assign an ample buffer pool to one of the temporary

table spaces, and leave any others with a smaller buffer pool. Such a buffer pool assignment will help ensure efficient utilization of main memory. For example, if your catalog table space uses 4 KB pages, and the remaining table spaces use 8 KB pages, the best temporary table space configuration may be a single 8 KB temporary table space with a large buffer pool, and a single 4 KB table space with a small buffer pool.

- There is generally no advantage to defining more than one temporary table space of any single page size.

Ensuring system temporary table spaces page sizes meet requirements:

The use of larger record identifiers (RID) increases the row size in your result sets from queries or positioned updates. If the row size in your result sets is close to the maximum row length limit for your existing system temporary table spaces, you might need to create a system temporary table space with a larger page size.

Prerequisite

Ensure that you have SYSCTRL or SYSADM authority to create a system temporary table space if required.

Procedure

To ensure that the maximum page size of your system temporary table space is large enough for your queries or positioned updates:

1. Determine the maximum row size in your result sets from queries or positioned updates. Monitor your queries or calculate the maximum row size using the DDL statement that you used to create your tables.
2. List your table spaces using the LIST TABLESPACES command, as shown in the following example:

```
db2 LIST TABLESPACES SHOW DETAIL
...
Tablespace ID           = 1
Name                    = TEMPSPACE1
Type                    = System managed space
Contents                = System Temporary data
State                   = 0x0000
Detailed explanation:
  Normal
Total pages             = 10
Useable pages           = 10
Used pages              = 10
Free pages              = Not applicable
High water mark (pages) = Not applicable
Page size (bytes)       = 4096
Extent size (pages)     = 32
Prefetch size (pages)   = 320
Number of containers    = 10
...
```

You can identify the system temporary table spaces in the output by looking for table spaces whose Contents fields have a value of System Temporary data. Take note of the page size for each of your system temporary table spaces and the page size of the table spaces where the tables referenced in the queries or updates were created.

3. Check whether the largest row size in your result sets fits into your system temporary table space page size:
 - maximum_row_size > maximum_row_length - 8 bytes (structure overhead in single partition)
 - maximum_row_size > maximum_row_length - 16 bytes (structure overhead in DPF)

where `maximum_row_size` is the maximum row size for your result sets, and `maximum_row_length` is the maximum length allowed based on the largest page size of all of your system temporary table spaces. Review the "SQL and XML limits" in *SQL Reference, Volume 1* to determine the maximum row length per table space page size.

If the maximum row size is less than the calculated value then your queries will run in the same manner that they did in DB2 UDB Version 8, and you do not need to continue with this task.

4. Create a system temporary table space that is at least one page size larger than the table space page size where the tables were created if you do not already have a system temporary table with that page size. For example, on the Windows operating systems, if you created your table in a table space with 4 KB page size, create the additional system temporary table space using an 8 KB page size:

```
db2 CREATE SYSTEM TEMPORARY TABLESPACE tmp_tbsp
      PAGESIZE 8K
      MANAGED BY SYSTEM
      USING ('d:\tmp_tbsp', 'e:\tmp_tbsp')
```

If your table space page size is 32 KB, you can reduce the information that you are selecting in your queries or split the queries to fit in the system temporary table space page. For example, if you select all columns from a table, you can instead select only the columns that you really required or a substring of certain columns to avoid exceeding the page size limitation.

Comparison of SMS and DMS table spaces

There are a number of trade-offs to consider when determining which type of table space you should use to store your data.

Advantages of an SMS Table Space:

- Space is not allocated by the system until it is required
- Creating a table space requires less initial work, because you do not have to predefine the containers
- Indexes created on range partitioned data can be stored in a different table space than the table data

Advantages of a DMS Table Space:

- The size of a table space can be increased by adding or extending containers, using the ALTER TABLESPACE statement. Existing data can be automatically rebalanced across the new set of containers to retain optimal I/O efficiency.
- A table can be split across multiple table spaces, based on the type of data being stored:
 - Long field (LF) and large object (LOB) data
 - Indexes
 - Regular table data

You might want to separate your table data for performance reasons, or to increase the amount of data stored for a table. For example, if you are using large table spaces with a 4-KB page size, you can have a table with 2 TB of regular table data, a separate table space with 2 TB of index data, and another separate table space with 2 TB of long data. If these three types of data were stored in one table space instead, the total space would be limited to 2 TB. Using larger page sizes allows you to store even more data. See the related links for the complete list of database manager page size limits.

- Indexes created on range partitioned data can be stored in a different table space than the table data.
- The location of the data on the disk can be controlled, if this is allowed by the operating system.
- In general, a well-tuned set of DMS table spaces will outperform SMS table spaces.

Note: For performance-sensitive applications, particularly those involving a large number of insert operations, it is suggested that you use DMS table spaces.

Also, placement of data can differ on the two types of table spaces. For example, consider the need for efficient table scans: it is important that the pages in an extent are physically contiguous. With SMS, the file system of the operating system decides where each logical file page is physically placed. The pages might be allocated contiguously depending on the level of other activity on the file system and the algorithm used to determine placement. With DMS, however, the database manager can ensure the pages are physically contiguous because it interfaces with the disk directly.

In general, small personal databases are easiest to manage with SMS table spaces. On the other hand, for large, growing databases you will probably only want to use SMS table spaces for the temporary table spaces and catalog table space, and separate DMS table spaces, with multiple containers, for each table. In addition, you will probably want to store long field (LF) data and indexes on their own table spaces.

If you choose to use DMS table spaces with device containers, you must be willing to tune and administer your environment.

SMS and DMS workload considerations

The primary type of workload being managed by the database manager in your environment can affect your choice of what table space type to use, and what page size to specify.

An online transaction processing (OLTP) workload is characterized by transactions that need random access to data, often involve frequent insert or update activity and queries which usually return small sets of data. Given that the access is random, and involves one or a few pages, prefetching is less likely to occur.

DMS table spaces using device containers perform best in this situation. DMS table spaces with file containers, or SMS table spaces, are also reasonable choices for OLTP workloads if maximum performance is not required. Note that using DMS table spaces with file containers, where FILE SYSTEM CACHING is turned off, can perform at a level comparable to DMS raw table space containers. With little or no sequential I/O expected, the settings for the EXTENTSIZE and the PREFETCHSIZE parameters on the CREATE TABLESPACE statement are not important for I/O efficiency. However, setting a sufficient number of page cleaners, using the *chnpggs_thresh* configuration parameter, is important.

A query workload is characterized by transactions that need sequential or partially sequential access to data, which usually return large sets of data. A DMS table space using multiple device containers (where each container is on a separate disk) offers the greatest potential for efficient parallel prefetching. The value of the PREFETCHSIZE parameter on the CREATE TABLESPACE statement should be set to the value of the EXTENTSIZE parameter, multiplied by the number of device containers. Alternatively, you can specify a prefetch size of -1 and the database

manager automatically chooses an appropriate prefetch size. This allows the database manager to prefetch from all containers in parallel. If the number of containers changes, or there is a need to make prefetching more or less aggressive, the PREFETCHSIZE value can be changed accordingly by using the ALTER TABLESPACE statement.

A reasonable alternative for a query workload is to use files, if the file system has its own prefetching. The files can be either of DMS type using file containers, or of SMS type. Note that if you use SMS, you need to have the directory containers map to separate physical disks to achieve I/O parallelism.

Your goal for a mixed workload is to make single I/O requests as efficient as possible for OLTP workloads, and to maximize the efficiency of parallel I/O for query workloads.

The considerations for determining the page size for a table space are as follows:

- For OLTP applications that perform random row read and write operations, a smaller page size is usually preferable because it does not waste buffer pool space with unwanted rows.
- For decision-support system (DSS) applications that access large numbers of consecutive rows at a time, a larger page size is usually better because it reduces the number of I/O requests that are required to read a specific number of rows.
- Larger page sizes may allow you to reduce the number of levels in the index.
- Larger pages support rows of greater length.
- On default 4 KB pages, tables are restricted to 500 columns, while the larger page sizes (8 KB, 16 KB, and 32 KB) support 1012 columns.
- The maximum size of the table space is proportional to the page size of the table space.

SMS and DMS device considerations

There are a few options to consider when choosing to use file system files versus devices for table space containers: the buffering of data and whether to use LOB or LOG data.

- **Buffering of data**

Table data read from disk is usually available in the database buffer pool. In some cases, a data page might be freed from the buffer pool before the application has actually used the page, particularly if the buffer pool space is required for other data pages. For table spaces that use system managed space (SMS) or database managed space (DMS) file containers, file system caching above can eliminate I/O that would otherwise have been required.

Table spaces using database managed space (DMS) device containers do not use the file system or its cache. As a result, you might increase the size of the database buffer pool and reduce the size of the file system cache to offset the fact DMS table spaces that use device containers do not use double buffering.

If system-level monitoring tools show that I/O is higher for a DMS table space using device containers compared to the equivalent SMS table space, this difference might be because of double buffering.

- **Using LOB or LONG data**

When an application retrieves either LOB or LONG data, the database manager does not cache the data in its buffers. Each time an application needs one of these pages, the database manager must retrieve it from disk. However, if LOB or LONG data is stored in SMS or DMS file containers, file system caching might provide buffering and, as a result, better performance.

Because system catalogs contain some LOB columns, you should keep them in SMS table spaces or in DMS-file table spaces.

Considerations when choosing table spaces for your tables

When determining how to map tables to table spaces, you should consider the distribution of your tables, the amount and type of data in the table, and administrative issues.

The distribution of your tables

At a minimum, you should ensure that the table space you choose is in a database partition group with the distribution you want.

The amount of data in the table

If you plan to store many small tables in a table space, consider using SMS for that table space. The DMS advantages with I/O and space management efficiency are not as important with small tables. The SMS advantages, and only when needed, are more attractive with smaller tables. If one of your tables is larger, or you need faster access to the data in the tables, a DMS table space with a small extent size should be considered.

You may wish to use a separate table space for each very large table, and group all small tables together in a single table space. This separation also allows you to select an appropriate extent size based on the table space usage.

The type of data in the table

You may, for example, have tables containing historical data that is used infrequently; the end-user may be willing to accept a longer response time for queries executed against this data. In this situation, you could use a different table space for the historical tables, and assign this table space to less expensive physical devices that have slower access rates.

Alternatively, you may be able to identify some essential tables for which the data has to be readily available and for which you require fast response time. You may want to put these tables into a table space assigned to a fast physical device that can help support these important data requirements.

Using DMS table spaces, you can also distribute your table data across three different table spaces: one for index data; one for large object (LOB) and long field (LF) data; and one for regular table data. This allows you to choose the table space characteristics and the physical devices supporting those table spaces to best suit the data. For example, you could put your index data on the fastest devices you have available, and as a result, obtain significant performance improvements. If you split a table across DMS table spaces, you should consider backing up and restoring those table spaces together if roll-forward recovery is enabled. SMS table spaces do not support this type of data distribution across table spaces.

Administrative issues

Some administrative functions can be performed at the table space level instead of the database or table level. For example, taking a backup of a table space instead of a database can help you make better use of your time and resources. It allows you to frequently back up table spaces with large volumes of changes, while only occasionally backing up table spaces with very low volumes of changes.

You can restore a database or a table space. If unrelated tables do not share table spaces, you have the option to restore a smaller portion of your database and reduce costs.

A good approach is to group related tables in a set of table spaces. These tables could be related through referential constraints, or through other defined business constraints.

If you need to drop and redefine a particular table often, you may want to define the table in its own table space, because it is more efficient to drop a DMS table space than it is to drop a table.

Automatic re-sizing of table spaces

Enabling automatic storage table spaces for automatic resizing allows the database manager to handle the full file system condition automatically by adding a new stripe set of containers.

Two table space types can exist within a database system: system-managed space (SMS) and database-managed space (DMS). The containers associated with SMS table spaces are file system directories, and the files within these directories grow as the objects in the table space grow. The files grow until a file system limit is reached for one of the containers or until the table space size limit of the database is reached (see SQL and XML limits).

DMS table spaces are made up of file containers or raw device containers, and their sizes are set when the containers are assigned to the table space. The table space is considered to be full when all of the space within the containers has been used. However, unlike for SMS table spaces, you can add or extend containers using the ALTER TABLESPACE statement, allowing more storage space to be given to the table space. DMS table spaces also have a feature called *auto-resize*: as space is consumed in a DMS table space that can be automatically resized, the database system might extend the table space by one or more file containers. SMS table spaces have similar capabilities for growing automatically, but the term *auto-resize* is used exclusively for DMS.

Automatic resizing of table spaces has the following implications:

- Table spaces that are enabled for automatic resizing have metadata associated with them that is not recognized by Version 8.2.1 or earlier releases. Any attempt to use a database with table spaces enabled for automatic resizing with these versions results in a failure (most likely, returning an SQL0980C or SQL0902C error). An error might be sent if you try to connect to a database or try to restore a database. If you enabled table spaces for automatic resizing, disabling the *auto-resize* feature for these table spaces removes the metadata, allowing the database to be used with Version 8.2.1 or earlier releases.
- If you disable the *auto-resize* feature, the values that are associated with INCREASESIZE and MAXSIZE are lost if you subsequently enable this feature.
- You cannot this feature for table spaces that use raw device containers, and you cannot add raw device containers to a table space that can be automatically resized. Attempting these operations results in errors (SQL0109N). If you need to add raw device containers, you must disable the feature first.
- A redirected restore operation cannot change the container definitions to include a raw device container. Attempting this kind of operation results in an error (SQL0109N).
- Because the maximum size limits how the database manager automatically increases a table space, the maximum size also limits how you can increase a

table space. That is, when you perform an operation that adds space to a table space, the resulting size must be less than or equal to the maximum size. You can add space by using the ADD, EXTEND, RESIZE, or BEGIN NEW STRIPE SET clause of the ALTER TABLESPACE statement.

Enabling and disabling the auto-resize feature

By default, the auto-resize feature is not enabled for a DMS table space. The following statement creates a DMS table space without enabling auto-resize:

```
CREATE TABLESPACE DMS1 MANAGED BY DATABASE
  USING (FILE '/db2files/DMS1' 10 M)
```

To enable the auto-resize feature, specify the AUTORESIZE YES clause for the CREATE TABLESPACE statement:

```
CREATE TABLESPACE DMS1 MANAGED BY DATABASE
  USING (FILE '/db2files/DMS1' 10 M) AUTORESIZE YES
```

You can also enable or disable the auto-resize feature after creating a DMS table space by using ALTER TABLESPACE statement with the AUTORESIZE clause:

```
ALTER TABLESPACE DMS1 AUTORESIZE YES
ALTER TABLESPACE DMS1 AUTORESIZE NO
```

Two other attributes, MAXSIZE and INCREASESIZE, are associated with auto-resize table spaces:

Maximum size (MAXSIZE)

The MAXSIZE clause of the CREATE TABLESPACE statement defines the maximum size for the table space. For example, the following statement creates a table space that can grow to 100 megabytes (per database partition if the database has multiple database partitions):

```
CREATE TABLESPACE DMS1 MANAGED BY DATABASE
  USING (FILE '/db2files/DMS1' 10 M)
  AUTORESIZE YES MAXSIZE 100 M
```

The MAXSIZE NONE clause specifies that there is no maximum limit for the table space. The table space can grow until a file system limit or table space limit is reached (see SQL and XML limits). If you do not specify the MAXSIZE clause, there is no maximum limit when the auto-resize feature is enabled.

Use the ALTER TABLESPACE statement to change the value of MAXSIZE for a table space that has auto-resize already enabled, as shown in the following examples:

```
ALTER TABLESPACE DMS1 MAXSIZE 1 G
ALTER TABLESPACE DMS1 MAXSIZE NONE
```

If you specify a maximum size, the actual value that the database manager enforces might be slightly smaller than the value specified because the database manager attempts to keep container growth consistent.

Increase size (INCREASESIZE)

The INCREASESIZE clause of the CREATE TABLESPACE statement defines the amount of space used to increase the table space when there are no free extents

within the table space but a request for one or more extents was made. You can specify the value as an explicit size or as a percentage, as shown in the following examples:

```
CREATE TABLESPACE DMS1 MANAGED BY DATABASE
  USING (FILE '/db2files/DMS1' 10 M)
  AUTORESIZE YES INCREASESIZE 5 M
```

```
CREATE TABLESPACE DMS1 MANAGED BY DATABASE
  USING (FILE '/db2files/DMS1' 10 M)
  AUTORESIZE YES INCREASESIZE 50 PERCENT
```

A percentage value means that the amount by which to increase, as specified by the INCREASESIZE value, is calculated every time that the table space needs to grow and that growth is based on a percentage of the table space size at that time. For example, if the table space is 20 MB in size and the INCREASESIZE value is 50% , the table space grows by 10 MB the first time (to a size of 30 MB) and by 15 MB the next time.

If you do not specify the INCREASESIZE clause when you enable the auto-resize feature, the database manager determines an appropriate value to use, which might change over the life of the table space. As with AUTORESIZE and MAXSIZE, you can change the value of INCREASESIZE using the ALTER TABLESPACE statement.

If you specify a size increase, the actual value that the database manager will use might be slightly different than the value that you provide. This adjustment in the value used is done to keep growth consistent across the containers in the table space.

How table spaces are extended

For table spaces that can be automatically resized, the database manager attempts to increase the size of the table space when all of the existing space has been used and a request for more space is made. The database manager determines which of the containers can be extended in the table space so that a rebalance does not occur. The database manager extends only those containers that exist within the last range of the table space map (the map describes the storage layout for the table space) and extends them by an equal amount.

For example, consider the following statement:

```
CREATE TABLESPACE TS1 MANAGED BY DATABASE
  USING (FILE 'C:\TS1CONT' 1000, FILE 'D:\TS1CONT' 1000,
        FILE 'E:\TS1CONT' 2000, FILE 'F:\TS1CONT' 2000)
  EXTENTSIZE 4
  AUTORESIZE YES
```

Keeping in mind that the database manager uses a small portion (one extent) of each container for metadata, following is the table space map that is created for the table space based on the CREATE TABLESPACE statement. (The table space map is part of the output from a table space snapshot.)

Table space map:

Range Number	Stripe Set	Stripe Offset	Max Extent	Max Page	Start Stripe	End Stripe	Adj.	Containers
[0]	[0]	0	995	3983	0	248	0	4 (0,1,2,3)
[1]	[0]	0	1495	5983	249	498	0	2 (2,3)

The table space map shows that the containers with an identifier of 2 or 3 (E:\TS1CONT and F:\TS1CONT) are the only containers in the last range of the map. Therefore, when the database manager automatically extends the containers in this table space, it extends only those two containers.

Note: If you create a table space with all of the containers having the same size, there is only one range in the map. In such a case, the database manager extends each of the containers. To prevent restricting extensions to only a subset of the containers, create a table space with containers of equal size.

As discussed previously, you can specify a limit on the maximum size of the table space, or you can specify a value of NONE, which does not limit growth. If you specify NONE or no limit, the upper limit is defined by the file system limit or by the table space limit; the database manager does not attempt to increase the table space size past the upper limit. However, before that limit is reached, an attempt to increase a container might fail due to a full file system. In this case, the database manager does not increase the table space size any further and returns an out-of-space condition to the application. There are two ways to resolve this situation:

- Increase the amount of space available on the file system that is full.
- Perform container operations on the table space such that the container in question is no longer in the last range of the table space map. The easiest way to do this is to add a new stripe set to the table space with a new set of containers, and the best practice is to ensure that the containers are all the same size. You can add new stripe sets by using the ALTER TABLESPACE statement with the BEGIN NEW STRIPE SET clause. By adding a new stripe set, a new range is added to the table space map. With a new range, the containers that the database manager automatically attempts to extend are within this new stripe set, and the older containers remain unchanged.

Note: When a user-initiated container operation is pending or a subsequent rebalance is in progress, the auto-resize feature is disabled until the operation is committed or the rebalance is complete.

For example, for DMS table spaces, suppose that a table space has three containers that are the same size and that each resides on its own file system. As work is done on the table space, the database manager automatically extends these three containers. Eventually, one of the file systems becomes full, and the corresponding container can no longer grow. If more free space cannot be made available on the file system, you must perform container operations on the table space such that the container in question is no longer in the last range of the table space map. In this case, you could add a new stripe set specifying two containers (one on each of the file systems that still has space), or you could specify more or fewer containers (again, making sure that each container being added is the same size and that there is sufficient room for growth on each of the file systems being used). When the database manager attempts to increase the size of the table space, it now attempts to extend the containers in this new stripe set instead of attempting to extend the older containers.

Monitoring

Information about automatic resizing for DMS table spaces is displayed as part of the table space monitor snapshot output. The increase size and maximum size values are included in the output, as shown in the following sample:

Auto-resize enabled	= Yes or No
Current tablespace size (bytes)	= ###
Maximum tablespace size (bytes)	= ### or NONE
Increase size (bytes)	= ###
Increase size (percent)	= ###
Time of last successful resize	= DD/MM/YYYY HH:MM:SS.SSSSSS
Last resize attempt failed	= Yes or No

Automatic prefetchsize adjustment after adding or dropping containers

The database manager is set up so that the automatic prefetch size is the default for any table spaces created using Version 8.2 (and later).

You should not have to worry about remembering to adjust the prefetch size after you add or drop containers. The default for the database manager to automatically adjust the prefetch size removes this worry from you. If there is the possibility that you might forget to update the prefetch size of a table space after either adding or dropping containers, do not change the default but allow the prefetch size to be determined by the database manager automatically. If you do not allow automatic prefetch size adjustment by changing the default and you forget to update the prefetch size, then there might be a noticeable degradation in the performance of the database.

There are three ways not to have the prefetch size of the table space set at AUTOMATIC:

- Create the table space with a specific prefetch size. Manually choosing a value for the prefetch size indicates that you will remember to adjust, if necessary, the prefetch size whenever there is an adjustment in the number of containers associated with the table space.
- Do not use prefetch size when creating the table space, and have the *dft_prefetch_sz* database configuration parameter set to a non-AUTOMATIC value. The database manager checks this parameter when there is no explicit mention of the prefetch size when creating the table space. If a value other than AUTOMATIC is found, then that value is what is used as the default prefetch size. And you will need to remember to adjust, if necessary, the prefetch size whenever there is an adjustment in the number of containers associated with the table space.
- Alter the prefetch size manually by using the ALTER TABLESPACE statement.

Use of DB2_PARALLEL_IO

Prefetch requests are broken down into several smaller prefetch requests based on the parallelism of a table space, and before the requests are submitted to the prefetch queues. The DB2_PARALLEL_IO registry variable is used to define the number of physical spindles per container as well as influencing the parallel I/O on the table space. With parallel I/O off, the parallelism of a table space is equal to the number of containers. With parallel I/O on, the parallelism of a table space is equal to the number of container multiplied by the value given in the DB2_PARALLEL_IO registry variable. (Another way of saying this is, the parallelism of the table space is equal to the prefetch size divided by the extent size of the table space.)

Here are several examples of how the DB2_PARALLEL_IO registry variable influences the prefetch size. (Assume all of the following table spaces have been defined with an AUTOMATIC prefetch size.)

- `DB2_PARALLEL_IO=*`
 - All table spaces will use the default where the number of spindles equals 6 for each container. The prefetch size will be 6 times larger with parallel I/O on.
 - All table spaces will have parallel I/O on. The prefetch request is broken down to several smaller requests, each equal to the prefetch size divided by the extent size (or equal to the number of containers times the number of spindles).
- `DB2_PARALLEL_IO=*:3`
 - All table spaces will use 3 as the number of spindles per container.
 - All table spaces will have parallel I/O on.
- `DB2_PARALLEL_IO=*:3,1:1`
 - All table spaces will use 3 as the number of spindles per container except for table space 1 which will use 1.
 - All table spaces will have parallel I/O on.

Table spaces without file system caching

The recommended method of enabling or disabling non-buffered I/O on UNIX, Linux, and Windows is at the table space level.

This allows you to enable or disable non-buffered I/O on specific table spaces while avoiding any dependency on the physical layout of the database. It also allows the database manager to determine which I/O is best suited for each file, buffered or non-buffered.

The `NO FILE SYSTEM CACHING` clause is used to enable non-buffered I/O, thus disabling file caching for a particular table space. Once enabled, based on platform, the database manager automatically determines which of the Direct I/O (DIO) or Concurrent I/O (CIO) is to be used. Given the performance improvement in CIO, the database manager uses it whenever it is supported; there is no user interface to specify which one is to be used.

In order to obtain the maximum benefits of non-buffered I/O, it might be necessary to increase the size of buffer pools. However, if the self-tuning memory manager is enabled and the buffer pool size is set to `AUTOMATIC`, the database manager will self-tune the buffer pool size for optimal performance. Note that this feature is not available prior to Version 9.

To disable or enable file system caching, specify the `NO FILE SYSTEM CACHING` or the `FILE SYSTEM CACHING` clause in the `CREATE TABLESPACE` or `ALTER TABLESPACE` statement, respectively. The default setting is used if neither clause is specified. In the case of `ALTER TABLESPACE`, existing connections to the database must be terminated before the new caching policy takes effect.

Note: If an attribute is altered from the default to either `FILE SYSTEM CACHING` or `NO FILE SYSTEM CACHING`, there is no mechanism to change it back to the default.

This method of enabling and disabling file system caching provides control of the I/O mode, buffered or non-buffered, at the table space level.

Note: I/O access to long field (LF) data and large object (LOB) data will be buffered for both SMS and DMS containers, regardless of the setting for the table space in question.

The GET SNAPSHOT FOR TABLESPACES command can be used to query the current setting of the file system caching clause. For example, the following is a snippet from the DB2 GET SNAPSHOT FOR TABLESPACES ON db1 output:

```
Tablespace name           = USERSPACE1
Tablespace ID             = 2
Tablespace Type           = database managed space
Tablespace Content Type   = All permanent data. Large table space.
Tablespace Page size (bytes) = 4096
Tablespace Extent size (pages) = 32
Automatic Prefetch size enabled = Yes
Buffer pool ID currently in use = 1
Buffer pool ID next startup = 1
Using automatic storage = Yes
Auto-resize enabled = Yes
File system caching = No
Tablespace State = 0x'00000000'
  Detailed explanation:
    Normal
Tablespace Prefetch size (pages) = 32
Total number of pages = 256
```

Alternate methods to enable/disable non-buffered I/O on UNIX, Linux, and Windows

Some UNIX platforms support the disabling of file system caching at a file system level by using the MOUNT option. Consult your operating system documentation for more information. However, it is important to understand the difference between disabling file system caching at the table space level and at the file system level. At the table space level, the database manager controls which files are to be opened with and without file system caching. At the file system level, every file residing on that particular file system will be opened without file system caching. Some platforms such as AIX have certain requirements before you can use this feature, such as serialization of read and write access. While the database manager adheres to these requirements, if the target file system contains non-DB2 files, before enabling this feature, consult your operating system documentation for any requirements.

Note: The now-deprecated registry variable DB2_DIRECT_IO, introduced in Version 8.1 FixPak 4, enables no file system caching for all SMS containers except for long field data, large object data, and temporary table spaces on AIX JFS2. Setting this registry variable in Version 9.1 or later is equivalent to altering all table spaces, SMS and DMS, with the NO FILE SYSTEM CACHING clause. However, using DB2_DIRECT_IO is not recommended, and this variable will be removed in a later release. Instead, you should enable NO FILE SYSTEM CACHING at the table space level.

Alternate methods to enable/disable non-buffered I/O on Windows

In previous releases, the performance registry variable DB2NTNOCACHE could be used to disable file system caching for all DB2 files in order to make more memory available to the database so that the buffer pool or sortheap can be increased. In Version 9.5, DB2NTNOCACHE is deprecated and might be removed in a future release. The difference between DB2NTNOCACHE and using the NO FILE SYSTEM CACHING clause is the ability to disable caching for selective table spaces. Starting in Version 9.5, since the NO FILE SYSTEM CACHING is used as the default, unless

FILE SYSTEM CACHING is specified explicitly, there is no need to set this registry variable to disable file system caching across the entire instance if the instance includes only newly created table spaces.

Performance considerations

Non-buffered I/O is essentially used for performance improvements. In some cases, however, performance degradation might be due to, but is not limited to, a combination of a small buffer pool size and a small file system cache. Suggestions for improving performance include:

- If self-tuning memory manager is not enabled, enable it and set the buffer pool size to automatic using ALTER BUFFERPOOL <name> SIZE AUTOMATIC. This allows the database manager to self-tune the buffer pool size.
- If self-tuning memory manager is not to be enabled, increase the buffer pool size in increments of 10 or 20 percent until performance is improved.
- If self-tuning memory manager is not to be enabled, alter the table space to use "FILE SYSTEM CACHING". This essentially disables the non-buffered I/O and reverts back to buffered I/O for container access.

Performance tuning should be tested in a controlled environment before implementing it on the production system.

When choosing to use file system files versus devices for table space containers, you should consider file system caching, which is performed as follows:

- For DMS file containers (and all SMS containers), the operating system might cache pages in the file system cache (unless the table space is defined with NO FILESYSTEM CACHING).
- For DMS device container table spaces, the operating system does not cache pages in the file system cache.

Using CIO/DIO as the default file system caching mechanism for new table space containers

The default I/O mechanism for newly created table space containers on most AIX, Linux, Solaris, and Windows platforms is CIO/DIO (concurrent I/O or Direct I/O). This default provides an increase of throughput over buffered I/O on heavy transaction processing workloads and rollbacks.

The FILE SYSTEM CACHING or NO FILE SYSTEM CACHING attribute specifies whether or not I/O operations are to be cached at the file system level:

- FILE SYSTEM CACHING specifies that all I/O operations in the target table space are to be cached at the file system level.
- NO FILE SYSTEM CACHING specifies that all I/O operations are to bypass the file system-level cache.

Note: When using DMS table spaces, you should use a separate table space for long field (LF) data and for large object (LOB) data so that the regular table spaces are not affected. (For SMS table spaces, the CIO/DIO (NO FILE SYSTEM CACHING) attribute is disabled.)

The following interfaces contain the FILE SYSTEM CACHING attribute:

- CREATE TABLESPACE statement
- CREATE DATABASE command
- sqlcrea() API (using the *sqlfscaching* field of the SQLETSDESC structure)

When this attribute is not specified on the CREATE TABLESPACE statement, or on the CREATE DATABASE command, the database manager processes the request using the default behaviour based on the platform and file system type. See “File system caching configurations” for the exact behavior. For the sqlcrea() API, a value of 0x2 for the field *sqlfscaching* field, instructs the database manager to use the default setting.

Note that the following tools currently interpret the value for FILE SYSTEM CACHING attribute:

- GET SNAPSHOT FOR TABLESPACES command
- db2pd -tablespaces command
- db2look -d <dbname> -l command

For db2look, if the FILE SYSTEM CACHING attribute is not specified, the output does not contain this attribute.

Example

Suppose that the database and all related table space containers reside on an AIX JFS file system and the following statement was issued:

```
DB2 CREATE TABLESPACE JFS2
```

In previous versions, if the attribute was not specified, the database manager would have used buffered I/O (FILE SYSTEM CACHING) for the I/O mechanism; with Version 9.5, the database manager uses NO FILE SYSTEM CACHING.

File system caching configurations

The operating system, by default, caches file data that is read from and written to disk.

A typical read operation involves physical disk access to read the data from disk into the file system cache, and then to copy the data from the cache to the application buffer. Similarly, a write operation involves physical disk access to copy the data from the application buffer into the file system cache, and then to copy it from the cache to the physical disk. This behavior of caching data at the file system level is reflected in the FILE SYSTEM CACHING clause of the CREATE TABLESPACE statement. Since the database manager manages its own data caching using buffer pools, the caching at the file system level is not needed if the size of the buffer pool is tuned appropriately.

Note: The database manager already prevents caching of most DB2 data, except temporary data and LOBs on AIX, by invalidating the pages from the cache.

In some cases, caching at the file system level and in the buffer pools causes performance degradation because of the extra CPU cycles required for the double caching. To avoid this double caching, most file systems have a feature that disables caching at the file system level. This is generically referred to as *non-buffered I/O*. On UNIX, this feature is commonly known as *Direct I/O (or DIO)*. On Windows, this is equivalent to opening the file with the FILE_FLAG_NO_BUFFERING flag. In addition, some file systems such as IBM JFS2 or Symantec VERITAS VxFS also support enhanced Direct I/O, that is, the higher-performing *Concurrent I/O (CIO)* feature. The database manager supports this feature with the NO FILE SYSTEM CACHING table space clause. When this is set, the database manager automatically takes advantage of CIO on file systems

where this feature exists. This feature might help to reduce the memory requirements of the file system cache, thus making more memory available for other uses.

Prior to Version 9.5, the keyword FILE SYSTEM CACHING was implied if neither NO FILE SYSTEM CACHING nor FILE SYSTEM CACHING was specified. With Version 9.5, if neither keyword is specified, the default, NO FILE SYSTEM CACHING, is used. This change affects only newly created table spaces. Existing table spaces created prior to Version 9.5 are not affected. This change applies to AIX, Linux, Solaris, and Windows with the following exceptions, where the default behavior remains to be FILE SYSTEM CACHING:

- AIX JFS
- Solaris non-VxFS
- Linux for System z™
- All SMS temporary table space files
- Long Field (LF) and Large object (LOB) data files in SMS permanent table space files

To override the default setting, specify FILE SYSTEM CACHING or NO FILE SYSTEM CACHING.

Supported configurations

Table 46 shows the supported configuration for using table spaces without file system caching. It also indicates: (a) whether DIO or enhanced DIO will be used in each case, and (b) the default behavior when neither NO FILE SYSTEM CACHING nor FILE SYSTEM CACHING is specified for a table space based on the platform and file system type.

Table 46. Supported configurations for table spaces without file system caching

Platforms	File system type and minimum level required	DIO or CIO requests submitted by the database manager when NO FILE SYSTEM CACHING is specified	Default behavior when neither NO FILE SYSTEM CACHING nor FILE SYSTEM CACHING is specified
AIX 5.3+	Journal File System (JFS)	DIO	FILE SYSTEM CACHING (See Note 1.)
AIX 5.3+	Concurrent Journal File System (JFS2)	CIO	NO FILE SYSTEM CACHING
AIX 5.3+	VERITAS Storage Foundation for DB2 4.1 (VxFS)	CIO	NO FILE SYSTEM CACHING
HP-UX 11i (PA-RISC)	VERITAS Storage Foundation 4.1 (VxFS)	CIO	FILE SYSTEM CACHING
HP-UX Version 11i v2 (Itanium®)	VERITAS Storage Foundation 4.1 (VxFS)	CIO	FILE SYSTEM CACHING
Solaris 9	UNIX File System (UFS)	DIO	FILE SYSTEM CACHING (See Note 2.)
Solaris 10	UNIX File System (UFS)	CIO	FILE SYSTEM CACHING (See Note 2.)
Solaris 9, 10	VERITAS Storage Foundation for DB2 4.1 (VxFS)	CIO	NO FILE SYSTEM CACHING

Table 46. Supported configurations for table spaces without file system caching (continued)

Platforms	File system type and minimum level required	DIO or CIO requests submitted by the database manager when NO FILE SYSTEM CACHING is specified	Default behavior when neither NO FILE SYSTEM CACHING nor FILE SYSTEM CACHING is specified
Linux distributions SLES 9+ and RHEL 4+ (on these architectures: x86, x86_64, IA64, POWER™)	ext2, ext3, reiserfs	DIO	NO FILE SYSTEM CACHING
Linux distributions SLES 9+ and RHEL 4+ (on these architectures: x86, x86_64, IA64, POWER)	VERITAS Storage Foundation 4.1 (VxFS)	CIO	NO FILE SYSTEM CACHING
Linux distributions SLES 9+ and RHEL 4+ (on this architecture: zSeries®)	ext2, ext3 or reiserfs on a Small Computer System Interface (SCSI) disks using Fibre Channel Protocol (FCP)	DIO	FILE SYSTEM CACHING
Windows	No specific requirement, works on all DB2 supported file systems	DIO	NO FILE SYSTEM CACHING

Note:

1. On AIX JFS, FILE SYSTEM CACHING is the default.
2. On Solaris UFS, NO FILE SYSTEM CACHING is the default.
3. The VERITAS Storage Foundation for the database manager might have different operating system prerequisites. The platforms listed above are the supported platforms for the current release. Consult the VERITAS Storage Foundation for DB2 support for prerequisite information.
4. If SFDB2 5.0 is used instead of the above minimum levels, the SFDB2 5.0 MP1 RP1 release must be used. This release includes fixes that are specific to the 5.0 version.
5. If you do not want the database manager to choose NO FILE SYSTEM CACHING for the default setting, specify FILE SYSTEM CACHING in the relevant SQL, commands, or APIs.

Examples

Example 1: By default, this new table space will be created using non-buffered I/O; the NO FILE SYSTEM CACHING clause is implied:

```
CREATE TABLESPACE table space name ...
```

Example 2: On the following statement, the NO FILE SYSTEM CACHING clause indicates that file system level caching will be OFF for this particular table space:

```
CREATE TABLESPACE table space name ... NO FILE SYSTEM CACHING
```

Example 3: The following statement disables file system level caching for an existing table space:

```
ALTER TABLESPACE table space name ... NO FILE SYSTEM CACHING
```


Example 4: The following statement enables file system level caching for an existing table space:

```
ALTER TABLESPACE table space name ... FILE SYSTEM CACHING
```

Table space extent sizes

The extent size for a table space represents the number of pages of table data that will be written to a container before data will be written to the next container.

When selecting an extent size, you should consider:

- The size and type of tables in the table space.

Space in DMS table spaces is allocated to a table one extent at a time. As the table is populated and an extent becomes full, a new extent is allocated. DMS table space container storage is prereserved which means that new extents are allocated until the container is completely used.

Space in SMS table spaces is allocated to a table either one extent at a time or one page at a time. As the table is populated and an extent or page becomes full, a new extent or page is allocated until all of the extents or pages in the file system are used. When using SMS table spaces, multipage file allocation is allowed. Multipage file allocation allows extents to be allocated instead of a page at a time.

Multipage file allocation is enabled by default. The value of the *multipage_alloc* database configuration parameter will indicate if multipage file allocation is enabled.

Note: Multipage file allocation is not applicable to temporary table spaces.

A table is made up of the following separate table objects:

- A data object. This is where the regular column data is stored.
- An index object. This is where all indexes defined on the table are stored.
- A long field (LF) data object. This is where long field data, if your table has one or more LONG columns, is stored.
- Two large object (LOB) data objects. If your table has one or more LOB columns, they are stored in these two table objects:
 - One table object for the LOB data
 - A second table object for metadata describing the LOB data.
- A block map object for multidimensional clustering (MDC) tables.
- An extra XDA object, which stores XML documents.

Each table object is stored separately, and each object allocates new extents as needed. Each DMS table object is also paired with a metadata object called an extent map, which describes all of the extents in the table space that belong to the table object. Space for extent maps is also allocated one extent at a time. Therefore, the initial allocation of space for an object in a DMS table space is two extents. (The initial allocation of space for an object in an SMS table space is one page.)

If you have many small tables in a DMS table space, you may have a relatively large amount of space allocated to store a relatively small amount of data. In such a case, you should specify a small extent size. On the other hand, if you have a very large table that has a high growth rate, and you are using a DMS table space with a small extent size, you could have unnecessary overhead related to the frequent allocation of additional extents.

- The type of access to the tables.

If access to the tables includes many queries or transactions that process large quantities of data, prefetching data from the tables may provide significant performance benefits.

- The minimum number of extents required.

If there is not enough space in the containers for five extents of the table space, the table space will not be created.

Table space page sizes

When designing table spaces, you need to consider page sizes.

You can use a 4K, 8K, 16K or 32K page size limit. Each of these page sizes also has maximums for each of the table space types that you must adhere to.

Table 47 shows the page size-specific limits for the different types of table spaces:

Table 47. Table space page size-specific limits

Table space type (in gigabytes)	4K page size limit	8K page size limit	16K page size limit	32K page size limit
SMS table spaces	64	128	256	512
DMS table spaces (regular)	64	128	256	512
DMS table spaces (large)	2048	4096	8192	16 384
Automatic storage table spaces (regular)	64	128	256	512
Automatic storage table spaces (large)	2048	4096	8192	16 384
Temporary table spaces	64	128	256	512

For database and index page size limits for the different types of table spaces, see the database manager page size-specific limits in SQL and XML limits.

To ensure that the maximum page size of your system temporary table space is large enough for your queries or positioned updates, see “Ensuring system temporary table space page sizes meet requirements” in *Migration Guide*.

Table space disk I/O

The type and design of your table space determines the efficiency of the I/O performed against that table space.

You should understand the following concepts before considering other issues concerning table space design and use:

Big-block reads

A read where several pages (usually an extent) are retrieved in a single request. Reading several pages at once is more efficient than reading each page separately.

Prefetching

The reading of pages in advance of those pages being referenced by a query. The overall objective is to reduce response time. This can be achieved if the prefetching of pages can occur asynchronously to the

execution of the query. The best response time is achieved when either the CPU or the I/O subsystem is operating at maximum capacity.

Page cleaning

As pages are read and modified, they accumulate in the database buffer pool. When a page is read in, it is read into a buffer pool page. If the buffer pool is full of modified pages, one of these modified pages must be written out to the disk before the new page can be read in. To prevent the buffer pool from becoming full, page cleaner agents write out modified pages to guarantee the availability of buffer pool pages for future read requests.

Whenever it is advantageous to do so, the database manager performs big-block reads. This typically occurs when retrieving data that is sequential or partially sequential in nature. The amount of data read in one read operation depends on the extent size — the bigger the extent size, the more pages can be read at one time.

Sequential prefetching performance can be further enhanced if pages can be read from disk into contiguous pages within a buffer pool. Since buffer pools are page-based by default, there is no guarantee of finding a set of contiguous pages when reading in contiguous pages from disk. Block-based buffer pools can be used for this purpose because they not only contain a page area, they also contain a block area for sets of contiguous pages. Each set of contiguous pages is named a block and each block contains a number of pages referred to as blocksize. The size of the page and block area, as well as the number of pages in each block is configurable.

How the extent is stored on disk affects I/O efficiency. In a DMS table space using device containers, the data tends to be contiguous on disk, and can be read with a minimum of seek time and disk latency. If files are being used, a large file that has been pre-allocated for use by a DMS table space also tends to be contiguous on disk, especially if the file was allocated in a clean file space. However, the data may have been broken up by the file system and stored in more than one location on disk. This occurs most often when using SMS table spaces, where files are extended one page at a time, making fragmentation more likely.

You can control the degree of prefetching by changing the PREFETCHSIZE option on the CREATE TABLESPACE or ALTER TABLESPACE statements, or you can set the prefetch size to AUTOMATIC to have the database manager automatically choose the best size to use. (The default value for all table spaces in the database is set by the *dft_prefetch_sz* database configuration parameter.) The PREFETCHSIZE parameter tells the database manager how many pages to read whenever a prefetch is triggered. By setting PREFETCHSIZE to be a multiple of the EXTENTSIZE parameter on the CREATE TABLESPACE statement, you can cause multiple extents to be read in parallel. (The default value for all table spaces in the database is set by the *dft_extent_sz* database configuration parameter.) The EXTENTSIZE parameter specifies the number of 4 KB pages that will be written to a container before skipping to the next container.

For example, suppose you had a table space that used three devices. If you set the PREFETCHSIZE to be three times the EXTENTSIZE, the database manager can do a big-block read from each device in parallel, thereby significantly increasing I/O throughput. This assumes that each device is a separate physical device, and that the controller has sufficient bandwidth to handle the data stream from each device.

Note that the database manager may have to dynamically adjust the prefetch parameters at run time based on query speed, buffer pool utilization, and other factors.

Some file systems use their own prefetching method (such as the Journaled File System on AIX). In some cases, file system prefetching is set to be more aggressive than the database manager prefetching. This may cause prefetching for SMS and DMS table spaces with file containers to appear to outperform prefetching for DMS table spaces with devices. This is misleading, because it is likely the result of the additional level of prefetching that is occurring in the file system. DMS table spaces should be able to outperform any equivalent configuration.

For prefetching (or even reading) to be efficient, a sufficient number of clean buffer pool pages must exist. For example, there could be a parallel prefetch request that reads three extents from a table space, and for each page being read, one modified page is written out from the buffer pool. The prefetch request may be slowed down to the point where it cannot keep up with the query. Page cleaners should be configured in sufficient numbers to satisfy the prefetch request.

Defining initial table spaces

When a database is created, three table spaces are defined: (1) SYSCATSPACE for the system catalog tables, (2) TEMPSPACE1 for system temporary tables created during database processing, and (3) USERSPACE1 for user-defined tables and indexes.

Note: When you first create a database no user temporary table space is created.

Unless otherwise specified, the three default table spaces are managed by Automatic Storage.

Using the CREATE DATABASE command, you can specify the page size for the default buffer pool and the initial table spaces. This default also represents the default page size for all future CREATE BUFFERPOOL and CREATE TABLESPACE statements. If you do not specify the page size when creating the database, the default page size is 4 KB.

To define initial table spaces using the command line, enter:

```
CREATE DATABASE <name>
  CATALOG TABLESPACE
    MANAGED BY SYSTEM USING ('<path>')
    EXTENTSIZE <value> PREFETCHSIZE <value>
  USER TABLESPACE
    MANAGED BY DATABASE USING (FILE'<path>' 5000,
                               FILE'<path>' 5000)
    EXTENTSIZE <value> PREFETCHSIZE <value>
  TEMPORARY TABLESPACE
    MANAGED BY SYSTEM USING ('<path>')
  WITH "<comment>"
```

If you do not want to use the default definition for these table spaces, you might specify their characteristics on the CREATE DATABASE command. For example, the following command could be used to create your database on Windows:

```
CREATE DATABASE PERSONL
  CATALOG TABLESPACE
    MANAGED BY SYSTEM USING ('d:\pcatalog','e:\pcatalog')
    EXTENTSIZE 16 PREFETCHSIZE 32
  USER TABLESPACE
```

```

MANAGED BY DATABASE USING (FILE'd:\db2data\person1' 5000,
                           FILE'd:\db2data\person1' 5000)
EXTENTSIZE 32 PREFETCHSIZE 64
TEMPORARY TABLESPACE
MANAGED BY SYSTEM USING ('f:\db2temp\person1')
WITH "Personnel DB for BSchiefer Co"

```

In this example, the definition for each of the initial table spaces is explicitly provided. You only need to specify the table space definitions for those table spaces for which you do not want to use the default definition.

Note: When working in a partitioned database environment, you cannot create or assign containers to specific database partitions. First, you must create the database with default user and temporary table spaces. Then you should use the CREATE TABLESPACE statement to create the required table spaces. Finally, you can drop the default table spaces.

The coding of the MANAGED BY phrase on the CREATE DATABASE command follows the same format as the MANAGED BY phrase on the CREATE TABLESPACE statement.

You can add additional user and temporary table spaces if you want. You cannot drop the catalog table space SYSCATSPACE, or create another one; and there must always be at least one system temporary table space with a page size of 4 KB. You can create other system temporary table spaces. You also cannot change the page size or the extent size of a table space after it has been created.

Attaching DMS direct disk access devices

When working with containers to store data, the database manager supports direct disk access (raw I/O).

This type of support allows you to attach a direct disk access (raw) device to any DB2 database system.

You must know the device or file names of the containers you are going to reference when creating your table spaces. You must know the amount of space associated with each device or file name that is to be allocated to the table space. You will need the correct permissions to read and write to the container.

The physical and logical methods for identifying direct disk access differs based on operating system:

- On the Windows operating systems:

To specify a physical hard drive, use the following syntax:

```
\\.\PhysicalDriveN
```

where N represents one of the physical drives in the system. In this case, N could be replaced by 0, 1, 2, or any other positive integer:

```
\\.\PhysicalDrive5
```

To specify a logical drive, that is, an unformatted database partition, use the following syntax:

```
\\.\N:
```

where N: represents a logical drive letter in the system. For example, N: could be replaced by E: or any other drive letter. To overcome the limitation imposed by using a letter to identify the drive, you can use a globally unique identifier (GUID) with the logical drive.

For Windows, there is a new method for specifying DMS raw table space containers. Volumes (that is, basic disk database partitions or dynamic volumes) are assigned a globally unique identifier (GUID) when they are created. The GUID can be used as a device identifier when specifying the containers in a table space definition. The GUIDs are unique across systems which means that in a multi-partition database, GUIDs are different for each database partition even if the disk partition definitions are the same.

A tool called *db2listvolumes.exe* is available (only on Windows operating systems) to make it easy to display the GUIDs for all the disk volumes defined on a Windows system. This tool creates two files in the current directory where the tool is run. One file, called *volumes.xml*, contains information about each disk volume encoded in XML for easy viewing on any XML-enabled browser. The second file, called *tablespace.ddl*, contains the required syntax for specifying table space containers. This file must be updated to fill in the remaining information needed for a table space definition. The *db2listvolumes* command does not require any command line arguments.

- On Linux and UNIX platforms, a logical volume can appear to users and applications as a single, contiguous, and extensible disk volume. Although it appears this way, it can reside on noncontiguous physical database partitions or even on more than one physical volume. The logical volume must also be contained within a single volume group. There is a limit of 256 logical volumes per volume group. There is a limit of 32 physical volumes per volume group. You can create additional logical volumes using the *mklv* command. This command allows you to specify the name of the logical volume and to define its characteristics, including the number and location of logical partitions to allocate for it.

After you create a logical volume, you can change its name and characteristics with the *chlv* command, and you can increase the number of logical partitions allocated to it with the *extendlv* command. The default maximum size for a logical volume at creation is 512 logical partitions, unless specified to be larger. The *chlv* command is used to override this limitation.

Within AIX, the set of operating system commands, library subroutines, and other tools that allow you to establish and control logical volume storage is called the Logical Volume Manager (LVM). The LVM controls disk resources by mapping data between a simpler and flexible logical view of storage space and the actual physical disks.

For more information on the *mklv* and other logical volume commands, and the LVM, refer to *AIX 5L Version 5.2 System Management Concepts: Operating System and Devices*.

Configuring and setting up DMS direct disk access (Linux)

When working with containers to store data, the database manager supports direct disk (raw) access using the block device interface (that is, raw I/O).

Before setting up raw I/O on Linux, one or more free IDE or SCSI disk database partitions are required. In order to reference the disk partition when creating the table space, you must know the name of the disk partition and the amount of space associated with the disk partition that is to be allocated to the table space.

The following information should be used when working in a Linux environment. On Linux/390, the database manager does not support direct disk access devices.

To configure or raw I/O on Linux:

In this example, the raw database partition to be used is /dev/sda5. It should not contain any valuable data.

1. Calculate the number of 4 096-byte pages in this database partition, rounding down if necessary. For example:

```
# fdisk /dev/sda
Command (m for help): p

Disk /dev/sda: 255 heads, 63 sectors, 1106 cylinders
Units = cylinders of 16065 * 512 bytes
```

Table 48. Linux raw I/O calculations.

Device boot	Start	End	Blocks	Id	System
/dev/sda1	1	523	4200997	83	Linux
/dev/sda2	524	1106	4682947+	5	Extended
/dev/sda5	524	1106	4682947	83	Linux

```
Command (m for help): q
#
```

The number of pages in /dev/sda5 is:

```
num_pages = floor( (4682947 * 1024)/4096 )
num_pages = 1170736
```

2. Create the table space by specifying the disk partition name. For example:

```
CREATE TABLESPACE dms1
MANAGED BY DATABASE
USING (DEVICE '/dev/sda5' 1170736)
```

3. To specify logical partitions by using junction points (or volume mount points), mount the RAW partition to another NTFS-formatted volume as a junction point, then specify the path to the junction point on the NTFS volume as the container path. For example:

```
CREATE TABLESPACE TS4
MANAGED BY DATABASE USING (DEVICE 'C:\JUNCTION\DISK_1' 10000,
DEVICE 'C:\JUNCTION\DISK_2' 10000)
```

The database manager first queries the partition to see whether there is a file system R on it; if yes, the partition is not treated as a RAW device, and performs normal file system I/O operations on the partition.

Table spaces on raw devices are also supported for all other page sizes supported by the database manager.

Prior to Version 9, direct disk access using a raw controller utility on Linux was used. This method is now deprecated, and its use is discouraged. The database manager will still allow you to use this method if the Linux operating system still supports it, however, there will be a message in the db2diag.log that will indicate that its use is deprecated.

The prior method would have required you to "bind" a disk partition to a raw controller, then specify that raw controller to the database manager using the CREATE TABLESPACE command:

```
CREATE TABLESPACE dms1
MANAGED BY DATABASE
USING (DEVICE '/dev/raw/raw1' 1170736)
```

Creating table spaces

For non-automatic storage table spaces you must know the device or file names of the containers that you will reference when creating your table spaces.

In addition, you must know the space associated with each device or file name that you will allocate to the table space. For automatic storage table spaces, the database manager assigns containers to the table space based on the storage paths associated with the database.

Table spaces establish the relationship between the physical storage devices used by your database system and the logical containers or tables used to store data.

Creating a table space within a database assigns containers to the table space and records its definitions and attributes in the database system catalog. You can then create tables within this table space. You must know the device or file names of the containers that you will reference when creating your table spaces. In addition, you must know the space associated with each device or file name that you will allocate to the table space.

When you create a database, three initial table spaces are created. The page size for the three initial table spaces is based on the default that is established or accepted when you use the CREATE DATABASE command. This default also represents the default page size for all future CREATE BUFFERPOOL and CREATE TABLESPACE statements. If you do not specify the page size when creating the database, the default page size is 4 KB. If you do not specify the page size when creating a table space, the default page size is the one set when you created the database.

To create an SMS table space using the command line, enter:

```
CREATE TABLESPACE <NAME>
MANAGED BY SYSTEM
USING ('<path>')
```

To create a DMS table space using the command line, enter:

```
CREATE TABLESPACE <NAME>
MANAGED BY DATABASE
USING (FILE'<path>' <size>)
```

To create an automatic storage table space using the command line, enter either of the following statements:

```
CREATE TABLESPACE <NAME>
CREATE TABLESPACE <NAME>
MANAGED BY AUTOMATIC STORAGE
```

The following SQL statement creates an SMS table space on Windows using three directories on three separate drives:

```
CREATE TABLESPACE RESOURCE
MANAGED BY SYSTEM
USING ('d:\acc_tbsp', 'e:\acc_tbsp', 'f:\acc_tbsp')
```

The following SQL statement creates a DMS table space using two file containers, each with 5,000 pages:

```

CREATE TABLESPACE RESOURCE
  MANAGED BY DATABASE
  USING (FILE'd:\db2data\acc_tbsp' 5000,
        FILE'e:\db2data\acc_tbsp' 5000)

```

In the previous two examples, explicit names are provided for the containers. However, if you specify relative container names, the container is created in the subdirectory created for the database.

When creating table space containers, the database manager creates any directory levels that do not exist. For example, if a container is specified as `/project/user_data/container1`, and the directory `/project` does not exist, then the database manager creates the directories `/project` and `/project/user_data`.

Any directories created by the database manager are created with `PERMISSION 700`. This means that only the instance owner has read, write, and execute access. Because only the instance owner has this access, the following scenario might occur when multiple instances are being created:

- Using the same directory structure as described above, suppose that directory levels `/project/user_data` do not exist.
- user1 creates an instance, named user1 by default, then creates a database, and then creates a table space with `/project/user_data/container1` as one of its containers.
- user2 creates an instance, named user2 by default, then creates a database, and then attempts to create a table space with `/project/user_data/container2` as one of its containers.

Because the database manager created directory levels `/project/user_data` with `PERMISSION 700` from the first request, user2 does not have access to these directory levels and cannot create container2 in those directories. In this case, the `CREATE TABLESPACE` operation fails.

There are two methods to resolve this conflict:

1. Create the directory `/project/user_data` before creating the table spaces and set the permission to whatever access is needed for both user1 and user2 to create the table spaces. If all levels of table space directory exist, the database manager does not modify the access.
2. After user1 creates `/project/user_data/container1`, set the permission of `/project/user_data` to whatever access is needed for user2 to create the table space.

If a subdirectory is created by the database manager, it might also be deleted by the database manager when the table space is dropped.

The assumption in this scenario is that the table spaces are not associated with a specific database partition group. The default database partition group `IBMDEFAULTGROUP` is used when the following parameter is not specified in the statement:

```

IN database_partition_group_name

```

The following SQL statement creates a DMS table space on an AIX system using three logical volumes of 10 000 pages each, and specifies their I/O characteristics:

```

CREATE TABLESPACE RESOURCE
  MANAGED BY DATABASE
  USING (DEVICE '/dev/rdb1v6' 10000,

```



```

        DEVICE '/dev/rdb1v7' 10000,
        DEVICE '/dev/rdb1v8' 10000)
OVERHEAD 7.5
TRANSFERRATE 0.06

```

The UNIX devices mentioned in this SQL statement must already exist, and the instance owner and the SYSADM group must be able to write to them.

The following example creates a DMS table space on a database partition group called ODDGROUP in a UNIX multi-partition database. ODDGROUP must be previously created with a CREATE DATABASE PARTITION GROUP statement. In this case, the ODDGROUP database partition group is assumed to be made up of database partitions numbered 1, 3, and 5. On all database partitions, use the device /dev/hdisk0 for 10 000 4 KB pages. In addition, declare a device for each database partition of 40 000 4 KB pages.

```

CREATE TABLESPACE PLANS IN ODDGROUP
MANAGED BY DATABASE
USING (DEVICE '/dev/HDISK0' 10000, DEVICE '/dev/n1hd01' 40000)
      ON DBPARTITIONNUM 1
      (DEVICE '/dev/HDISK0' 10000, DEVICE '/dev/n3hd03' 40000)
      ON DBPARTITIONNUM 3
      (DEVICE '/dev/HDISK0' 10000, DEVICE '/dev/n5hd05' 40000)
      ON DBPARTITIONNUM 5

```

The database manager can greatly improve the performance of sequential I/O using the sequential prefetch facility, which uses parallel I/O.

You can also create a table space that uses a page size larger than the default 4 KB size. The following SQL statement creates an SMS table space on a Linux and UNIX system with an 8 KB page size.

```

CREATE TABLESPACE SMS8K
PAGESIZE 8192
MANAGED BY SYSTEM
USING ('FSMS_8K_1')
BUFFERPOOL BUFPPOOL8K

```

Notice that the associated buffer pool must also have the same 8 KB page size.

The created table space cannot be used until the buffer pool it references is activated.

You can use the ALTER TABLESPACE statement to add, drop, or resize containers to a DMS table space and modify the PREFETCHSIZE, OVERHEAD, and TRANSFERRATE settings for a table space. You should commit the transaction issuing the table space statement as soon as possible following the ALTER TABLESPACE SQL statement to prevent system catalog contention.

Note: The PREFETCHSIZE value should be a multiple of the EXTENTSIZE value. For example if the EXTENTSIZE is 10, the PREFETCHSIZE should be 20 or 30. You should use the following equation to set your prefetch size manually when creating a table space:

$$\text{prefetch size} = (\text{number of containers}) \times (\text{number of physical spindles per container}) \times \text{extent size}$$

You should also consider letting the database manager automatically determine the prefetch size, by setting PREFETCHSIZE to AUTOMATIC.

Direct I/O (DIO) improves memory performance because it bypasses caching at the file system level. This process reduces CPU overhead and makes more memory available to the database instance.

Concurrent I/O (CIO) includes the advantages of DIO and also relieves the serialization of write accesses.

DIO and CIO are supported on AIX; DIO is supported on HP-UX, Solaris, Linux, and Windows operating systems.

The keywords `NO FILE SYSTEM CACHING` and `FILE SYSTEM CACHING` are part of the `CREATE` and `ALTER TABLESPACE` SQL statements to allow you to specify whether DIO or CIO is to be used with each table space. When `NO FILE SYSTEM CACHING` is in effect, the database manager attempts to use Concurrent I/O (CIO) wherever possible. In cases where CIO is not supported (for example, if JFS is used), DIO is used instead.

When you issue the `CREATE TABLESPACE` statement, the dropped table recovery feature is turned on by default. This feature lets you recover dropped table data using table space-level restore and rollforward operations. This is useful because it is faster than database-level recovery, and your database can remain available to users.

However, the dropped table recovery feature can have some performance impact on forward recovery when there are many drop table operations to recover or when the history file is very large.

You might want to disable this feature if you plan to run numerous drop table operations, and you either use circular logging or you do not think you will want to recover any of the dropped tables. To disable this feature, you can explicitly set the `DROPPED TABLE RECOVERY` option to `OFF` when you issue the `CREATE TABLESPACE` statement. Alternatively, you can turn off the dropped table recovery feature for an existing table space using the `ALTER TABLESPACE` statement.

Creating system temporary table spaces

A system temporary table space is used to store system temporary tables. A database must always have at least one system temporary table space since system temporary tables can only be stored in such a table space.

When a database is created, one of the three default table spaces defined is a system temporary table space called "TEMPSPACE1".

To create another system temporary table space, use the `CREATE TABLESPACE` statement. For example,

```
CREATE SYSTEM TEMPORARY TABLESPACE tmp_tbsp
MANAGED BY SYSTEM
USING ('d:\tmp_tbsp', 'e:\tmp_tbsp')
```

You should have at least one table space of each page size.

The only database partition group that can be specified when creating a system temporary table space is `IBMTEMPGROUP`.

Creating user temporary table spaces

User temporary table spaces are not created by default when a database is created. If your application programs need to use temporary tables, you need to create a user temporary table space where the temporary tables will reside.

Like regular table spaces, user temporary table spaces can be created in any database partition group other than IBMTEMPGROUP. IBMDEFAULTGROUP is the default database partition group that is used when creating a user temporary table.

The DECLARE GLOBAL TEMPORARY TABLE statement defines declared temporary tables for use within a user temporary table space.

To create a user temporary table space, use the CREATE TABLESPACE statement:

```
CREATE USER TEMPORARY TABLESPACE usr_tbsp
MANAGED BY DATABASE
USING (FILE 'd:\db2data\user_tbsp' 5000,
FILE 'e:\db2data\user_tbsp' 5000)
```

Altering table spaces

To alter a table space using the command line, use the ALTER TABLESPACE statement.

You can alter SMS, DMS, and automatic storage containers. You can also rename a table space, and switch it from offline to online mode.

Altering SMS table spaces

For SMS table spaces, you can only add one containers to single-partitioned databases, or one or more containers to partitioned databases.

The process for doing this is the same as described in: “Adding or extending DMS containers.”

Altering DMS table spaces

For DMS table spaces, you can add, extend, rebalance, resize, drop or reduce containers.

Adding or extending DMS containers

You can increase the size of a DMS table space (that is, one created with the MANAGED BY DATABASE clause) by adding one or more containers to the table space.

When new containers are added to a table space, or existing containers are extended, a rebalance of the table space might occur. The process of rebalancing involves moving table space extents from one location to another. During this process, an attempt is made to keep data striped within the table space. Rebalancing does not necessarily occur across all containers but depends on many factors, such as the existing container configuration, the size of the new containers, and how full the table space is.

When containers are added to an existing table space, they might be added such that they do not start in stripe 0, as described in “DMS table space maps” on page 158. Where they start in the map is determined by the database manager and is based on the size of the containers being added. If the container being added is not large enough, it is positioned such that it ends in the last stripe of the map. If it is large enough, it is positioned to start in stripe 0.

No rebalancing occurs if you are adding new containers and creating a new stripe set. A new stripe set is created using the BEGIN NEW STRIPE SET clause on the

ALTER TABLESPACE statement. You can also add containers to existing stripe sets using the ADD TO STRIPE SET clause on the ALTER TABLESPACE statement.

Access to the table space is not restricted during the rebalancing. If you need to add more than one container, you should add them at the same time.

To add a container to a DMS table space using the command line, enter:

```
ALTER TABLESPACE <name>
  ADD (DEVICE '<path>' <size>, FILE '<filename>' <size>)
```

The following example illustrates how to add two new device containers (each with 10 000 pages) to a table space on a Linux and UNIX system:

```
ALTER TABLESPACE RESOURCE
  ADD (DEVICE '/dev/rhd9' 10000,
       DEVICE '/dev/rhd10' 10000)
```

Note that the ALTER TABLESPACE statement allows you to change other properties of the table space that can affect performance.

Rebalancing DMS containers

The ALTER TABLESPACE statement lets you add a container to an existing table space or extend a container to increase its storage capacity.

Containers cannot be manually added to automatic storage table spaces. The database manager will automatically extend or add containers as necessary.

When a table space is created, its table space map is created and all of the initial containers are lined up such that they all start in stripe 0. This means that data is striped evenly across all of the table space containers until the individual containers fill up. (See Example 1.)

Adding a container that is smaller than existing containers results in a uneven distribution of data. This can cause parallel I/O operations, such as prefetching data, to perform less efficiently than they could on containers of equal size.

When new containers are added to a table space or existing containers are extended, a rebalance of the table space data *may* occur.

Rebalancing

The process of rebalancing when adding or extending containers involves moving table space extents from one location to another, and it is done in an attempt to keep data striped within the table space.

Access to the table space is not restricted during rebalancing; objects can be dropped, created, populated, and queried as usual. However, the rebalancing operation can have a significant impact on performance. If you need to add more than one container, and you plan to rebalance the containers, you should add them at the same time within a single ALTER TABLESPACE statement to prevent the database manager from having to rebalance the data more than once.

The table space *high-water mark* plays a key part in the rebalancing process. The high-water mark is the page number of the highest allocated page in the table space. For example, a table space has 1000 pages and an extent size of 10, resulting in 100 extents. If the 42nd extent is the highest allocated extent in the table space, then the high-water mark is $42 * 10 = 420$ pages. This is not the same as used

pages because some of the extents below the high-water mark may have been freed up so that they are available for reuse.

Before the rebalance starts, a new table space map is built based on the container changes made. The rebalancer moves extents from their location determined by the current map into the location determined by the new map. The rebalancer starts at extent 0, moving one extent at a time until the extent holding the high-water mark has been moved. As each extent is moved, the current map is altered, one piece at a time, to look like the new map. When the rebalance is complete, the current map and new map should look identical up to the stripe holding the high-water mark. The current map is then made to look completely like the new map and the rebalancing process is complete. If the location of an extent in the current map is the same as its location in the new map, then the extent is not moved and no I/O takes place.

When adding a new container, the placement of that container within the new map depends on its size and the size of the other containers in its stripe set. If the container is large enough such that it can start at the first stripe in the stripe set and end at (or beyond) the last stripe in the stripe set, then it will be placed that way (see Example 2). If the container is not large enough to do this, it will be positioned in the map such that it ends in the last stripe of the stripe set (see Example 4.) This is done to minimize the amount of data that needs to be rebalanced.

Note: In the following examples, the container sizes do not take the size of the container tag into account. The container sizes are very small, and are just used for the purpose of illustration, they are not recommended container sizes. The examples show containers of different sizes within a table space, but you are advised to use containers of the same size.

Example 1:

If you create a table space with three containers and an extent size of 10, and the containers are 60, 40, and 80 pages respectively (6, 4, and 8 extents), the table space is created with a map that can be diagrammed as shown in Figure 13 on page 196.

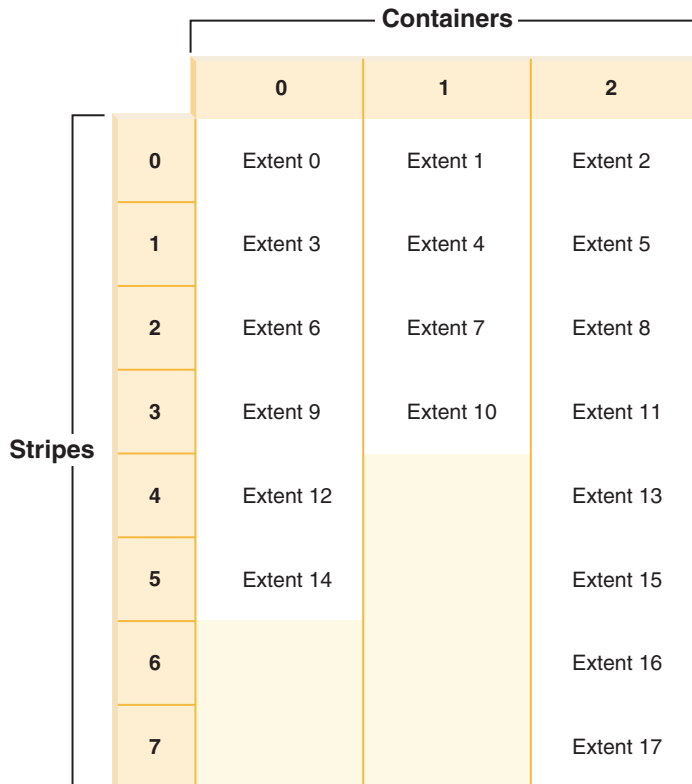


Figure 13. Table space with three containers and 18 extents

The corresponding table space map, as shown in a table space snapshot, looks like this:

Range Number	Stripe Set	Stripe Offset	Max Extent	Max Page	Start Stripe	End Stripe	Adj.	Containers
[0]	[0]	0	11	119	0	3	0	3 (0, 1, 2)
[1]	[0]	0	15	159	4	5	0	2 (0, 2)
[2]	[0]	0	17	179	6	7	0	1 (2)

The headings in the table space map are Range Number, Stripe Set, Stripe Offset, Maximum extent number addressed by the range, Maximum page number addressed by the range, Start Stripe, End Stripe, Range adjustment, and Container list.

Example 2:

If an 80-page container is added to the table space in Example 1, the container is large enough to start in the first stripe (stripe 0) and end in the last stripe (stripe 7). It is positioned such that it starts in the first stripe. The resulting table space can be diagrammed as shown in Figure 14 on page 197.




Figure 14. Table space with four containers and 26 extents

The corresponding table space map, as shown in a table space snapshot, will look like this:

Range Number	Stripe Set	Stripe Offset	Max Extent	Max Page	Start Stripe	End Stripe	Adj.	Containers
[0]	[0]	0	15	159	0	3	0	4 (0, 1, 2, 3)
[1]	[0]	0	21	219	4	5	0	3 (0, 2, 3)
[2]	[0]	0	25	259	6	7	0	2 (2, 3)

If the high-water mark is within extent 14, the rebalancer starts at extent 0 and moves all of the extents up to and including 14. The location of extent 0 within both of the maps is the same so this extent does not need to move. The same is true for extents 1 and 2. Extent 3 does need to move so the extent is read from the old location (second extent within container 0) and is written to the new location (first extent within container 3). Every extent after this up to and including extent 14 is moved. Once extent 14 is moved, the current map looks like the new map and the rebalancer terminates.

If the map is altered such that all of the newly added space comes after the high-water mark, then a rebalance is not necessary and all of the space is available immediately for use. If the map is altered such that some of the space comes after the high-water mark, then the space in the stripes above the high-water mark is available for use. The rest is not available until the rebalance is complete.

If you decide to extend a container, the function of the rebalancer is similar. If a container is extended such that it extends beyond the last stripe in its stripe set, the stripe set will expand to fit this and the following stripe sets will be shifted out accordingly. The result is that the container will not extend into any stripe sets following it.

Example 3:

Consider the table space from Example 1. If you extend container 1 from 40 pages to 80 pages, the new table space looks like Figure 15.




Figure 15. Table space with three containers and 22 extents

The corresponding table space map, as shown in a table space snapshot, looks like this:

Range Number	Stripe Set	Stripe Offset	Max Extent	Max Page	Start Stripe	End Stripe	Adj.	Containers
[0]	[0]	0	17	179	0	5	0	3 (0, 1, 2)
[1]	[0]	0	21	219	6	7	0	2 (1, 2)

Example 4:

Consider the table space from Example 1. If a 50-page (5-extent) container is added to it, the container will be added to the new map in the following way. The container is not large enough to start in the first stripe (stripe 0) and end at or beyond the last stripe (stripe 7), so it is positioned such that it ends in the last stripe. (See Figure 16 on page 199.)




Figure 16. Table space with four containers and 23 extents

The corresponding table space map, as shown in a table space snapshot, will look like this:

Range Number	Stripe Set	Stripe Offset	Max Extent	Max Page	Start Stripe	End Stripe	Adj.	Containers
[0]	[0]	0	8	89	0	2	0	3 (0, 1, 2)
[1]	[0]	0	12	129	3	3	0	4 (0, 1, 2, 3)
[2]	[0]	0	18	189	4	5	0	3 (0, 2, 3)
[3]	[0]	0	22	229	6	7	0	2 (2, 3)

To extend a container, use the EXTEND or RESIZE clause on the ALTER TABLESPACE statement. To add containers and rebalance the data, use the ADD clause on the ALTER TABLESPACE statement. If you are adding a container to a table space that already has more than one stripe set, you can specify which stripe set you want to add to. To do this, you use the ADD TO STRIPE SET clause on the ALTER TABLESPACE statement. If you do not specify a stripe set, the default behavior will be to add the container to the current stripe set. The current stripe set is the most recently created stripe set, not the one that last had space added to it.

Any change to a stripe set may cause a rebalance to occur to that stripe set and any others following it.

You can monitor the progress of a rebalance by using table space snapshots. A table space snapshot can provide information about a rebalance such as the start time of the rebalance, how many extents have been moved, and how many extents need to move.

Without rebalancing (using stripe sets)

If you add or extend a container, and the space added is above the table space high-water mark, a rebalance will not occur.

Adding a container will almost always add space below the high-water mark. In other words, a rebalance is often necessary when you add a container. There is an option to force new containers to be added above the high-water mark, which allows you to choose not to rebalance the contents of the table space. An advantage of this method is that the new container will be available for immediate use. The option not to rebalance applies only when you add containers, not when you extend existing containers. When you extend containers you can only avoid rebalancing if the space you add is above the high-water mark. For example, if you have a number of containers that are the same size, and you extend each of them by the same amount, the relative positions of the extents will not change, and a rebalance will not occur.

Adding containers to a table space without rebalancing is done by adding a new *stripe set*. A stripe set is a set of containers in a table space that has data striped across it separately from the other containers that belong to that table space. The existing containers in the existing stripe sets remain untouched, and the containers you add become part of a new stripe set.

To add containers without rebalancing, use the `BEGIN NEW STRIPE SET` clause on the `ALTER TABLESPACE` statement.

Example 5:

If you have a table space with three containers and an extent size of 10, and the containers are 30, 40, and 40 pages (3, 4, and 4 extents respectively), the table space can be diagrammed as shown in Figure 17.

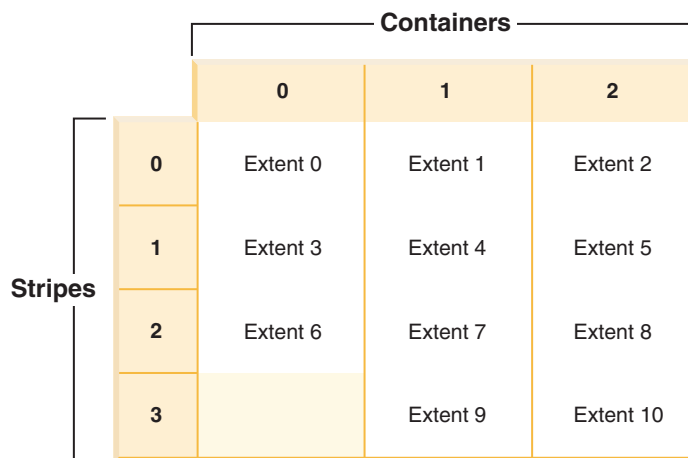


Figure 17. Table space with three containers and 11 extents

The corresponding table space map, as shown in a table space snapshot, will look like this:

Range Number	Stripe Set	Stripe Offset	Max Extent	Max Page	Start Stripe	End Stripe	Adj.	Containers
[0]	[0]	0	8	89	0	2	0	3 (0, 1, 2)
[1]	[0]	0	10	109	3	3	0	2 (1, 2)

Example 6:

When you add two new containers that are 30 pages and 40 pages (3 and 4 extents respectively) with the `BEGIN NEW STRIPE SET` clause, the existing ranges are not affected; instead, a new set of ranges is created. This new set of ranges is a stripe set and the most recently created one is called the current stripe set. After the two new containers is added, the table space looks like Figure 18.

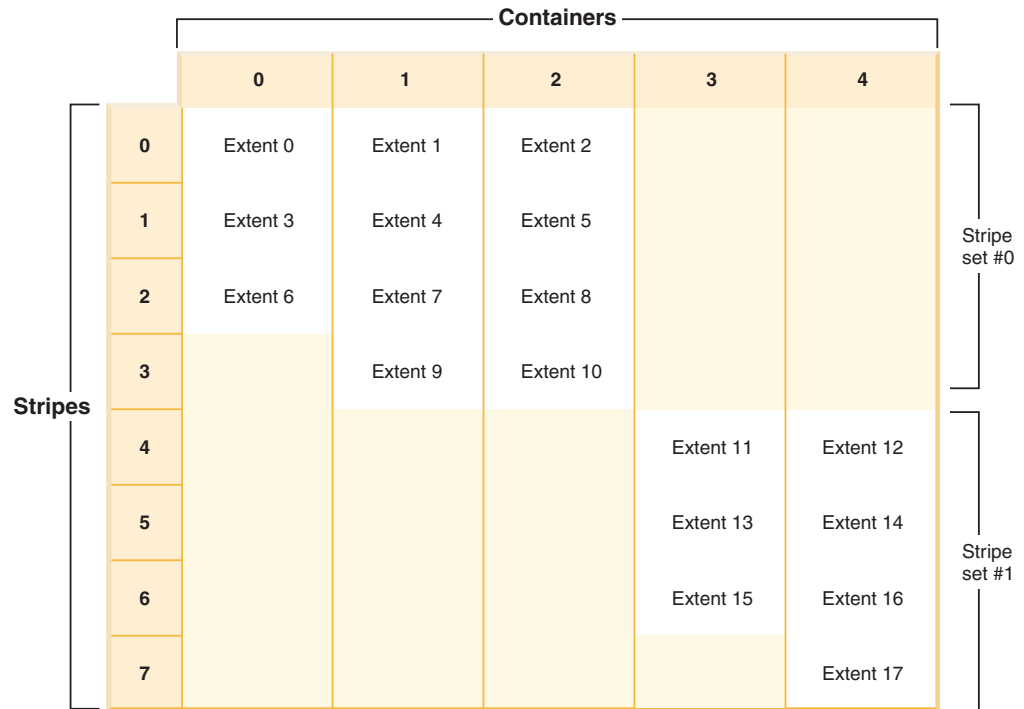


Figure 18. Table space with two stripe sets

The corresponding table space map, as shown in a table space snapshot, looks like this:

Range Number	Stripe Set	Stripe Offset	Max Extent	Max Page	Start Stripe	End Stripe	Adj.	Containers
[0]	[0]	0	8	89	0	2	0	3 (0, 1, 2)
[1]	[0]	0	10	109	3	3	0	2 (1, 2)
[2]	[1]	4	16	169	4	6	0	2 (3, 4)
[3]	[1]	4	17	179	7	7	0	1 (4)

If you add new containers to a table space, and you do *not* use the `TO STRIPE SET` clause with the `ADD` clause, the containers are added to the current stripe set (the highest stripe set). You can use the `ADD TO STRIPE SET` clause to add containers to any stripe set in the table space. You must specify a valid stripe set.

The database manager tracks the stripe sets using the table space map, and adding new containers without rebalancing generally causes the map to grow faster than when containers are rebalanced. When the table space map becomes too large, you will receive error `SQL0259N` when you try to add more containers.

Resizing DMS containers

Containers cannot be manually resized in automatic storage table spaces.

The database manager will automatically extend or add containers as necessary. However, you can resize the containers in a DMS table space (that is, one created with the `MANAGED BY DATABASE` clause).

Each raw device can only be used as one container. The raw device size is fixed after its creation. When you are considering to use the `RESIZE` or `EXTEND` options to increase a raw device container, you should check the raw device size first to ensure that you do not attempt to increase the device container size larger than the raw device size.

You can also drop existing containers from a DMS table space, reduce the size of existing containers in a DMS table space, and add new containers to a DMS table space without requiring a rebalance of the data across all of the containers.

The dropping of existing table space containers as well as the reduction in size of existing containers is only allowed if the number of extents being dropped or reduced in size is less than or equal to the number of free extents above the “high-water mark” in the table space. The high-water mark is the page number of the highest allocated page in the table space. This mark is not the same as the number of used pages in the table space because some of the extents below the high-water mark might have been made available for reuse.

The number of free extents above the high-water mark in the table space is important because all extents up to and including the high-water mark must sit in the same logical position within the table space. The resulting table space must have enough space to hold all of the data. If there is not enough free space, an error message (`SQL20170N, SQLSTATE 57059`) will result.

To drop containers, the `DROP` option is used on the `ALTER TABLESPACE` statement. For example:

```
ALTER TABLESPACE TS1 DROP (FILE 'file1', DEVICE '/dev/rdisk1')
```

To reduce the size of existing containers, you can use either the `RESIZE` option or the `REDUCE` option. When using the `RESIZE` option, all of the containers listed as part of the statement must either be increased in size, or decreased in size. You cannot increase some containers and decrease other containers in the same statement. You should consider the resizing method if you know the new lower limit for the size of the container. You should consider the reduction method if you do not know (or care about) the current size of the container.

To decrease the size of one or more containers in a DMS table space using the command line, enter:

```
ALTER TABLESPACE <name>  
  REDUCE (FILE '<filename>' <size>)
```

The following example illustrates how to reduce a file container (which already exists with 1 000 pages) in a table space on a Windows-based system:

```
ALTER TABLESPACE PAYROLL  
  REDUCE (FILE 'd:\hldr\finance' 200)
```

Following this action, the file is decreased from 1 000 pages in size to 800 pages.

To increase the size of one or more containers in a DMS table space using the command line, enter:

```
ALTER TABLESPACE <name>  
  RESIZE (DEVICE '<path>' <size>)
```

The following example illustrates how to increase two device containers (each already existing with 1 000 pages) in a table space on a Linux and UNIX system:

```
ALTER TABLESPACE HISTORY
  RESIZE (DEVICE '/dev/rhd7' 2000,
         DEVICE '/dev/rhd8' 2000)
```

Following this action, the two devices have increased from 1 000 pages in size to 2 000 pages. The contents of the table space might be rebalanced across the containers. Access to the table space is not restricted during the rebalancing.

To extend one or more containers in a DMS table space using the command line, enter:

```
ALTER TABLESPACE <name>
  EXTEND (FILE '<filename>' <size>)
```

The following example illustrates how to increase file containers (each already existing with 1 000 pages) in a table space on a Windows-based system:

```
ALTER TABLESPACE PERSNEL
  EXTEND (FILE 'e:\wrkhist1' 200
         FILE 'f:\wrkhist2' 200)
```

Following this action, the two files have increased from 1 000 pages in size to 1 200 pages. The contents of the table space might be rebalanced across the containers. Access to the table space is not restricted during the re-balancing.

The addition or modification of DMS containers (both file and raw device containers) is performed in parallel through prefetchers. To achieve an increase in parallelism of these create or resize container operations, you can increase the number of prefetchers running in the system. The only process which is not done in parallel is the logging of these actions and, in the case of creating containers, the tagging of the containers.

Note: To maximize the parallelism of the CREATE TABLESPACE or ALTER TABLESPACE statements (with respect to adding new containers to an existing table space) ensure the number of prefetchers is greater than or equal to the number of containers being added. The number of prefetchers is controlled by the *num_ioservers* database configuration parameter. The database has to be stopped for the new parameter value to take effect. In other words, all applications and users must disconnect from the database for the change to take affect.

Note that the ALTER TABLESPACE statement allows you to change other properties of the table space that can affect performance.

Dropping or reducing DMS containers

With a DMS table space, you can drop a container from the table space or reduce the size of a container using the ALTER TABLESPACE statement.

Dropping or reducing a container will only be allowed if the number of extents being dropped by the operation is less than or equal to the number of free extents above the high-water mark in the table space. This is necessary because page numbers cannot be changed by the operation and therefore all extents up to and including the high-water mark must sit in the same logical position within the table space. Therefore, the resulting table space must have enough space to hold all of the data up to and including the high-water mark. In the situation where there is not enough free space, you will receive an error immediately upon execution of the statement.

The high-water mark is the page number of the highest allocated page in the table space. For example, a table space has 1000 pages and an extent size of 10, resulting in 100 extents. If the 42nd extent is the highest allocated extent in the table space that means that the high-water mark is $42 * 10 = 420$ pages. This is not the same as used pages because some of the extents below the high-water mark may have been freed up such that they are available for reuse.

When containers are dropped or reduced, a rebalance will occur if data resides in the space being dropped from the table space. Before the rebalance starts, a new table space map is built based on the container changes made. The rebalancer will move extents from their location determined by the current map into the location determined by the new map. The rebalancer starts with the extent that contains the high-water mark, moving one extent at a time until extent 0 has been moved. As each extent is moved, the current map is altered one piece at a time to look like the new map. If the location of an extent in the current map is the same as its location in the new map, then the extent is not moved and no I/O takes place. Because the rebalance moves extents starting with the highest allocated one, ending with the first extent in the table space, it is called a *reverse rebalance* (as opposed to the *forward rebalance* that occurs when space is added to the table space after adding or extending containers).

When containers are dropped, the remaining containers are renumbered such that their container IDs start at 0 and increase by 1. If all of the containers in a stripe set are dropped, the stripe set will be removed from the map and all stripe sets following it in the map will be shifted down and renumbered such that there are no gaps in the stripe set numbers.

Note: In the following examples, the container sizes do not take the size of the container tag into account. The container sizes are very small, and are just used for the purpose of illustration, they are not recommended container sizes. The examples show containers of different sizes within a table space, but this is just for the purpose of illustration; you are advised to use containers of the same size.

For example, consider a table space with three containers and an extent size of 10. The containers are 20, 50, and 50 pages respectively (2, 5, and 5 extents). The table space diagram is shown in Figure 19 on page 205.

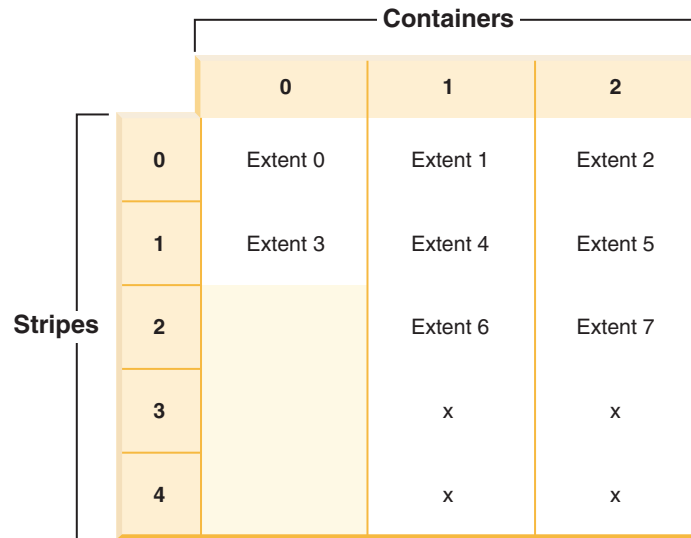


Figure 19. Table space with 12 extents, including four extents with no data

An X indicates that there is an extent but there is no data in it.

If you want to drop container 0, which has two extents, there must be at least two free extents above the high-water mark. The high-water mark is in extent 7, leaving four free extents, therefore you can drop container 0.

The corresponding table space map, as shown in a table space snapshot, will look like this:

Range Number	Stripe Set	Stripe Offset	Max Extent	Max Page	Start Stripe	End Stripe	Adj.	Containers
[0]	[0]	0	5	59	0	1	0	3 (0, 1, 2)
[1]	[0]	0	11	119	2	4	0	2 (1, 2)

After the drop, the table space will have just Container 0 and Container 1. The new table space diagram is shown in Figure 20 on page 206.

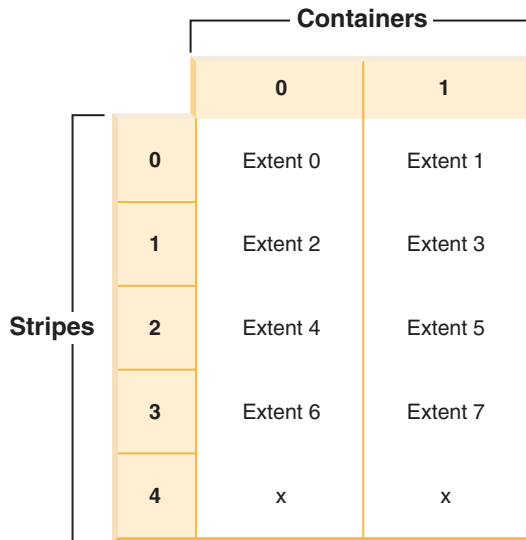


Figure 20. Table space after a container is dropped

The corresponding table space map, as shown in a table space snapshot, will look like this:

Range Number	Stripe Set	Stripe Offset	Max Extent	Max Page	Start Stripe	End Stripe	Adj.	Containers
[0]	[0]	0	9	99	0	4	0	2 (0, 1)

If you want to reduce the size of a container, the rebalancer works in a similar way.

To reduce a container, use the REDUCE or RESIZE option on the ALTER TABLESPACE statement. To drop a container, use the DROP option on the ALTER TABLESPACE statement.

Altering automatic storage table spaces

For automatic storage table spaces, you can only reduce the size of the containers.

The process for doing this is the same as described in: “Resizing DMS containers” on page 201.

Renaming a table space

Use the RENAME TABLESPACE statement to rename a table space.

You cannot rename the SYSCATSPACE table space. You cannot rename a table space that is in a rollforward pending or rollforward-in-progress state.

When restoring a table space that has been renamed since it was backed up, you must use the new table space name in the RESTORE DATABASE command. If you use the previous table space name, it will not be found. Similarly, if you are rolling forward the table space with the ROLLFORWARD DATABASE command, ensure that you use the new name. If the previous table space name is used, it will not be found.

You can give an existing table space a new name without being concerned with the individual objects within the table space. When renaming a table space, all the catalog records referencing that table space are changed.

Switching table spaces from offline to online

The SWITCH ONLINE clause of the ALTER TABLESPACE statement can be used to remove the OFFLINE state from a table space if the containers associated with that table space have become accessible.

The table space has the OFFLINE state removed while the rest of the database is still up and being used.

An alternative to the use of this clause is to disconnect all applications from the database and then to have the applications connect to the database again. This removes the OFFLINE state from the table space.

To remove the OFFLINE state from a table space using the command line, enter:

```
db2 ALTER TABLESPACE <name>
    SWITCH ONLINE
```

Optimizing table space performance when data is on RAID devices

Follow these guidelines to optimize performance when data is stored on Redundant Array of Independent Disks (RAID) devices.

1. When creating a table space on a set of RAID devices, create the containers for a given table space (SMS or DMS) on separate devices.

Consider an example where you have fifteen 146 GB disks configured as three RAID-5 arrays with five disks in each array. After formatting, each disk can hold approximately 136 GB of data. Each array can therefore store approximately 544 GB (4 active disks × 136 GB). If you have a table space that requires 300 GB of storage, create three containers, and put each container on a separate device. Each container uses 100 GB (300 GB/3) on a device, and there are 444 GB (544 GB - 100 GB) left on each device for additional table spaces.

2. Select an appropriate extent size for the table spaces. The extent size for a table space is the amount of data that the database manager writes to a container before writing to the next container. Ideally, the extent size should be a multiple of the underlying segment size of the disks, where the segment size is the amount of data that the disk controller writes to one physical disk before writing to the next physical disk. Choosing an extent size that is a multiple of the segment size ensures that extent-based operations, such as parallel sequential read in prefetching, do not compete for the same physical disks. Also, choose an extent size that is a multiple of the page size.

In the example, if the segment size is 64 KB and the page size is 16 KB, an appropriate extent size might be 256 KB.

3. Use the DB2_PARALLEL_IO registry variable to enable parallel I/O for all table spaces and to specify the number of physical disks per container.

For the situation in the example, set DB2_PARALLEL_IO = *:4.

If you set the prefetch size of a table space to AUTOMATIC, the database manager uses the number of physical disks value that you specified for DB2_PARALLEL_IO to determine the prefetch size value. If the prefetch size is not set to AUTOMATIC, you can set it manually, taking into account the RAID stripe size, which is the value of the segment size multiplied by the number of active disks. Choose a prefetch size value that meets the following conditions:

- It is equal to the RAID stripe size multiplied by the number of RAID parallel devices (or a whole number representation of this product).
- It is a whole number representation of the extent size.

In the example, you might set the prefetch size to 768 KB. This value is equal to the RAID stripe size (256 KB) multiplied by the number of RAID parallel devices (3). It is also a multiple of the extent size (256 KB). Choosing this prefetch size means that a single prefetch will engage all the disks in all the arrays. If you want the prefetchers to work more aggressively because your workload involves mainly sequential scans, you can instead use a multiple of this value, such as 1536 KB (768 KB x 2).

4. Do not set the `DB2_USE_PAGE_CONTAINER_TAG` registry variable. As described earlier, you should create a table space with an extent size that is equal to, or a multiple of, the RAID stripe size. However, when you set `DB2_USE_PAGE_CONTAINER_TAG` to ON, a one-page container tag is used, and the extents do not line up with the RAID stripes. As a result, it might be necessary during an I/O request to access more physical disks than would be optimal.

Dropping table spaces

When you drop a table space, you delete all the data in that table space, free the containers, remove the catalog entries, and cause all objects defined in the table space to be either dropped or marked as invalid.

You can reuse the containers in an empty table space by dropping the table space, but you must commit the `DROP TABLESPACE` statement before attempting to reuse the containers.

Note: You cannot drop a table space without dropping all table spaces that are associated with it. Example, if you have a table in one table space and its index created in another table space, you need to drop both index and data table spaces in one `DROP TABLESPACE` command.

Dropping user table spaces

You can drop a user table space that contains all of the table data including index and LOB data within that single user table space. You can also drop a user table space that might have tables spanned across several table spaces. That is, you might have table data in one table space, indexes in another, and any LOBs in a third table space. You must drop all three table spaces at the same time in a single statement. All of the table spaces that contain tables that are spanned must be part of this single statement or the drop request will fail.

To drop a user table space using the command line, enter:

```
DROP TABLESPACE <name>
```

The following SQL statement drops the table space `ACCOUNTING`:

```
DROP TABLESPACE ACCOUNTING
```

Dropping user temporary table spaces

You can only drop a user temporary table space if there are no declared temporary tables currently defined in that table space. When you drop the table space, no attempt is made to drop all of the declared temporary tables in the table space.

Note: A declared temporary table is implicitly dropped when the application that declared it disconnects from the database.

Dropping system temporary table spaces

You cannot drop a system temporary table space that has a page size of 4 KB without first creating another system temporary table space. The new system temporary table space must have a page size of 4 KB because the database must always have at least one system temporary table space that has a page size of 4 KB. For example, if you have a single system temporary table space with a page size of 4 KB, and you want to add a container to it, and it is an SMS table space, you must first add a new 4 KB page size system temporary table space with the proper number of containers, and then drop the old system temporary table space. (If you were using DMS, you could add a container without having to drop and recreate the table space.)

The default table space page size is the page size that the database was created with (which is 4 KB by default, but could also be 8 KB, 16 KB, or 32 KB).

This is the statement to create a system temporary table space:

```
CREATE SYSTEM TEMPORARY TABLESPACE <name>  
MANAGED BY SYSTEM USING ('<directories>')
```

Then, to drop a system table space using the command line, enter:

```
DROP TABLESPACE <name>
```

The following SQL statement creates a new system temporary table space called TEMPSPACE2:

```
CREATE SYSTEM TEMPORARY TABLESPACE TEMPSPACE2  
MANAGED BY SYSTEM USING ('d:\systemp2')
```

Once TEMPSPACE2 is created, you can then drop the original system temporary table space TEMPSPACE1 with the command:

```
DROP TABLESPACE TEMPSPACE1
```

Chapter 10. Schemas

A *schema* is a collection of named objects; it provides a way to group those objects logically. A schema is also a name qualifier; it provides a way to use the same natural name for several objects, and to prevent ambiguous references to those objects.

For example, the schema names 'INTERNAL' and 'EXTERNAL' make it easy to distinguish two different SALES tables (INTERNAL.SALES, EXTERNAL.SALES).

Schemas also enable multiple applications to store data in a single database without encountering namespace collisions.

A schema is distinct from, and should not be confused with, an *XML schema*, which is a standard that describes the structure and validates the content of XML documents.

A schema can contain tables, views, nicknames, triggers, functions, packages, and other objects. A schema is itself a database object. It is explicitly created using the CREATE SCHEMA statement, with the current user or a specified authorization ID recorded as the schema owner. It can also be implicitly created when another object is created, if the user has IMPLICIT_SCHEMA authority.

A *schema name* is used as the high order part of a two-part object name. If the object is specifically qualified with a schema name when created, the object is assigned to that schema. If no schema name is specified when the object is created, the default schema name is used (specified in the CURRENT SCHEMA special register).

For example, a user with DBADM authority creates a schema called C for user A:

```
CREATE SCHEMA C AUTHORIZATION A
```

User A can then issue the following statement to create a table called X in schema C (provided that user A has the CREATETAB database authority):

```
CREATE TABLE C.X (COL1 INT)
```

Some schema names are reserved. For example, built-in functions belong to the SYSIBM schema, and the pre-installed user-defined functions belong to the SYSFUN schema.

When a database is created, if it is not created with the RESTRICTIVE option, all users have IMPLICIT_SCHEMA authority. With this authority, users implicitly create a schema whenever they create an object with a schema name that does not already exist. When schemas are implicitly created, CREATEIN privileges are granted which allows any user to create other objects in this schema. The ability to create objects such as aliases, distinct types, functions, and triggers is extended to implicitly-created schemas. The default privileges on an implicitly-created schema provide backward compatibility with previous versions.

If IMPLICIT_SCHEMA authority is revoked from PUBLIC, schemas can be explicitly created using the CREATE SCHEMA statement, or implicitly created by users (such as those with DBADM authority) who have been granted IMPLICIT_SCHEMA authority. Although revoking IMPLICIT_SCHEMA authority

from PUBLIC increases control over the use of schema names, it can result in authorization errors when existing applications attempt to create objects.

Schemas also have privileges, allowing the schema owner to control which users have the privilege to create, alter, copy, and drop objects in the schema. This provides a way to control the manipulation of a subset of objects in the database. A schema owner is initially given all of these privileges on the schema, with the ability to grant the privileges to others. An implicitly-created schema is owned by the system, and all users are initially given the privilege to create objects in such a schema. A user with SYSADM or DBADM authority can change the privileges that are held by users on any schema. Therefore, access to create, alter, copy, and drop objects in any schema (even one that was implicitly created) can be controlled.

Designing schemas

While organizing your data into tables, it might be beneficial to group the tables and other related objects together. This is done by defining a schema through the use of the CREATE SCHEMA statement.

Information about the schema is kept in the system catalog tables of the database to which you are connected. As other objects are created, they can be placed within the schemas you create, however, note that an object can exist in only one schema.


Schemas can be compared to directories, with the current schema being the current directory. Using this analogy, SET SCHEMA is equivalent to the change directory command.

Important: It is important to understand that there is no relation between authorization IDs and schemas except for the default CURRENT SCHEMA setting (described below).

While designing your databases and tables, you should also consider the schemas in your system, including their names and the objects that will be associated with each of them.

Most objects in a database are assigned a unique name that consists of two parts. The first (leftmost) part is called the qualifier or schema, and the second (rightmost) part is called the simple (or unqualified) name. Syntactically, these two parts are concatenated as a single string of characters separated by a period. When any object that can be qualified by a schema name (such as a table, index, view, user-defined data type, user-defined function, nickname, package, or trigger) is first created, it is assigned to a particular schema based on the qualifier in its name.

For example, the following diagram illustrates how a table is assigned to a particular schema during the table creation process:



You should also be familiar with how schema access is granted, in order to give your users the correct authority and instructions:

Schema names

When creating a new schema, the name must not identify a schema name already described in the catalog and the name cannot begin with "SYS". For other restrictions and recommendations, see "Schema name restrictions and recommendations" on page 215.

Access to schemas

Unqualified access to objects within a schema is not allowed since the schema is used to enforce uniqueness in the database. This becomes clear when considering the possibility that two users could create two tables (or other objects) with the same name. Without a schema to enforce uniqueness, ambiguity would exist if a third user attempted to query the table. It is not possible to determine which table to use without some further qualification.

The definer of any objects created as part of the CREATE SCHEMA statement is the schema owner. This owner can GRANT and REVOKE schema privileges to other users.

If a user has SYSADM or DBADM authority, then that user can create a schema with any valid name. When a database is created, IMPLICIT_SCHEMA authority is granted to PUBLIC (that is, to all users).

If users do not have IMPLICIT_SCHEMA or DBADM authority, the only schema they can create is one that has the same name as their own authorization ID.

Default schema

If a schema or qualifier is not specified as part of the name of the object to be created, that object is assigned to the default schema as indicated in the CURRENT_SCHEMA special register. The default value of this special register is the value of the session authorization ID.

A default schema is needed by unqualified object references in dynamic statements. You can set a default schema for a specific DB2 connection by setting the CURRENT SCHEMA special register to the schema that you want as the default. No designated authorization is required to set this special register, so any user can set the CURRENT SCHEMA.

The syntax of the SET SCHEMA statement is:

```
SET SCHEMA = <schema-name>
```

You can issue this statement interactively or from within an application. The initial value of the CURRENT SCHEMA special register is equal to the authorization ID of the current session user. For more information, see the SET SCHEMA statement.

Note:

- There are other ways to set the default schema upon connection. For example, by using the cli.ini file for CLI/ODBC applications, or by using the connection properties for the JDBC application programming interface.
- The default schema record is not created in the system catalogs, but it exists only as a value that the database manager can obtain (from the CURRENT SCHEMA special register) whenever a schema or qualifier is not specified as part of the name of the object to be created.

Implicit creation

You can implicitly create schemas if you have IMPLICIT_SCHEMA authority. With this authority, you can implicitly create a schema whenever you create an object with a schema name that does not already exist. Often schemas are implicitly created the first time a data object in the schema is created, provided the user creating the object holds the IMPLICIT_SCHEMA authority.

Explicit creation

Schemas can also be explicitly created and dropped by executing the CREATE SCHEMA and DROP SCHEMA statements from the command line or from an application program. For more information, see the CREATE SCHEMA and DROP SCHEMA statements.

Table and view aliases by schema

To allow another user to access a table or view without entering the schema name as part of the qualification on the table or view name requires that an alias be established for that user. The definition of the alias would define the fully-qualified table or view name including the user's schema; then the user simply queries using the alias name. The alias would be fully-qualified by the user's schema as part of the alias definition.

Grouping objects by schema

Database object names might be made up of a single identifier or they might be *schema-qualified objects* made up of two identifiers. The schema, or high-order part, of a schema-qualified object provides a means to classify or group objects in the database. When an object such as a table, view, alias, distinct type, function, index, package or trigger is created, it is assigned to a schema. This assignment is done either explicitly or implicitly.

Explicit use of the schema occurs when you use the high-order part of a two-part object name when referring to that object in a statement. For example, USER A issues a CREATE TABLE statement in schema C as follows:

```
CREATE TABLE C.X (COL1 INT)
```

Implicit use of the schema occurs when you do not use the high-order part of a two-part object name. When this happens, the CURRENT SCHEMA special register is used to identify the schema name used to complete the high-order part of the object name. The initial value of CURRENT SCHEMA is the authorization ID of the current session user. If you want to change this during the current session, you can use the SET SCHEMA statement to set the special register to another schema name.

Some objects are created within certain schemas and stored in the system catalog tables when the database is created.

You do not have to explicitly specify in which schema an object is to be created; if not specified, the authorization ID of the statement is used. For example, for the following CREATE TABLE statement, the schema name defaults to the authorization ID that is currently logged on (that is, the CURRENT SCHEMA special register value):

```
CREATE TABLE X (COL1 INT)
```

Dynamic SQL and XQuery statements typically use the CURRENT SCHEMA special register value to implicitly qualify any unqualified object name references.

Before creating your own objects, you need to consider whether you want to create them in your own schema or by using a different schema that logically groups the objects. If you are creating objects that will be shared, using a different schema name can be very beneficial.

Schema name restrictions and recommendations

There are some restrictions and recommendations that you need to be aware of when naming schemas.

- User-defined types (UDTs) cannot have schema names longer than the schema length listed in: SQL and XML limits.
- The following schema names are reserved words and must not be used: SYSCAT, SYSFUN, SYSIBM, SYSSTAT, SYSPROC.
- To avoid potential migration problems in the future, do not use schema names that begin with SYS. The database manager will not allow you to create triggers, user-defined types or user-defined functions using a schema name beginning with SYS.
- It is recommended that you not use SESSION as a schema name. Declared temporary tables must be qualified by SESSION. It is therefore possible to have an application declare a temporary table with a name identical to that of a persistent table, in which case the application logic can become overly complicated. Avoid the use of the schema SESSION, except when dealing with declared temporary tables.

Creating schemas

You can use schemas to group objects as you create those objects. An object can belong to only one schema. Use the CREATE SCHEMA statement to create schemas. Information about the schemas is kept in the system catalog tables of the database to which you are connected.

To create a schema and optionally make another user the owner of the schema, you need SYSADM or DBADM authority. If you do not hold either of these two authorities, you can still create a schema using your own authorization ID. The definer of any objects created as part of the CREATE SCHEMA statement is the schema owner. This owner can GRANT and REVOKE schema privileges to other users.

To create a schema from the command line, enter the following statement:

```
CREATE SCHEMA <schema-name> [ AUTHORIZATION <schema-owner-name> ]
```

Where <schema-name> is the name of the schema. This name must be unique within the schemas already recorded in the catalog, and the name cannot begin with SYS. If the optional AUTHORIZATION clause is specified, the <schema-owner-name> becomes the owner of the schema. If this clause is not specified, the authorization ID that issued this statement becomes the owner of the schema.

For more information, see the CREATE SCHEMA statement. See also “Schema name restrictions and recommendations” on page 215.

Copying schemas

The db2move utility and the ADMIN_COPY_SCHEMA procedure allow you to quickly make copies of a database schema. Once a model schema is established, you can use it as a template for creating new versions.

Use the ADMIN_COPY_SCHEMA procedure to copy a single schema within the same database or the db2move utility with the -co COPY action to copy a single schema or multiple schemas from a source database to a target database. Most database objects from the source schema are copied to the target database under the new schema.

Troubleshooting tips

Both the ADMIN_COPY_SCHEMA procedure and the db2move utility invoke the LOAD command. While the load is processing, the table spaces wherein the database target objects reside are put into backup pending state.

ADMIN_COPY_SCHEMA procedure

Using this procedure with the COPYNO option places the table spaces wherein the target object resides into backup pending state, as described in the note above. To get the table space out of the set integrity pending state, this procedure issues a SET INTEGRITY statement. In situations where a target table object has referential constraints defined, the target table is also placed in the set integrity pending state. Because the table spaces are already in backup pending state, the ADMIN_COPY_SCHEMA procedure’s attempt to issue a SET INTEGRITY statement will fail.

To resolve this situation, issue a BACKUP DATABASE command to get the affected table spaces out of backup pending state. Next, look at the **Statement_text** column of the error table generated by this procedure to

find a list of tables in the set integrity pending state. Then issue the SET INTEGRITY statement for each of the tables listed to take each table out of the set integrity pending state.

db2move utility

This utility attempts to copy all allowable schema objects with the exception of the following types:

- table hierarchy
- staging tables (not supported by the load utility in multiple partition database environments)
- jars (Java™ routine archives)
- nicknames
- packages
- view hierarchies
- object privileges (All new objects are created with default authorizations)
- statistics (New objects do not contain statistics information)
- index extensions (user-defined structured type related)
- user-defined structured types and their transform functions

Unsupported type errors

If an object of one of the unsupported types is detected in the source schema, an entry is logged to an error file, indicating that an unsupported object type is detected. The COPY operation will still succeed—the logged entry is meant to inform you of objects not copied by this operation.

Objects not coupled with schemas

Objects that are not coupled with a schema, such as table spaces and event monitors, are not operated on during a copy schema operation. You should create them on the target database before the copy schema operation is invoked.

Replicated tables

When copying a replicated table, the new copy of the table is not enabled for replication. The table is re-created as a regular table.

Different instances

The source database must be cataloged if it does not reside in the same instance as the target database.

SCHEMA_MAP option

When using the SCHEMA_MAP option to specify a different schema name on the target database, the copy schema operation will perform only minimal parsing of the object definition statements to replace the original schema name with the new schema name. For example, any instances of the original schema that appear inside the contents of an SQL procedure are not replaced with the new schema name. Thus the copy schema operation may fail to recreate these objects. You can use the DDL in the error file to manually recreate these failed objects after the copy operation completes.

Interdependencies between objects

The copy schema operation attempts to recreate objects in an order that satisfies the interdependencies between these objects. For example, if a table T1 contains a column that references a user-defined function U1, then it will recreate U1 before recreating T1. However, dependency information for procedures is not readily available in the catalogs, so when recreating procedures, the copy schema operation will first attempt to recreate all

procedures, then retry to recreate those that failed (on the assumption that if they depended on a procedure that was successfully created during the previous attempt, then during a subsequent attempt they will be recreated successfully). The operation will continually try to recreate these failed procedures as long as it is able to successfully recreate one or more during a subsequent attempt. During every attempt at recreating a procedure, an error (and DDL) is logged into the error file. You might see many entries in the error file for the same procedures, but these procedures might have even been successfully recreated during a subsequent attempt. You should query the SYSCAT.PROCEDURES table upon completion of the copy schema operation to determine if these procedures listed in the error file were successfully recreated.

For more information, see the ADMIN_COPY_SCHEMA procedure and the db2move utility.

Example of schema copy using the ADMIN_COPY_SCHEMA procedure

Use the ADMIN_COPY_SCHEMA procedure as shown below to copy a single schema within the same database.

```
DB2 "SELECT SUBSTR(OBJECT_SCHEMA,1, 8)
      AS OBJECT_SCHEMA, SUBSTR(OBJECT_NAME,1, 15)
      AS OBJECT_NAME, SQLCODE, SQLSTATE, ERROR_TIMESTAMP, SUBSTR(DIAGTEXT,1, 80)
      AS DIAGTEXT, SUBSTR(STATEMENT,1, 80)
      AS STATEMENT FROM COPYERRSCH.COPYERRTAB"

CALL SYSPROC.ADMIN_COPY_SCHEMA('SOURCE_SCHEMA', 'TARGET_SCHEMA',
                                'COPY', NULL, 'SOURCETS1', 'SOURCETS2', 'TARGETTS1', 'TARGETTS2',
                                SYS_ANY', 'ERRORSCHEMA', 'ERRORNAME')
```

The output from this SELECT statement is shown below:

```
OBJECT_SCHEMA OBJECT_NAME      SQLCODE      SQLSTATE ERROR_TIMESTAMP
-----
SALES          EXPLAIN_STREAM      -290 55039    2006-03-18-03.22.34.810346

DIAGTEXT
-----
[IBM][CLI Driver][DB2/LINUX8664] SQL0290N Table space access is not allowed.

STATEMENT
-----
set integrity for "SALES"."ADVISE_INDEX", "SALES"."ADVISE_MQT", "SALES"."

1 record(s) selected.
```

Examples of schema copy using the db2move utility

Use the db2move utility with the -co COPY action to copy one or more schemas from a source database to a target database. Once a model schema is established, you can use it as a template for creating new versions.

Example 1: Using the -c COPY options

The following example of the db2move -co COPY options copies the schema BAR and renames it FOO from the sample database to the target database:

```
db2move sample COPY -sn BAR -co target_db target schema_map
"((BAR,FOO))" -u userid -p password
```

The new (target) schema objects are created using the same object names as the objects in the source schema, but with the target schema qualifier. It is possible to create copies of tables with or without the data from the source table. The source and target databases can be on different systems.

Example 2: Specifying table space name mappings during the COPY operation

The following example shows how to specify specific table space name mappings to be used instead of the table spaces from the source system during a db2move COPY operation. You can specify the SYS_ANY keyword to indicate that the target table space should be chosen using the default table space selection algorithm. In this case, the db2move utility chooses any available table space to be used as the target:

```
db2move sample COPY -sn BAR -co target_db target schema_map  
"((BAR,F00))" tablespace_map "(SYS_ANY)" -u userid -p password
```

The SYS_ANY keyword can be used for all table spaces, or you can specify specific mappings for some table spaces, and the default table space selection algorithm for the remaining:

```
db2move sample COPY -sn BAR -co target_db target schema_map "  
((BAR,F00))" tablespace_map "(TS1, TS2), (TS3, TS4), SYS_ANY)"  
-u userid -p password
```

This indicates that table space TS1 is mapped to TS2, TS3 is mapped to TS4, but the remaining table spaces use a default table space selection algorithm.

Example 3: Changing the object owners after the COPY operation

You can change the owner of each new object created in the target schema after a successful COPY. The default owner of the target objects is the connect user. If this option is specified, ownership is transferred to a new owner as demonstrated:

```
db2move sample COPY -sn BAR -co target_db target schema_map  
"((BAR,F00))" tablespace_map "(SYS_ANY)" owner jrichards  
-u userid -p password
```

The new owner of the target objects is jrichards.

The db2move utility must be invoked on the target system if source and target schemas reside on different systems. For copying schemas from one database to another, this action requires a list of schema names to be copied from a source database, separated by commas, and a target database name.

To copy a schema, issue db2move from an OS command prompt as follows:

```
db2move <dbname> COPY -co <COPY- options>  
-u <userid> -p <password>
```

Restarting a failed copy schema operation

Errors occurring during a db2move COPY operation can be handled in various ways depending on the type of object being copied, or the phase during which the COPY operation failed (that is, either the recreation of objects phase, or the loading of data phase).

The db2move utility reports errors and messages to the user using message and error files. Copy schema operations use the COPYSHEMA_<timestamp>.MSG message file, and the COPYSHEMA_<timestamp>.err error file. These files are

created in the current working directory. The current time is appended to the filename to ensure uniqueness of the files. It is up to the user to delete these message and error files when they are no longer required.

Note: It is possible to have multiple db2move instances running simultaneously. The COPY option does not return any SQLCODES. This is consistent with db2move behavior.

Object types

The type of object being copied can be categorized as one of two types : physical objects and business objects.

A physical object refers to an object that physically resides in a container, such as tables, indexes and user-defined structured types. A business object refers to cataloged objects that do not reside in containers, such as views, user-defined structured types (UDTs), and aliases.

Errors occurring during the recreation of a physical object cause the utility to rollback, whereas, errors during the recreation of a logical object do not.

Restarting the copy schema operation

After addressing the issues causing the load operations to fail (described in the error file), you can reissue the db2move -COPY command using the -tf option to specify which tables to copy and to populate with data (passing in the LOADTABLE.err filename) as shown in the following syntax:

```
db2move sourcedb COPY -tf LOADTABLE.err -co TARGET_DB mytarget_db
-mode load_only
```

You can also input the table names manually using the -tn option, as shown in the following syntax:

```
db2move sourcedb COPY -tn "F00"."TABLE1","F00 1"."TAB 444",
-co TARGET_DB mytarget_db -mode load_only
```

Note: The load_only mode requires inputting at least one table using the -tn or -tf option.

Examples

Errors occurring during a db2move COPY schema operation can be handled in various ways depending on the type of object being copy copied, or the phase of the COPY operation failure.

The db2move utility reports schema copy errors and messages in the following message and error files:

- COPYSHEMA <timestamp>.MSG message file
- COPYSHEMA_<timestamp>.err error file

These files are created in the current working directory. The current time is appended to the filename to ensure uniqueness of the files. These message and error files should be deleted when they are no longer required.

Note: It is possible to have multiple db2move instances running simultaneously. The COPY option does not return any SQLCODES. This is consistent with db2move behavior.

Example 1: Schema copy errors related to physical objects

Failures which occur during the recreation of physical objects on the target database, are logged in the error file `COPYSCHEMA_<timestamp>.err`. For each failing object, the error file contains information such as object name, object type, DDL text, time stamp, and a string formatted sqlca (sqlca field names, followed by their data values).

Sample output for the `COPYSCHEMA_<timestamp>.err` error file:

```
1. schema: F00.T1
   Type:      TABLE
   Error Msg: SQL0104N An unexpected token 'F00.T1'...
   Timestamp: 2005-05-18-14.08.35.65
   DDL:       create view F00.v1

2. schema:  F00.T3
   Type:      TABLE
   Error Msg: SQL0204N F00.V1 is an undefined name.
   Timestamp: 2005-05-18-14.08.35.68
   DDL:       create table F00.T3
```

If any errors creating physical objects are logged at the end of the recreation phase and before attempting the load phase, the `db2move` utility fails and an error is returned. All object creation on the target database is rolled back, and all internally created tables are cleaned up on the source database. The rollback occurs at the end of the recreation phase after attempting to recreate each object, rather than after the first failure, in order to gather all possible errors into the error file. This allows you the opportunity to fix any problems before restarting the `db2move` operation. If there are no failures, the error file is deleted.

Example 2: Schema copy errors related to business objects

Failures that occur during the recreation of business objects on the target database, do not cause the `db2move` utility to fail. Instead, these failures are logged in the `COPYSCHEMA_<timestamp>.err` error file. Upon completion of the `db2move` utility, you can examine the failures, address any issues, and manually recreate each failed object (the DDL is provided in the error file for convenience).

If an error occurs while `db2move` is attempting to repopulate table data using the load utility, the `db2move` utility will not fail. Rather, generic failure information is logged to the `COPYSCHEMA_<timestamp>.err` file (object name, object type, DDL text, time stamp, sqlca, and so on), and the fully qualified name of the table is logged into another file, `LOADTABLE_<timestamp>.err`. Each table is listed per line to satisfy the `db2move -tf` option format, similar to the following:

```
"F00"."TABLE1"
"F00 1"."TAB 444"
```

Example 3: Other types of db2move failures

Internal operations such as memory errors, or file system errors can cause the `db2move` utility to fail.

Should the internal operation failure occur during the ddl recreation phase, all successfully created objects are rolled back from the target schema, and all internally created tables such as the DMT table and the `db2look` table, are cleaned up on the source database.

Should the internal operation failure occur during the load phase, all successfully created objects remain on the target schema. All tables that experience a failure during a load operation, and all tables, which have not

yet been loaded are logged in the LOADTABLE.err error file. You can then issue the db2move COPY command using the LOADTABLE.err as discussed in Example 2. If the db2move utility abends (for example a system crash, the utility traps, the utility is killed, and so on), then the information regarding which tables still need to be loaded is lost. In this case, you can drop the target schema using the ADMIN_DROP_SCHEMA procedure and reissue the db2move COPY command.

Regardless of what error you might encounter during an attempted copy schema operation, you always have the option of dropping the target schema using the ADMIN_DROP_SCHEMA procedure and reissuing the db2move COPY command.

Dropping schemas

Before dropping a schema, all objects that were in that schema must be dropped or moved to another schema. The schema name must be in the catalog when attempting the DROP statement; otherwise an error is returned.

To drop a schema using the command line, enter:

```
DROP SCHEMA <name> RESTRICT
```

In the following example, the schema "joeschma" is dropped:

```
DROP SCHEMA joeschma RESTRICT
```

The RESTRICT keyword enforces the rule that no objects can be defined in the specified schema for the schema to be deleted from the database, and it must be specified.

Part 3. Database objects

Logical database design consists of defining database objects.

The following database objects can be created within a DB2 database:

- Tables
- Constraints
- Indexes
- Triggers
- Sequences
- Views

These database objects can be created using graphical user interfaces or by explicitly executing statements. The statements used to create these database objects are called Data Definition Language (DDL) statements and are generally prefixed by the keywords `CREATE` or `ALTER`.

Understanding the features and functionality that each of these database objects provides is important to implement a good database design that meets your current business's data storage needs while remaining flexible enough to accommodate expansion and growth over time.

Chapter 11. Tables

Tables are logical structures maintained by the database manager. Tables are made up of columns and rows.

At the intersection of every column and row is a specific data item called a *value*. A *column* is a set of values of the same type or one of its subtypes. A *row* is a sequence of values arranged so that the *n*th value is a value of the *n*th column of the table.

An application program can determine the order in which the rows are populated into the table, but the actual order of rows is determined by the database manager, and typically cannot be controlled. Multidimensional clustering (MDC) provides some sense of clustering, but not actual ordering between the rows.

Types of tables

Depending on your environment, you will need to create one or more tables in your DB2 databases to store your data. When you create tables, you specify the type of content in each of the columns in the tables, and you define other characteristics, such as the primary key and check constraints for the enforcement of business rules. When you create tables, you also need to consider the type that is best suited for your needs.

You can create all of the following types of tables using the CREATE TABLE statement, except for (Declared) Global temporary tables:

Base tables

These types of tables hold persistent data.

Regular tables

These types of tables are implemented as a heap. Regular tables with indexes are the "general purpose" table choice.

Append mode tables

These types of tables are regular tables that are optimized primarily for INSERTs. Regular tables are placed into append mode through an ALTER TABLE statement. Append mode is best used for tables where clustering to any specific index is not important, and the insert rate is both high and also there are not many or no deletes to the table.

Result tables

These types of tables are made up of sets of rows that the database manager selects or generates from one or more tables to satisfy a query.

Summary tables

These types of tables are defined by a query that is also used to determine the data in the table. Summary tables can be used to improve the performance of queries. If the database manager determines that a portion of a query can be resolved using a summary table, the database manager can rewrite the query to use the summary table. This decision is based on database configuration settings, such as the CURRENT REFRESH AGE and the CURRENT QUERY OPTIMIZATION special registers.

Typed tables

A table can define the data type of each column separately, or base the

types on the attributes of a user-defined structured type. This is called a *typed table*. A user-defined structured type may be part of a type hierarchy. A *subtype* inherits attributes from its *supertype*. Similarly, a typed table can be part of a table hierarchy. A *subtable* inherits columns from its *supertable*. Note that the term *subtype* applies to a user-defined structured type and all user-defined structured types that are below it in the type hierarchy. A *proper subtype* of a structured type T is a structured type below T in the type hierarchy. Similarly, the term *subtable* applies to a typed table and all typed tables that are below it in the table hierarchy. A *proper subtable* of a table T is a table below T in the table hierarchy.

(Declared) Global temporary tables

These types of tables, also referred to as user-defined temporary tables, are used by applications that work with data in the database. Results from manipulation of the data need to be stored temporarily in a table. A user temporary table space must exist before creating global temporary tables.

Note: The description of global temporary tables does not appear in the system catalog thus making it not persistent for, and not able to be shared with, other applications. When the application using this table terminates or disconnects from the database, any data in the table is deleted and the table is implicitly dropped.

Global temporary tables do not support:

- LOB-type columns (or distinct-type columns based on LOBs)
- User-defined type columns
- LONG VARCHAR columns
- XML columns

These types of tables are created with a DECLARE GLOBAL TEMPORARY TABLE statement and are used to hold temporary data on behalf of a single application. This table is dropped implicitly when the application disconnects from the database.

Multidimensional clustering (MDC) tables

These types of tables are implemented as tables that are physically clustered on more than one key, or dimension, at the same time. MDC tables are used in data warehousing and large database environments. Clustering indexes on regular tables support single-dimensional clustering of data. MDC tables provide the benefits of data clustering across more than one dimension.

Note: MDC is a type of table, but can coexist with a partitioned table or can be partitioned table as well. Thus, it is not mutually exclusive to other types of tables. A partitioned table can also be an APPEND table. Some other combinations are not allowed, for example, MDC and APPEND, RCT and any other table, or other combinations). Also, MDC provides *guaranteed clustering* within the composite dimension, while for regular tables with a clustering index, clustering is attempted by the database manager, but not guaranteed and it typically degrades over time.

Range-clustered tables (RCT)

These types of tables are implemented as sequential clusters of data that provide fast, direct access. Each record in the table has a predetermined record ID (RID) which is an internal identifier used to locate a record in a table. RTC tables are used where the data is tightly clustered across one or more columns in the table. The largest and smallest values in the columns

define the range of possible values. You use these columns to access records in the table; this is the most optimal method of utilizing the predetermined record identifier (RID) aspect of RCT tables.

Partitioned tables

These types of tables are implemented with data divided across multiple data partitions according to values in the table partitioning key columns for the table. Partitioned tables allow easier roll-in and rollout of table data, easier administration, flexible index placement, and better query processing than regular tables.

For each table that holds your data, consider which of the possible table types would best suit your needs. For example, if you have data records that will be loosely clustered (not monotonically increasing), consider using a regular table and indexes. If you have data records that will have duplicate (not unique) values in the key, you should not use a range-clustered table. If you cannot afford to preallocate a fixed amount of storage on disk for the range-clustered tables you might want, you should not use this type of table. These factors will help you to determine whether you have data that can be used as a range-clustered table.

Designing tables

When designing tables, you need to be familiar with certain concepts, determine the space requirements for tables and user data, and determine whether you will take advantage of certain features, such as space compression and optimistic locking.

When designing partitioned tables, you need to be familiar with the partitioning concepts, such as:

- Data organization schemes
- Table partitioning keys
- Keys used for distributing data across data partitions
- Keys used for MDC dimensions

For these and other partitioning concepts, see “Table partitioning and data organization schemes” on page 253.

Table design concepts

When designing tables, you need to be familiar with some related concepts.

Specifying column data types

When defining columns, you need to name the columns, define the type of data that will be included in those columns (called *data types*), and define the length of the data for each column in the table you are creating.

Character data stored as binary data

Small integer

This data type is used to store binary integer values that have a precision of 15 bits. The range for a small integer value is -32 768 to +32 767. The small integer data type uses the smallest amount of storage space possible to store numerical values (2 bytes of space is required for each value stored). The term SMALLINT is used to declare a small integer column in a table definition.

Integer

This data type is used to store binary integer values that have a

precision of 31 bytes. Although the integer data type requires twice as much storage space as the small integer data type (4 bytes of space is required for each value stored), its range of values is much greater. The range for an integer value is -2 147 483 648 to +2 147 483 647. The terms INTEGER and INT can be used to declare an integer column in a table definition.

Big integer

This data type is used to store binary integer values that have a precision of 63 bits on platforms that provide support for 64 bit integers. Processing large numbers that are stored as big integers is more efficient than processing similar numbers that have been stored as decimal values. In addition, calculations performed with big integer values are more precise than calculations performed with real or double values.

This data type requires four times as much storage space as the small integer data type (8 bytes of space is required for each value stored.) The range for a big integer is -9 223 372 036 854 775 808 to +9 223 372 036 854 775 807. The term BIGINT is used to declare a big integer column in a table definition.

Decimal

This data type is used to store numbers that contain both whole and fractional parts; the parts are combined and stored in packed decimal format. A precision (the total number of digits) and a scale (the number of digits to use for the fractional part of the number) must be specified whenever a decimal data type is declared. The range for the precision of a decimal is 1 to 31. The amount of storage space needed to store a decimal value can be calculated by using the following equation: *precision* divided by 2 (truncated) + 1 = *bytes of space required*.

For example, a DECIMAL(8,2) value would require 5 bytes of storage space (8 divided by 2 = 4; 4 + 1 = 5), whereas a DECIMAL(7,2) value would require 4 bytes of storage space (7 divided by 2 = 3.5 (truncated to 3); 3 + 1 = 4).

The terms DECIMAL, DEC, NUMERIC, and NUM can all be used to declare a decimal column in a table definition.

Note: If the precision and scale values are not provided for a decimal column definition, by default, a precision value of 5 and a scale value of 0 are used (therefore, 3 bytes of storage space is needed).

Single-precision floating point

This data type is used to store a 32-bit approximation of a real number. Although the single-precision floating-point data type and the integer data type require the same amount of storage space (4 bytes of space is required for each value stored), the range for a single-precision floating-point number is much greater: $10E^{-38}$ to $10E^{+38}$.

The terms REAL and FLOAT can be used to declare a single-precision floating-point column in a table definition. However, if the term FLOAT is used, the length specified for the column must be between 1 and 24—the FLOAT can be used to

represent both single- and double-precision floating-point data types; the length specified determines which actual data type is to be used.

Double-precision floating point

The double-precision floating-point data type is used to store a 64-bit approximation of a real number. Although the double-precision floating-point data type requires the same amount of storage space as the big integer data type (8 bytes of space is required for each value stored), the range for a double-precision floating-point number is the largest possible: -1.79760^{+308} to $-2.225E^{-307}$, 0, and $2.225E^{-307}$ to $-1.79769E^{+308}$.

Fixed-length character string

This data type is used to store character and character string data that has a specific length that does not exceed 254 characters. The terms CHARACTER and CHAR can be used to declare a fixed-length character string column in a table definition; the length of the character string data to be stored must be specified whenever a fixed-length character string data type is declared. The amount of storage space needed to store a fixed-length character string value can be determined by using the following equation: *fixed length* × 1 = *bytes of space required*. For example, a CHAR(8) value would require 8 bytes of storage.

Note: When fixed-length character string data types are used, storage space can be wasted if the actual length of the data is significantly smaller than the length specified when the column was defined. For example, if the values YES and NO were to be stored in a column that was defined as CHAR(20). Therefore, the fixed length specified for a fixed-length character string column should be as close as possible to the actual length of the data that will be stored in the column.

Variable length character data

This data type is used to store character string data that varies in length. Varying-length character string data can be up to 32 672 characters long; however, the actual length allowed is governed by one restriction: the data must fit on a single table space page. This means that for a table that resides in a table space that used 4K pages, varying-length character string data cannot be more than 4 092 characters long; for a table that resides in a table space that used 8K pages, varying-length character string data cannot be more than 8 188 characters long and so on, up to 32K. Because table spaces are created with 4K pages by default, you must explicitly create a table space with a larger page size if you want to use a varying length character string data type to store strings that contain more than 4 092 characters.

Note:

- You must also have sufficient space in the table row to accommodate the character string data. In other words, the storage requirements for other columns in the table must be added to the storage requirements of the character string data and the total amount of storage space needed must not exceed the size of the table space's page.
- When a varying-length string data value is updated and the new value is larger than the original value, the record containing the value will be moved to another page in the table. Such records are known as pointer

records. Too many pointer records can cause a significant decrease in performance because multiple pages must be retrieved in order to process a single data record.

The terms CHARACTER VARYING, CHAR VARYING, and VARCHAR can be used to declare a varying-length character string column in a table definition. When a varying length character string column is defined, the maximum number of characters that are expected to be stored in that column must be specified as part of the declaration. Subsequent character string data values that are stored in the column can be shorter than or equal to the maximum length specified; if they are longer, they will not be stored and error is returned.

The amount of storage space needed to store a varying-length character string value can be determined by using the following equation: $(string\ length \times 1) + 4 = bytes\ of\ space\ required$. Thus, if a character string containing 30 characters were stored using a VARCHAR(30) data type, that particular value would require 34 bytes of storage space. (All character strings using this data type would have to be less than or equal to 30 characters in length.)

Variable length long character data

The varying-length long character string data type is also used to store string data that varies in length. This data type is used to store character string data that is less than or equal to 32 700 characters long in a table that resides in a table space that uses 4K pages. In other words, when the varying-length long character string is used, the page size/character string data length restrictions that apply to varying-length character string data are not applicable.

The term LONG VARCHAR is used to declare a varying-length long character string column in a table definition. The amount of storage space needed to store a varying-length character string value can be determined by using this equation: $(string\ length \times 1) + 24 = bytes\ of\ space\ required$. The LONG VARCHAR and LONG VARGRAPHIC data types are deprecated and might be removed in a future release. When choosing a data type for a column, use data types such as VARCHAR, VARGRAPHIC, CLOB, or DBCLOB since these will continue to be supported in future releases and are recommended for portable applications.

Note: The FOR BIT DATA clause can be used with any character string data type when declaring a column in a table definition. If this clause is used, code page conversions will not be performed during data exchange operations and the data itself will be treated and compared as binary (bit) data.

Character large objects (CLOBs)

A CLOB (character large object) value can be up to 2 gigabytes (2 147 483 647 bytes) long. A CLOB is used to store large SBCS or mixed (SBCS and MBCS) character-based data (such as documents written with a single character set) and, therefore, has an SBCS or mixed code page associated with it.

Variable length character stored as binary data (Large objects–LOBs and Binary large objects–BLOBs)

The term *large object* and the generic acronym LOB refer to the BLOB, CLOB, or DBCLOB data type. LOB values are subject to restrictions that apply to LONG VARCHAR values, as described in the section “Variable length character data”. These restrictions apply even if the length attribute

of the LOB string is 254 bytes or less. This data type is used to store binary string data that varies in length. It is frequently used to store nontraditional data such as documents, graphic images, pictures, audio, and video.

Note: Binary large objects data cannot be manipulated by SQL the same way that other data can. For example, binary large object values cannot be sorted.

Unicode data

All data types supported are also supported in a Unicode database. In particular, graphic string data is supported for a Unicode database, and is stored in UCS-2 encoding. Every client, including SBCS clients, can work with graphic string data types in UCS-2 encoding when connected to a Unicode database.

Date and time data (Timestamps)

The *date* data type is used to store a three-part value (year, month, and day) that designates a valid calendar data. The range for the year part is 0001 to 9999; the range for the month part is 1 to 12; and the range for the day part is 1 to n (28, 29, 30, or 31) where n is dependent upon the month part and whether the year part corresponds to a leap year. Externally, the date data type appears to be a fixed-length character string data type that has a length of 10. However, internally, the date data type requires much less storage space—4 bytes of space is required for each value stored, because date values are stored as packed strings. The term DATE is used to declare a date column in a table definition.

The time data type is used to store a three-part value (hours, minutes, and seconds) that designates a valid time of day under a 24-hour clock. The range for the hours part is 0 to 24; the range for the minutes part is 0 to 59; and the range for the seconds part is also 0 to 59. (If the hours part is set to 24, the minutes and seconds parts must be set to 0.) Externally, the time data type appears to be a fixed-length character string data type that has a length of 8. However, like date values, time values are stored as packed strings—in this case, 3 bytes of space is required for each time value stored. The term TIME is used to declare a time column in a table definition.

Like dates, the representation of time varies in different parts of the world. Thus the format of a time value is also determined by the territory code associated with the database being used. Table 49 shows the time formats that are available, along with an example of their string representation:

Table 49. Date formats (YYYY = Year, MM = Month, DD = Day)

Format name	Abbreviation	Date string format
International Standards Organization	ISO	YYYY-MM-DD
IBM USA Standard	USA	MM/DD/YYYY
IBM European Standard	EUR	MM/DD/YYYY
Japanese Industrial Standard	JIS	YYYY-MM-DD
Site specific	LOC	Based on database's territory code

Numeric data

All numbers have a sign and a precision. The sign is considered positive if

the value of a number is zero. The precision is the number of bits or digits excluding the sign. See the data type section in the description of the CREATE TABLE statement.

Monetary data

Version 9.5 introduces DECFLOAT, a decimal floating-point data type that is useful in business applications (for example, financial applications) that deal with exact decimal values. Binary floating-point data types (REAL and DOUBLE), which provide binary approximations for decimal data, are not appropriate in such applications. DECFLOAT combines the accuracy of DECIMAL with some of the performance advantage of FLOAT, which is beneficial in applications where monetary values are being manipulated.

XML data

The XML data type is used to define columns of a table that store XML values, where all stored XML values must be well-formed XML documents. The introduction of this native XML data type provides the ability to store well-formed XML documents in their native hierarchical format in the database alongside other relational data.

Generated columns

A generated column is defined in a table where the stored value is computed using an expression, rather than being specified through an insert or update operation.

When creating a table where it is known that certain expressions or predicates will be used all the time, you can add one or more generated columns to that table. By using a generated column there is opportunity for performance improvements when querying the table data.

For example, there are two ways in which the evaluation of expressions can be costly when performance is important:

1. The evaluation of the expression must be done many times during a query.
2. The computation is complex.

To improve the performance of the query, you can define an additional column that would contain the results of the expression. Then, when issuing a query that includes the same expression, the generated column can be used directly; or, the query rewrite component of the optimizer can replace the expression with the generated column.

Where queries involve the joining of data from two or more tables, the addition of a generated column can allow the optimizer a choice of possibly better join strategies.

Generated columns will be used to improve performance of queries. As a result, generated columns will likely be added after the table has been created and populated.

Examples

The following is an example of defining a generated column on the CREATE TABLE statement:

```
CREATE TABLE t1 (c1 INT,
                 c2 DOUBLE,
                 c3 DOUBLE GENERATED ALWAYS AS (c1 + c2)
                 c4 GENERATED ALWAYS AS
                   (CASE WHEN c1 > c2 THEN 1 ELSE NULL END))
```

After creating this table, indexes can be created using the generated columns. For example,

```
CREATE INDEX i1 ON t1(c4)
```

Queries can take advantage of the generated columns. For example,

```
SELECT COUNT(*) FROM t1 WHERE c1 > c2
```

can be written as:

```
SELECT COUNT(*) FROM t1 WHERE c4 IS NOT NULL
```

Another example:

```
SELECT c1 + c2 FROM t1 WHERE (c1 + c2) * c1 > 100
```

can be written as:

```
SELECT c3 FROM t1 WHERE c3 * c1 > 100
```

Auto numbering and identifier columns

An identity column provides a way for DB2 to automatically generate a unique numeric value for each row that is added to the table.

When creating a table in which you need to uniquely identify each row that will be added to the table, you can add an identity column to the table. To guarantee a unique numeric value for each row that is added to a table, you should define a unique index on the identity column or declare it a primary key.

Other uses of an identity column are an order number, an employee number, a stock number, or an incident number. The values for an identity column can be generated by the DB2 database manager: ALWAYS or BY DEFAULT.

An identity column defined as GENERATED ALWAYS is given values that are always generated by the DB2 database manager. Applications are not allowed to provide an explicit value. An identity column defined as GENERATED BY DEFAULT gives applications a way to explicitly provide a value for the identity column. If the application does not provide a value, then DB2 will generate one. Since the application controls the value, DB2 cannot guarantee the uniqueness of the value. The GENERATED BY DEFAULT clause is meant for use for data propagation where the intent is to copy the contents of an existing table; or, for the unload and reloading of a table.

Once created, you cannot alter the table description to include an identity column.

If rows are inserted into a table with explicit identity column values specified, the next internally generated value is not updated, and might conflict with existing values in the table. Duplicate values will generate an error message if the uniqueness of the values in the identity column is being enforced by a primary-key or a unique index that has been defined on the identity column.

To define an identity column on a new table, use the AS IDENTITY clause on the CREATE TABLE statement.

Example

The following is an example of defining an identity column on the CREATE TABLE statement:

```
CREATE TABLE table (col1 INT,
                    col2 DOUBLE,
                    col3 INT NOT NULL GENERATED ALWAYS AS IDENTITY
                    (START WITH 100, INCREMENT BY 5))
```

In this example the third column is the identity column. You can also specify the value used in the column to uniquely identify each row when added. Here the first row entered has the value of “100” placed in the column; every subsequent row added to the table has the associated value increased by five.

Constraining column data with constraints, defaults, and null settings

Data often must adhere to certain restrictions or rules. Such restrictions may apply to single pieces of information, such as the format and sequence numbers, or they may apply to several pieces of information.

Nullability of column data values

Null values represent unknown states. By default, all of the built-in data types support the presence of null values. However, some business rules might dictate that a value must always be provided for some columns, for example, emergency information. For this condition, you can use the NOT NULL constraint to ensure that a given column of a table is never assigned the null value. Once a NOT NULL constraint has been defined for a particular column, any insert or update operation that attempts to place a null value in that column will fail.

Default column data values

Just as some business rules dictate that a value must always be provided, other business rules may dictate what that value should be, for example, the gender of an employee must be either M or F. The column default constraint is used to ensure that a given column of a table is always assigned a predefined value whenever a row that does not have a specific value for that column is added to the table. The default value provided for a column can be null, a constraint value that is compatible with the data type of the column, or a value that is provided by the database manager. For more information, see: “Default column and data type definitions” on page 235.

Keys A key is a single column or a set of columns in a table or index that can be used to identify or access a specific row of data. Any column can be part of a key and the same column can be part of more than one key. A key that consists of a single column is called an atomic key; a key that is composed of more than one column is called a composite key. In addition to having atomic or composite attributes, keys are classified according to how they are used to implement constraints:

- A unique key is used to implement unique constraints.
- A primary key is used to implement entity integrity constraints. (A primary key is a special type of unique key that does not support null values.)
- A foreign key is used to implement referential integrity constraints. (Foreign keys must reference primary keys or unique keys; foreign keys do not have corresponding indexes.)

Keys are normally specified during the declaration of a table, an index, or a referential constraint definition.

Constraints

Constraints are rules that limit the values that can be inserted, deleted, or

updated in a table. There are check constraints, primary key constraints, referential constraints, unique constraints, unique key constraints, foreign key constraints, and informational constraints. For details about each of these types of constraints, see: Chapter 12, “Constraints,” on page 267 or “Types of constraints” on page 267.

Default column and data type definitions:

Certain columns and data types have predefined or assigned default values.

For example, default column values for the various data types are as follows:

- *NULL*
- *0* Used for small integer, integer, decimal, single-precision floating point, and double-precision floating point.
- *Blank*: Used for fixed-length and fixed-length double-byte character strings.
- *Zero-length string*: Used for varying-length character strings, binary large objects, character large objects, and double-byte character large objects.
- *Date*: This is the system date at the time the row is inserted (obtained from the `CURRENT_DATE` special register). When a date column is added to an existing table, existing rows are assigned the date January, 01, 0001.
- *Time or Timestamp*: This is the system time or system date/time of the at the time the statement is inserted (obtained from the `CURRENT_TIME` special register). When a time column is added to an existing table, existing rows are assigned the time 00:00:00 or a timestamp that contains the date January, 01, 0001 and the time 00:00:00.

Note: All the rows get the same default time/timestamp value for a given statement.

- *Distinct user-defined data type*: This is the system-defined default value for the base data type of the distinct user-defined data type (cast to the distinct user-defined data type).

Primary key, referential integrity, check, and unique constraints

Constraints are rules that limit the values that can be inserted, deleted, or updated in a table.

Primary key constraints

A primary key constraint is a column or combination of columns that has the same properties as a unique constraint. You can use a primary key and foreign key constraints to define relationships between tables.

Referential integrity (or foreign key) constraints

A foreign key constraint (also referred to as a referential constraint or a referential integrity constraint) is a logical rule about values in one or more columns in one or more tables. For example, a set of tables shares information about a corporation’s suppliers. Occasionally, a supplier’s name changes. You can define a referential constraint stating that the ID of the supplier in a table must match a supplier ID in the supplier information. This constraint prevents insert, update, or delete operations that would otherwise result in missing supplier information.

Check constraints

A (table) check constraint sets restrictions on data added to a specific table.

Unique constraints

A unique constraint (also referred to as a unique key constraint) is a rule

that forbids duplicate values in one or more columns within a table. Unique and primary keys are the supported unique constraints.

Unicode table and data considerations

The Unicode character encoding standard is a fixed-length, character encoding scheme that includes characters from almost all of the living languages of the world.

For more information on Unicode table and data considerations, see:

- “Unicode character encoding” in *Internationalization Guide*
- “Character comparisons based on collating sequences” in *Internationalization Guide*
- “Date and time formats by territory code” in *Internationalization Guide*
- “Conversion table files for euro-enabled code pages” in *Internationalization Guide*

Additional information on Unicode can be found in the latest edition of *The Unicode Standard*, and from the Unicode Consortium web site at www.unicode.org.

Space requirements for tables

When designing tables, you need to take space requirements into considerations

Long field (LF) data

Long field (LF) data is stored in a separate table object that is structured differently than the storage space for other data types.

Data is stored in 32-KB areas that are broken up into segments whose sizes are “powers of two” times 512 bytes. (Hence these segments can be 512 bytes, 1024 bytes, 2048 bytes, and so on, up to 32 768 bytes.)

Long field data types (LONG VARCHAR or LONG VARGRAPHIC) are stored in a way that enables free space to be reclaimed easily. Allocation and free space information is stored in 4 KB allocation pages, which appear infrequently throughout the object.

The amount of unused space in the object depends on the size of the long field data, and whether this size is relatively constant across all occurrences of the data. For data entries larger than 255 bytes, this unused space can be up to 50 percent of the size of the long field data.

If character data is less than the page size, and it fits into the record along with the rest of the data, the CHAR, GRAPHIC, VARCHAR, or VARGRAPHIC data types should be used instead of LONG VARCHAR or LONG VARGRAPHIC.

Large object (LOB) data

Large object (LOB) data is stored in two separate table objects that are structured differently than the storage space for other data types. To estimate the space required by LOB data, you need to consider the two table objects used to store data defined with these data types:

- *LOB Data Objects*: Data is stored in 64 MB areas that are broken up into segments whose sizes are “powers of two” times 1024 bytes. (Hence these segments can be 1024 bytes, 2048 bytes, 4096 bytes, and so on, up to 64 MB.)

To reduce the amount of disk space used by LOB data, you can specify the COMPACT option on the lob-options clause of the CREATE TABLE and the ALTER TABLE statements. The COMPACT option minimizes the amount of disk space required by allowing the LOB data to be split into

smaller segments. This process does not involve data compression, but simply uses the minimum amount of space, to the nearest 1 KB boundary. Using the COMPACT option may result in reduced performance when appending to LOB values.

The amount of free space contained in LOB data objects is influenced by the amount of update and delete activity, as well as the size of the LOB values being inserted.

- *LOB Allocation Objects:* Allocation and free space information is stored in 4 KB allocation pages that are separated from the actual data. The number of these 4 KB pages is dependent on the amount of data, including unused space, allocated for the large object data. The overhead is calculated as follows: one 4 KB page for every 64 GB, plus one 4 KB page for every 8 MB.

If character data is less than the page size, and it fits into the record along with the rest of the data, the CHAR, GRAPHIC, VARCHAR, or VARGRAPHIC data types should be used instead of BLOB, CLOB, or DBCLOB.

System catalog tables

System catalog tables are created when a database is created. The system tables grow as database objects and privileges are added to the database.

Initially, they use approximately 3.5 MB of disk space.

The amount of space allocated for the catalog tables depends on the type of table space, and the extent size of the table space containing the catalog tables. For example, if a DMS table space with an extent size of 32 is used, the catalog table space is initially allocated 20 MB of space. Note: For databases with multiple partitions, the catalog tables reside only on the database partition from which the CREATE DATABASE command was issued. Disk space for the catalog tables is only required for that database partition.

Temporary tables

Some statements require temporary tables for processing (such as a work file for sorting operations that cannot be done in memory). These temporary tables require disk space; the amount of space required is dependent upon the size, number, and nature of the queries, and the size of returned tables.

Your work environment is unique which makes the determination of your space requirements for temporary tables difficult to estimate. For example, more space may appear to be allocated for system temporary table spaces than is actually in use due to the longer life of various system temporary tables. This could occur when DB2_SMS_TRUNC_TMPTABLE_THRESH registry variable is used.

You can use the database system monitor and the table space query APIs to track the amount of work space being used during the normal course of operations.

You can use the DB2_OPT_MAX_TEMP_SIZE registry variable to limit the amount of temporary table space used by queries.

Table page sizes

Rows of table data are organized into blocks called pages. Pages can be four sizes: 4, 8, 16, and 32 kilobytes. Table data pages do not contain the data for columns

defined with LONG VARCHAR, LONG VARGRAPHIC, BLOB, CLOB or DCLOB data types. The rows in a table data page do, however, contain a descriptor of these columns.

You can create buffer pools or table spaces that have page sizes of 4 KB, 8 KB, 16 KB, or 32 KB. All tables created within a table space of a particular size have a matching page size. A single table or index object can be as large as 512 GB, assuming a 32 KB page size.

You can have a maximum of 1012 columns when you are using an 8 KB, 16 KB, or 32 KB page size. You can have a maximum of 500 columns for a 4 KB page size. Larger record identifiers (RIDs), which allow more data pages per table object and more records per page, change the required amount of memory and space used by log files and system temporary table spaces. The maximum possible number of rows you can have for a 32KB page size is approximately 2335, if the minimum row size (12) is being used.

Maximum row lengths vary, depending on page size used:

- When the page size is 4 KB, the row length can be up to 4 005 bytes.
- When the page size is 8 KB, the row length can be up to 8 101 bytes.
- When the page size is 16 KB, the row length can be up to 16 293 bytes.
- When the page size is 32 KB, the row length can be up to 32 677 bytes.

To determine the page size for a table space you must consider the following:

- For OLTP applications that perform random row read and write operations, a smaller page size is usually preferable, because it wastes less buffer pool space with unwanted rows.
- For DSS applications that access large numbers of consecutive rows at a time, a larger page size is usually better, because it reduces the number of I/O requests that are required to read a specific number of rows. There is, however, an exception to this. If your row size is smaller than $\text{pagesize} / \text{maximum rows}$, there will be wasted space on each page. In this situation, a smaller page size may be more appropriate.

Larger page sizes may allow you to reduce the number of levels in the index. Larger pages support rows of greater length. Using the default of 4 KB pages, tables are restricted to 500 columns. Larger page sizes (8 KB, 16 KB, and 32 KB) support 1012 columns. The maximum size of the table space is proportional to the page size of the table space.

Space requirements for user table data

By default, table data is stored on 4-KB pages. Each page (regardless of page size) contains 68 bytes of overhead for the database manager. This leaves 4028 bytes to hold user data (or rows), although no row on a 4-KB page can exceed 4005 bytes in length. A row will *not* span multiple pages. You can have a maximum of 500 columns when using a 4-KB page size.

Table data pages *do not* contain the data for columns defined with LONG VARCHAR, LONG VARGRAPHIC, BLOB, CLOB, or DBCLOB data types. The rows in a table data page do, however, contain a descriptor for these columns.

Rows are usually inserted into a regular table in first-fit order. The file is searched (using a free space map) for the first available space that is large enough to hold the new row. When a row is updated, it is updated in place, unless there is

insufficient space left on the page to contain it. If this is the case, a record is created in the original row location that points to the new location in the table file of the updated row.

If the ALTER TABLE APPEND ON statement is issued, data is always appended, and information about any free space on the data pages is not kept.

If the table has a clustering index defined on it, the database manager will attempt to physically cluster the data according to the key order of that clustering index. When a row is inserted into the table, the database manager will first look up its key value in the clustering index. If the key value is found, the database manager attempts to insert the record on the data page pointed to by that key; if the key value is not found, the next higher key value is used, so that the record is inserted on the page containing records having the next higher key value. If there is insufficient space on the target page in the table, the free space map is used to search neighboring pages for space. Over time, as space on the data pages is completely used up, records are placed further and further from the target page in the table. The table data would then be considered unclustered, and a table reorganization can be used to restore clustered order.

If the table is a multidimensional clustering (MDC) table, the database manager will guarantee that records are always physically clustered along one or more defined dimensions, or clustering indexes. When an MDC table is defined with certain dimensions, a block index is created for each of the dimensions, and a composite block index is created which maps cells (unique combinations of dimension values) to blocks. This composite block index is used to determine to which cell a particular record belongs, and exactly which blocks or extents in the table contains records belonging to that cell. As a result, when inserting records, the database manager searches the composite block index for the list of blocks containing records having the same dimension values, and limits the search for space to those blocks only. If the cell does not yet exist, or if there is insufficient space in the cell's existing blocks, then another block is assigned to the cell and the record is inserted into it. A free space map is still used within blocks to quickly find available space in the blocks.

The number of 4-KB pages for each user table in the database can be estimated by calculating:

$$\text{ROUND DOWN}(4028/(\text{average row size} + 10)) = \text{records_per_page}$$

and then inserting the result into:

$$(\text{number_of_records}/\text{records_per_page}) * 1.1 = \text{number_of_pages}$$

where the average row size is the sum of the average column sizes, and the factor of "1.1" is for overhead.

Note: This formula provides only an estimate. The estimate's accuracy is reduced if the record length varies because of fragmentation and overflow records.

You also have the option to create buffer pools or table spaces that have an 8 KB, 16 KB, or 32 KB page size. All tables created within a table space of a particular size have a matching page size. A single table or index object can be as large as 512 GB, assuming a 32 KB page size. You can have a maximum of 1012 columns when using an 8 KB, 16 KB, or 32 KB page size. The maximum number of columns is 500 for a 4-KB page size. Maximum row lengths also vary, depending on page size:

- When the page size is 4-KB, the row length can be up to 4005 bytes.

- When the page size is 8 KB, the row length can be up to 8101 bytes.
- When the page size is 16 KB, the row length can be up to 16 293 bytes.
- When the page size is 32 KB, the row length can be up to 32 677 bytes.

A larger page size facilitates a reduction in the number of levels in any index. If you are working with OLTP (online transaction processing) applications, that perform random row reads and writes, a smaller page size is better, because it wastes less buffer space with undesired rows. If you are working with DSS (decision support system) applications, which access large numbers of consecutive rows at a time, a larger page size is better because it reduces the number of I/O requests required to read a specific number of rows.

You cannot restore a backup image to a different page size.

You cannot import IXF data files that represent more than 755 columns.

Declared temporary tables can be created only in their own user temporary table space type. There is no default user temporary table space. Temporary tables cannot have LONG data. The tables are dropped implicitly when an application disconnects from the database, and estimates of the space requirements for their tables should take this into account.

Space compression for tables

It might be possible for tables to occupy less space when stored on disk by utilizing such features as compression for data rows, NULL values, and system default values. Through data compression, you might be able to save disk storage space by using fewer database pages to store data. Since more logical data can be stored per page, fewer pages will need to be read in order to access the same amount of logical data. This means that compression can also result in disk I/O savings. I/O speed might also increase because more logical data can be cached in the buffer pool.

To implement data compression in a database system, there are two methods you can employ:

(Space) Value compression

This method optimizes space usage for the representation of data, and the storage structures used internally by the database management system (DBMS) to store data. Value compression involves removing duplicate entries for a value, and only storing one copy. The stored copy keeps track of the location of any references to the stored value.

When creating a table, you can use the optional `VALUE COMPRESSION` clause to specify that the table is using the space saving row format at the table level and possibly at the column level.

When `VALUE COMPRESSION` is used, NULLs and zero-length data that has been assigned to defined variable-length data types (`VARCHAR`, `VARGRAPHIC`, `LONG VARCHAR`, `LONG VARGRAPHIC`, `BLOB`, `CLOB`, and `DBCLOB`) will not be stored on disk. Only overhead values associated with these data types will take up disk space.

If `VALUE COMPRESSION` is used then the optional `COMPRESS SYSTEM DEFAULT` option can also be used to further reduce disk space usage. Minimal disk space is used if the inserted or updated value is equal to the system default value for the data type of the column. The default value will not be stored on disk. Data types that support `COMPRESS SYSTEM`

DEFAULT include all numerical type columns, fixed-length character, and fixed-length graphic string data types. This means that zeros and blanks can be compressed.

(Data) Row compression

This method compresses data rows by replacing repeating patterns that span multiple column values within a row with shorter symbol strings. Data row compression aims to achieve disk storage space savings. It can also lead to disk I/O savings. Also, more data can be cached in the buffer pool, thereby leading to increased bufferpool hit ratios. However, there is an associated cost in the form of extra CPU cycles needed to compress and decompress data. The storage savings and performance impact of data row compression are tied to the characteristics of the data within the database, the layout and tuning of the database, and application workload. Only the data on a data page or in log records is compressed.

Data row compression uses a static dictionary-based compression algorithm to compress data by row. Compressing data at the row level allows repeating patterns that span multiple column values within a row to be replaced with shorter symbol strings. In order to compress table data, the table COMPRESS attribute must be set to YES and a compression dictionary must exist for the table.

To build a compression dictionary (and subsequently compress a table), perform a classic (offline) table reorganization. A compression dictionary is also built with the following operations: INSERT, including IMPORT, LOAD INSERT and LOAD REPLACE, and also from some REDISTRIBUTE operations. All the data rows that exist in a table participate in the building of the compression dictionary. The dictionary is stored along with the table data rows in the data object portions of the table.

To decompress a table, set the table COMPRESS attribute to NO and then perform a classic (offline) table reorganization.

A data row that is inserted into a page may be compressed if the COMPRESS attribute for the table is YES and a dictionary exists. This applies to any insert row operation, including an insert through the import or LOAD operations. Compression is enabled for an entire table, however, each row is compressed individually. Therefore, a table could contain both compressed and non-compressed rows at the same time.

Data row compression is not applicable to index, long, LOB, and XML data objects.

Row compression is not compatible with table data replication support.

Row compression statistics can be generated using the RUNSTATS command and are stored in the system catalog table SYSCAT.TABLES. A compression estimation option is available with the INSPECT utility. The query optimizer includes decompression cost in its costing model.

Depending upon UPDATE activity and the positioning of update changes within a data row, there could be an increase in log consumption. For information about update logging and column ordering, see “Ordering columns to minimize update logging” on page 242.

For rows undergoing a size-increasing update, the new version of the row might not fit on the current data page. Rather, the new image of the row is stored on an overflow page. To minimize the creation of such pointer-overflow records, more free space can be added within a data page. For example, if 5% free space was used without compression, then allocate

10% free space with compression. This recommendation is especially important for data that is heavily updated.

Space compression for existing tables

By specifying `VALUE COMPRESSION` clause, the row format of existing table can be changed to allow space compression. Note that the sum of the byte counts of the columns allowing space compression might exceed the sum of the byte counts of the columns that do not allow space compression. This is acceptable as long as the sum of the byte counts does not exceed allowable row length of the table in the table space. For example, the allowable row length is 4005 bytes in a table space with 4 KB page size. If the allowable row length is exceeded, the error message `SQL0670N` is returned. The byte count formula is documented as part of the `CREATE TABLE` statement.

Similarly, by removing the `VALUE COMPRESSION` clause, table rows that previously allowed space compression can be changed to no longer allow space compression. The same condition regarding the sum of the byte counts of the columns applies; and the error message `SQL0670N` is returned as necessary.

To determine if you should consider space compression for your table, you should know that a table with the majority of values equal to the system default values, or `NULL`, would benefit from the new row format. For example, where there is an `INTEGER` column and 90% of the column has values of 0 (the default value for the data type `INTEGER`), or `NULL`, compressing this table plus this column would benefit from the new row format and save a lot of disk space.

When altering a table, you can use the `VALUE COMPRESSION` clause to specify that the table is using the space row format at the table level and possibly at the column level. You would use `ACTIVATE VALUE COMPRESSION` to specify that the table will use the space saving techniques or you would use `DEACTIVATE VALUE COMPRESSION` to specify that the table will no longer use space saving techniques for data in the table.

If you use `DEACTIVATE VALUE COMPRESSION`, this will implicitly disable any `COMPRESS SYSTEM DEFAULT` options associated with columns in that table.

After modifying the table to a new row format, all subsequent rows inserted, loaded, or updated will have the new row format. To have every row modified to the new row format, you should run a reorganization of the table or perform an update operation on existing rows before changing the row format.

Ordering columns to minimize update logging

When you define columns using the `CREATE TABLE` statement, consider the order of the columns, particularly for update-intensive workloads. Columns which are updated frequently should be grouped together, and defined towards or at the end of the table definition. This results in better performance, fewer bytes logged, and fewer log pages written, as well as a smaller active log space requirement for transactions performing a large number of updates.

The database manager does not automatically assume that columns specified in the `SET` clause of an `UPDATE` statement are changing in value. In order to limit index maintenance and the amount of the row which needs to be logged, the database compares the new column value against the old column value to determine if the

column is changing. Only the columns that are changing in value are treated as being updated. Exceptions to this UPDATE behavior occur for columns where the data is stored outside of the data row (long, LOB, ADT, and XML column types), or for fixed-length columns when the registry variable DB2ASSUMEUPDATE is enabled. For these exceptions, the column value is assumed to be changing so no comparison will be made between the new and old column value.

There are three different types of UPDATE log records.

- Full before and after row image logging. The entire before and after image of the row is logged. This is the only type of logging performed on tables enabled with DATA CAPTURE CHANGES, and results in the most number of bytes being logged for an update to a row.
- Full XOR logging. The XOR differences between the before and after row images, from the first byte that is changing until the end of the smaller row, then any residual bytes in the longer row. This results in less logged bytes than the full before and after image logging, with the number of bytes of data beyond the log record header information being the size of the largest row image.
- Partial XOR logging. The XOR differences between the before and after row images, from the first byte that is changing until the last byte that is changing. Byte positions may be first or last bytes of a column. This results in the least number of bytes being logged and the most efficient type of log record for an update to a row.

When DATA CAPTURE CHANGES is not enabled on the table, the amount of data that is logged for an update depends on:

- the proximity of the updated columns (COLNO)
- whether the updated columns are fixed in length or variable length
- whether row compression (COMPRESS YES) is enabled

When the total length of the row is not changing, even when row compression is enabled, the database manager computes and writes the optimal partial XOR log record.

When the total length of the row is changing, which is common when variable-length columns are updated and row compression is enabled, the database manager determines which byte is first to be changed and write a full XOR log record.

Data row compression

The purpose of data row compression is to achieve disk storage space savings, and it can also lead to disk I/O savings. Also, more data can be cached in the buffer pool, thereby increasing bufferpool hit ratios. Data row compression uses a static dictionary-based compression algorithm to compress data by row.

Compressing data at the row level allows repeating patterns that span multiple column values within a row to be replaced with shorter symbol strings.

Note: There is an associated cost in the form of extra CPU cycles needed to compress and decompress data. The storage savings and performance impact of data row compression are tied to the characteristics of the data within the database, the layout and tuning of the database, and the application workload. Only the data on a data page or in log records is compressed.

To compress table data, a compression dictionary must exist for the table, you must set the COMPRESS attribute of the CREATE TABLE or ALTER TABLE

statement to YES, and there needs to be sufficient data in the table. If these compression conditions exist for the table, then when you issue an INSERT statement or a LOAD INSERT, IMPORT INSERT, or REDISTRIBUTE command, data added to the table is compressed.

In Version 9.5, data row compression is automatically enabled if a table has the COMPRESS attribute set to YES and after the data compression dictionary has been created. If you created or altered a table with a COMPRESS attribute set to YES, no manual operation or database request is required on your part: that is, you do not need to perform an explicit classic (offline) table reorganization to create the data compression dictionary.

Note: If you set the COMPRESS attribute to YES and a compression dictionary exists, compression applies to any insert row operation, including an insert through an import or a load operation. Compression is enabled for an entire table; however, each row is compressed individually. Therefore, a table could contain both compressed and non-compressed rows at the same time.

To explicitly build a compression dictionary (and subsequently compress a table), perform a classic (offline) table reorganization. All of the data rows that exist in a table participate in the building of the compression dictionary. The dictionary is stored with the table data rows in the data object portions of the table.

To decompress a table, set the table COMPRESS attribute to NO, and then perform a classic (offline) table reorganization.

Restrictions

- Data row compression is not applicable to index, long, LOB, and XML data objects.
- Row compression is not compatible with table data replication support.
- You can generate row compression statistics by using the RUNSTATS command. These are stored in the system catalog table SYSCAT.TABLES. A compression estimation option, which estimates the effectiveness of row compression for a table, is available with the INSPECT utility. The query optimizer includes decompression cost in its costing model.
- Depending upon update activity and the positioning of update changes within a data row, there might be an increase in log space consumption.
- If a row is increasing in size, the new version of the row might not fit on the current data page. In that case, the new image of the row is stored on an overflow page. To minimize the creation of such pointer-overflow records, you can add more free space within a data page. For example, if 5% free space was used without compression, allocate 10% free space with compression. This recommendation is especially important for data that is heavily updated.

Optimistic locking

With Version 9.5, enhanced optimistic locking support provides a technique for SQL database applications that does not hold row locks between selecting, and updating or deleting rows.

Applications can be written to optimistically assume that unlocked rows are unlikely to change before the update or delete. If the rows do change, the updates or deletes will fail and the application's logic can handle such failures, for example, by retrying the select.

The advantage of this enhanced optimistic locking is improved concurrency, since other applications can read and write those same rows. In three-tier environments where business transactions have no correlation to database transactions, this optimistic locking technique is used, since locks cannot be maintained across business transactions.

Optimistic locking

Optimistic locking is a technique for SQL database applications that does not hold row locks between selecting and updating or deleting a row.

The application is written to optimistically assume that unlocked rows are unlikely to change before the update or delete operation. If the row does change, the update or delete will fail and the application logic handles such failures by, for example, retrying the select. One advantage of optimistic locking is improved concurrency, because other applications can read and write that row. In a three tier environment where business transactions have no correlation to database transaction, the optimistic locking technique is used, because locks cannot be maintained across the business transaction.

However, optimistic locking by values has some disadvantages:

- Can result in *false positives* without additional data server support, a condition when using optimistic locking whereby a row that is *changed* since it was selected cannot be updated without first being selected again. (This can be contrasted with *false negatives*, the condition whereby a row that is *unchanged* since it was selected cannot be updated without first being selected again.)
- Requires more retry logic in applications
- It is complicated for applications to build the UPDATE search conditions
- It is inefficient for the DB2 server to search for the target row based on values
- Data type mismatches between some client types and database types, for example, timestamps, prevent all columns from being used in the searched update

Version 9.5 adds support for easier and faster optimistic locking with no *false positives*. This support is added using the following new SQL functions, expressions, and features:

- Row Identifier (RID_BIT or RID) built-in function
- ROW CHANGE TOKEN expression
- Time-based update detection
- Implicitly hidden columns

DB2 applications can enable *optimistic locking by values* by building a searched UPDATE statement that finds the row with the exact same values that were selected. The searched UPDATE fails if the row's column values have changed.

Applications using this programming model will benefit from the enhanced optimistic locking feature. Note that applications that do not use this programming model are not considered optimistic locking applications, and they will continue to work as before.

Row Identifier (RID_BIT or RID) built-in function

This built-in function can be used in the SELECT list or predicates statement. In a predicate, for example, WHERE RID_BIT(tab)=?, the RID_BIT equals predicate is implemented as a new direct access method in order to to efficiently locate the row. Previously, so called values optimistic

locking with values was done by adding all the selected column values to the predicates and relying on some unique column combinations to qualify only a single row, with a less efficient access method.

ROW CHANGE TOKEN expression

This new expression returns a token as BIGINT. The token represents a relative point in the modification sequence of a row. An application can compare the current row change token value of a row with the row change token value that was stored when the row was last fetched to determine whether the row has changed.

Time-based update detection:

This feature is added to SQL using the ROW CHANGE TIMESTAMP expression. To support this feature, the table needs to have a new generated row change timestamp column defined to store the timestamp values. This can be added to existing tables using the ALTER TABLE statement, or the row change timestamp column can be defined when creating a new table. The row change timestamp column's existence also affects the behavior of optimistic locking in that the column is used to improve the granularity of the row change token from page level to row level, which could greatly benefit optimistic locking applications. This feature has also been added to DB2 for z/OS®.

Implicitly hidden columns:

For compatibility, this feature eases the adoption of the row change timestamp columns to existing tables and applications. Implicitly hidden columns are not externalized when implicit column lists are used. For example:

- A SELECT * against the table does not return a implicitly hidden columns in the result table
- An INSERT statement without a column list does not expect a value for implicitly hidden columns, but the column should be defined to allow nulls or have another default value.

Note: Refer to the DB2 Glossary for the definition of optimistic locking terms, such as *optimistic concurrency control*, *pessimistic locking*, *ROWID*, and *update detection*.

Optimistic locking restrictions and considerations

This topic lists optimistic locking restrictions that you need to be aware of.

- ROW CHANGE TIMESTAMP columns are not supported in the following keys, columns, and names (sqlstate 429BV is returned if used):
 - Primary keys
 - Foreign keys
 - Multidimensional clustered (MDC) columns
 - Range partition columns
 - Database hashed partitioning keys
 - DETERMINED BY constraint columns
 - Nicknames
- The RID() function is not supported in Database Partitioning Feature (DPF) configurations.
- Online or offline table reorg performed between the fetch and update operations in an optimistic locking scenario may cause the update to fail, but this should be handled by normal application retry logic.

- In Version 9.5, the `IMPLICITLY HIDDEN` attribute is restricted to only `ROW CHANGE TIMESTAMP` columns for optimistic locking.
- Inplace reorg is restricted for tables where a `ROW CHANGE TIMESTAMP` column was added to an existing table until all rows are guaranteed to have been materialized (SQL2219, reason code 13, is returned for this error). This can be accomplished with a `LOAD REPLACE` command or with a classic table reorg. This will prevent *false positives*. Tables created with the `ROW CHANGE TIMESTAMP` column have no restrictions.

Considerations for implicitly hidden columns

A column defined as `IMPLICITLY HIDDEN` is not part of the result table of a query that specifies `*` in a `SELECT` list. However, an implicitly hidden column can be explicitly referenced in a query.

If a column list is not specified on the insert, then the `VALUES` clause or the `SELECT LIST` for the insert should not include this column (in general, it must be a generated, defaultable, or nullable column).

For example, an implicitly hidden column can be referenced in the `SELECT` list, or in a predicate in a query. Additionally, an implicitly hidden column can be explicitly referenced in a `CREATE INDEX` statement, `ALTER TABLE` statement, `INSERT` statement, `MERGE` statement, or `UPDATE` statement. An implicitly hidden column can be referenced in a referential constraint. A `REFERENCES` clause that does not contain a column list refers implicitly to the primary key of the parent table. It is possible that the primary key of the parent table includes a column defined as implicitly hidden. Such a referential constraint is allowed.

- If the `SELECT` list of the fullselect of the materialized query definition explicitly refers to an implicitly hidden column, that column will be part of the materialized query table. Otherwise, an implicitly hidden column is not part of a materialized query table that refers to a table containing an implicitly hidden column.
- If the `SELECT` list of the fullselect of a view definition (`CREATE VIEW` statement) explicitly refers to an implicitly hidden column, that column will be part of the view, (however the view column is not considered to be 'hidden'). Otherwise, an implicitly hidden column is not part of a view that refers to a table containing an implicitly hidden column.

Considerations for Label Based Access Control (LBAC)

When a column is protected under LBAC, access by a user to that column is determined by the LBAC policies and the security label of the user. This protection, if applied to a row change timestamp column, extends to the reference to that column via both the `ROW CHANGE TIMESTAMP` and `ROW CHANGE TOKEN` expressions which are derived from that column.

Therefore when determining the security policies for a table, ensure that the access to the row change timestamp column is available for all users which need to use optimistic locking or time based update detection as appropriate. Note that if there is no row change timestamp column then the `ROW CHANGE TOKEN` expression cannot be blocked by LBAC. However, if the table is altered to add a row change timestamp column then any LBAC considerations will then apply.

Granularity of row change tokens and *false negatives*

The RID_BIT() built-in function and the row change token are the only requirements for optimistic locking. However, the schema of the table also affects the behavior of optimistic locking.

For example, a row change timestamp column, defined using either of the following statement clauses shown below, causes the DB2 server to store the time when a row is last changed (or initially inserted). This provides a way to capture the timestamp of the most recent change to a row. This is a timestamp column and it is maintained by the database manager, unless the GENERATED BY DEFAULT clause is used to accept a user-provided input value.

```
GENERATED ALWAYS FOR EACH ROW ON UPDATE AS ROW CHANGE TIMESTAMP  
GENERATED BY DEFAULT FOR EACH ROW ON UPDATE AS ROW CHANGE TIMESTAMP
```

Therefore, when an application uses the new ROW CHANGE TOKEN expression on a table, there are two possibilities to consider:

- *The table does not have a row change timestamp column:* A ROW CHANGE TOKEN expression returns a derived BIGINT value that is shared by all rows located on the same page. If one row on a page is updated, the row change token is changed for all the rows on the same page. This means an update can fail when changes are made to other rows, a property referred to as a false negative.

Note: Use this mode only if the application can tolerate *false negatives* and does not want to add additional storage to each row for a ROW CHANGE TIMESTAMP column.

- *The table has a row change timestamp column:* A ROW CHANGE TOKEN expression returns a BIGINT value derived from the timestamp value in the column. In this case, *false negatives* may occur but are more infrequent: If the table is reorganized or redistributed, *false negatives* can occur if the row is moved and an application uses the prior RID_BIT() value.

Time-based update detection

Some applications need to know database updates for certain time ranges, which may be used for replication of data, auditing scenarios, and so forth. The new ROW CHANGE TIMESTAMP expression provides this information.

It returns a timestamp representing the time when a row was last changed, expressed in local time similar to CURRENT TIMESTAMP. For a row that has been updated, this reflects the most recent update to the row. Otherwise, the value corresponds to the original insert of the row.

The value of the ROW CHANGE TIMESTAMP is unique for every row per table of a database or table partition, that is, not all rows per database partition are unique, only the rows in the same table. The value represents the modification sequence of the row. Rows which are more lately modified always have later values than rows which were earlier modified. Since the value is always growing from earlier to later, it may become out of sync with the system clock if:

- The system clock is changed
- The row change timestamp column is GENERATED BY DEFAULT (intended for data propagation only) and a row is provided with an out of sync value

The prerequisite for using the ROW CHANGE TIMESTAMP expression is that the table must have a row change timestamp column defined. Every row returns the timestamp of when it was inserted or last updated. There are two methods in which the row change timestamp column can be part of the table:

- *The table was created with a row change timestamp column. A ROW CHANGE TIMESTAMP expression returns the value of the column. For this category, the timestamp is precise. The row change timestamp in general when generated by the database is limited by speed of inserts and possible clock manipulations including DST adjustment.*
- *The table was not created with a row change timestamp column, but one was later added through an ALTER TABLE statement. A ROW CHANGE TIMESTAMP expression returns the value of the column. For this category, the old (pre-alter) rows do not contain the actual timestamp until they are first updated or an offline table reorganization is performed.*

Note: The timestamp is an approximate time that the actual update occurred in the database, as of the system clock at the time and taking into account the limitation that no timestamps can be repeated within a database/table partition. In practice this is normally a very accurate representation of the time of the update. The row change timestamp, in general, when generated by the database, is limited by speed of inserts and possible clock manipulations including DST adjustments.

Rows that have not been updated since the ALTER TABLE statement will return the type default value for the column, which is midnight Jan 01, year 1. Only rows that have been updated will have a unique timestamp. Rows which have the timestamp materialized via an offline table reorganization will return a unique timestamp generated during the reorganization of the table. Reorg using the INPLACE option is not sufficient as it does not materialize schema changes.

In either case, the timestamp of a row may also be updated if a redistribute is performed. If the row is moved from one database partition to another during a redistribute then a new timestamp must be generated which is guaranteed to be unique at the target.

Time values generated for ROW CHANGE TIMESTAMPS

There are some boundary conditions on the exact values generated for the row change timestamp columns due to the enforcement of unique values per partition.

Whenever the system clock is adjusted into the past for clock correction or for a daylight saving time policy on the DB2 server, it is possible that timestamps will appear to be in the future relative to the current value of the system clock, or the value of the CURRENT TIMESTAMP special register. This occurs when a timestamp was generated prior to the system clock adjustment, that is, later than the adjusted time, as the timestamps are always generated in an ascending fashion to maintain uniqueness.

When timestamps are generated for columns which were added to the table by a REORG operation or as part of a LOAD operation, the timestamps will be sequentially generated at some point in the processing of the utility starting from an initial timestamp value. If the utility is able to process rows faster than the timestamp granularity (that is, more than 1 million rows per second), then the values generated for some of the rows may also appear to be in the future relative to the system clock or the CURRENT TIMESTAMP special register.

In each case, once the system clock catches up to the row change timestamp values, there will be a close approximation of the time that the row was inserted. Until such time, timestamps will be generated in ascending sequence by the finest granularity allowed by the timestamp type.

RID_BIT() and RID() built-in function

The RID_BIT() and *row change token* can be selected for every row in a table. The SELECT can occur at any isolation level that the application requires.

The application can UPDATE the same (unchanged) row with optimistic locking by searching on:

- The RID_BIT() to directly access (not scan) the update target row
- The row change token to ensure this is the same unchanged row

This update (or delete) can occur at any point after the select, within the same unit of work, or even across connection boundaries; the only requirement is having obtained the two values above for a given row at some point in time.

Optimistic locking is used in the “WebSphere-Oriented Programming Model”. For example, Microsoft .NET uses this model to process SELECT statements followed by UPDATE or DELETE statements as follows:

- Connect to the database server and SELECT the desired rows from a table
- Disconnect from the database, or release the row locks so that other applications can read, update, delete, and insert data without any concurrency conflicts due to locks and resources held by the application (isolation “Uncommitted Read” allows higher concurrency AND assuming other applications COMMIT their update and delete transactions, then this optimistic locking application will read the updated values and the optimistic searched update/delete will succeed)
- Perform some local calculations on the SELECTed row data
- Reconnect to the database server, and search for UPDATE or DELETE on one or more particular targeted rows (and, if the target row has changed, handle failed UPDATE or DELETE statements)

Applications using this programming model will benefit from the enhanced optimistic locking feature. Note that applications that do not use this programming model are not considered optimistic locking applications, and they will continue to work as before.

RID_BIT() and RID() built-in function features

Following are the new features that will be implemented for enhanced optimistic locking and for update detection:

RID_BIT(<table designator>)

A new built-in function that returns the Record identifier (RID) of a row as VARCHAR(16) FOR BIT DATA.

Note: DB2 for z/OS implements a built-in function RID with a return type of BIGINT, but that is not large enough for Linux, UNIX, and Windows RIDs. For compatibility, this RID() built-in function returns BIGINT, in addition to RID_BIT().

This RID() built-in function does not work in DPF environments, and does not include table version information. Otherwise, it works the same as RID_BIT. You should use it only when coding applications that will be ported to z/OS servers. Except where necessary, this topic refers only to RID_BIT.

RID_BIT() built-in function

This built-in function can be used in the SELECT list or predicates statement. In a predicate, for example, WHERE RID_BIT(tab)=?, the RID_BIT

equals predicate is implemented as a new direct access method in order to efficiently locate the row. Previously, so called values *optimistic locking with values* was done by adding all the selected column values to the predicates and relying on some unique column combinations to qualify only a single row, with a less efficient access method.

ROW CHANGE TOKEN FOR <table designator>

A new expression that returns a token as BIGINT. The token represents a relative point in the modification sequence of a row. An application can compare the current row change token value of a row with the row change token value that was stored when the row was last fetched to determine whether the row has changed.

ROW CHANGE TIMESTAMP column

A GENERATED column with default type of TIMESTAMP which can be defined as either:

```
GENERATED ALWAYS FOR EACH ROW ON UPDATE  
AS ROW CHANGE TIMESTAMP
```

or (suggested only for data propagation or unload and reload operations):

```
GENERATED BY DEFAULT FOR EACH ROW ON UPDATE  
AS ROW CHANGE TIMESTAMP
```

The data in this column changes every time the row is changed. When this column is defined, the ROW CHANGE TOKEN value will be derived from it. Note that when GENERATED ALWAYS is used, the database manager ensures that this value is unique within a database partition or within table partition to ensure that no *false positives* are possible.

To use the first two elements, RID_BIT and ROW CHANGE TOKEN, no other changes are need to the database schema. Note, however, that without the ROW CHANGE TIMESTAMP column, the row change token is shared by every row on the same page. Updates to any row on the page may cause *false negatives* for other rows stored on the same page. With this column, the ROW CHANGE TOKEN is derived from the timestamp and is not shared with any other rows in the table or database partition. See “Granularity of row change tokens and *false negatives*” on page 248.

Time-based update detection feature

A new expression that returns a timestamp value that represents the time when the row in the table identified by the table designator was last changed.

```
ROW CHANGE TIMESTAMP FOR <table designator>
```

The ROW CHANGE TIMESTAMP expression is not supported for tables that do not have a ROW CHANGE TIMESTAMP column.

The ROW CHANGE TIMESTAMP expression is used for time-based update detection scenarios only, and requires that a row change timestamp column is defined for the table identified by the table designator. This column is managed by the database manager and is used to store the timestamp value that will be returned by the ROW CHANGE TIMESTAMP expression. This timestamp differs from the CURRENT TIMESTAMP in that it is guaranteed unique when assigned by the database per row per database partition. It is a local timestamp approximation of the modification time of each individual row inserted or updated.

Note: Despite the inter-relation of these two features, that is, the RID_BIT() and RID() built-in function and the time-based update detection feature, it is important to note that the usage of ROW CHANGE TOKEN and ROW CHANGE TIMESTAMP expressions are not interchangeable; specifically, that ROW CHANGE TIMESTAMP expression is not part of the optimistic locking usage.

Planning the enablement of optimistic locking

Since the new SQL expressions and attributes for optimistic locking can be used with no DDL changes to the tables involved, you can easily try optimistic locking in your test applications.

Note that without DDL changes, optimistic locking applications may get more *false negatives* than with DDL changes. An application that does get false negatives may not scale well in a production environment because the false negatives may cause too many retries. Therefore, to avoid false negatives, optimistic locking target table(s) should be either:

- Created with a ROW CHANGE TIMESTAMP column
- Altered to contain the ROW CHANGE TIMESTAMP column

If the recommended DDL changes are done, there will be many fewer false negatives. The only false negatives will occur due to table level operations like reorg, not concurrent applications operating on different rows.

In general, the database manager allows false negatives (online or offline reorg, for example) and the presence of a row change timestamp column is sufficient to determine whether page or row level granularity is being used. You can also query the SYSCAT.COLUMNS for a table that has rows with a YES in the ROWCHANGESTAMP column.

A thorough analysis of the application and database may indicate that this DDL is not required, for example, if there is one row per page, or if the update and delete operations are very infrequent and rarely, or never, on the same data page. Such analysis is the exception.

For the update timestamp detection usage, you must make changes to the DDL for the table, and possibly reorganize the table to materialize the values. If there is concern that these changes could have a negative impact on the production database, you should first prototype the changes in a test environment. For instance, the extra columns can affect the row size limitations and plan selection.

Conditions to be aware of

- You should be aware of conditions relating to the system clock and the granularity of the timestamp values. If a table has a ROW CHANGE TIMESTAMP column, after an insert or update, the new row will have a unique ROW CHANGE TIMESTAMP value in that table on that database partition.
- To ensure uniqueness, the generated timestamp of a row will always increase, regardless if the system clock is adjusted backwards or if the update or insertion of data is happening faster than timestamp granularity. Therefore, the ROW CHANGE TIMESTAMP may be in the future compared with the system time and DB2's CURRENT TIMESTAMP special register. Unless the system clock is gets completely out of sync, or the database manager is inserting or updating at more than one million rows per second, then this should normally be very close to the actual time. In contrast to the CURRENT TIMESTAMP, this value is also generated per row at the time of the update, therefore, it is normally much closer than the CURRENT TIMESTAMP, which is generated once for an entire

statement that could take a very long time to complete, depending on the complexity and number of rows affected.

Enabling optimistic locking in applications

There are a number of steps that you need to perform in order to enable optimistic locking support in your applications.

1. In the initial query, SELECT the row identifier (using the “RID_BIT() and RID() built-in function” on page 250) and row change token for each row that you need to process.
2. Release the row locks so that other applications can SELECT, INSERT, UPDATE and DELETE from the table.
3. Perform a searched UPDATE or DELETE on the target rows, using the row identifier and row change token in the search condition, optimistically assuming that the unlocked row has not changed since the original SELECT statement
4. If the row has changed, the UPDATE operation will fail and the application logic must handle the failure. For instance, the application retries the SELECT and UPDATE operations.

After running the above steps:

- If the number of retries performed by your application seems higher than expected or is desired, then adding a row change timestamp column to your table to ensure that only changes to the row identified by the RID_BIT function will invalidate only the row change token, and not other activity on the same data page.
- To see rows which have been inserted or updated in a given time range, create or alter the table to contain a row change timestamp column. This column will be maintained by the database manager automatically and can be queried using either the column name or the ROW CHANGE TIMESTAMP expression.
- For row change timestamp columns only, if the column is defined with the IMPLICITLY HIDDEN attribute, then it is not externalized when there is an implicit reference to the columns of the table. However, an implicitly hidden column can always be referenced explicitly in SQL statements. This can be useful when adding a column to a table can cause existing applications using implicit column lists to fail.

Table partitioning and data organization schemes

Table partitioning is a data organization scheme in which table data is divided across multiple data partitions according to values in one or more partitioning columns of the table. Data from a given table is partitioned into multiple storage objects, which can be in different table spaces.

For complete details about table partitioning and data organization schemes, see the *Partitioning and Clustering Guide*.

Creating tables

The database manager controls changes and access to the data stored in the tables. You can create tables using the CREATE TABLE statement. Complex statements can be used to define all the attributes and qualities of tables. However, if all the defaults are used, the statement to create a table is quite simple.

```
CREATE TABLE <table name> (<column name> <data type> <column options>,  
                             <column name> <data type> <column options>, ...)
```


The <table name> may or may not include a qualifier. The name must be unique when compared to all table, view, and alias names in the system catalog. The name must also not be SYSIBM, SYSCAT, SYSFUN, or SYSSTAT.

The <column name> names the columns in the table. This name cannot be qualified and must be unique within the other columns of the table.

Any <column options> that exist for a column further define the attributes of the column. The options include NOT NULL in order to prevent the column from containing null values, specific options for LOB data types, and the SCOPE of the reference type columns, any constraints on the columns, and any defaults for the columns. For more information, see the CREATE TABLE statement.

Declaring global temporary tables

To create global temporary tables from within your applications, use the DECLARE GLOBAL TEMPORARY TABLE statement.

Global temporary tables, also referred to as user-defined temporary tables, are used by applications that work with data in the database. Results from manipulation of the data need to be stored temporarily in a table. A user temporary table space must exist before creating global temporary tables.

Note: The description of global temporary tables does not appear in the system catalog thus making it not persistent for, and not able to be shared with, other applications. When the application using this table terminates or disconnects from the database, any data in the table is deleted and the table is implicitly dropped. Global temporary tables do not support:

- LOB-type columns (or distinct-type columns based on LOBs)
- User-defined type columns
- LONG VARCHAR columns
- XML columns

Example

```
DECLARE GLOBAL TEMPORARY TABLE gbl_temp
  LIKE emp1tabl
  ON COMMIT DELETE ROWS
  NOT LOGGED
  IN usr_tbsp
```

This statement creates a global temporary table called gbl_temp. This table is defined with columns that have exactly the same name and description as the columns of the emp1tabl. The implicit definition only includes the column name, data type, nullability characteristic, and column default value attributes. All other column attributes including unique constraints, foreign key constraints, triggers, and indexes are not defined. When a COMMIT operation is performed, all data in the table is deleted if no WITH HOLD cursor is open on the table. Changes made to the user temporary table are not logged. The global temporary table is placed in the specified user temporary table space. This table space must exist or the declaration of this table will fail.

If a ROLLBACK or ROLLBACK TO SAVEPOINT is specified when creating this table, either you can specify to delete all the rows in the table (DELETE ROWS, which is the default), or you can specify that the rows of the table are to be preserved (PRESERVE ROWS).

The table is dropped implicitly when the application disconnects from the database.

Creating tables like existing tables

Creating a new source table might be necessary when the characteristics of the target table do not sufficiently match the characteristics of the source when issuing the ALTER TABLE statement with the ATTACH PARTITION clause. Before creating a new source table, you can attempt to correct the mismatch between the existing source table and the target table.

To create a table, the privileges held by the authorization ID of the statement must include at least one of the following authorities and privileges:

- CREATETAB authority on the database and USE privilege on the table space, as well as one of:
 - IMPLICIT_SCHEMA authority on the database, if the implicit or explicit schema name of the table does not exist
 - CREATEIN privilege on the schema, if the schema name of the table refers to an existing schema
- SYSADM or DBADM authority

If attempts to correct the mismatch fail, error SQL20408N or SQL20307N is returned.

To create a new source table:

1. Use the db2look command to produce the CREATE TABLE statement to create a table identical to the target table:

```
db2look -d <source database name> -t <target database name> -e -p
```

2. Remove the partitioning clause from the db2look output and change the name of the table created to a new name (in this example, referred to here as sourceC).
3. Next, load all of the data from the original source table to the newly created source table, sourceC using a LOAD FROM CURSOR command:

```
DECLARE mycurs CURSOR FOR SELECT * FROM source
LOAD FROM mycurs OF CURSOR REPLACE INTO sourceC
```

If this command fails because the original data is incompatible with the definition of table sourceC, you must transform the data in the original table as it is being transferred to sourceC.

4. After the data has been successfully copied to sourceC, submit the ALTER TABLE target ...ATTACH sourceC statement.

Creating tables for staging data

A *staging table* allows incremental maintenance support for deferred materialized query table. The staging table collects changes that need to be applied to the materialized query table to synchronize it with the contents of underlying tables. The use of staging tables eliminates the high lock contention caused by immediate maintenance content when an immediate refresh of the materialized query table is requested. Also, the materialized query tables no longer need to be entirely regenerated whenever a REFRESH TABLE is performed.

Materialized query tables are a powerful way to improve response time for complex queries, especially queries that might require some of the following operations:

- Aggregated data over one or more dimensions
- Joins and aggregates data over a group of tables
- Data from a commonly accessed subset of data
- Repartitioned data from a table, or part of a table, in a partitioned database environment

Here are some of the key restrictions regarding staging tables:

1. The query used to define the materialized query table, for which the staging table is created, must be incrementally maintainable; that is, it must adhere to the same rules as a materialized query table with an immediate refresh option.
2. Only a deferred refresh can have a supporting staging table. The query also defines the materialized query table associated with the staging table. The materialized query table must be defined with REFRESH DEFERRED.
3. When refreshing using the staging tables, only a refresh to the current point in time is supported.
4. Partitioned hierarchy tables and partitioned typed tables are not supported. (Partitioned tables are tables where data is partitioned into multiple storage objects based on the specifications provided in the PARTITION BY clause of the CREATE TABLE statement.)

An inconsistent, incomplete, or pending state staging table cannot be used to incrementally refresh the associated materialized query table unless some other operations occur. These operations will make the content of the staging table consistent with its associated materialized query table and its underlying tables, and to bring the staging table out of pending. Following a refresh of a materialized query table, the content of its staging table is cleared and the staging table is set to a normal state. A staging table might also be pruned intentionally by using the SET INTEGRITY statement with the appropriate options. Pruning will change the staging table to an inconsistent state. For example, the following statement forces the pruning of a staging table called STAGTAB1:

```
SET INTEGRITY FOR STAGTAB1 PRUNE;
```

When a staging table is created, it is put in a pending state and has an indicator that shows that the table is inconsistent or incomplete with regard to the content of underlying tables and the associated materialized query table. The staging table needs to be brought out of the pending and inconsistent state in order to start collecting the changes from its underlying tables. While in a pending state, any attempts to make modifications to any of the staging table's underlying tables will fail, as will any attempts to refresh the associated materialized query table.

There are several ways a staging table might be brought out of a pending state; for example:

- SET INTEGRITY FOR <staging table name> STAGING IMMEDIATE UNCHECKED
- SET INTEGRITY FOR <staging table name> IMMEDIATE CHECKED

Modifying tables

This section provides topics on how you can modify tables.

Altering materialized query table properties

With some restrictions, you can change a materialized query table to a regular table or a regular table to a materialized query table. You cannot change other table types; only regular and materialized query tables can be changed. For example, you cannot change a replicated materialized query table to a regular table, nor the reverse.

Once a regular table has been altered to a materialized query table, the table is placed in a set integrity pending state. When altering in this way, the `fullselect` in the materialized query table definition must match the original table definition, that is:

- The number of columns must be the same.
- The column names and positions must match.
- The data types must be identical.

If the materialized query table is defined on an original table, then the original table cannot itself be altered into a materialized query table. If the original table has triggers, check constraints, referential constraints, or a defined unique index, then it cannot be altered into a materialized query table. If altering the table properties to define a materialized query table, you are not allowed to alter the table in any other way in the same `ALTER TABLE` statement.

When altering a regular table into a materialized query table, the `fullselect` of the materialized query table definition cannot reference the original table directly or indirectly through views, aliases, or materialized query tables.

To change a materialized query table to a regular table, use the following:

```
ALTER TABLE sumtable
SET SUMMARY AS DEFINITION ONLY
```

To change a regular table to a materialized query table, use the following:

```
ALTER TABLE regtable
SET SUMMARY AS <fullselect>
```

The restrictions on the `fullselect` when altering the regular table to a materialized query table are very much like the restrictions when creating a summary table using the `CREATE SUMMARY TABLE` statement.

Refreshing the data in a materialized query table

You can refresh the data in one or more materialized query tables by using the `REFRESH TABLE` statement. The statement can be embedded in an application program, or issued dynamically. To use this statement, you must have either `SYSADM` or `DBADM` authority, or `CONTROL` privilege on the table to be refreshed.

The following example shows how to refresh the data in a materialized query table:

```
REFRESH TABLE SUMTAB1
```

Changing column properties

Use the `ALTER TABLE` statement to change column properties, such as nullability, LOB options, scope, constraints and compression attributes, data types and so forth. For complete details, see the `ALTER TABLE` statement.

To alter a table, you must have at least one of the following privileges on the table to be altered:

- ALTER privilege
- CONTROL privilege
- SYSADM or DBADM authority
- ALTERIN privilege on the schema of the table

To change the definition of a existing column, to edit and test SQL when changing table columns, or to validate related objects when changing table columns, you must have DBADM authority.

For example, from the command line, enter:

```
ALTER TABLE EMPLOYEE
ALTER COLUMN WORKDEPT
SET DEFAULT '123'
```

Adding and dropping columns

To add columns to existing tables, or to drop columns from existing tables, use the ALTER TABLE statement with the ADD COLUMN, or DROP COLUMN, clause, respectively. The table must not be a typed table.

For all existing rows in the table, the value of the new column is set to its default value. The new column is the last column of the table; that is, if initially there are n columns, the added column is column $n+1$. Adding the new column must not make the total byte count of all columns exceed the row size limit.

To add a column, issue the following statement:

```
ALTER TABLE SALES
ADD COLUMN SOLD_QTY
SMALLINT NOT NULL DEFAULT 0
```

To delete or drop a column, issue the following statement:

```
ALTER TABLE SALES
DROP COLUMN SOLD_QTY
```

Modifying DEFAULT clause column definitions

The DEFAULT clause provides a default value for a column in the event that a value is not supplied on INSERT or is specified as DEFAULT on INSERT or UPDATE. If a specific default value is not specified following the DEFAULT keyword, the default value depends on the data type. If a column is defined as an XML or structured type, then a DEFAULT clause cannot be specified.

Omission of DEFAULT from a column-definition results in the use of the null value as the default for the column, as described in: “Default column and data type definitions” on page 235.

Specific types of values that can be specified with the DEFAULT keyword, see the ALTER TABLE statement.

Modifying the generated or identity property of a column

You can add and drop the generated or identity property of a column in a table using the ALTER COLUMN clause in the ALTER TABLE statement.

You can do one of the following actions:

- When working with an existing non-generated column, you can add a generated expression attribute. The modified column then becomes a generated column.
- When working with an existing generated column, you can drop a generated expression attribute. The modified column then becomes a normal, non-generated column.
- When working with an existing non-identity column, you can add a identity attribute. The modified column then becomes an identity column.
- When working with an existing identity column, you can drop the identity attribute. The modified column then becomes a normal, non-generated, non-identity column.
- When working with an existing generated column, you can alter a generated column from being GENERATED ALWAYS to GENERATED BY DEFAULT. The reverse is also true; that is, you can alter a generated column from being GENERATED BY DEFAULT to GENERATED ALWAYS. This is only possible when working with a generated column.
- You can drop the default attribute from the user-defined default column. When you do this, the new default value is null.
- You can drop the default, identity, or generation attribute and then set a new default, identity, or generation attribute in the same ALTER COLUMN statement.
- For both the CREATE TABLE and ALTER TABLE statements, the “ALWAYS” is an optional word in the GENERATED clause. This means that GENERATED ALWAYS is equivalent to GENERATED when used in the ALTER TABLE statement.

Modifying column definitions

Use the ALTER TABLE statement to drop columns, or change their types and attributes. For example, you can increase the length of an existing VARCHAR or VARCHAR column. The number of characters might increase up to a value dependent on the page size used.

To modify the default value associated with a column, once you have defined the new default value, the new value is used for the column in any subsequent SQL operations where the use of the default is indicated. The new value must follow the rules for assignment and have the same restrictions as documented under the CREATE TABLE statement.

Note: Generate columns cannot have their default value altered by this statement.

When changing these table attributes using SQL, it is no longer necessary to drop the table and then recreate it, a time consuming process that can be complex when object dependencies exist.

To modify the length and type of a column of an existing table using the command line, enter:

```
ALTER TABLE <table_name>
  ALTER COLUMN <column_name>
  <modification_type>
```

For example, to increase a column up to 4000 characters, use something similar to the following:

```
ALTER TABLE t1
  ALTER COLUMN colnam1
  SET DATA TYPE VARCHAR(4000)
```

In another example, to allow a column to have a new VARGRAPHIC value, use an statement similar to the following:

```
ALTER TABLE t1
  ALTER COLUMN colnam2
  SET DATA TYPE VARGRAPHIC(2000)
```

You cannot alter the column of a typed table. However, you can add a scope to an existing reference type column that does not already have a scope defined. For example:

```
ALTER TABLE t1
  ALTER COLUMN colnamt1
  ADD SCOPE tytab1
```

To modify the default value of a column of an existing table using the command line, enter:

```
ALTER TABLE <table_name>
  ALTER COLUMN <column_name>
  SET DEFAULT 'new_default_value'
```

For example, to change the default value for a column, use something similar to the following:

```
ALTER TABLE t1
  ALTER COLUMN colnam1
  SET DEFAULT '123'
```

Renaming tables

You can use the RENAME statement to rename an existing table.

When renaming tables, the source table must not be referenced in any existing definitions (view or materialized query table), triggers, SQL functions, or constraints. It must also not have any generated columns (other than identity columns), or be a parent or dependent table. Catalog entries are updated to reflect the new table name. For more information and examples, see the RENAME statement.

To change the definition of existing columns, see “Changing column properties” on page 257 and the ALTER TABLE statement.

Recovering inoperative summary tables

Summary tables can become *inoperative* as a result of a revoked SELECT privilege on an underlying table.

The following steps can help you recover an inoperative summary table:

- Determine the statement that was initially used to create the summary table. You can obtain this information from the TEXT column of the SYSCAT.VIEW catalog view.
- Re-create the summary table by using the CREATE SUMMARY TABLE statement with the same summary table name and same definition.
- Use the GRANT statement to re-grant all privileges that were previously granted on the summary table. (Note that all privileges granted on the inoperative summary table are revoked.)

If you do not want to recover an inoperative summary table, you can explicitly drop it with the `DROP TABLE` statement, or you can create a new summary table with the same name but a different definition.

An inoperative summary table only has entries in the `SYSCAT.TABLES` and `SYSCAT.VIEWS` catalog views; all entries in the `SYSCAT.TABDEP`, `SYSCAT.TABAUTH`, `SYSCAT.COLUMNS` and `SYSCAT.COLAUTH` catalog views are removed.

Viewing table definitions

You can use the `SYSCAT.COLUMNS` catalog view, to view table definitions. Each row represents a column defined for a table, view, or nickname. To see the data in the columns, use the `SELECT` statement.

Table or view aliases

An *alias* is an alternative name for a table or a view. It can be used to reference a table or a view if an existing table or view *can* be referenced.

An alias cannot be used in all contexts; for example, it cannot be used in the check condition of a check constraint. An alias cannot reference a declared temporary table.

Like tables or views, an alias can be created, dropped, and have comments associated with it. However, unlike tables, aliases can refer to each other in a process called *chaining*. Aliases are publicly referenced names, so no special authority or privilege is required to use them. Access to the table or the view referred to by an alias, however, does require the authorization associated with these objects.

There are other types of aliases, such as database and network aliases. Aliases can also be created for *nicknames* that refer to data tables or views located on federated systems.

Dropping tables

A table can be dropped with a `DROP TABLE` statement. When a table is dropped, the row in the `SYSCAT.TABLES` system catalog view that contains information about that table is dropped, and any other objects that depend on the table are affected.

For example:

- All column names are dropped.
- Indexes created on any columns of the table are dropped.
- All views based on the table are marked inoperative.
- All privileges on the dropped table and dependent views are implicitly revoked.
- All referential constraints in which the table is a parent or dependent are dropped.
- All packages and cached dynamic SQL and XQuery statements dependent on the dropped table are marked invalid, and remain so until the dependent objects are re-created. This includes packages dependent on any supertable above the subtable in the hierarchy that is being dropped.
- Any reference columns for which the dropped table is defined as the scope of the reference become “unscoped”.

- An alias definition on the table is not affected, because an alias can be undefined
- All triggers dependent on the dropped table are marked inoperative.

To drop a table using the command line, enter:

```
DROP TABLE <table_name>
```

The following statement drops the table called DEPARTMENT:

```
DROP TABLE DEPARTMENT
```

An individual table cannot be dropped if it has a subtable. However, all the tables in a table hierarchy can be dropped by a single DROP TABLE HIERARCHY statement, as in the following example:

```
DROP TABLE HIERARCHY person
```

The DROP TABLE HIERARCHY statement must name the root table of the hierarchy to be dropped.

There are differences when dropping a table hierarchy compared to dropping a specific table:

- DROP TABLE HIERARCHY does not activate deletion-triggers that would be activated by individual DROP table statements. For example, dropping an individual subtable would activate deletion-triggers on its supertables.
- DROP TABLE HIERARCHY does not make log entries for the individual rows of the dropped tables. Instead, the dropping of the hierarchy is logged as a single event.

Dropping materialized query or staging tables

You cannot alter a materialized query or staging table, but you can drop it. All indexes, primary keys, foreign keys, and check constraints referencing the table are dropped. All views and triggers that reference the table are made inoperative. All packages depending on any object dropped or marked inoperative will be invalidated.

To drop a materialized query or staging table using the command line, enter:

```
DROP TABLE <table_name>
```

The following statement drops the materialized query table XT:

```
DROP TABLE XT
```

A materialized query table might be explicitly dropped with the DROP TABLE statement, or it might be dropped implicitly if any of the underlying tables are dropped.

A staging table might be explicitly dropped with the DROP TABLE statement, or it might be dropped implicitly when its associated materialized query table is dropped.

Scenarios and examples of tables

This section provides scenarios and examples of tables.

Scenarios: Optimistic locking and time-based detection

Three scenarios are provided that show you how to enable and implement optimistic locking in your applications, with and without time-based detection, and with and without implicitly hidden columns.

Scenario: Using optimistic locking in an application program

This scenario demonstrates how optimistic locking is implemented in an application program, covering six different scenarios.

Consider the following sequence of events in an application designed and enabled for optimistic locking:

```
SELECT QUANTITY, row change token FOR STOCK, RID_BIT(STOCK)
INTO :h_quantity, :h_rct, :h_rid
FROM STOCK WHERE PARTNUM = 3500
```

In this scenario, the application logic reads each row. Since this application is enabled for optimistic locking as described in “Enabling optimistic locking in applications” on page 253, the select list includes the RID_BIT() value saved in the :h_rid host variable and the row change token value saved in the :h_rct host variable.

With optimistic locking enabled, the application optimistically assumes any rows targeted for update or delete will remain unchanged, even if they are unprotected by locks. To improve database concurrency, the application removes the row lock(s) using one of the following methods:

- Committing the unit of work, in which case the row locks are removed
- Closing the cursor using the WITH RELEASE clause, in which case the row locks are removed
- Using a lower isolation level:
 - CURSOR STABILITY (CS) in which case the row is not locked after the cursor fetches to the next row, or to the end of the result table.
 - UNCOMMITTED READ (UR) in which case any uncommitted data has a new (uncommitted) row change token value. If the uncommitted data is rolled back, then the old committed row change token will be a different value.

Note: Assuming updates are not normally rolled back, using UR allows the most concurrency.

- Disconnecting from the database, thus releasing all DB2 server resources for the application. (.NET applications often use this mode).

The application processes the rows and decides it wants to optimistically update one of them:

```
UPDATE STOCK SET QUANTITY = QUANTITY - 1
WHERE row change token FOR STOCK = :h_rct AND
RID_BIT(STOCK) = :h_rid
```

The UPDATE statement updates the row identified in the SELECT statement shown above.

The searched UPDATE predicate is planned as a direct fetch to the table:

```
RID_BIT(STOCK) = :h_rid
```

Direct fetch is a very efficient access plan, that is simple for the DB2 optimizer to cost. If the RID_BIT() predicate does not find a row, the row was deleted and the update fails with row not found.

Assuming that the RID_BIT() predicate finds a row, the predicate row change token FOR STOCK = :h_rct will find the row if the row change token has not changed. If the row change token has changed since the SELECT, the searched UPDATE fails with row not found.

Table 50 lists the possible scenarios that could occur when optimistic locking is enabled.

Table 50. Scenarios that could occur when optimistic locking is enabled

Scenario ID	Action	Result
Scenario 1	There is a row change timestamp column defined on the table and no other application has changed the row.	The update succeeds as the row change token predicate succeeds for the row identified by :h_rid.
Scenario 2	There is a ROW CHANGE TIMESTAMP defined on the table. Another application updates the row after the select and before the update (and commits), updating the row change timestamp column.	The row change token predicate fails comparing the token generated from the timestamp in the row at the time of the select and the token value of the timestamp currently in the row. So the UPDATE statement fails to find a row.
Scenario 3	There is a ROW CHANGE TIMESTAMP defined on the table. Another application updates the row and so the row has a new row change token. This application selects the row at isolation UR and gets the new uncommitted row change token.	This application runs the UPDATE, which will lock wait until the other application releases its row lock. The row change token predicate will succeed if the other application commits the change with the new token, so the UPDATE succeeds. The row change token predicate will fail if the other application rolls back to the old token, so the UPDATE fails to find a row.
Scenario 4	There is no row change timestamp column defined on the table. Another row is updated, deleted or inserted on the same page, after the select and before the update.	The row change token predicate fails comparing the token because the row change token value for all rows on the page has changed, so the UPDATE statement fails to find a row even though our row has not actually changed. This false negative scenario would not result in an UPDATE failure if a row change timestamp column was added.
Scenario 5	The table has been altered to contain a row change timestamp column, and the row returned in the select has not been modified since the time of the alter. Another application updates the row, adding the row change timestamp column to that row in the process with the current timestamp.	The row change token predicate fails comparing the token generated from before with the token value created from the row change timestamp column so the UPDATE statement fails to find a row. Since the row of interest has actually been changed this is not a false negative scenario.

Table 50. Scenarios that could occur when optimistic locking is enabled (continued)

Scenario ID	Action	Result
Scenario 6	The table is reorganized after the select and before the update. The row ID identified by :h_rid does not find a row, or contains a row with a different token so the update fails. This is the form of false negative that cannot be avoided even with the existence of a row change timestamp column in the row.	The row itself is not updated by the reorganization but the RID_BIT portion of the predicate cannot identify the original row after the reorganization.

Scenarios: Optimistic locking using implicitly hidden columns

The following scenarios demonstrate how optimistic locking is implemented in an application program using implicitly hidden columns, that is, a column defined with the IMPLICITLY HIDDEN attribute.

For these scenarios, assume that table SALARY_INFO is defined with three columns, and the first column is an implicitly hidden ROW CHANGE TIMESTAMP column whose values are always generated.

Scenario 1:

In the following statement, the implicitly hidden column is explicitly referenced in the column list and a value is provided for it in the VALUES clause:

```
INSERT INTO SALARY_INFO (UPDATE_TIME, LEVEL, SALARY)
VALUES (DEFAULT, 2, 30000)
```

Scenario 2:

The following INSERT statement uses an implicit column list. An implicit column list does not include implicitly hidden columns, therefore, the VALUES clause only contains values for the other two columns:

```
INSERT INTO SALARY_INFO
VALUES (2, 30000)
```

In this case, column UPDATE_TIME must be defined to have a default value, and that default value is used for the row that is inserted.

Scenario 3:

In the following statement, the implicitly hidden column is explicitly referenced in the select list and a value for it appears in the result set:

```
SELECT UPDATE_TIME, LEVEL, SALARY FROM SALARY_INFO
WHERE LEVEL = 2
```

```
UPDATE_TIME      LEVEL    SALARY
-----
2006-11-28-10.43.27.560841      2      30000
```

Scenario 4:

In the following statement the column list is generated implicitly through use of the * notation, and the implicitly hidden column does not appear in the result set:

```
SELECT * FROM SALARY_INFO
WHERE LEVEL = 2
```

```
LEVEL    SALARY
-----
2        30000
```

Scenario 5:

In the following statement, the column list is generated implicitly through use of the * notation, and the implicitly hidden column value also appears by using the ROW CHANGE TIMESTAMP FOR expression:

```
SELECT ROW CHANGE TIMESTAMP FOR SALARY_INFO
       AS ROW_CHANGE_STAMP, SALARY_INFO.*
FROM SALARY_INFO WHERE LEVEL = 2
```

The result table will be similar to scenario 3 (column UPDATE_TIME will be ROW_CHANGE_STAMP).

Scenario: Time-based update detection

This scenario demonstrates how optimistic locking is implemented in an application program using update detection by timestamp, covering three different scenarios.

In this scenario, the application selects all rows that have changed in the last 30 days.

```
SELECT * FROM TAB WHERE
       ROW CHANGE TIMESTAMP FOR TAB <=
       CURRENT_TIMESTAMP AND
       ROW CHANGE TIMESTAMP FOR TAB >=
       CURRENT_TIMESTAMP - 30 days;
```

Scenario 1:

No row change timestamp column is defined on the table. Statement fails with SQL20431N. This SQL expression is only supported for tables with a row change timestamp column defined.

Note: This scenario will work on z/OS.

Scenario 2:

A row change timestamp column was defined when the table was created:

```
CREATE TABLE TAB ( ..., RCT TIMESTAMP NOT NULL
                    GENERATED ALWAYS
                    FOR EACH ROW ON UPDATE AS
                    ROW CHANGE TIMESTAMP)
```

This statement returns all rows inserted or updated in the last 30 days.

Scenario 3:

A row change timestamp column was added to the table using the ALTER TABLE statement at some point in the last 30 days:

```
ALTER TABLE TAB ADD COLUMN RCT TIMESTAMP NOT NULL
                    GENERATED ALWAYS
                    FOR EACH ROW ON UPDATE AS
                    ROW CHANGE TIMESTAMP
```

This statement returns all the rows in the table. Any rows that have not been modified since the ALTER TABLE statement will use the default value of the timestamp of the ALTER TABLE statement itself, and all other rows that have been modified since then will have a unique timestamp.

Chapter 12. Constraints

Within any business, data must often adhere to certain restrictions or rules. For example, an employee number must be unique. The database manager provides *constraints* as a way to enforce such rules.

The following types of constraints are available:

- NOT NULL constraints
- Unique (or unique key) constraints
- Primary key constraints
- Foreign key (or referential integrity) constraints
- (Table) Check constraints
- Informational constraints

Constraints are only associated with tables and are either defined as part of the table creation process (using the CREATE TABLE statement) or are added to a table's definition after the table has been created (using the ALTER TABLE statement). You can use the ALTER TABLE statement to modify constraints. In most cases, existing constraints can be dropped at any time; this action does not affect the table's structure or the data stored in it.

Note: Unique and primary constraints are only associated with table objects, they are often enforced through the use of one or more unique or primary key indexes.

Types of constraints

A *constraint* is a rule that is used for optimization purposes.

There are five types of constraints:

- A *NOT NULL constraint* is a rule that prevents null values from being entered into one or more columns within a table.
- A *unique constraint* (also referred to as a *unique key constraint*) is a rule that forbids duplicate values in one or more columns within a table. Unique and primary keys are the supported unique constraints. For example, a unique constraint can be defined on the supplier identifier in the supplier table to ensure that the same supplier identifier is not given to two suppliers.
- A *primary key constraint* is a column or combination of columns that has the same properties as a unique constraint. You can use a primary key and foreign key constraints to define relationships between tables.
- A *foreign key constraint* (also referred to as a *referential constraint* or a *referential integrity constraint*) is a logical rule about values in one or more columns in one or more tables. For example, a set of tables shares information about a corporation's suppliers. Occasionally, a supplier's name changes. You can define a referential constraint stating that the ID of the supplier in a table must match a supplier ID in the supplier information. This constraint prevents insert, update, or delete operations that would otherwise result in missing supplier information.
- A *(table) check constraint* (simply called a *check constraint*) sets restrictions on data added to a specific table. For example, a table check constraint can ensure that the salary level for an employee is at least \$20,000 whenever salary data is added or updated in a table containing personnel information.

An *informational constraint* is an attribute of a certain type of constraint, but one that is not enforced by the database manager.

NOT NULL constraints

NOT NULL constraints prevent null values from being entered into a column.

The null value is used in databases to represent an unknown state. By default, all of the built-in data types provided with the database manager support the presence of null values. However, some business rules might dictate that a value must always be provided (for example, every employee is required to provide emergency contact information). The NOT NULL constraint is used to ensure that a given column of a table is never assigned the null value. Once a NOT NULL constraint has been defined for a particular column, any insert or update operation that attempts to place a null value in that column will fail.

Because constraints only apply to a particular table, they are usually defined along with a table's attributes, during the table creation process. The following CREATE TABLE statement shows how the NOT NULL constraint would be defined for a particular column:

```
CREATE TABLE EMPLOYEES (. . .
                        EMERGENCY_PHONE CHAR(14) NOT NULL,
                        . . .
                        );
```

Unique constraints

Unique constraints ensure that the values in a set of columns are unique and not null for all rows in the table. The columns specified in a unique constraint must be defined as NOT NULL. The database manager uses a unique index to enforce the uniqueness of the key during changes to the columns of the unique constraint.

Unique constraints can be defined in the CREATE TABLE or ALTER TABLE statement using the UNIQUE clause. For example, a typical unique constraint in a DEPARTMENT table might be that the department number is unique and not null.

Figure 21 shows that a duplicate record is prevented from being added to a table when a unique constraint exists for the table:

Department number	
001	
002	
003	
004	
005	

Invalid record

003	
-----	--

Figure 21. Unique constraints prevent duplicate data

The database manager enforces the constraint during insert and update operations, ensuring data integrity.

A table can have an arbitrary number of unique constraints, with at most one unique constraint defined as the primary key. A table cannot have more than one unique constraint on the same set of columns.

A unique constraint that is referenced by the foreign key of a referential constraint is called the *parent key*.

- When a unique constraint is defined in a CREATE TABLE statement, a unique index is automatically created by the database manager and designated as a primary or unique system-required index.
- When a unique constraint is defined in an ALTER TABLE statement and an index exists on the same columns, that index is designated as unique and system-required. If such an index does not exist, the unique index is automatically created by the database manager and designated as a primary or unique system-required index.

Note: There is a distinction between defining a unique constraint and creating a unique index. Although both enforce uniqueness, a unique index allows nullable columns and generally cannot be used as a parent key.

Primary key constraints

You can use primary key and foreign key constraints to define relationships between tables.

A primary key is a column or combination of columns that has the same properties as a unique constraint. Because the primary key is used to identify a row in a table, it must be unique, and must have the NOT NULL attribute. A table cannot have more than one primary key, but it can have multiple unique keys. Primary keys are optional, and can be defined when a table is created or altered. They are also beneficial, because they order the data when data is exported or reorganized.

(Table) Check constraints

A *check constraint* (also referred to as a *table check constraint*) is a database rule that specifies the values allowed in one or more columns of every row of a table. Specifying check constraints is done through a restricted form of a search condition.

Foreign key (referential) constraints

Foreign key constraints (also known as *referential constraints* or *referential integrity constraints*) enable you to define required relationships between and within tables.

For example, a typical foreign key constraint might state that every employee in the EMPLOYEE table must be a member of an existing department, as defined in the DEPARTMENT table.

Referential integrity is the state of a database in which all values of all foreign keys are valid. A *foreign key* is a column or a set of columns in a table whose values are required to match at least one primary key or unique key value of a row in its parent table. A *referential constraint* is the rule that the values of the foreign key are valid only if one of the following conditions is true:

- They appear as values of a parent key.
- Some component of the foreign key is null.

To establish this relationship, you would define the department number in the EMPLOYEE table as the foreign key, and the department number in the DEPARTMENT table as the primary key.

Figure 22 shows how a record with an invalid key is prevented from being added to a table when a foreign key constraint exists between two tables:




Figure 22. Foreign and primary key constraints

The table containing the parent key is called the *parent table* of the referential constraint, and the table containing the foreign key is said to be a *dependent* of that table.

Referential constraints can be defined in the CREATE TABLE statement or the ALTER TABLE statement. Referential constraints are enforced by the database manager during the execution of INSERT, UPDATE, DELETE, ALTER TABLE, MERGE, ADD CONSTRAINT, and SET INTEGRITY statements.

Referential integrity rules involve the following terms:

Table 51. Referential integrity terms

Concept	Terms
Parent key	A primary key or a unique key of a referential constraint.
Parent row	A row that has at least one dependent row.
Parent table	A table that contains the parent key of a referential constraint. A table can be a parent in an arbitrary number of referential constraints. A table that is the parent in a referential constraint can also be the dependent in a referential constraint.
Dependent table	A table that contains at least one referential constraint in its definition. A table can be a dependent in an arbitrary number of referential constraints. A table that is the dependent in a referential constraint can also be the parent in a referential constraint.
Descendent table	A table is a descendent of table T if it is a dependent of T or a descendent of a dependent of T.
Dependent row	A row that has at least one parent row.
Descendent row	A row is a descendent of row r if it is a dependent of r or a descendent of a dependent of r.
Referential cycle	A set of referential constraints such that each table in the set is a descendent of itself.
Self-referencing table	A table that is a parent and a dependent in the same referential constraint. The constraint is called a <i>self-referencing constraint</i> .
Self-referencing row	A row that is a parent of itself.

The purpose of a referential constraint is to guarantee that table relationships are maintained and that data entry rules are followed. This means that as long as a referential constraint is in effect, the database manager guarantees that for each row in a child table that has a non-null value in its foreign key columns, a row exists in a corresponding parent table that has a matching value in its parent key.

When an SQL operation attempts to change data in such a way that referential integrity will be compromised, a foreign key (or referential) constraint could be violated. The database manager handles these types of situations by enforcing a set of rules that are associated with each referential constraint. This set of rules consist of:

- An insert rule
- An update rule
- A delete rule

When an SQL operation attempts to change data in such a way that referential integrity will be compromised, a referential constraint could be violated. For example,

- An insert operation could attempt to add a row of data to a child table that has a value in its foreign key columns that does not match a value in the corresponding parent table's parent key.
- An update operation could attempt to change the value in a child table's foreign key columns to a value that has no matching value in the corresponding parent table's parent key.

- An update operation could attempt to change the value in a parent table's parent key to a value that does not have a matching value in a child table's foreign key columns.
- A delete operation could attempt to remove a record from a parent table that has a matching value in a child table's foreign key columns.

The database manager handles these types of situations by enforcing a set of rules that are associated with each referential constraint. This set of rules consists of:

- An insert rule
- An update rule
- A delete rule

Insert rule

The insert rule of a referential constraint is that a non-null insert value of the foreign key must match some value of the parent key of the parent table. The value of a composite foreign key is null if any component of the value is null. This rule is implicit when a foreign key is specified.

Update rule

The update rule of a referential constraint is specified when the referential constraint is defined. The choices are `NO ACTION` and `RESTRICT`. The update rule applies when a row of the parent or a row of the dependent table is updated.

In the case of a parent row, when a value in a column of the parent key is updated, the following rules apply:

- If any row in the dependent table matches the original value of the key, the update is rejected when the update rule is `RESTRICT`.
- If any row in the dependent table does not have a corresponding parent key when the update statement is completed (excluding `AFTER` triggers), the update is rejected when the update rule is `NO ACTION`.

The value of the parent unique keys cannot be changed if the update rule is `RESTRICT` and there are one or more dependent rows. However, if the update rule is `NO ACTION`, parent unique keys can be updated as long as every child has a parent key by the time the update statement completes. A non-null update value of a foreign key must be equal to a value of the primary key of the parent table of the relationship.

Also, the use of `NO ACTION` or `RESTRICT` as update rules for referential constraints determines when the constraint is enforced. An update rule of `RESTRICT` is enforced before all other constraints, including those referential constraints with modifying rules such as `CASCADE` or `SET NULL`. An update rule of `NO ACTION` is enforced after other referential constraints. Note that the `SQLSTATE` returned is different depending on whether the update rule is `RESTRICT` or `NO ACTION`.

In the case of a dependent row, the `NO ACTION` update rule is implicit when a foreign key is specified. `NO ACTION` means that a non-null update value of a foreign key must match some value of the parent key of the parent table when the update statement is completed.

The value of a composite foreign key is null if any component of the value is null.

Delete rule

The delete rule of a referential constraint is specified when the referential constraint is defined. The choices are NO ACTION, RESTRICT, CASCADE, or SET NULL. SET NULL can be specified only if some column of the foreign key allows null values.

If the identified table or the base table of the identified view is a parent, the rows selected for delete must not have any dependents in a relationship with a delete rule of RESTRICT, and the DELETE must not cascade to descendent rows that have dependents in a relationship with a delete rule of RESTRICT.

If the delete operation is not prevented by a RESTRICT delete rule, the selected rows are deleted. Any rows that are dependents of the selected rows are also affected:

- The nullable columns of the foreign keys of any rows that are their dependents in a relationship with a delete rule of SET NULL are set to the null value.
- Any rows that are their dependents in a relationship with a delete rule of CASCADE are also deleted, and the above rules apply, in turn, to those rows.

The delete rule of NO ACTION is checked to enforce that any non-null foreign key refers to an existing parent row after the other referential constraints have been enforced.

The delete rule of a referential constraint applies only when *a row* of the parent table is deleted. More precisely, the rule applies only when *a row* of the parent table is the object of a delete or propagated delete operation (defined below), and that row has dependents in the dependent table of the referential constraint. Consider an example where P is the parent table, D is the dependent table, and p is a parent row that is the object of a delete or propagated delete operation. The delete rule works as follows:

- With RESTRICT or NO ACTION, an error occurs and no rows are deleted.
- With CASCADE, the delete operation is propagated to the dependents of p in table D.
- With SET NULL, each nullable column of the foreign key of each dependent of p in table D is set to null.

Any table that can be involved in a delete operation on P is said to be *delete-connected* to P. Thus, a table is delete-connected to table P if it is a dependent of P, or a dependent of a table to which delete operations from P cascade.

The following restrictions apply to delete-connected relationships:

- When a table is delete-connected to itself in a referential cycle of more than one table, the cycle must not contain a delete rule of either RESTRICT or SET NULL.
- A table must not both be a dependent table in a CASCADE relationship (self-referencing or referencing another table) and have a self-referencing relationship with a delete rule of either RESTRICT or SET NULL.
- When a table is delete-connected to another table through multiple relationships where such relationships have overlapping foreign keys, these relationships must have the same delete rule and none of these can be SET NULL.
- When a table is delete-connected to another table through multiple relationships where one of the relationships is specified with delete rule SET NULL, the foreign key definition of this relationship must not contain any distribution key or MDC key column, or add data partitioning key column, or RCT key column.

- When two tables are delete-connected to the same table through CASCADE relationships, the two tables must not be delete-connected to each other where the delete connected paths end with delete rule RESTRICT or SET NULL.

Informational constraints

An *informational constraint* is a constraint attribute that can be used by the SQL compiler to improve the access to data. Informational constraints are not enforced by the database manager, and are not used for additional verification of data; rather, they are used to improve query performance.

Informational constraints are defined using the CREATE TABLE or ALTER TABLE statements. You first add referential integrity or check constraints and then associate constraint attributes to them specifying whether the database manager is to enforce the constraint or not; and, whether the constraint is to be used for query optimization or not.

Designing constraints

When designing and creating constraints, it is a good idea to use a naming convention that properly identifies the different types constraints. This is particularly important for diagnosing errors that might occur.

You can design the following types of constraints:

- NOT NULL constraints
- Unique constraints
- Primary key constraints
- (Table) Check constraints
- Foreign key (referential) constraints
- Information constraints

Designing unique constraints

Unique constraints ensure that every value in the specified key is unique. A table can have any number of unique constraints, with one unique constraint defined as a primary key.

Restrictions

- A unique constraint might not be defined on a subtable.
- There can be only one primary key per table.

You define a unique constraint with the UNIQUE clause in the CREATE TABLE or ALTER TABLE statements. The unique key can consist of more than one column. More than one unique constraint is allowed on a table.

Once established, the unique constraint is enforced automatically by the database manager when an INSERT or UPDATE statement modifies the data in the table. The unique constraint is enforced through a unique index.

When a unique constraint is defined in an ALTER TABLE statement and an index exists on the same set of columns of that unique key, that index becomes the unique index and is used by the constraint.

You can take any one unique constraint and use it as the primary key. The primary key can be used as the parent key in a referential constraint (along with other

unique constraints). You define a primary key with the PRIMARY KEY clause in the CREATE TABLE or ALTER TABLE statement. The primary key can consist of more than one column.

A primary index forces the value of the primary key to be unique. When a table is created with a primary key, the database manager creates a primary index on that key.

Some performance tips for indexes used as unique constraints include:

When performing an initial load of an empty table with indexes, LOAD gives better performance than IMPORT. This is true no matter whether you are using the INSERT or REPLACE modes of LOAD. When appending a substantial amount of data to an existing table with indexes (using IMPORT INSERT, or LOAD INSERT), LOAD gives slightly better performance than IMPORT. If you are using the IMPORT command for an initial large load of data, create the unique key after the data has been imported or loaded. This avoids the overhead of maintaining the index while the table is being loaded. It also results in the index using the least amount of storage. If you are using the load utility in REPLACE mode, create the unique key before loading the data. In this case, creation of the index during the load is more efficient than using the CREATE INDEX statement after the load.

Designing primary key constraints

Each table can have one primary key. A primary key is a column or combination of columns that has the same properties as a unique constraint. You can use a primary key and foreign key constraints to define relationships between tables.

Because the primary key is used to identify a row in a table, it should be unique and have very few additions or deletions. A table cannot have more than one primary key, but it can have multiple unique keys. Primary keys are optional, and can be defined when a table is created or altered, using the PRIMARY KEY clause. They are also beneficial, because they order the data when data is exported or reorganized.

Primary key constraints are designed like unique constraints, as described in “Designing unique constraints” on page 274. The only difference is that you can have only one primary key constraint per table, whereas, you can have many unique constraints.

Note: You can have primary key constraints based on composite primary keys.

Designing check constraints

When creating check constraints, one of two things can happen: (i) all the rows meet the check constraint, or (ii) some or all the rows do not meet the check constraint.

All the rows meet the check constraint

When all the rows meet the check constraint, the check constraint will be created successfully. Future attempts to insert or update data that does not meet the constraint business rule will be rejected.

Some or all the rows do not meet the check constraint

When there are some rows that do not meet the check constraint, the check constraint will not be created (that is, the ALTER TABLE statement will fail). The ALTER TABLE statement, which adds a new constraint to the EMPLOYEE table, is shown below. The check constraint is named

CHECK_JOB. The database manager will use this name to inform you about which constraint was violated if an INSERT or UPDATE statement fails. The CHECK clause is used to define a table-check constraint.

```
ALTER TABLE EMPLOYEE
  ADD CONSTRAINT check_job
  CHECK (JOB IN ('Engineer', 'Sales', 'Manager'));
```

An ALTER TABLE statement was used because the table had already been defined. If there are values in the EMPLOYEE table that conflict with the constraint being defined, the ALTER STATEMENT will not be completed successfully.

As check constraints and other types of constraints are used to implement business rules, you may need to change them from time to time. This could happen when the business rules change in your organization. Whenever a check constraint needs to be changed, you must drop it and recreate a new one. Check constraints can be dropped at any time, and this action will not affect your table or the data within it. When you drop a check constraint, you must be aware that data validation performed by the constraint will no longer be in effect.

Comparison of check constraints and BEFORE triggers

You need to consider the difference between check constraints when considering whether to use triggers or check constraints to preserve the integrity of your data.

The integrity of the data in a relational database must be maintained as multiple users access and change the data. Whenever data is shared, there is a need to ensure the accuracy of the values within databases.

Check constraints

A (table) check constraint sets restrictions on data added to a specific table. You can use a table check constraint to define restrictions, beyond those of the data type, on the values that are allowed for a column in the table. Table check constraints take the form of range checks or checks against other values in the same row of the same table.

If the rule applies for all applications that use the data, use a table check constraint to enforce your restriction on the data allowed in the table. Table check constraints make the restriction generally applicable and easier to maintain.

The enforcement of check constraints is important for maintaining data integrity, but it also carries a certain amount of overhead that can impact performance whenever large volumes of data are modified.

BEFORE triggers

By using triggers that run before an update or insert, values that are being updated or inserted can be modified *before* the database is actually modified. These can be used to transform input from the application (user view of the data) to an internal database format where desired. BEFORE triggers can also be used to cause other non-database operations to be activated through user-defined functions.

In addition to modification, a common use of the BEFORE triggers is for data verification using the SIGNAL clause.

There are two differences between BEFORE triggers and check constraints when used for data verification:

1. BEFORE triggers, unlike check constraints, are not restricted to access other values in the same row of the same table.

2. During a SET INTEGRITY operation on a table after a LOAD operation, triggers (including BEFORE triggers) are not executed. Check constraints, however, are verified.

Designing foreign key (referential) constraints

Referential integrity is imposed by adding foreign key (or referential) constraints to table and column definitions, and to create an index on all the foreign key columns. Once the index and foreign key constraints are defined, changes to the data within the tables and columns is checked against the defined constraint. Completion of the requested action depends on the result of the constraint checking.

Referential constraints are established with the FOREIGN KEY clause, and the REFERENCES clause in the CREATE TABLE or ALTER TABLE statements. There are effects from a referential constraint on a typed table or to a parent table that is a typed table that you should consider before creating a referential constraint.

The identification of foreign keys enforces constraints on the values within the rows of a table or between the rows of two tables. The database manager checks the constraints specified in a table definition and maintains the relationships accordingly. The goal is to maintain integrity whenever one database object references another, without performance degradation.

For example, primary and foreign keys each have a department number column. For the EMPLOYEE table, the column name is WORKDEPT, and for the DEPARTMENT table, the name is DEPTNO. The relationship between these two tables is defined by the following constraints:

- There is only one department number for each employee in the EMPLOYEE table, and that number exists in the DEPARTMENT table.
- Each row in the EMPLOYEE table is related to no more than one row in the DEPARTMENT table. There is a unique relationship between the tables.
- Each row in the EMPLOYEE table that has a non-null value for WORKDEPT is related to a row in the DEPTNO column of the DEPARTMENT table.
- The DEPARTMENT table is the parent table, and the EMPLOYEE table is the dependent table.

The statement defining the parent table, DEPARTMENT, is:

```
CREATE TABLE DEPARTMENT
  (DEPTNO  CHAR(3)    NOT NULL,
   DEPTNAME VARCHA(29) NOT NULL,
   MGRNO   CHAR(6),
   ADMRDEPT CHAR(3)   NOT NULL,
   LOCATION CHAR(16),
   PRIMARY KEY (DEPTNO))
IN RESOURCE
```

The statement defining the dependent table, EMPLOYEE, is:

```
CREATE TABLE EMPLOYEE
  (EMPNO   CHAR(6)    NOT NULL PRIMARY KEY,
   FIRSTNAME VARCHA(12) NOT NULL,
   LASTNAME VARCHA(15) NOT NULL,
   WORKDEPT CHAR(3),
   PHONENO  CHAR(4),
   PHOTO    BLOB(10m) NOT NULL,
   FOREIGN KEY DEPT (WORKDEPT)
   REFERENCES DEPARTMENT ON DELETE NO ACTION)
IN RESOURCE
```

By specifying the DEPTNO column as the primary key of the DEPARTMENT table and WORKDEPT as the foreign key of the EMPLOYEE table, you are defining a referential constraint on the WORKDEPT values. This constraint enforces referential integrity between the values of the two tables. In this case, any employees that are added to the EMPLOYEE table must have a department number that can be found in the DEPARTMENT table.

The delete rule for the referential constraint in the employee table is NO ACTION, which means that a department cannot be deleted from the DEPARTMENT table if there are any employees in that department.

Although the previous examples use the CREATE TABLE statement to add a referential constraint, the ALTER TABLE statement can also be used.

Another example: The same table definitions are used as those in the previous example. Also, the DEPARTMENT table is created before the EMPLOYEE table. Each department has a manager, and that manager is listed in the EMPLOYEE table. MGRNO of the DEPARTMENT table is actually a foreign key of the EMPLOYEE table. Because of this referential cycle, this constraint poses a slight problem. You could add a foreign key later. You could also use the CREATE SCHEMA statement to create both the EMPLOYEE and DEPARTMENT tables at the same time.

See also, “Foreign keys in referential constraints” on page 280.

Examples of interaction between triggers and referential constraints

Update operations may cause the interaction of triggers with referential constraints and check constraints.

Figure 23 and the associated description are representative of the processing that is performed for an statement that updates data in the database.

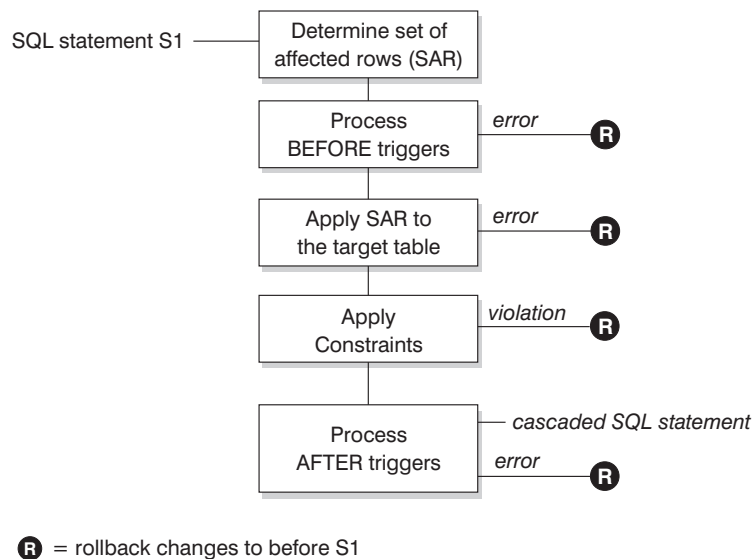


Figure 23. Processing an statement with associated triggers and constraints

Figure 23 shows the general order of processing for an statement that updates a table. It assumes a situation where the table includes BEFORE triggers, referential

constraints, check constraints and AFTER triggers that cascade. The following is a description of the boxes and other items found in Figure 23 on page 278.

- statement S_1

This is the DELETE, INSERT, or UPDATE statement that begins the process. The statement S_1 identifies a table (or an updatable view over some table) referred to as the *subject table* throughout this description.

- Determine set of affected rows

This step is the starting point for a process that repeats for referential constraint delete rules of CASCADE and SET NULL and for cascaded statements from AFTER triggers.

The purpose of this step is to determine the *set of affected rows* for the statement. The set of rows included is based on the statement:

- for DELETE, all rows that satisfy the search condition of the statement (or the current row for a positioned DELETE)
- for INSERT, the rows identified by the VALUES clause or the fullselect
- for UPDATE, all rows that satisfy the search condition (or the current row for a positioned UPDATE).

If the set of affected rows is empty, there will be no BEFORE triggers, changes to apply to the subject table, or constraints to process for the statement.

- Process BEFORE triggers

All BEFORE triggers are processed in ascending order of creation. Each BEFORE trigger will process the triggered action once for each row in the set of affected rows.

An error may occur during the processing of a triggered action in which case all changes made as a result of the original statement S_1 (so far) are rolled back.

If there are no BEFORE triggers or the set of affected is empty, this step is skipped.

- Apply the set of affected rows to the subject table

The actual delete, insert, or update is applied using the set of affected rows to the subject table in the database.

An error may occur when applying the set of affected rows (such as attempting to insert a row with a duplicate key where a unique index exists) in which case all changes made as a result of the original statement S_1 (so far) are rolled back.

- Apply Constraints

The constraints associated with the subject table are applied if set of affected rows is not empty. This includes unique constraints, unique indexes, referential constraints, check constraints and checks related to the WITH CHECK OPTION on views. Referential constraints with delete rules of cascade or set null may cause additional triggers to be activated.

A violation of any constraint or WITH CHECK OPTION results in an error and all changes made as a result of S_1 (so far) are rolled back.

- Process AFTER triggers

All AFTER triggers activated by S_1 are processed in ascending order of creation. FOR EACH STATEMENT triggers will process the triggered action exactly once, even if the set of affected rows is empty. FOR EACH ROW triggers will process the triggered action once for each row in the set of affected rows.

An error may occur during the processing of a triggered action in which case all changes made as a result of the original S_1 (so far) are rolled back.

The triggered action of a trigger may include triggered statements that are DELETE, INSERT or UPDATE statements. For the purposes of this description, each such statement is considered a *cascaded statement*.

A cascaded statement is a DELETE, INSERT, or UPDATE statement that is processed as part of the triggered action of an AFTER trigger. This statement starts a cascaded level of trigger processing. This can be thought of as assigning the triggered statement as a new S_1 and performing all of the steps described here recursively.

Once all triggered statements from all AFTER triggers activated by each S_1 have been processed to completion, the processing of the original S_1 is completed.

- R = roll back changes to before S_1

Any error (including constraint violations) that occurs during processing results in a roll back of all the changes made directly or indirectly as a result of the original statement S_1 . The database is therefore back in the same state as immediately prior to the execution of the original statement S_1

Foreign keys in referential constraints

A foreign key references a primary key or a unique key in the same or another table. A foreign key assignment indicates that referential integrity is to be maintained according to the specified referential constraints.

You define a foreign key with the FOREIGN KEY clause in the CREATE TABLE or ALTER TABLE statement. A foreign key makes its table dependent on another table called a parent table. The values in the column or set of columns that make up the foreign key in one table must match the unique key or primary key values of the parent table.

The number of columns in the foreign key must be equal to the number of columns in the corresponding primary or unique constraint (called a parent key) of the parent table. In addition, corresponding parts of the key column definitions must have the same data types and lengths. The foreign key can be assigned a *constraint name*. If you do not assign a name, one is automatically assigned. For ease of use, it is recommended that you assign a *constraint name* and do not use the system-generated name.

The value of a composite foreign key matches the value of a parent key **if** the value of each column of the foreign key is equal to the value of the corresponding column of the parent key. A foreign key containing null values cannot match the values of a parent key, since a parent key by definition can have no null values. However, a null foreign key value is always valid, regardless of the value of any of its non-null parts.

The following rules apply to foreign key definitions:

- A table can have many foreign keys
- A foreign key is nullable if any part is nullable
- A foreign key value is null if any part is null.

When working with foreign keys you can do the following:

- Create a table with zero or more foreign keys.
- Define foreign keys when a table is created or altered.
- Drop foreign keys when a table is altered.

Table constraint implications for utility operations

If the table being loaded into has referential integrity constraints, the load utility places the table into the set integrity pending state to inform you that the SET INTEGRITY statement is required to be run on the table, in order to verify the referential integrity of the loaded rows. After the load utility has completed, you will need to issue the SET INTEGRITY statement to carry out the referential integrity checking on the loaded rows and to bring the table out of the set integrity pending state.

For example, if the DEPARTMENT and EMPLOYEE tables are the only tables that have been placed in set integrity pending state, you can execute the following statement:

```
SET INTEGRITY FOR DEPARTMENT ALLOW WRITE ACCESS,  
EMPLOYEE ALLOW WRITE ACCESS,  
IMMEDIATE CHECKED FOR EXCEPTION IN DEPARTMENT,  
USE DEPARTMENT_EX,  
IN EMPLOYEE USE EMPLOYEE_EX
```

The import utility is affected by referential constraints in the following ways:

- The REPLACE and REPLACE CREATE functions are not allowed if the object table has any dependents other than itself.

To use these functions, first drop all foreign keys in which the table is a parent. When the import is complete, re-create the foreign keys with the ALTER TABLE statement.

- The success of importing into a table with self-referencing constraints depends on the order in which the rows are imported.

Statement dependencies when changing objects

Statement dependencies include package and cached dynamic SQL and XQuery statements. A *package* is a database object that contains the information needed by the database manager to access data in the most efficient way for a particular application program. *Binding* is the process that creates the package the database manager needs in order to access the database when the application is executed.

Packages and cached dynamic SQL and XQuery statements can be dependent on many types of objects.

These objects could be explicitly referenced, for example, a table or user-defined function that is involved in an SQL SELECT statement. The objects could also be implicitly referenced, for example, a dependent table that needs to be checked to ensure that **referential constraints** are not violated when a row in a parent table is deleted. Packages are also dependent on the privileges which have been granted to the package creator.

If a package or cached dynamic query statement depends on an object and that object is dropped, the package or cached dynamic query statement is placed in an “invalid” state. If a package depends on a user-defined function and that function is dropped, the package is placed in an “inoperative” state, with the following conditions:

- A cached dynamic SQL or XQuery statement that is in an invalid state is automatically re-optimized on its next use. If an object required by the statement has been dropped, execution of the dynamic SQL or XQuery statement might fail with an error message.

- A package that is in an invalid state is implicitly rebound on its next use. Such a package can also be explicitly rebound. If a package was marked invalid because a trigger was dropped, the rebound package no longer invokes the trigger.
- A package that is in an inoperative state must be explicitly rebound before it can be used.

Federated database objects have similar dependencies. For example, dropping a server invalidates any packages or cached dynamic SQL referencing nicknames associated with that server.

In some cases, it is not possible to rebound the package. For example, if a table has been dropped and not re-created, the package cannot be rebound. In this case, you need to either re-create the object or change the application so it does not use the dropped object.

In many other cases, for example if one of the **constraints** was dropped, it is possible to rebound the package.

The following system catalog views help you to determine the state of a package and the package's dependencies:

- SYSCAT.PACKAGEAUTH
- SYSCAT.PACKAGEDEP
- SYSCAT.PACKAGES

Designing informational constraints

Constraints that are enforced by the database manager when records are inserted or updated can lead to high amounts of system overhead, especially when loading large quantities of records that have referential integrity constraints. If an application has already verified information before inserting a record into the table, it may be more efficient to use informational constraints, rather than normal constraints.

Informational constraints tell the database manager what rules the data conforms to, but the rules are not enforced by the database manager. However, this information can be used by the DB2 optimizer and could result in better performance of SQL queries.

The following example illustrates the use of information constraints and how they work. This simple table contains information on applicants' age and gender:

```
CREATE TABLE APPLICANTS
(
  AP_NO INT NOT NULL,
  GENDER CHAR(1) NOT NULL,
  CONSTRAINT GENDEROK
  CHECK (GENDER IN ('M', 'F'))
  NOT ENFORCED
  ENABLE QUERY OPTIMIZATION,
  AGE INT NOT NULL,
  CONSTRAINT AGEOK
  CHECK (AGE BETWEEN 1 AND 80)
  NOT ENFORCED
  ENABLE QUERY OPTIMIZATION,
);
```

This example contains two clauses that change the behavior of the column constraints. The first option is NOT ENFORCED, which instructs the database manager not to enforce the checking of this column when data is inserted or updated.

The second option is ENABLE QUERY OPTIMIZATION which is used by the database manager when SELECT statements are run against this table. When this value is specified, the database manager will use the information in the constraint when optimizing the SQL.

If the table contains the NOT ENFORCED option, the behavior of insert statements may appear odd. The following SQL will not result in any errors when run against the APPLICANTS table:

```
INSERT INTO APPLICANTS VALUES
(1, 'M', 54),
(2, 'F', 38),
(3, 'M', 21),
(4, 'F', 89),
(5, 'C', 10),
(6, 'S', 100),
```

Applicant number five has a gender (C), for child, and applicant number six has both an unusual gender and exceeds the age limits of the AGE column. In both cases the database manager will allow the insert to occur since the constraints are NOT ENFORCED. The result of a select statement against the table is shown below:

```
SELECT * FROM APPLICANTS
WHERE GENDER = 'C';

APPLICANT GENDER AGE
-----
0 record(s) selected.
```

The database manager returned the incorrect answer to the query, even though the value 'C' is found within the table, but the constraint on this column tells the database manager that the only valid values are either 'M' or 'F'. The ENABLE QUERY OPTIMIZATION keyword also allowed the database manager to use this constraint information when optimizing the statement. If this is not the behavior that you want, then the constraint needs to be changed through the use of the ALTER TABLE statement, as shown below:

```
ALTER TABLE APPLICANTS
ALTER CHECK AGEOK DISABLE QUERY OPTIMIZATION
```

If the query is reissued, the database manager will return the following correct results:

```
SELECT * FROM APPLICANTS
WHERE SEC = 'C';

APPLICANT GENDER AGE
-----
5 C 10

1 record(s) selected.
```

The best scenario for using informational constraints occurs when you can guarantee that the application program is the only application inserting and updating the data. If the application already checks all of the information beforehand (like gender and age) then using informational constraints can result in

faster performance and no duplication of effort. Another possible use of informational constraints is in the design of data warehouses.

Creating and modifying constraints

Constraints can be added to existing tables with the ALTER TABLE statement.

The constraint name cannot be the same as any other constraint specified within an ALTER TABLE statement, and must be unique within the table (this includes the names of any referential integrity constraints that are defined). Existing data is checked against the new condition before the statement succeeds.

Creating and modifying unique constraints

Unique constraints can be added to an existing table. The constraint name cannot be the same as any other constraint specified within the ALTER TABLE statement, and must be unique within the table (this includes the names of any referential integrity constraints that are defined). Existing data is checked against the new condition before the statement succeeds.

To define unique constraints using the command line, use the ADD CONSTRAINT option of the ALTER TABLE statement. For example, the following statement adds a unique constraint to the EMPLOYEE table that represents a new way to uniquely identify employees in the table:

```
ALTER TABLE EMPLOYEE
  ADD CONSTRAINT NEWID UNIQUE(EMPNO,HIREDATE)
```

To modify this constraint, you would have to drop it, and then recreate it.

Creating and modifying primary key constraints

A primary key constraint can be added to an existing table. The constraint name must be unique within the table (this includes the names of any referential integrity constraints that are defined). Existing data is checked against the new condition before the statement succeeds.

To add primary keys using the command line, enter:

```
ALTER TABLE <name>
  ADD CONSTRAINT <column_name>
  PRIMARY KEY <column_name>
```

An existing constraint cannot be modified. To define another column, or set of columns, as the primary key, the existing primary key definition must first be dropped, and then recreated.

Creating and modifying check constraints

When a table check constraint is added, packages and cached dynamic SQL that insert or update the table might be marked as invalid.

To add a table check constraint using the command line, enter:

```
ALTER TABLE EMPLOYEE
  ADD CONSTRAINT REVENUE CHECK (SALARY + COMM > 25000)
```

To modify this constraint, you would have to drop it, and then recreate it.

Creating and modifying foreign key (referential) constraints

A foreign key is a reference to the data values in another table. There are different types of foreign key constraints.

When a foreign key is added to a table, packages and cached dynamic SQL containing the following statements might be marked as invalid:

- Statements that insert or update the table containing the foreign key

- Statements that update or delete the parent table.

To add foreign keys using the command line, enter:

```
ALTER TABLE <name>
  ADD CONSTRAINT <column_name>
  FOREIGN KEY <column_name>
  ON DELETE <action_type>
  ON UPDATE <action_type>
```

The following examples show the ALTER TABLE statement to add primary keys and foreign keys to a table:

```
ALTER TABLE PROJECT
  ADD CONSTRAINT PROJECT_KEY
  PRIMARY KEY (PROJNO)
ALTER TABLE EMP_ACT
  ADD CONSTRAINT ACTIVITY_KEY
  PRIMARY KEY (EMPNO, PROJNO, ACTNO)
  ADD CONSTRAINT ACT_EMP_REF
  FOREIGN KEY (EMPNO)
  REFERENCES EMPLOYEE
  ON DELETE RESTRICT
  ADD CONSTRAINT ACT_PROJ_REF
  FOREIGN KEY (PROJNO)
  REFERENCES PROJECT
  ON DELETE CASCADE
```

To modify this constraint, you would have to drop it and then recreate it.

Creating and modifying informational constraints

To improve the performance of queries, you can add informational constraints to your tables. You add informational constraints using the CREATE TABLE or ALTER TABLE statement when you specify the NOT ENFORCED option on the DDL.

Restriction: After you define informational constraints on a table, you can only alter the column names for that table after you remove the informational constraints.

To specify informational constraints on a table using the command line, enter the following command for a new table:

```
ALTER TABLE <name> <constraint attributes> NOT ENFORCED
```

ENFORCED or NOT ENFORCED: Specifies whether the constraint is enforced by the database manager during normal operations such as insert, update, or delete.

- ENFORCED cannot be specified for a functional dependency (SQLSTATE 42621).
- NOT ENFORCED should only be specified if the table data is independently known to conform to the constraint. Query results might be unpredictable if the data does not actually conform to the constraint.

To modify this constraint, you would have to drop it and then recreate it.

Viewing constraint definitions for a table

Constraint definitions on a table can be found in the SYSCAT.INDEXES and SYSCAT.REFERENCES catalog views.

The UNIQUERULE column of the SYSCAT.INDEXES view indicates the characteristic of the index. If the value of this column is P, the index is a primary key, and if it is U, the index is a unique index (but not a primary key).

The SYSCAT.REFERENCES catalog view contains referential integrity (foreign key) constraint information.

Dropping constraints

You can explicitly drop a table check constraint using the ALTER TABLE statement, or implicitly drop it as the result of a DROP TABLE statement.

To drop constraints, use the ALTER TABLE statement with the DROP or DROP CONSTRAINT clauses. This allows you to BIND and continue accessing the tables that contain the affected columns. The name of all unique constraints on a table can be found in the SYSCAT.INDEXES system catalog view.

Dropping unique constraints

You can explicitly drop a unique constraint using the ALTER TABLE statement.

The DROP UNIQUE clause of the ALTER TABLE statement drops the definition of the unique constraint constraint-name and all referential constraints that are dependent upon this unique constraint. The constraint-name must identify an existing unique constraint.

```
ALTER TABLE <table-name>
  DROP UNIQUE <constraint-name>
```

Dropping this unique constraint invalidates any packages or cached dynamic SQL that used the constraint.

Dropping primary key constraints

Use the DROP PRIMARY KEY clause of the ALTER TABLE statement to drop primary key constraints.

The DROP PRIMARY KEY clause of the ALTER TABLE statement drops the definition of the primary key and all referential constraints that are dependent upon this primary key. The table must have a primary key. To drop a primary key using the command line, enter:

```
ALTER TABLE <table-name>
  DROP PRIMARY KEY
```

Dropping (table) check constraints

When you drop a check constraint, all packages and cached dynamic statements with INSERT or UPDATE dependencies on the table are invalidated. The name of all check constraints on a table can be found in the SYSCAT.CHECKS catalog view. Before attempting to drop a table check constraint having a system-generated name, look for the name in the SYSCAT.CHECKS catalog view.

The following statement drops the check constraint constraint-name. The constraint-name must identify an existing check constraint defined on the table. To drop a table check constraint using the command line:

```
ALTER TABLE <table_name>
  DROP <check_constraint_name>
```

Dropping foreign key (referential) constraints

Use the DROP CONSTRAINT clause of the ALTER TABLE statement to drop foreign key constraints.

The DROP CONSTRAINT clause of the ALTER TABLE statement drops the constraint *constraint-name*. The *constraint-name* must identify an existing foreign key constraint, primary key, or unique constraint defined on the table. To drop foreign keys using the command line, enter:

```
ALTER TABLE <table-name>
  DROP FOREIGN KEY <foreign_key_name>
```

The following examples use the DROP PRIMARY KEY and DROP FOREIGN KEY clauses in the ALTER TABLE statement to drop primary keys and foreign keys on a table:

```
ALTER TABLE EMP_ACT
  DROP PRIMARY KEY
  DROP FOREIGN KEY ACT_EMP_REF
  DROP FOREIGN KEY ACT_PROJ_REF
ALTER TABLE PROJECT
  DROP PRIMARY KEY
```

When a foreign key constraint is dropped, packages or cached dynamic statements containing the following might be marked as invalid:

- Statements that insert or update the table containing the foreign key
- Statements that update or delete the parent table.

Chapter 13. Indexes

An *index* is a set of one or more keys, each key pointing to a row in a table. The SQL *optimizer* automatically chooses the most efficient way to access data in tables. The optimizer takes indexes into consideration when determining the fastest access path to data.

Note: Not all indexes point to rows in a table. MDC block indexes point to extents (or blocks) of the data. XML indexes for XML data use particular XML pattern expressions to index paths and values in XML documents stored within a single column. The data type of that column must be XML. Both MDC block indexes and XML indexes are system generated indexes.

Indexes are used by the database manager to:

- Improve performance. In most cases, access to data is faster with an index. Although an index cannot be created for a view, an index created for the table on which a view is based can sometimes improve the performance of operations on that view.
- Ensure uniqueness. A table with a unique index cannot have rows with identical keys.

As data is added to a table, unless other actions have been carried out on the table or the data being added, the data is simply appended to the bottom of the table. There is no order to the data. When searching for a particular row of data, each row of the table from first to last must be checked. Indexes are used as a means to access the data within the table in an order that might otherwise not be available.

A column value in a row of data can be used to identify the entire row. One or more columns might be needed to identify the row. Such columns are known as a *key*. A column can be used in more than one key.

An index is ordered by the values within a key. Keys can be unique or non-unique. Each table should have at least one unique key; but can also have other, non-unique keys. Each index has exactly one key. For example, you might use the employee ID number (unique) as the key for one index and the department number (non-unique) as the key for a different index.

Example

Table A in Figure 24 on page 290 has an index based on the employee numbers in the table. This key value provides a pointer to the rows in the table. For example, employee number 19 points to employee KMP. An index allows efficient access to rows in a table by creating a path to the data through pointers.

Unique indexes can be created to ensure uniqueness of the index key. An *index key* is a column or an ordered collection of columns on which an index is defined. Using a unique index will ensure that the value of each index key in the indexed column or columns is unique.

Figure 24 on page 290 shows the relationship between an index and a table.

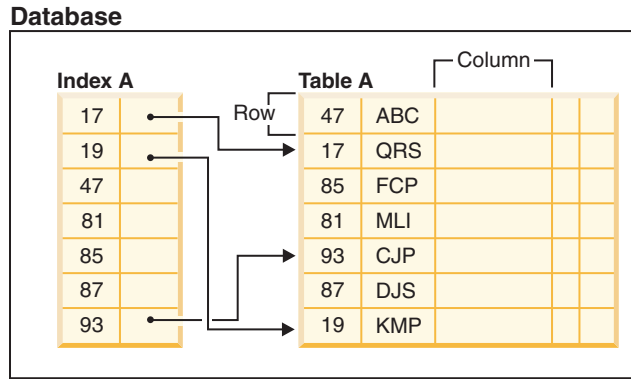


Figure 24. Relationship between an index and a table

Figure 25 illustrates the relationships among some database objects. It also shows that tables, indexes, and long data are stored in table spaces.

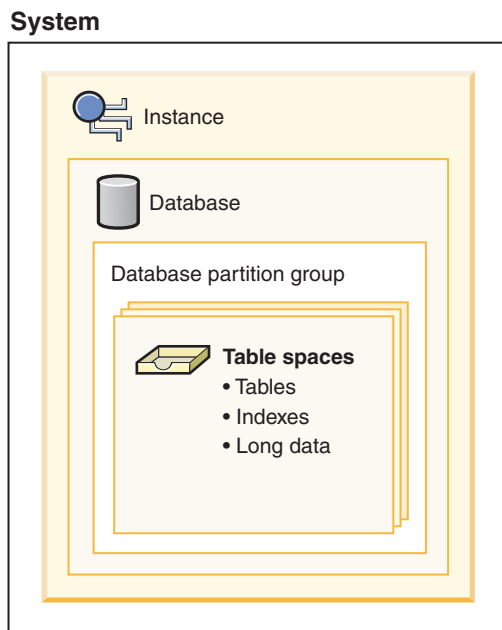


Figure 25. Relationships among selected database objects

Types of indexes

There are five types of indexes: unique and non-unique indexes, and clustered and non-clustered indexes, and system generated block indexes for multidimensional clustered (MDC) tables .

Unique and non-unique indexes

Unique indexes are indexes that help maintain data integrity by ensuring that no two rows of data in a table have identical key values.

When attempting to create a unique index for a table that already contains data, values in the column or columns that comprise the index are checked for uniqueness; if the table contains rows with duplicate key values, the index creation process fails. Once a unique index has been defined for a table, uniqueness is

enforced whenever keys are added or changed within the index. (This includes insert, update, load, import, and set integrity, to name a few.) In addition to enforcing the uniqueness of data values, a unique index can also be used to improve data retrieval performance during query processing.

Non-unique indexes, on the other hand, are not used to enforce constraints on the tables with which they are associated. Instead, non-unique indexes are used solely to improve query performance by maintaining a sorted order of data values that are used frequently.

Clustered and non-clustered indexes

Index architectures are classified as clustered or non-clustered. Clustered indexes are indexes whose order of the rows in the data pages correspond to the order of the rows in the index. This is why only one clustered index can exist in a given table, whereas, many non-clustered indexes can exist in the table. In some relational database management systems, the leaf node of the clustered index corresponds to the actual data, not a pointer to data that resides elsewhere.

Both clustered and non-clustered indexes contain only keys and record IDs in the index structure. The record IDs always point to rows in the data pages. The only difference between clustered and non-clustered indexes is that the database manager attempts to keep the data in the data pages in the same order as the corresponding keys appear in the index pages. Thus the database manager will attempt to insert rows with similar keys onto the same pages. If the table is reorganized, it will be inserted into the data pages in the order of the index keys.

Reorganizing a table with respect to a chosen index re-clusters the data. A clustered index is most useful for columns that have range predicates because it allows better sequential access of data in the table. This results in fewer page fetches, since like values are on the same data page.

In general, only one of the indexes in a table can have a high degree of clustering.

Improving performance with clustering indexes

Clustering indexes can improve the performance of most query operations because they provide a more linear access path to data, which has been stored in pages. In addition, because rows with similar index key values are stored together, prefetching is usually more efficient when clustering indexes are used.

However, clustering indexes cannot be specified as part of the table definition used with the CREATE TABLE statement. Instead, clustering indexes are only created by executing the CREATE INDEX statement with the CLUSTER option specified. Then the ALTER TABLE statement should be used to add a primary key that corresponds to the clustering index created to the table. This clustering index will then be used as the table's primary key index.

Note: Setting PCTFREE in the table to an appropriate value using the ALTER TABLE statement can help the table remain clustered by leaving adequate free space to insert rows in the pages with similar values. For more information, see the ALTER TABLE statement and "Reducing the need to reorganize tables and indexes" in *Tuning Database Performance*.

Generally, clustering is more effectively maintained if the clustering index is unique.

Differences between primary key or unique key constraints and unique indexes

It is important to understand that there is no significant difference between a primary unique key constraint and a unique index. The database manager uses a combination of a unique index and the NOT NULL constraint to implement the relational database concept of primary and unique key constraints. Therefore, unique indexes do not enforce primary key constraints by themselves because they allow null values. (Although null values represent unknown values, when it comes to indexing, a null value is treated as being equal to other null values.)

Therefore, if a unique index consists of a single column, only one null value is allowed—more than one null value would violate the unique constraint. Similarly, if a unique index consists of multiple columns, a specific combination of values and nulls can be used only once.

Bi-directional indexes

By default, bi-directional indexes allow scans in both the forward and reverse directions. The ALLOW REVERSE SCANS clause of the CREATE INDEX statement enables both forward and reverse index scans, that is, in the order defined at index creation time and in the opposite (or reverse) order. This option allows you to:

- Facilitate MIN and MAX functions
- Fetch previous keys
- Eliminate the need for the database manager to create a temporary table for the reverse scan
- Eliminate redundant reverse order indexes

If DISALLOW REVERSE SCANS is specified then the index cannot be scanned in reverse order. (But physically it will be exactly the same as an ALLOW REVERSE SCANS index.)

Designing indexes

Indexes are typically used to speed up access to a table. However, they can also serve a logical data design purpose.

For example, a unique index does not allow entry of duplicate values in the columns, thereby guaranteeing that no two rows of a table are the same. Indexes can also be created to order the values in a column in ascending or descending sequence.

Note: When creating indexes, keep in mind that while they may improve read performance, they will negatively impact write performance. This is because for every row that the database manager writes to a table, it must also update any affected indexes. Therefore, you should create indexes only when there is a clear overall performance advantage.

When creating indexes, you must also take into account the structure of the tables and the type of queries that are most frequently performed on them. For example, columns appearing in the WHERE clause of a frequently issued query are good candidates for indexes. In less frequently run queries, however, the cost that an index incurs for performance in INSERT and UPDATE statement might outweigh the benefits.

Similarly, columns that figure in a GROUP BY clause of a frequent query might benefit from the creation of an index, particularly if the number of values used to group the rows is small relative to the number of rows being grouped.

Guidelines and considerations when designing indexes

- An *index* is defined by columns in the table. It can be defined by the creator of a table, or by a user who knows that certain columns require direct access. A primary index key is automatically created on the primary key, unless a user-defined index already exists.
- An *index key* is a column or collection of columns on which an index is defined, and determines the usefulness of an index. Although the order of the columns making up an index key does not make a difference to index key creation, it might make a difference to the optimizer when it is deciding whether or not to use an index.
- Any number of indexes can be defined on a particular table, and they can have a beneficial effect on the performance of queries. The index manager must maintain the indexes during update, delete and insert operations. Creating a large number of indexes for a table that receives many updates can slow down processing of requests. Similarly, large index keys can also slow down processing of requests. Therefore, use indexes only where a clear advantage for frequent access exists.
- Column data which is not part of the unique index key but which is to be stored or maintained in the index is called an include column. Include columns can be specified for unique indexes only. When creating an index with include columns, only the unique key columns are sorted and considered for uniqueness. The use of include columns may enable index only access for data retrieval, thus improving performance.
- If the table being indexed is empty, an index is still created, but no index entries are made until the table is loaded or rows are inserted. If the table is not empty, the database manager creates the index entries while processing the CREATE INDEX statement.
- For a *clustering index*, the database manager attempts to place new rows for the table physically close to existing rows with similar key values (as defined by the index).
- If you want a *primary key index* to be a clustering index, a primary key should not be specified on the CREATE TABLE statement. Once a primary key is created, the associated index cannot be modified. Instead, issue a CREATE TABLE without a primary key clause. Then issue a CREATE INDEX statement, specifying clustering attributes. Finally, use the ALTER TABLE statement to add a primary key that corresponds to the index just created. This index will be used as the primary key index.
- Indexes consume disk space. The amount of disk space varies depending on the length of the key columns and the number of rows being indexed. The size of the index increases as more data is inserted into the table. Therefore, consider the amount of data being indexed when planning the size of the database. Some of the indexing sizing considerations include:
 - Primary and unique key constraints will always create a system-generated unique index.
 - The creation of an MDC table will also create system-generated block indexes.
 - XML columns will always cause system-generated indexes to be created.
 - It is usually beneficial to create indexes on foreign key constraint columns.

Note: The maximum number of columns in an index is 64. However, if you are indexing a typed table, the maximum number of columns in an index is 63. The maximum length of an index key, including all overhead, is $indexpagesize/4$. The maximum indexes allowed on a table is 32 767. The maximum length of an index key must not be greater than the index key length limit for the page size. For column stored lengths, see the “CREATE TABLE statement”. For the key length limits, see the “SQL and XQuery limits” topic.

Note:

Tools for designing indexes

Once you have created your tables, you need to consider how rapidly the database manager will be able to retrieve data from them. You can use the Design Advisor or the `db2adv` command to help you design your indexes.

Creating useful indexes on your tables can significantly improve query performance. Like indexes of a book, indexes on tables allow specific information to be located rapidly, with minimal searching. Using an index to retrieve particular rows from a table can reduce the number of expensive input/output operations that the database manager needs to perform. This is because an index allows the database manager to locate a row by reading in a relatively small number of data pages, rather than by performing an exhaustive search of all data pages until all matches are found.

The DB2 Design Advisor is a tool that can help you significantly improve your workload performance. The task of selecting which indexes, MQTs, clustering dimensions, or database partitions to create for a complex workload can be quite daunting. The Design Advisor identifies all of the objects needed to improve the performance of your workload. Given a set of SQL statements in a workload, the Design Advisor will generate recommendations for:

- New indexes
- New materialized query tables (MQTs)
- Conversion to multidimensional clustering (MDC) tables
- Redistribution of tables
- Deletion of indexes and MQTs unused by the specified workload (through the GUI tool)

You can have the Design Advisor implement some or all of these recommendations immediately or schedule them for a later time.

Using either the Design Advisor GUI or the command-line tool, the Design Advisor can help simplify the following tasks:

- Planning for or setting up a new database
- Workload performance tuning

Space requirements for indexes

When designing indexes, you need to be aware of their space requirements.

For each index, the space needed can be estimated as:

$$(\text{average index key size} + \text{index key overhead}) * \text{number of rows} * 2$$

where:

- The “average index key size” is the byte count of each column in the index key. (When estimating the average column size for VARCHAR and VARGRAPHIC columns, use an average of the current data size, plus two bytes. Do not use the maximum declared size.)
- The “index key overhead” depends on the type of table on which the index is created. For large tables (with or without XML indexes), the value is 11, unless the table is partitioned in which case the value is 13. For regular tables, the value is 9 without XML indexes and 11 with XML indexes. For all regular tables that are partitioned, the value is 11.
- The factor of “2” is for overhead, such as non-leaf pages and free space.

Note:

1. For every column that allows null values, add one extra byte for the null indicator.
2. For block indexes created internally for multidimensional clustering (MDC) tables, the “number of rows” would be replaced by the “number of blocks”.

For each index on an XML column, the space needed can be estimated as:

$$(\text{average index key size} + \text{index key overhead}) * \text{number of indexed nodes} * 2$$

where:

- The “average index key size” is the sum of the key parts that make up the index. The XML index is made up of several XML key parts plus a value (sql-data-type):

$$\text{fixed overhead} + \text{variable overhead} + \text{byte count of sql-data-type}$$

where:

- The “fixed overhead” is 14 bytes.
- The “variable overhead” is the average depth of the indexed node plus 4 bytes.
- The byte count of the sql-data-type value follows the same rules as SQL.
- The “number of indexed nodes” is the number of documents to be inserted multiplied by the number of nodes in a sample document that satisfy the XML pattern expression (XMLPATTERN) in the index definition.

Indexes created before Version 8 (type-1 indexes) are different from those created at version 8 (type-2 indexes) and following. To find out what type of index exists for a table, use the ADMIN_GET_TAB_INFO table function. To convert type-1 indexes to type-2 indexes, use the REORG INDEXES CONVERT command.

When using the REORG INDEXES command, ensure that you have sufficient free space in the table space where the indexes are stored. The amount of free space should be equal to the current size of the index. Additional space may be required if you choose to reorganize the indexes with the ALLOW WRITE ACCESS option. The additional space is for the logs of the activity affecting the indexes during the reorganization of the indexes.

Temporary space is required when creating the index. The maximum amount of temporary space required during index creation can be estimated as:

$$(\text{average index key size} + \text{index key overhead}) * \text{number of rows} * 3.2$$

or

$$(\text{average index key size} + \text{index key overhead}) * \text{number of indexed nodes} * 3.2$$

where the factor of "3.2" is for index overhead, and space required for sorting during index creation.

Note: In the case of non-unique indexes, only five bytes are required to store duplicate key entries. The estimate shown above assumes no duplicates. The space required to store an index may be over-estimated by the formula shown above.

Temporary space is required when inserting if the number of index nodes exceeds 64 KB of data. The amount of temporary space can be estimated as:

$$(\text{average index key size}) * \text{number of indexed nodes} * 1.2$$

The following two formulas can be used to estimate the number of keys per leaf page (the second provides a more accurate estimate). The accuracy of these estimates depends largely on how well the averages reflect the actual data.

Note: For SMS table spaces, the minimum required space for leaf pages is 12 KB. For DMS table spaces, the minimum is an extent.

- A rough estimate of the average number of keys per leaf page is:

$$\frac{(.9 * (U - (M*2))) * (D + 1)}{K + 7 + (5 * D)}$$

where:

- U, the usable space on a page, is approximately equal to the page size minus 100. For a page size of 4096, U is 3996.
- $M = U / (9 + \text{minimumKeySize})$
- D = average number of duplicates per key value
- K = *averageKeySize*

Remember that *minimumKeySize* and *averageKeySize* must have an extra byte for each nullable key part, and an extra two bytes for the length of each variable length key part.

If there are include columns, they should be accounted for in *minimumKeySize* and *averageKeySize*.

The "minimumKeySize" is the sum of the key parts that make up the index:

$$\text{fixed overhead} + \text{variable overhead} + \text{byte count of sql-data-type}$$

where:

- The "fixed overhead" is 13 bytes.
- The "variable overhead" is the minimum depth of the indexed node plus 4 bytes.
- The byte count of the sql-data-type value follows the same rules as SQL.

The .9 can be replaced by any $(100 - \text{pctfree})/100$ value, if a percent free value other than the default value of ten percent is specified during index creation.

- A more accurate estimate of the average number of keys per leaf page is:

$$L = \text{number of leaf pages} = X / (\text{avg number of keys on leaf page})$$

where X is the total number of rows in the table.

For the index on an XML column, X is the total number of indexed nodes in the column.

You can estimate the original size of an index as:

$$(L + 2L/(\text{average number of keys on leaf page})) * \text{pagesize}$$

For DMS table spaces, add the sizes of all indexes on a table and round up to a multiple of the extent size for the table space on which the index resides.

You should provide additional space for index growth due to INSERT/UPDATE activity, from which page splits may result.

Use the following calculation to obtain a more accurate estimate of the original index size, as well as an estimate of the number of levels in the index. (This may be of particular interest if include columns are being used in the index definition.) The average number of keys per non-leaf page is roughly:

$$\frac{(.9 * (U - (M*2))) * (D + 1)}{K + 13 + (9 * D)}$$

where:

- U, the usable space on a page, is approximately equal to the page size minus 100. For a page size of 4096, U is 3996.
- D is the average number of duplicates per key value on non-leaf pages (this will be much smaller than on leaf pages, and you may want to simplify the calculation by setting the value to 0).
- M = U / (9 + *minimumKeySize* for non-leaf pages)
- K = *averageKeySize* for non-leaf pages

The *minimumKeySize* and the *averageKeySize* for non-leaf pages will be the same as for leaf pages, except when there are include columns. Include columns are not stored on non-leaf pages.

You should not replace .9 with (100 - pctfree)/100, unless this value is greater than .9, because a maximum of 10 percent free space will be left on non-leaf pages during index creation.

The number of non-leaf pages can be estimated as follows:

```

if L > 1 then {P++; Z++;}
While (Y > 1)
{
    P = P + Y
    Y = Y / N
    Z++
}

```

where:

- P is the number of pages (0 initially).
- L is the number of leaf pages.
- N is the number of keys for each non-leaf page.
- Y = L / N
- Z is the number of levels in the index tree (1 initially).

Total number of pages is:

$$T = (L + P + 2) * 1.0002$$

The additional 0.02 percent is for overhead, including space map pages.

The amount of space required to create the index is estimated as:

$$T * \text{pagesize}$$

Creating indexes

Indexes can be created to order the values in a column in ascending or descending sequence. You can use the CREATE INDEX statement, the DB2 Design Advisor, or the db2advis Design Advisor command to create the indexes.

For example, to create an index using the CREATE INDEX statement from the command line, enter:

```
CREATE UNIQUE INDEX EMP_IX
ON EMPLOYEE(EMPNO)
INCLUDE(FIRSTNAME, JOB)
```

The INCLUDE clause, applicable only on unique indexes, specifies additional columns to be appended to the set of index key columns. Any columns included with this clause are not used to enforce uniqueness. These included columns may improve the performance of some queries through index only access. This option might:

- Eliminate the need to access data pages for more queries
- Eliminate redundant indexes

If SELECT EMPNO, FIRSTNAME, JOB FROM EMPLOYEE is issued to the table on which this index resides, all of the required data can be retrieved from the index without reading data pages. This improves performance.

Note: In indexes created in Version 8 or later, known as type 2 indexes, keys are just marked as deleted when a row is deleted or updated. There are referred to as pseudo-deleted keys. These keys are not physically removed from a page until clean up is done some time after the deletion or update has committed. Such a clean up might be done by a subsequent transaction which is changing the page where the key is marked deleted. Clean up of pseudo-deleted keys can be explicitly triggered using the CLEANUP ONLY ALL option of the REORG INDEXES utility.

Indexes for tables in a partitioned database environment are built using the same CREATE INDEX statement. Data in the indexes is distributed based on the distribution key of the table. When this is done, a B+ tree is created on each database partition in the database partition group. Each B+ tree indexes the part of the table belonging to that database partition. Columns in a unique index defined on a multi-partition database must be a superset of the columns in the distribution key.

Note: In Version 9.5, on Solaris platforms, the CREATE INDEX statement will hang if a RAW device is used. Sun will be fixing this problem and releasing the fix in a kernel patch.

Modifying indexes

If you want to modify your index, you have to drop the index first and then create the index again. There is no ALTER INDEX statement.

For example, you cannot add a column to the list of key columns without dropping the previous definition and creating a new index. You can add a comment to describe the purpose of the index using the COMMENT statement.

Renaming indexes

You can use the RENAME statement to rename an existing index.

To rename an existing index, issue the following statement from the command line:

```
RENAME INDEX <source index name> TO <target index name>
```

- <source index name> is the name of the existing index that is to be renamed. The name, including the schema name, must identify an index that already

exists in the database. It must not be the name of an index on a declared global temporary table. The schema name must not be SYSIBM, SYSCAT, SYSFUN, or SYSSTAT.

- <target index name> specifies the new name for the index without a schema name. The schema name of the source object is used to qualify the new name for the object. The qualified name must not identify an index that already exists in the database.

When renaming an index, the source index must not be a system-generated index. If the statement is successful, the system catalog tables are updated to reflect the new index name.

Rebuilding indexes

Certain database operations, such as a rollforward through a create index that was not fully logged, can cause an index object to become invalid because the index is not created during the rollforward operation. The index object can be recovered by recreating the indexes in it.

When the database manager detects that an index is no longer valid, it automatically attempts to rebuild it. When the rebuild takes place, it is controlled by the **indexrec** parameter of the database or database manager configuration file. There are five possible settings for this:

- SYSTEM
- RESTART
- RESTART_NO_REDO
- ACCESS
- ACCESS_NO_REDO

RESTART_NO_REDO and ACCESS_NO_REDO are similar to RESTART and ACCESS.

The NO_REDO options mean that even if the index was fully logged during the original operation, such as CREATE INDEX, the index will not be recreated during rollforward, but will instead be created either at restart time or first access. See the **indexrec** parameter for more information.

If database restart time is not a concern, it is better for invalid indexes to be rebuilt as part of the process of returning a database to a consistent state. When this approach is used, the time needed to restart a database will be longer due to the index recreation process; however, normal processing will not be impacted once the database has been returned to a consistent state.

On the other hand, when indexes are rebuilt as they are accessed, the time taken to restart a database is faster, but an unexpected degradation in response time may occur as a result of an index being recreated; for example, users accessing a table that has an invalid index would have to wait for the index to be rebuilt. In addition, unexpected locks may be acquired and held long after an invalid index has been recreated, especially if the transaction that caused the index recreation to occur never terminates (that is, commits or rolls back the changes made).

Dropping indexes

You cannot change any clause of an index definition; you must drop the index and create it again. (Dropping an index does not cause any other objects to be dropped but might cause some packages to be invalidated.). Use the DROP statement to drop indexes.

A primary key or unique key index cannot be explicitly dropped. You must use one of the following methods to drop it:

- If the primary index or unique constraint was created automatically for the primary key or unique key, dropping the primary key or unique key will cause the index to be dropped. Dropping is done through the ALTER TABLE statement.
- If the primary index or the unique constraint was user-defined, the primary key or unique key must be dropped first, through the ALTER TABLE statement. After the primary key or unique key is dropped, the index is no longer considered the primary index or unique index, and it can be explicitly dropped.

To drop an index using the command line, enter:

```
DROP INDEX <index_name>
```

The following statement drops the index called PH:

```
DROP INDEX PH
```

Any packages and cached dynamic SQL and XQuery statements that depend on the dropped indexes are marked invalid. The application program is not affected by changes resulting from adding or dropping indexes.

Chapter 14. Triggers

A *trigger* defines a set of actions that are performed in response to an insert, update, or delete operation on a specified table. When such an SQL operation is executed, the trigger is said to have been *activated*. Triggers are optional and are defined using the CREATE TRIGGER statement.

Triggers can be used, along with referential constraints and check constraints, to enforce data integrity rules. Triggers can also be used to cause updates to other tables, automatically generate or transform values for inserted or updated rows, or invoke functions to perform tasks such as issuing alerts.

Triggers are a useful mechanism for defining and enforcing *transitional* business rules, which are rules that involve different states of the data (for example, a salary that cannot be increased by more than 10 percent).

Using triggers places the logic that enforces business rules inside the database. This means that applications are not responsible for enforcing these rules. Centralized logic that is enforced on all of the tables means easier maintenance, because changes to application programs are not required when the logic changes.

The following are specified when creating a trigger:

- The *subject table* specifies the table for which the trigger is defined.
- The *trigger event* defines a specific SQL operation that modifies the subject table. The event can be an insert, update, or delete operation.
- The *trigger activation time* specifies whether the trigger should be activated before or after the trigger event occurs.

The statement that causes a trigger to be activated includes a *set of affected rows*. These are the rows of the subject table that are being inserted, updated, or deleted. The *trigger granularity* specifies whether the actions of the trigger are performed once for the statement or once for each of the affected rows.

The *triggered action* consists of an optional search condition and a set of statements that are executed whenever the trigger is activated. The statements are only executed if the search condition evaluates to true. If the trigger activation time is before the trigger event, triggered actions can include statements that select, set transition variables, or signal SQL states. If the trigger activation time is after the trigger event, triggered actions can include statements that select, insert, update, delete, or signal SQL states.

The triggered action can refer to the values in the set of affected rows using *transition variables*. Transition variables use the names of the columns in the subject table, qualified by a specified name that identifies whether the reference is to the old value (before the update) or the new value (after the update). The new value can also be changed using the SET Variable statement in before, insert, or update triggers.

Another means of referring to the values in the set of affected rows is to use *transition tables*. Transition tables also use the names of the columns in the subject table, but specify a name to allow the complete set of affected rows to be treated as

a table. Transition tables can only be used in AFTER triggers (that is, not with BEFORE and INSTEAD OF triggers), and separate transition tables can be defined for old and new values.

Multiple triggers can be specified for a combination of table, event (INSERT, UPDATE, DELETE, INSTEAD OF), or activation time (BEFORE, AFTER). When more than one trigger exists for a particular table, event, and activation time, the order in which the triggers are activated is the same as the order in which they were created. Thus, the most recently created trigger is the last trigger to be activated.

The activation of a trigger might cause *trigger cascading*, which is the result of the activation of one trigger that executes statements that cause the activation of other triggers or even the same trigger again. The triggered actions might also cause updates resulting from the application of referential integrity rules for deletions that can, in turn, result in the activation of additional triggers. With trigger cascading, a chain of triggers and referential integrity delete rules can be activated, causing significant change to the database as a result of a single INSERT, UPDATE, or DELETE statement.

When multiple triggers have insert, update, or delete actions against the same object, conflict resolution mechanism, like temporary tables, are used to resolve access conflicts, and this can have a noticeable impact on performance, particularly in partitioned database environments.

Types of triggers

A *trigger* defines a set of actions that are performed in response to an insert, update, or delete operation on a specified table. When such an SQL operation is executed, the trigger is said to have been *activated*. Triggers are optional and are defined using the CREATE TRIGGER statement.

Triggers can be used, along with referential constraints and check constraints, to enforce data integrity rules. Triggers can also be used to cause updates to other tables, automatically generate or transform values for inserted or updated rows, or invoke functions to perform tasks such as issuing alerts.

The following types of triggers are supported:

BEFORE triggers

Run before an update, or insert. Values that are being updated or inserted can be modified before the database is actually modified. You can use triggers that run before an update or insert in several ways:

- To check or modify values before they are actually updated or inserted in the database. This is useful if you need to transform data from the way the user sees it to some internal database format.
- To run other non-database operations coded in user-defined functions.

BEFORE DELETE triggers

Run before a delete. Checks values (a raises an error, if necessary).

AFTER triggers

Run after an update, insert, or delete. You can use triggers that run after an update or insert in several ways:

- To update data in other tables. This capability is useful for maintaining relationships between data or in keeping audit trail information.

- To check against other data in the table or in other tables. This capability is useful to ensure data integrity when referential integrity constraints aren't appropriate, or when table check constraints limit checking to the current table only.
- To run non-database operations coded in user-defined functions. This capability is useful when issuing alerts or to update information outside the database.

INSTEAD OF triggers

Describe how to perform insert, update, and delete operations against views that are too complex to support these operations natively. They allow applications to use a view as the sole interface for all SQL operations (insert, delete, update and select).

BEFORE triggers

By using triggers that run before an update or insert, values that are being updated or inserted can be modified before the database is actually modified. These can be used to transform input from the application (user view of the data) to an internal database format where desired.

These BEFORE triggers can also be used to cause other non-database operations to be activated through user-defined functions.

BEFORE DELETE triggers run before a delete operation. They check the values and raise an error, if necessary.

Examples

The following example defines a DELETE TRIGGER with a complex default:

```
CREATE TRIGGER trigger1
  BEFORE UPDATE ON table1
  REFERENCING NEW AS N
  WHEN (N.expected_delivery_date IS NULL)
  SET N.expected_delivery_date = N.order_date + 5 days;
```

The following example defines a DELETE TRIGGER with a cross table constraint that is not a referential integrity constraint:

```
CREATE TRIGGER trigger2
  BEFORE UPDATE ON table2
  REFERENCING NEW AS N
  WHEN (n.salary > (SELECT maxsalary FROM salaryguide WHERE rank = n.position))
  SIGNAL SQLSTATE '78000' SET MESSAGE_TEXT = 'Salary out of range';
```

AFTER triggers

Triggers that run after an update, insert, or delete can be used in several ways.

- Triggers can update, insert, or delete data in the same or other tables. This is useful to maintain relationships between data or to keep audit trail information.
- Triggers can check data against values of data in the rest of the table or in other tables. This is useful when you cannot use referential integrity constraints or check constraints because of references to data from other rows from this or other tables.
- Triggers can use user-defined functions to activate non-database operations. This is useful, for example, for issuing alerts or updating information outside the database.

Example

The following example presents an AFTER trigger that increases the number of employees when a new employee is hired.

```
CREATE TRIGGER NEW_HIRE
  AFTER INSERT ON EMPLOYEE
  FOR EACH ROW
  UPDATE COMPANY_STATS SET NBEMP = NBEMP + 1
```

INSTEAD OF triggers

INSTEAD OF triggers describe how to perform insert, update, and delete operations against complex views. INSTEAD OF triggers allow applications to use a view as the sole interface for all SQL operations (insert, delete, update and select).

Usually, INSTEAD OF triggers contain the inverse of the logic applied in a view body. For example, consider a view that decrypts columns from its source table. The INSTEAD OF trigger for this view encrypts data and then inserts it into the source table, thus performing the symmetrical operation.

Using an INSTEAD OF trigger, the requested modify operation against the view gets replaced by the trigger logic, which performs the operation on behalf of the view. From the perspective of the application this happens transparently, as it perceives that all operations are performed against the view. Only one INSTEAD OF trigger is allowed for each kind of operation on a given subject view.

The view itself must be an untyped view or an alias that resolves to an untyped view. Also, it cannot be a view that is defined using WITH CHECK OPTION (a symmetric view) or a view on which a symmetric view has been defined directly or indirectly.

Example

The following example presents three INSTEAD OF triggers that provide logic for INSERTs, UPDATEs, and DELETEs to the defined view (EMPV). The view EMPV contains a join in its from clause and therefore cannot natively support any modify operation.

```
CREATE VIEW EMPV(EMPNO, FIRSTNME, MIDINIT, LASTNAME, PHONENO,
  HIREDATE, DEPTNAME)
AS SELECT EMPNO, FIRSTNME, MIDINIT, LASTNAME, PHONENO,
  HIREDATE, DEPTNAME
  FROM EMPLOYEE, DEPARTMENT WHERE
    EMPLOYEE.WORKDEPT = DEPARTMENT.DEPTNO

CREATE TRIGGER EMPV_INSERT INSTEAD OF INSERT ON EMPV
REFERENCING NEW AS NEWEMP FOR EACH ROW
INSERT INTO EMPLOYEE (EMPNO, FIRSTNME, MIDINIT, LASTNAME,
  WORKDEPT, PHONENO, HIREDATE)
VALUES(EMPNO, FIRSTNME, MIDINIT, LASTNAME,
  COALESCE((SELECT DEPTNO FROM DEPARTMENT AS D
    WHERE D.DEPTNAME = NEWEMP.DEPTNAME),
    RAISE_ERROR('70001', 'Unknown dept name')),
  PHONENO, HIREDATE)

CREATE TRIGGER EMPV_UPDATE INSTEAD OF UPDATE ON EMPV
REFERENCING NEW AS NEWEMP OLD AS OLDEMP
FOR EACH ROW
BEGIN ATOMIC
VALUES(CASE WHEN NEWEMP.EMPNO = OLDEMP.EMPNO THEN 0
```

```

        ELSE RAISE_ERROR('70002', 'Must not change EMPNO') END);
UPDATE EMPLOYEE AS E
  SET (FIRSTNME, MIDINIT, LASTNAME, WORKDEPT, PHONENO, HIREDATE)
    = (NEWEMP.FIRSTNME, NEWEMP.MIDINIT, NEWEMP.LASTNAME,
      COALESCE((SELECT DEPTNO FROM DEPARTMENT AS D
                WHERE D.DEPTNAME = NEWEMP.DEPTNAME),
              RAISE_ERROR ('70001', 'Unknown dept name')),
      NEWEMP.PHONENO, NEWEMP.HIREDATE)
  WHERE NEWEMP.EMPNO = E.EMPNO;
END

CREATE TRIGGER EMPV_DELETE INSTEAD OF DELETE ON EMPV
REFERENCING OLD AS OLDEMP FOR EACH ROW
DELETE FROM EMPLOYEE AS E WHERE E.EMPNO = OLDEMP.EMPNO

```

Designing triggers

When creating a trigger, you must associate it with a table; when creating an `INSTEAD OF` trigger, you must associate it with a view. This table or view is called the *target table* of the trigger. The term modify operation refers to any change in the state of the target table.

A *modify operation* is initiated by:

- an `INSERT` statement
- an `UPDATE` statement, or a referential constraint which performs an `UPDATE`
- a `DELETE` statement, or a referential constraint which performs a `DELETE`
- a `MERGE` statement

You must associate each trigger with one of these three types of modify operations. The association is called the *trigger event* for that particular trigger.

You must also define the action, called the *triggered action*, that the trigger performs when its trigger event occurs. The triggered action consists of one or more statements which can execute either before or after the database manager performs the trigger event. Once a trigger event occurs, the database manager determines the set of rows in the subject table that the modify operation affects and executes the trigger.

Guidelines when creating triggers:

When creating a trigger, you must declare the following attributes and behavior:

- The name of the trigger.
- The name of the subject table.
- The trigger activation time (`BEFORE` or `AFTER` the modify operation executes).
- The trigger event (`INSERT`, `DELETE`, or `UPDATE`).
- The old transition variable value, if any.
- The new transition variable value, if any.
- The old transition table value, if any.
- The new transition table value, if any.
- The granularity (`FOR EACH STATEMENT` or `FOR EACH ROW`).
- The triggered action of the trigger (including a triggered action condition and triggered statement(s)).
- If the trigger event is `UPDATE` a trigger-column list if the trigger should only fire when specific columns are specified in the update statement.

Designing multiple triggers:

When triggers are defined using the CREATE TRIGGER statement, their creation time is registered in the database in form of a timestamp. The value of this timestamp is subsequently used to order the activation of triggers when there is more than one trigger that should be run at the same time. For example, the timestamp is used when there is more than one trigger on the same subject table with the same event and the same activation time. The timestamp is also used when there are one or more AFTER or INSTEAD OF triggers that are activated by the trigger event and referential constraint actions caused directly or indirectly (that is, recursively by other referential constraints) by the triggered action.

Consider the following two triggers:

```
CREATE TRIGGER NEW_HIRED
  AFTER INSERT ON EMPLOYEE
  FOR EACH ROW
  BEGIN ATOMIC
    UPDATE COMPANY_STATS
    SET NBEMP = NBEMP + 1;
  END

CREATE TRIGGER NEW_HIRED_DEPT
  AFTER INSERT ON EMPLOYEE
  REFERENCING NEW AS EMP
  FOR EACH ROW
  BEGIN ATOMIC
    UPDATE DEPTS
    SET NBEMP = NBEMP + 1
    WHERE DEPT_ID = EMP.DEPT_ID;
  END
```

The above triggers are activated when you run an INSERT operation on the employee table. In this case, the timestamp of their creation defines which of the above two triggers is activated first.

The activation of the triggers is conducted in ascending order of the timestamp value. Thus, a trigger that is newly added to a database runs after all the other triggers that are previously defined.

Old triggers are activated before new triggers to ensure that new triggers can be used as *incremental* additions to the changes that affect the database. For example, if a triggered statement of trigger T1 inserts a new row into a table T, a triggered statement of trigger T2 that is run after T1 can be used to update the same row in T with specific values. Because the activation order of triggers is predictable, you can have multiple triggers on a table and still know that the newer ones will be acting on a table that has already been modified by the older ones.

Trigger interactions with referential constraints:

A trigger event can occur as a result of changes due to referential constraint enforcement. For example, given two tables DEPT and EMP, if deleting or updating DEPT causes propagated deletes or updates to EMP by means of referential integrity constraints, then delete or update triggers defined on EMP become activated as a result of the referential constraint defined on DEPT. The triggers on EMP are run either BEFORE or AFTER the deletion (in the case of ON DELETE CASCADE) or update of rows in EMP (in the case of ON DELETE SET NULL), depending on their activation time.

Specifying what makes a trigger fire (triggering statement or event)

Every trigger is associated with an event. Triggers are activated when their corresponding event occurs in the database. This trigger event occurs when the specified action, either an UPDATE, INSERT, or DELETE statement (including those caused by actions of referential constraints), is performed on the target table.

For example:

```
CREATE TRIGGER NEW_HIRE
AFTER INSERT ON EMPLOYEE
FOR EACH ROW
UPDATE COMPANY_STATS SET NBEMP = NBEMP + 1
```

The above statement defines the trigger new_hire, which activates when you perform an insert operation on table employee.

You associate every trigger event, and consequently every trigger, with exactly one target table and exactly one modify operation. The modify operations are:

Insert operation

An insert operation can only be caused by an INSERT statement or the insert operation of a MERGE statement. Therefore, triggers are not activated when data is loaded using utilities that do not use INSERT, such as the LOAD command.

Delete operation

A delete operation can be caused by a DELETE statement, or the delete operation of a MERGE statement, or as a result of a referential constraint rule of ON DELETE CASCADE.

Update operation

An update operation can be caused by an UPDATE statement, or the update operation of a MERGE statement, or as a result of a referential constraint rule of ON DELETE SET NULL.

If the trigger event is an update operation, the event can be associated with specific columns of the target table. In this case, the trigger is only activated if the update operation attempts to update any of the specified columns. This provides a further refinement of the event that activates the trigger.

For example, the following trigger, REORDER, activates only if you perform an update operation on the columns ON_HAND or MAX_STOCKED, of the table PARTS:

```
CREATE TRIGGER REORDER
AFTER UPDATE OF ON_HAND, MAX_STOCKED ON PARTS
REFERENCING NEW AS N_ROW
FOR EACH ROW
WHEN (N_ROW.ON_HAND < 0.10 * N_ROW.MAX_STOCKED)
BEGIN ATOMIC
VALUES(ISSUE_SHIP_REQUEST(N_ROW.MAX_STOCKED -
                          N_ROW.ON_HAND,
                          N_ROW.PARTNO));
END
```

When a trigger is activated, it runs according to its level of granularity as follows:

FOR EACH ROW

It runs as many times as the number of rows in the set of affected rows. If you need to refer to the specific rows affected by the triggered action, use

FOR EACH ROW granularity. An example of this is the comparison of the new and old values of an updated row in an AFTER UPDATE trigger.

FOR EACH STATEMENT

It runs once for the entire trigger event.

If the set of affected rows is empty (that is, in the case of a searched UPDATE or DELETE in which the WHERE clause did not qualify any rows), a FOR EACH ROW trigger does not run. But a FOR EACH STATEMENT trigger still runs once.

For example, keeping a count of number of employees can be done using FOR EACH ROW.

```
CREATE TRIGGER NEW_HIRED
AFTER INSERT ON EMPLOYEE
FOR EACH ROW
UPDATE COMPANY_STATS SET NBEMP = NBEMP + 1
```

You can achieve the same affect with one update by using a granularity of FOR EACH STATEMENT.

```
CREATE TRIGGER NEW_HIRED
AFTER INSERT ON EMPLOYEE
REFERENCING NEW TABLE AS NEWEMPS
FOR EACH STATEMENT
UPDATE COMPANY_STATS
SET NBEMP = NBEMP + (SELECT COUNT(*) FROM NEWEMPS)
```

Note:

- A granularity of FOR EACH STATEMENT is not supported for BEFORE triggers.
- The maximum nesting level for triggers is 16. That is, the maximum number of cascading trigger activations is 16. A trigger activation refers to the activation of a trigger upon a triggering event, such as insert, update, or delete of data in a column of a table, or generally to a table.

Specifying when a trigger fires (BEFORE, AFTER, and INSTEAD OF clauses)

The *trigger activation time* specifies when the trigger should be activated, relative to the trigger event.

There are three activation times that you can specify: BEFORE, AFTER, or INSTEAD OF:

- If the activation time is BEFORE, the triggered actions are activated for each row in the set of affected rows before the trigger event executes. Hence, the subject table will only be modified after the BEFORE trigger has completed execution for each row. Note that BEFORE triggers must have a granularity of FOR EACH ROW.
- If the activation time is AFTER, the triggered actions are activated for each row in the set of affected rows or for the statement, depending on the trigger granularity. This occurs after the trigger event has been completed, and after the database manager checks all constraints that the trigger event might affect, including actions of referential constraints. Note that AFTER triggers can have a granularity of either FOR EACH ROW or FOR EACH STATEMENT.

For example, the activation time of the following trigger is AFTER the INSERT operation on employee:

```

CREATE TRIGGER NEW_HIRE
AFTER INSERT ON EMPLOYEE
FOR EACH ROW
UPDATE COMPANY_STATS SET NBEMP = NBEMP + 1

```

- If the activation time is **INSTEAD OF**, the triggered actions are activated for each row in the set of affected rows instead of executing the trigger event. **INSTEAD OF** triggers must have a granularity of **FOR EACH ROW**, and the subject table must be a view. No other triggers are able to use a view as the subject table.

The following diagram illustrates the execution model of **BEFORE** and **AFTER** triggers:




Figure 26. Trigger execution model

For a given table with both before and AFTER triggers, and a modifying event that is associated with these triggers, all the BEFORE triggers are activated first. The first activated BEFORE trigger for a given event operates on the set of rows targeted by the operation and makes any update modifications to the set that its logic prescribes. The output of this BEFORE trigger is accepted as input by the next before-trigger. When all of the BEFORE triggers that are activated by the

event have been fired, the intermediate result set, the result of the BEFORE trigger modifications to the rows targeted by the trigger event operation, is applied to the table. Then each AFTER trigger associated with the event is fired. The AFTER triggers might modify the same table, another table, or perform an action external to the database.

The different activation times of triggers reflect different purposes of triggers. Basically, BEFORE triggers are an extension to the constraint subsystem of the database management system. Therefore, you generally use them to:

- Perform validation of input data
- Automatically generate values for newly inserted rows
- Read from other tables for cross-referencing purposes

BEFORE triggers are not used for further modifying the database because they are activated before the trigger event is applied to the database. Consequently, they are activated before integrity constraints are checked.

Conversely, you can view AFTER triggers as a module of application logic that runs in the database every time a specific event occurs. As a part of an application, AFTER triggers always see the database in a consistent state. Note that they are run after the integrity constraint validations. Consequently, you can use them mostly to perform operations that an application can also perform. For example:

- Perform follow on modify operations in the database.
- Perform actions outside the database, for example, to support alerts. Note that actions performed outside the database are not rolled back if the trigger is rolled back.

In contrast, you can view an INSTEAD OF trigger as a description of the inverse operation of the view it is defined on. For example, if the select list in the view contains an expression over a table, the INSERT statement in the body of its INSTEAD OF INSERT trigger will contain the reverse expression.

Because of the different nature of BEFORE, AFTER, and INSTEAD OF triggers, a different set of SQL operations can be used to define the triggered actions of BEFORE and AFTER, INSTEAD OF triggers. For example, update operations are not allowed in BEFORE triggers because there is no guarantee that integrity constraints will not be violated by the triggered action. Similarly, different trigger granularities are supported in BEFORE, AFTER, and INSTEAD OF triggers.

The triggered SQL statement of all triggers can be a dynamic compound statement. However, BEFORE triggers face some restrictions; they cannot contain the following SQL statements:

- UPDATE
- DELETE
- INSERT
- MERGE

Defining conditions for when trigger-action will fire (WHEN clause)

The activation of a trigger results in the running of its associated triggered action. Every trigger has exactly one triggered action which, in turn, has two components: an optional *triggered action condition* or WHEN clause, and a set of *triggered statement(s)*.

The *triggered action condition* is an optional clause of the triggered action which specifies a search condition that must evaluate to true to run statements within the triggered action. If the WHEN clause is omitted, then the statements within the triggered action are always executed.

The triggered action condition is evaluated once for each row if the trigger is a FOR EACH ROW trigger, and once for the statement if the trigger is a FOR EACH STATEMENT trigger.

This clause provides further control that you can use to fine tune the actions activated on behalf of a trigger. An example of the usefulness of the WHEN clause is to enforce a data dependent rule in which a triggered action is activated only if the incoming value falls inside or outside of a certain range.

The activation of a trigger results in the running of its associated triggered action. Every trigger has exactly one triggered action which, in turn, has two components:

The triggered action condition defines whether or not the set of triggered statements are performed for the row or for the statement for which the triggered action is executing. The set of triggered statements define the set of actions performed by the trigger in the database as a consequence of its event having occurred.

For example, the following trigger action specifies that the set of triggered statements should only be activated for rows in which the value of the on_hand column is less than ten per cent of the value of the max_stocked column. In this case, the set of triggered statements is the invocation of the issue_ship_request function.

```
CREATE TRIGGER REORDER
  AFTER UPDATE OF ON_HAND, MAX_STOCKED ON PARTS
  REFERENCING NEW AS N_ROW
  FOR EACH ROW

  WHEN (N_ROW.ON_HAND < 0.10 * N_ROW.MAX_STOCKED)
  BEGIN ATOMIC
    VALUES(ISSUE_SHIP_REQUEST(N_ROW.MAX_STOCKED -
                               N_ROW.ON_HAND,
                               N_ROW.PARTNO));
  END
```

The set of triggered statements carries out the real actions caused by activating a trigger. Not every SQL operation is meaningful in every trigger. Depending on whether the trigger activation time is BEFORE or AFTER, different kinds of operations might be appropriate as a triggered statement.

In most cases, if any triggered statement returns a negative return code, the triggering statement together with all trigger and referential constraint actions are rolled back. The trigger name, SQLCODE, SQLSTATE and many of the tokens from the failing triggered statement are returned in the error message.

Supported SQL PL statements in triggers

The triggered SQL statement of all triggers can be a dynamic compound statement.

That is, triggered SQL statements can contain one or more of the following elements:

- CALL statement

- DECLARE variable statement
- SET variable statement
- WHILE loop
- FOR loop
- IF statement
- SIGNAL statement
- ITERATE statement
- LEAVE statement
- GET DIAGNOSTIC statement
- fullselect

However, only AFTER and INSTEAD OF triggers can contain one or more of the following SQL statements:

- UPDATE statement
- DELETE statement
- INSERT statement
- MERGE statement

Accessing old and new column values in triggers using transition variables

When you implement a FOR EACH ROW trigger, it might be necessary to refer to the value of columns of the row in the set of affected rows, for which the trigger is currently executing. Note that to refer to columns in tables in the database (including the subject table), you can use regular SELECT statements.

A FOR EACH ROW trigger can refer to the columns of the row for which it is currently executing by using two transition variables that you can specify in the REFERENCING clause of a CREATE TRIGGER statement. There are two kinds of transition variables, which are specified as OLD and NEW, together with a correlation-name. They have the following semantics:

OLD AS correlation-name

Specifies a correlation name which captures the original state of the row, that is, before the triggered action is applied to the database.

NEW AS correlation-name

Specifies a correlation name which captures the value that is, or was, used to update the row in the database when the triggered action is applied to the database.

Consider the following example:

```
CREATE TRIGGER REORDER
AFTER UPDATE OF ON_HAND, MAX_STOCKED ON PARTS
REFERENCING NEW AS N_ROW
FOR EACH ROW
WHEN (N_ROW.ON_HAND < 0.10 * N_ROW.MAX_STOCKED
AND N_ROW.ORDER_PENDING = 'N')
BEGIN ATOMIC
VALUES(ISSUE_SHIP_REQUEST(N_ROW.MAX_STOCKED -
                          N_ROW.ON_HAND,
                          N_ROW.PARTNO));
UPDATE PARTS SET PARTS.ORDER_PENDING = 'Y'
WHERE PARTS.PARTNO = N_ROW.PARTNO;
END
```


Based on the definition of the OLD and NEW transition variables given above, it is clear that not every transition variable can be defined for every trigger. Transition variables can be defined depending on the kind of trigger event:

UPDATE

An UPDATE trigger can refer to both OLD and NEW transition variables.

INSERT

An INSERT trigger can only refer to a NEW transition variable because before the activation of the INSERT operation, the affected row does not exist in the database. That is, there is no original state of the row that would define old values before the triggered action is applied to the database.

DELETE

A DELETE trigger can only refer to an OLD transition variable because there are no new values specified in the delete operation.

Note: Transition variables can only be specified for FOR EACH ROW triggers. In a FOR EACH STATEMENT trigger, a reference to a transition variable is not sufficient to specify to which of the several rows in the set of affected rows the transition variable is referring. Instead, refer to the set of new and old rows by using the OLD TABLE and NEW TABLE clauses of the CREATE TRIGGER statement. For more information on these clauses, see the CREATE TRIGGER statement.

Referencing old and new table result sets using transition tables

In both FOR EACH ROW and FOR EACH STATEMENT triggers, it might be necessary to refer to the whole set of affected rows. This is necessary, for example, if the trigger body needs to apply aggregations over the set of affected rows (for example, MAX, MIN, or AVG of some column values).

A trigger can refer to the set of affected rows by using two transition tables that can be specified in the REFERENCING clause of a CREATE TRIGGER statement. Just like the transition variables, there are two kinds of transition tables, which are specified as OLD_TABLE and NEW_TABLE together with a table-name, with the following semantics:

OLD_TABLE AS table-name

Specifies the name of the table which captures the original state of the set of affected rows (that is, before the triggering SQL operation is applied to the database).

NEW_TABLE AS table-name

Specifies the name of the table which captures the value that is used to update the rows in the database when the triggered action is applied to the database.

For example:

```
CREATE TRIGGER REORDER
AFTER UPDATE OF ON_HAND, MAX_STOCKED ON PARTS
REFERENCING NEW_TABLE AS N_TABLE
NEW AS N_ROW
FOR EACH ROW
WHEN ((SELECT AVG (ON_HAND) FROM N_TABLE) > 35)
BEGIN ATOMIC
```

```

VALUES(INFORM_SUPERVISOR(N_ROW.PARTNO,
                          N_ROW.MAX_STOCKED,
                          N_ROW.ON_HAND));
END

```

Note that `NEW_TABLE` always has the full set of updated rows, even on a `FOR EACH ROW` trigger. When a trigger acts on the table on which the trigger is defined, `NEW_TABLE` contains the changed rows from the statement that activated the trigger. However, `NEW_TABLE` does not contain the changed rows that were caused by statements within the trigger, as that would cause a separate activation of the trigger.

The transition tables are read-only. The same rules that define the kinds of transition variables that can be defined for which trigger event, apply for transition tables:

UPDATE

An `UPDATE` trigger can refer to both `OLD_TABLE` and `NEW_TABLE` transition tables.

INSERT

An `INSERT` trigger can only refer to a `NEW_TABLE` transition table because before the activation of the `INSERT` operation the affected rows do not exist in the database. That is, there is no original state of the rows that defines old values before the triggered action is applied to the database.

DELETE

A `DELETE` trigger can only refer to an `OLD_TABLE` transition table because there are no new values specified in the delete operation.

Note: It is important to observe that transition tables can be specified for both granularities of `AFTER` triggers: `FOR EACH ROW` and `FOR EACH STATEMENT`.

The scope of the `OLD_TABLE` and `NEW_TABLE` table-name is the trigger body. In this scope, this name takes precedence over the name of any other table with the same unqualified *table-name* that might exist in the schema. Therefore, if the `OLD_TABLE` or `NEW_TABLE` table-name is for example, `X`, a reference to `X` (that is, an unqualified `X`) in the `FROM` clause of a `SELECT` statement will always refer to the transition table even if there is a table named `X` in the in the schema of the trigger creator. In this case, the user has to make use of the fully qualified name in order to refer to the table `X` in the schema.

Creating triggers

A trigger defines a set of actions that are executed in conjunction with, or triggered by, an `INSERT`, `UPDATE`, or `DELETE` clause on a specified table or a typed table.

Use triggers to:

- Validate input data
- Generate a value for a newly-inserted row
- Read from other tables for cross-referencing purposes
- Write to other tables for audit-trail purposes

You can use triggers to support general forms of integrity or business rules. For example, a trigger can check a customer's credit limit before an order is accepted or update a summary data table.

Benefits:

- **Faster application development:** Because a trigger is stored in the database, you do not have to code the actions that it performs in every application.
- **Easier maintenance:** Once a trigger is defined, it is automatically invoked when the table that it is created on is accessed.
- **Global enforcement of business rules:** If a business policy changes, you only need to change the trigger and not each application program.

Restrictions:

- You cannot use triggers with nicknames.
- If the trigger is a BEFORE trigger, the column name specified by the triggered action must not be a generated column other than an identity column. That is, the generated identity value is visible to BEFORE triggers.

When creating an atomic trigger, care must be taken with the end-of-statement character. The command line processor, by default, considers a ";" the end-of-statement marker. You should manually edit the end-of-statement character in your script to create the atomic trigger so that you are using a character other than ";". For example, the ";" could be replaced by another special character like "#". You can also precede the CREATE TRIGGER DDL with:

```
--#SET TERMINATOR @
```

To change the terminator in the CLP on the fly, the following syntax will set it back:

```
--#SET TERMINATOR
```

To create a trigger from the command line, enter:

```
db2 -td <delimiter> -vf <script>
```

where the <delimiter> is the alternative end-of-statement character and the <script> is the modified script with the new <delimiter> in it.

To create a trigger from the command line, enter:

```
CREATE TRIGGER <name>  
<action> ON <table_name>  
<operation>  
<triggered_action>
```

The following statement creates a trigger that increases the number of employees each time a new person is hired, by adding 1 to the number of employees (NBEMP) column in the COMPANY_STATS table each time a row is added to the EMPLOYEE table.

```
CREATE TRIGGER NEW_HIRED  
AFTER INSERT ON EMPLOYEE  
FOR EACH ROW  
UPDATE COMPANY_STATS SET NBEMP = NBEMP+1;
```

A trigger body can include one or more of the following statements: INSERT, searched UPDATE, searched DELETE, fullselect, SET Variable, and SIGNAL SQLSTATE. The trigger can be activated before or after the INSERT, UPDATE, or DELETE statement to which it refers.

Modifying and dropping triggers

Triggers cannot be modified. They must be dropped and then created again according to the new definitions you require.

Trigger dependencies

- All dependencies of a trigger on some other object are recorded in the SYSCAT.TRIGDEP system catalog view. A trigger can depend on many objects.
- If an object that a trigger is dependent on is dropped, the trigger becomes inoperative but its definition is retained in the system catalog view. To re-validate this trigger, you must retrieve its definition from the system catalog view and submit a new CREATE TRIGGER statement.
- If a trigger is dropped, its description is deleted from the SYSCAT.TRIGGERS system catalog view and all of its dependencies are deleted from the SYSCAT.TRIGDEP system catalog view. All packages having UPDATE, INSERT, or DELETE dependencies on the trigger are invalidated.
- If the view is dependent on the trigger and it is made inoperative, the trigger is also marked inoperative. Any packages dependent on triggers that have been marked inoperative are invalidated.

A trigger object can be dropped using the DROP TRIGGER statement, but this procedure will cause dependent packages to be marked invalid, as follows:

- If an update trigger without an explicit column list is dropped, then packages with an update usage on the target table are invalidated.
- If an update trigger with a column list is dropped, then packages with update usage on the target table are only invalidated if the package also had an update usage on at least one column in the column-name list of the CREATE TRIGGER statement.
- If an insert trigger is dropped, packages that have an insert usage on the target table are invalidated.
- If a delete trigger is dropped, packages that have a delete usage on the target table are invalidated.

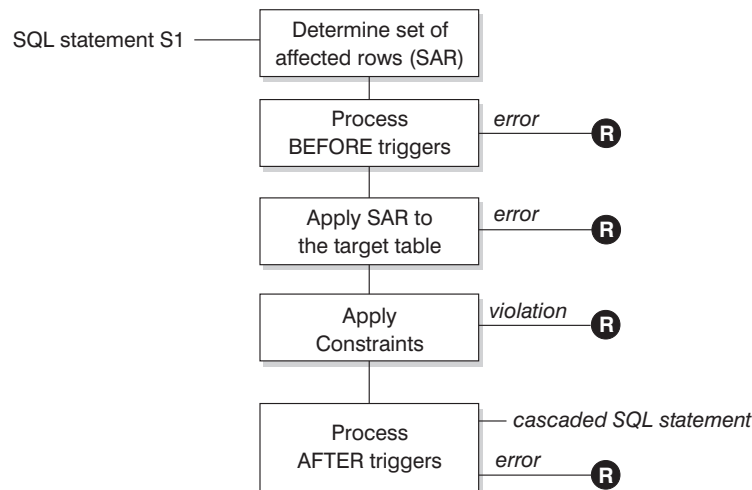
A package remains invalid until the application program is explicitly bound or rebound, or it is run and the database manager automatically rebinds it.

Examples of triggers and trigger use

Examples of interaction between triggers and referential constraints

Update operations may cause the interaction of triggers with referential constraints and check constraints.

Figure 23 on page 278 and the associated description are representative of the processing that is performed for an statement that updates data in the database.



R = rollback changes to before S1

Figure 27. Processing an statement with associated triggers and constraints

Figure 23 on page 278 shows the general order of processing for an statement that updates a table. It assumes a situation where the table includes BEFORE triggers, referential constraints, check constraints and AFTER triggers that cascade. The following is a description of the boxes and other items found in Figure 23 on page 278.

- statement S_1

This is the DELETE, INSERT, or UPDATE statement that begins the process. The statement S_1 identifies a table (or an updatable view over some table) referred to as the *subject table* throughout this description.

- Determine set of affected rows

This step is the starting point for a process that repeats for referential constraint delete rules of CASCADE and SET NULL and for cascaded statements from AFTER triggers.

The purpose of this step is to determine the *set of affected rows* for the statement. The set of rows included is based on the statement:

- for DELETE, all rows that satisfy the search condition of the statement (or the current row for a positioned DELETE)
- for INSERT, the rows identified by the VALUES clause or the fullselect
- for UPDATE, all rows that satisfy the search condition (or the current row for a positioned UPDATE).

If the set of affected rows is empty, there will be no BEFORE triggers, changes to apply to the subject table, or constraints to process for the statement.

- Process BEFORE triggers

All BEFORE triggers are processed in ascending order of creation. Each BEFORE trigger will process the triggered action once for each row in the set of affected rows.

An error may occur during the processing of a triggered action in which case all changes made as a result of the original statement S_1 (so far) are rolled back.

If there are no BEFORE triggers or the set of affected is empty, this step is skipped.

- Apply the set of affected rows to the subject table

The actual delete, insert, or update is applied using the set of affected rows to the subject table in the database.

An error may occur when applying the set of affected rows (such as attempting to insert a row with a duplicate key where a unique index exists) in which case all changes made as a result of the original statement S_1 (so far) are rolled back.

- Apply Constraints

The constraints associated with the subject table are applied if set of affected rows is not empty. This includes unique constraints, unique indexes, referential constraints, check constraints and checks related to the WITH CHECK OPTION on views. Referential constraints with delete rules of cascade or set null may cause additional triggers to be activated.

A violation of any constraint or WITH CHECK OPTION results in an error and all changes made as a result of S_1 (so far) are rolled back.

- Process AFTER triggers

All AFTER triggers activated by S_1 are processed in ascending order of creation. FOR EACH STATEMENT triggers will process the triggered action exactly once, even if the set of affected rows is empty. FOR EACH ROW triggers will process the triggered action once for each row in the set of affected rows.

An error may occur during the processing of a triggered action in which case all changes made as a result of the original S_1 (so far) are rolled back.

The triggered action of a trigger may include triggered statements that are DELETE, INSERT or UPDATE statements. For the purposes of this description, each such statement is considered a *cascaded statement*.

A cascaded statement is a DELETE, INSERT, or UPDATE statement that is processed as part of the triggered action of an AFTER trigger. This statement starts a cascaded level of trigger processing. This can be thought of as assigning the triggered statement as a new S_1 and performing all of the steps described here recursively.

Once all triggered statements from all AFTER triggers activated by each S_1 have been processed to completion, the processing of the original S_1 is completed.

- R = roll back changes to before S_1

Any error (including constraint violations) that occurs during processing results in a roll back of all the changes made directly or indirectly as a result of the original statement S_1 . The database is therefore back in the same state as immediately prior to the execution of the original statement S_1

Examples of defining actions using triggers

Assume that your general manager wants to keep the names of customers who have sent three or more complaints in the last 72 hours in a separate table. The general manager also wants to be informed whenever a customer name is inserted in this table more than once.

To define such actions, you define:

- An UNHAPPY_CUSTOMERS table:

```
CREATE TABLE UNHAPPY_CUSTOMERS (  
  NAME          VARCHAR (30),  
  EMAIL_ADDRESS VARCHAR (200),  
  INSERTION_DATE DATE)
```

- A trigger to automatically insert a row in UNHAPPY_CUSTOMERS if 3 or more messages were received in the last 3 days (assumes the existence of a CUSTOMERS table that includes a NAME column and an E_MAIL_ADDRESS column):

```

CREATE TRIGGER STORE_UNHAPPY_CUST
AFTER INSERT ON ELECTRONIC_MAIL
REFERENCING NEW AS N
FOR EACH ROW
WHEN (3 <= (SELECT COUNT(*)
           FROM ELECTRONIC_MAIL
           WHERE SENDER = N.SENDER
           AND SENDING_DATE(MESSAGE) > CURRENT DATE - 3 DAYS)
)
BEGIN ATOMIC
  INSERT INTO UNHAPPY_CUSTOMERS
  VALUES ((SELECT NAME
           FROM CUSTOMERS
           WHERE EMAIL_ADDRESS = N.SENDER), N.SENDER, CURRENT DATE);
END

```

- A trigger to send a note to the general manager if the same customer is inserted in UNHAPPY_CUSTOMERS more than once (assumes the existence of a SEND_NOTE function that takes 2 character strings as input):

```

CREATE TRIGGER INFORM_GEN_MGR
AFTER INSERT ON UNHAPPY_CUSTOMERS
REFERENCING NEW AS N
FOR EACH ROW
WHEN (1 <(SELECT COUNT(*)
         FROM UNHAPPY_CUSTOMERS
         WHERE EMAIL_ADDRESS = N.EMAIL_ADDRESS)
)
BEGIN ATOMIC
  VALUES(SEND_NOTE('Check customer:' CONCAT N.NAME,
                  'bigboss@vnet.ibm.com'));
END

```

Example of defining business rules using triggers

Suppose your company has the policy that all e-mail dealing with customer complaints must have Mr. Nelson, the marketing manager, in the carbon copy (CC) list.

Because this is a rule, you might want to express it as a constraint such as one of the following (assuming the existence of a CC_LIST UDF to check it):

```

ALTER TABLE ELECTRONIC_MAIL ADD
CHECK (SUBJECT <> 'Customer complaint' OR
      CONTAINS (CC_LIST(MESSAGE), 'nelson@vnet.ibm.com') = 1)

```

However, such a constraint prevents the insertion of e-mail dealing with customer complaints that do not have the marketing manager in the cc list. This is certainly not the intent of your company's business rule. The intent is to forward to the marketing manager any e-mail dealing with customer complaints that were not copied to the marketing manager. Such a business rule can only be expressed with a trigger because it requires taking actions that cannot be expressed with declarative constraints. The trigger assumes the existence of a SEND_NOTE function with parameters of type E_MAIL and character string.

```

CREATE TRIGGER INFORM_MANAGER
AFTER INSERT ON ELECTRONIC_MAIL
REFERENCING NEW AS N
FOR EACH ROW
WHEN (N.SUBJECT = 'Customer complaint' AND
      CONTAINS (CC_LIST(MESSAGE), 'nelson@vnet.ibm.com') = 0)
BEGIN ATOMIC
  VALUES(SEND_NOTE(N.MESSAGE, 'nelson@vnet.ibm.com'));
END

```


Example of preventing operations on tables using triggers

Suppose you want to prevent undeliverable e-mail from being stored in a table named `ELECTRONIC_MAIL`. To do so, you need to prevent the execution of certain SQL `INSERT` statements.

There are two ways to do this:

- Define a `BEFORE` trigger that returns an error whenever the subject of an e-mail is *undelivered mail*:

```
CREATE TRIGGER BLOCK_INSERT
NO CASCADE BEFORE INSERT ON ELECTRONIC_MAIL
REFERENCING NEW AS N
FOR EACH ROW
WHEN (SUBJECT(N.MESSAGE) = 'undelivered mail')
BEGIN ATOMIC
  SIGNAL SQLSTATE '85101'
  SET MESSAGE_TEXT = ('Attempt to insert undelivered mail');
END
```

- Define a check constraint forcing values of the new column `SUBJECT` to be different from *undelivered mail*:

```
ALTER TABLE ELECTRONIC_MAIL
ADD CONSTRAINT NO_UNDELIVERED
CHECK (SUBJECT <> 'undelivered mail')
```

Chapter 15. Sequences

A *sequence* is a database object that allows the automatic generation of values, such as cheque numbers. Sequences are ideally suited to the task of generating unique key values. Applications can use sequences to avoid possible concurrency and performance problems resulting from column values used to track numbers. The advantage that sequences have over numbers created outside the database is that the database server keeps track of the numbers generated. A crash and restart will not cause duplicate numbers from being generated.

The sequence numbers generated have the following properties:

- Values can be any exact numeric data type with a scale of zero. Such data types include: SMALLINT, BIGINT, INTEGER, and DECIMAL.
- Consecutive values can differ by any specified integer increment. The default increment value is 1.
- Counter value is recoverable. The counter value is reconstructed from logs when recovery is required.
- Values can be cached to improve performance. Pre-allocating and storing values in the cache reduces synchronous I/O to the log when values are generated for the sequence. In the event of a system failure, all cached values that have not been used are considered lost. The value specified for CACHE is the maximum number of sequence values that could be lost.

There are two expressions that can be used with sequences:

- **NEXT VALUE expression:** returns the next value for the specified sequence. A new sequence number is generated when a NEXT VALUE expression specifies the name of the sequence. However, if there are multiple instances of a NEXT VALUE expression specifying the same sequence name within a query, the counter for the sequence is incremented only once for each row of the result, and all instances of NEXT VALUE return the same value for each row of the result.
- **PREVIOUS VALUE expression:** returns the most recently generated value for the specified sequence for a previous statement within the current application process. That is, for any given connection, the PREVIOUS VALUE remains constant even if another connection invokes NEXT VALUE.

For complete details and examples of these expressions, see “Sequence reference” in *SQL Reference, Volume 1*.

Designing sequences

When designing sequences you need to consider the differences between identity columns and sequences, and which is more appropriate for your environment. If you decide to use sequences, you need to be familiar with the available options and parameters.

Before designing sequences, see “Sequences compared to identity columns” on page 324.

In addition to being simple to set up and create, the sequence has a variety of additional options that allows you more flexibility in generating the values:

- Choose from a variety of data types (SMALLINT, INTEGER, BIGINT, DECIMAL)
- Change starting values (START WITH)
- Change the sequence increment, including specifying increasing or decreasing values (INCREMENT BY)
- Set minimum and maximum values where the sequence numbering starts and stops (MINVALUE/MAXVALUE)
- Allow wrapping of values so that sequences can start over again, or disallow cycling (CYCLE/NO CYCLE)
- Allow caching of sequence values to improve performance, or disallow caching(CACHE/NO CACHE)

Even after the sequence has been generated, many of these values can be altered. For instance, you may want to set a different starting value depending on the day of the week. Another practical example of using sequences is for the generation and processing of bank checks. The sequence of bank check numbers is extremely important, and there are serious consequences if a batch of sequence numbers is lost or corrupted.

For improved performance, you should also be aware of and make use of the CACHE option. This option tells the database manager how many sequence values should be generated by the system before going back to the catalog to generate another set of sequences. The default CACHE value is 20, if not specified. Using the default as an example, the database manager automatically generates 20 sequential values in memory (1, 2, ..., 20) when the first sequence value is requested. Whenever a new sequence number is required, this memory cache of values is used to return the next value. Once this cache of values is used up, the database manager will generate the next twenty values (21, 22, ..., 40).

By implementing caching of sequence numbers, the database manager does not have to continually go to the catalog tables to get the next value. This reduces the overhead associated with retrieving sequence numbers, but it also leads to possible gaps in the sequences if a system failure occurs, or if the system is shut down. For instance, if you decide to set the sequence cache to 100, the database manager will cache 100 values of these numbers and also set the system catalog to show that the next sequence of values should begin at 201. In the event that the database is shut down, the next set of sequence numbers will begin at 201. The numbers that were generated from 101 to 200 will be lost from the set of sequences if they were not used. If gaps in generated values cannot be tolerated in your application, you need to set the caching value to NO CACHE despite the higher system overhead this will cause.

For more information on all available options and associated values, see the CREATE SEQUENCE statement.

Managing sequence behavior

You can tailor the behavior of sequences to meet the needs of your application. You change the attributes of a sequence when you issue the CREATE SEQUENCE statement to create a new sequence, and when you issue the ALTER SEQUENCE statement for an existing sequence.

Following are some of the attributes of a sequence that you can specify:

Data type

The AS clause of the CREATE SEQUENCE statement specifies the numeric

data type of the sequence. The data type determines the possible minimum and maximum values of the sequence. The minimum and maximum values for a data type are listed in SQL and XML limits. You cannot change the data type of a sequence; instead, you must drop the sequence by issuing the DROP SEQUENCE statement and issue a CREATE SEQUENCE statement with the new data type.

Start value

The START WITH clause of the CREATE SEQUENCE statement sets the initial value of the sequence. The RESTART WITH clause of the ALTER SEQUENCE statement resets the value of the sequence to a specified value.

Minimum value

The MINVALUE clause sets the minimum value of the sequence.

Maximum value

The MAXVALUE clause sets the maximum value of the sequence.

Increment value

The INCREMENT BY clause sets the value that each NEXT VALUE expression adds to the current value of the sequence. To decrement the value of the sequence, specify a negative value.

Sequence cycling

The CYCLE clause causes the value of a sequence that reaches its maximum or minimum value to generate its respective minimum value or maximum value on the following NEXT VALUE expression.

Note: CYCLE should only be used if unique numbers are not required or if it can be guaranteed that older sequence values are not in use anymore once the sequence cycles.

For example, to create a sequence called `id_values` that starts with a minimum value of 0, has a maximum value of 1000, increments by 2 with each NEXT VALUE expression, and returns to its minimum value when the maximum value is reached, issue the following statement:

```
CREATE SEQUENCE id_values
  START WITH 0
  INCREMENT BY 2
  MAXVALUE 1000
  CYCLE
```

Application performance and sequences

Like identity columns, using sequences to generate values generally improves the performance of your applications in comparison to alternative approaches. The alternative to sequences is to create a single-column table that stores the current value and to increment that value with either a trigger or under the control of the application. However, in a distributed environment where applications concurrently access the single-column table, the locking required to force serialized access to the table can seriously affect performance.

Sequences avoid the locking issues that are associated with the single-column table approach and can cache sequence values in memory to improve response time. To maximize the performance of applications that use sequences, ensure that your sequence caches an appropriate amount of sequence values. The CACHE clause of the CREATE SEQUENCE and ALTER SEQUENCE statements specifies the maximum number of sequence values that the database manager generates and stores in memory.

If your sequence must generate values in order, without introducing gaps in that order because of a system failure or database deactivation, use the ORDER and NO CACHE clauses in the CREATE SEQUENCE statement. The NO CACHE clause guarantees that no gaps appear in the generated values at the cost of some of your application's performance because it forces your sequence to write to the database log every time it generates a new value. Note that gaps can still appear due to transactions that rollback and do not actually use that sequence value that they requested.

Sequences compared to identity columns

Although sequences and identity columns appear to serve similar purposes for DB2 applications, there is an important difference. An identity column automatically generates values for a column in a single table using the LOAD utility. A sequence generates sequential values upon request that can be used in any SQL statement using the CREATE SEQUENCE statement.

Identity columns

Allow the database manager to automatically generate a unique numeric value for each row that is added to the table. If you are creating a table and you know you will need to uniquely identify each row that is added to that table, then you can add an identity column to the table definition as part of the CREATE TABLE statement:

```
CREATE TABLE <table name>
(<column name 1> INT,
 <column name 2>, DOUBLE,
 <column name 3> INT NOT NULL GENERATED ALWAYS AS IDENTITY
 (START WITH <value 1>, INCREMENT BY <value 2>))
```

In this example, the third column identifies the identity column. One of the attributes that you can define is the value used in the column to uniquely define each row when a row is added. The value following the INCREMENT BY clause shows by how much subsequent values of the identity column contents will be increased for every row added to the table.

Once created, the identity properties can be changed or removed using the ALTER TABLE statement. You can also use the ALTER TABLE statement to add identity properties on other columns.

Sequences

Allow the automatic generation of values. Sequences are ideally suited to the task of generating unique key values. Applications can use sequences to avoid possible concurrency and performance problems resulting from the generation of a unique counter through other means. Unlike an identity column, a sequence is not tied to a particular table column, nor is it bound to a unique table column and only accessible through that table column.

A sequence can be created, and later altered, so that it generates values by incrementing or decrementing values either without a limit; or to a user-defined limit, and then stopping; or to a user-defined limit, then cycling back to the beginning and starting again. Sequences are only supported in single partition databases.

The following example shows how to create a sequence called orderseq:

```
CREATE SEQUENCE orderseq
  START WITH 1
  INCREMENT BY 1
  NOMAXVALUE
  NOCYCLE
  CACHE 50
```

In this example, the sequence starts at 1 and increases by 1 with no upper limit. There is no reason to cycle back to the beginning and restart from 1 because there is no assigned upper limit. The CACHE parameter specifies the maximum number of sequence values that the database manager preallocates and keeps in memory.

Creating sequences

To create sequences, use the CREATE SEQUENCE statement. Unlike an identity column attribute, a sequence is not tied to a particular table column nor is it bound to a unique table column and only accessible through that table column.

There are several restrictions on where NEXT VALUE or PREVIOUS VALUE expressions can be used. A sequence can be created, or altered, so that it generates values in one of these ways:

- Increment or decrement monotonically (changing by a constant amount) without bound
- Increment or decrement monotonically to a user-defined limit and stop
- Increment or decrement monotonically to a user-defined limit and cycle back to the beginning and start again

Note: Use caution when recovering databases that use sequences: For sequence values that are used outside the database, for example sequence numbers used for bank checks, if the database is recovered to a point in time before the database failure, then this could cause the generation of duplicate values for some sequences. To avoid possible duplicate values, databases that use sequence values outside the database should not be recovered to a prior point in time.

To create a sequence called order_seq using defaults for all the options, issue the following statement in an application program or through the use of dynamic SQL statements:

```
CREATE SEQUENCE order_seq
```

This sequence starts at 1 and increases by 1 with no upper limit.

This example could represent processing for a batch of bank checks starting from 101 to 200. The first order would have been from 1 to 100. The sequence starts at 101 and increase by 1 with an upper limit of 200. NOCYCLE is specified so that duplicate cheque numbers are not produced. The number associated with the CACHE parameter specifies the maximum number of sequence values that the database manager preallocates and keeps in memory.

```
CREATE SEQUENCE order_seq
  START WITH 101
  INCREMENT BY 1
  MAXVALUE 200
  NOCYCLE
  CACHE 25
```

For more information about these and other options, and authorization requirements, see the CREATE SEQUENCE statement.

Generating sequential values

Generating sequential values is a common database application development problem. The best solution to that problem is to use sequences and sequence expressions in SQL. Each *sequence* is a uniquely named database object that can be accessed only by sequence expressions.

There are two *sequence expressions*: the PREVIOUS VALUE expression and the NEXT VALUE expression. The PREVIOUS VALUE expression returns the value most recently generated in the application process for the specified sequence. Any NEXT VALUE expressions occurring in the same statement as the PREVIOUS VALUE expression have no effect on the value generated by the PREVIOUS VALUE expression in that statement. The NEXT VALUE sequence expression increments the value of the sequence and returns the new value of the sequence.

To create a sequence, issue the CREATE SEQUENCE statement. For example, to create a sequence called `id_values` using the default attributes, issue the following statement:

```
CREATE SEQUENCE id_values
```

To generate the first value in the application session for the sequence, issue a VALUES statement using the NEXT VALUE expression:

```
VALUES NEXT VALUE FOR id_values
```

```
1
-----
1
```

1 record(s) selected.

To update the value of a column with the next value of the sequence, include the NEXT VALUE expression in the UPDATE statement, as follows:

```
UPDATE staff
  SET id = NEXT VALUE FOR id_values
  WHERE id = 350
```

To insert a new row into a table using the next value of the sequence, include the NEXT VALUE expression in the INSERT statement, as follows:

```
INSERT INTO staff (id, name, dept, job)
  VALUES (NEXT VALUE FOR id_values, 'Kandil', 51, 'Mgr')
```

Determining when to use identity columns or sequences

While there are similarities between identity columns and sequences, there are also differences. The characteristics of each can be used when designing your database and applications.

Depending on your database design and the applications using the database, the following characteristics will assist you in determining when to use identity columns and when to use sequences.

Identity column characteristics

- An identity column automatically generates values for a single table.
- When an identity column is defined as GENERATED ALWAYS, the values used are always generated by the database manager. Applications are not allowed to provide their own values during the modification of the contents of the table.

- After inserting a row, the generated identity value can be retrieved either by using the `IDENTITY_VAL_LOCAL()` function or by selecting the identity column back from the insert by using the `SELECT FROM INSERT` statement.
- The `LOAD` utility can generate `IDENTITY` values.

Sequence characteristics

- Sequences are not tied to any one table.
- Sequences generate sequential values that can be used in any SQL or XQuery statement.

Since sequences can be used by any application, there are two expressions used to control the retrieval of the next value in the specified sequence and the value generated previous to the statement being executed. The `PREVIOUS VALUE` expression returns the most recently generated value for the specified sequence for a previous statement within the current session. The `NEXT VALUE` expression returns the next value for the specified sequence. The use of these expressions allows the same value to be used across several SQL and XQuery statements within several tables.

Modifying sequences

Modify the attributes of an existing sequence with the `ALTER SEQUENCE` statement.

The attributes of the sequence that can be modified include:

- Changing the increment between future values
- Establishing new minimum or maximum values
- Changing the number of cached sequence numbers
- Changing whether the sequence will cycle or not
- Changing whether sequence numbers must be generated in order of request
- Restarting the sequence

There are two tasks that are not found as part of the creation of the sequence. They are:

- **RESTART:** Resets the sequence to the value specified implicitly or explicitly as the starting value when the sequence was created.
- **RESTART WITH <numeric-constant>:** Resets the sequence to the exact numeric constant value. The numeric constant can be any positive or negative value with no non-zero digits to the right of any decimal point.

After restarting a sequence or changing to `CYCLE`, it is possible to generate duplicate sequence numbers. Only future sequence numbers are affected by the `ALTER SEQUENCE` statement.

The data type of a sequence cannot be changed. Instead, you must drop the current sequence and then create a new sequence specifying the new data type.

All cached sequence values not used by the database manager are lost when a sequence is altered.

Viewing sequence definitions

Use the VALUES statement using the PREVIOUS VALUE option to view the reference information associated with a sequence or to view the sequence itself.

To display the current value of the sequence, issue a VALUES statement using the PREVIOUS VALUE expression:

```
VALUES PREVIOUS VALUE FOR id_values  
  
1  
-----  
1  
  
1 record(s) selected.
```

You can repeatedly retrieve the current value of the sequence, and the value that the sequence returns does not change until you issue a NEXT VALUE expression. In the following example, the PREVIOUS VALUE expression returns a value of 1, until the NEXT VALUE expression in the current connection increments the value of the sequence:

```
VALUES PREVIOUS VALUE FOR id_values  
  
1  
-----  
1  
  
1 record(s) selected.
```

```
VALUES PREVIOUS VALUE FOR id_values  
  
1  
-----  
1  
  
1 record(s) selected.
```

```
VALUES NEXT VALUE FOR id_values  
  
1  
-----  
2  
  
1 record(s) selected.
```

```
VALUES PREVIOUS VALUE FOR id_values  
  
1  
-----  
2  
  
1 record(s) selected.
```

This is even true if another connection consumes sequence values at the same time.

Dropping sequences

To delete a sequence, use the DROP statement.

When dropping sequences, the authorization ID of the statement must hold SYSADM or DBADM authority.

A specific sequence can be dropped by using:

```
DROP SEQUENCE <sequence_name>
```

where the <sequence_name> is the name of the sequence to be dropped and includes the implicit or explicit schema name to exactly identify an existing sequence.

Sequences that are system-created for IDENTITY columns cannot be dropped using the DROP SEQUENCE statement.

Once a sequence is dropped, all privileges on the sequence are also dropped.

Examples of how to code sequences

Many applications that are written require the use of sequence number to track invoice numbers, customer numbers, and other objects which get incremented by one whenever a new item is required. The database manager can auto-increment values in a table through the use of identity columns. Although this technique works well for individual tables, it may not be the most convenient way of generating unique values that need to be used across multiple tables.

The *sequence* object lets you create a value that gets incremented under programmer control and can be used across many tables. The following example shows a sequence number being created for customer numbers using a data type of integer:

```
CREATE SEQUENCE customer_no AS INTEGER
```

By default the sequence number starts at one and increments by one at a time and is of an INTEGER data type. The application needs to get the next value in the sequence by using the NEXT VALUE function. This function generates the next value for the sequence which can then be used for subsequent SQL statements:

```
VALUES NEXT VALUE FOR customer_no
```

Instead of generating the next number with the VALUES function, the programmer could have used this function within an INSERT statement. For instance, if the first column of the customer table contained the customer number, an INSERT statement could be written as follows:

```
INSERT INTO customers VALUES  
(NEXT VALUE FOR customer_no, 'comment', ...)
```

If the sequence number needs to be used for inserts into other tables, the PREVIOUS VALUE function can be used to retrieve the previously generated value. For instance, if the customer number just created needs to be used for a subsequent invoice record, the SQL would include the PREVIOUS VALUE function:

```
INSERT INTO invoices  
(34,PREVIOUS VALUE FOR customer_no, 234.44, ...)
```

The PREVIOUS VALUE function can be used multiple times within the application and it will only return the last value generated by that application. It may be possible that subsequent transactions have already incremented the sequence to another value, but you will always see the last number that is generated.

Sequence reference

sequence-reference:

```
| nextval-expression |
| prevval-expression |
```

nextval-expression:

```
| NEXT VALUE FOR sequence-name |
```

prevval-expression:

```
| PREVIOUS VALUE FOR sequence-name |
```

NEXT VALUE FOR *sequence-name*

A NEXT VALUE expression generates and returns the next value for the sequence specified by *sequence-name*.

PREVIOUS VALUE FOR *sequence-name*

A PREVIOUS VALUE expression returns the most recently generated value for the specified sequence for a previous statement within the current application process. This value can be referenced repeatedly by using PREVIOUS VALUE expressions that specify the name of the sequence. There may be multiple instances of PREVIOUS VALUE expressions specifying the same sequence name within a single statement; they all return the same value. In a partitioned database environment, a PREVIOUS VALUE expression may not return the most recently generated value.

A PREVIOUS VALUE expression can only be used if a NEXT VALUE expression specifying the same sequence name has already been referenced in the current application process, in either the current or a previous transaction (SQLSTATE 51035).

Notes

- A new value is generated for a sequence when a NEXT VALUE expression specifies the name of that sequence. However, if there are multiple instances of a NEXT VALUE expression specifying the same sequence name within a query, the counter for the sequence is incremented only once for each row of the result, and all instances of NEXT VALUE return the same value for a row of the result.
- The same sequence number can be used as a unique key value in two separate tables by referencing the sequence number with a NEXT VALUE expression for the first row (this generates the sequence value), and a PREVIOUS VALUE expression for the other rows (the instance of PREVIOUS VALUE refers to the sequence value most recently generated in the current session), as shown below:

```
INSERT INTO order(orderno, cutno)
VALUES (NEXT VALUE FOR order_seq, 123456);
```

```
INSERT INTO line_item (orderno, partno, quantity)
VALUES (PREVIOUS VALUE FOR order_seq, 987654, 1);
```

- NEXT VALUE and PREVIOUS VALUE expressions can be specified in the following places:

- select-statement or SELECT INTO statement (within the select-clause, provided that the statement does not contain a DISTINCT keyword, a GROUP BY clause, an ORDER BY clause, a UNION keyword, an INTERSECT keyword, or EXCEPT keyword)
- INSERT statement (within a VALUES clause)
- INSERT statement (within the select-clause of the fullselect)
- UPDATE statement (within the SET clause (either a searched or a positioned UPDATE statement), except that NEXT VALUE cannot be specified in the select-clause of the fullselect of an expression in the SET clause)
- SET Variable statement (except within the select-clause of the fullselect of an expression; a NEXT VALUE expression can be specified in a trigger, but a PREVIOUS VALUE expression cannot)
- VALUES INTO statement (within the select-clause of the fullselect of an expression)
- CREATE PROCEDURE statement (within the routine-body of an SQL procedure)
- CREATE TRIGGER statement within the triggered-action (a NEXT VALUE expression may be specified, but a PREVIOUS VALUE expression cannot)
- NEXT VALUE and PREVIOUS VALUE expressions cannot be specified (SQLSTATE 428F9) in the following places:
 - Join condition of a full outer join
 - DEFAULT value for a column in a CREATE or ALTER TABLE statement
 - Generated column definition in a CREATE OR ALTER TABLE statement
 - Summary table definition in a CREATE TABLE or ALTER TABLE statement
 - Condition of a CHECK constraint
 - CREATE TRIGGER statement (a NEXT VALUE expression may be specified, but a PREVIOUS VALUE expression cannot)
 - CREATE VIEW statement
 - CREATE METHOD statement
 - CREATE FUNCTION statement
 - An argument list of an XMLQUERY, XMLEXISTS, or XMLTABLE expression
- In addition, a NEXT VALUE expression cannot be specified (SQLSTATE 428F9) in the following places:
 - CASE expression
 - Parameter list of an aggregate function
 - Subquery in a context other than those explicitly allowed above
 - SELECT statement for which the outer SELECT contains a DISTINCT operator
 - Join condition of a join
 - SELECT statement for which the outer SELECT contains a GROUP BY clause
 - SELECT statement for which the outer SELECT is combined with another SELECT statement using the UNION, INTERSECT, or EXCEPT set operator
 - Nested table expression
 - Parameter list of a table function
 - WHERE clause of the outer-most SELECT statement, or a DELETE or UPDATE statement
 - ORDER BY clause of the outer-most SELECT statement

- select-clause of the fullselect of an expression, in the SET clause of an UPDATE statement
- IF, WHILE, DO ... UNTIL, or CASE statement in an SQL routine
- When a value is generated for a sequence, that value is consumed, and the next time that a value is requested, a new value will be generated. This is true even when the statement containing the NEXT VALUE expression fails or is rolled back.

If an INSERT statement includes a NEXT VALUE expression in the VALUES list for the column, and if an error occurs at some point during the execution of the INSERT (it could be a problem in generating the next sequence value, or a problem with the value for another column), then an insertion failure occurs (SQLSTATE 23505), and the value generated for the sequence is considered to be consumed. In some cases, reissuing the same INSERT statement might lead to success.

For example, consider an error that is the result of the existence of a unique index for the column for which NEXT VALUE was used and the sequence value generated already exists in the index. It is possible that the next value generated for the sequence is a value that does not exist in the index and so the subsequent INSERT would succeed.

- If in generating a value for a sequence, the maximum value for the sequence is exceeded (or the minimum value for a descending sequence) and cycles are not permitted, then an error occurs (SQLSTATE 23522). In this case, the user could ALTER the sequence to extend the range of acceptable values, or enable cycles for the sequence, or DROP and CREATE a new sequence with a different data type that has a larger range of values.

For example, a sequence may have been defined with a data type of SMALLINT, and eventually the sequence runs out of assignable values. DROP and re-create the sequence with the new definition to redefine the sequence as INTEGER.

- A reference to a NEXT VALUE expression in the select statement of a cursor refers to a value that is generated for a row of the result table. A sequence value is generated for a NEXT VALUE expression for each row that is fetched from the database. If blocking is done at the client, the values may have been generated at the server prior to the processing of the FETCH statement. This can occur when there is blocking of the rows of the result table. If the client application does not explicitly FETCH all the rows that the database has materialized, then the application will not see the results of all the generated sequence values (for the materialized rows that were not returned).
- A reference to a PREVIOUS VALUE expression in the select statement of a cursor refers to a value that was generated for the specified sequence prior to the opening of the cursor. However, closing the cursor can affect the values returned by PREVIOUS VALUE for the specified sequence in subsequent statements, or even for the same statement in the event that the cursor is reopened. This would be the case when the select statement of the cursor included a reference to NEXT VALUE for the same sequence name.
- **Compatibilities**
 - For compatibility with previous versions of DB2:
 - NEXTVAL and PREVVAl can be specified in place of NEXT VALUE and PREVIOUS VALUE
 - For compatibility with IBM IDS:
 - *sequence-name*.NEXTVAL can be specified in place of NEXT VALUE FOR *sequence-name*

- *sequence-name*.CURRVAL can be specified in place of PREVIOUS VALUE FOR *sequence-name*

Examples

Assume that there is a table called "order", and that a sequence called "order_seq" is created as follows:

```
CREATE SEQUENCE order_seq
  START WITH 1
  INCREMENT BY 1
  NO MAXVALUE
  NO CYCLE
  CACHE 24
```

Following are some examples of how to generate an "order_seq" sequence number with a NEXT VALUE expression:

```
INSERT INTO order(orderno, custno)
  VALUES (NEXT VALUE FOR order_seq, 123456);
```

or

```
UPDATE order
  SET orderno = NEXT VALUE FOR order_seq
  WHERE custno = 123456;
```

or

```
VALUES NEXT VALUE FOR order_seq INTO :hv_seq;
```


Chapter 16. Views

A *view* is an efficient way of representing data without the need to maintain it. A view is not an actual table and requires no permanent storage. A “virtual table” is created and used.

A *view* provides a different way of looking at the data in one or more tables; it is a named specification of a result table. The specification is a SELECT statement that is run whenever the view is referenced in an SQL statement. A view has columns and rows just like a table. All views can be used just like tables for data retrieval. Whether a view can be used in an insert, update, or delete operation depends on its definition.

A view can include all or some of the columns or rows contained in the tables on which it is based. For example, you can join a department table and an employee table in a view, so that you can list all employees in a particular department.

Figure 28 shows the relationship between tables and views.




Figure 28. Relationship between tables and views

You can use views to control access to sensitive data, because views allow multiple users to see different presentations of the same data. For example, several users might be accessing a table of data about employees. A manager sees data about his or her employees but not employees in another department. A recruitment officer sees the hire dates of all employees, but not their salaries; a financial officer sees the salaries, but not the hire dates. Each of these users works with a view derived from the table. Each view appears to be a table and has its own name.

When the column of a view is directly derived from the column of a base table, that view column inherits any constraints that apply to the table column. For

example, if a view includes a foreign key of its table, insert and update operations using that view are subject to the same referential constraints as is the table. Also, if the table of a view is a parent table, delete and update operations using that view are subject to the same rules as are delete and update operations on the table.

A view can derive the data type of each column from the result table, or base the types on the attributes of a user-defined structured type. This is called a *typed view*. Similar to a typed table, a typed view can be part of a view hierarchy. A *subview* inherits columns from its *superview*. The term *subview* applies to a typed view and to all typed views that are below it in the view hierarchy. A *proper subview* of a view *V* is a view below *V* in the typed view hierarchy.

A view can become inoperative (for example, if the table is dropped); if this occurs, the view is no longer available for SQL operations.

Designing views

A *view* provides a different way of looking at the data in one or more tables; it is a named specification of a result table.

The specification is a SELECT statement that is run whenever the view is referenced in an SQL statement. A view has columns and rows just like a base table. All views can be used just like tables for data retrieval. Whether a view can be used in an insert, update, or delete operation depends on its definition.

Views are classified by the operations they allow. They can be:

- Deletable
- Updatable
- Insertable
- Read-only

The view type is established according to its update capabilities. The classification indicates the kind of SQL operation that is allowed against the view.

Referential and check constraints are treated independently. They do not affect the view classification.

For example, you might not be able to insert a row into a table due to a referential constraint. If you create a view using that table, you also cannot insert that row using the view. However, if the view satisfies all the rules for an insertable view, it will still be considered an insertable view. This is because the insert restriction is on the table, not on the view definition.

For more information, see the CREATE VIEW statement.

System catalog views

The database manager maintains a set of tables and views that contain information about the data under its control. These tables and views are collectively known as the *system catalog*.

The system catalog contains information about the logical and physical structure of database objects such as tables, views, indexes, packages, and functions. It also contains statistical information. The database manager ensures that the descriptions in the system catalog are always accurate.

The system catalog views are like any other database view. SQL statements can be used to query the data in the system catalog views. A set of updatable system catalog views can be used to modify certain values in the system catalog.

Views with the check option

A view that is defined WITH CHECK OPTION enforces any rows that are modified or inserted against the SELECT statement for that view. Views with the check option are also called *symmetric views*. For example, a symmetric view that only returns only employees in department 10 will not allow insertion of employees in other departments. This option, therefore, ensures the integrity of the data being modified in the database, returning an error if the condition is violated during an INSERT or UPDATE operation.

If your application cannot define the desired rules as table check constraints, or the rules do not apply to all uses of the data, there is another alternative to placing the rules in the application logic. You can consider creating a view of the table with the conditions on the data as part of the WHERE clause and the WITH CHECK OPTION clause specified. This view definition restricts the retrieval of data to the set that is valid for your application. Additionally, if you can update the view, the WITH CHECK OPTION clause restricts updates, inserts, and deletes to the rows applicable to your application.

The WITH CHECK OPTION must not be specified for the following views:

- Views defined with the read-only option (a read-only view)
- View that reference the NODENUMBER or PARTITION function, a nondeterministic function (for example, RAND), or a function with external action
- Typed views

Example 1

Following is an example of a view definition using the WITH CHECK OPTION. This option is required to ensure that the condition is always checked. The view ensures that the DEPT is always 10. This will restrict the input values for the DEPT column. When a view is used to insert a new value, the WITH CHECK OPTION is always enforced:

```
CREATE VIEW EMP_VIEW2
  (EMPNO, EMPNAME, DEPTNO, JOBTITLE, HIREDATE)
AS SELECT ID, NAME, DEPT, JOB, HIREDATE FROM EMPLOYEE
  WHERE DEPT=10
WITH CHECK OPTION;
```

If this view is used in an INSERT statement, the row will be rejected if the DEPTNO column is not the value 10. It is important to remember that there is no data validation during modification if the WITH CHECK OPTION is not specified.

If this view is used in a SELECT statement, the conditional (WHERE clause) would be invoked and the resulting table would only contain the matching rows of data. In other words, the WITH CHECK OPTION does not affect the result of a SELECT statement.

Example 2

With a view, you can make a subset of table data available to an application program and validate data that is to be inserted or updated. A view can have column names that are different from the names of corresponding columns in the original tables. For example:

```
CREATE VIEW <name> (<column>, <column>, <column>)
  SELECT <column_name> FROM <table_name>
  WITH CHECK OPTION
```

Example 3

The use of views provides flexibility in the way your programs and end-user queries can look at the table data.

The following SQL statement creates a view on the EMPLOYEE table that lists all employees in Department A00 with their employee and telephone numbers:

```
CREATE VIEW EMP_VIEW (DA00NAME, DA00NUM, PHONENO)
  AS SELECT LASTNAME, EMPNO, PHONENO FROM EMPLOYEE
  WHERE WORKDEPT = 'A00'
  WITH CHECK OPTION
```

The first line of this statement names the view and defines its columns. The name EMP_VIEW must be unique within its schema in SYSCAT.TABLES. The view name appears as a table name although it contains no data. The view will have three columns called DA00NAME, DA00NUM, and PHONENO, which correspond to the columns LASTNAME, EMPNO, and PHONENO from the EMPLOYEE table. The column names listed apply one-to-one to the select list of the SELECT statement. If column names are not specified, the view uses the same names as the columns of the result table of the SELECT statement.

The second line is a SELECT statement that describes which values are to be selected from the database. It might include the clauses ALL, DISTINCT, FROM, WHERE, GROUP BY, and HAVING. The name or names of the data objects from which to select columns for the view must follow the FROM clause.

Example 4

The WITH CHECK OPTION clause indicates that any updated or inserted row to the view must be checked against the view definition, and rejected if it does not conform. This enhances data integrity but requires additional processing. If this clause is omitted, inserts and updates are not checked against the view definition.

The following SQL statement creates the same view on the EMPLOYEE table using the SELECT AS clause:

```
CREATE VIEW EMP_VIEW
  SELECT LASTNAME AS DA00NAME,
         EMPNO AS DA00NUM,
         PHONENO
  FROM EMPLOYEE
  WHERE WORKDEPT = 'A00'
  WITH CHECK OPTION
```

For this example, the EMPLOYEE table might have salary information in it, which should not be made available to everyone. The employee's phone number, however, should be generally accessible. In this case, a view could be created from the LASTNAME and PHONENO columns only. Access to the view could be

granted to PUBLIC, while access to the entire EMPLOYEE table could be restricted to those who have the authorization to see salary information.

Nested view definitions

If a view is based on another view, the number of predicates that must be evaluated is based on the WITH CHECK OPTION specification.

If a view is defined without WITH CHECK OPTION, the definition of the view is not used in the data validity checking of any insert or update operations. However, if the view directly or indirectly depends on another view defined with the WITH CHECK OPTION, the definition of that super view is used in the checking of any insert or update operation.

If a view is defined with the WITH CASCADED CHECK OPTION or just the WITH CHECK OPTION (CASCADED is the default value of the WITH CHECK OPTION), the definition of the view is used in the checking of any insert or update operations. In addition, the view inherits the search conditions from any updatable views on which the view depends. These conditions are inherited even if those views do not include the WITH CHECK OPTION. Then the inherited conditions are multiplied together to conform to a constraint that is applied for any insert or update operations for the view or any views depending on the view.

As an example, if a view V2 is based on a view V1, and the check option for V2 is defined with the WITH CASCADED CHECK OPTION, the predicates for both views are evaluated when INSERT and UPDATE statements are performed against the view V2:

```
CREATE VIEW EMP_VIEW2 AS
  SELECT EMPNO, EMPNAME, DEPTNO FROM EMP
     WHERE DEPTNO = 10
  WITH CHECK OPTION;
```

The following example shows a CREATE VIEW statement using the WITH CASCADED CHECK OPTION. The view EMP_VIEW3 is created based on a view EMP_VIEW2, which has been created with the WITH CHECK OPTION. If you want to insert or update a record to EMP_VIEW3, the record should have the values DEPTNO=10 and EMPNO=20.

```
CREATE VIEW EMP_VIEW3 AS
  SELECT EMPNO, EMPNAME, DEPTNO FROM EMP_VIEW2
     WHERE EMPNO > 20
  WITH CASCADED CHECK OPTION;
```

Note: The condition DEPTNO=10 is enforced for inserting or updating operations to EMP_VIEW3 even if EMP_VIEW2 does not include the WITH CHECK OPTION.

The WITH LOCAL CHECK OPTION can also be specified when creating a view. If a view is defined with the LOCAL CHECK OPTION, the definition of the view is used in the checking of any insert or update operations. However, the view does not inherit the search conditions from any updatable views on which it depends.

Deletable views

Depending on how a view is defined, the view can be deletable. A deletable view is a view against which you can successfully issue a DELETE statement.

There are a few rules that need to be followed for a view to be considered deletable:

- Each FROM clause of the outer fullselect identifies only one table (with no OUTER clause), deletable view (with no OUTER clause), deletable nested table expression, or deletable common table expression.
- The database manager needs to be able to derive the rows to be deleted in the table using the view definition. Certain operations make this impossible
 - A grouping of multiple rows into one using a GROUP BY clause or column functions result in a loss of the original row and make the view non deletable.
 - Similarly when the rows are derived from a VALUES there is no table to delete from. Again the view is not deletable.
- The outer fullselect doesn't use the GROUP BY or HAVING clauses.
- The outer fullselect doesn't include column functions in its select list.
- The outer fullselect doesn't use set operations (UNION, EXCEPT, or INTERSECT) with the exception of UNION ALL
- The tables in the operands of a UNION ALL must not be the same table, and each operand must be deletable.
- The select list of the outer fullselect does not include DISTINCT.

A view must meet all the rules listed above to be considered a deletable view. For example, the following view is deletable. It follows all the rules for a deletable view.

```
CREATE VIEW deletable_view
  (number, date, start, end)
AS
  SELECT number, date, start, end
  FROM employee.summary
  WHERE date='01012007'
```

Insertable views

Insertable views allow you to insert rows using the view definition. A view is insertable if an INSTEAD OF trigger for the insert operation has been defined for the view, or at least one column of the view is updatable (independent of an INSTEAD OF trigger for update), and the fullselect of the view does not include UNION ALL. A given row can be inserted into a view (including a UNION ALL) if, and only if, it fulfills the check constraints of exactly one of the underlying tables. To insert into a view that includes non-updatable columns, those columns must be omitted from the column list.

The view shown below is an insertable view. However, in this example, an attempt to insert the view will fail. This is because there are columns in the table that do not accept null values. Some of these columns are not present in the view definition. When you try to insert a value using the view, the database manager will try to insert a null value into a NOT NULL column. This action is not permitted.

```
CREATE VIEW insertable_view
  (number, name, quantity)
AS
  SELECT number, name, quantify FROM ace.supplies
```

Note: The constraints defined on the table are independent of the operations that can be performed using a view based on that table.

Updatable views

An updatable view is a special case of a deletable view. A deletable view becomes an updatable view when at least one of its columns is updatable.

A column of a view is updatable when all of the following rules are true:

- The view is deletable.
- The column resolves to a column of a table (not using a dereference operation) and the READ ONLY option is not specified.
- All the corresponding columns of the operands of a UNION ALL have exactly matching data types (including length or precision and scale) and matching default values if the fullselect of the view includes a UNION ALL.

The following example uses constant values that cannot be updated. However, the view is a deletable view and at least one of its columns is updatable. Therefore, it is an updatable view.

```
CREATE VIEW updatable_view
  (number, current_date, current_time, temperature)
AS
  SELECT number, CURRENT DATE, CURRENT TIME, temperature)
FROM weather.forecast
WHERE number = 300
```

Read-only views

A view is *read-only* if it is *not* deletable, updatable, or insertable. A view can be read-only if it is a view that does not comply with at least one of the rules for deletable views.

The READONLY column in the SYSCAT.VIEWS catalog view indicates a view is read-only (R).

The example shown below is not a deletable view as it uses the DISTINCT clause and the SQL statement involves more than one table:

```
CREATE VIEW read_only_view
  (name, phone, address)
AS
  SELECT DISTINCT viewname, viewphone, viewaddress
FROM employee.history adam, employer.dept sales
WHERE adam.id = sales.id
```

Creating views

Views are derived from one or more tables, nicknames, or views, and can be used interchangeably with tables when retrieving data. When changes are made to the data shown in a view, the data is changed in the table itself. The table, nickname, or view on which the view is to be based must already exist before the view can be created.

A view can be created to limit access to sensitive data, while allowing more general access to other data.

When inserting into a view where the select list of the view definition directly or indirectly includes the name of an identity column of a table, the same rules apply as if the INSERT statement directly referenced the identity column of the table.

In addition to using views as described above, a view can also be used to:

- Alter a table without affecting application programs. This can happen by creating a view based on an underlying table. Applications that use the underlying table are not affected by the creation of the new view. New applications can use the created view for different purposes than those applications that use the underlying table.

- Sum the values in a column, select the maximum values, or average the values.
- Provide access to information in one or more data sources. You can reference nicknames within the CREATE VIEW statement and create multi-location/global views (the view could join information in multiple data sources located on different systems).

When you create a view that references nicknames using standard CREATE VIEW syntax, you will see a warning alerting you to the fact that the authentication ID of view users will be used to access the underlying object or objects at data sources instead of the view creator authentication ID. Use the FEDERATED keyword to suppress this warning.

A typed view is based on a predefined structured type. You can create a typed view using the CREATE VIEW statement.

An alternative to creating a view is to use a nested or common table expression to reduce catalog lookup and improve performance.

A sample CREATE VIEW statement is shown below. The underlying table, EMPLOYEE, has columns named SALARY and COMM. For security reasons this view is created from the ID, NAME, DEPT, JOB, and HIREDATE columns. In addition, access on the DEPT column is restricted. This definition will only show the information of employees who belong to the department whose DEPTNO is 10.

```
CREATE VIEW EMP_VIEW1
  (EMPID, EMPNAME, DEPTNO, JOBTITLE, HIREDATE)
AS SELECT ID, NAME, DEPT, JOB, HIREDATE FROM EMPLOYEE
WHERE DEPT=10;
```

After the view has been defined, the access privileges can be specified. This provides data security since a restricted view of the table is accessible. As shown above, a view can contain a WHERE clause to restrict access to certain rows or can contain a subset of the columns to restrict access to certain columns of data.

The column names in the view do not have to match the column names of the base table. The table name has an associated schema as does the view name.

Once the view has been defined, it can be used in statements such as SELECT, INSERT, UPDATE, and DELETE (with restrictions). The DBA can decide to provide a group of users with a higher level privilege on the view than the table.

Creating views that use user-defined functions (UDFs)

Once you create a view that uses a UDF, the view will always use this same UDF as long as the view exists even if you create other UDFs with the same names later. If you want to pick up a new UDF you need to recreate the view.

The following SQL statement creates a view with a function in its definition:

```
CREATE VIEW EMPLOYEE_PENSION (NAME, PENSION)
  AS SELECT NAME, PENSION(HIREDATE, BIRTHDATE, SALARY, BONUS)
  FROM EMPLOYEE
```

The UDF function PENSION calculates the current pension an employee is eligible to receive, based on a formula involving their HIREDATE, BIRTHDATE, SALARY, and BONUS.

Modifying typed views

Certain properties of a typed view can be changed without requiring the view to be dropped and recreated. One such property is the adding of a scope to a reference column of a typed view.

The ALTER VIEW statement modifies an existing typed view definition by altering a reference type column to add a scope. The DROP statement deletes a typed view. You can also:

- Modify the contents of a typed view through INSTEAD OF triggers
- Alter a typed view to enable statistics collection

Changes you make to the underlying content of a typed view require that you use triggers. Other changes to a typed view require that you drop and then re-create the typed view.

The data type of the column-name in the ALTER VIEW statement must be REF (type of the typed table name or typed view name).

Other database objects such as tables and indexes are not affected although packages and cached dynamic statements are marked invalid.

To alter a typed view using the command line, enter:

```
ALTER VIEW <view_name> ALTER <column_name>
ADD SCOPE <typed table or view name>
```

Recovering inoperative views

An inoperative view is a view that is no longer available for SQL statements.

Views can become *inoperative*:

- As a result of a revoked privilege on an underlying table
- If a table, alias, or function is dropped.
- If the superview becomes inoperative. (A superview is a typed view upon which another typed view, a subview, is based.)
- When the views they are dependent on are dropped.

The following steps can help you recover an inoperative view:

1. Determine the SQL statement that was initially used to create the view. You can obtain this information from the TEXT column of the SYSCAT.VIEW catalog view.
2. Set the current schema to the content of the QUALIFIER column.
3. Set the function path to the content of the FUNC_PATH column.
4. Re-create the view by using the CREATE VIEW statement with the same view name and same definition.
5. Use the GRANT statement to re-grant all privileges that were previously granted on the view. (Note that all privileges granted on the inoperative view are revoked.)

If you do not want to recover an inoperative view, you can explicitly drop it with the DROP VIEW statement, or you can create a new view with the same name but a different definition.

An inoperative view only has entries in the SYSCAT.TABLES and SYSCAT.VIEWS catalog views; all entries in the SYSCAT.TABDEP, SYSCAT.TABAUTH, SYSCAT.COLUMNS and SYSCAT.COLAUTH catalog views are removed.

Dropping views

Use the DROP VIEW statement to drop views. Any views that are dependent on the view being dropped will be made inoperative.

To drop a view using the command line, enter:

```
DROP VIEW <view_name>
```

The following example shows how to drop a view named EMP_VIEW:

```
DROP VIEW EMP_VIEW
```

As in the case of a table hierarchy, it is possible to drop an entire view hierarchy in one statement by naming the root view of the hierarchy, as in the following example:

```
DROP VIEW HIERARCHY VPerson
```

Part 4. Reference

Chapter 17. Conforming to naming rules

Naming rules

Rules exist for the naming of all objects, users and groups. Some of these rules are specific to the platform you are working on.

For example, there is a rule regarding the use of upper and lowercase letters in a name.

- On UNIX platforms, names must be in lowercase.
- On Windows platforms, names can be in upper, lower, and mixed-case.

Unless otherwise specified, all names can include the following characters:

- A through Z. When used in most names, characters A through Z are converted from lowercase to uppercase.
- 0 through 9.
- ! % () { } . - ^ ~ _ (underscore) @, #, \$, and space.
- \ (backslash).

Names cannot begin with a number or with the underscore character.

Do not use SQL reserved words to name tables, views, columns, indexes, or authorization IDs.

There are other special characters that might work separately depending on your operating system and where you are working with the DB2 database. However, while they might work, there is no guarantee that they will work. It is not recommended that you use these other special characters when naming objects in your database.

User and group names also need to follow the rules forced on specific operation systems by the related systems. For example, on Linux and UNIX platforms, allowed characters for user names and primary group names must be lowercase a through z, 0 through 9, and _ (underscore) for names not starting with 0 through 9.

Lengths must be less than or equal to the lengths listed in: SQL and XML limits.

You also need to consider object naming rules, naming rules in an NLS environment, and naming rules in a Unicode environment.

Restrictions on the AUTHID identifier: Version 9.5, and later, of the DB2 database system allows you to have an 128-byte authorization ID, but when the authorization ID is interpreted as an operating system user ID or group name, the operating system naming restrictions apply (for example, a limitation to 8 or 30 character user IDs and group names). Therefore, while you can grant an 128-byte authorization ID, it is not possible to connect as a user that has that authorization ID. If you write your own security plugin, you should be able to take full advantage of the extended sizes for the authorization ID. For example, you can give your security plugin a 30-byte user ID and it can return an 128-byte authorization ID during authentication that you are able to connect with.

DB2 object naming rules

All objects follow the General Naming Rules. In addition, some objects have additional restrictions shown in the accompanying tables.

Table 52. Database, database alias and instance naming rules

Objects	Guidelines
<ul style="list-style-type: none"> • Databases • Database aliases • Instances 	<ul style="list-style-type: none"> • Database names must be unique within the location in which they are cataloged. On Linux and UNIX implementations, this location is a directory path, while on Windows implementations, it is a logical disk. • Database alias names must be unique within the system database directory. When a new database is created, the alias defaults to the database name. As a result, you cannot create a database using a name that exists as a database alias, even if there is no database with that name. • Database, database alias and instance name lengths must be less than or equal to 8 bytes • On Windows, no instance can have the same name as a service name. <p>Note: To avoid potential problems, do not use the special characters @, #, and \$ in a database name if you intend to use the database in a communications environment. Also, because these characters are not common to all keyboards, do not use them if you plan to use the database in another language.</p>

Table 53. Database object naming rules

Objects	Guidelines
<ul style="list-style-type: none"> • Aliases • Audit policies • Buffer pools • Columns • Event monitors • Indexes • Methods • Nodegroups • Packages • Package versions • Roles • Schemas • Stored procedures • Tables • Table spaces • Triggers • Trusted contexts • UDFs • UDTs • Views 	<p>Lengths for these objects must be less than or equal to the lengths listed in: SQL and XML limits. Object names can also include:</p> <ul style="list-style-type: none"> • Valid accented characters (such as ö) • Multibyte characters, except multibyte spaces (for multibyte environments) <p>Package names and package versions can also include periods (.), hyphens (-), and colons (:).</p>

Table 54. Federated database object naming rules

Objects	Guidelines
<ul style="list-style-type: none"> • Function mappings • Index specifications • Nicknames • Servers • Type mappings • User mappings • Wrappers 	<p>Lengths for these objects must be less than or equal to the lengths listed in: SQL and XML limits. Names for federated database objects can also include:</p> <ul style="list-style-type: none"> • Valid accented letters (such as ö) • Multibyte characters, except multibyte spaces (for multibyte environments)

Delimited identifiers and object names

Keywords can be used. If a keyword is used in a context where it could also be interpreted as an SQL keyword, it must be specified as a delimited identifier.

Using delimited identifiers, it is possible to create an object that violates these naming rules; however, subsequent use of the object could result in errors. For example, if you create a column with a + or - sign included in the name and you subsequently use that column in an index, you will experience problems when you attempt to reorganize the table.

Additional schema names information

- User-defined types (UDTs) cannot have schema names longer than the lengths listed in: SQL and XML limits.
- The following schema names are reserved words and must not be used: SYSCAT, SYSFUN, SYSIBM, SYSSTAT.
- To avoid potential migration problems in the future, do not use schema names that begin with SYS. The database manager will not allow you to create triggers, user-defined types or user-defined functions using a schema name beginning with SYS.
- It is recommended that you not use SESSION as a schema name. Declared temporary tables must be qualified by SESSION. It is therefore possible to have an application declare a temporary table with a name identical to that of a persistent table, in which case the application logic can become overly complicated. Avoid the use of the schema SESSION, except when dealing with declared temporary tables.

Delimited identifiers and object names

Keywords can be used. If a keyword is used in a context where it could also be interpreted as an SQL keyword, it must be specified as a delimited identifier.

Using delimited identifiers, it is possible to create an object that violates these naming rules; however, subsequent use of the object could result in errors. For example, if you create a column with a + or - sign included in the name and you subsequently use that column in an index, you will experience problems when you attempt to reorganize the table.

User, user ID and group naming rules

User, user ID and group names must follow naming guidelines.

Table 55. User, user ID and group naming rules

Objects	Guidelines
<ul style="list-style-type: none">• Group names• User names• User IDs	<ul style="list-style-type: none">• Group names must be less than or equal to the group name length listed in: SQL and XML limits.• User IDs on Linux and UNIX operating systems can contain up to 8 characters.• User names on Windows can contain up to 30 characters.• When not using Client authentication, non-Windows 32-bit clients connecting to Windows with user names longer than the user name length listed in listed in SQL and XML limits are supported when the user name and password are specified explicitly.• Names and IDs cannot:<ul style="list-style-type: none">– Be USERS, ADMINS, GUESTS, PUBLIC, LOCAL or any SQL reserved word– Begin with IBM, SQL or SYS.

Note:

1. Some operating systems allow case sensitive user IDs and passwords. You should check your operating system documentation to see if this is the case.
2. The authorization ID returned from a successful CONNECT or ATTACH is truncated to the authorization name length listed in SQL and XML limits. An ellipsis (...) is appended to the authorization ID and the SQLWARN fields contain warnings to indicate truncation.
3. Trailing blanks from user IDs and passwords are removed.

Naming rules in an NLS environment

The basic character set that can be used in database names consists of the single-byte uppercase and lowercase Latin letters (A...Z, a...z), the Arabic numerals (0...9) and the underscore character (_).

This list is augmented with three special characters (#, @, and \$) to provide compatibility with host database products. Use special characters #, @, and \$ with care in an NLS environment because they are not included in the NLS host (EBCDIC) invariant character set. Characters from the extended character set can also be used, depending on the code page that is being used. If you are using the database in a multiple code page environment, you must ensure that all code pages support any elements from the extended character set you plan to use.

When naming database objects (such as tables and views), program labels, host variables, cursors, and elements from the extended character set (for example, letters with diacritical marks) can also be used. Precisely which characters are available depends on the code page in use.

Extended Character Set Definition for DBCS Identifiers: In DBCS environments, the extended character set consists of all the characters in the basic character set, plus the following:

- All double-byte characters in each DBCS code page, except the double-byte space, are valid letters.
- The double-byte space is a special character.
- The single-byte characters available in each mixed code page are assigned to various categories as follows:

Category	Valid Code Points within each Mixed Code Page
Digits	x30-39
Letters	x23-24, x40-5A, x61-7A, xA6-DF (A6-DF for code pages 932 and 942 only)
Special Characters	All other valid single-byte character code points

Naming rules in a Unicode environment

In a Unicode database, all identifiers are in multibyte UTF-8. Therefore, it is possible to use any UCS-2 character in identifiers where the use of a character in the extended character set (for example, an accented character, or a multibyte character) is allowed by the DB2 database system.

Clients can enter any character that is supported by their environment, and all the characters in the identifiers will be converted to UTF-8 by the database manager. Two points must be taken into account when specifying national language characters in identifiers for a Unicode database:

- Each non-ASCII character requires two to four bytes. Therefore, an n -byte identifier can only hold somewhere between $n/4$ and n characters, depending on the ratio of ASCII to non-ASCII characters. If you have only one or two non-ASCII (for example, accented) characters, the limit is closer to n characters, while for an identifier that is completely non-ASCII (for example, in Japanese), only $n/4$ to $n/3$ characters can be used.
- If identifiers are to be entered from different client environments, they should be defined using the common subset of characters available to those clients. For example, if a Unicode database is to be accessed from Latin-1, Arabic, and Japanese environments, all identifiers should realistically be limited to ASCII.

Chapter 18. SQL and XML limits

The following tables describe certain SQL and XML limits. Adhering to the most restrictive case can help you to design application programs that are easily portable.

Table 56 lists limits in bytes. These limits are enforced after conversion from the application code page to the database code page when creating identifiers. The limits are also enforced after conversion from the database code page to the application code page when retrieving identifiers from the database. If, during either of these processes, the identifier length limit is exceeded, truncation occurs or an error is returned.

Character limits vary depending on the code page of the database and the code page of the application. For example, because the width of a UTF-8 character can range from 1 to 4 bytes, the character limit for an identifier in a Unicode table whose limit is 128 bytes will range from 32 to 128 characters, depending on which characters are used. If an attempt is made to create an identifier whose name is longer than the limit for this table after conversion to the database code page, an error is returned.

Applications that store identifier names must be able to handle the potentially increased size of identifiers after code page conversion has occurred. When identifiers are retrieved from the catalog, they are converted to the application code page. Conversion from the database code page to the application code page can result in an identifier becoming longer than the byte limit for the table. If a host variable declared by the application cannot store the entire identifier after code page conversion, it is truncated. If that is unacceptable, the host variable can be increased in size to be able to accept the entire identifier name.

The same rules apply to DB2 utilities retrieving data and converting it to a user-specified code page. If a DB2 utility, such as export, is retrieving the data and forcing conversion to a user-specified code page (using the export CODEPAGE modifier or the **DB2CODEPAGE** registry variable), and the identifier expands beyond the limit that is documented in this table because of code page conversion, an error might be returned or the identifier might be truncated.

Table 56. Identifier Length Limits

Description	Maximum in Bytes
Alias name	128
Attribute name	128
Audit policy name	128
Authorization name (can only be single-byte characters)	128
Buffer pool name	18
Column name ²	128
Constraint name	128
Correlation name	128
Cursor name	128
Data partition name	128

Table 56. Identifier Length Limits (continued)

Description	Maximum in Bytes
Data source column name	255
Data source index name	128
Data source name	128
Data source table name (<i>remote-table-name</i>)	128
Database partition group name	128
Database partition name	128
Event monitor name	128
External program name	128
Function mapping name	128
Group name	128
Host identifier ¹	255
Identifier for a data source user (<i>remote-authorization-name</i>)	128
Identifier in an SQL procedure (condition name, for loop identifier, label, result set locator, statement name, variable name)	128
Index name	128
Index extension name	18
Index specification name	128
Label name	128
Namespace uniform resource identifier (URI)	1000
Nickname	128
Package name	128
Package version ID	64
Parameter name	128
Password to access a data source	32
Procedure name	128
Role name	128
Savepoint name	128
Schema name ²	128
Security label component name	128
Security label name	128
Security policy name	128
Sequence name	128
Server (database alias) name	8
Specific name	128
SQL condition name	128
SQL variable name	128
Statement name	128
Table name	128
Table space name	18

Table 56. Identifier Length Limits (continued)

Description	Maximum in Bytes
Transform group name	18
Trigger name	128
Trusted context name	128
Type mapping name	18
User-defined function name	128
User-defined method name	128
User-defined type name ²	128
View name	128
Wrapper name	128
XML element name, attribute name, or prefix name	1000
XML schema location uniform resource identifier (URI)	1000
Note:	
<ol style="list-style-type: none"> Individual host language compilers might have a more restrictive limit on variable names. The SQLDA structure is limited to storing 30-byte column names, 18-byte user-defined type names, and 8-byte schema names for user-defined types. Because the SQLDA is used in the DESCRIBE statement, embedded SQL applications that use the DESCRIBE statement to retrieve column or user-defined type name information must conform to these limits. 	

Table 57. Numeric Limits

Description	Limit
Smallest SMALLINT value	-32 768
Largest SMALLINT value	+32 767
Smallest INTEGER value	-2 147 483 648
Largest INTEGER value	+2 147 483 647
Smallest BIGINT value	-9 223 372 036 854 775 808
Largest BIGINT value	+9 223 372 036 854 775 807
Largest decimal precision	31
Maximum exponent (E_{max}) for REAL values	38
Smallest REAL value	-3.402E+38
Largest REAL value	+3.402E+38
Minimum exponent (E_{min}) for REAL values	-37
Smallest positive REAL value	+1.175E-37
Largest negative REAL value	-1.175E-37
Maximum exponent (E_{max}) for DOUBLE values	308
Smallest DOUBLE value	-1.79769E+308
Largest DOUBLE value	+1.79769E+308
Minimum exponent (E_{min}) for DOUBLE values	-307

Table 58. String Limits (continued)

Description	Limit
Maximum length of CLOB (in bytes)	2 147 483 647
Maximum length of serialized XML (in bytes)	2 147 483 647
Maximum length of GRAPHIC (in double-byte characters)	127
Maximum length of VARGRAPHIC (in double-byte characters)	16 336
Maximum length of LONG VARGRAPHIC (in double-byte characters) ¹	16 350
Maximum length of DBCLOB (in double-byte characters)	1 073 741 823
Maximum length of BLOB (in bytes)	2 147 483 647
Maximum length of character constant	32 672
Maximum length of graphic constant	16 336
Maximum length of concatenated character string	2 147 483 647
Maximum length of concatenated graphic string	1 073 741 823
Maximum length of concatenated binary string	2 147 483 647
Maximum number of hexadecimal constant digits	32 672
Largest instance of a structured type column object at run time (in gigabytes)	1
Maximum size of a catalog comment (in bytes)	254
Note:	
1. The LONG VARCHAR and LONG VARGRAPHIC data types are deprecated, not recommended, and might be removed in a future release.	

Table 59. XML Limits

Description	Limit
Maximum depth of an XML document (in levels)	125
Maximum size of an XML schema document (in bytes)	31 457 280

Table 60. Datetime Limits

Description	Limit
Smallest DATE value	0001-01-01
Largest DATE value	9999-12-31
Smallest TIME value	00:00:00
Largest TIME value	24:00:00
Smallest TIMESTAMP value	0001-01-01-00.00.00.000000
Largest TIMESTAMP value	9999-12-31-24.00.00.000000

Table 61. Database Manager Limits

Description	Limit
Applications	
Maximum number of host variable declarations in a precompiled program ³	storage
Maximum length of a host variable value (in bytes)	2 147 483 647

Table 61. Database Manager Limits (continued)

Description	Limit
Maximum number of declared cursors in a program	storage
Maximum number of rows changed in a unit of work	storage
Maximum number of cursors opened at one time	storage
Maximum number of connections per process within a DB2 client	512
Maximum number of simultaneously opened LOB locators in a transaction	4 194 304
Maximum size of an SQLDA (in bytes)	storage
Maximum number of prepared statements	storage
Buffer Pools	
Maximum NPAGES in a buffer pool for 32-bit releases	1 048 576
Maximum NPAGES in a buffer pool for 64-bit releases	2 147 483 647
Maximum total size of all buffer pool slots (4K)	2 147 483 646
Concurrency	
Maximum number of concurrent users of a server ⁴	64 000
Maximum number of concurrent users per instance	64 000
Maximum number of concurrent applications per database	60 000
Maximum number of databases per instance concurrently in use	256
Constraints	
Maximum number of constraints on a table	storage
Maximum number of columns in a UNIQUE constraint (supported through a UNIQUE index)	64
Maximum combined length of columns in a UNIQUE constraint (supported through a UNIQUE index, in bytes) ⁹	8192
Maximum number of referencing columns in a foreign key	64
Maximum combined length of referencing columns in a foreign key (in bytes) ⁹	8192
Maximum length of a check constraint specification (in bytes)	65 535
Databases	
Maximum database partition number	999
Indexes	
Maximum number of indexes on a table	32 767 or storage
Maximum number of columns in an index key	64
Maximum length of an index key including all overhead ^{7 9}	<i>indexpagesize/4</i>
Maximum length of a variable index key part (in bytes) ⁸	1022 or storage
Maximum size of an index per database partition in an SMS table space (in gigabytes) ⁷	16 384
Maximum size of an index per database partition in a regular DMS table space (in gigabytes) ⁷	512
Maximum size of an index per database partition in a large DMS table space (in gigabytes) ⁷	16 384

Table 61. Database Manager Limits (continued)

Description	Limit
Maximum size of an index over XML data per database partition (in terabytes)	2
Maximum length of a variable index key part for an index over XML data (in bytes) ⁷	<i>pagesize/4 - 207</i>
Log records	
Maximum Log Sequence Number	281 474 976 710 655
Monitoring	
Maximum number of simultaneously active event monitors	32
Routines	
Maximum number of parameters in a procedure	32 767
Maximum number of parameters in a user-defined function	90
Maximum number of nested levels for routines	64
Maximum number of schemas in the SQL path	64
Maximum length of the SQL path (in bytes)	2048
Security	
Maximum number of elements in a security label component of type set or tree	64
Maximum number of elements in a security label component of type array	65 535
Maximum number of security label components in a security policy	16
SQL	
Maximum total length of an SQL statement (in bytes)	2 097 152
Maximum number of tables referenced in an SQL statement or a view	storage
Maximum number of host variable references in an SQL statement	32 767
Maximum number of constants in a statement	storage
Maximum number of elements in a select list ⁷	1012
Maximum number of predicates in a WHERE or HAVING clause	storage
Maximum number of columns in a GROUP BY clause ⁷	1012
Maximum total length of columns in a GROUP BY clause (in bytes) ⁷	32 677
Maximum number of columns in an ORDER BY clause ⁷	1012
Maximum total length of columns in an ORDER BY clause (in bytes) ⁷	32 677
Maximum level of subquery nesting	storage
Maximum number of subqueries in a single statement	storage
Maximum number of values in an insert operation ⁷	1012
Maximum number of SET clauses in a single update operation ⁷	1012
Tables and Views	

Table 61. Database Manager Limits (continued)

Description	Limit
Maximum number of columns in a table ⁷	1012
Maximum number of columns in a view ¹	5000
Maximum number of columns in a data source table or view that is referenced by a nickname	5000
Maximum number of columns in a distribution key ⁵	500
Maximum length of a row including all overhead ^{2 7}	32 677
Maximum number of rows in a non-partitioned table, per database partition	128 x 10 ¹⁰
Maximum number of rows in a data partition, per database partition	128 x 10 ¹⁰
Maximum size of a table per database partition in a regular table space (in gigabytes) ^{3 7}	512
Maximum size of a table per database partition in a large DMS table space (in gigabytes) ⁷	16 384
Maximum number of data partitions for a single table	32 767
Maximum number of table partitioning columns	16
Table Spaces	
Maximum size of a LOB object (in terabytes)	4
Maximum size of a LF object (in terabytes)	2
Maximum number of table spaces in a database	32 768
Maximum number of tables in an SMS table space	65 534
Maximum size of a regular DMS table space (in gigabytes) ^{3 7}	512
Maximum size of a large DMS table space (in terabytes) ^{3 7}	16
Maximum size of a temporary DMS table space (in terabytes) ³⁷	16
Maximum number of table objects in a DMS table space ⁶	51 000
Maximum number of storage paths in an automatic storage database	128
Maximum length of a storage path that is associated with an automatic storage database (in bytes)	175
Triggers	
Maximum run-time depth of cascading triggers	16
User-defined Types	
Maximum number of attributes in a structured type	4082

Table 61. Database Manager Limits (continued)

Description	Limit
Note:	
1. This maximum can be achieved using a join in the CREATE VIEW statement. Selecting from such a view is subject to the limit of most elements in a select list.	
2. The actual data for BLOB, CLOB, LONG VARCHAR, DBCLOB, and LONG VARGRAPHIC columns is not included in this count. However, information about the location of that data does take up some space in the row.	
3. The numbers shown are architectural limits and approximations. The practical limits may be less.	
4. The actual value is controlled by the max_connections and max_coordagents database manager configuration parameters.	
5. This is an architectural limit. The limit on the most columns in an index key should be used as the practical limit.	
6. Table objects include data, indexes, LONG VARCHAR or VARGRAPHIC columns, and LOB columns. Table objects that are in the same table space as the table data do not count extra toward the limit. However, each table object that is in a different table space than the table data does contribute one toward the limit for each table object type per table in the table space in which the table object resides.	
7. For page size-specific values, see Table 62.	
8. This is limited only by the longest index key, including all overhead (in bytes). As the number of index key parts increases, the maximum length of each key part decreases.	
9. The maximum can be less, depending on index options.	

Table 62. Database Manager Page Size-specific Limits

Description	4K page size limit	8K page size limit	16K page size limit	32K page size limit
Maximum number of columns in a table	500	1012	1012	1012
Maximum length of a row including all overhead	4005	8101	16 293	32 677
Maximum size of a table per database partition in a regular table space (in gigabytes)	64	128	256	512
Maximum size of a table per database partition in a large DMS table space (in gigabytes)	2048	4096	8192	16 384
Maximum length of an index key including all overhead (in bytes)	1024	2048	4096	8192
Maximum size of an index per database partition in an SMS table space (in gigabytes)	2048	4096	8192	16 384
Maximum size of an index per database partition in a regular DMS table space (in gigabytes)	64	128	256	512

Table 62. Database Manager Page Size-specific Limits (continued)

Description	4K page size limit	8K page size limit	16K page size limit	32K page size limit
Maximum size of an index per database partition in a large DMS table space (in gigabytes)	2048	4096	8192	16 384
Maximum size of an index over XML data per database partition (in terabytes)	2	2	2	2
Maximum size of a regular DMS table space (in gigabytes)	64	128	256	512
Maximum size of a large DMS table space (in gigabytes)	2048	4096	8192	16 384
Maximum size of a temporary DMS table space (in terabytes)	2	4	8	16
Maximum number of elements in a select list	500	1012	1012	1012
Maximum number of columns in a GROUP BY clause	500	1012	1012	1012
Maximum total length of columns in a GROUP BY clause (in bytes)	4005	8101	16 293	32 677
Maximum number of columns in an ORDER BY clause	500	1012	1012	1012
Maximum total length of columns in an ORDER BY clause (in bytes)	4005	8101	16 293	32 677
Maximum number of values in an insert operation	500	1012	1012	1012
Maximum number of SET clauses in a single update operation	500	1012	1012	1012

Chapter 19. Registry and environment variables

Environment variables and the profile registry

Environment and registry variables control your database environment.

You can use the Configuration Assistant (db2ca) to configure registry variables and configuration parameters.

Prior to the introduction of the DB2 database profile registry, changing your environment variables on Windows workstations (for example) required you to change an environment variable and restart. Now, your environment is controlled, with a few exceptions, by registry variables stored in the DB2 profile registries. Users on UNIX operating systems with system administration (SYSADM) authority for a given instance can update registry values for that instance. On Windows, updating profile registry variables requires local Administrator authority or SYSADM authority according to the following conditions:

- If extended security is enabled, SYSADM users must belong to the DB2ADMNS group.
- If extended security is not enabled, SYSADM users can make updates provided that the appropriate permissions have been granted to them in the Windows registry.

Use the db2set command to update registry variables without restarting; this information is stored immediately in the profile registries. However, changes do not affect the currently running DB2 applications or users. The DB2 registry applies the updated information to DB2 server instances and DB2 applications started after the changes are made.

Note: There are DB2 environment variables **DB2INSTANCE**, and **DB2NODE**, which might not be stored in the DB2 profile registries. On some operating systems the set command must be used in order to update these environment variables. These changes are in effect until the next time the system is restarted. On Linux and UNIX platforms, the export command might be used instead of the set command.

Using the profile registry allows for centralized control of the environment variables. Different levels of support are now provided through the different profiles. Remote administration of the environment variables is also available when using the DB2 Administration Server.

There are four profile registries:

- The DB2 Instance-Level Profile Registry. The majority of the DB2 environment variables are placed within this registry. The environment variable settings for a particular instance are kept in this registry. Values defined in this level override their settings in the global level.
- The DB2 Global-Level Profile Registry. If an environment variable is not set for a particular instance, this registry is used. This registry is visible to all instances pertaining to a particular copy of DB2 ESE, one global-level profile exists in the installation path.

- The DB2 Instance Node-Level Profile Registry. This registry level contains variable settings that are specific to a database partition in a partitioned database environment. Values defined in this level override their settings at the instance and global levels.
- The DB2 Instance Profile Registry. This registry contains a list of all instance names associated with the current copy. Each installation has its own list. You can see the complete list of all the instances available on the system by running `db2ilist`.

DB2 configures the operating environment by checking for registry values and environment variables and resolving them in the following order:

1. Environment variables set with the `set` command. (Or the `export` command on UNIX platforms.)
2. Registry values set with the instance node-level profile (using the `db2set -i <instance name> <nodenum>` command).
3. Registry values set with the instance-level profile (using the `db2set -i` command).
4. Registry values set with the global-level profile (using the `db2set -g` command).

Instance-level profile registry

There are a couple of UNIX and Windows differences when working with a partitioned database environment. These differences are shown in the following example.

Assume that there is a partitioned database environment with three physical database partitions that are identified as “red”, “white”, and “blue”. On UNIX platforms, if the instance owner runs the following from any of the database partitions:

```
db2set -i FOO=BAR
```

or

```
db2set FOO=BAR      ('-i' is implied)
```

the value of FOO will be visible to all nodes of the current instance (that is, “red”, “white”, and “blue”).

On UNIX platforms, the instance level profile registry is stored in a text file inside the `sql1ib` directory. In partitioned database environments, the `sql1ib` directory is located on the filesystem shared by all physical database partitions.

On Windows platforms, if the user performs the same command from “red”, the value of FOO will only be visible on “red” of the current instance. The DB2 database manager stores the instance level profile registry inside the Windows registry. There is no sharing across physical database partitions. To set the registry variables on all the physical computers, use the “rah” command as follows:

```
rah db2set -i FOO=BAR
```

rah will remotely run the `db2set` command on “red”, “white”, and “blue”.

It is possible to use **DB2REMOTEPREG** so that the registry variables on non-instance-owning computers are configured to refer to those on the instance

owning computer. This effectively creates an environment where the registry variables on the instance-owning computer are shared amongst all computers in the instance.

Using the example shown above, and assuming that “red” is the owning computer, then one would set **DB2REMOTEPREG** on “white” and “blue” computers to share the registry variables on “red” by doing the following:

```
(on red) do nothing
(on white and blue) db2set DB2REMOTEPREG=\\red
```

The setting for **DB2REMOTEPREG** must not be changed after it is set.

Here is how **REMOTEPREG** works:

When the DB2 database manager reads the registry variables on Windows, it first reads the **DB2REMOTEPREG** value. If **DB2REMOTEPREG** is set, it then opens the registry on the remote computer whose computer name is specified in the **DB2REMOTEPREG** variable. Subsequent reading and updating of the registry variables will be redirected to the specified remote computer.

Accessing the remote registry requires that the Remote Registry Service is running on the target computer. Also, the user logon account and all DB2 service logon accounts have sufficient access to the remote registry. Therefore, to use **DB2REMOTEPREG**, you should operate in a Windows domain environment so that the required registry access can be granted to the domain account.

There are Microsoft Cluster Server (MSCS) considerations. You should not use **DB2REMOTEPREG** in an MSCS environment. When running in an MSCS configuration where all computers belong to the same MSCS cluster, the registry variables are maintained in the cluster registry. Therefore, they are already shared between all computers in the same MSCS cluster and there is no need to use **DB2REMOTEPREG** in this case.

When running in a multi-partitioned failover environment where database partitions span across multiple MSCS clusters, you cannot use **DB2REMOTEPREG** to point to the instance-owning computer because the registry variables of the instance-owning computer reside in the cluster registry.

Declaring, showing, changing, resetting, and deleting registry and environment variables

It is strongly recommended that all specific registry variables be defined in the DB2 database profile registry. If DB2 variables are set outside of the registry, remote administration of those variables is not possible, and the workstation must be restarted in order for the variable values to take effect.

The `db2set` command supports the local declaration of the registry and environment variables.

To display help information for the command, use:

```
db2set -?
```

To list the complete set of all supported registry variables, use:

```
db2set -lr
```

To list all defined registry variables for the current or default instance, use:

```
db2set
```

To list all defined registry variables in the profile registry, use:

```
db2set -all
```

To show the value of a registry variable in the current or default instance, use:

```
db2set registry_variable_name
```

To show the value of a registry variable at all levels, use:

```
db2set registry_variable_name -all
```

To change a registry variable for in the current or default instance, use:

```
db2set registry_variable_name=new_value
```

To change a registry variable default for all databases in the instance, use:

```
db2set registry_variable_name=new_value  
-i instance_name
```

To change a registry variable default for a particular database partition in an instance, use:

```
db2set registry_variable_name=new_value  
-i instance_name database_partition_number
```

To change a registry variable default for all instances pertaining to a particular installation in the system, use:

```
db2set registry_variable_name=new_value -g
```

If you use an aggregate registry variable such as **DB2_WORKLOAD** to configure your registry variables for an SAP environment, you can set that variable using:

```
db2set DB2_WORKLOAD=SAP
```

If you use the Lightweight Directory Access Protocol (LDAP), you can set registry variables in LDAP using:

- To set registry variables at the user level within LDAP, use:

```
db2set -u1
```
- To set registry variables at the global level within LDAP, use:

```
db2set -g1 user_name
```

When running in an LDAP environment, you can set a DB2 registry variable value so that its scope is global to all servers and all users that belong to a directory partition or to a Windows domain. Currently, there are only two DB2 registry variables that can be set at the LDAP global level:

DB2LDAP_KEEP_CONNECTION and **DB2LDAP_SEARCH_SCOPE**.

For example, to set the search scope value at the global level in LDAP, use:

```
db2set -g1 db2ldap_search_scope = value
```

where the value can be local, domain, or global.

Note:

1. When the DB2 profile.env file is updated by two or more users with the db2set command at the same time, or very close to the same time, the size of the profile.env file is reduced to zero. Also, the output from db2set -all displays inconsistent values.
2. There is a difference between the -g option, which is used to set DB2 registry variables that apply to all instances pertaining to the same installation of DB2 ESE, and the -gl option which is specifically used at the LDAP global level.
3. The user level registry variable is only supported on Windows when running in an LDAP environment.
4. Variable settings at the user level contains user specific variable settings. Any changes to the user level are written to the LDAP directory.
5. The parameters "-i", "-g", "-gl", and "-ul" cannot be used at the same time in the same command.
6. Some variables will always default to the global level profile (global means the variables are shared between all instances running on the same DB2 copy). They cannot be set at the instance or database partition level profiles; for example, **DB2SYSTEM** and **DB2INSTDEF**.
7. On UNIX, you must have system administration (SYSADM) authority to change registry values for an instance. Only users with root authority can change parameters in global-level registries.

To reset a registry variable for an instance back to the default found in the Global Profile Registry, use:

```
db2set -r registry_variable_name
```

To reset a registry variable for a database partition in an instance back to the default found in the Global Profile Registry, use:

```
db2set -r registry_variable_name database_partition_number
```

To delete a variable's value at a specified level, you can use the same command syntax to set the variable but specify nothing for the variable value. For example, to delete the variable's setting at the database partition level, enter:

```
db2set registry_variable_name= -i instance_name
database_partition_number
```

To delete a variable's value and to restrict its use, if it is defined at a higher profile level, enter:

```
db2set registry_variable_name= -null -r instance_name
```

This command deletes the setting for the parameter you specify and restricts high level profiles from changing this variable's value (in this case, DB2 global-level profile). However, the variable you specify could still be set by a lower level profile (in this case, the DB2 database partition-level profile).

Setting environment variables on Windows

Windows operating systems have one system environment variable, **DB2INSTANCE**, that can only be set outside the profile registry; however, you are not required to set **DB2INSTANCE**. The DB2 profile registry variable **DB2INSTDEF** might be set in the global level profile to specify the instance name to use if **DB2INSTANCE** is not defined.

DB2 Enterprise Server Edition servers on Windows have two system environment variables, **DB2INSTANCE** and **DB2NODE**, that can only be set outside the profile

registry. You are not required to set **DB2INSTANCE**. The DB2 profile registry variable **DB2INSTDEF** might be set in the global level profile to specify the instance name to use if **DB2INSTANCE** is not defined.

The **DB2NODE** environment variable is used to route requests to a target logical node within a computer. This environment variable must be set in the session in which the application or command is issued and not in the DB2 profile registry. If this variable is not set, the target logical node defaults to the logical node which is defined as zero (0) on the computer.

To determine the settings of an environment variable, use the echo command. For example, to check the value of the **DB2PATH** environment variable, enter:

```
echo %db2path%
```

You can set the DB2 environment variables **DB2INSTANCE** and **DB2NODE** as follows (using **DB2INSTANCE** in this description):

- Right click on My Computer and select Properties .
- Select the Advanced tab, click Environment Variables, and do the following:
 1. If the **DB2INSTANCE** variable does not exist:
 - a. Click New.
 - b. Fill in the Variable Name field with **DB2INSTANCE**.
 - c. Fill in the Variable Value field with the instance name, for example db2inst.
 2. If the **DB2INSTANCE** variable already exists, append a new value:
 - a. Select the **DB2INSTANCE** environment variable.
 - b. Change the Value field to the instance name, for example db2inst.
 3. Restart your system for these changes to take effect.

Note: The environment variable **DB2INSTANCE** can also be set at the session (process) level. For example, if you want to start a second DB2 instance called TEST, issue the following commands in a command window:

```
set DB2INSTANCE=TEST
db2start
```

When working in C Shell, issue the following commands in a command window:

```
setenv DB2INSTANCE TEST
```

The profile registries are located as follows:

- The DB2 Instance-Level Profile Registry in the Windows operating system registry, with the path:

```
\HKEY_LOCAL_computer\SOFTWARE\IBM\DB2\PROFILES\instance_name
```

Note: The *instance_name* is the name of the DB2 instance.

- The DB2 Global-Level Profile Registry in the Windows registry, with the path:

```
\HKEY_LOCAL_computer\SOFTWARE\IBM\DB2\GLOBAL_PROFILE
```

- The DB2 Instance Node-Level Profile Registry in the Windows registry, with the path:

```
...\SOFTWARE\IBM\DB2\PROFILES\instance_name\NODES\node_number
```

Note: The *instance_name* and the *node_number* are specific to the database partition you are working with.

- There is no DB2 Instance Profile Registry required. For each of the DB2 instances in the system, a key is created in the path:

```
\HKEY_LOCAL_computer\SOFTWARE\IBM\DB2\PROFILES\instance_name
```

The list of instances can be obtained by counting the keys under the PROFILES key.

Setting environment variables on Linux and UNIX operating systems

On UNIX operating systems, you must set the system environment variable **DB2INSTANCE**.

The scripts `db2profile` (for Bourne or Korn shell) and `db2cshrc` (for C shell) are provided as examples to help you set up the database environment. You can find these files in `insthome/sqllib`, where `insthome` is the home directory of the instance owner.

These scripts include statements to:

- Update a user's path with the following directories:
 - `insthome/sqllib/bin`
 - `insthome/sqllib/adm`
 - `insthome/sqllib/misc`
- Set **DB2INSTANCE** to the default local `instance_name` for execution.

Note: Except for `PATH` and **DB2INSTANCE**, all other supported variables must be set in the DB2 profile registry. To set variables that are not supported by the DB2 database manager, define them in your script files, `userprofile` and `usercshrc`.

An instance owner or `SYSADM` user might customize these scripts for all users of an instance. Alternatively, users can copy and customize a script, then invoke a script directly or add it to their `.profile` or `.login` files.

To change the environment variable for the current session, issue commands similar to the following:

- For Korn shell:


```
DB2INSTANCE=<inst1>
export DB2INSTANCE
```
- For Bourne shell:


```
export DB2INSTANCE=<inst1>
```
- For C shell:


```
setenv DB2INSTANCE <inst1>
```

In order for the DB2 profile registry to be administered properly, the following file ownership rules must be followed on UNIX operating systems.

- The DB2 Instance-Level Profile Registry file is located under:


```
INSTHOME/sqllib/profile.env
```

The access permissions and ownership of this file should be:

```
-rw-rw-r-- <db2inst1> <db2iadm1> profile.env
```

where <db2inst1> is the instance owner, and <db2iadm1> is the instance owner's group.

The INSTHOME is the home path of the instance owner.

- The DB2 Global-Level Profile Registry is located under:
<installation path>/default.env

for all Linux and UNIX platforms.

The access permissions and ownership of this file should be:

```
-rw-rw-r-- root <group> default.env
```

where <group> is the group name for group ID 0; for example, on AIX, it is "system".

To modify a global registry variable in root installations, you must be logged on with root authority.

- The DB2 Instance Node-Level Profile Registry is located under:
INSTHOME/sqllib/nodes/<node_number>.env

The access permissions and ownership of the directory and this file should be:

```
drwxrwsr-w <Instance_Owner> <Instance_Owner_Group> nodes
```

```
-rw-rw-r-- <Instance_Owner> <Instance_Owner_Group> <node_number>.env
```

The INSTHOME is the home path of the instance owner.

- The DB2 Instance Profile Registry is located under:
<installation path>/profiles.reg

for all Linux and UNIX platforms

The access permissions and ownership of this file should be:

```
-rw-r--r-- root system profiles.reg
```

Setting the current instance environment variables

When you run commands to start or stop an instance's database manager, DB2 applies the command to the current instance. DB2 determines the current instance as follows:

- If the **DB2INSTANCE** environment variable is set for the current session, its value is the current instance. To set the **DB2INSTANCE**, enter:
set db2instance=<new_instance_name>
- If **DB2INSTANCE** is not set for the current session, the DB2 database manager uses the setting for the **DB2INSTANCE** environment variable from the system environment variables. On Windows, system environment variables are set in the System Environment registry.
- If **DB2INSTANCE** is not set at all, the DB2 database manager uses the registry variable, **DB2INSTDEF**.

To set the **DB2INSTDEF** registry variable at the global level of the registry, enter:

```
db2set db2instdef=<new_instance_name> -g
```

To determine which instance applies to the current session, enter:

```
db2 get instance
```

Aggregate registry variables

An aggregate registry variable allows several registry variables to be grouped as a configuration that is identified by another registry variable name. Each registry variable that is part of the group has a predefined setting. The aggregate registry variable is given a value that is interpreted as declaring several registry variables.

The intention of an aggregate registry variable is to ease registry configuration for broad operational objectives.

The only valid aggregate registry variable is **DB2_WORKLOAD**.

Valid values for this variable are:

- SAP
- 1C
- TPM
- CM
- WC

Any registry variable that is implicitly configured through an aggregate registry variable might also be explicitly defined. Explicitly setting a registry variable that was previously given a value through the use of an aggregate registry variable is useful when doing performance or diagnostic testing. Explicitly setting a variable that is configured implicitly by an aggregate is referred to as overriding the variable.

If you attempt to modify an explicitly set registry variable by using an aggregate registry variable, a warning is issued and the explicitly set value is kept. This warning tells you that the explicit value is maintained. If the aggregate registry variable is used first and then you specify an explicit registry variable, a warning is not given.

None of the registry variables that are configured through setting an aggregate registry variable are shown unless you explicitly make that request for each variable. When you query the aggregate registry variable, only the value assigned to that variable is shown. Most users should not care about the values for each individual variable.

The following example shows the interaction between using the aggregate registry variable and explicitly setting a registry variable. For example, you might have set the **DB2_WORKLOAD** aggregate registry variable to SAP and have overridden the **DB2_SKIPDELETED** registry variable to NO. By entering db2set, you would receive the following results:

```
DB2_WORKLOAD=SAP
DB2_SKIPDELETED=NO
```

In another situation, you might have set **DB2ENVLIST**, set the **DB2_WORKLOAD** aggregate registry variable to SAP, and overridden the **DB2_SKIPDELETED** registry variable to NO. (This assumes that the **DB2_SKIPDELETED** registry variable is part of the group making up the SAP environment.) In addition, those registry variables that were configured automatically through setting the aggregate registry variable will show the name of the aggregate displayed within square brackets, adjacent to its value. The **DB2_SKIPDELETED** registry variable will show a NO value and will show [0] displayed adjacent to its value.

When you no longer require the configuration associated with **DB2_WORKLOAD**, you can disable the implicit values of each registry variable in the group by deleting the aggregate registry variable's value using the command:

```
db2set DB2_WORKLOAD=
```

After deleting the **DB2_WORKLOAD** aggregate registry variable value, restart the database. After the database is restarted, the registry variables that were implicitly configured by the aggregate registry variable are no longer in effect. The method used to delete an aggregate registry variable's value is the same as deleting an individual registry variable.

Deleting an aggregate registry variable's value does not delete a registry variable's value that has been explicitly set. It does not matter that the registry variable is a member of the group definition being disabled. The explicit setting for the registry variable is maintained.

You might need to see the values for each registry variable that is a member of the **DB2_WORKLOAD** aggregate registry variable. For instance, you might want to see the values that would be used if you configured **DB2_WORKLOAD** to SAP. To find the values that would be used if **DB2_WORKLOAD=SAP**, run `db2set -gd DB2_WORKLOAD=SAP`.

DB2 registry and environment variables

DB2 database products provide a number of registry variables and environment variables that you might need to know about to get up and running.

To view a list of all supported registry variables, execute the following command:

```
db2set -lr
```

To change the value for a variable in the current or default instance, execute the following command:

```
db2set registry_variable_name=new_value
```

Whether the DB2 environment variables **DB2INSTANCE**, **DB2NODE**, **DB2PATH**, and **DB2INSTPROF** are stored in the DB2 profile registries depends on your operating system. To update these environment variables, use the set command. These changes are only effective in the local (current) command prompt and are in effect until the next time the system is rebooted. On Linux and UNIX operating systems, you can use the export command instead of the set command.

You must set the values for the changed registry variables before you execute the `db2start` command.

Note: If a registry variable requires Boolean values as arguments, the values YES, 1, and ON are all equivalent and the values NO, 0, and OFF are also equivalent. For any variable, you can specify any of the appropriate equivalent values.

The following table lists all registry variables per category:

Table 63. Registry and environment variables summary

Variable category	Registry or environment variable name
General	DB2ACCOUNT DB2BIDI DB2_CAPTURE_LOCKTIMEOUT DB2CODEPAGE DB2_COLLECT_TS_REC_INFO DB2_CONNRETRIES_INTERVAL DB2CONSOLECP DB2COUNTRY DB2DBDFT DB2DBMSADDR DB2DISCOVERYTIME DB2FFDC DB2FODC DB2_FORCE_APP_ON_MAX_LOG DB2GRAPHICUNICODESERVER DB2INCLUDE DB2INSTDEF DB2INSTOWNER DB2_LIC_STAT_SIZE DB2LOCALE DB2_MAX_CLIENT_CONNRETRIES DB2_OBJECT_TABLE_ENTRIES DB2_SYSTEM_MONITOR_SETTINGS DB2TERRITORY DB2_VIEW_REOPT_VALUES
System environment	DB2_ALTERNATE_GROUP_LOOKUP DB2_CLP_EDITOR DB2_CLP_HISTSIZ DB2CONNECT_IN_APP_PROCESS DB2_COPY_NAME DB2DBMSADDR DB2_DIAGPATH DB2DOMAINLIST DB2ENVLIST DB2INSTANCE DB2INSTPROF DB2LDAPSecurityConfig DB2LIBPATH DB2LOGINRESTRICTIONS DB2NODE DB2OPTIONS DB2_PARALLEL_IO DB2PATH DB2PROCESSORS DB2RCMD_LEGACY_MODE DB2SYSTEM DB2_UPDDBCFG_SINGLE_DBPARTITION DB2_USE_PAGE_CONTAINER_TAG DB2_WORKLOAD

Table 63. Registry and environment variables summary (continued)

Variable category	Registry or environment variable name
Communications	DB2CHECKCLIENTINTERVAL DB2COMM DB2FCMCOMM DB2_FORCE-NLS_CACHE DB2RSHCMD DB2RSHTIMEOUT DB2SORCVBUF DB2SOSNDBUF DB2TCP_CLIENT_CONTIMEOUT DB2TCP_CLIENT_RCVTIMEOUT DB2TCPCONNMGRS
Command-line	DB2BQTIME DB2BQTRY DB2_CLPPROMPT DB2IQTIME DB2RQTIME
Partitioned database environment	DB2CHGPWD_EEE DB2_FCM_SETTINGS DB2_NUM_FAILOVER_NODES DB2_PARTITIONEDLOAD_DEFAULT DB2PORTRANGE
Query compiler	DB2_ANTIJOIN DB2_INLIST_TO_NLJN DB2_LIKE_VARCHAR DB2_MINIMIZE_LISTPREFETCH DB2_NEW_CORR_SQ_FF DB2_OPT_MAX_TEMP_SIZE DB2_REDUCED_OPTIMIZATION DB2_SELECTIVITY DB2_SQLROUTINE_PREPOPTS

Table 63. Registry and environment variables summary (continued)

Variable category	Registry or environment variable name
Performance	DB2_ALLOCATION_SIZE DB2_APM_PERFORMANCE DB2ASSUMEUPDATE DB2_ASYNC_IO_MAXFILOP DB2_AVOID_PREFETCH DB2BPVARS DB2CHKPTR DB2CHKSQLDA DB2_EVALUNCOMMITTED DB2_EXTENDED_IO_FEATURES DB2_EXTENDED_OPTIMIZATION DB2_IO_PRIORITY_SETTING DB2_IO_PRIORITY_SETTING DB2_KEEP_AS_AND_DMS_CONTAINERS_OPEN DB2_KEEPTABLELOCK DB2_LARGE_PAGE_MEM DB2_LOGGER_NON_BUFFERED_IO DB2MAXFSCRSEARCH DB2_MAX_INACT_STMTS DB2_MAX_NON_TABLE_LOCKS DB2_MDC_ROLLOUT DB2MEMDISCLAIM DB2MEMMAXFREE DB2_MEM_TUNING_RANGE DB2_MMAP_READ DB2_MMAP_WRITE DB2_NO_FORK_CHECK DB2NTMEMSIZE DB2NTNOCACHE DB2NTPRICLASS DB2NTWORKSET DB2_OVERRIDE_BPF DB2_PINNED_BP DB2PRIORITIES DB2_RESOURCE_POLICY DB2_RCT_FEATURES DB2_SET_MAX_CONTAINER_SIZE DB2_SKIPDELETED DB2_SKIPINSERTED DB2_SMS_TRUNC_TMPTABLE_THRESH DB2_SORT_AFTER_TQ DB2_SELUDI_COMM_BUFFER DB2_TRUSTED_BINDIN DB2_USE_ALTERNATE_PAGE_CLEANING DB2_USE_IOCP

Table 63. Registry and environment variables summary (continued)

Variable category	Registry or environment variable name
Miscellaneous	DB2ADMINSERVER DB2_ATS_ENABLE DB2AUTH DB2CLIINIPATH DB2_COMMIT_ON_EXIT DB2_CREATE_DB_ON_PATH DB2DEFPREP DB2_DISABLE_FLUSH_LOG DB2_DISPATCHER_PEEKTIMEOUT DB2_DJ_INI DB2DMNBCKCTLR DB2_DOCHOST DB2_DOCPORT DB2_ENABLE_AUTOCONFIG_DEFAULT DB2_ENABLE_LDAP DB2_EVMON_EVENT_LIST_SIZE DB2_EVMON_STMT_FILTER DB2_EXTSECURITY DB2_FALLBACK DB2_FMP_COMM_HEAPSZ DB2_GRP_LOOKUP DB2_HADR_BUF_SIZE DB2_HADR_NO_IP_CHECK DB2_HADR_PEER_WAIT_LIMIT DB2_HADR_SORCVBUF DB2_HADR_SOSNDBUF DB2LDAP_BASEDN DB2LDAPCACHE DB2LDAP_CLIENT_PROVIDER DB2LDAPHOST DB2LDAP_KEEP_CONNECTION DB2LDAP_SEARCH_SCOPE DB2_LOAD_COPY_NO_OVERRIDE DB2LOADREC DB2LOCK_TO_RB DB2_MAP_XML_AS_CLOB_FOR_DLC DB2_MAX_LOB_BLOCK_SIZE DB2_MEMORY_PROTECT DB2NOEXITLIST DB2_NUM_CKPW_DAEMONS DB2_OPTSTATS_LOG DB2REMOTEPEG DB2_RESOLVE_CALL_CONFLICT DB2ROUTINE_DEBUG DB2SATELLITEID DB2_SERVER_CONTIMEOUT DB2_SERVER_ENCALG DB2SORT DB2_THREAD_SUSPENSION DB2_TRUNCATE_REUSESTORAGE DB2_USE_DB2JCCT2_JROUTINE DB2_UTIL_MSGPATH DB2_VENDOR_INI DB2_XBSA_LIBRARY

General registry variables

DB2ACCOUNT

- Operating system: All
- Default: NULL
- This variable defines the accounting string that is sent to the remote host. Refer to the DB2 Connect User's Guide for details.

DB2BIDI

- Operating system: All
- Default: NO, Values: YES or NO
- This variable enables bidirectional support and the **DB2CODEPAGE** variable is used to declare the code page to be used.

DB2_CAPTURE_LOCKTIMEOUT

- Operating system: All
- Default: NULL, Values: ON or NULL
- This variable specifies to log descriptive information about lock timeouts at the time that they occur. The logged information identifies: the key applications involved in the lock contention that resulted in the lock timeout, the details about what these applications were running at the time of the lock timeout, and the details about the lock causing the contention. Information is captured for both the lock requestor (the application that received the lock timeout error) and the current lock owner. A text report is written and stored in a file for each lock timeout.

The files are created using the following naming convention:

db2locktimeout.par.AGENTID.yyyy-mm-dd-hh-mm-ss, where *par* is the database partition number; *AGENTID* is the Agent ID; *yyyy-mm-dd-hh-mm-ss* is the timestamp consisting of the year, month, day, hour, minute and second. In non-partitioned database environments, *par* is set to 0.

The location of the file is based on the value set in the **diagpath** database configuration parameter. If **diagpath** is not set, then the file is located in one of the following directories:

- In Windows environments:
 - If you do not set the **DB2INSTPROF** environment variable, information is written to *x:\SQLLIB\DB2INSTANCE*, where *x* is the drive reference, *SQLLIB* is the directory that you specified for the **DB2PATH** registry variable, and *DB2INSTANCE* is the name of the instance owner.
 - If you set the **DB2INSTPROF** environment variable, information is written to *x:\DB2INSTPROF\DB2INSTANCE*, where *x* is the drive reference, *DB2INSTPROF* is the name of the instance profile directory, and *DB2INSTANCE* is the name of the instance owner.
- In Linux and UNIX environments: information is written to *INSTHOME/sqllib/db2dump*, where *INSTHOME* is the home directory of the instance.

Delete lock timeout report files when you no longer need them. Because the report files are in the same location as other diagnostics logs, the DB2 system could shutdown if the directory is allowed to get full. If you need to keep some lock timeout report files, move them to a directory or folder different than where the DB2 logs are stored.

DB2CODEPAGE

- Operating system: All
- Default: derived from the language ID, as specified by the operating system.
- This variable specifies the code page of the data presented to DB2 for database client application. The user should not set **DB2CODEPAGE** unless explicitly stated in DB2 documents, or asked to do so by DB2 service. Setting **DB2CODEPAGE** to a value not supported by the operating system can produce unexpected results. Normally, you do not need to set **DB2CODEPAGE** because DB2 automatically derives the code page information from the operating system.

Note: Because Windows does not report a Unicode code page (in the Windows regional settings) instead of the ANSI code page, a Windows application will not behave as a Unicode client. To override this behavior, set the **DB2CODEPAGE** registry variable to 1208 (for the Unicode code page) to cause the application to behave as a Unicode application.

DB2_COLLECT_TS_REC_INFO

- Operating system: All
- Default: ON, Values: ON or OFF
- This variable specifies whether DB2 will process all log files when rolling forward a table space, regardless of whether the log files contain log records that affect the table space. To skip the log files known not to contain any log records affecting the table space, set this variable to ON. **DB2_COLLECT_TS_REC_INFO** must be set before the log files are created and used so that the information required for skipping log files is collected.

DB2_CONNRETRIES_INTERVAL

- Operating system: All
- Default: Not set, Values: an integer number of seconds
- This variable specifies the sleep time between consecutive connection retries, in seconds, for the automatic client reroute feature. You can use this variable in conjunction with **DB2_MAX_CLIENT_CONNRETRIES** to configure the retry behavior for automatic client reroute.

If **DB2_MAX_CLIENT_CONNRETRIES** is set, but **DB2_CONNRETRIES_INTERVAL** is not, **DB2_CONNRETRIES_INTERVAL** defaults to 30. If **DB2_MAX_CLIENT_CONNRETRIES** is not set, but **DB2_CONNRETRIES_INTERVAL** is set, **DB2_MAX_CLIENT_CONNRETRIES** defaults to 10. If neither **DB2_MAX_CLIENT_CONNRETRIES** nor **DB2_CONNRETRIES_INTERVAL** is set, the automatic client reroute feature reverts to its default behavior of retrying the connection to a database repeatedly for up to 10 minutes.

DB2CONSOLECP

- Operating system: Windows
- Default: NULL, Values: all valid code page values
- Specifies the code page for displaying DB2 message text. When specified, this value overrides the operating system code page setting.

DB2COUNTRY

- Operating system: Windows

- Default: NULL, Values: all valid numeric country, territory, or region codes
- This variable specifies the country, territory, or region code of the client application. When specified, this value overrides the operating system setting.

Note: **DB2COUNTRY** is deprecated and may be removed in a future release. Instead, use **DB2TERRITORY**, which accepts the same values as **DB2COUNTRY**

DB2DBDFT

- Operating system: All
- Default: NULL
- This variable specifies the database alias name of the database to be used for implicit connects. If an application has no database connection but SQL or XQuery statements are issued, an implicit connect will be made if the **DB2DBDFT** environment variable has been defined with a default database.

DB2DBMSADDR

- Operating system: Windows 32-bit
- Default: 0x20000000, Values: 0x20000000 to 0xB0000000 in increments of 0x10000
- This variable specifies the default database manager shared memory address in hexadecimal format. If db2start fails due to a shared memory address collision, this registry variable can be modified to force the database manager instance to allocate its shared memory at a different address.

DB2DISCOVERYTIME

- Operating system: Windows
- Default: 40 seconds, Minimum: 20 seconds
- This variable specifies the amount of time that SEARCH discovery will search for DB2 systems.

DB2_EXPRESSION_RULES

- Operating system: All
- Default: Empty, Values: RAISE_ERROR_PERMIT_SKIP or RAISE_ERROR_PERMIT_DROP
- The settings for the **DB2_EXPRESSION_RULES** registry variable control how the DB2 Optimizer determines the access plan for queries which involve a RAISE_ERROR function. The default behaviour of the RAISE_ERROR function is that no filtering may be pushed beyond the expression containing this function. This can result in no predicates being applied during the table accesses which can lead to excessive computation of expressions, excessive locking and poor query performance.

In certain cases this behaviour is too strict, depending on the particular business requirements of the application, it may not matter if predicates and joins are applied before the application of RAISE_ERROR. For example in the context of a row level security implementation, there is typically an expression of the form:

```

CASE WHEN <conditions for validatin access to this row>
      THEN NULL
      ELSE RAISE_ERROR(...)
END

```

The application may only be concerned with validating access to the rows which are selected by the query and not in validating access to every row in the table. Thus predicates could be applied in the base table access and the expression containing the RAISE_ERROR only needs to be executed after all the filtering is performed. In this case a value of **DB2_EXPRESSION_RULES=RAISE_ERROR_PERMIT_SKIP** may be appropriate.

Another alternative is in the context of COLUMN LEVEL security. In this case there are typically expressions of the form:

```

CASE WHEN <conditions for validating access to this row and column>
      THEN <table.column>
      ELSE RAISE_ERROR(...)
END

```

In this case the application may only want errors to be raised if the user attempts to receive the data for a particular row and column contains a value that the user is not allowed to retrieve. In this case a setting of **DB2_EXPRESSION_RULES=RAISE_ERROR_PERMIT_DROP** will only cause the expression containing the RAISE_ERROR function to be evaluated if the particular column is used by a predicate, column function or if it is returned as output from the query.

DB2FFDC

- Operating system: All
- Default: ON, Values: ON, CORE:OFF
- This variable provides the ability to deactivate core file generation. By default, this registry variable is set to ON. If this registry variable is not set, or is set to a value other than CORE:OFF, core files may be generated if the DB2 database system abends.

Core files, which are used for problem determination and are created in the **diagpath** directory, contain the entire process image of the terminating DB2 process. Consideration should be given to the available file system space because core files can be quite large. The size is dependent on the DB2 database system configuration and the state of the process at the time the problem occurs.

On Linux operating systems, the default core file size limit is set to 0 (that is, `ulimit -c`). With this setting, core files are not generated. To allow core files to be created on Linux platforms, set the value to unlimited.

Note: **DB2FFDC** is being deprecated in version 9.5, and will be removed in a later release. The new registry variable **DB2FODC** incorporates **DB2FFDC**'s functionality.

DB2FODC

- Operating system: All
- Default: The concatenation of all First Occurrence Data Collection (FODC) parameters (see below)
 - for Linux and UNIX: `"CORELIMIT=val DUMPCORE=ON DUMPPDIR=diagpath"`

- for Windows: "DUMPCORE=ON DUMPPDIR=*diagpath*"

Note that the parameters are separated by spaces.

- This registry variable controls a set of troubleshooting-related parameters used in FODC. Use **DB2FODC** to control different aspects of data collection in outage situations.

This registry variable is read once, during the DB2 instance startup. To perform updates to the FODC parameters online, use db2pdcfg tool. Use the **DB2FODC** registry variable to sustain the configuration across reboots. You do not need to specify all of the parameters, nor do you need to specify them in a particular order. The default value is assigned to any parameter that is not specified. For example, if you don't want the core files dumped, but you do want the other parameters' default behaviors, you would issue the command:

```
db2set DB2FODC="DUMPCORE=OFF"
```

Parameters:

CORELIMIT

- Operating system: Linux and UNIX
- Default: Current ulimit setting, Values: 0 to unlimited
- This option specifies the maximum size, in gigabytes, of core files created. This value overrides the current core file size limit setting. Consideration should be given to the available file system space because core files can be quite large. The size is dependent on the DB2 configuration and the state of the process at the time the problem occurs.

If **CORELIMIT** is set, DB2 will use this value override current user core limit (ulimit) setting to generate the core file.

If **CORELIMIT** is not set, DB2 will set the core file size to the value equal to the current ulimit setting. One exception is AIX where an ulimit setting of "unlimited" will be overridden with a value of 8 GB for DB2 server processes only. If you require a core dump larger than 8 GB, set ulimit to an appropriately large value, such as the size of RAM, or set **CORELIMIT** with a sufficiently large value.

Note: Any changes to the user core limit or **CORELIMIT** are not effective until the next recycling of the DB2 instance.

DUMPCORE

- Operating system: Linux, Solaris, AIX
- Default: ON, Values: ON or OFF
- This option specifies if core file generation is to take place. Core files, which are used for problem determination and are created in the **diagpath** directory, contain the entire process image of the terminating DB2 process. However, whether or not an actual core file dump occurs depends on the current ulimit setting and value of the **CORELIMIT** parameter. Some operating systems also have configuration settings for core dumps, which may dictate the behavior of application core dumping.

The recommended method for disabling core file dumps is to set **DUMPCORE** to OFF.

DUMPDIR

- Operating system: All
- Default: **diagpath** directory, or the default diagnostic directory if **diagpath** is not defined, Values: *path to directory*
- This option specifies the absolute path name of the directory for core file creation. This option may be used for other large binary dumps that have to be stored outside of FODC package, not for only core files.

DB2_FORCE_APP_ON_MAX_LOG

- Operating system: All
- Default: TRUE, Values: TRUE or FALSE
- Specifies what happens when the **max_log** configuration parameter value is exceeded. If set to TRUE, the application is forced off the database and the unit of work is rolled back.

If FALSE, the current statement fails. The application can still commit the work completed by previous statements in the unit of work, or it can roll back the work completed to undo the unit of work.

Note: This DB2 registry variable affects the ability of the import utility to recover from log full situations. If **DB2_FORCE_APP_ON_MAX_LOG** is set to TRUE and you issue an IMPORT command with the **COMMITCOUNT** command option, the import utility will not be able to perform a commit in order to avoid running out of active log space. When the import utility encounters an SQL0964C (Transaction Log Full), it will be forced off the database and the current unit of work will be rolled back.

DB2GRAPHICUNICODESERVER

- Operating system: All
- Default: OFF, Values: ON or OFF
- This registry variable is used to accommodate existing applications written to insert graphic data into a Unicode database. Its use is only needed for applications that specifically send sqldbchar (graphic) data in Unicode instead of the code page of the client. (sqldbchar is a supported SQL data type in C and C++ that can hold a single double-byte character.) When set to ON, you are telling the database that graphic data is coming in Unicode, and the application expects to receive graphic data in Unicode.

DB2INCLUDE

- Operating system: All
- Default: Current directory
- Specifies a path to be used during the processing of the SQL INCLUDE text-file statement during DB2 PREP processing. It provides a list of directories where the INCLUDE file might be found. Refer to Developing Embedded SQL Applications for descriptions of how **DB2INCLUDE** is used in the different precompiled languages.

DB2INSTDEF

- Operating system: Windows
- Default: DB2
- This variable sets the value to be used if **DB2INSTANCE** is not defined.

DB2INSTOWNER

- Operating system: Windows
- Default: NULL
- The registry variable created in the DB2 profile registry when the instance is first created. This variable is set to the name of the instance-owning machine.

DB2_LIC_STAT_SIZE

- Operating system: All
- Default: NULL, Range: 0 to 32767
- This variable determines the maximum size (in MBs) of the file containing the license statistics for the system. A value of zero turns the license statistic gathering off. If the variable is not recognized or not defined, the variable defaults to unlimited. The statistics are displayed using the License Center.

DB2LOCALE

- Operating system: All
- Default: NO, Values: YES or NO
- This variable specifies whether the default "C" locale of a process is restored to the default "C" locale after calling DB2 and whether to restore the process locale back to the original 'C' after calling a DB2 function. If the original locale was not 'C', then this registry variable is ignored.

DB2_MAX_CLIENT_CONNRETRIES

- Operating system: All
- Default: Not set, Values: an integer number of maximum times to retry the connection
- This variable specifies the maximum number of connection retries that the automatic client reroute feature will attempt. You can use this variable in conjunction with **DB2_CONNRETRIES_INTERVAL** to configure the retry behavior for automatic client reroute.

If **DB2_MAX_CLIENT_CONNRETRIES** is set, but **DB2_CONNRETRIES_INTERVAL** is not, **DB2_CONNRETRIES_INTERVAL** defaults to 30. If **DB2_MAX_CLIENT_CONNRETRIES** is not set, but **DB2_CONNRETRIES_INTERVAL** is set, **DB2_MAX_CLIENT_CONNRETRIES** defaults to 10. If neither **DB2_MAX_CLIENT_CONNRETRIES** nor **DB2_CONNRETRIES_INTERVAL** is set, the automatic client reroute feature reverts to its default behavior of retrying the connection to a database repeatedly for up to 10 minutes.

DB2_OBJECT_TABLE_ENTRIES

- Operating system: All
 - Default: 0, Values: 0–65532
- The actual maximum value possible on your system depends on the page size and extent size, but it cannot exceed 65532.
- This variable specifies the expected number of objects in a table space. If you know that a large number of objects (for example, 1000 or more) will be created in a DMS table space, you should set this registry variable to the approximate number before creating the table space. This will reserve contiguous storage for object metadata during table space

creation. Reserving contiguous storage reduces the chance that an online backup will block operations which update entries in the metadata (for example, CREATE INDEX, IMPORT REPLACE). It will also make resizing the table space easier because the metadata will be stored at the start of the table space.

If the initial size of the table space is not large enough to reserve the contiguous storage, the table space creation will continue without the additional space reserved.

DB2_SYSTEM_MONITOR_SETTINGS

- Operating system: All
- The registry variable controls a set of parameters which allow you to modify the behavior of various aspects of DB2 monitoring. Separate each parameter by a semicolon, as in the following example:

```
db2set DB2_SYSTEM_MONITOR_SETTINGS=OLD_CPU_USAGE:TRUE;  
DISABLE_CPU_USAGE:TRUE
```

Every time you set **DB2_SYSTEM_MONITOR_SETTINGS**, each parameter must be set explicitly. Any parameter that you do not specify when setting this variable reverts back to its default value. So in the following example:

```
db2set DB2_SYSTEM_MONITOR_SETTINGS=DISABLE_CPU_USAGE:TRUE
```

Note: Currently, this registry variable only has settings for Linux; additional settings for other operating systems will be added in future releases.

OLD_CPU_USAGE will be restored to its default setting.

- Parameters:

OLD_CPU_USAGE

- Operating system: Linux
- Values: TRUE/ON, FALSE/OFF
- Default value on RHEL4 and SLES9: TRUE (Note: a setting of FALSE for OLD_CPU_USAGE will be ignored—only the old behavior will be used.)
- Default value on RHEL5, SLES10, and others: FALSE
- This parameter controls how the instance obtains CPU usage times on Linux platforms. If set to TRUE, the older method of getting CPU usage time is used. This method returns both system and user CPU usage times, but consumes more CPU in doing so (that is, it has a higher overhead). If set to FALSE, the newer method of getting CPU usage is used. This method returns only the user CPU usage value, but is faster because it has less overhead.

DISABLE_CPU_USAGE

- Operating system: Linux
- Values: TRUE/ON, FALSE/OFF
- Default value on RHEL4 and SLES9: TRUE
- Default value on RHEL5, SLES10, and others: FALSE
- This parameter allows you to determine whether CPU usage is read or not. When DISABLE_CPU_USAGE is enabled (set

to TRUE), CPU usage is not read, allowing you to avoid the overhead that can sometimes occur during the retrieval of CPU usage.

DB2TERRITORY

- Operating system: All
- Default: Derived from the language ID, as specified by the operating system.
- This variable specifies the region, or territory code of the client application, which influences date and time formats.

DB2_VIEW_REOPT_VALUES

- Operating system: All
- Default: NO, Values: YES, NO
- This variable enables all users to store the cached values of a reoptimized SQL or XQuery statement in the EXPLAIN_PREDICATE table when the statement is explained. When this variable is set to NO, only DBADM is allowed to save these values in the EXPLAIN_PREDICATE table.

System environment variables

DB2_ALTERNATE_GROUP_LOOKUP

- Operating system: AIX
- Default: NULL, Values: NULL or GETGRSET
- This variable allows DB2 to obtain group information from an alternative source provided by the operating system. On AIX, the function getgrset is used. This provides the ability to obtain groups from somewhere other than local files via Loadable Authentication Modules.

DB2_CLP_EDITOR

- Operating system: All
- Default: Notepad(Windows), vi (UNIX), Values: Any valid editor that is located in the operating system path

Note: This registry variable is not set to the default value during installation. Instead, the code that makes use of this variable uses a default value if the registry variable is not set.

- This variable determines the editor to be used when executing the EDIT command. From a CLP interactive session, the EDIT command launches an editor preloaded with a user-specified command which can then be edited and run.

DB2_CLP_HISTSZ

- Operating system: All
- Default: 20, Values: 1–500 inclusive

Note: This registry variable is not set to the default value during installation. Instead, the code that makes use of this variable uses a default value of 20 if the registry variable is not set or if it is set to a value outside of the valid range.

- This variable determines the number of commands stored in the command history during CLP interactive sessions. Because the command

history is held in memory, a very high value for this variable might result in a performance impact depending on the number and length of commands run in a session.

DB2CONNECT_IN_APP_PROCESS

- Operating system: All
- Default: YES, Values: YES or NO
- When you set this variable to NO, local DB2 Connect clients on a DB2 Enterprise Server Edition machine are forced to run within an agent. Some advantages of running within an agent are that local clients can be monitored and that they can use SYSPLEX support.

DB2_COPY_NAME

- Operating system: Windows
- Default: The name of the default copy of DB2 installed on your machine. Values: the name of a copy of DB2 installed on your machine. The name can be up to 128 characters long.
- The **DB2_COPY_NAME** variable stores the name of the copy of DB2 currently in use. If you have multiple DB2 copies installed on your machine, you cannot use **DB2_COPY_NAME** to switch to a different copy of DB2, you must run the command `INSTALLPATH\bin\db2envar.bat` to change the copy currently in use.

DB2DBMSADDR

- Operating system: Linux on x86 and Linux on zSeries (31-bit)
- Default: NULL, Values: virtual addresses in the range 0x09000000 to 0xB0000000 in increments of 0x10000
- The **DB2DBMSADDR** registry variable specifies the default database shared memory address in hexadecimal format.

Note: An incorrect address can cause severe issues with the DB2 system, ranging from an inability to start a DB2 instance, to an ability to connect to the database. An incorrect address is one that collides with an area in memory that is already in use, or is predestined to be used for something else. To address this problem, reset the **DB2DBMSADDR** registry variable to NULL by using the following command:

```
db2set DB2DBMSADDR=
```

This variable can be used to fine tune the address space layout of DB2 processes. This variable changes the location of the instance shared memory from its current location at virtual address 0x10000000 to the new value.

DB2_DIAGPATH

- Operating system: All
- Default: The default value is the instance db2dump directory on UNIX and Linux operating systems, and the instance DB2 directory on Windows operating systems.
- This parameter applies to ODBC and DB2 CLI applications only. This parameter allows you to specify the fully qualified path for DB2 diagnostic information. This directory could possibly contain dump files, trap files, an error log, a notification file, and an alert log file, depending on your platform. Setting this environment variable has the same effect for ODBC and CLI applications in the scope of that environment as setting the DB2

database manager configuration parameter **diagpath**, and as setting the CLI/ODBC configuration keyword **DiagPath**.

DB2DOMAINLIST

- Operating system: All
- Default: NULL, Values: A list of Windows domain names separated by commas (",")
- This variable defines one or more Windows domains. The list, which is maintained on the server, defines the domains that the requesting user ID is authenticated against. Only users belonging to these domains have their connection or attachment requests accepted.

This variable is effective only when CLIENT authentication is set in the database manager configuration. It is needed if a single sign-on from a Windows desktop is required in a Windows domain environment.

DB2 servers versions 7.1 or later support **DB2DOMAINLIST**, but only in a pure Windows domain environment. Starting with Version 8 FixPak 15 and Version 9.1 Fix Pack 3, **DB2DOMAINLIST** is supported if either the client or the server is running in a Windows environment.

DB2ENVLIST

- Operating system: UNIX
- Default: NULL
- This variable lists specific variable names for either stored procedures or user-defined functions. By default, the db2start command filters out all user environment variables except those prefixed with **DB2** or **db2**. If specific environment variables must be passed to either stored procedures or user-defined functions, you can list the variable names in the **DB2ENVLIST** environment variable. Separate each variable name by one or more spaces.

DB2INSTANCE

- Operating system: All
- Default: **DB2INSTDEF** on Windows 32-bit operating systems.
- This environment variable specifies the instance that is active by default. On UNIX, users must specify a value for **DB2INSTANCE**.

DB2INSTPROF

- Operating system: Windows
- Default: Documents and Settings\All Users\Application Data\IBM\DB2*Copy Name* (Windows XP, Windows 2003), ProgramData\IBM\DB2*Copy Name* (Windows Vista)
- This environment variable specifies the location of the instance directory on Windows operating systems. Beginning with version 9.5, the instance directory (and other user data files) cannot be under the sqllib directory.

DB2LDAPSecurityConfig

- Operating system: All
- Default: NULL, Values: valid name and path to the IBM LDAP security plug-in configuration file
- This variable is used to specify the location of the IBM LDAP security plug-in configuration file. If the variable is not set, the IBM LDAP security plug-in configuration file is named **IBMLDAPSecurity.ini** and is in one of the following locations:
 - On Linux and UNIX operating systems: *INSTHOME*/sqllib/cfg/

- On Windows operating systems: %DB2PATH%\cfg\

On Windows operating systems, this variable should be set in the global system environment to ensure it is picked up by the DB2 service.

DB2LIBPATH

- Operating system: UNIX
- Default: NULL
- DB2 constructs its own shared library path. If you want to add a PATH into the engine's library path (for example, on AIX, a user-defined function requires a specific entry in **LIBPATH**), you must set **DB2LIBPATH**. The actual value of **DB2LIBPATH** is appended to the end of the DB2 constructed shared library path.

DB2LOGINRESTRICTIONS

- Operating system: AIX
- Default: LOCAL, Values: LOCAL, REMOTE, SU, NONE
- This registry variable allows you to use an AIX operating system API called `loginrestrictions()`. This API determines whether a user is allowed to access the system. By calling this API, DB2 database security is able to enforce the login restrictions that are specified by the operating system. There are different values that can be submitted to this API when using this registry variable. The values are:
 - REMOTE
DB2 only enforces login restrictions to verify that the account can be used for remote logins through the *rlogind* or *telnetd* programs.
 - SU
DB2 Version 9.1 only enforces su restrictions to verify that the su command is permitted, and that the current process has a group ID that can invoke the su command to switch to the account.
 - NONE
DB2 does not enforce any login restrictions.
 - LOCAL (or the variable is not set)
DB2 only enforces login restrictions to verify that local logins are permitted for this account. This is the normal behavior when logging in.

No matter which one of these options you set, user accounts or IDs that have the specified privileges are able to use DB2 successfully both locally on the server and from remote clients. For a description of the `loginrestrictions()` API, please refer to AIX documentation.

DB2NODE

- Operating system: All
- Default: NULL, Values: 1 to 999
- Used to specify the target logical node of a database partition server that you want to attach to or connect to. If this variable is not set, the target logical node defaults to the logical node which is defined with port 0 on the machine. In a partitioned database environment, the connection settings could have an impact on acquiring trusted connections. For example, if the **DB2NODE** variable is set to a node such that the establishment of a connection on that node requires going through an intermediate node (a hop node), it is the IP address of that intermediate node and the communication protocol used to communicate between the hop node and the connection node that are considered when evaluating

this connection in order to determine whether or not it can be marked as a trusted connection. In other words, it is not the original node from which the connection was initiated that is considered. Rather, it is the hop node that is considered.

DB2OPTIONS

- Operating system: All
- Default: NULL
- Used to set the command line processor options.

DB2_PARALLEL_IO

- Operating system: All
- Default: NULL, Values: *TablespaceID*:*[n]*,... – a comma-separated list of defined table spaces (identified by their numeric table space ID). If the prefetch size of a table space is AUTOMATIC, you can indicate to the DB2 database manager the number of disks per container for that table space by specifying the table space ID, followed by a colon, followed by the number of disks per container, *n*. If *n* is not specified, the default is 6.

You can replace *TablespaceID* with an asterisk (*) to specify all table spaces. For example, if **DB2_PARALLEL_IO** =*, all table spaces use six as the number of disks per container. If you specify both an asterisk (*) and a table space ID, the table space ID setting takes precedence. For example, if **DB2_PARALLEL_IO** =*,1:3, all table spaces use six as the number of disks per container, except for table space 1, which uses three.

- This registry variable is used to change the way DB2 calculates the I/O parallelism of a table space. When I/O parallelism is enabled (either implicitly, by the use of multiple containers, or explicitly, by setting **DB2_PARALLEL_IO**), it is achieved by issuing the correct number of prefetch requests. Each prefetch request is a request for an extent of pages. For example, a table space has two containers and the prefetch size is four times the extent size. If the registry variable is set, a prefetch request for this table space will be broken into four requests (one extent per request) with a possibility of four prefetchers servicing the requests in parallel.

You might want to set the registry variable if the individual containers in the table space are striped across multiple physical disks or if the container in a table space is created on a single RAID device that is composed of more than one physical disk.

If this registry variable is not set, the degree of parallelism of any table space is the number of containers of the table space. For example, if **DB2_PARALLEL_IO** is set to NULL and a table space has four containers, four extent-sized prefetch requests are issued; or if a tablespace has two containers and the prefetch size is four times the extent size, the prefetch request for this table space will be broken into two requests (each request will be for two extents).

If this registry variable is set, and the prefetch size of the table is not AUTOMATIC, the degree of parallelism of the table space is the prefetch size divided by the extent size. For example, if **DB2_PARALLEL_IO** is set for a table space that has a prefetch size of 160 and an extent size of 32 pages, five extent-sized prefetch requests are issued.

If this registry variable is set, and the prefetch size of the table space is AUTOMATIC, DB2 automatically calculates the prefetch size of a table space using the following equation:

Prefetch size =
 (number of containers) * (number of disks per container)
 * extent size

The number after the colon is used by DB2 for the *number of disks per container* in the equation. If only an asterisk is used but a number is not specified, a default of 6 disks per container is used.

The following table summarizes the different options available and how parallelism is calculated for each situation:

Table 64. How Parallelism is Calculated

Prefetch size of table space	DB2_PARALLEL_IO Setting	Parallelism is equal to:
AUTOMATIC (Prefetchsize=number of containers * 1 * extent size)	Not set	Number of containers
AUTOMATIC (Prefetchsize=number of containers * 6 * extent size)	Table space ID	Number of containers * 6
AUTOMATIC (Prefetchsize=number of containers * n * extent size)	Table space ID:n	Number of containers * n
Not AUTOMATIC	Not set	Number of containers
Not AUTOMATIC	Table space ID	Prefetch size/extent size
Not AUTOMATIC	Table space ID:n	Prefetch size/extent size

For example, consider you have three table spaces and their ID's are 3,4, and 5 respectively. Their extent sizes are all 4096 bytes, and they all have two containers each. The prefetch sizes of table spaces 3 and 4 are both AUTOMATIC, and that of table space 5 is 16384 bytes. Suppose you set **DB2_PARALLEL_IO**=*:5,4:10, the table spaces will have parallelism derived as follows:

- Tablespace 3: The value of *n* (number of disks per container) is 5, extent size=4096, number of containers=2, and prefetch size is AUTOMATIC. Therefore, prefetch size is 2*5*4096 and parallelism=number of containers * *n* =2*5=10.
- Tablespace 4: Note the value of *n* (number of disks per container) is set specifically to 10 for this tablespace. Extent size=4096, number of containers=2 *n*=10, and prefetch size is AUTOMATIC. Therefore, prefetch size =2*10*4096, and parallelism=number of containers*n=2*10=20.
- Tablespace 5: The value of *n* is still 5, but it has no influence, as prefetch size is not AUTOMATIC. Extent size=4096 number of containers=2 and prefetch size=16384. Therefore, parallelism=prefetch size/extent size=16384/4096=4.

Disk contention might result using this variable in some scenarios. For example, if a table space has two containers and each of the two containers have each a single disk dedicated to it, setting the registry variable might result in contention on those disks because the two prefetchers will be accessing each of the two disks at once. However, if each of the two containers was striped across multiple disks, setting the registry variable would potentially allow access to four different disks at once.

To activate changes to this registry variable, issue a db2stop command and then enter a db2start command.

DB2PATH

- Operating system: Windows
- Default: Varies by operating system
- This environment variable is used to specify the directory where the product is installed on Windows 32-bit operating systems.

DB2PROCESSORS

- Operating system: Windows
- Default: NULL, Values: 0– $n-1$ (where n = the number of processors)
- This variable sets the process affinity mask for a particular db2syscs process. In environments running multiple logical nodes, this variable is used to associate a logical node to a processor or set of processors. When specified, DB2 issues the SetProcessAffinityMask() api. If unspecified, the db2syscs process is associated with all processors on the server.

DB2RCMD_LEGACY_MODE

- Operating system: Windows,
- Default: NULL, Values: YES, ON, TRUE, or 1, or NO, OFF, FALSE, or 0
- This variable allows users to enable or disable the DB2 Remote Command Service's enhanced security. To run the DB2 Remote Command Service in a secure manner, set **DB2RCMD_LEGACY_MODE** to NO, OFF, FALSE, 0, or NULL. To run in legacy mode (without enhanced security), set **DB2RCMD_LEGACY_MODE** to YES, ON, TRUE, or 1. The secure mode is only available if your domain controller is running Windows 2000 or later.

Note: If **DB2RCMD_LEGACY_MODE** is set to YES, ON, TRUE, or 1, all requests sent to the DB2 Remote Command Service are processed under the context of the requestor. To facilitate this, you must allow either or both the machine and service logon account to impersonate the client by enabling the machine and service logon accounts at the domain controller.

Note: If **DB2RCMD_LEGACY_MODE** is set to NO, OFF, FALSE, or 0, you must have SYSADM authority in order to have the DB2 Remote Command Service execute commands on your behalf.

DB2SYSTEM

- Operating system: Windows and UNIX
- Default: NULL
- Specifies the name that is used by your users and database administrators to identify the DB2 server system. If possible, this name should be unique within your network.

This name is displayed in the system level of the Control Center's object tree to aid administrators in the identification of server systems that can be administered from the Control Center.

When using the Search the Network function of the Configuration Assistant, DB2 discovery returns this name and it is displayed at the system level in the resulting object tree. This name aids users in identifying the system that contains the database they wish to access. A value for **DB2SYSTEM** is set at installation time as follows:

- On Windows the setup program sets it equal to the computer name specified for the Windows system.
- On UNIX systems, it is set equal to the UNIX system's TCP/IP hostname.

DB2_UPDDBCFG_SINGLE_DBPARTITION

- Operating system: All
- Default: Not set, Values: 0/FALSE/NO, 1/TRUE/YES
- When set to 1, TRUE, or, YES, this registry variable allows you to specify that any updates and resets to your database affect only a specific partition. If the variable is not set, updates and requests follow the version 9.5 behavior.
- Beginning with version 9.5, updates or changes to a database configuration act across all database partitions, when you do not specify a partition clause. **DB2_UPDDBCFG_SINGLE_DBPARTITION** enables you to revert to the behavior of previous versions of DB2, in which updates to a database configuration apply only to the local database partition or the database partition that is set by the **DB2NODE** registry variable. This allows for backward compatibility support for any existing command scripts or applications that require this behavior.

Note: This variable does not apply to update or reset requests made by calling ADMIN_CMD routines.

DB2_USE_PAGE_CONTAINER_TAG

- Operating system: All
- Default: NULL, Values: ON, NULL
- By default, DB2 stores a container tag in the first extent of each DMS container, whether it is a file or a device. The container tag is the metadata for the container. Before DB2 Version 8.1, the container tag was stored in a single page, and it thus required less space in the container. To continue to store the container tag in a single page, set **DB2_USE_PAGE_CONTAINER_TAG** to ON.

However, if you set this registry variable to ON when you use RAID devices for containers, I/O performance might degrade. Because for RAID devices you create table spaces with an extent size equal to or a multiple of the RAID stripe size, setting the **DB2_USE_PAGE_CONTAINER_TAG** to ON causes the extents not to line up with the RAID stripes. As a result, an I/O request might need to access more physical disks than would be optimal. Users are strongly advised against enabling this registry variable unless you have very tight space constraints, or you require behavior consistent with pre-Version 8 databases..

To activate changes to this registry variable, issue a db2stop command and then enter a db2start command.

DB2_WORKLOAD

- Operating system: All
- Default: Not set, Values: 1C, CM, SAP, TPM, WC
- Each value for **DB2_WORKLOAD** represents a specific grouping of several registry variables with predefined settings.
- These are the valid values:

- 1C** Use this setting when you want to configure a set of registry variables in your database for 1C applications.
- CM** Use this setting when you want to configure a set of registry variables in your database for IBM Content Manager. This value is available in DB2 Version 9.5 Fix Pack 3 and later.
- SAP** Use this setting when you want to configure a set of registry variables in your database for the SAP environment.

When you have set **DB2_WORKLOAD=SAP**, the user table space **SYSTOOLSPACE** and the user temporary table space **SYSTOOLSTMPSPACE** are not automatically created. These table spaces are used for tables created automatically by the following wizards, utilities, or functions:

- Automatic maintenance
- Design Advisor
- Control Center database information panel
- **SYSINSTALLOBJECTS** stored procedure, if the table space input parameter is not specified
- **GET_DBSIZE_INFO** stored procedure

Without the **SYSTOOLSPACE** and **SYSTOOLSTMPSPACE** table spaces, you cannot use these wizards, utilities, or functions.

To be able to use these wizards, utilities, or functions, do either of the following:

- Manually create the **SYSTOOLSPACE** table space to hold the objects that the tools need (in a partitioned database environment, create this table space on the catalog partition).

For example:

```
CREATE REGULAR TABLESPACE SYSTOOLSPACE
  IN IBMCATGROUP
  MANAGED BY SYSTEM
  USING ('SYSTOOLSPACE')
```

- Specifying a valid table space, call the **SYSINSTALLOBJECTS** stored procedure to create the objects for the tools, and specify the identifier for the particular tool. **SYSINSTALLOBJECTS** will create a table space for you. If you do not want to use **SYSTOOLSPACE** for the objects, specify a different user-defined table space.

After completing at least one of these choices, create the **SYSTOOLSTMPSPACE** temporary table space (also on the catalog partition, if you're working in a partitioned database environment). For example:

```
CREATE USER TEMPORARY TABLESPACE SYSTOOLSTMPSPACE
  IN IBMCATGROUP
  MANAGED BY SYSTEM
  USING ('SYSTOOLSTMPSPACE')
```

Once the table space **SYSTOOLSPACE** and the temporary table space **SYSTOOLSTMPSPACE** are created, you can use the wizards, utilities, or functions mentioned earlier.

- TPM** Use this setting when you want to configure a set of registry variables in your database for the IBM Tivoli Provisioning Manager.

- WC Use this setting when you want to configure a set of registry variables in your database for Websphere Commerce. This value is available in DB2 Version 9.5 Fix Pack 4 and later.

Communications variables

DB2CHECKCLIENTINTERVAL

- Operating system: All, server only
- Default=50, Values: A numeric value that is greater than or equal to zero.
- This variable specifies the frequency of TCP/IP client connection verifications. It permits early detection of client termination, instead of waiting until after the completion of the query. If this variable is set to 0, no verification is performed.

Lower values cause more frequent checks. As a guide, for low frequency, use 100; for medium frequency, use 50; for high frequency use 10. The value is measured in an internal DB2 metric. The values represent a linear scale, that is, increasing the value from 50 to 100 doubles the interval. Checking more frequently for client status while executing a database request lengthens the time taken to complete queries. If the DB2 workload is heavy (that is, it involves many internal requests), setting **DB2CHECKCLIENTINTERVAL** to a low value has a greater impact on performance than in a situation where the workload is light. Since DB2[®] Universal Database[™], Version 8.1.4, the default value for **DB2CHECKCLIENTINTERVAL** has been 50. Prior to version 8.1.4, the default value was 0.

DB2COMM

- Operating system: All, server only
- Default=NULL, Values: NPIPE, TCPIP, SSL
- This variable specifies the communication managers that are started when the database manager is started. If this variable is not set, no DB2 communications managers are started at the server.

DB2FCMCOMM

- Operating system: All supported DB2 Enterprise Server Edition platforms
- Default=TCPIP4, Values: TCPIP4 or TCPIP6
- This variable specifies how the host names in the db2nodes.cfg file are resolved. All host names are resolved as IPv4 or IPv6. If an IP address instead of a host name is specified in db2nodes.cfg, the form of the IP determines if IPv4 or IPv6 is used. If **DB2FCMCOMM** is not set, its default setting of IPv4 means that only IPv4 hosts can be started.

Note: If the IP format resolved from the hostname specified in db2nodes.cfg, or the IP format directly specified in db2nodes.cfg does not match the setting of **DB2FCMCOMM**, db2start will fail.

DB2_FORCE-NLS_CACHE

- Operating system: AIX, HP_UX, Solaris
- Default=FALSE, Values: TRUE or FALSE
- This variable is used to eliminate the chance of lock contention in multi-threaded applications. When this registry variable is TRUE, the code page and territory code information is saved the first time a thread accesses it. From that point, the cached information is used for any other

thread that requests this information. This eliminates lock contention and results in a performance benefit in certain situations. This setting should not be used if the application changes locale settings between connections. It is probably not needed in such a situation because multi-threaded applications typically do not change their locale settings because it is not *thread safe* to do so.

DB2RSHCMD

- Operating system: UNIX
- Default=rsh (remsh on HP-UX), Values are a full path name to rsh, remsh, or ssh
- By default, DB2 database system uses rsh as the communication protocol when starting remote database partitions and with the db2_all script to run utilities and commands on all database partitions. For example, setting this registry variable to the full path name for ssh causes DB2 database products to use ssh as the communication protocol for the requested running of the utilities and commands. It may also be set to the full path name of a script that invokes the remote command program with appropriate default parameters. This variable is only required for partitioned databases, or for single-partition environments where the db2start command is run from a different server than where the DB2 product was installed. The instance owner must be able to use the specified remote shell program to log in from each DB2 database node to each other DB2 database node, without being prompted for any additional verification or authentication (that is, passwords or password phrases).

For detailed instructions on setting the DB2RSHCMD registry variable to use a ssh shell with DB2, see the white paper “Configure DB2 Universal Database for UNIX to use OpenSSH.”

DB2RSHTIMEOUT

- Operating system: UNIX
- Default=30 seconds, Values: 1 - 120
- This variable is only applicable if **DB2RSHCMD** is set to a non-null value. This registry variable is used to control the timeout period that the DB2 database system will wait for any remote command. After this timeout period, if no response is received, the assumption is made that the remote database partition is not reachable and the operation has failed.

Note: The time value given is not the time required to run the remote command, it is the time needed to authenticate the request.

DB2SORCVBUF

- Operating system: All
- Default=65 536
- Specifies the value of TCP/IP receive buffers.

DB2SOSNDBUF

- Operating system: All
- Default=65 536
- Specifies the value of TCP/IP send buffers.

DB2TCP_CLIENT_CONTIMEOUT

- Operating system: All, client only

- Default=0 (no timeout), Values: 0 - 32 767 seconds
- The **DB2TCP_CLIENT_CONTIMEOUT** registry variable specifies the number of seconds a client waits for the completion on a TCP/IP connect operation. If a connection is not established in the seconds specified, then the DB2 database manager returns the error -30081 `selectForConnectTimeout`.

There is no timeout if the registry variable is not set or is set to 0.

Note: Operating systems also have a connection timeout value that may take effect prior to the timeout you set using **DB2TCP_CLIENT_CONTIMEOUT**. For example, AIX has a default `tcp_keepinit=150` (in half seconds) that would terminate the connection after 75 seconds.

DB2TCP_CLIENT_RCVTIMEOUT

- Operating system: All, client only
- Default=0 (no timeout), Values: 0 - 32 767 seconds
- The **DB2TCP_CLIENT_RCVTIMEOUT** registry variable specifies the number of seconds a client waits for data on a TCP/IP receive operation. If data from the server is not received in the seconds specified, then the DB2 database manager returns the error -30081 `selectForRecvTimeout`.

There is no timeout if the registry variable is not set or is set to 0.

Note: The value of the **DB2TCP_CLIENT_RCVTIMEOUT** can be overridden by the CLI, using the `db2cli.ini` keyword `ReceiveTimeout` or the connection attribute `SQL_ATTR_RECEIVE_TIMEOUT`.

DB2TCPCONNMGRS

- Operating system: All
- Default=1 on serial machines; square root of the number of processors rounded up to a maximum of sixteen connection managers on symmetric multiprocessor machines. Values: 1 to 16
- The default number of connection managers is created if the registry variable is not set. If the registry variable is set, the value assigned here overrides the default value. The number of TCP/IP connection managers specified up to a maximum of 16 is created. If less than 1 is specified then **DB2TCPCONNMGRS** is set to a value of 1 and a warning is logged that the value is out of range. If greater than 16 is specified then **DB2TCPCONNMGRS** is set to a value of 16 and a warning is logged that the value is out of range. Values between 1 and 16 are used as given. When there is greater than one connection manager created, connection throughput should improve when multiple client connections are received simultaneously. There may be additional TCP/IP connection manager processes (on UNIX) or threads (on Windows operating systems) if the user is running on a SMP machine, or has modified the **DB2TCPCONNMGRS** registry variable. Additional processes or threads require additional storage.

Note: Having the number of connection managers set to 1 causes a drop in performance on remote connections in systems with a lot of users, frequent connects and disconnects, or both.

Command-line variables

DB2BQTIME

- Operating system: All
- Default=1 second, Minimum value: 1 second
- This variable specifies the amount of time the command-line processor front end sleeps before it checks whether the back-end process is active and establishes a connection to it.

DB2BQTRY

- Operating system: All
- Default=60 retries, Minimum value: 0 retries
- This variable specifies the number of times the command-line processor front-end process tries to determine whether the back-end process is already active. It works in conjunction with **DB2BQTIME**.

DB2_CLPPROMPT

- Operating system: All
- Default=None (if it is not defined, “db2 => ” will be used as the default CLP interactive prompt), Values: Any text string of length less than 100 that contains zero or more of the following tokens %i, %d, %ia, %da, or %n. Users need not set this variable unless they explicitly wish to change the default CLP interactive prompt (db2 =>).
- This registry variable allows a user to define the prompt to be used in the Command Line Processor (CLP) interactive mode. The variable can be set to any text string of length less than 100 characters containing zero or more of the optional tokens %i, %d, %ia, %da, or %n. When running in CLP interactive mode, the prompt to be used is constructed by taking the text-string specified in the **DB2_CLPPROMPT** registry variable and replacing all occurrences of the tokens %i, %d, %ia, %da, or %n by the local alias of the current attached instance, the local alias of the current database connection, the authorization ID of the current attached instance, the authorization ID of the current database connection, and newline (that is, a carriage-return) respectively.

Note:

1. If the **DB2_CLPPROMPT** registry variable is changed within CLP interactive mode, the new value for **DB2_CLPPROMPT** will not take effect until the CLP interactive mode has been closed and reopened.
2. If no instance attachment exists, %ia is replaced by the empty string and %i is replaced by the value of the **DB2INSTANCE** registry variable. On Windows platforms only, if the **DB2INSTANCE** variable is not set, %i is replaced by the value of the **DB2INSTDEF** registry variable. If neither of these variables are set, %i is replaced by the empty string.
3. If no database connection exists, %da is replaced by the empty string and %d is replaced by the value of the **DB2DBDFT** registry variable. If the **DB2DBDFT** variable is not set, %d is replaced by the empty string.
4. The interactive input prompt will always present the values for the authorization IDs, database names, and instance names in upper case.

DB2IQTIME

- Operating system: All
- Default=5 seconds, Minimum value: 1 second
- This variable specifies the amount of time the command line processor back end process waits on the input queue for the front end process to pass commands.

DB2RQTIME

- Operating system: All
- Default=5 seconds, Minimum value: 1 second
- This variable specifies the amount of time the command line processor back end process waits for a request from the front end process.

Partitioned database environment variables

DB2CHGPWD_EEE

- Operating system: DB2 ESE on AIX, Linux, and Windows
- Default=NULL, Values: YES or NO
- This variable specifies whether you allow other users to change passwords on AIX or Windows ESE systems. You must ensure that the passwords for all database partitions or nodes are maintained centrally using either a Windows domain controller on Windows, or LDAP on AIX. If not maintained centrally, passwords may not be consistent across all database partitions or nodes. This could result in a password being changed only at the database partition to which the user connects to make the change.

DB2_FCM_SETTINGS

- Operating system: Linux
- Default=NULL, Values: FCM_MAXIMIZE_SET_SIZE:[YES|TRUE|NO|FALSE]. The default value for FCM_MAXIMIZE_SET_SIZE is NO.
- Starting from Version 9.5 fix pack 4, you can set the **DB2_FCM_SETTINGS** registry variable with the FCM_MAXIMIZE_SET_SIZE token to preallocate a default 2 GB space for the fast communication manager (FCM) buffer. The token needs have a value of either YES or TRUE to enable this feature.

DB2_NUM_FAILOVER_NODES

- Operating system: All
- Default=2, Values: 0 to the required number of database partitions
- On each machine where a given database exists, set **DB2_NUM_FAILOVER_NODES** to specify the total number of database partitions in that database.

In a DB2 database high availability solution, if a database server fails, the database partitions on the failed machine can be restarted on another machine. The fast communication manager (FCM) uses **DB2_NUM_FAILOVER_NODES** to calculate how much memory to reserve on each machine to facilitate this failover.

For example, consider the following configuration:

- Machine A has two database partitions: 1 and 2.
- Machine B has two database partitions: 3 and 4.
- **DB2_NUM_FAILOVER_NODES** is set to 4 on both A and B.

At DB2START, FCM will reserve enough memory on both A and B to manage up to four database partitions so that if one machine fails, the two database partitions on the failed machine can be restarted on the other machine. If machine A fails, database partitions 1 and 2 can be restarted on machine B. If machine B fails, database partitions 3 and 4 can be restarted on machine A.

DB2_PARTITIONEDLOAD_DEFAULT

- Operating system: All supported ESE platforms
- Default=YES, Values: YES or NO
- The **DB2_PARTITIONEDLOAD_DEFAULT** registry variable lets users change the default behavior of the load utility in an ESE environment when no ESE-specific load options are specified. The default value is YES, which specifies that in an ESE environment if you do not specify ESE-specific load options, loading is attempted on all database partitions on which the target table is defined. When the value is NO, loading is attempted only on the database partition to which the load utility is currently connected.

Note: This variable is deprecated and may be removed in a later release. The LOAD command has various options that can be used to achieve the same behavior. You can achieve the same results as the NO setting for this variable by specifying the following with the LOAD command: PARTITIONED DB CONFIG MODE LOAD_ONLY OUTPUT_DBPARTNUMS x, where x is the partition number of the partition into which you want to load data.

DB2PORTRANGE

- Operating system: Windows
- Values: nnnn:nnnn
- This value is set to the TCP/IP port range used by FCM so that any additional database partitions created on another machine will also have the same port range.

Query compiler variables

DB2_ANTIJOIN

- Operating system: All
- Default=NO in a ESE environment, Default=YES in a non-ESE environment, Values: YES, NO, or EXTEND
- For DB2 Enterprise Server Edition: when YES is specified, the optimizer searches for opportunities to transform "NOT EXISTS" subqueries into anti-joins which can be processed more efficiently by DB2. For non-ESE environments: when NO is specified, the optimizer limits the opportunities to transform "NOT EXISTS" subqueries into anti-joins. In both ESE and NON-ESE environments, when EXTEND is specified, the optimizer searches for opportunities to transform both "NOT IN" and "NOT EXISTS" subqueries into anti-joins.

DB2_INLIST_TO_NLJN

- Operating system: All
- Default=NO, Values: YES or NO
- In some situations, the SQL and XQuery compiler can rewrite an IN list predicate to a join. For example, the following query:

```

SELECT *
FROM EMPLOYEE
WHERE DEPTNO IN ('D11', 'D21', 'E21')

```

could be written as:

```

SELECT *
FROM EMPLOYEE, (VALUES 'D11', 'D21', 'E21') AS V(DNO)
WHERE DEPTNO = V.DNO

```

This revision might provide better performance if there is an index on DEPTNO. The list of values would be accessed first and joined to EMPLOYEE with a nested loop join using the index to apply the join predicate.

Sometimes the optimizer does not have accurate information to determine the best join method for the rewritten version of the query. This can occur if the IN list contains parameter markers or host variables which prevent the optimizer from using catalog statistics to determine the selectivity. This registry variable causes the optimizer to favor nested loop joins to join the list of values, using the table that contributes the IN list as the inner table in the join.

Note: When either or both of the DB2 query compiler variables **DB2_MINIMIZE_LISTPREFETCH** and **DB2_INLIST_TO_NLJN**, are set to YES, they remain active even if REOPT(ONCE) is specified.

DB2_LIKE_VARCHAR

- Operating system: All
- Default=Y,Y,
- Controls the use of sub-element statistics. These are statistics about the content of data in columns when the data has a structure in the form of a series of sub-fields or sub-elements delimited by blanks. Collection of sub-element statistics is optional and controlled by options in the RUNSTATS command or API.

This registry variable affects how the optimizer deals with a predicate of the form:

```
COLUMN LIKE '%xxxxxx%'
```

where the xxxxxx is any string of characters.

The syntax showing how this registry variable is used is:

```
db2set DB2_LIKE_VARCHAR=[Y|N|S|num1] [,Y|N|S|num2]
```

where

- The term preceding the comma, or the only term to the right of the predicate, means the following but only if the second term is specified as N or the column does not have positive sub-element statistics:
 - S – The optimizer estimates the length of each element in a series of elements concatenated together to form a column based on the length of the string enclosed in the % characters.
 - Y – The default. Use a default value of 1.9 for the algorithm parameter. Use a variable-length sub-element algorithm with the algorithm parameter.
 - N – Use a fixed-length sub-element algorithm.
 - num1 – Use the value of num1 as the algorithm parameter with the variable length sub-element algorithm.

- The term following the comma means the following, but only for columns that do have positive sub-element statistics:
 - N – Do not use sub-element statistics. The first term takes effect
 - Y – The default. Use a variable-length sub-element algorithm that uses sub-element statistics together with the 1.9 default value for the algorithm parameter in the case of columns with positive sub-element statistics.
 - num2 – Use a variable-length sub-element algorithm that uses sub-element statistics together with the value of num2 as the algorithm parameter in the case of columns with positive sub-element statistics.

DB2_MINIMIZE_LISTPREFETCH

- Operating system: All
- Default=NO, Values: YES or NO
- List prefetch is a special table access method that involves retrieving the qualifying RIDs from the index, sorting them by page number and then prefetching the data pages. Sometimes the optimizer does not have accurate information to determine if list prefetch is a good access method. This might occur when predicate selectivities contain parameter markers or host variables that prevent the optimizer from using catalog statistics to determine the selectivity.

This registry variable prevents the optimizer from considering list prefetch in such situations.

Note: When either or both of the DB2 query compiler variables **DB2_MINIMIZE_LISTPREFETCH** and **DB2_INLIST_TO_NLJN**, are set to YES, they remain active even if REOPT(ONCE) is specified.

DB2_NEW_CORR_SQ_FF

- Operating system: All
- Default=OFF, Values: ON or OFF
- Affects the selectivity value computed by the query optimizer for certain subquery predicates when it is set to ON. It can be used to improve the accuracy of the selectivity value of equality subquery predicates that use the MIN or MAX aggregate function in the SELECT list of the subquery. For example:

```
SELECT * FROM T WHERE
T.COL = (SELECT MIN(T.COL)
FROM T WHERE ...)
```

DB2_OPT_MAX_TEMP_SIZE

- Operating system: All
- Default=NULL, Values: amount of space in megabytes that can be used by a query in all temporary table spaces
- Limits the amount of space that queries can use in the temporary table spaces. Setting **DB2_OPT_MAX_TEMP_SIZE** can cause the optimizer to choose a plan that is more expensive than would otherwise be chosen, but which uses less space in the temporary table spaces. If you set **DB2_OPT_MAX_TEMP_SIZE**, be sure to balance your need to limit use of temporary table space against the efficiency of the plan your setting causes to be chosen.

If **DB2_WORKLOAD=SAP** is set, **DB2_OPT_MAX_TEMP_SIZE** is automatically set to 10 240 (10 GB).

If you run a query that uses temporary table space in excess of the value set for **DB2_OPT_MAX_TEMP_SIZE**, the query does not fail, but you receive a warning that its performance may be suboptimal, as not all resources may be available.

The operations considered by the optimizer that are affected by the limit set by **DB2_OPT_MAX_TEMP_SIZE** are:

- Explicit sorts for operations such as ORDER BY, DISTINCT, GROUP BY, merge scan joins, and nested loop joins.
- Explicit temporary tables
- Implicit temporary tables for hash joins and duplicate merge joins

DB2_REDUCED_OPTIMIZATION

- Operating system: All
- Default=NO, Values: NO, YES, any integer, DISABLE, NO_SORT_NLJOIN, or NO_SORT_MGJOIN
- This registry variable lets you request either reduced optimization features or rigid use of optimization features at the specified optimization level. If you reduce the number of optimization techniques used, you also reduce time and resource use during optimization.

Note: Although optimization time and resource use might be reduced, the risk of producing a less than optimal data access plan is increased. Use this registry variable only when advised by IBM or one of its partners.

- If set to NO

The optimizer does not change its optimization techniques.

- If set to YES

If the optimization level is 5 (the default) or lower, the optimizer disables some optimization techniques that might consume significant prepare time and resources but do not usually produce a better access plan.

If the optimization level is exactly 5, the optimizer scales back or disables some additional techniques, which might further reduce optimization time and resource use, but also further increase the risk of a less than optimal access plan. For optimization levels lower than 5, some of these techniques might not be in effect in any case. If they are, however, they remain in effect.

- If set to any integer

The effect is the same as YES, with the following additional behavior for dynamically prepared queries optimized at level 5. If the total number of joins in any query block exceeds the setting, then the optimizer switches to greedy join enumeration instead of disabling additional optimization techniques as described above for level 5 optimization levels. which implies that the query will be optimized at a level similar to optimization level 2.

- If set to DISABLE

The behavior of the optimizer when unconstrained by this **DB2_REDUCED_OPTIMIZATION** variable is sometimes to dynamically reduce the optimization for dynamic queries at optimization level 5. This setting disables this behavior and requires the optimizer to perform full level 5 optimization.

- If set to NO_SORT_NLJOIN

The optimizer does not generate query plans that force sorts for nested loop joins (NLJN). These types of sorts can be useful for improving performance; therefore, be careful when using the NO_SORT_NLJOIN option, as performance can be severely impacted.

- If set to NO_SORT_MGJOIN

The optimizer does not generate query plans that force sorts for merge scan joins (MSJN). These types of sorts can be useful for improving performance; therefore, be careful when using the NO_SORT_MGJOIN option, as performance can be severely impacted.

Note that the dynamic optimization reduction at optimization level 5 takes precedence over the behavior described for optimization level of exactly 5 when **DB2_REDUCED_OPTIMIZATION** is set to YES as well as the behavior described for the integer setting.

DB2_SELECTIVITY

- Operating system: All
- Default=NO, Values: YES or NO
- This registry variable controls where the SELECTIVITY clause can be used in search conditions in SQL statements.

When this registry variable is set to YES, the SELECTIVITY clause can be specified for the following predicates:

- A basic predicate in which at least one expression contains host variables
- A LIKE predicate in which the MATCH expression, predicate expression, or escape expression contains host variables

DB2_SQLROUTINE_PREPOPTS

- Operating system: All
- Default=Empty string, Values:
 - BLOCKING {UNAMBIG | ALL | NO}
 - DATETIME {DEF | USA | EUR | ISO | JIS | LOC}
 - DEGREE {1 | *degree-of-parallelism* | ANY}
 - DYNAMICRULES {BIND | INVOKEBIND | DEFINEBIND | RUN | INVOKERUN | DEFINERUN}
 - EXPLAIN {NO | YES | ALL}
 - EXPLSNAP {NO | YES | ALL}
 - FEDERATED {NO | YES}
 - INSERT {DEF | BUF}
 - ISOLATION {CS | RR | UR | RS | NC}
 - QUERYOPT *optimization-level*
 - REOPT {NONE | ONCE | ALWAYS}
 - VALIDATE {RUN | BIND}
- The **DB2_SQLROUTINE_PREPOPTS** registry variable can be used to customize the precompile and bind options for SQL and XQuery procedures. When setting this variable, separate each of the options with a space, as follows:

```
db2set DB2_SQLROUTINE_PREPOPTS="BLOCKING ALL VALIDATE RUN"
```

For a complete description of each option and its settings, see "BIND command."

If you want to achieve the same results as **DB2_SQLROUTINE_PREOPTS** for select individual procedures, but without restarting the instance, use the **SET_ROUTINE_OPTS** procedure.

Performance variables

DB2_ALLOCATION_SIZE

- Operating system: All
- Default: 128 KB, Range: 64 KB - 256 MB
- Specifies the size of memory allocations for buffer pools.

The potential advantage of setting a higher value for this registry variable is fewer allocations will be required to reach a desired amount of memory for a buffer pool.

The potential cost of setting a higher value for this registry variable is wasted memory if the buffer pool is altered by a non-multiple of the allocation size. For example, if the value for **DB2_ALLOCATION_SIZE** is 8 MB and a buffer pool is reduced by 4 MB, this 4 MB will be wasted because an entire 8 MB segment cannot be freed.

Note: **DB2_ALLOCATION_SIZE** is deprecated and may be removed in a later release.

DB2_APM_PERFORMANCE

- Operating system: All
- Default: OFF, Values: ON or OFF
- Set this variable to ON to enable performance-related changes in the access plan manager (APM) that affect the behavior of the query cache (package cache). These settings are not usually recommended for production systems. They introduce some limitations, such as the possibility of out-of-package cache errors or increased memory use, or both.

Setting **DB2_APM_PERFORMANCE** to ON also enables the NO PACKAGE LOCK mode. This mode allows the global query cache to operate without the use of package locks, which are internal system locks that protect cached package entries from being removed. The NO PACKAGE LOCK mode might result in somewhat improved performance, but certain database operations are not allowed. These prohibited operations might include: operations that invalidate packages, operations that inoperate packages, and PRECOMPILE, BIND, and REBIND.

DB2ASSUMEUPDATE

- Operating system: All
- Default: OFF, Values: ON or OFF
- When enabled, this variable allows the DB2 database system to assume that all fixed-length columns provided in an UPDATE statement are being changed. This eliminates the need for the DB2 database system to compare the existing column values to the new values to determine if the column is actually changing. Using this registry variable can cause additional logging and index maintenance when columns are provided for update (for example, in a SET clause) but are not actually being modified.

The activation of the **DB2ASSUMEUPDATE** registry variable is effective on the db2start command.

DB2_ASYNC_IO_MAXFILOP

- Operating system: All
- Default: The value of the **maxfilop** configuration parameter, Values: from the value of **maxfilop** to the value of **max_int**
- **DB2_ASYNC_IO_MAXFILOP** is deprecated and may be removed in a later release. This variable is obsolete because of the shared file handle table maintained by the threaded database manager. For more information, see "Shared file handle table" in *Data Servers, Databases, and Database Objects Guide*.

DB2_ASYNC_IO_MAXFILOP can still be set in Version 9.5, but it will have no effect. If you want to limit the number of file handles that can be open for each database, use the **maxfilop** configuration parameter.

DB2_AVOID_PREFETCH

- Operating system: All
- Default: OFF, Values: ON or OFF
- Specifies whether prefetch should be used during crash recovery. If **DB2_AVOID_PREFETCH** =ON, prefetch is not used.

DB2BPVARS

- Operating system: As specified for each parameter
- Default: Path
- Two sets of parameters are available to tune buffer pools. One set of parameters, available only on Windows, specify that buffer pools should use scatter read for specific types of containers. The other set of parameters, available on all platforms, affect prefetching behavior.

Parameters are specified in an ASCII file, one parameter on each line, in the form **parameter=value**. For example, a file named **bpvars.vars** might contain the following lines:

```
NO_NT_SCATTER = 1
NUMPREFETCHQUEUES = 2
```

Assuming that **bpvars.vars** is stored in **F:\vars**, to set these variables you execute the following command:

```
db2set DB2BPVARS=F:\vars\bpvars.vars
```

Scatter-read parameters

The scatter-read parameters are recommended for systems with a large amount of sequential prefetching against the respective type of containers and for which you have already set **DB2NTNOCACHE** to ON. These parameters, available only on Windows platforms, are **NT_SCATTER_DMSFILE**, **NT_SCATTER_DMSDEVICE**, and **NT_SCATTER_SMS**. Specify the **NO_NT_SCATTER** parameter to explicitly disallow scatter read for any container. Specific parameters are used to turn scatter read on for all containers of the indicated type. For each of these parameters, the default is zero (or OFF); and the possible values include: zero (or OFF) and 1 (or ON).

Note: You can turn on scatter read only if **DB2NTNOCACHE** is set to ON to turn Windows file caching off. If **DB2NTNOCACHE** is set to OFF or not set, a warning message is written to the administration notification log if you attempt to turn on scatter read for any container, and scatter read remains disabled.

Prefetch-adjustment parameters

The prefetch-adjustment parameters are **NUMPREFETCHQUEUES** and **PREFETCHQUEUESIZE**. These parameters are available on all platforms and can be used to improve bufferpool data prefetching. For example, consider sequential prefetching in which the desired **PREFETCHSIZE** is divided into **PREFETCHSIZE/EXTENTSIZE** prefetch requests. In this case, requests are placed on prefetch queues from which I/O servers are dispatched to perform asynchronous I/O. By default, the DB2 database manager maintains one queue of size $\max(200, 2 * \text{NUM_IOSERVERS})$ for each database partition. In some environments, performance improves with either more queues or queues of a different size, or both. The number of prefetch queues should be at most one half of the number of I/O servers. When you set these parameters, consider other parameters such as **PREFETCHSIZE**, **EXTENTSIZE**, **NUM_IOSERVERS**, and buffer pool size, as well as workload characteristics such as the number of current users.

If you think the default values are too small for your environment, first increase the values only slightly. For example, you might set **NUMPREFETCHQUEUES=4** and **PREFETCHQUEUESIZE=200**. Make changes to these parameters in a controlled manner so that you can monitor and evaluate the effects of the change.

For **NUMPREFETCHQUEUES**, the default is 1, and the range of values is 1 to **NUM_IOSERVERS**. If you set **NUMPREFETCHQUEUES** to less than 1, it is adjusted to 1. If you set it greater than **NUM_IOSERVERS**, it is adjusted to **NUM_IOSERVERS**.

For **PREFETCHQUEUESIZE**, the default value is $\max(200, 2 * \text{NUM_IOSERVERS})$. The range of values is 1 to 32767. If you set **PREFETCHQUEUESIZE** to less than 1, it is adjusted to the default. If set greater than 32767, it is adjusted to 32767.

Note: **DB2BPVARS** is deprecated and may be removed in a later release.

DB2CHKPTR

- Operating system: All
- Default: OFF, Values: ON or OFF
- Specifies whether or not pointer checking for input is required.

DB2CHKSQLDA

- Operating system: All
- Default: ON, Values: ON or OFF
- Specifies whether or not SQLDA checking for input is required.

DB2_EVALUNCOMMITTED

- Operating system: All
- Default: OFF, Values: ON, OFF
- When enabled, this variable allows, where possible, table or index access scans to defer or avoid row locking until a data record is known to satisfy predicate evaluation.

With this variable enabled, predicate evaluation may occur on uncommitted data.

DB2_EVALUNCOMMITTED is applicable only to statements using either Cursor Stability or Read Stability isolation levels. For index scans, the index must be a type-2 index.

Furthermore, deleted rows are skipped unconditionally on table scan access while deleted keys are not skipped for type-2 index scans unless the registry variable **DB2_SKIPDELETED** is also set.

The activation of the **DB2_EVALUNCOMMITTED** registry variable is effective on the `db2start` command. The decision as to whether deferred locking is applicable is made at statement compile or bind time.

DB2_EXTENDED_IO_FEATURES

- Operating system: AIX
- Default: OFF, Values: ON, OFF
- Set this variable to ON to enable features that enhance I/O performance. This enhancement includes improving the hit rate of memory caches as well as reducing the latency on high priority I/O. These features are only available on certain combinations of software and hardware configuration; setting this variable to ON for other configurations will be ignored by either the DB2 database management system or by the operating system. The minimum configuration requirements are:
 - Database version: DB2 V9.1
 - RAW device must be used for database containers (container on file systems is not supported)
 - Operating system: AIX 5.3 TL4
 - Storage subsystem: Shark DS8000™ supports all the enhanced I/O performance features. Refer to the Shark DS8000 documentation for setup and prerequisite information.

The default I/O priority settings for HIGH, MEDIUM, and LOW are 3, 8, and 12, respectively; you can use the **DB2_IO_PRIORITY_SETTING** registry variable to change these settings.

DB2_EXTENDED_OPTIMIZATION

- Operating system: All
- Default: OFF, Values: ON, OFF, or ENHANCED_MULTIPLE_DISTINCT
- This variable specifies whether or not the query optimizer uses optimization extensions to improve query performance. The ON and ENHANCED_MULTIPLE_DISTINCT values specify different optimization extensions. Use a comma-separated list when you want to use both.

The ENHANCED_MULTIPLE_DISTINCT value might improve the performance of queries where multiple distinct aggregate operations are involved in one single select operation and where the ratio of processors to the number of database partitions is low (for example, the ratio is less than or equal to 1). This setting should be used in DPF (Database Partitioning Feature) environments without symmetric multiprocessors (SMPs).

The optimization extensions might not improve query performance in all environments. Testing should be done to determine individual query performance improvements.

DB2_HASH_JOIN

- Operating system: All
- Default: YES, Values: YES or NO
- Specifies hash join as a possible join method when compiling an access plan. **DB2_HASH_JOIN** needs to be tuned to get the best performance. Hash join performance is best if you can avoid hash loops and overflow

to disk. To tune hash join performance, estimate the maximum amount of memory available for the **sheapthres** configuration parameter, and then tune the **sortheap** configuration parameter. Increase its value until you avoid as many hash loops and disk overflows as possible, but do not reach the limit specified by the **sheapthres** configuration parameter.

Note: **DB2_HASH_JOIN** is deprecated in Version 9.5 and might be removed in a future release.

DB2_IO_PRIORITY_SETTING

- Operating system: AIX
- Values: HIGH:#, MEDIUM:#, LOW:#, where # can be 1 to 15
- This variable is used in combination with the **DB2_EXTENDED_IO_FEATURES** registry variable. This registry variable provides a means to override the default HIGH, MEDIUM, and LOW I/O priority settings for the DB2 database system, which are 3, 8, and 12, respectively. This registry variable must be set prior to the start of an instance; any modification requires an instance restart. Note that setting this registry variable alone does not enable the enhanced I/O features, **DB2_EXTENDED_IO_FEATURES** must be set to enable them. All system requirements for **DB2_EXTENDED_IO_FEATURES** also apply to this registry variable.

DB2_KEEP_AS_AND_DMS_CONTAINERS_OPEN

- Operating system: All
- Default: NO, Values: YES or NO
- When you set this variable to ON, each DMS table space container has a file handle opened until the database is deactivated. Query performance might improve because the overhead to open the containers is eliminated. You should use this registry only in pure DMS environments, otherwise performance of queries against SMS table spaces might be impacted negatively.

DB2_KEEPTABLELOCK

- Operating system: All
- Default: OFF, Values: ON, TRANSACTION, OFF, CONNECTION
- When this variable is set to ON or TRANSACTION, this variable allows the DB2 database system to maintain the table lock when an Uncommitted Read or Cursor Stability isolation level is closed. The table lock that is kept is released at the end of the transaction, just as it would be released for Read Stability and Repeatable Read scans.

When this variable is set to CONNECTION, a table lock is released for an application until the application either rolls back the transaction or the connection is reset. The table lock continues to be held across commits and application requests to drop the table lock are ignored by the database. The table lock remains allocated to the application. Thus, when the application re-requests the table lock, the lock is already available.

For application workloads that can leverage this optimization, performance should improve. However, the workloads of other application executing concurrently might be impacted. Other applications might get blocked from accessing a given table resulting in poor concurrency. DB2 SQL catalog tables are not impacted by this setting. The CONNECTION setting also includes the behavior described with the ON or TRANSACTION setting.

This registry variable is checked at statement compile or bind time.

DB2_LARGE_PAGE_MEM

- Operating system: AIX, Linux, Windows Server 2003
- Default: NULL, Values: Use * to denote all applicable memory regions should use large page memory, or a comma-separated list of specific memory regions that should use large page memory. Available regions vary by operating system. On AIX, the following regions can be specified: DB, DBMS, FCM, or PRIVATE. On Linux, the following region can be specified: DB. On Windows Server 2003, the following region can be specified: DB. Huge page memory is only available on AIX.
- The **DB2_LARGE_PAGE_MEM** registry variable is used to enable large page or huge page support. Setting **DB2_LARGE_PAGE_MEM=DB** enables large page memory for the database shared memory region, and if **database_memory** is set to AUTOMATIC, disables automatic tuning of this shared memory region by STMM. On AIX, setting **DB2_LARGE_PAGE_MEM=DB:16GB** enables huge page memory for the database shared memory region.

Memory access-intensive applications that use large amounts of virtual memory might obtain performance improvements by using large or huge pages. To enable the DB2 database system to use them, you must first configure the operating system to use large or huge pages.

To enable large pages for agent private memory on 64-bit DB2 for AIX (the **DB2_LARGE_PAGE_MEM=PRIVATE** setting), you have to configure large pages on the operating system and the instance owner must possess the **CAP_BYPASS_RAC_VMM** and **CAP_PROPAGATE** capabilities.

On AIX 5L™, you can set this variable to FCM. FCM memory resides in its own memory set, so you must add the FCM keyword to the value of the **DB2_LARGE_PAGE_MEM** registry variable to enable large pages for FCM memory.

On Linux, there is an additional requirement for the availability of the **libcap.so.1** library. This library must be installed for this option to work. If this option is turned on and the library is not on the system, the DB2 database disables the large kernel pages and continues to function as it would without them.

On Linux, to verify that large kernel pages are available, issue the following command:

```
cat /proc/meminfo
```

If large kernel pages are available, the following three lines should appear (with different numbers depending on the amount of memory configured on your server):

```
HugePages_Total: 200
HugePages_Free: 200
Hugepagesize: 16384 kB
```

If you do not see these lines, or if the **HugePages_Total** is 0, you need to configure the operating system or kernel.

On Windows, the amount of large page memory that is available on the system is less than the total available memory. After the system has been running for some time, memory can become fragmented, and the amount of large page memory decreases.

DB2_LOGGER_NON_BUFFERED_IO

- Operating system: All

- Default: OFF, Values: ON, OFF
- This variable enables direct I/O on the log file system .

DB2_LOGGER_NON_BUFFERED_IO registry variable is available starting in the DB2 Version 9.5 Fix Pack 1 release.

DB2MAXFSCRSEARCH

- Operating system: All
- Default: 5, Values: -1, 1 to 33554
- Specifies the number of free space control record (FSCRs) to search when adding a record to a table. The default is to search five FSCRs. Modifying this value allows you to balance insert speed with space reuse. Use large values to optimize for space reuse. Use small values to optimize for insert speed. Setting the value to -1 forces the database manager to search all FSCRs.

DB2_MAX_INACT_STMTS

- Operating system: All
- Default: Not set, Values: up to 4 GB
- This variable overrides the default limit on the number of inactive statements kept by any one application. You can choose a different value in order to increase or reduce the amount of system monitor heap used for inactive statement information. The default limit is 250.
The system monitor heap can become exhausted if an application contains a very high number of statements in a unit of work, or if there are a large number of applications executing concurrently.

DB2_MAX_NON_TABLE_LOCKS

- Operating system: All
- Default: YES, Values: See description
- This variable defines the maximum number of NON table locks a transaction can have before it releases all of these locks. NON table locks are table locks that are kept in the hash table and transaction chain even when the transaction has finished using them. Because transactions often access the same table more than once, retaining locks and changing their state to NON can improve performance.

For best results, the recommended value for this variable is the maximum number of tables expected to be accessed by any connection. If no user-defined value is specified, the default value is as follows: If the **locklist** size is greater than or equal to

`SQLP_THRESHOLD_VAL_OF_LRG_LOCKLIST_SZ_FOR_MAX_NON_LOCKS`

(currently 8000), the default value is

`SQLP_DEFAULT_MAX_NON_TABLE_LOCKS_LARGE`

(currently 150). Otherwise, the default value is

`SQLP_DEFAULT_MAX_NON_TABLE_LOCKS_SMALL`

(currently 0).

DB2_MDC_ROLLOUT

- Operating system: All
- Default: IMMEDIATE, Values: IMMEDIATE, OFF, or DEFER

- This variable enables a performance enhancement known as “rollout” for deletions from MDC tables. Rollout is a faster way of deleting rows in an MDC table, when entire cells (intersections of dimension values) are deleted in a search DELETE statement. The benefits are reduced logging and more efficient processing.
- There are three possible outcomes of the variable setting:
 - No rollout - if OFF is specified
 - Immediate rollout - if IMMEDIATE is specified.
 - Rollout with deferred index cleanup - if DEFER is specified
- If the value is changed after startup, any new compilations of a statement will respect the new registry value setting. For statements that are in the package cache, no change in delete processing will be made until the statement is recompiled. The SET CURRENT MDC ROLLOUT MODE statement overrides the value of **DB2_MDC_ROLLOUT** at the application connection level.

DB2MEMDISCLAIM

- Operating system: ALL
- Default: YES, Values: YES or NO
- Memory used by DB2 database system processes might have some associated paging space. This paging space might remain reserved even when the associated memory has been freed. Whether or not this is so depends on the operating system’s (tunable) virtual memory management allocation policy. The **DB2MEMDISCLAIM** registry variable controls whether DB2 agents explicitly request that the operating system disassociate the reserved paging space from the freed memory.

A **DB2MEMDISCLAIM** setting of YES results in smaller paging space requirements, and possibly less disk activity from paging. A **DB2MEMDISCLAIM** setting of NO results in larger paging space requirements, and possibly more disk activity from paging. In some situations, such as if paging space is plentiful and real memory is so plentiful that paging never occurs, a setting of NO provides a minor performance improvement.

DB2MEMMAXFREE

- Operating system: All
- Default: NULL, Values: 0 to $2^{32}-1$ bytes
- Specifies the maximum number of bytes of unused private memory that is retained by DB2 database system processes before unused memory is returned to the operating system.

If **DB2MEMMAXFREE** is not set, DB2 database system processes retain up to 20% of unused private memory (based on the amount of private memory currently consumed), before freeing memory back to the operating system.

Note: **DB2MEMMAXFREE** is deprecated and will be removed in a future release. This variable is no longer necessary because the database manager now uses a threaded engine model. Do not set this variable. Doing so will likely hurt performance and may lead to unexpected behavior.

DB2_MEM_TUNING_RANGE

- Operating system: AIX, Windows

- Default: NULL, Values: a sequence of percentages n , m where $n=minfree$ and $m=maxfree$
- The amount of physical memory that the DB2 instance leaves free is important because this dictates how much memory other applications running on the same machine are able to use. When self tuning of database shared memory is enabled, the amount of physical memory left free by a given instance depends on the need for memory by the instance (and its active databases). When an instance is in urgent need of additional memory, it will allocate memory until the free physical memory on the system reaches the percentage specified by *minfree*. When the instance is less in need of memory, it will maintain a larger amount of free physical memory, specified as a percentage by *maxfree*. As a result, it is a requirement that the value set for *minfree* must be less than the value of *maxfree*.

If this variable is not set, the DB2 database manager will calculate values for *minfree* and *maxfree* based on the amount of memory on the server. It is recommended that this variable *not* be set, unless you are running the self-tuning memory manager (STMM), have **database_memory** set to AUTOMATIC, and are experiencing problems related to an insufficient amount of free physical memory.

DB2_MMAP_READ

- Operating system: AIX
- Default: OFF, Values: ON or OFF
- This variable is used in conjunction with **DB2_MMAP_WRITE** to allow the DB2 database system to use mmap as an alternate method of I/O. When these variables are set to ON, data that is read to and written from the DB2 buffer pools bypasses the AIX memory cache. If you have a relatively small DB2 buffer pool, and you cannot or choose not to increase the size of this buffer pool, you should take advantage of AIX memory caching by setting **DB2_MMAP_READ** and **DB2_MMAP_WRITE** to OFF.

DB2_MMAP_WRITE

- Operating system: AIX
- Default: OFF, Values: ON or OFF
- This variable is used in conjunction with **DB2_MMAP_READ** to allow the DB2 database system to use mmap as an alternate method of I/O. When these variables are set to ON, data that is read to and written from the DB2 buffer pools bypasses the AIX memory cache. If you have a relatively small DB2 buffer pool, and you cannot or choose not to increase the size of this buffer pool, you should take advantage of AIX memory caching by setting **DB2_MMAP_READ** and **DB2_MMAP_WRITE** to OFF.

DB2_NO_FORK_CHECK

- Operating system: UNIX
- Default: OFF, Values: ON or OFF
- When this variable is enabled, the DB2 runtime client minimizes checks to determine if the current process is a result of a fork call. This can improve performance of DB2 applications that do not use the fork() api.

Note: This variable is deprecated and will be removed in a future release. It is unnecessary because the current process id (pid) is cached when the process is started or newly forked.

DB2NTMEMSIZE

- Operating system: Windows
- Default: (varies by memory segment)
- Windows requires that all shared memory segments be reserved at DLL initialization time in order to guarantee matching addresses across processes. **DB2NTMEMSIZE** permits the user to override the DB2 defaults on Windows if necessary. In most situations, the default values should be sufficient. The memory segments, default sizes, and override options are:
 1. Parallel FCM Buffers: default size is 512 MB on 32-bit platforms, 4.5 GB on 64-bit platforms; override option is FCM:<number of bytes>
 2. Fenced Mode Communication: default size is 80 MB on 32-bit platforms, 512 MB on 64-bit platforms; override option is APLD:<number of bytes>

More than one segment may be overridden by separating the override options with a semicolon (;). For example, on a 32-bit version of DB2, to limit the FCM buffers to 1 GB, and the fenced stored procedures limit to 256 MB, use:

```
db2set DB2NTMEMSIZE=FCM:1073741824;APLD:268435456
```

DB2NTNOCACHE

- Operating system: Windows
- Default: OFF, Values: ON or OFF
- The **DB2NTNOCACHE** registry variable specifies whether the DB2 database system opens database files with a NOCACHE option. If **DB2NTNOCACHE=ON**, file system caching is eliminated. If **DB2NTNOCACHE=OFF**, the operating system caches DB2 files. This applies to all data except for files that contain long fields or LOBs. Eliminating system caching allows more memory to be available to the database so that the buffer pool or sort heap can be increased.

In Windows, files are cached when they are opened, which is the default behavior. One MB is reserved from a system pool for every 1 GB in the file. Use this registry variable to override the undocumented 192 MB limit for the cache. When the cache limit is reached, an out-of-resource error is given.

Note: **DB2NTNOCACHE** has been deprecated since Version 8.2 and will be removed in a future release. You can achieve the same benefit for table space containers by using the CREATE TABLESPACE and ALTER TABLESPACE SQL statements.

DB2NTPRICLASS

- Operating system: Windows
- Default: NULL, Values: R, H, (any other value)
- Sets the priority class for the DB2 instance (program DB2SYSCS.EXE). There are three priority classes:
 - NORMAL_PRIORITY_CLASS (the default priority class)
 - REALTIME_PRIORITY_CLASS (set by using R)
 - HIGH_PRIORITY_CLASS (set by using H)

This variable is used in conjunction with individual thread priorities (set using **DB2PRIORITIES**) to determine the absolute priority of DB2 threads relative to other threads in the system.

Note: **DB2NTPRICLASS** is deprecated and should only be used at the recommendation of service. Use DB2 service classes to adjust agent priority and prefetch priority. Care should be taken when using this variable. Misuse could adversely affect overall system performance.

For more information, please refer to the `SetPriorityClass()` API in the Win32 documentation.

DB2NTWORKSET

- Operating system: Windows
- Default: 1,1
- Used to modify the minimum and maximum working-set size available to the DB2 database manager. By default, when Windows is not in a paging situation, the working set of a process can grow as large as needed. However, when paging occurs, the maximum working set that a process can have is approximately 1 MB. **DB2NTWORKSET** allows you to override this default behavior.

Specify **DB2NTWORKSET** using the syntax **DB2NTWORKSET**=*min*, *max*, where *min* and *max* are expressed in megabytes.

DB2_OVERRIDE_BPF

- Operating system: All
- Default: Not set, Values: a positive numeric number of pages OR <entry>[:<entry>...] where <entry>=<buffer pool ID>,<number of pages>
- This variable specifies the size of the buffer pool, in pages, to be created at database activation, rollforward recovery, or crash recovery. It is useful when memory constraints cause failures to occur during database activation, rollforward recovery, or crash recovery. The memory constraint could arise either in the rare case of a real memory shortage or, because of the attempt by the database manager to allocate a large buffer pool, in the case where there were inaccurately configured buffer pools. For example, when even a minimal buffer pool of 16 pages is not brought up by the database manager, try specifying a smaller number of pages using this environment variable. The value given to this variable overrides the current buffer pool size.

You can also use <entry>[:<entry>...] where <entry>=<buffer pool ID>,<number of pages> to temporarily change the size of all or a subset of the buffer pools so that they can start up.

DB2_PINNED_BP

- Operating system: AIX, HP-UX, Linux
- Default: NO, Values: YES or NO
- Setting this variable to YES causes DB2 to request that the operating system pins DB2's Database Shared Memory. When configuring DB2 to pin Database Shared Memory, care should be taken to ensure that the system is not overcommitted, as the operating system will have reduced flexibility in managing memory.

On Linux, in addition to modifying this registry variable, the library, `libcap.so.1` is also required.

Setting this variable to YES means that self tuning for database shared memory (activated by setting the **database_memory** configuration parameter to AUTOMATIC) cannot be enabled.

For AIX operating systems, to use database memory pinning with medium page size support (which is a default behavior), ensure that the instance owner has the CAP_BYPASS_RAC_VMM and CAP_PROPAGATE capabilities by logging on as root authority and issuing the following command:

```
chuser capabilities=CAP_BYPASS_RAC_VMM,CAP_PROPAGATE <instance_owner_user_id>
```

For HP-UX in a 64-bit environment, in addition to modifying this registry variable, the DB2 instance group must be given the MLOCK privilege. To do this, a user with root access rights performs the following actions:

1. Adds the DB2 instance group to the /etc/privgroup file. For example, if the DB2 instance group belongs to db2iadm1 group then the following line must be added to the /etc/privgroup file:

```
db2iadm1 MLOCK
```

2. Issues the following command:

```
setprivgrp -f /etc/privgroup
```

DB2PRIORITIES

- Operating system: All
- Values setting is platform dependent
- Controls the priorities of DB2 processes and threads.

Note: **DB2PRIORITIES** is deprecated and should only be used at the recommendation of service. Use DB2 service classes to adjust agent priority and prefetch priority.

DB2_RCT_FEATURES

- Operating system: All
- Default: NULL. Values: GROUPUPDATE=[ON|OFF]. The default value for GROUPUPDATE is OFF.
- This variable allows for optimized and reduced update processing for a searched UPDATE statement which targets multiple rows in an range clustered table when only equal predicates on the leading and subset of key sequence columns are specified. Logging is also reduced due to a single log record for all rows updated on a page, instead of a log record for each row updated.

Usage :

```
db2set DB2_RCT_FEATURES=GROUPUPDATE=ON
```

Note: This registry variable is available with Version 9.5 fix pack 3. It will not be possible to reverse migrate to older fix packs once this registry variable is enabled. In addition, group update will not be performed when the target table of the UPDATE statement uses DATA CAPTURE CHANGES, triggers, variable length columns, has any secondary indexes, or the update modifies any primary or foreign key columns.

DB2_RESOURCE_POLICY

- Operating system: AIX 5 or higher, all Linux except zSeries (32-bit), Windows Server 2003 or higher
- Default: Not set, Values: valid path to configuration file

- Defines a resource policy which can be used to limit what operating system resources are used by the DB2 database or it contains rules for assigning specific operating system resources to specific DB2 database objects. For example, on AIX, Linux, or Windows operating systems, this registry variable can be used to limit the set of processors that the DB2 database system uses. The extent of resource control varies depending on the operating system.

On AIX NUMA and Linux NUMA enabled machines, a policy can be defined which specifies what resource sets the DB2 database system uses. When resource set binding is used, each individual DB2 process is bound to a particular resource set. This can be beneficial in some performance tuning scenarios.

You can set the registry variable to indicate the path to a configuration file which defines a policy for binding DB2 processes to operating system resources. The resource policy allows you to specify a set of operating system resources to restrict the DB2 database system. Each DB2 process is bound to a single resource of the set. Resource assignment occurs in a circular round robin fashion.

Sample configuration files:

Example 1: Bind all DB2 processes to either CPU 1 or 3.

```
<RESOURCE_POLICY>
  <GLOBAL_RESOURCE_POLICY>
    <METHOD>CPU</METHOD>
    <RESOURCE_BINDING>
      <RESOURCE>1</RESOURCE>
    </RESOURCE_BINDING>
    <RESOURCE_BINDING>
      <RESOURCE>3</RESOURCE>
    </RESOURCE_BINDING>
  </GLOBAL_RESOURCE_POLICY>
</RESOURCE_POLICY>
```

Example 2: (AIX only) Bind DB2 processes to one of the following resource sets: sys/node.03.00000, sys/node.03.00001, sys/node.03.00002, sys/node.03.00003

```
<RESOURCE_POLICY>
  <GLOBAL_RESOURCE_POLICY>
    <METHOD>RSET</METHOD>
    <RESOURCE_BINDING>
      <RESOURCE>sys/node.03.00000</RESOURCE>
    </RESOURCE_BINDING>
    <RESOURCE_BINDING>
      <RESOURCE>sys/node.03.00001</RESOURCE>
    </RESOURCE_BINDING>
    <RESOURCE_BINDING>
      <RESOURCE>sys/node.03.00002</RESOURCE>
    </RESOURCE_BINDING>
    <RESOURCE_BINDING>
      <RESOURCE>sys/node.03.00003</RESOURCE>
    </RESOURCE_BINDING>
  </GLOBAL_RESOURCE_POLICY>
</RESOURCE_POLICY>
```

Note: For AIX only, use of the RSET method requires CAP_NUMA_ATTACH capability.

Example 3: (Linux only) Bind all memory from bufferpool IDs 2 and 3 which are associated with the SAMPLE database to NUMA node 3. Also

use 80 percent of the total database memory for the binding to NUMA node 3 and leave 20 percent to be striped across all nodes for non-bufferpool specific memory.

```
<RESOURCE_POLICY>
  <DATABASE_RESOURCE_POLICY>
    <DBNAME>sample</DBNAME>
    <METHOD>NODEMASK</METHOD>
    <RESOURCE_BINDING>
      <RESOURCE>3</RESOURCE>
      <DBMEM_PERCENTAGE>80</DBMEM_PERCENTAGE>
    <BUFFERPOOL_BINDING>
      <BUFFERPOOL_ID>2</BUFFERPOOL_ID>
      <BUFFERPOOL_ID>3</BUFFERPOOL_ID>
    </BUFFERPOOL_BINDING>
  </RESOURCE_BINDING>
</DATABASE_RESOURCE_POLICY>
</RESOURCE_POLICY>
```

Example 4: (For Linux and Windows only) Define two distinct processor sets specified by CPU masks 0x0F and 0xF0. Bind DB2 processes and bufferpool ID 2 to processor set 0x0F and DB2 processes and bufferpool ID 3 to processor set 0xF0. For each processor set, use 50 percent of the total database memory for the binding.

This resource policy is useful when a mapping between processors and NUMA nodes is desired. An example of such a scenario is a system with 8 processors and 2 NUMA nodes where processors 0 to 3 belong to NUMA node 0 and processors 4 to 7 belong to NUMA node 1. This resource policy allows for processor binding while implicitly maintaining memory locality (ie. a hybrid of CPU method and NODEMASK method).

```
<RESOURCE_POLICY>
  <DATABASE_RESOURCE_POLICY>
    <DBNAME>sample</DBNAME>
    <METHOD>CPUMASK</METHOD>
    <RESOURCE_BINDING>
      <RESOURCE>0x0F</RESOURCE>
      <DBMEM_PERCENTAGE>50</DBMEM_PERCENTAGE>
    <BUFFERPOOL_BINDING>
      <BUFFERPOOL_ID>2</BUFFERPOOL_ID>
    </BUFFERPOOL_BINDING>
  </RESOURCE_BINDING>
  <RESOURCE>0xF0</RESOURCE>
  <DBMEM_PERCENTAGE>50</DBMEM_PERCENTAGE>
  <BUFFERPOOL_BINDING>
    <BUFFERPOOL_ID>3</BUFFERPOOL_ID>
  </BUFFERPOOL_BINDING>
</RESOURCE_BINDING>
</DATABASE_RESOURCE_POLICY>
</RESOURCE_POLICY>
```

Note: Use of the RSET method requires CAP_NUMA_ATTACH capability and is not supported on Linux.

The configuration file specified by the **DB2_RESOURCE_POLICY** registry variable accepts a **SCHEDULING_POLICY** element. You can use the **SCHEDULING_POLICY** element on some platforms to select

- The operating system scheduling policy used by the DB2 server
You can set an operating system scheduling policy for DB2 on AIX, and for DB2 on Windows using the **DB2NTPRICLASS** registry variable.
- The operating system priorities used by individual DB2 server agents

Alternatively, you can use the registry variables **DB2PRIORITIES** and **DB2NTPRICLASS** to control the operating system scheduling policy and set DB2 agent priorities. However, the specification of a **SCHEDULING_POLICY** element in the resource policy configuration file provides a single place to specify both the scheduling policy and the associated agent priorities.

Example 1: Selection of the AIX **SCHED_FIFO2** scheduling policy with a priority boost for the db2 log writer and reader processes.

```
<RESOURCE_POLICY>
<SCHEDULING_POLICY>
<POLICY_TYPE>SCHED_FIFO2</POLICY_TYPE>
<PRIORITY_VALUE>60</PRIORITY_VALUE>

<EDU_PRIORITY>
<EDU_NAME>db2loggr</EDU_NAME>
<PRIORITY_VALUE>56</PRIORITY_VALUE>
</EDU_PRIORITY>

<EDU_PRIORITY>
<EDU_NAME>db2logwr</EDU_NAME>
<PRIORITY_VALUE>56</PRIORITY_VALUE>
</EDU_PRIORITY>
</SCHEDULING_POLICY>
</RESOURCE_POLICY>
```

Example 2: Replacement for **DB2NTPRICLASS=H** on Windows.

```
<RESOURCE_POLICY>
<SCHEDULING_POLICY>
<POLICY_TYPE>HIGH_PRIORITY_CLASS</POLICY_TYPE>
</SCHEDULING_POLICY>
</RESOURCE_POLICY>
```

DB2_SET_MAX_CONTAINER_SIZE

- Operating system: All
- Default: Not set, Values: -1, any positive integer greater than 65536 bytes
- This registry variable allows you to limit the size of individual containers for automatic storage table spaces with the AutoResize feature enabled.

Note: Although you can specify **DB2_SET_MAX_CONTAINER_SIZE** in bytes, kilobytes, or megabytes, db2set indicates its value in bytes.

- If the value is set to -1, there will be no limit to the size of a container.

DB2_SKIPDELETED

- Operating system: All
- Default: OFF, Values: ON or OFF
- When enabled, this variable allows statements using either Cursor Stability or Read Stability isolation levels to unconditionally skip deleted keys during index access and deleted rows during table access. With **DB2_EVALUNCOMMITTED** enabled, deleted rows are automatically skipped, but uncommitted pseudo-deleted keys in type-2 indexes are not skipped unless **DB2_SKIPDELETED** is also enabled.

This registry variable does not impact the behavior of cursors on the DB2 catalog tables.

This registry variable is activated with the db2start command.

DB2_SKIPINSERTED

- Operating system: All

- Default: OFF, Values: ON or OFF
- When the **DB2_SKIPINSERTED** registry variable is enabled, it allows statements using either Cursor Stability or Read Stability isolation levels to skip uncommitted inserted rows as if they had not been inserted. This registry variable does not impact the behavior of cursors on the DB2 catalog tables. This registry variable is activated at database startup, while the decision to skip uncommitted inserted rows is made at statement compile or bind time.

Note: Skip inserted behavior is not compatible with tables that have pending rollout cleanup. As a result, scanners might wait for locks on a RID only to discover that the RID is part of a rolled out block.

DB2_SMS_TRUNC_TMPTABLE_THRESH

- Operating system: All
- Default: 0, Values: -1, 0-*n*, where *n* = the number of extents per temporary table in the SMS table space container that are to be maintained
- This variable specifies a minimum file size threshold at which the file representing a temporary table is maintained in SMS table spaces. By default, this variable is set to 0, which means no special threshold handling is done. Instead, once a temporary table is no longer needed, that file is truncated to 0 extent.

When the value of this variable is greater than 0, a larger file is maintained. This reduces some of the system overhead involved in dropping and recreating the file each time a temporary table is used.

If this variable is set to -1, the file is not truncated and the file is allowed to grow indefinitely, restricted only by system resources.

DB2_SORT_AFTER_TQ

- Operating system: All
- Default: NO, Values: YES or NO
- Specifies how the optimizer works with directed table queues in a partitioned database environment when the receiving end requires the data to be sorted and the number of receiving nodes is equal to the number of sending nodes.

When **DB2_SORT_AFTER_TQ=NO**, the optimizer tends to sort at the sending end and merge the rows at the receiving end.

When **DB2_SORT_AFTER_TQ=YES**, the optimizer tends to transmit the rows unsorted, not merge at the receiving end, and sort the rows at the receiving end after receiving all the rows.

DB2_SELUDI_COMM_BUFFER

- Operating system: All
- Default: OFF, Values: ON or OFF
- This variable is used during the processing of blocking cursors over SELECT from UPDATE, INSERT, or DELETE (UDI) queries. When enabled, this registry variable prevents the result of a query from being stored in a temporary table. Instead, during the OPEN processing of a blocking cursor for a SELECT from UDI query, the DB2 database system attempts to buffer the entire result of the query directly into the communications buffer memory area.

Note: If the communications buffer space is not large enough to hold the entire result of query, an SQLCODE -906 error is issued, and the transaction is rolled back. See the **aslheapsz** and **rqrioblk** database manager configuration parameters for information on adjusting the size of the communication buffer memory area for local and remote applications respectively.

This registry variable is not supported in partitioned database environments or when intrapartition parallelism is enabled.

DB2_TRUSTED_BINDIN

- Operating system: All
- Default: OFF, Values: OFF, ON, or CHECK
- When **DB2_TRUSTED_BINDIN** is enabled, it speeds up the execution of query statements containing host variables within an embedded unfenced stored procedure.

When this variable is enabled, there is no conversion from the external SQLDA format to an internal DB2 format during the binding of SQL and XQuery statements contained within an embedded unfenced stored procedure. This will speed up the processing of the embedded SQL and XQuery statements.

The following data types are not supported in embedded unfenced stored procedures when this variable is enabled:

- SQL_TYP_DATE
- SQL_TYP_TIME
- SQL_TYP_STAMP
- SQL_TYP_CGSTR
- SQL_TYP_BLOB
- SQL_TYP_CLOB
- SQL_TYP_DBCLOB
- SQL_TYP_CSTR
- SQL_TYP_LSTR
- SQL_TYP_BLOB_LOCATOR
- SQL_TYP_CLOB_LOCATOR
- SQL_TYP_DCLOB_LOCATOR
- SQL_TYP_BLOB_FILE
- SQL_TYP_CLOB_FILE
- SQL_TYP_DCLOB_FILE
- SQL_TYP_BLOB_FILE_OBSOLETE
- SQL_TYP_CLOB_FILE_OBSOLETE
- SQL_TYP_DCLOB_FILE_OBSOLETE

If these data types are encountered, an SQLCODE -804, SQLSTATE 07002 error is returned.

Note: The data type and length of the input host variable must match the internal data type and length of the corresponding element exactly. For host variables, this requirement will always be met. However, for parameter markers, care must be taken to ensure that matching data types are used. The CHECK option can be used to ensure that the data types and lengths match for all input host variables, but this option negates most of the performance improvements.

Note: `DB2_TRUSTED_BINDIN` is deprecated and will be removed in a later release.

DB2_USE_ALTERNATE_PAGE_CLEANING

- Operating system: All
- Default: Not set, Values: ON or OFF
- This variable specifies whether a DB2 database uses the alternate method of page cleaning algorithms or the default method of page cleaning. When this variable is set to ON, the DB2 system writes changed pages to disk, keeping ahead of `LSN_GAP` and proactively finding victims. Doing this allows the page cleaners to better utilize available disk I/O bandwidth. When this variable is set to ON, the `chngpgs_thresh` database configuration parameter is no longer relevant because it does not control page cleaner activity.

DB2_USE_IOCP

- Operating system: AIX 5.3 TL9
- Default: OFF Values: OFF or ON
- Configure IOCP before enabling this registry variable.
- This variable enables the use of AIX I/O completion ports (IOCP) when submitting and collecting asynchronous I/O (AIO) requests. This feature is used to enhance performance in a non-uniform memory access (NUMA) environment by avoiding remote memory access. This performance variable is available with Version 9.5 fix pack 3.

Miscellaneous variables

DB2ADMINSERVER

- Operating system: Windows and UNIX
- Default: NULL
- Specifies the DB2 Administration Server.

DB2_ATS_ENABLE

- Operating system: All
- Default: NULL, Values: YES/TRUE/ON/1 or NO/FALSE/OFF/0
- This variable controls whether the administrative task scheduler is running. The administrative task scheduler is disabled by default. When the scheduler is disabled, you can use the built-in procedures and views to define and modify tasks but the scheduler will not execute the tasks.

DB2AUTH

- Operating system: All
- Values: `TRUSTEDCLIENT_SRVRENC`, `TRUSTEDCLIENT_DATAENC`, `DISABLE_CHGPASS`
- This variable allows you to tune the behavior of user authentication.
 - `TRUSTEDCLIENT_SRVRENC`: This value forces untrusted clients to use `SERVER_ENCRYPT`.
 - `TRUSTEDCLIENT_DATAENC`: This value forces untrusted clients to use `DATA_ENCRYPT`.
 - `DISABLE_CHGPASS`: This value disables the ability to change the password from the client.
 - `OSAUTHDB`: Starting in DB2 Version 9.5 Fix Pack 4, this value instructs DB2 to use the authentication and group setting for a user

on the AIX operating system. By setting **DB2AUTH** to the **OSAUTHDB** value, you can configure DB2 to authenticate users and acquire their groups through the operating system. If the AIX operating system settings for a user's authentication and group information are configured to use LDAP, the AIX operating system will, in turn, perform authentication through an LDAP server. The LDAP server can be any one of the following:

- IBM Tivoli Directory Server (ITDS)
- Microsoft Active Directory (MSAD)
- Sun One Directory Server

DB2CLIINIPATH

- Operating system: All
- Default: NULL
- Used to override the default path of the DB2 CLI/ODBC configuration file (db2cli.ini) and specify a different location on the client. The value specified must be a valid path on the client system.

DB2_COMMIT_ON_EXIT

- Operating system: UNIX
- Default: OFF, Values: OFF/NO/0 or ON/YES/1
- On UNIX operating systems, prior to DB2 UDB Version 8, DB2 committed any remaining in-flight transactions on successful application exit. In DB2 UDB Version 8, the behavior was changed so that in-flight transactions were rolled back on exit. This registry variable allows users with embedded SQL applications which depend on the earlier behavior to continue to enable it in version 9. This registry variable does not affect JDBC, CLI, and ODBC applications.

Note that this registry variable is deprecated, and the commit-on-exit behavior will no longer be supported in future release. Users should determine whether any of their applications developed prior to version 9 continue to depend on this functionality, and add the appropriate explicit COMMIT or ROLLBACK statements to the application as required. If the registry variable is turned on, care should be taken not to implement new applications which fail to explicitly COMMIT before exit.

Most users should leave this registry variable at the default setting.

DB2_CREATE_DB_ON_PATH

- Operating system: Windows
- Default: NULL, Values: YES or NO
- Set this registry variable to YES to enable support for the use of a path (as well as a drive) as a database path. The setting of **DB2_CREATE_DB_ON_PATH** is checked when a database is created, when the database manager configuration parameter **dftdbpath** is set, and when a database is restored. The fully qualified database path can be up to 215 characters in length.

If **DB2_CREATE_DB_ON_PATH** is not set (or is set to NO) and you specify a path for the database path when creating or restoring a database, error SQL1052N is returned.

If **DB2_CREATE_DB_ON_PATH** is not set (or is set to NO) and you update the **dftdbpath** database manager configuration parameter, error SQL5136N is returned.

CAUTION:

If path support is used to create new databases, applications written prior to Version 9.1 using the `db2DbDirGetNextEntry()` API or an older version of it, might not work correctly. Refer to http://www.ibm.com/software/data/db2/support/db2_9/ for details on various scenarios and the proper course of action.

DB2DEFPREP

- Operating system: All
- Default: NO, Values: ALL, YES, or NO
- Simulates the runtime behavior of the DEFERRED_PREPARE precompile option for applications that were precompiled before this option was available. For example, if a DB2 v2.1.1 or earlier application were run in a DB2 v2.1.2 or later environment, **DB2DEFPREP** could be used to indicate the desired “deferred prepare” behavior.

Note: **DB2DEFPREP** is deprecated and will be removed in a future release. This variable is only needed by users using old versions of DB2 where the DEFERRED_PREPARE precompile option is not available.

DB2_DISABLE_FLUSH_LOG

- Operating system: All
- Default: OFF, Values: ON or OFF
- Specifies whether to disable closing the active log file when the online backup is completed.

When an online backup completes, the last active log file is truncated, closed, and made available to be archived. This ensures that your online backup has a complete set of archived logs available for recovery. You might want to disable closing the last active log file if you are concerned that you are wasting portions of the Log Sequence Number (LSN) space. Each time an active log file is truncated, the LSN is incremented by an amount proportional to the space truncated. If you perform a large number of online backups each day, you might disable closing the last active log file.

You might also want to disable closing the last active log file if you find you are receiving log full messages a short time after the completion of the online backup. When a log file is truncated, the reserved active log space is incremented by the amount proportional to the size of the truncated log. The active log space is freed once the truncated log file is reclaimed. The reclamation occurs a short time after the log file becomes inactive. During the short interval between these two events, you may receive log full messages.

During any backup which includes logs, this registry variable is ignored, since the active log file must be truncated and closed in order for the backup to include the logs.

DB2CONNECT_DISCONNECT_ON_INTERRUPT

- Operating system: All
- Default: NO, Values: YES/TRUE/1 or NO/FALSE/0
- When set to YES (TRUE or 1), this variable specifies that the connection to a Version 8 (or higher) DB2 Universal Database z/OS server should be broken immediately when an interrupt occurs. You can use this variable in the following configurations:

- If you are running a DB2 client with a Version 8 (or higher) DB2 UDB z/OS server, set **DB2CONNECT_DISCONNECT_ON_INTERRUPT** to YES on the client.
- If you are running a DB2 client through a DB2 Connect gateway to a Version 8 (or higher) DB2 UDB z/OS server, set **DB2CONNECT_DISCONNECT_ON_INTERRUPT** to YES on the gateway.

DB2_DISPATCHER_PEEKTIMEOUT

- Operating system: All
- Default: 1, Values: 0 to 32767 seconds; 0 denotes that timeout is immediate
- **DB2_DISPATCHER_PEEKTIMEOUT** allows you to adjust the time, in seconds, that a dispatcher waits for a client's connection request before handing the client off to an agent. In most cases, you should not need to adjust this registry variable. This registry variable only affects instances that have DB2 Connect connection concentrator enabled.

This registry variable and the **DB2_SERVER_CONTIMEOUT** registry variable both configure the handling of a new client during connect time. If there are many slow clients connecting to an instance, the dispatcher may be held up for up to 1 second to timeout each client, causing the dispatcher to become a bottle neck, if many clients are connecting simultaneously. If an instance with multiple active databases is experiencing very slow connection times, **DB2_DISPATCHER_PEEKTIMEOUT** may be lowered to 0. Lowering **DB2_DISPATCHER_PEEKTIMEOUT** causes the dispatcher to only look into the client's connect request if it is already there; the dispatcher will not wait for the connect request to arrive. If an invalid value is set, the default value is used. This registry variable is not dynamic.

DB2_DJ_INI

- Operating system: All
- Default:
 - UNIX: *db2_instance_directory*/cfg/db2dj.ini
 - Windows: *db2_install_directory*\cfg\db2dj.ini
- Specifies the absolute path name of the federation configuration file, for example: `db2set DB2_DJ_INI=$HOME/sql1lib/cfg/my_db2dj.ini` This file contains the settings for data source environment variables. These environment variables are used by the Informix® wrapper and by the wrappers provided by WebSphere® Federation Server.

Here is a sample federation configuration file:

```
INFORMIXDIR=/informix/client_sdk
INFORMIXSERVER=inf93
ORACLE_HOME=/usr/oracle9i
SYBASE=/sybase/V12
SYBASE_OCS=OCS-12_5
```

The following restrictions apply to the db2dj.ini file:

- Entries must follow the format *evname=value* where *evname* is the name of the environment variable and *value* is its value.
- The environment variable name has a maximum length of 255 bytes.
- The environment variable value has a maximum length of 765 bytes.

This variable is ignored unless the database manager parameter **FEDERATED** is set to YES.

DB2DMNBCKCTLR

- Operating system: Windows
- Default: NULL, Values: ? or a domain name
- If you know the name of the domain for which the DB2 server is the backup domain controller, set **DB2DMNBCKCTLR=DOMAIN_NAME**. The DOMAIN_NAME must be in upper case. To have DB2 determine the domain for which the local machine is a backup domain controller, set **DB2DMNBCKCTLR=?**. If the **DB2DMNBCKCTLR** profile variable is not set or is set to blank, DB2 performs authentication at the primary domain controller.

Note: DB2 does not use an existing backup domain controller by default because a backup domain controller can get out of synchronization with the primary domain controller, causing a security exposure. Getting out of synchronization can occur when the primary domain controller's security database is updated but the changes are not propagated to a backup domain controller. This could occur if there are network latencies or if the computer browser service is not operational.

Note: **DB2DMNBCKCTLR** is deprecated and will be removed in a later release. This variable is no longer necessary because there are no more backup domain controllers in the Active Directory.

DB2_DOCHOST

- Operating system: All
- Default: Not set (but DB2 will still try to access the Information Center from the IBM Web site at publib.boulder.ibm.com/infocenter/db2luw/v9r5), Values: <http://hostname> where *hostname* is a valid host name or IP address
- Specifies the host name on which the *DB2 Information Center* is installed. This variable can be automatically set during the installation of the *DB2 Information Center* if the automatic configuration option is selected in the DB2 Setup wizard.

DB2_DOCPORT

- Operating system: All
- Default: NULL, Values: any valid port number
- Specifies the port number through which the DB2 help system serves the DB2 documentation. This variable can be automatically set during the installation of the *DB2 Information Center* if the automatic configuration option is selected in the DB2 Setup wizard.

DB2_ENABLE_AUTOCONFIG_DEFAULT

- Operating system: All
- Default: NULL, Values: YES or NO
- This variable controls whether the Configuration Advisor is run automatically at database creation. If **DB2_ENABLE_AUTOCONFIG_DEFAULT** is not set (null), the effect is the same as if the variable was set to YES and the Configuration Advisor is run at database creation. You do not need to restart the instance after you set this variable. If you execute the AUTOCONFIGURE command or run CREATE DB AUTOCONFIGURE, these commands override the setting of **DB2_ENABLE_AUTOCONFIG_DEFAULT**.

DB2_ENABLE_LDAP

- Operating system: All
- Default: NO, Values: YES or NO
- Specifies whether or not the Lightweight Directory Access Protocol (LDAP) is used. LDAP is an access method to directory services.

DB2_EVMON_EVENT_LIST_SIZE

- Operating system: All
- Default: 0 (no limit), Values: A value specified in KB/Kb/kb, MB/Mb/mb, or GB/Gb/gb; While there is no fixed upper limit for this variable, it is limited by the amount of available memory from the monitor heap.
- This registry variable specifies the maximum number of bytes that can be queued up waiting to be written to a particular event monitor. Once this limit is reached, agents attempting to send event monitor records will wait until the queue size drops below this threshold.

Note: If activity records cannot be allocated from the monitor heap, they will be dropped. To prevent this from happening, set the **mon_heap_sz** configuration parameter to AUTOMATIC. If you have **mon_heap_sz** set to a specific value, ensure that **DB2_EVMON_EVENT_LIST_SIZE** is set to a smaller value. These actions, however, cannot guarantee that activity records will not be dropped, as the monitor heap is also used for tracking other monitor elements.

DB2_EVMON_STMT_FILTER

- Operating system: All
- Default: Not set; Values:
 - ALL: Indicates that the output for all statement event monitors is to be filtered. This option is exclusive.
 - 'nameA nameB nameC': Where each name in the string represents the name of an event monitor for which records are to be filtered. If more than one name is supplied, each name must be separated by a single blank. All input names will be made uppercase by DB2. The maximum number of event monitors you can specify is 32. Each monitor name can be up to a maximum of 18 characters long.
 - 'nameA:op1,op2 nameB:op1,op2 nameC:op1': Where each name in the string represents the name of an event monitor for which records are to be filtered. Each option (*op1*, *op2*, etc.) represents an integer value mapping to a specific SQL operation. Specifying integer values allows users to determine which rules to apply to which event monitor.
- **DB2_EVMON_STMT_FILTER** can be used to reduce the number of records written by a statement event monitor. When set, this registry variable causes only the records for the following SQL operations to be written to the specified event monitor:

Table 65. List of integer values that can be specified in **DB2_EVMON_STMT_FILTER** to represent specific SQL operations

SQL operation	Integer value mapping
SELECT	15
EXECUTE	2
EXECUTE_IMMEDIATE	3
CLOSE	6

Table 65. List of integer values that can be specified in **DB2_EVMON_STMT_FILTER** to represent specific SQL operations (continued)

SQL operation	Integer value mapping
STATIC COMMIT	8
STATIC ROLLBACK	9
CALL	12
PRE_EXEC	17

All other operations will not appear in the output of the statement event monitor. To customize the set of operations for which records are written to the event monitor, use integer values.

Example 1:

```
db2set DB2_EVMON_STMT_FILTER= 'mon1 monitor3'
```

In this example, mon1 and monitor3 event monitors will receive a record for a restricted list of application requests. For example, if an application being monitored by the mon1 statement event monitor prepares a dynamic SQL statement, opens a cursor based on that statement, fetches 10,000 rows from that cursor, and then issues a cursor close request, only a record for a close request will appear in the mon1 event monitor output.

Example 2:

```
db2set DB2_EVMON_STMT_FILTER='evmon1:3,8 evmon2:9,15'
```

In this example, evmon1 and evmon2 will receive a record for a restricted list of application requests. For example, if an application being monitored by the evmon1 statement event monitor issues a create statement, only the execute immediate and static commit operations will appear in the evmon1 event monitor output. If an application being monitored by the evmon2 statement event monitor performs SQL involving both a select and a static rollback only these two operations will appear in the evmon2 event monitor output.

Note: Refer to the sqlmon.h header file for definitions of database system monitor constants.

DB2_EXTSECURITY

- Operating system: Windows
- Default: ON, Values: ON or OFF
- Prevents unauthorized access to DB2 by locking DB2 system files. To avoid potential problems, this registry variable should not be turned off.

DB2_FALLBACK

- Operating system: Windows
- Default: OFF, Values: ON or OFF
- This variable allows you to force all database connections off during the fallback processing. It is used in conjunction with the failover support in the Windows environment with Microsoft Cluster Server (MSCS). If **DB2_FALLBACK** is not set or is set to OFF, and a database connection exists during the fall back, the DB2 resource cannot be brought offline. This will mean the fallback processing will fail.

DB2_FMP_COMM_HEAPSZ

- Operating system: Windows, UNIX
- Default: 20 MB, or enough space to run 10 fenced routines (whichever is larger). On AIX, the default is 256 MB
- This variable specifies, in 4 KB pages, the size of the pool used for fenced routine invocations, such as stored procedure or user-defined function calls. The space used by each fenced routine is twice the value of the **aslheapsz** configuration parameter.

If you are running a large number of fenced routines on your system, you may need to increase the value of this variable. If you are running a very small number of fenced routines, you can reduce it.

Setting this value to 0 means that no set is created, and as a result no fenced routines can be invoked. It also means that the health monitor and the automatic database maintenance functionality (such as automatic backups, statistics collection, and REORG) will be disabled since this functionality relies on the fenced routine infrastructure.

DB2_GRP_LOOKUP

- Operating system: Windows
- Default: NULL, Values: LOCAL, DOMAIN, TOKEN, TOKENLOCAL, TOKENDOMAIN
- This variable specifies which Windows security mechanism is used to enumerate the groups to which a user belongs.

DB2_HADR_BUF_SIZE

- Operating system: All
- Default: 2*logbufsz
- This variable specifies the standby log receiving buffer size in unit of log pages. If not set, DB2 will use two times the primary **logbufsz** configuration parameter value for the standby receiving buffer size. This variable should be set in the standby instance. It is ignored by the primary database.

If HADR synchronization mode (the **hadr_syncmode** database configuration parameter) is set to ASYNC, during peer state, a slow standby may cause the send operation on the primary to stall and therefore block transaction processing on the primary. A larger than default log-receiving buffer can be configured on a standby database to allow it to hold more unprocessed log data. This may allow for brief periods where the primary generates log data faster than the standby can consume it, without blocking transaction processing at the primary.

Note: A larger log receiving buffer size can help absorb peak transaction loads on the primary database, but the buffer will still fill up if the average replay rate on the standby database is slower than the log rate on the primary database.

DB2_HADR_NO_IP_CHECK

- Operating system: All
- Default: OFF, Values: ON | OFF
- Specifies whether to bypass IP check for HADR connections
- This variable is primarily used in Network Address Translation (NAT) environments to bypass IP cross check for HADR connections. Use of this variable is not recommended in other environments because it weakens the sanity check of the HADR configuration. By default, configuration consistency for the local and remote host parameters is

verified when an HADR connection is established. Hostnames are mapped to IP addresses for the cross check. Two checks are performed:

- **HADR_LOCAL_HOST** parameter on primary = **HADR_REMOTE_HOST** parameter on standby
- **HADR_REMOTE_HOST** parameter on primary = **HADR_LOCAL_HOST** parameter on standby

The connection will be closed if the check fails.

When this parameter is turned on, no IP check occurs.

DB2_HADR_PEER_WAIT_LIMIT

- Operating system: All
- Default: 0 (meaning no limit) Values: 0 to max unsigned 32 bit integer, inclusive
- Starting DB2 Version 9.5 Fix Pack 1, when registry variable **DB2_HADR_PEER_WAIT_LIMIT** is set, the HADR primary database will break out of peer state if logging on the primary database has been blocked for the specified number of seconds because of log replication to the standby. When this limit is reached, the primary database will break the connection to the standby database. If the peer window is disabled, the primary database will enter disconnected state and logging resumes. If the peer window is enabled, the primary database will enter disconnected peer state, in which logging continues to be blocked. The primary database leaves disconnected peer state upon re-connection or peer window expiration. Logging resumes once the primary leaves disconnected peer state. This parameter has no effect on a standby database. It is recommended that the same value be used on both primary and standby databases though. Invalid values (not a number or negative numbers) will be interpreted as "0", meaning no limit. This parameter is static. Database instance needs to be restarted to make this parameter effective.

DB2_HADR_SORCVBUF

- Operating system: All
- Default: Operating system TCP socket receive buffer size, Values: 1024 to 4294967295
- This variable specifies the operating system (OS) TCP socket receive buffer size for the HADR connection, which allows users to customize the HADR TCP/IP behavior distinctly from other connections. Some operating systems will automatically round or silently cap the user specified value. The actual buffer size used for the HADR connection is logged in the db2diag.log file. Consult your operating system network tuning guide for the optimal setting for this parameter based on your network traffic. This variable should be used in conjunction with **DB2_HADR_SOSNDBUF**.

DB2_HADR_SOSNDBUF

- Operating system: All
- Default: Operating system TCP socket send buffer size, Values: 1024 to 4294967295
- This variable specifies the operating system (OS) TCP socket send buffer size for the HADR connection, which allows users to customize the HADR TCP/IP behavior distinctly from other connections. Some operating systems will automatically round or silently cap the user specified value. The actual buffer size used for the HADR connection is

logged in the db2diag.log file. Consult your operating system network tuning guide for the optimal setting for this parameter based on your network traffic. This variable should be used in conjunction with **DB2_HADR_SORCVBUF**.

DB2LDAP_BASEDN

- Operating system: All
- Default: NULL, Values: Any valid base distinguished domain name.
- When this is set, the LDAP objects for DB2 will be stored in the LDAP directory under
 - CN=System
 - CN=IBM
 - CN=DB2

under the base distinguished name specified.

When this is set for the Microsoft Active Directory Server, ensure that CN=DB2, CN=IBM, and CN=System are defined under this distinguished name.

DB2LDAPCACHE

- Operating system: All
- Default: YES, Values: YES or NO
- Specifies that the LDAP cache is to be enabled. This cache is used to catalog the database, node, and DCS directories on the local machine. To ensure that you have the latest entries in the cache, do the following:
 - REFRESH LDAP IMMEDIATE ALL

This command updates and removes incorrect entries from the database directory and the node directory.

DB2LDAP_CLIENT_PROVIDER

- Operating system: Windows
- Default: NULL (Microsoft, if available, is used; otherwise IBM is used.) Values: IBM or Microsoft
- When running in a Windows environment, DB2 supports using either Microsoft LDAP clients or IBM LDAP clients to access the LDAP directory. This registry variable is used to explicitly select the LDAP client to be used by DB2.

Note: To display the current value of this registry variable, use the db2set command:

```
db2set DB2LDAP_CLIENT_PROVIDER
```

DB2LDAPHOST

- Operating system: All
- Default: NULL, Values: Any valid hostname
- Specifies the hostname of the location for the LDAP directory.

DB2LDAP_KEEP_CONNECTION

- Operating system: All
- Default: YES, Values: YES or NO
- Specifies whether DB2 caches its internal LDAP connection handles. When this variable is set to NO, DB2 will not cache its LDAP connection handles to the directory server. This will likely result in a negative performance impact, but it might be desirable to set

DB2LDAP_KEEP_CONNECTION to NO if the number of simultaneously active LDAP client connections to the directory server needs to be minimized.

To maximize performance, this variable is set to YES by default.

The **DB2LDAP_KEEP_CONNECTION** registry variable is only implemented as a global level profile registry variable in LDAP, so you must set it by specifying the -gl option with the db2set command as follows:

```
db2set -gl DB2LDAP_KEEP_CONNECTION=NO
```

DB2LDAP_SEARCH_SCOPE

- Operating system: All
- Default: DOMAIN, Values: LOCAL, DOMAIN, or GLOBAL
- Specifies the search scope for information found in database partitions or domains in the Lightweight Directory Access Protocol (LDAP). LOCAL disables searching in the LDAP directory. DOMAIN only searches in LDAP for the current directory partition. GLOBAL searches in LDAP in all directory partitions until the object is found.

DB2_LOAD_COPY_NO_OVERRIDE

- Operating system: All
- Default: NONRECOVERABLE, Values: COPY YES or NONRECOVERABLE
- This variable will convert any LOAD COPY NO to either LOAD COPY YES or NONRECOVERABLE, depending on the value of the variable. This variable is applicable to HADR primary databases as well as to standard (non-HADR) databases; it is ignored on an HADR standby database. On an HADR primary database, if this variable is not set, LOAD COPY NO is converted to LOAD NONRECOVERABLE. The value of this variable either specifies a nonrecoverable load or the copy destination, using the same syntax as a COPY YES clause.

DB2LOADREC

- Operating system: All
- Default: NULL
- Used to override the location of the load copy during roll forward. If the user has changed the physical location of the load copy, **DB2LOADREC** must be set before issuing the roll forward.

DB2LOCK_TO_RB

- Operating system: All
- Default: NULL, Values: STATEMENT
- Specifies whether lock timeouts cause the entire transaction to be rolled back, or only the current statement. If **DB2LOCK_TO_RB** is set to STATEMENT, locked timeouts cause only the current statement to be rolled back. Any other setting results in transaction rollback.

DB2_MAP_XML_AS_CLOB_FOR_DLC

- Operating system: All
- Default: NO, Values: YES or NO
- The **DB2_MAP_XML_AS_CLOB_FOR_DLC** registry variable provides the ability to override the default DESCRIBE and FETCH behavior of XML values for clients (or DRDA Application Requestors) that do not support XML as a data type. The default value is NO, which specifies

that for these clients a DESCRIBE of XML values will return BLOB(2GB), and a FETCH of XML values will result in an implicit XML serialization to BLOB that includes an XML declaration indicating an encoding of UTF-8.

When the value is YES, a DESCRIBE of XML values will return CLOB(2GB), and a FETCH of XML values will result in an implicit XML serialization to CLOB that does not contain an XML declaration.

Note: **DB2_MAP_XML_AS_CLOB_FOR_DLC** is deprecated and will be removed in a future release. This variable is no longer necessary because most existing DB2 applications that access XML values do so with an XML capable client.

DB2_MAX_LOB_BLOCK_SIZE

- Operating system: All
- Default: 0 (no limit), Values: 0 to 21487483647
- Sets the maximum amount of LOB or XML data to be returned in a block. This is not a hard maximum; if this maximum is reached on the server during data retrieval, the server finishes writing out the current row before generating a reply for the command, such as FETCH, to the client.

DB2_MEMORY_PROTECT

- Operating system: AIX with storage key support
- Default: NO, Values: NO or YES
- This registry variable enables a memory protection feature that uses storage keys to prevent data corruption in the buffer pool caused by invalid memory access. Memory protection works by identifying at which times the DB2 engine threads should have access to the buffer pool memory and at which times they should not have access. When **DB2_MEMORY_PROTECT** is set to YES, any time a DB2 engine thread tries to illegally access buffer pool memory, that engine thread traps. **DB2_MEMORY_PROTECT** needs to be enabled for the trap resilience feature **DB2_THREAD_SUSPENSION** to work.

Note: You will not be able to use the memory protection if **DB2_LGPAGE_BP** is set to YES. Even if **DB2_MEMORY_PROTECT** is set to YES, DB2 will fail to protect the buffer pool memory and disable the feature.

DB2NOEXITLIST

- Operating system: All
- Default: OFF, Values: ON or OFF
- This variable indicates that DB2 should not load an exit list handler and that it should not perform a commit when the application exits, regardless of the setting of the **DB2_COMMIT_ON_EXIT** registry variable.

When **DB2NOEXITLIST** is turned off and **DB2_COMMIT_ON_EXIT** is turned on, any in-flight transactions for embedded SQL applications are automatically committed. It is recommended to explicitly add COMMIT or ROLLBACK statements when an application exits.

Applications that dynamically load and unload the DB2 library before the application terminates might crash when calling the DB2 exit handler. This crash might happen because the application attempts to

call a function that does not exist in memory. To avoid this situation, set the **DB2NOEXITLIST** registry variable.

DB2_NUM_CKPW_DAEMONS

- Operating system: UNIX
- Default: 3, Values: 1[:FORK] to 100[:FORK]
- You can use the **DB2_NUM_CKPW_DAEMONS** registry variable to start a configurable number of check password daemons. The daemons are created during db2start and handle check password requests when the default IBMOSauthserver security plugin is in use. Increasing the setting for **DB2_NUM_CKPW_DAEMONS** can decrease the time required to establish a database connection, but this is only effective in scenarios where many connections are being made simultaneously and where authentication is expensive.

DB2_NUM_CKPW_DAEMONS can be set to a value between 1 and 100. The database manager will create the number of daemons specified by **DB2_NUM_CKPW_DAEMONS**. Each daemon can handle check password requests directly.

An optional FORK parameter can be added to enable the check password daemons to explicitly spawn an external check password program (db2ckpw) to handle check password requests. This is similar to setting **DB2_NUM_CKPW_DAEMONS** to zero in previous releases. In FORK mode, each check password daemon will spawn the check password program for each request to check a password. The daemons in FORK mode are started as the instance owner.

If **DB2_NUM_CKPW_DAEMONS** is set to zero, the effective value is set to 3:FORK, where 3 check password daemons are started in FORK mode.

DB2_OPTSTATS_LOG

- Operating system: All
- Default: Not set (see details below), Values = OFF, ON {NUM | SIZE | NAME | DIR}
- **DB2_OPTSTATS_LOG** specifies the attributes of the statistics event logging files which are used to monitor and analyze statistics collection related activities. When **DB2_OPTSTATS_LOG** is not set or set to ON, statistics event logging is enabled, allowing you to monitor system performance and keep a history for better problem determination. Log records are written to the first log file until it is full. Subsequent records are written to the next available log file. If the maximum number of files is reached, the oldest log file will be overwritten with the new records. If system resource consumption is of great concern to you, disable this registry variable by setting it to OFF.

When statistics event logging is explicitly enabled (set to ON), there are a number of options you can modify:

- **NUM**: the maximum number of rotating log files. Default = 5, Values 1 - 15
- **SIZE**: the maximum size of rotating log files. (The size of each rotating file is SIZE/NUM.) Default = 100 Mb, Values 1 Mb - 4096 Mb
- **NAME**: the base name for rotating log files. Default = db2optstats.number.log, for instance db2optstats.0.log, db2optstats.1.log, etc.

- DIR: the base directory for rotating log files. Default = \$DIAGPATH/events

You can specify a value for as many of these options as you want, just ensure that ON is the first value when you want to enable statistics logging. For instance, to enable statistics logging with maximum of 6 log files, a maximum file size of 25 Mb, a base file name of mystatslog, and the directory mystats, issue the following command:

```
db2set DB2_OPTSTATS_LOG=ON,NUM=6,SIZE=25,NAME=mystatslog,DIR=mystats
```

DB2REMOTEPREG

- Operating system: Windows
- Default: NULL, Values: Any valid Windows machine name
- Specifies the remote machine name that contains the Win32 registry list of DB2 instance profiles and DB2 instances. The value for **DB2REMOTEPREG** should only be set once after DB2 is installed, and should not be modified. Use this variable with extreme caution.

DB2_RESOLVE_CALL_CONFLICT

- Operating system: AIX, HP-UX, Solaris, Linux, Windows
- Default: YES, Values: YES, NO, or ALL
- Eliminates SQLCODE SQL0746 runtime errors when CALL statements are issued in two situations:
 - In triggers, when you set **DB2_RESOLVE_CALL_CONFLICT** to YES (the default value) or ALL
 - In SQL table functions, when you set **DB2_RESOLVE_CALL_CONFLICT** to ALL

To comply with the SQL standard rules for order of execution, the database manager might not read or modify the subject table of a trigger before modifications caused by a triggering operation are completed. The database manager also might not read or modify the subject table of a table function accessed elsewhere in the invoking statement. When there are multiple table access requests for these subject tables within the same invoking statements, these conflicting requests are described as *mutating tables conflicts* because the tables are potentially changing at the same time that a trigger or an SQL table function tries to reference them.

When you set this variable to YES, the default, the database manager strictly enforces the order of execution for accessing tables in procedures called within triggers through the use of temporary tables, as needed.

When you set this variable to ALL, the database manager strictly enforces the order of execution for accessing tables in procedures called within triggers and in table functions through the use of temporary tables, as needed. The ALL setting is available starting in Fix Pack 2.

When you set this variable to NO, the database manager assumes that mutating table conflicts should not occur, but if they do, it prevents them and generates SQLCODE SQL0746 error messages.

- Before changing the value of **DB2_RESOLVE_CALL_CONFLICT**, stop the instance. After changing the value, restart the instance, then rebind any packages which cause the invocation of triggers or table functions. To rebind SQL procedures, use the following statement:

```
CALL SYSPROC.REBIND_ROUTINE_PACKAGE
('P','procedureschema.procedurename','CONSERVATIVE');
```

- **DB2_RESOLVE_CALL_CONFLICT** has performance impacts:

- On triggers, when you set the registry variable to YES or ALL. Typically, the impact is small. The performance in OLTP environments is usually not affected because in most cases only one or a small number of rows are modified by the triggering statement. When you follow the general recommendation to use system-managed space (SMS) for temporary table spaces, the performance impact from using **DB2_RESOLVE_CALL_CONFLICT** is expected to be low.
- On SQL table functions, when you set the registry variable to ALL. The amount of impact depends on the type of conflict:
 - When both a CALL statement and another statement in the same SQL table function try to access the same table, the impact is small.
 - When both a CALL statement within an SQL table function and a SELECT statement invoking the same SQL table function try to access the same table, the impact depends on the amount of data inserted into the temporary table.

DB2ROUTINE_DEBUG

- Operating system: AIX and Windows
- Default: OFF, Values: ON or OFF
- Specifies whether to enable the debug capability for Java stored procedures. If you are not debugging Java stored procedures, use the default, OFF. There is a performance impact to enable debugging.

Note: **DB2ROUTINE_DEBUG** is deprecated and will be removed in a future release. This stored procedure debugger has been replaced by the Unified Debugger.

DB2SATELLITEID

- Operating system: All
- Default: NULL, Values: a valid satellite ID declared in the Satellite Control Database
- Specifies the satellite ID that is passed to the satellite control server when a satellite synchronizes. If a value is not specified for this variable, the logon ID is used as the satellite ID.

DB2_SERVER_CONTIMEOUT

- Operating system: All
- Default: 180, Values: 0 to 32767 seconds
- This registry variable and the **DB2_DISPATCHER_PEEKTIMEOUT** registry variable both configure the handling of a new client during connect time. **DB2_SERVER_CONTIMEOUT** allows you to adjust the time, in seconds, that an agent waits for a client's connection request before terminating the connection. In most cases, you should not need to adjust this registry variable, but if DB2 clients are constantly being timed out by the server at connect time, you can set a higher value for **DB2_SERVER_CONTIMEOUT** to extend the timeout period. If an invalid value is set, the default value is used. This registry variable is not dynamic.

DB2_SERVER_ENCALG

- Operating system: All
- Default: NULL, Values: AES_CMP or AES_ONLY
- The Advanced Encryption Standard (AES) can be used to encrypt user IDs and passwords when you connect to DB2 for Linux, UNIX and

Windows Version 9.5 Fix Pack 3 and later database servers. When the database manager configuration parameter **authentication** is set to `SERVER_ENCRYPT`, the default behavior is for the database server to accept whatever encryption algorithm the client proposes. The **DB2_SERVER_ENCALG** registry variable can be used to change that behavior. It affects all connections to the DB2 instance, regardless of the application language.

If this variable is set to `AES_ONLY`, the database server will only accept connections that use AES encryption. If the client does not support AES encryption, then the connection will be rejected.

If this variable is set to `AES_CMP`, the database server will accept user IDs and passwords that are encrypted using either AES or DES, but it will negotiate for AES if the client supports AES encryption.

You do not need to set **DB2_SERVER_ENCALG** if your application is programmatically enabled to use AES encryption.

DB2SORT

- Operating system: All, server only
- Default: NULL
- This variable specifies the location of a library to be loaded at runtime by the load utility. The library contains the entry point for functions used in sorting indexing data. Use **DB2SORT** to exploit vendor-supplied sorting products for use with the load utility in generating table indexes. The path supplied must be relative to the database server.

DB2_THREAD_SUSPENSION

- Operating system: AIX with storage key support
- Default: OFF, Values: ON or OFF
- This registry variable enables or disables the DB2 thread suspension feature. It allows you to control whether a DB2 instance sustains a trap by suspending a faulty engine thread (a thread which has tried to illegally access memory protected with storage keys).

Note: **DB2_THREAD_SUSPENSION** can only be enabled if the registry variable **DB2_MEMORY_PROTECT** is set to YES.

DB2_TRUNCATE_REUSESTORAGE

- Operating system: All
- Default: NULL (not set), Values: IMPORT, import
- You can use this variable to resolve lock contention between the `IMPORT` with `REPLACE` command and the `BACKUP ... ONLINE` command. In some situations, online backup and truncate operations are unable to execute concurrently. When this occurs, you can set **DB2_TRUNCATE_REUSESTORAGE** to `IMPORT` or `import`, and physical truncation of the object, including data, indexes, long fields, large objects and block maps (for multidimensional clustering tables), is skipped and only logical truncation is performed. That is, the `IMPORT` with `REPLACE` command empties the table, causing the object's logical size to decrease, but the storage on disk remains allocated.

This registry variable is dynamic; you can set it or unset it without having to stop and start instance. You can set **DB2_TRUNCATE_REUSESTORAGE** before an online backup starts and then unset it after online backup completes. For multi-partitioned

environments, the registry variable will only be active on the nodes on which the variable is set. **DB2_TRUNCATE_REUSESTORAGE** is only effective on DMS permanent objects.

In SAP environments, when **DB2_WORKLOAD=SAP** is set, the default value of this registry variable is **IMPORT**.

DB2_USE_DB2JCCT2_JROUTINE

- Operating system: All
- Default: Not set, Values: ON/YES/1/TRUE or OFF/NO/0/FALSE
- The default driver for Java stored procedures and user-defined functions is the IBM Data Server Driver for JDBC and SQLJ. If you want to use the deprecated driver DB2 JDBC Type 2 Driver for Linux, UNIX, and Windows to serve SQL requests for Java routines, set **DB2_USE_DB2JCCT2_JROUTINE** to any of OFF, NO, 0, or FALSE.

DB2_UTIL_MSGPATH

- Operating system: All
- Default: *instanceName*/tmp directory
- The **DB2_UTIL_MSGPATH** registry variable is used in conjunction with the **SYSPROC.ADMIN_CMD** procedure, the **SYSPROC.ADMIN_REMOVE_MSGS** procedure, and the **SYSPROC.ADMIN_GET_MSGS** UDF. It applies on the instance level. **DB2_UTIL_MSGPATH** can be set to indicate a directory path on the server where the fenced user ID can read, write and delete files. This directory must be accessible from all coordinator partitions, and must have sufficient space to contain utility message files.

If this path is not set, the *instanceName*/tmp directory is used as the default (note that *instanceName*/tmp is cleaned up when DB2 is uninstalled).

If this path is changed, the files that existed in the directory pointed to by the previous setting are not automatically moved or deleted. If you want to retrieve the contents of the messages created under the old path, you must manually move these messages (which are prefixed with the utility name and suffixed with the user ID) to the new directory to which **DB2_UTIL_MSGPATH** points. The next utility message file is created, read, and cleaned up in the new location.

The files under the **DB2_UTIL_MSGPATH** directory are utility specific, not transaction dependent. They are not part of the backup image. The files under the **DB2_UTIL_MSGPATH** directory are user managed; that means a user can delete the message files using the **SYSPROC.ADMIN_REMOVE_UTILMSG** procedure. These files are not cleaned up by uninstalling DB2.

DB2_VENDOR_INI

- Operating system: AIX, HP-UX, Solaris, and Windows
- Default: NULL, Values: Any valid path and file.
- Points to a file containing all vendor-specific environment settings. The value is read when the database manager starts.

Note: **DB2_VENDOR_INI** is deprecated in Version 9.5 and might be removed in a future release. You can put the environment variable settings that it contains into the file specified by the **DB2_DJ_INI** variable instead.

DB2_XBSA_LIBRARY

- Operating system: AIX, HP-UX, Solaris, and Windows
- Default: NULL, Values: Any valid path and file.
- Points to the vendor-supplied XBSA library. On AIX, the setting must include the shared object if it is not named shr.o. HP-UX, Solaris, and Windows do not require the shared object name. For example, to use Legato's NetWorker Business Suite Module for DB2, the registry variable must be set as follows:

```
db2set DB2_XSBA_LIBRARY="/usr/lib/libxdb2.a(bsashr10.o)"
```

The XBSA interface can be invoked through the BACKUP DATABASE or the RESTORE DATABASE commands. For example:

```
db2 backup db sample use XBSA
db2 restore db sample use XBSA
```

Chapter 20. Configuration parameters

When a DB2 database instance or a database is created, a corresponding configuration file is created with default parameter values. You can modify these parameter values to improve performance and other characteristics of the instance or database.

The disk space and memory allocated by the database manager on the basis of default values of the parameters might be sufficient to meet your needs. In some situations, however, you might not be able to achieve maximum performance using these default values.

Configuration files contain parameters that define values such as the resources allocated to the DB2 database products and to individual databases, and the diagnostic level. There are two types of configuration files:

- The database manager configuration file for each DB2 instance
- The database configuration file for each individual database.

The *database manager configuration file* is created when a DB2 instance is created. The parameters it contains affect system resources at the instance level, independent of any one database that is part of that instance. Values for many of these parameters can be changed from the system default values to improve performance or increase capacity, depending on your system's configuration.

There is one database manager configuration file for each client installation as well. This file contains information about the client enabler for a specific workstation. A subset of the parameters available for a server are applicable to the client.

Database manager configuration parameters are stored in a file named `db2system`. This file is created when the instance of the database manager is created. In Linux and UNIX environments, this file can be found in the `sql1ib` subdirectory for the instance of the database manager. In Windows, the default location of this file is the instance subdirectory of the `sql1ib` directory. If the `DB2INSTPROF` variable is set, the file is in the instance subdirectory of the directory specified by the `DB2INSTPROF` variable.

In Version 9.5, the implicit default value of `DB2INSTPROF` at the global level (`db2set -g`) will be stored at the new location shown below, even if `DB2INSTPROF` is not set:

- On Vista environments: `ProgramData\IBM\DB2\<DB2COPYNAME>`
- On Windows 2003/XP environments: `Documents and Settings\All Users\Application Data\IBM\DB2\<Copy Name>`

Other profile-registry variables that specify where run-time data files should go should query the value of `DB2INSTPROF`. This includes the following variables:

- `DB2CLINIPATH`
- `DIAGPATH`
- `SPM_LOG_PATH`

Database configuration parameters are stored in a file named `SQLDBC0N` for databases created before Version 8.2; all database configuration parameters are stored in a file named `SQLDBC0NF` for databases created in Version 8.2 and later.

These files cannot be directly edited, and can only be changed or viewed via a supplied API or by a tool which calls that API.

In a partitioned database environment, this file resides on a shared file system so that all database partition servers have access to the same file. The configuration of the database manager is the same on all database partition servers.

Most of the parameters either affect the amount of system resources that will be allocated to a single instance of the database manager, or they configure the setup of the database manager and the different communications subsystems based on environmental considerations. In addition, there are other parameters that serve informative purposes only and cannot be changed. All of these parameters have global applicability independent of any single database stored under that instance of the database manager.

A *database configuration file* is created when a database is created, and resides where that database resides. There is one configuration file per database. Its parameters specify, among other things, the amount of resource to be allocated to that database. Values for many of the parameters can be changed to improve performance or increase capacity. Different changes may be required, depending on the type of activity in a specific database.

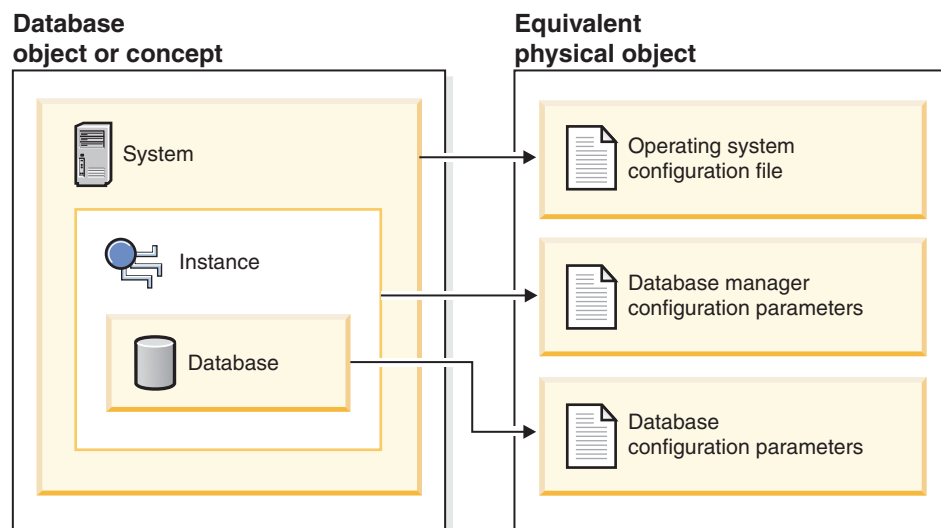


Figure 29. Relationship between database objects and configuration files

Configuring the DB2 database manager with configuration parameters

The disk space and memory allocated by the database manager on the basis of default values of the parameters might be sufficient to meet your needs. In some situations, however, you might not be able to achieve maximum performance using these default values.

Since the default values are oriented towards machines that have relatively small memory resources and are dedicated as database servers, you might need to modify these values if your environment has:

- Large databases
- Large numbers of connections
- High performance requirements for a specific application

- Unique query or transaction loads or types

Each transaction processing environment is unique in one or more aspects. These differences can have a profound impact on the performance of the database manager when using the default configuration. For this reason, you are strongly advised to tune your configuration for your environment.

A good starting point for tuning your configuration is to use the Configuration Advisor or the AUTOCONFIGURE command which will generate values for parameters based on your responses to questions about workload characteristics.

Some configuration parameters can be set to AUTOMATIC, allowing the database manager to automatically manage these parameters to reflect the current resource requirements. To turn off the AUTOMATIC setting of a configuration parameter while maintaining the current internal setting, use the MANUAL keyword with the UPDATE DATABASE CONFIGURATION command. If the database manager updates the value of these parameters, the get db/dbm cfg show detail commands will show the new value.

Parameters for an individual database are stored in a configuration file named SQLDBCONF. This file is stored along with other control files for the database in the SQLnnnnn directory, where nnnnn is a number assigned when the database was created. Each database has its own configuration file, and most of the parameters in the file specify the amount of resources allocated to that database. The file also contains descriptive information, as well as flags that indicate the status of the database.

Attention: If you edit db2system, SQLDBCON, or SQLDBCONF using a method other than those provided by the database manager, you might make the database unusable. Do not change these files using methods other than those documented and supported by the database manager.

In a partitioned database environment, a separate SQLDBCONF file exists for each database partition. The values in the SQLDBCONF file may be the same or different at each database partition, but the recommendation is that in a homogeneous environment, the configuration parameter values should be the same on all database partitions. Typically, there could be a catalog node needing different database configuration parameters setting, while the other data partitions have different values again, depending on their machine types, and other information.

Updating configuration parameters using the command line processor:

Commands to change the settings can be entered as follows:

For database manager configuration parameters:

- GET DATABASE MANAGER CONFIGURATION (or GET DBM CFG)
- UPDATE DATABASE MANAGER CONFIGURATION (or UPDATE DBM CFG)
- RESET DATABASE MANAGER CONFIGURATION (or RESET DBM CFG) to reset *all* database manager parameters to their default values
- AUTOCONFIGURE.

For database configuration parameters:

- GET DATABASE CONFIGURATION (or GET DB CFG)
- UPDATE DATABASE CONFIGURATION (or UPDATE DB CFG)

- RESET DATABASE CONFIGURATION (or RESET DB CFG) to reset *all* database parameters to their default values
- AUTOCONFIGURE.

Updating configuration parameters using application programming interfaces (APIs).

The APIs can be called from an application or a host-language program. Call the following DB2 APIs to view or update configuration parameters:

- db2AutoConfig - Access the Configuration Advisor
- db2CfgGet - Get the database manager or database configuration parameters
- db2CfgSet - Set the database manager or database configuration parameters

Updating configuration parameters using the Configuration Assistant

The Configuration Assistant can also be used to set the database manager configuration parameters on a client. Other parameters can be changed online; these are called *configurable online configuration parameters*.

Viewing updated configuration values

For some database manager configuration parameters, the database manager must be stopped (db2stop) and then restarted (db2start) for the new parameter values to take effect.

For some database parameters, changes will only take effect when the database is reactivated, or switched from offline to online. In these cases, all applications must first disconnect from the database. (If the database was activated, or switched from offline to online, then it must be deactivated and reactivated.) Then, at the first new connect to the database, the changes will take effect.

If you change the setting of a configurable online database manager configuration parameter while you are attached to an instance, the default behavior of the UPDATE DBM CFG command will be to apply the change immediately. If you do not want the change applied immediately, use the DEFERRED option on the UPDATE DBM CFG command.

To change a database manager configuration parameter online:

```
db2 attach to <instance-name>
db2 update dbm cfg using <parameter-name> <value>
db2 detach
```

For clients, changes to the database manager configuration parameters take effect the next time the client connects to a server.

If you change a configurable online database configuration parameter while connected, the default behavior is to apply the change online, wherever possible. You should note that some parameter changes might take a noticeable amount of time to take effect due to the overhead associated with allocating space. To change configuration parameters online from the command line processor, a connection to the database is required. To change a database configuration parameter online:

```
db2 connect to <dbname>
db2 update db cfg using <parameter-name> <parameter-value>
db2 connect reset
```


Each configurable online configuration parameter has a *propagation class* associated with it. The propagation class indicates when you can expect a change to the configuration parameter to take effect. There are three propagation classes:

- **Immediate:** Parameters that change immediately upon command or API invocation. For example, *diaglevel* has a propagation class of immediate.
- **Statement boundary:** Parameters that change on statement and statement-like boundaries. For example, if you change the value of *sortheap*, all new requests will start using the new value.
- **Transaction boundary:** Parameters that change on transaction boundaries. For example, a new value for *dl_expint* is updated after a COMMIT statement.

While new parameter values might not be immediately effective, viewing the parameter settings (using the GET DATABASE MANAGER CONFIGURATION or GET DATABASE CONFIGURATION command) will always show the latest updates. Viewing the parameter settings using the SHOW DETAIL clause on these commands will show both the latest updates and the values in memory.

Rebinding applications after updating database configuration parameters

Changing some database configuration parameters can influence the access plan chosen by the SQL and XQuery optimizer. After changing any of these parameters, you should consider rebinding your applications to ensure the best access plan is being used for your SQL and XQuery statements. Any parameters that were modified online (for example, by using the UPDATE DATABASE CONFIGURATION IMMEDIATE command) will cause the SQL and XQuery optimizer to choose new access plans for new query statements. However, the query statement cache will not be purged of existing entries. To clear the contents of the query cache, use the FLUSH PACKAGE CACHE statement.

Note: A number of configuration parameters (for example, *userexit*) are described as having acceptable values of either “Yes” or “No”, or “On” or “Off” in the help and other DB2 documentation. To clarify, “Yes” should be considered equivalent to “On” and “No” should be considered equivalent to “Off”.

Configuration parameters summary

The following tables list the parameters in the database manager and database configuration files for database servers. When changing the database manager and database configuration parameters, consider the detailed information for each parameter. Specific operating environment information including defaults is part of each parameter description.

Database Manager Configuration Parameter Summary

For some database manager configuration parameters, the database manager must be stopped (db2stop) and restarted (db2start) for the new parameter values to take effect. Other parameters can be changed online; these are called *configurable online configuration parameters*. If you change the setting of a configurable online database manager configuration parameter while you are attached to an instance, the default behavior of the UPDATE DBM CFG command applies the change immediately. If you do not want the change applied immediately, use the DEFERRED option on the UPDATE DBM CFG command.

The column “Auto” in the following table indicates whether the parameter supports the AUTOMATIC keyword on the UPDATE DBM CFG command.

When updating a parameter to automatic, it is also possible to specify a starting value as well as the AUTOMATIC keyword. Note that the value can mean something different for each parameter, and in some cases it is not applicable. Before specifying a value, read the parameter’s documentation to determine what it represents. In the following example, **num_poolagents** will be updated to AUTOMATIC and the database manager will use 20 as the minimum number of idle agents to pool:

```
db2 update dbm cfg using num_poolagents 20 automatic
```

To unset the AUTOMATIC feature, the parameter can be updated to a value or the MANUAL keyword can be used. When a parameter is updated to MANUAL, the parameter is no longer automatic and is set to its current value (as displayed in the Current Value column from the GET DBM CFG SHOW DETAIL or GET DB CFG SHOW DETAIL commands).

The column “Perf. Impact” provides an indication of the relative importance of each parameter as it relates to system performance. It is impossible for this column to apply accurately to all environments; you should view this information as a generalization.

- **High** — Indicates the parameter can have a significant impact on performance. You should consciously decide the values of these parameters, which, in some cases, means that you will accept the default values provided.
- **Medium** — Indicates that the parameter can have some impact on performance. Your specific environment and needs will determine how much tuning effort should be focused on these parameters.
- **Low** — Indicates that the parameter has a less general or less significant impact on performance.
- **None** — Indicates that the parameter does not directly impact performance. Although you do not have to tune these parameters for performance enhancement, they can be very important for other aspects of your system configuration, such as communication support, for example.

The columns “Token”, “Token Value”, and “Data Type” provide information that you will need when calling the db2CfgGet or the db2CfgSet API. This information includes configuration parameter identifiers, entries for the **token** element in the db2CfgParam data structure, and data types for values that are passed to the structure.

Table 66. Configurable Database Manager Configuration Parameters

Parameter	Cfg. Online	Auto.	Perf. Impact	Token	Token Value	Data Type	Additional Information
agent_stack_sz	No	No	Low	SQLF_KTN_AGENT_STACK_SZ	61	UInt16	“agent_stack_sz - Agent stack size” on page 460
agentpri	No	No	High	SQLF_KTN_AGENTPRI	26	Sint16	“agentpri - Priority of agents” on page 461
aslheapsz	No	No	High	SQLF_KTN_ASLHEAPSZ	15	UInt32	“aslheapsz - Application support layer heap size” on page 463
audit_buf_sz	No	No	High	SQLF_KTN_AUDIT_BUF_SZ	312	Sint32	“audit_buf_sz - Audit buffer size” on page 464
authentication ¹	No	No	Low	SQLF_KTN_AUTHENTICATION	78	UInt16	“authentication - Authentication type” on page 465

Table 66. Configurable Database Manager Configuration Parameters (continued)

Parameter	Cfg. Online	Auto.	Perf. Impact	Token	Token Value	Data Type	Additional Information
catalog_noauth	Yes	No	None	SQLF_KTN_CATALOG_NOAUTH	314	Uint16	"catalog_noauth - Cataloging allowed without authority" on page 466
clnt_krb_plugin	No	No	None	SQLF_KTN_CLNT_KRB_PLUGIN	812	char(33)	"clnt_krb_plugin - Client Kerberos plug-in" on page 467
clnt_pw_plugin	No	No	None	SQLF_KTN_CLNT_PW_PLUGIN	811	char(33)	"clnt_pw_plugin - Client userid-password plug-in" on page 467
cluster_mgr	No	No	None	SQLF_KTN_CLUSTER_MGR	920	char(262)	"cluster_mgr - Cluster manager name" on page 468
comm_bandwidth	Yes	No	Medium	SQLF_KTN_COMM_BANDWIDTH	307	float	"comm_bandwidth - Communications bandwidth" on page 468
conn_elapse	Yes	No	Medium	SQLF_KTN_CONN_ELAPSE	508	Uint16	"conn_elapse - Connection elapse time" on page 469
cpuspeed	Yes	No	High	SQLF_KTN_CPUSPEED	42	float	"cpuspeed - CPU speed" on page 469
dft_account_str	Yes	No	None	SQLF_KTN_DFT_ACCOUNT_STR	28	char(25)	"dft_account_str - Default charge-back account" on page 470
dft_monswitches • dft_mon_bufpool • dft_mon_lock • dft_mon_sort • dft_mon_stmt • dft_mon_table • dft_mon_timestamp • dft_mon_uow	Yes	No	Medium	SQLF_KTN_DFT_MONSWITCHES ² • SQLF_KTN_DFT_MON_BUFPOOL • SQLF_KTN_DFT_MON_LOCK • SQLF_KTN_DFT_MON_SORT • SQLF_KTN_DFT_MON_STMT • SQLF_KTN_DFT_MON_TABLE • SQLF_KTN_DFT_MON_TIMESTAMP • SQLF_KTN_DFT_MON_UOW	29 • 33 • 34 • 35 • 31 • 32 • 36 • 30	Uint16 • Uint16 • Uint16 • Uint16 • Uint16 • Uint16 • Uint16 • Uint16	"dft_monswitches - Default database system monitor switches" on page 471
dftdbpath	Yes	No	None	SQLF_KTN_DFTDBPATH	27	char(215)	"dftdbpath - Default database path" on page 472
diaglevel	Yes	No	Low	SQLF_KTN_DIAGLEVEL	64	Uint16	"diaglevel - Diagnostic error capture level" on page 473
diagpath	Yes	No	None	SQLF_KTN_DIAGPATH	65	char(215)	"diagpath - Diagnostic data directory path" on page 474
dir_cache	No	No	Medium	SQLF_KTN_DIR_CACHE	40	Uint16	"dir_cache - Directory cache support" on page 475
discover ³	No	No	Medium	SQLF_KTN_DISCOVER	304	Uint16	"discover - Discovery mode" on page 476
discover_inst	Yes	No	Low	SQLF_KTN_DISCOVER_INST	308	Uint16	"discover_inst - Discover server instance" on page 477
fcm_num_buffers	Yes	Yes	Medium	SQLF_KTN_FCM_NUM_BUFFERS	503	Uint32	"fcm_num_buffers - Number of FCM buffers" on page 477
fcm_num_channels	Yes	Yes	Medium	SQLF_KTN_FCM_NUM_CHANNELS	902	Uint32	"fcm_num_channels - Number of FCM channels" on page 478
fed_noauth	Yes	No	None	SQLF_KTN_FED_NOAUTH	806	Uint16	"fed_noauth - Bypass federated authentication" on page 479
federated	Yes	No	Medium	SQLF_KTN_FEDERATED	604	Sint16	"federated - Federated database system support" on page 479
federated_async	Yes	Yes	Medium	SQLF_KTN_FEDERATED_ASYNC	849	Sint32	"federated_async - Maximum asynchronous TQs per query configuration parameter" on page 480

Table 66. Configurable Database Manager Configuration Parameters (continued)

Parameter	Cfg. Online	Auto.	Perf. Impact	Token	Token Value	Data Type	Additional Information
fenced_pool	Yes	Yes	Medium	SQLF_KTN_FENCED_POOL	80	Sint32	"fenced_pool - Maximum number of fenced processes" on page 480
group_plugin	No	No	None	SQLF_KTN_GROUP_PLUGIN	810	char(33)	"group_plugin - Group plug-in" on page 482
health_mon	Yes	No	Low	SQLF_KTN_HEALTH_MON	804	UInt16	"health_mon - Health monitoring" on page 482
indexrec ⁴	Yes	No	Medium	SQLF_KTN_INDEXREC	20	UInt16	"indexrec - Index re-creation time" on page 483
instance_memory	Yes	Yes	Medium	SQLF_KTN_INSTANCE_MEMORY	803	UInt64	"instance_memory - Instance memory" on page 485
intra_parallel	No	No	High	SQLF_KTN_INTRA_PARALLEL	306	Sint16	"intra_parallel - Enable intra-partition parallelism" on page 487
java_heap_sz	No	No	High	SQLF_KTN_JAVA_HEAP_SZ	310	Sint32	"java_heap_sz - Maximum Java interpreter heap size" on page 487
jdk_path	No	No	None	SQLF_KTN_JDK_PATH	311	char(255)	"jdk_path - Software Developer's Kit for Java installation path" on page 488
keepfenced	No	No	Medium	SQLF_KTN_KEEPFENCED	81	UInt16	"keepfenced - Keep fenced process" on page 489
local_gssplugin	No	No	None	SQLF_KTN_LOCAL_GSSPLUGIN	816	char(33)	"local_gssplugin - GSS API plug-in used for local instance level authorization" on page 490
max_connections	Yes	Yes	Medium	SQLF_KTN_MAX_CONNECTIONS	802	Sint32	"max_connections - Maximum number of client connections" on page 490
max_connretries	Yes	No	Medium	SQLF_KTN_MAX_CONNRETRIES	509	UInt16	"max_connretries - Node connection retries" on page 491
max_coordagents	Yes	Yes	Medium	SQLF_KTN_MAX_COORDAGENTS	501	Sint32	"max_coordagents - Maximum number of coordinating agents" on page 491
max_querydegree	Yes	No	High	SQLF_KTN_MAX_QUERYDEGREE	303	Sint32	"max_querydegree - Maximum query degree of parallelism" on page 492
max_time_diff	No	No	Medium	SQLF_KTN_MAX_TIME_DIFF	510	UInt16	"max_time_diff - Maximum time difference among nodes" on page 493
mon_heap_sz	Yes	Yes	Low	SQLF_KTN_MON_HEAP_SZ	79	UInt16	"mon_heap_sz - Database system monitor heap size" on page 495
notifylevel	Yes	No	Low	SQLF_KTN_NOTIFYLEVEL	605	Sint16	"notifylevel - Notify level" on page 496
num_initagents	No	No	Medium	SQLF_KTN_NUM_INITAGENTS	500	UInt32	"num_initagents - Initial number of agents in pool" on page 497
num_initfenced	No	No	Medium	SQLF_KTN_NUM_INITFENCED	601	Sint32	"num_initfenced - Initial number of fenced processes" on page 498
num_poolagents	Yes	Yes	High	SQLF_KTN_NUM_POOLAGENTS	502	Sint32	"num_poolagents - Agent pool size" on page 498
numdb	No	No	Low	SQLF_KTN_NUMDB	6	UInt16	"numdb - Maximum number of concurrently active databases including host and System i databases" on page 499
query_heap_sz	No	No	Medium	SQLF_KTN_QUERY_HEAP_SZ	49	Sint32	"query_heap_sz - Query heap size" on page 500
resync_interval	No	No	None	SQLF_KTN_RESYNC_INTERVAL	68	UInt16	"resync_interval - Transaction resync interval" on page 501

Table 66. Configurable Database Manager Configuration Parameters (continued)

Parameter	Cfg. Online	Auto.	Perf. Impact	Token	Token Value	Data Type	Additional Information
rqrioblk	No	No	High	SQLF_KTN_RQRIOBLK	1	Uint16	"rqrioblk - Client I/O block size" on page 502
sheapthres	No	No	High	SQLF_KTN_SHEAPTHRES	21	Uint32	"sheapthres - Sort heap threshold" on page 503
spm_log_file_sz	No	No	Low	SQLF_KTN_SPM_LOG_FILE_SZ	90	Sint32	"spm_log_file_sz - Sync point manager log file size" on page 504
spm_log_path	No	No	Medium	SQLF_KTN_SPM_LOG_PATH	313	char(226)	"spm_log_path - Sync point manager log file path" on page 505
spm_max_resync	No	No	Low	SQLF_KTN_SPM_MAX_RESYNC	91	Sint32	"spm_max_resync - Sync point manager resync agent limit" on page 506
spm_name	No	No	None	SQLF_KTN_SPM_NAME	92	char(8)	"spm_name - Sync point manager name" on page 506
srvcon_auth	No	No	None	SQLF_KTN_SRVCON_AUTH	815	Uint16	"srvcon_auth - Authentication type for incoming connections at the server" on page 506
srvcon_gssplugin_list	No	No	None	SQLF_KTN_SRVCON_GSSPLUGIN_LIST	814	char(256)	"srvcon_gssplugin_list - List of GSS API plug-ins for incoming connections at the server" on page 507
srv_plugin_mode	No	No	None	SQLF_KTN_SRV_PLUGIN_MODE	809	Uint16	"srv_plugin_mode - Server plug-in mode" on page 508
srvcon_pw_plugin	No	No	None	SQLF_KTN_SRVCON_PW_PLUGIN	813	char(33)	"srvcon_pw_plugin - Userid-password plug-in for incoming connections at the server" on page 507
start_stop_time	Yes	No	Low	SQLF_KTN_START_STOP_TIME	511	Uint16	"start_stop_time - Start and stop timeout" on page 508
svcname	No	No	None	SQLF_KTN_SVCENAME	24	char(14)	"svcname - TCP/IP service name" on page 509
sysadm_group	No	No	None	SQLF_KTN_SYSADM_GROUP	39	char(128)	"sysadm_group - System administration authority group name" on page 510
sysctrl_group	No	No	None	SQLF_KTN_SYSCTRL_GROUP	63	char(128)	"sysctrl_group - System control authority group name" on page 510
sysmaint_group	No	No	None	SQLF_KTN_SYSMaint_GROUP	62	char(128)	"sysmaint_group - System maintenance authority group name" on page 511
sysmon_group	No	No	None	SQLF_KTN_SYSMON_GROUP	808	char(128)	"sysmon_group - System monitor authority group name" on page 512
tm_database	No	No	None	SQLF_KTN_TM_DATABASE	67	char(8)	"tm_database - Transaction manager database name" on page 512
tp_mon_name	No	No	None	SQLF_KTN_TP_MON_NAME	66	char(19)	"tp_mon_name - Transaction processor monitor name" on page 513
trust_allclnts ⁵	No	No	None	SQLF_KTN_TRUST_ALLCLNTS	301	Uint16	"trust_allclnts - Trust all clients" on page 514
trust_clntauth	No	No	None	SQLF_KTN_TRUST_CLNTAUTH	302	Uint16	"trust_clntauth - Trusted clients authentication" on page 515
util_impact_lim	Yes	No	High	SQLF_KTN_UTIL_IMPACT_LIM	807	Uint32	"util_impact_lim - Instance impact policy" on page 516

Table 66. Configurable Database Manager Configuration Parameters (continued)

Parameter	Cfg. Online	Auto.	Perf. Impact	Token	Token Value	Data Type	Additional Information
Note:							
1. The valid values are defined in sqlenv.h.							
2.							
Bit 1 (xxxx xxx1): dft_mon_uow Bit 2 (xxxx xx1x): dft_mon_stmt Bit 3 (xxxx x1xx): dft_mon_table Bit 4 (xxxx lxxx): dft_mon_buffpool Bit 5 (xxx1 xxxx): dft_mon_lock Bit 6 (xx1x xxxx): dft_mon_sort Bit 7 (x1xx xxxx): dft_mon_timestamp							
3. Valid values (defined in sqlutil.h):							
SQLF_DSCVR_KNOWN (1) SQLF_DSCVR_SEARCH (2)							
4. Valid values (defined in sqlutil.h):							
SQLF_INX_REC_SYSTEM (0) SQLF_INX_REC_REFERENCE (1)							
5. Valid values (defined in sqlutil.h):							
SQLF_TRUST_ALLCLNTS_NO (0) SQLF_TRUST_ALLCLNTS_YES (1) SQLF_TRUST_ALLCLNTS_DRDAONLY (2)							

Table 67. Informational Database Manager Configuration Parameters

Parameter	Token	Token Value	Data Type	Additional Information
nodetype ¹	SQLF_KTN_NODETYPE	100	UInt16	"nodetype - Machine node type" on page 496
release	SQLF_KTN_RELEASE	101	UInt16	"release - Configuration file release level" on page 501
Note:				
1. Valid values (defined in sqlutil.h):				
SQLF_NT_STANDALONE (0) SQLF_NT_SERVER (1) SQLF_NT_REQUESTOR (2) SQLF_NT_STAND_REQ (3) SQLF_NT_MPP (4) SQLF_NT_SATELLITE (5)				

Database Configuration Parameter Summary

The following table lists the parameters in the database configuration file. When changing the database configuration parameters, consider the detailed information for the parameter.

For some database configuration parameters, changes only take effect when the database is reactivated. In these cases, all applications must first disconnect from the database. (If the database was activated, then it must be deactivated and reactivated.) The changes take effect at the next connection to the database. Other parameters can be changed online; these are called *configurable online configuration parameters*.

Refer to the Database Manager Configuration Parameter Summary section above for a description of the "Auto.", "Perf. Impact", "Token", "Token Value", and "Data Type" columns.

The AUTOMATIC keyword is also supported on the UPDATE DB CFG command. In the following example, **database_memory** will be updated to AUTOMATIC and the database manager will use 20000 as a starting value when making further changes to this parameter:

```
db2 update db cfg using for sample using database_memory 20000 automatic
```

Starting with Version 9.5, you can update and reset database configuration parameter values across some or all platforms without having to issue the db2_all command, or without having to update or reset each partition individually. For details, see Configuring databases across multiple partitions.

Table 68. Configurable Database Configuration Parameters

Parameter	Cfg. Online	Auto.	Perf. Impact	Token	Token Value	Data Type	Additional Information
alt_collate	No	No	None	SQLF_DBTN_ALT_COLLATE	809	Uint32	"alt_collate - Alternate collating sequence" on page 517
applheapsz	Yes	Yes	Medium	SQLF_DBTN_APPLHEAPSZ	51	Uint16	"applheapsz - Application heap size" on page 520
appl_memory	Yes	Yes	Medium	SQLF_DBTN_APPL_MEMORY	904	Uint64	"appl_memory - Application Memory configuration parameter" on page 519
archretrydelay	Yes	No	None	SQLF_DBTN_ARCHRETRYDELAY	828	Uint16	"archretrydelay - Archive retry delay on error" on page 521
<ul style="list-style-type: none"> • auto_maint • auto_db_backup • auto_tbl_maint • auto_runstats • auto_stats_prof • auto_stmt_stats • auto_prof_upd • auto_reorg 	Yes	No	Medium	<ul style="list-style-type: none"> • SQLF_DBTN_AUTO_MAINT • SQLF_DBTN_AUTO_DB_BACKUP • SQLF_DBTN_AUTO_TBL_MAINT • SQLF_DBTN_AUTO_RUNSTATS • SQLF_DBTN_AUTO_STATS_PROF • SQLF_DBTN_AUTO_STMT_STATS • SQLF_DBTN_AUTO_PROF_UPD • SQLF_DBTN_AUTO_REORG 	<ul style="list-style-type: none"> • 831 • 833 • 835 • 837 • 839 • 905 • 844 • 841 	Uint16	"auto_maint - Automatic maintenance" on page 522
auto_del_rec_obj	Yes	No	Medium	SQLF_DBTN_AUTO_DEL_REC_OBJ	912	Uint16	"auto_del_rec_obj - Automated deletion of recovery objects configuration parameter" on page 521
autorestart	Yes	No	Low	SQLF_DBTN_AUTO_RESTART	25	Uint16	"autorestart - Auto restart enable" on page 524
avg_appls	Yes	Yes	High	SQLF_DBTN_AVG_APPLS	47	Uint16	"avg_appls - Average number of active applications" on page 524
blk_log_dsk_ful	Yes	No	None	SQLF_DBTN_BLK_LOG_DSK_FUL	804	Uint16	"blk_log_dsk_ful - Block on log disk full" on page 525
catalogcache_sz	Yes	No	Medium	SQLF_DBTN_CATALOGCACHE_SZ	56	Sint32	"catalogcache_sz - Catalog cache size" on page 526
chnpggs_thresh	No	No	High	SQLF_DBTN_CHNGPGS_THRESH	38	Uint16	"chnpggs_thresh - Changed pages threshold" on page 527
database_memory	Yes	Yes	Medium	SQLF_DBTN_DATABASE_MEMORY	803	Uint64	"database_memory - Database shared memory size" on page 530
dbheap	Yes	Yes	Medium	SQLF_DBTN_DB_HEAP	58	Uint64	"dbheap - Database heap" on page 533
db_mem_thresh	Yes	No	Low	SQLF_DBTN_DB_MEM_THRESH	849	Uint16	"db_mem_thresh - Database memory threshold" on page 532

Table 68. Configurable Database Configuration Parameters (continued)

Parameter	Cfg. Online	Auto.	Perf. Impact	Token	Token Value	Data Type	Additional Information
decflt_rounding	No	No	None	SQLF_DBTN_DECFLT_ROUNDING	913	Unit16	"decflt_rounding - Decimal floating point rounding configuration parameter" on page 534
dft_degree	Yes	No	High	SQLF_DBTN_DFT_DEGREE	301	Sint32	"dft_degree - Default degree" on page 536
dft_extent_sz	Yes	No	Medium	SQLF_DBTN_DFT_EXTENT_SZ	54	Uint32	"dft_extent_sz - Default extent size of table spaces" on page 536
dft_loadrec_ses	Yes	No	Medium	SQLF_DBTN_DFT_LOADREC_SES	42	Sint16	"dft_loadrec_ses - Default number of load recovery sessions" on page 537
dft_mttb_types	No	No	None	SQLF_DBTN_DFT_MTTB_TYPES	843	Uint32	"dft_mttb_types - Default maintained table types for optimization" on page 537
dft_prefetch_sz	Yes	Yes	Medium	SQLF_DBTN_DFT_PREFETCH_SZ	40	Sint16	"dft_prefetch_sz - Default prefetch size" on page 538
dft_queryopt	Yes	No	Medium	SQLF_DBTN_DFT_QUERYOPT	57	Sint32	"dft_queryopt - Default query optimization class" on page 539
dft_refresh_age	No	No	Medium	SQLF_DBTN_DFT_REFRESH_AGE	702	char(22)	"dft_refresh_age - Default refresh age" on page 539
dft_sqlmathwarn	No	No	None	SQLF_DBTN_DFT_SQLMATHWARN	309	Sint16	"dft_sqlmathwarn - Continue upon arithmetic exceptions" on page 540
discover_db	Yes	No	Medium	SQLF_DBTN_DISCOVER	308	Uint16	"discover_db - Discover database" on page 541
dlchktime	Yes	No	Medium	SQLF_DBTN_DLCHKTIME	9	Uint32	"dlchktime - Time interval for checking deadlock" on page 541
dyn_query_mgmt	No	No	Low	SQLF_DBTN_DYN_QUERY_MGMT	604	Uint16	"dyn_query_mgmt - Dynamic SQL and XQuery query management" on page 542
enable_xmlchar	Yes	No	None	SQLF_DBTN_ENABLE_XMLCHAR	853	Uint32	"enable_xmlchar - Enable conversion to XML configuration parameter" on page 543
failarchpath	Yes	No	None	SQLF_DBTN_FAILARCHPATH	826	char(243)	"failarchpath - Failover log archive path" on page 543
hadr_local_host	No	No	None	SQLF_DBTN_HADR_LOCAL_HOST	811	char(256)	"hadr_local_host - HADR local host name" on page 544
hadr_local_svc	No	No	None	SQLF_DBTN_HADR_LOCAL_SVC	812	char(41)	"hadr_local_svc - HADR local service name" on page 545
hadr_peer_window	No	No	Low (see Note 4)	SQLF_DBTN_HADR_PEER_WINDOW	914	Uint32	"hadr_peer_window - HADR peer window configuration parameter" on page 545
hadr_remote_host	No	No	None	SQLF_DBTN_HADR_REMOTE_HOST	813	char(256)	"hadr_remote_host - HADR remote host name" on page 546
hadr_remote_inst	No	No	None	SQLF_DBTN_HADR_REMOTE_INST	815	char(9)	"hadr_remote_inst - HADR instance name of the remote server" on page 546
hadr_remote_svc	No	No	None	SQLF_DBTN_HADR_REMOTE_SVC	814	char(41)	"hadr_remote_svc - HADR remote service name" on page 547
hadr_syncmode	No	No	None	SQLF_DBTN_HADR_SYNCMODE	817	Uint32	"hadr_syncmode - HADR synchronization mode for log write in peer state" on page 547
hadr_timeout	No	No	None	SQLF_DBTN_HADR_TIMEOUT	816	Uint32	"hadr_timeout - HADR timeout value" on page 548
indexrec ²	Yes	No	Medium	SQLF_DBTN_INDEXREC	30	Uint16	"indexrec - Index re-creation time" on page 483

Table 68. Configurable Database Configuration Parameters (continued)

Parameter	Cfg. Online	Auto.	Perf. Impact	Token	Token Value	Data Type	Additional Information
locklist	Yes	Yes	High when it affects escalation	SQLF_DBTN_LOCK_LIST	704	UInt64	"locklist - Maximum storage for lock list" on page 551
locktimeout	No	No	Medium	SQLF_DBTN_LOCKTIMEOUT	34	Sint16	"locktimeout - Lock timeout" on page 553
logarchmeth1	Yes	No	None	SQLF_DBTN_LOGARCHMETH1	822	char(252)	"logarchmeth1 - Primary log archive method" on page 555
logarchmeth2	Yes	No	None	SQLF_DBTN_LOGARCHMETH2	823	char(252)	"logarchmeth2 - Secondary log archive method" on page 556
logarchopt1	Yes	No	None	SQLF_DBTN_LOGARCHOPT1	824	char(243)	"logarchopt1 - Primary log archive options" on page 557
logarchopt2	Yes	No	None	SQLF_DBTN_LOGARCHOPT2	825	char(243)	"logarchopt2 - Secondary log archive options" on page 557
logbufsz	No	No	High	SQLF_DBTN_LOGBUFSZ	33	UInt16	"logbufsz - Log buffer size" on page 558
logfilsiz	No	No	Medium	SQLF_DBTN_LOGFIL_SIZ	92	UInt32	"logfilsiz - Size of log files" on page 558
logindexbuild	Yes	Yes	None	SQLF_DBTN_LOGINDEXBUILD	818	UInt32	"logindexbuild - Log index pages created" on page 560
logprimary	No	No	Medium	SQLF_DBTN_LOGPRIMARY	16	UInt16	"logprimary - Number of primary log files" on page 560
logretain ³	No	No	Low	SQLF_DBTN_LOG_RETAIN	23	UInt16	"logretain - Log retain enable" on page 562
logsecond	Yes	No	Medium	SQLF_DBTN_LOGSECOND	17	UInt16	"logsecond - Number of secondary log files" on page 562
max_log	Yes	Yes		SQLF_DBTN_MAX_LOG	807	UInt16	"max_log - Maximum log per transaction" on page 564
maxappls	Yes	Yes	Medium	SQLF_DBTN_MAXAPPLS	6	UInt16	"maxappls - Maximum number of active applications" on page 564
maxfilop	Yes	No	Medium	SQLF_DBTN_MAXFILOP	3	UInt16	"maxfilop - Maximum database files open per application" on page 565
maxlocks	Yes	Yes	High when it affects escalation	SQLF_DBTN_MAXLOCKS	15	UInt16	"maxlocks - Maximum percent of lock list before escalation" on page 566
min_dec_div_3	No	No	High	SQLF_DBTN_MIN_DEC_DIV_3	605	Sint32	"min_dec_div_3 - Decimal division scale to 3" on page 568
mincommit	Yes	No	High	SQLF_DBTN_MINCOMMIT	32	UInt16	"mincommit - Number of commits to group" on page 569
mirrorlogpath	No	No	Low	SQLF_DBTN_MIRRORLOGPATH	806	char(242)	"mirrorlogpath - Mirror log path" on page 570
newlogpath	No	No	Low	SQLF_DBTN_NEWLOGPATH	20	char(242)	"newlogpath - Change the database log path" on page 571
num_db_backups	Yes	No	None	SQLF_DBTN_NUM_DB_BACKUPS	601	UInt16	"num_db_backups - Number of database backups" on page 573
num_freqvalues	Yes	No	Low	SQLF_DBTN_NUM_FREQVALUES	36	UInt16	"num_freqvalues - Number of frequent values retained" on page 573
num_iocleaners	No	Yes	High	SQLF_DBTN_NUM_IOCLEANERS	37	UInt16	"num_iocleaners - Number of asynchronous page cleaners" on page 574

Table 68. Configurable Database Configuration Parameters (continued)

Parameter	Cfg. Online	Auto.	Perf. Impact	Token	Token Value	Data Type	Additional Information
num_ioservers	No	Yes	High	SQLF_DBTN_NUM_IOSERVERS	39	Uint16	"num_ioservers - Number of I/O servers" on page 576
num_log_span	Yes	Yes		SQLF_DBTN_NUM_LOG_SPAN	808	Uint16	"num_log_span - Number log span" on page 576
num_quantiles	Yes	No	Low	SQLF_DBTN_NUM_QUANTILES	48	Uint16	"num_quantiles - Number of quantiles for columns" on page 577
numarchretry	Yes	No	None	SQLF_DBTN_NUMARCHRETRY	827	Uint16	"numarchretry - Number of retries on error" on page 578
overflowlogpath	No	No	Medium	SQLF_DBTN_OVERFLOWLOGPATH	805	char(242)	"overflowlogpath - Overflow log path" on page 579
pckcachesz	Yes	Yes	High	SQLF_DBTN_PCKCACHE_SZ	505	Uint32	"pckcachesz - Package cache size" on page 580
rec_his_retentn	No	No	None	SQLF_DBTN_REC_HIS_RETENTN	43	Sint16	"rec_his_retentn - Recovery history retention period" on page 583
self_tuning_mem	Yes	No	High	SQLF_DBTN_SELF_TUNING_MEM	848	Uint16	"self_tuning_mem- Self-tuning memory" on page 584
seqdetect	Yes	No	High	SQLF_DBTN_SEQDETECT	41	Uint16	"seqdetect - Sequential detection flag" on page 585
sheapthres_shr	Yes	Yes	High	SQLF_DBTN_SHEAPTHRES_SHR	802	Uint32	"sheapthres_shr - Sort heap threshold for shared sorts" on page 586
softmax	No	No	Medium	SQLF_DBTN_SOFTMAX	5	Uint16	"softmax - Recovery range and soft checkpoint interval" on page 587
sortheap	Yes	Yes	High	SQLF_DBTN_SORT_HEAP	52	Uint32	"sortheap - Sort heap size" on page 589
stat_heap_sz	Yes	Yes	Low	SQLF_DBTN_STAT_HEAP_SZ	45	Uint32	"stat_heap_sz - Statistics heap size" on page 590
stmtheap	Yes	Yes	Medium	SQLF_DBTN_STMT_HEAP	821	Uint32	"stmtheap - Statement heap size" on page 591
trackmod	No	No	Low	SQLF_DBTN_TRACKMOD	703	Uint16	"trackmod - Track modified pages enable" on page 592
tsm_mgmtclass	Yes	No	None	SQLF_DBTN_TSM_MGMTCLASS	307	char(30)	"tsm_mgmtclass - Tivoli Storage Manager management class" on page 592
tsm_nodename	Yes	No	None	SQLF_DBTN_TSM_NODENAME	306	char(64)	"tsm_nodename - Tivoli Storage Manager node name" on page 592
tsm_owner	Yes	No	None	SQLF_DBTN_TSM_OWNER	305	char(64)	"tsm_owner - Tivoli Storage Manager owner name" on page 593
tsm_password	Yes	No	None	SQLF_DBTN_TSM_PASSWORD	501	char(64)	"tsm_password - Tivoli Storage Manager password" on page 593
userexit	No	No	Low	SQLF_DBTN_USER_EXIT	24	Uint16	"userexit - User exit enable" on page 594
util_heap_sz	Yes	No	Low	SQLF_DBTN_UTIL_HEAP_SZ	55	Uint32	"util_heap_sz - Utility heap size" on page 594
vendoropt	Yes	No	None	SQLF_DBTN_VENDOROPT	829	char(242)	"vendoropt - Vendor options" on page 595<
wlm_collect_int	Yes	No	Low	SQLF_DBTN_WLM_COLLECT_INT	907	Sint32	"wlm_collect_int - Workload management collection interval configuration parameter" on page 595

Table 68. Configurable Database Configuration Parameters (continued)

Parameter	Cfg. Online	Auto.	Perf. Impact	Token	Token Value	Data Type	Additional Information
<p>Note: The bits of SQLF_DBTN_AUTONOMIC_SWITCHES indicate the default settings for a number of auto-maintenance configuration parameters. The individual bits making up this composite parameter are:</p> <p>1.</p> <pre> Default => Bit 1 on (xxxx xxxx xxxx xxx1): auto_maint Bit 2 off (xxxx xxxx xxxx xx0x): auto_db_backup Bit 3 on (xxxx xxxx xxxx x1xx): auto_tbl_maint Bit 4 on (xxxx xxxx xxxx lxxx): auto_runstats Bit 5 off (xxxx xxxx xxx0 xxxx): auto_stats_prof Bit 6 off (xxxx xxxx xx0x xxxx): auto_prof_upd Bit 7 off (xxxx xxxx x0xx xxxx): auto_reorg Bit 8 off (xxxx xxxx 0xxx xxxx): auto_storage Bit 9 off (xxxx xxx0 xxxx xxxx): auto_stmt_stats 0 0 0 0 Maximum => Bit 1 on (xxxx xxxx xxxx xxx1): auto_maint Bit 2 off (xxxx xxxx xxxx xx1x): auto_db_backup Bit 3 on (xxxx xxxx xxxx l1xx): auto_tbl_maint Bit 4 on (xxxx xxxx xxxx lxxx): auto_runstats Bit 5 off (xxxx xxxx xxx1 xxxx): auto_stats_prof Bit 6 off (xxxx xxxx xx1x xxxx): auto_prof_upd Bit 7 off (xxxx xxxx x1xx xxxx): auto_reorg Bit 8 off (xxxx xxxx lxxx xxxx): auto_storage Bit 9 off (xxxx xxx1 xxxx xxxx): auto_stmt_stats 0 1 F F </pre> <p>2. Valid values (defined in sqlutil.h):</p> <pre> SQLF_INX_REC_SYSTEM (0) SQLF_INX_REC_REFERENCE (1) SQLF_INX_REC_RESTART (2) </pre> <p>3. Valid values (defined in sqlutil.h):</p> <pre> SQLF_LOGRETAIN_NO (0) SQLF_LOGRETAIN_RECOVERY (1) SQLF_LOGRETAIN_CAPTURE (2) </pre> <p>4. If you set the hadr_peer_window parameter to a non-zero time value, then the primary database might seem to hang on transactions when it is in disconnected peer state, because it is waiting for confirmation from the standby database even though it is not connected to the standby database.</p>							

Table 69. Informational Database Configuration Parameters

Parameter	Token	Token Value	Data Type	Additional Information
backup_pending	SQLF_DBTN_BACKUP_PENDING	112	UInt16	"backup_pending - Backup pending indicator" on page 525
codepage	SQLF_DBTN_CODEPAGE	101	UInt16	"codepage - Code page for the database" on page 528
codeset	SQLF_DBTN_CODESET	120	char(9) ¹	"codeset - Codeset for the database" on page 528
collate_info	SQLF_DBTN_COLLATE_INFO	44	char(260)	"collate_info - Collating information" on page 528
country/region	SQLF_DBTN_COUNTRY	100	UInt16	"country/region - Database territory code" on page 529
database_consistent	SQLF_DBTN_CONSISTENT	111	UInt16	"database_consistent - Database is consistent" on page 529
database_level	SQLF_DBTN_DATABASE_LEVEL	124	UInt16	"database_level - Database release level" on page 530
hadr_db_role	SQLF_DBTN_HADR_DB_ROLE	810	UInt32	"hadr_db_role - HADR database role" on page 544
log_retain_status	SQLF_DBTN_LOG_RETAIN_STATUS	114	UInt16	"log_retain_status - Log retain status indicator" on page 554
loghead	SQLF_DBTN_LOGHEAD	105	char(12)	"loghead - First active log file" on page 560
logpath	SQLF_DBTN_LOGPATH	103	char(242)	"logpath - Location of log files" on page 560
multipage_alloc	SQLF_DBTN_MULTIPAGE_ALLOC	506	UInt16	"multipage_alloc - Multipage file allocation enabled" on page 571
numsegs	SQLF_DBTN_NUMSEGS	122	UInt16	"numsegs - Default number of SMS containers" on page 578

Table 69. Informational Database Configuration Parameters (continued)

Parameter	Token	Token Value	Data Type	Additional Information
pagesize	SQLF_DBTN_PAGESIZE	846	Uint32	"pagesize - Database default page size" on page 580
release	SQLF_DBTN_RELEASE	102	Uint16	"release - Configuration file release level" on page 501
restore_pending	SQLF_DBTN_RESTORE_PENDING	503	Uint16	"restore_pending - Restore pending" on page 583
restrict_access	SQLF_DBTN_RESTRICT_ACCESS	852	Sint32	"restrict_access - Database has restricted access configuration parameter" on page 583
rollfwd_pending	SQLF_DBTN_ROLLFWD_PENDING	113	Uint16	"rollfwd_pending - Roll forward pending indicator" on page 584
territory	SQLF_DBTN_TERRITORY	121	char(5) ²	"territory - Database territory" on page 591
user_exit_status	SQLF_DBTN_USER_EXIT_STATUS	115	Uint16	"user_exit_status - User exit status indicator" on page 594
<p>Note:</p> <ol style="list-style-type: none"> 1. char(17) on HP-UX, Linux and Solaris Operating Environment. 2. char(33) on HP-UX, Linux and Solaris. 				

DB2 Administration Server (DAS) Configuration Parameter Summary

Table 70. DAS Configuration Parameters

Parameter	Parameter Type	Additional Information
authentication	Configurable	"authentication - Authentication type DAS" on page 596
contact_host	Configurable Online	"contact_host - Location of contact list" on page 597
das_codepage	Configurable Online	"das_codepage - DAS code page" on page 597
das_territory	Configurable Online	"das_territory - DAS territory" on page 598
dasadm_group	Configurable	"dasadm_group - DAS administration authority group name" on page 598
db2system	Configurable Online	"db2system - Name of the DB2 server system" on page 599
discover	Configurable Online	"discover - DAS discovery mode" on page 599
exec_exp_task	Configurable	"exec_exp_task - Execute expired tasks" on page 600
jdk_64_path	Configurable Online	"jdk_64_path - 64-Bit Software Developer's Kit for Java installation path DAS" on page 550
jdk_path	Configurable Online	"jdk_path - Software Developer's Kit for Java installation path DAS" on page 600
sched_enable	Configurable	"sched_enable - Scheduler mode" on page 601
sched_userid	Informational	"sched_userid - Scheduler user ID" on page 601
smtp_server	Configurable Online	"smtp_server - SMTP server" on page 601
toolscat_db	Configurable	"toolscat_db - Tools catalog database" on page 602
toolscat_inst	Configurable	"toolscat_inst - Tools catalog database instance" on page 602
toolscat_schema	Configurable	"toolscat_schema - Tools catalog database schema" on page 602

Configuration parameter section headings

Each of the configuration parameter descriptions contain some or all of the following section headings, as applicable. In some cases they are mutually exclusive, for example, valid values are not needed if the [range] is specified. In most cases, these headings are self-explanatory.

Table 71. Description of the configuration parameter section headings

Section heading	Description and possible values
Configuration type	Possible values are: <ul style="list-style-type: none"> • Database manager • Database • DB2 Administration Server
Applies to	If applicable, lists the data server types that the configuration parameter applies to. Possible values are: <ul style="list-style-type: none"> • Client • Database server with local and remote clients • Database server with local clients • DB2 Administration Server • OLAP functions • Partitioned database server with local and remote clients • Partitioned database server with local and remote clients when federation is enabled. • Satellite database server with local clients
Parameter type	Possible values are: <ul style="list-style-type: none"> • Configurable (the database manager must be restarted to have the changes take effect) • Configurable online (can be dynamically updated online without having to restart the database manager) • Informational (values are for your information only and cannot be updated)
Default [range]	If applicable, lists the default value and the possible ranges, including NULL values or automatic settings. If the range differs by platform, then the values are listed by platform or platform type, for example, 32-bit or 64-bit platforms. Note that in most cases the default value is not listed as part of the range.
Unit of measure	If applicable, lists the unit of measure. Possible values are: <ul style="list-style-type: none"> • Bytes • Counter • Megabytes per second • Milliseconds • Minutes • Pages (4 KB) • Percentage • Seconds
Valid values	If applicable, lists the valid value. This heading is mutually exclusive with the default [range] heading.
Examples	If applicable, lists examples.
Propagation class	If applicable, possible values are: <ul style="list-style-type: none"> • Immediate • Statement boundary
When allocated	If applicable, indicates when the configuration parameter is allocated by the database manager.
When freed	If applicable, indicates when the configuration parameter is freed by the database manager.
Restrictions	If applicable, lists any restrictions that apply to the configuration parameter.
Limitations	If applicable, lists any limitations that apply to the configuration parameter.
Recommendations	If applicable, lists any recommendations that apply to the configuration parameter.
Usage notes	If applicable, lists any usage notes that apply to the configuration parameter.

Configuration parameters that affect the number of agents

There are a number of database manager configuration parameters related to database agents and how they are managed.

The following database manager configuration parameters determine how many database agents are created and how they are managed:

- Agent Pool Size (*num_poolagents*): The total number of idle agents to pool that are kept available in the system. The default value for this parameter is 100, AUTOMATIC.
- Initial Number of Agents in Pool (*num_initagents*): When the database manager is started, a pool of worker agents is created based on this value. This speeds up performance for initial queries. The worker agents all begin as idle agents.
- Maximum Number of Connections (*max_connections*): specifies the maximum number of connections allowed to the database manager system on each database partition.
- Maximum Number of Coordinating Agents (*max_coordagents*): For partitioned database environments and environments with intra-partition parallelism enabled when **Connection concentrator** is enabled, this value limits the number of coordinating agents.

Configuration parameters that affect query optimization

Several configuration parameters affect the access plan chosen by the SQL or XQuery compiler. Many of these are appropriate to a single-partition database environment and some are only appropriate to a partitioned database environment. Assuming a homogeneous partitioned database environment, where the hardware is the same, the values used for each parameter should be the same on all database partitions.

Note: When you change a configuration parameter dynamically, the optimizer might not read the changed parameter values immediately because of older access plans in the package cache. To reset the package cache, execute the FLUSH PACKAGE CACHE command.

In a federated system, if the majority of your queries access nicknames, evaluate the types of queries that you send before you change your environment. For example, in a federated database the buffer pool does not cache pages from data sources, which are the DBMSs and data within the federated system. For this reason, increasing the size of the buffer does not guarantee that the optimizer will consider additional access-plan alternatives when it chooses an access plan for queries that contain nicknames. However, the optimizer might decide that local materialization of data source tables is the least-cost route or a necessary step for a sort operation. In that case, increasing the resources available might improve performance.

The following configuration parameters or factors affect the access plan chosen by the SQL or XQuery compiler:

- The size of the buffer pools that you specified when you created or altered them. When the optimizer chooses the access plan, it considers the I/O cost of fetching pages from disk to the buffer pool and estimates the number of I/Os required to satisfy a query. The estimate includes a prediction of buffer-pool usage, because additional physical I/Os are not required to read rows in a page that is already in the buffer pool.

The optimizer considers the value of the *npages* column in the SYSCAT.BUFFERPOOLS system catalog tables and, in partitioned database environments, the SYSCAT.BUFFERPOOLDBPARTITIONS system catalog tables.

The I/O costs of reading the tables can have an impact on:

- How two tables are joined

- Whether an unclustered index will be used to read the data
- Default Degree (*dft_degree*)

The *dft_degree* configuration parameter specifies parallelism by providing a default value for the CURRENT DEGREE special register and the DEGREE bind option. A value of one (1) means no intra-partition parallelism. A value of minus one (-1) means the optimizer determines the degree of intra-partition parallelism based on the number of processors and the type of query.

Note: Intra-parallel processing does not occur unless you enable it by setting the *intra_parallel* database manager configuration parameter.
- Default Query Optimization Class (*dft_queryopt*)

Although you can specify a query optimization class when you compile SQL or XQuery queries, you can also set a default query optimization class.
- Average Number of Active Applications (*avg_appls*)

The optimizer uses the *avg_appls* parameter to help estimate how much of the buffer pool might be available at run-time for the access plan chosen. Higher values for this parameter can influence the optimizer to choose access plans that are more conservative in buffer pool usage. If you specify a value of 1, the optimizer considers that the entire buffer pool will be available to the application.
- Sort Heap Size (*sortheap*)

If the rows to be sorted occupy more than the space available in the sort heap, several sort passes are performed, where each pass sorts a subset of the entire set of rows. Each sort pass is stored in a system temporary table in the buffer pool, which might be written to disk. When all the sort passes are complete, these sorted subsets are merged into a single sorted set of rows. A sort is considered to be “piped” if it does not require a system temporary table to store the final, sorted list of data. That is, the results of the sort can be read in a single, sequential access. Piped sorts result in better performance than non-piped sorts and will be used if possible.

When choosing an access plan, the optimizer estimates the cost of the sort operations, including evaluating whether a sort can be piped, by:

 - Estimating the amount of data to be sorted
 - Looking at the *sortheap* parameter to determine if there is enough space for the sort to be piped.
- Maximum Storage for Lock List (*locklist*) and Maximum Percent of Lock List Before Escalation (*maxlocks*)

When the isolation level is **repeatable read (RR)**, the optimizer considers the values of the *locklist* and *maxlocks* parameters to determine whether row level locks might be escalated to a table level lock. If the optimizer estimates that lock escalation will occur for a table access, then it chooses a table level lock for the access plan, instead of incurring the overhead of lock escalation during the query execution.
- CPU Speed (*cpuspeed*)

The optimizer uses the CPU speed to estimate the cost of performing certain operations. CPU cost estimates and various I/O cost estimates help select the best access plan for a query.

The CPU speed of a machine can have a significant influence on the access plan chosen. This configuration parameter is automatically set to an appropriate value when the database is installed or migrated. Do not adjust this parameter unless you are modelling a production environment on a test system or assessing the impact of a hardware change. Using this parameter to model a different

hardware environment allows you to find out the access plans that might be chosen for that environment. To have the database manager recompute the value of this automatic configuration parameter, set it to -1.

- Statement Heap Size (*stmtheap*)

Although the size of the statement heap does not influence the optimizer in choosing different access paths, it can affect the amount of optimization performed for complex SQL or XQuery statements.

If the *stmtheap* parameter is not set large enough, you might receive a warning indicating that there is not enough memory available to process the statement. For example, SQLCODE +437 (SQLSTATE 01602) might indicate that the amount of optimization that has been used to compile a statement is less than the amount that you requested.

- Communications Bandwidth (*comm_bandwidth*)

Communications bandwidth is used by the optimizer to determine access paths. The optimizer uses the value in this parameter to estimate the cost of performing certain operations between the database partition servers in a partitioned database environment.

- Application Heap Size (*applheapsz*)

Large schemas require sufficient space in the application heap.

Restrictions and behavior when configuring *max_coordagents* and *max_connections*

The Version 9.5 default for the *max_coordagents* and *max_connections* parameters will be AUTOMATIC, with *max_coordagents* set to 200 and *max_connections* set to -1 (that is, set to the value of *max_coordagents*). These settings set Concentrator to OFF.

While configuring *max_coordagents* or *max_connections* online, there will be some restrictions and behavior to be aware of:

- If the value of *max_coordagents* is increased, the setting takes effect immediately and new requests will be allowed to create new coordinating agents. If the value is decreased, the number of coordinating agents will not be reduced immediately. Rather, the number of coordinating agents will no longer increase, and existing coordinating agents might terminate after finishing their current set of work, in order to reduce the overall number of coordinating agents. New requests for work that require a coordinating agent will not be serviced until the total number of coordinating agents falls below the new value and a coordinating agent becomes free.
- If the value for *max_connections* is increased, the setting takes effect immediately and new connections previously blocked because of this parameter will be allowed. If the value is decreased, the database manager will not actively terminate existing connections; instead, new connections will not be allowed until enough of the existing connections are terminated to bring the value down below the new maximum.
- If *max_connections* is set to -1 (default), then the maximum number of connections allowed is the same as *max_coordagents*, and when *max_coordagents* is updated offline or online; the maximum number of connections allowed will be updated as well.

While changing the value of *max_coordagents* or *max_connections* online, you cannot change it such that connection Concentrator will be turned either ON, if it's off, or OFF, if it's ON. For example, if at DB2START time *max_coordagents* is less than *max_connections* (Concentrator is ON), then all updates done online to these two

parameters must maintain the relationship $max_coordagents < max_connections$. Similarly, if at DB2START time, $max_coordagents$ is greater than or equal to $max_connections$ (Concentrator is OFF), then all updates done online must maintain this relationship.

When you perform this type of update online, the database manager does not fail the operation, instead it defers the update. The warning SQL1362W message is returned, similar to any case when updating the database manager configuration parameters where IMMEDIATE is specified, but is not possible.

When setting $max_coordagents$ or $max_connections$ to AUTOMATIC, the following behavior can be expected:

- Both of these parameters can be configured with a starting value and an AUTOMATIC setting. For example, the following command associates a value of 200 and AUTOMATIC to the $max_coordagents$ parameter:

```
UPDATE DBM CONFIG USING max_coordagents 200 AUTOMATIC
```

These parameters will always have a value associated with them, either the value set as default, or some value that you specified. If only AUTOMATIC is specified when updating either parameter, that is, no value is specified, and the parameter previously had a value associated with it, that value would remain. Only the AUTOMATIC setting would be affected.

Note: When Concentrator is ON, the values assigned to these two configuration parameters are important even when the parameters are set to AUTOMATIC.

- If both parameters are set to AUTOMATIC, the database manager allows the number of connections and coordinating agents to increase as needed to suit the workload. However, the following caveats apply:
 1. When Concentrator is OFF, the database manager maintains a one-to-one ratio: for every connection there will be only *one* coordinating agent.
 2. When Concentrator is ON, the database manager tries to maintain the ratio of coordinating agents to connections set by the values in the parameters.

Note:

- The approach used to maintain the ratio is designed to be unintrusive and does not guarantee the ratio will be maintained perfectly. New connections are always allowed in this scenario, though they may have to wait for an available coordinating agent. New coordinating agents will be created as needed to maintain the ratio. As connections are terminated, the database manager might also terminate coordinating agents to maintain the ratio
- The database manager will not reduce the ratio that you set. The initial values of $max_coordagents$ and $max_connections$ that you set are considered a lower bound.
- The current and delayed values of both these parameters can be displayed through various means, such as CLP or APIs. The values displayed will always be the values set by the user. For example, if the following command were issued, and then 30 concurrent connections performing work on the instance were started, the displayed values for $max_connections$ and $max_coordagents$ will still be 20, AUTOMATIC:

```
UPDATE DBM CFG USING max_connections 20 AUTOMATIC,  
max_coordagents 20 AUTOMATIC
```

To determine the real number of connections and coordinating agents currently running monitor elements, you can also use the Health Monitor.

- If *max_connections* is set to AUTOMATIC with a value greater than *max_coordagents* (so that Concentrator is ON), and *max_coordagents* is not set to AUTOMATIC, then the database manager allows an unlimited number of connections that will use only a limited number of coordinating agents.

Note: Connections might have to wait for available coordinating agents.

The use of the AUTOMATIC option for the *max_coordagents* and *max_connections* configuration parameters is only valid in the following two scenarios:".

1. Both parameters are set to AUTOMATIC
2. Concentrator is enabled with *max_connections* set to AUTOMATIC, while *max_coordagents* is not.

All other configurations using AUTOMATIC for these parameters will be blocked and will return SQL6112N, with a reason code that explains the valid settings of AUTOMATIC for these two parameters.

Database Manager configuration parameters

agent_stack_sz - Agent stack size

This parameter determines the virtual memory that is allocated by DB2 for each agent.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

Linux (32-bit)

256 [16 – 1024]

Linux (64-bit) and UNIX

1024 [256 – 32768]

Windows

16 [8 – 1000]

Unit of measure

Pages (4 KB)

When allocated

When an agent is initialized to do work for an application

When freed

When an agent completes the work to be done for an application

You can use this parameter to optimize memory utilization of the server for a given set of applications. More complex queries will use more stack space, compared to the space used for simple queries.

This parameter is used to set the initial committed stack size for each agent in a Windows environment. By default, each agent stack can grow up to the default reserve stack size of 256 KB (64 4-KB pages). This limit is sufficient for most database operations. On UNIX and Linux, *agent_stack_sz* will be rounded up to the next larger power-of-2 based value. The default setting for UNIX should be sufficient for most workloads

However, when preparing a large SQL or XQuery statement, the agent can run out of stack space and the system will generate a stack overflow exception (0xC00000FD). When this happens, the server will shut down because the error is non-recoverable.

Note: In Version 9.5 and later, sqlcode -973 will be returned instead of a stack overflow exception..

The agent stack size can be increased by setting *agent_stack_sz* to a value larger than the default reserve stack size of 64 pages. Note that the value for *agent_stack_sz*, when larger than the default reserve stack size, is rounded by the Windows operating system to the nearest multiple of 1 MB; setting the agent stack size to 128 4-KB pages actually reserves a 1 MB stack for each agent. Setting the value for *agent_stack_sz* less than the default reserve stack size will have no effect on the maximum limit because the stack still grows if necessary up to the default reserve stack size. In this case, the value for *agent_stack_sz* is the initial committed memory for the stack when an agent is created.

You can change the default reserve stack size by using the db2hdr utility to change the header information for the db2syscs.exe file. Changing the default reserve stack size will affect all threads while changing *agent_stack_sz* only affects the stack size for agents. The advantage of changing the default stack size using the db2hdr utility is that it provides a better granularity, therefore allowing the stack size to be set at the minimum required stack size. However, you will have to stop and restart DB2 for a change to db2syscs.exe to take effect.

Recommendation: If you will be working with large or complex XML data in a 32-bit environment, you should update *agent_stack_sz* to at least 256 4-KB pages. Very complex XML schemas might require *agent_stack_sz* to be set much closer to the limit in order to avoid stack overflow exceptions during schema registration or during XML document validation.

You might be able to reduce the stack size in order to make more address space available to other clients, if your environment matches the following:

- Contains only simple applications (for example light OLTP), in which there are never complex queries
- Requires a relatively large number of concurrent clients (for example, more than 100).

On Windows, the agent stack size and the number of concurrent clients are inversely related: a larger stack size reduces the potential number of concurrent clients that can be running. This occurs because address space is limited on Windows platforms.

agentpri - Priority of agents

This parameter is deprecated in Version 9.5, but is still being used by pre-Version 9.5 data servers and clients. Any value specified for this configuration parameter will continue to work exactly as it did in previous versions, and this parameter

will continue to be fully supported. If this parameter is used for workload management (WLM), then the WLM service class agent priority will be ignored.

Note: The following information applies only to pre-Version 9.5 data servers and clients.

This parameter controls the priority given both to all agents, and to other database manager instance processes and threads, by the operating system scheduler. This priority determines how CPU time is given to the database manager processes, agents, and threads relative to the other processes and threads running on the machine.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

AIX -1 (system) [41 - 125]

Other UNIX

-1 (system) [41 - 128]

Windows

-1 (system) [0 - 6]

Solaris

-1 (system) [0 - 59]

When the parameter is set to -1 or system, no special action is taken and the database manager is scheduled in the normal way that the operating system schedules all processes and threads. When the parameter is set to a value other than -1 or system, the database manager will create its processes and threads with a static priority set to the value of the parameter. Therefore, this parameter allows you to control the priority with which the database manager processes and threads (in a partitioned database environment, this also includes coordinating and subagents, the parallel system controllers, and the FCM daemons) will execute on your machine.

You can use this parameter to increase database manager throughput. The values for setting this parameter are dependent on the operating system on which the database manager is running. For example, in a Linux or UNIX environment, numerically low values yield high priorities. When the parameter is set to a value between 41 and 125, the database manager creates its agents with a UNIX static priority set to the value of the parameter. This is important in Linux and UNIX environments because numerically low values yield high priorities for the database manager, but other processes (including applications and users) might experience delays because they cannot obtain enough CPU time. You should balance the setting of this parameter with the other activity expected on the machine.

Restrictions:

- If you set this parameter to a non-default value on Linux and UNIX platforms, you cannot use the governor to alter agent priorities.
- On the Solaris operating system, you should not change the default value (-1). Changing the default value sets the priority of DB2 processes to real-time, which can monopolize all available resources on the system.

Recommendation: The default value should be used initially. This value provides a good compromise between response time to other users/applications and database manager throughput.

If database performance is a concern, you can use benchmarking techniques to determine the optimum setting for this parameter. You should take care when increasing the priority of the database manager because performance of other user processes can be severely degraded, especially when the CPU utilization is very high. Increasing the priority of the database manager processes and threads can have significant performance benefits.

aslheapsz - Application support layer heap size

The application support layer heap represents a communication buffer between the local application and its associated agent. This buffer is allocated as shared memory by each database manager agent that is started.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

15 [1 - 524 288]

Unit of measure

Pages (4 KB)

When allocated

When the database manager agent process is started for the local application

When freed

When the database manager agent process is terminated

If the request to the database manager, or its associated reply, do not fit into the buffer they will be split into two or more send-and-receive pairs. The size of this buffer should be set to handle the majority of requests using a single send-and-receive pair. The size of the request is based on the storage required to hold:

- The input SQLDA
- All of the associated data in the SQLVARs
- The output SQLDA
- Other fields which do not generally exceed 250 bytes.

In addition to this communication buffer, this parameter is also used for two other purposes:

- It is used to determine the I/O block size when a blocking cursor is opened. This memory for blocked cursors is allocated out of the application's private address space, so you should determine the optimal amount of private memory to allocate for each application program. If the Data Server Runtime Client cannot allocate space for a blocking cursor out of an application's private memory, a non-blocking cursor will be opened.
- It is used to determine the communication size between agents and db2fmp processes. (A db2fmp process can be a user-defined function or a fenced stored procedure.) The number of bytes is allocated from shared memory for each db2fmp process or thread that is active on the system.

The data sent from the local application is received by the database manager into a set of contiguous memory allocated from the query heap. The *aslheapsz* parameter is used to determine the initial size of the query heap (for both local and remote clients). The maximum size of the query heap is defined by the *query_heap_sz* parameter.

Recommendation: If your application's requests are generally small and the application is running on a memory constrained system, you might want to reduce the value of this parameter. If your queries are generally very large, requiring more than one send and receive request, and your system is not constrained by memory, you might want to increase the value of this parameter.

Use the following formula to calculate a minimum number of pages for *aslheapsz*:

```
aslheapsz >= ( sizeof(input SQLDA)
               + sizeof(each input SQLVAR)
               + sizeof(output SQLDA)
               + 250 ) / 4096
```

where `sizeof(x)` is the size of `x` in bytes that calculates the number of pages of a given input or output value.

You should also consider the effect of this parameter on the number and potential size of blocking cursors. Large row blocks might yield better performance if the number or size of rows being transferred is large (for example, if the amount of data is greater than 4096 bytes). However, there is a trade-off in that larger record blocks increase the size of the working set memory for each connection.

Larger record blocks might also cause more fetch requests than are actually required by the application. You can control the number of fetch requests using the OPTIMIZE FOR clause on the SELECT statement in your application.

audit_buf_sz - Audit buffer size

This parameter specifies the size of the buffer used when auditing the database.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type
Configurable

Default [range]
0 [0 - 65 000]

Unit of measure
Pages (4 KB)

When allocated
When DB2 is started

When freed
When DB2 is stopped

The default value for this parameter is zero (0). If the value is zero (0), the audit buffer is not used. If the value is greater than zero (0), space is allocated for the audit buffer where the audit records will be placed when they are generated by the audit facility. The value times 4 KB pages is the amount of space allocated for the audit buffer. The audit buffer cannot be allocated dynamically; DB2 must be stopped and then restarted before the new value for this parameter takes effect.

By changing this parameter from the default to some value larger than zero (0), the audit facility writes records to disk asynchronously compared to the execution of the statements generating the audit records. This improves DB2 performance over leaving the parameter value at zero (0). The value of zero (0) means the audit facility writes records to disk synchronously with (at the same time as) the execution of the statements generating the audit records. The synchronous operation during auditing decreases the performance of applications running in DB2.

authentication - Authentication type

Specifies and determines how and where authentication of a user takes place.

Configuration type
Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type
Configurable

Default [range]
SERVER [CLIENT; SERVER; SERVER_ENCRYPT; DATA_ENCRYPT;
DATA_ENCRYPT_CMP; KERBEROS; KRB_SERVER_ENCRYPT;
GSSPLUGIN; GSS_SERVER_ENCRYPT]

If **authentication** is SERVER, the user ID and password are sent from the client to the server so that authentication can take place on the server. The value SERVER_ENCRYPT provides the same behavior as SERVER, except that any user IDs and passwords sent over the network are encrypted.

A value of `DATA_ENCRYPT` means the server accepts encrypted `SERVER` authentication schemes and the encryption of user data. The authentication works exactly the same way as `SERVER_ENCRYPT`.

The following user data are encrypted when using this authentication type:

- SQL statements
- SQL program variable data
- Output data from the server processing an SQL statement and including a description of the data
- Some or all of the answer set data resulting from a query
- Large object (LOB) streaming
- SQLDA descriptors

A value of `DATA_ENCRYPT_CMP` means the server accepts encrypted `SERVER` authentication schemes and the encryption of user data. In addition, this authentication type allows compatibility with earlier products that do not support `DATA_ENCRYPT` authentication type. These products are permitted to connect with the `SERVER_ENCRYPT` authentication type and without encrypting user data. Products supporting the new authentication type must use it. This authentication type is only valid in the server's database manager configuration file and is not valid when used on the `CATALOG DATABASE` command.

Note: For a standards compliance (defined in the "Standards compliance" topic) configuration, `SERVER` is the only supported value.

A value of `CLIENT` indicates that all authentication takes place at the client. No authentication needs to be performed at the server.

A value of `KERBEROS` means that authentication is performed at a Kerberos server using the Kerberos security protocol for authentication. With an authentication type of `KRB_SERVER_ENCRYPT` at the server and clients that support the Kerberos security system, the effective system authentication type is `KERBEROS`. If the clients do not support the Kerberos security system, the system authentication type is effectively equivalent to `SERVER_ENCRYPT`.

A value of `GSSPLUGIN` means that authentication is performed using an external GSSAPI-based security mechanism. With an authentication type of `GSS_SERVER_ENCRYPT` at the server and clients that support the `GSSPLUGIN` security mechanism, the effective system authentication type is `GSSPLUGIN` (that is, if the clients support one of the server's plug-ins). If the clients do not support the `GSSPLUGIN` security mechanism, the system authentication type is effectively equivalent to `SERVER_ENCRYPT`.

Recommendation: Typically, the default value (`SERVER`) is adequate for local clients. If remote clients are connecting to the database server then `SERVER_ENCRYPT` is the suggested value to protect the user ID and password.

catalog_noauth - Cataloging allowed without authority

This parameter specifies whether users are able to catalog and uncatalog databases and nodes, or DCS and ODBC directories, without `SYSADM` authority.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type
Configurable Online

Propagation class
Immediate

Default [range]

Database server with local and remote clients
NO [NO (0) — YES (1)]

Client; Database server with local clients
YES [NO (0) — YES (1)]

The default value (0) for this parameter indicates that SYSADM authority is required. When this parameter is set to 1 (yes), SYSADM authority is not required.

clnt_krb_plugin - Client Kerberos plug-in

This parameter specifies the name of the default Kerberos plug-in library to be used for client-side authentication and local authorization.

Configuration type
Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type
Configurable

Default [range]
Null or IBMkrb5 [any valid string]

By default, the value is null on Linux and UNIX systems, and IBMkrb5 on Windows operating systems. The plug-in is used when the client is authenticated using KERBEROS authentication, or when local authorization is performed and the authentication type in the DBM CFG is KERBEROS.

clnt_pw_plugin - Client userid-password plug-in

This parameter specifies the name of the userid-password plug-in library to be used for client-side authentication and local authorization.

Configuration type
Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type
Configurable

Default [range]
Null [any valid string]

By default, the value is null and the DB2-supplied userid-password plug-in library is used. The plug-in is used when the client is authenticated using CLIENT authentication, or when local authorization is performed and the authentication type in the DBM CFG is CLIENT, SERVER, SERVER_ENCRYPT or DATA_ENCRYPT. For non-root installations, if the DB2 userid and password plug-in library is used, the db2rfe command must be run before using your DB2 product.

cluster_mgr - Cluster manager name

This parameter enables the database manager to communicate incremental cluster configuration changes to the specified cluster manager.

Configuration type
Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Multi-partitioned database server with local and remote clients

Parameter type
Informational

Default
No default

Valid values

- TSA

This parameter is set during high availability cluster configuration using the DB2 High Availability Instance Configuration Utility (db2haicu).

comm_bandwidth - Communications bandwidth

This parameter helps the query optimizer determine access paths by indicating the bandwidth between database partition servers.

Configuration type
Database manager

Applies to
Partitioned database server with local and remote clients

Parameter type
Configurable Online

Propagation class
Statement boundary

Default [range]
-1 [.1 - 100 000]

A value of -1 causes the parameter value to be reset to the default. The default value is calculated based on the speed of the underlying communications adapter. A value of 100 can be expected for systems using Gigabit Ethernet.

Unit of measure

Megabytes per second

The value calculated for the communications bandwidth, in megabytes per second, is used by the query optimizer to estimate the cost of performing certain operations between the database partition servers of a partitioned database system. The optimizer does not model the cost of communications between a client and a server, so this parameter should reflect only the nominal bandwidth between the database partition servers, if any.

You can explicitly set this value to model a production environment on your test system or to assess the impact of upgrading hardware.

Recommendation: You should only adjust this parameter if you want to model a different environment.

The communications bandwidth is used by the optimizer in determining access paths. You should consider rebinding applications (using the REBIND PACKAGE command) after changing this parameter.

conn_elapse - Connection elapse time

This parameter specifies the number of seconds within which a TCP/IP connection is to be established between two database partition servers.

Configuration type

Database manager

Applies to

Partitioned database server with local and remote clients

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

10 [0–100]

Unit of measure

Seconds

If the attempt to connect succeeds within the time specified by this parameter, communications are established. If it fails, another attempt is made to establish communications. If the connection is attempted the number of times specified by the *max_connretries* parameter and always times out, an error is issued.

cpuspeed - CPU speed

This parameter reflects the CPU speed of the machine(s) the database is installed on.

When the DB2 instance is created, a small unit of work is run multiple times against the operating system. CPU speed is the average time taken to complete this unit of work.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable online

Propagation class

Statement boundary

Default [range]

-1 [1×10^{-10} — 1] A value of -1 will cause the parameter value to be reset based on the running of the measurement program.

Unit of measure

Milliseconds

This program is executed if benchmark results are not available if the data for the IBM RS/6000 model 530H is not found in the file, or if the data for your machine is not found in the file.

You can explicitly set this value to model a production environment on your test system or to assess the impact of upgrading hardware. By setting it to -1, *cpuspeed* will be re-computed.

Recommendation: You should only adjust this parameter if you want to model a different environment.

The CPU speed is used by the optimizer in determining access paths. You should consider rebinding applications (using the REBIND PACKAGE command) after changing this parameter.

dft_account_str - Default charge-back account

This parameter acts as the default suffix of accounting identifiers.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

Null [any valid string]

With each application connect request, an accounting identifier consisting of a DB2 Connect-generated prefix and the user supplied suffix is sent from the application requester to a DRDA application server. This accounting information provides a mechanism for system administrators to associate resource usage with each user access.

Note: This parameter is only applicable to DB2 Connect.

The suffix is supplied by the application program calling the `sqlsact()` API or the user setting the environment variable `DB2ACCOUNT`. If a suffix is not supplied by either the API or environment variable, DB2 Connect uses the value of this parameter as the default suffix value. This parameter is particularly useful for earlier database clients (anything prior to version 2) that do not have the capability to forward an accounting string to DB2 Connect.

Recommendation: Set this accounting string using the following:

- Alphabets (A through Z)
- Numerics (0 through 9)
- Underscore (_).

dft_monswitches - Default database system monitor switches

This parameter allows you to set a number of switches which are each internally represented by a bit of the parameter.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Propagation class

Immediate

Note: The change takes effect immediately if you explicitly `ATTACH` to the instance before modifying the `dft_mon_xxxx` switch settings. Otherwise the setting takes effect the next time the instance is restarted.

Default

All switches turned off, except `dft_mon_timestamp`, which is turned on by default

The parameter is unique in that you can update each of these switches independently by setting the following parameters:

dft_mon_uow

Default value of the snapshot monitor's unit of work (UOW) switch

dft_mon_stmt

Default value of the snapshot monitor's statement switch

dft_mon_table

Default value of the snapshot monitor's table switch

dft_mon_bufpool

Default value of the snapshot monitor's buffer pool switch

dft_mon_lock

Default value of the snapshot monitor's lock switch

dft_mon_sort

Default value of the snapshot monitor's sort switch

dft_mon_timestamp

Default value of the snapshot monitor's timestamp switch

Recommendation: Any switch (except `dft_mon_timestamp`) that is turned ON instructs the database manager to collect monitor data related to that switch. Collecting additional monitor data increases database manager overhead which can impact system performance. Turning the `dft_mon_timestamp` switch OFF becomes important as CPU utilization approaches 100%. When this occurs, the CPU time required for issuing timestamps increases dramatically. Furthermore, if the timestamp switch is turned OFF, the overall cost of other data under monitor switch control is greatly reduced.

All monitoring applications inherit these default switch settings when the application issues its first monitoring request (for example, setting a switch, activating the event monitor, taking a snapshot). You should turn on a switch in the configuration file only if you want to collect data starting from the moment the database manager is started. (Otherwise, each monitoring application can set its own switches and the data it collects becomes relative to the time its switches are set.)

dftdbpath - Default database path

This parameter contains the default file path used to create databases under the database manager. If no path is specified when a database is created, the database is created under the path specified by the `dftdbpath` parameter.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

UNIX Home directory of instance owner [any existing path]

Windows

Drive on which DB2 is installed [any existing path]

In a partitioned database environment, you should ensure that the path on which the database is being created is not an NFS-mounted path (on Linux and UNIX

platforms), or a network drive (in a Windows environment). The specified path must physically exist on each database partition server. To avoid confusion, it is best to specify a path that is locally mounted on each database partition server. The maximum length of the path is 205 characters. The system appends the database partition name to the end of the path.

Given that databases can grow to a large size and that many users could be creating databases (depending on your environment and intentions), it is often convenient to be able to have all databases created and stored in a specified location. It is also good to be able to isolate databases from other applications and data both for integrity reasons and for ease of backup and recovery.

For Linux and UNIX environments, the length of the *dftdbpath* name cannot exceed 215 characters and must be a valid, absolute, path name. For Windows, the *dftdbpath* can be a drive letter, optionally followed by a colon.

Recommendation: If possible, put high volume databases on a different disk than other frequently accessed data, such as the operating system files and the database logs.

diaglevel - Diagnostic error capture level

This parameter specifies the type of diagnostic errors that will be recorded in the db2diag.log file.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

3 [0 — 4]

Valid values for this parameter are:

- 0 – No diagnostic data captured
- 1 – Severe errors only
- 2 – All errors
- 3 – All errors and warnings
- 4 – All errors, warnings and informational messages

The *diagpath* configuration parameter is used to specify the directory that will contain the error file, alert log file, and any dump files that might be generated, based on the value of the *diaglevel* parameter.

Recommendation: You might want to increase the value of this parameter to gather additional problem determination data to help resolve a problem.

diagpath - Diagnostic data directory path

This parameter allows you to specify the fully qualified path for DB2 diagnostic information.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable online

Propagation class

Immediate

Default [range]

Null [any valid path name]

This directory could possibly contain dump files, trap files, an error log, a notification file, an alert log file, and first occurrence data collection (FODC) packages, depending on your platform.

If this parameter is null, the diagnostic information will be written to files in one of the following directories or folders:

- In Windows environments:
 - User data files, for example, files under instance directories, are written to a location that is different from where the code is installed, as follows:
 - In Windows Vista environments, user data files are written to ProgramData\IBM\DB2\.
 - In Windows 2003 and XP environments, user data files are written to Documents and Settings\All Users\Application Data\IBM\DB2*Copy Name*, where *Copy Name* is the name of your DB2 copy.
- In Linux and UNIX environments: Information is written to *INSTHOME*/sqlib/db2dump, where *INSTHOME* is the home directory of the instance.

In Version 9.5, the default value of **DB2INSTPROF** at the global level is stored at the new location shown above. Other profile registry variables that specify the location of the runtime data files should query the value of **DB2INSTPROF**. The other variables include the following ones:

- **DB2CLIINIPATH**
- **DIAGPATH**
- **SPM_LOG_PATH**

Recommendation: Use the default setting for the *diagpath* configuration parameter or use a centralized location for the *diagpath* value of multiple instances.

In a partitioned database environment, the *diagpath* parameter should use local storage at the host to get the best performance from logging. This creates a separate logging and diagnostic directory for each physical partition. You can use

the PD_GET_DIAG_HIST table function to retrieve the log records from the different partitions, and the PD_GET_LOG_MSGS table function to retrieve the notification log from all partitions.

dir_cache - Directory cache support

This parameter determines whether the database, node and DCS directory files will be cached in memory.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

Yes [Yes; No]

When allocated

- When an application issues its first connect, the application directory cache is allocated
- When a database manager instance is started (db2start), the server directory cache is allocated.

When freed

- When an the application process terminates, the application directory cache is freed
- When a database manager instance is stopped (db2stop), the server directory cache is freed.

The use of the directory cache reduces connect costs by eliminating directory file I/O and minimizing the directory searches required to retrieve directory information. There are two types of directory caches:

- An application directory cache that is allocated and used for each application process on the machine at which the application is running.
- A server directory cache that is allocated and used for some of the internal database manager processes.

For application directory caches, when an application issues its first connect, each directory file is read and the information is cached in private memory for this application. The cache is used by the application process on subsequent connect requests and is maintained for the life of the application process. If a database is not found in the application directory cache, the directory files are searched for the information, but the cache is not updated. If the application modifies a directory entry, the next connect within that application will cause the cache for this application to be refreshed. The application directory cache for other applications will not be refreshed. When the application process terminates, the cache is freed. (To refresh the directory cache used by a command line processor session, issue a db2 terminate command.)

For server directory caches, when a database manager instance is started (db2start), each directory file is read and the information is cached in the server memory. This cache is maintained until the instance is stopped (db2stop). If a directory entry is not found in this cache, the directory files are searched for the information. This server directory cache is never refreshed during the time the instance is running.

Recommendation: Use directory caching if your directory files do not change frequently and performance is critical.

In addition, on remote clients, directory caching can be beneficial if your applications issue several different connection requests. In this case, caching reduces the number of times a single application must read the directory files.

Directory caching can also improve the performance of taking database system monitor snapshots. In addition, you should explicitly reference the database name on the snapshot call, instead of using database aliases.

Note: Errors might occur when performing snapshot calls if directory caching is turned on and if databases are cataloged, uncataloged, created, or dropped after the database manager is started.

discover - Discovery mode

This parameter determines what kind of discovery requests, if any, the client can make.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

SEARCH [DISABLE, KNOWN, SEARCH]

From a client perspective, one of the following will occur:

- If *discover* = SEARCH, the client can issue search discovery requests to find DB2 server systems on the network. Search discovery provides a superset of the functionality provided by KNOWN discovery. If *discover* = SEARCH, both search and known discovery requests can be issued by the client.
- If *discover* = KNOWN, only known discovery requests can be issued from the client. By specifying some connection information for the administration server on a particular system, all the instance and database information on the DB2 system is returned to the client.
- If *discover* = DISABLE, discovery is disabled at the client.

The default discovery mode is SEARCH.

discover_inst - Discover server instance

This parameter specifies whether this instance can be detected by DB2 discovery.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable online

Propagation class

Immediate

Default [range]

ENABLE [ENABLE, DISABLE]

The parameter's default, enable, specifies that the instance can be detected, while disable prevents the instance from being discovered.

fcm_num_buffers - Number of FCM buffers

This parameter specifies the number of 4 KB buffers that are used for internal communications (messages) both among and within database servers.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

32-bit platforms

Automatic [128 - 65 300]

64-bit platforms

Automatic [128 - 524 288]

- Database server with local and remote clients: the default is 1024
- Database server with local clients: the default is 512
- Partitioned database server with local and remote clients: the default is 4096

On single-partition database systems, this parameter is used only if intra-partition parallelism is enabled by changing the *intra_parallel* parameter from its default value of NO to YES.

It is possible to set both an initial value and the AUTOMATIC attribute.

When set to AUTOMATIC, FCM monitors resource usage and can either increase or decrease resources, if they are not used within 30 minutes. The amount by which resources can be increased or decreased depends on the platform, in particular, that on Linux it can only be increased 25% above the starting value. If the database manager cannot allocate the number of resources specified when an instance is started, it scales back the configuration values incrementally until it can start the instance.

If you have multiple logical nodes on the same machine, you might find it necessary to increase the value of this parameter. You might also find it necessary to increase the value of this parameter if you run out of message buffers because of the number of users on the system, the number of database partition servers on the system, or the complexity of the applications.

If you are using multiple logical nodes, one pool of *fcm_num_buffers* buffers is shared by all the multiple logical nodes on the same machine. The size of the pool will be determined by multiplying the *fcm_num_buffers* value times the number of logical nodes on that physical machine. Re-examine the value you are using; consider how many FCM buffers in total will be allocated on the machine (or machines) where the multiple logical nodes reside.

fcm_num_channels - Number of FCM channels

This parameter specifies the number of FCM channels for each database partition.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients
- Satellite database server with local clients

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

UNIX 32-bit platforms

Automatic, with starting values of 256, 512, 2 048 [128 - 120 000]

UNIX 64-bit platforms

Automatic, with starting values of 256, 512, 2 048 [128 - 524 288]

Windows 32-bit

Automatic, with a starting value 10 000 [128 - 120 000]

Windows 64-bit

Automatic, with starting values of 256, 512, 2 048 [128 - 524 288]

- For database server with local and remote clients, the starting value is 512.
- For database server with local clients, the starting value is 256.
- For partitioned database environment servers with local and remote clients, the starting value is 2 048.

On non-partitioned database environments, the *intra_parallel* parameter must be active before *fcf_num_channels* can be used.

An FCM channel represents a logical communication end point between EDUs running in the DB2 engine. Both control flows (request and reply) and data flows (table queue data) rely on channels to transfer data between partitions.

When set to AUTOMATIC, FCM monitors channel usage, incrementally allocating and releasing resources as requirements change.

fed_noauth - Bypass federated authentication

This parameter determines whether federated authentication will be bypassed at the instance.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable online

Propagation class

Immediate

Default [range]

No [Yes; No]

When *fed_noauth* is set to *yes*, *authentication* is set to *server* or *server_encrypt*, and *federated* is set to *yes*, then authentication at the instance is bypassed. It is assumed that authentication will happen at the data source. Exercise caution when *fed_noauth* is set to *yes*. Authentication is done at neither the client nor at DB2. Any user who knows the SYSADM authentication name can assume SYSADM authority for the federated server.

federated - Federated database system support

This parameter enables or disables support for applications submitting distributed requests for data managed by data sources (such as the DB2 Family and Oracle).

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

No [Yes; No]

federated_async - Maximum asynchronous TQs per query configuration parameter

This parameter determines the maximum number of asynchrony table queues (ATQs) in the access plan that the federated server supports.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients when federation is enabled.

Parameter type

Configurable online

Default [range]

0 [0 to 32 767 inclusive, -1, ANY]

When ANY or -1 is specified, the optimizer determines the number of ATQs for the access plan. The optimizer assigns an ATQ to all eligible SHIP or remote pushdown operators in the plan. The value that is specified for DB2_MAX_ASYNC_REQUESTS_PER_QUERY server option limits the number of asynchronous requests.

Recommendation

The *federated_async* configuration parameter supplies the default or starting value for the special register and the bind option. You can override the value of this parameter by setting the value of the CURRENT FEDERATED ASYNCHRONY special register, FEDERATED_ASYNC bind option, or FEDERATED_ASYNC precompile option to a higher or a lower number.

If the special register or the bind option do not override the *federated_async* configuration parameter, the value of the parameter determines the maximum number of ATQs in the access plan that the federated server allows. If the special register or the bind option overrides this parameter, the value of the special register or the bind option determines the maximum number of ATQs in the plan.

Any changes to the *federated_async* configuration parameter affect dynamic statements as soon as the current unit of work commits. Subsequent dynamic statements recognize the new value automatically. A restart of the federated database is not needed. Embedded SQL packages are not invalidated nor implicitly rebound when the value of the *federated_async* configuration parameter changes.

If you want the new value of the *federated_async* configuration parameter to affect static SQL statements, you need to rebind the package.

fenced_pool - Maximum number of fenced processes

This parameter represents the number of threads cached in each db2fmp process for threaded db2fmp processes (processes serving threadsafe stored procedures and UDFs). For nonthreaded db2fmp processes, this parameter represents the number of processes cached.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable online

Default [range]

-1 (*max_coordagents*), Automatic [-1; 0–64 000]

Unit of measure

Counter

Restrictions:

- If this parameter is updated dynamically, and the value is decreased, the database manager does not proactively terminate db2fmp threads or processes, instead it stops caching them as they are used in order to reduce the number of cached db2fmp's down to the new value.
- If this parameter is updated dynamically, and the value is increased, the database manager caches more db2fmp threads and processes when they are created.
- When this parameter is set to -1, the default, it assumes the value of the *max_coordagents* configuration parameter. Note that only the value of *max_coordagents* is assumed and not the automatic setting or behavior.
- When this parameter is set to AUTOMATIC, also the default:
 - The database manager allows the number of db2fmp threads and processes cached to increase based on the high water mark of coordinating agents. Specifically, the automatic behavior of this parameter allows it to grow depending on the maximum number of coordinating agents the database manager has ever registered, at the same time, since it started.
 - The value assigned to this parameter represents a lower bound for the number of db2fmp threads and process to cache.

Recommendation: If your environment uses fenced stored procedures or user defined functions, then this parameter can be used to ensure that an appropriate number of db2fmp processes are available to process the maximum number of concurrent stored procedures and UDFs that run on the instance, ensuring that no new fenced mode processes need to be created as part of stored procedure and UDF execution.

If you find that the default value is not appropriate for your environment because an inappropriate amount of system resource is being given to db2fmp processes and is affecting performance of the database manager, the following might be useful in providing a starting point for tuning this parameter:

```
fenced_pool = # of applications allowed to make stored procedure and
UDF calls at one time
```

If *keepfenced* is set to YES, then each db2fmp process that is created in the cache pool will continue to exist and will use system resources even after the fenced routine call has been processed and returned to the agent.

If *keepfenced* is set to NO, then nonthreaded db2fmp processes will terminate when they complete execution, and there is no cache pool. Multithreaded db2fmp processes will continue to exist, but no threads will be pooled in these processes.

This means that even when *keepfenced* is set to NO, you can have one threaded C db2fmp process and one threaded Java db2fmp process on your system.

In previous versions, this parameter was known as *maxdari*.

group_plugin - Group plug-in

This parameter specifies the name of the group plug-in library.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

Null [any valid string]

By default, this value is null, and DB2 uses the operating system group lookup. The plug-in will be used for all group lookups. For non-root installations, if the DB2 userid and password plug-in library is used, the db2rfe command must be run before using your DB2 product.

health_mon - Health monitoring

This parameter allows you to specify whether you want to monitor an instance, its associated databases, and database objects according to various health indicators.

Configuration type

Database manager

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

On [On; Off]

Related Parameters

If *health_mon* is turned on (the default), an agent will collect information about the health of the objects you have selected. If an object is considered to be in an unhealthy position, based on thresholds that you have set, notifications can be sent, and actions can be taken automatically. If *health_mon* is turned off, the health of objects will not be monitored.

You can use the Health Center or the CLP to select the instance and database objects that you want to monitor. You can also specify where notifications should be sent, and what actions should be taken, based on the data collected by the health monitor.

indexrec - Index re-creation time

This parameter indicates when the database manager will attempt to rebuild invalid indexes, and whether or not any index build will be redone during DB2 rollforward or HADR log replay on the standby database.

Configuration type

Database and Database Manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

UNIX Database Manager

restart [restart; restart_no_redo; access; access_no_redo]

Windows Database Manager

restart [restart; restart_no_redo; access; access_no_redo]

Database

Use system setting [system; restart; restart_no_redo; access; access_no_redo]

There are five possible settings for this parameter:

SYSTEM

use system setting specified in the database manager configuration file to decide when invalid indexes will be rebuilt, and whether any index build log records are to be redone during DB2 rollforward or HADR log replay. (Note: This setting is only valid for database configurations.)

ACCESS

Invalid indexes are rebuilt when the index is first accessed. Any fully logged index builds are redone during DB2 rollforward or HADR log replay. When HADR is started and an HADR takeover occurs, any invalid indexes are rebuilt after takeover when the underlying table is first accessed.

ACCESS_NO_REDO

Invalid indexes will be rebuilt when the underlying table is first accessed. Any fully logged index build will not be redone during DB2 rollforward or HADR log replay and those indexes will be left invalid. When HADR is started and an HADR takeover takes place, any invalid indexes will be rebuilt after takeover when the underlying table is first accessed.

RESTART

The default value for *indexrec*. Invalid indexes will be rebuilt when a RESTART DATABASE command is either explicitly or implicitly issued. Any fully logged index build will be redone during DB2 rollforward or HADR log replay. When HADR is started and an HADR takeover takes place, any invalid indexes will be rebuilt at the end of takeover.

Note that a RESTART DATABASE command is implicitly issued if the *autorestart* parameter is enabled.

RESTART_NO_REDO

Invalid indexes will be rebuilt when a RESTART DATABASE command is either explicitly or implicitly issued. (A RESTART DATABASE command is implicitly issued if the *autorestart* parameter is enabled.) Any fully logged index build will not be redone during DB2 rollforward or HADR log replay and instead those indexes will be rebuilt when rollforward completes or when HADR takeover takes place.

Indexes can become invalid when fatal disk problems occur. If this happens to the data itself, the data could be lost. However, if this happens to an index, the index can be recovered by re-creating it. If an index is rebuilt while users are connected to the database, two problems could occur:

- An unexpected degradation in response time might occur as the index file is re-created. Users accessing the table and using this particular index would wait while the index was being rebuilt.
- Unexpected locks might be held after index re-creation, especially if the user transaction that caused the index to be re-created never performed a COMMIT or ROLLBACK.

Recommendation: The best choice for this option on a high-user server and if restart time is not a concern, would be to have the index rebuilt at DATABASE RESTART time as part of the process of bringing the database back online after a crash.

Setting this parameter to “ACCESS” or to “ACCESS_NO_REDO” will result in a degradation of the performance of the database manager while the index is being re-created. Any user accessing that specific index or table would have to wait until the index is recreated.

If this parameter is set to “RESTART”, the time taken to restart the database will be longer due to index re-creation, but normal processing would not be impacted once the database has been brought back online.

Note: At database recovery time, all SQL procedure executables on the file system that belong to the database being recovered are removed. If *indexrec* is set to RESTART, all SQL procedure executables are extracted from the database catalog and put back on the file system at the next connection to the database. If *indexrec* is not set to RESTART, an SQL executable is extracted to the file system only on first execution of that SQL procedure.

The difference between the RESTART and the RESTART_NO_REDO values, or between the ACCESS and the ACCESS_NO_REDO values, is only significant when full logging is activated for index build operations, such as CREATE INDEX and REORG INDEX operations, or for an index rebuild. You can activate logging by enabling the *logindexbuild* database configuration parameter or by enabling LOG INDEX BUILD when altering a table. By setting *indexrec* to either RESTART or ACCESS, operations involving a logged index build can be rolled forward without leaving the index object in an invalid state, which would require the index to be rebuilt at a later time.

instance_memory - Instance memory

This parameter specifies the maximum amount of memory that can be allocated for a database partition.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable online

Default [range]

32-bit platforms

Automatic [0 - 1 000 000]

64-bit platforms

Automatic [0 - 68 719 476 736]

Unit of measure

Pages (4 KB)

When allocated

When the instance is started

When freed

When the instance is stopped

The default value of **instance_memory** is **AUTOMATIC**, meaning that its actual value is computed at database partition activation time (db2start). The actual value used is between 75 percent and 95 percent of the physical RAM on the system, divided by the number of configured local database partitions in the instance. This value should be suitable for dedicated database server systems.

Note:

- If the specified value of **instance_memory** is larger than the amount of physical memory on the system, then db2start fails with a SQL1220N (The database manager shared memory set cannot be allocated.)
- If **instance_memory** is dynamically updated to a value less than the amount of physical RAM, the request is processed and a new upper limit is set. Dynamic decreases to **instance_memory** are allowed only if the new setting is larger than the current amount of in-use **instance_memory**, otherwise, the request is deferred to the next db2start.
- If **instance_memory** is dynamically updated to a value greater than the amount of physical RAM while the instance is active, the request is deferred, and the next db2start fails with a SQL1220N (The database manager shared memory set cannot be allocated.)

When fast communication manager (FCM) shared memory is allocated, each local database partition's share of the overall FCM shared memory size for the system is accounted for in that database partition's **instance_memory** limit.

If memory is requested for a particular heap, and the database partition memory limit (**instance_memory**) has already been reached, then DB2 will first attempt to

reduce memory usage in all shared and private heaps by the requested amount of memory. If there is still insufficient free **instance_memory**, then the request fails, and the application that initiated the request receives an appropriate SQLCODE that describes which heap experienced an out-of-memory failure.

The exception to this behavior is for memory requests that are known to be critical to the operation of DB2 (that is, failing the memory request results the database is marked as invalid, or the instance is shut down). Note that the critical requests will first attempt to reduce current memory usage by the database partition. If there is still insufficient free **instance_memory**, DB2 still requests that memory from the operating system. If the operating system allows the memory request, then the current value of **instance_memory** will exceed the configured limit, however, all other non-critical memory requests will fail until enough memory has been freed.

Note: Restriction for DPF instances: although **instance_memory** specifies the amount of memory a single DB2 database partition might allocate, it is an instance-level configuration parameter, so all database partitions have the same **instance_memory** setting. However, if **instance_memory** is set to AUTOMATIC, the actual upper bound is computed individually on each separate machine based on the amount of RAM and the number of local partitions defined, so it is possible for different partitions to have different memory limits in effect.

Controlling DB2 Memory consumption:

When **instance_memory** is set to AUTOMATIC, a fixed upper bound on total memory consumption for the instance is set at instance startup (db2start). Actual memory consumption by DB2 varies depending on the workload. When STMM is enabled to perform **database_memory** tuning (by default for new databases), during run-time, STMM dynamically updates the size of performance-critical heaps within the database shared memory set according to the free physical memory on the system, while ensuring that there is sufficient free **instance_memory** available for functional memory requirements.

Depending on workload, DB2's default memory configuration adapts to the memory requirements of the instance without requiring explicit self-tuning of overall instance memory. For instance:

- For heavily-used instances, STMM increases the size of performance-critical heaps as needed. More functional memory is consumed, as there are more database agents servicing applications and consuming functional memory. If there is enough free **instance_memory** but very little free physical memory on the system, STMM starts decreasing the size of performance-critical heaps ensure that the system does not start paging. As functional memory requirements drop, free physical memory on the system should increase, and STMM will start increasing the performance-critical heaps again.
- For less heavily-used instances, there is less functional memory consumed by the instance, and if there is insufficient free physical memory left on the system, STMM shrinks performance-critical heaps.

If **instance_memory** is set to a specific value, and at least one active database has an AUTOMATIC value for **database_memory**, and STMM is enabled for that database, then STMM increases the **database_memory** size such that DB2 uses almost the entire amount of memory specified by **instance_memory**, ensuring only that enough free **instance_memory** is available for functional memory requests. In this scenario, STMM does not

monitor free physical memory on the machine, therefore, **instance_memory** must be configured properly to ensure that paging will not occur.

Use the new `admin_get_dbp_mem_usage` user-defined function (UDF) to get the total memory consumption by a DB2 instance for a specific database partition, or for all database partitions. This UDF also returns the current upper bound value.

Limitation on some Linux kernels:

Due to operating system limitations on some Linux kernels, STMM does not allow setting **database_memory** to AUTOMATIC unless **instance_memory** is set to a specific value (not AUTOMATIC). If **database_memory** is set to AUTOMATIC, and **instance_memory** is later set back to AUTOMATIC, the **database_memory** configuration parameter will be updated to COMPUTED during the next database activation. If some databases are already active, STMM stops tuning their overall **database_memory** sizes. This limitation is removed on Red Hat Enterprise Linux (RHEL) 5 and SUSE Linux Enterprise Server 10 SP1 or higher platforms.

intra_parallel - Enable intra-partition parallelism

This parameter specifies whether the database manager can use intra-partition parallelism.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

NO (0) [SYSTEM (-1), NO (0), YES (1)]

A value of -1 causes the parameter value to be set to “YES” or “NO” based on the hardware on which the database manager is running.

Some of the operations that can take advantage of parallel performance improvements when this parameter is “YES” include database queries and index creation.

Note:

- Parallel index creation does not use this configuration parameter.
- If you change this parameter value, packages might be rebound to the database, and some performance degradation might occur.

java_heap_sz - Maximum Java interpreter heap size

This parameter determines the maximum size of the heap that is used by the Java interpreter started to service Java DB2 stored procedures and UDFs.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

HP-UX

4096 [0 - 524 288]

All other operating systems

2048 [0 - 524 288]

Unit of measure

Pages (4 KB)

When allocated

When a Java stored procedure or UDF starts

When freed

When the db2fmp process (fenced) or the db2agent process (trusted) terminates.

There is one heap for each DB2 process (one for each agent or subagent on Linux and UNIX platforms, and one for each instance on other platforms). There is one heap for each fenced UDF and fenced stored procedure process. There is one heap per agent (not including sub-agents) for trusted routines. There is one heap per db2fmp process running a Java stored procedure. For multithreaded db2fmp processes, multiple applications using threadsafe fenced routines are serviced from a single heap. In all situations, only the agents or processes that run Java UDFs or stored procedures ever allocate this memory. On partitioned database systems, the same value is used at each database partition.

XML data is materialized when passed to stored procedures as IN, OUT, or INOUT parameters. When you are using Java stored procedures, the heap size might need to be increased based on the quantity and size of XML arguments, and the number of external stored procedures that are being executed concurrently.

jdk_path - Software Developer's Kit for Java installation path

This parameter specifies the directory under which the Software Developer's Kit (SDK) for Java, to be used for running Java stored procedures and user-defined functions, is installed. The CLASSPATH and other environment variables used by the Java interpreter are computed from the value of this parameter.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]
Null [Valid path]

If the SDK for Java was installed with your DB2 product, this parameter is set properly. However, if you reset the database manager (dbm cfg) parameter, you need to specify where the SDK for Java is installed.

keepfenced - Keep fenced process

This parameter indicates whether or not a fenced mode process is kept after a fenced mode routine call is complete. Fenced mode processes are created as separate system entities in order to isolate user-written fenced mode code from the database manager agent process. This parameter is only applicable on database servers.

Configuration type
Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type
Configurable

Default [range]
Yes [Yes; No]

If *keepfenced* is set to *no*, and the routine being executed is not threadsafe, a new fenced mode process is created and destroyed for each fenced mode invocation. If *keepfenced* is set to *no*, and the routine being executed is threadsafe, the fenced mode process persists, but the thread created for the call is terminated. If *keepfenced* is set to *yes*, a fenced mode process or thread is reused for subsequent fenced mode calls. When the database manager is stopped, all outstanding fenced mode processes and threads will be terminated.

Setting this parameter to *yes* will result in additional system resources being consumed by the database manager for each fenced mode process that is activated, up to the value contained in the *fenced_pool* parameter. A new process is only created when no existing fenced mode process is available to process a subsequent fenced routine invocation. This parameter is ignored if *fenced_pool* is set to 0.

Recommendation: In an environment in which the number of fenced mode requests is large relative to the number of non-fenced mode requests, and system resources are not constrained, then this parameter can be set to *yes*. This will improve the fenced mode process performance by avoiding the initial fenced mode process creation overhead since an existing fenced mode process will be used to process the call. In particular, for Java routines, this will save the cost of starting the Java Virtual Machine (JVM), a very significant performance improvement.

For example, in an OLTP debit-credit banking transaction application, the code to perform each transaction could be performed in a stored procedure which executes in a fenced mode process. In this application, the main workload is performed out of fenced mode processes. If this parameter is set to *no*, each transaction incurs the overhead of creating a new fenced mode process, resulting in a significant

performance reduction. If, however, this parameter is set to *yes*, each transaction would try to use an existing fenced mode process, which would avoid this overhead.

In previous versions of DB2, this parameter was known as *keepdari*.

local_gssplugin - GSS API plug-in used for local instance level authorization

This parameter specifies the name of the default GSS API plug-in library to be used for instance level local authorization when the value of the *authentication* database manager configuration parameter is set to GSSPLUGIN or GSS_SERVER_ENCRYPT.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

Null [any valid string]

max_connections - Maximum number of client connections

This parameter indicates the maximum number of client connections allowed per database partition.

Configuration type

Database manager

Parameter type

Configurable online

Applies to

- Database server with local and remote clients
- Database server with local clients
- Database Server or Connect Server with local and remote clients" (for *max_connections*, *max_coordagents*, *num_initagents*, *num_poolagents*, and also *federated_async*, if you are using a Federated environment)

Default [range]

-1 and AUTOMATIC (*max_coordagents*) [-1 and AUTOMATIC; 1–64 000]

A setting of -1 means that the value associated with *max_coordagents* will be used, not the automatic setting or behavior. AUTOMATIC means that the database manager picks the value for this parameter using whatever technique works best. AUTOMATIC is an ON/OFF switch in the configuration file and is independent of the value, hence both -1 and AUTOMATIC can be the default setting.

For details, see: "Restrictions and behavior when configuring *max_coordagents* and *max_connections*" on page 458.

The Concentrator

The Concentrator is OFF when *max_connections* is equal to or less than *max_coordagents*. The Concentrator is ON when *max_connections* is greater than *max_coordagents*.

This parameter controls the maximum number of client applications that can be connected to a database partition in the instance. Typically, each application is assigned a coordinator agent. The agent facilitates the operations between the application and the database. When the default value for this parameter is used, the Concentrator feature is not activated. As a result, each agent operates within its own private memory and shares database manager and database global resources, such as the buffer pool, with other agents. When the parameter is set to a value greater than the default, the Concentrator feature is activated.

max_connretries - Node connection retries

This parameter specifies the maximum number of times an attempt will be made to establish a TCP/IP connection between two database partition servers.

Configuration type

Database manager

Applies to

Partitioned database server with local and remote clients

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

5 [0–100]

If the attempt to establish communication between two database partition servers fails (for example, the value specified by the *conn_elapse* parameter is reached), *max_connretries* specifies the number of connection retries that can be made to a database partition server. If the value specified for this parameter is exceeded, an error is returned.

max_coordagents - Maximum number of coordinating agents

This parameter is used to limit the number of coordinating agents.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable online

Default [range]

200, Automatic [-1; 0–64 000]

A setting of -1 translates into a value of 200.

For details, see: “Restrictions and behavior when configuring `max_coordagents` and `max_connections`” on page 458.

The Concentrator

When the Concentrator is OFF, that is, when `max_connections` is equal to or less than `max_coordagents`, this parameter determines the maximum number of coordinating agents that can exist at one time on a server node.

One coordinating agent is acquired for each local or remote application that connects to a database or attaches to an instance. Requests that require an instance attachment include CREATE DATABASE, DROP DATABASE, and Database System Monitor commands.

When the Concentrator is ON, that is, when `max_connections` is greater than `max_coordagents`, there might be more connections than coordinator agents to service them. An application is in an active state only if there is a coordinator agent servicing it. Otherwise, the application is in an inactive state. Requests from an active application will be serviced by the database coordinator agent (and subagents in SMP or MPP configurations). Requests from an inactive application will be queued until a database coordinator agent is assigned to service the application, when the application becomes active. As a result, this parameter might be used to control the load on the system.

`max_querydegree` - Maximum query degree of parallelism

This parameter specifies the maximum degree of intra-partition parallelism that is used for any SQL statement executing on this instance of the database manager. An SQL statement will not use more than this number of parallel operations within a database partition when the statement is executed.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Propagation class

Statement boundary

Default [range]

-1 (ANY) [ANY, 1 - 32 767] (ANY means system-determined)

The `intra_parallel` configuration parameter must be set to “YES” to enable the database partition to use intra-partition parallelism for SQL statements. The `intra_parallel` parameter is no longer required for parallel index creation.

The default value for this configuration parameter is -1. This value means that the system uses the degree of parallelism determined by the optimizer; otherwise, the user-specified value is used.

Note: The degree of parallelism for an SQL statement can be specified at statement compilation time using the CURRENT DEGREE special register or the DEGREE bind option.

The maximum query degree of parallelism for an active application can be modified using the SET RUNTIME DEGREE command. The actual runtime degree used is the lower of:

- *max_querydegree* configuration parameter
- Application runtime degree
- SQL statement compilation degree

This configuration parameter applies to queries only.

max_time_diff - Maximum time difference among nodes

This parameter specifies the maximum time difference, in minutes, that is permitted among the database partition servers listed in the node configuration file.

Configuration type

Database manager

Applies to

Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

60 [1 - 1 440]

Unit of measure

Minutes

Each database partition server has its own system clock. If two or more database partition servers are associated with a transaction, and their clocks are not synchronized to within the time specified by this parameter, the transaction is rejected and an SQLCODE is returned. (The transaction is rejected only if data modification is associated with it.)

DB2 uses *Coordinated Universal Time* (UTC), so different time zones are not a consideration when you set this parameter. The Coordinated Universal Time is the same as Greenwich Mean Time.

maxagents - Maximum number of agents

This parameter is deprecated in Version 9.5, but is still being used by pre-Version 9.5 data servers and clients. Any value specified for this configuration parameter will be ignored by the DB2 Version 9.5 database manager

Note: The following information applies only to pre-Version 9.5 data servers and clients.

This parameter indicates the maximum number of database manager agents, whether coordinator agents or subagents, available at any given time to accept application requests

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

200 [1 - 64 000]

400[®] [1 - 64 000] on partitioned database server with local and remote clients

Unit of measure

Counter

If you want to limit the number of coordinating agents, use the *max_coordagents* parameter.

This parameter can be useful in memory constrained environments to limit the total memory usage of the database manager, because each additional agent requires additional memory.

Recommendation: The value of *maxagents* should be at least the sum of the values for *maxappls* in each database allowed to be accessed concurrently. If the number of databases is greater than the *numdb* parameter, then the safest course is to use the product of *numdb* with the largest value for *maxappls*.

Each additional agent requires some resource overhead that is allocated at the time the database manager is started.

If you are encountering memory errors when trying to connect to a database, try making the following configuration adjustments:

- In a non-partitioned database environment with no intra-query parallelism enabled, increase the value of the *maxagents* database configuration parameter.
- In a partitioned database environment or an environment where intra-query parallelism is enabled, increase the larger of *maxagents* or *max_coordagents*.

maxcagents - Maximum number of concurrent agents

This parameter is deprecated in Version 9.5, but is still being used by pre-Version 9.5 data servers and clients. Any value specified for this configuration parameter will be ignored by the DB2 Version 9.5 database manager.

Note: The following information applies only to pre-Version 9.5 data servers and clients.

This parameter is used to control the load on the system during periods of high simultaneous application activity by limiting the maximum number of database manager agents that can be concurrently executing a database manager transaction

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients

- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

-1 (*max_coordagents*) [-1; 1 - *max_coordagents*]

Unit of measure

Counter

This parameter does not limit the number of applications that can have connections to a database. It only limits the number of database manager agents that can be processed concurrently by the database manager at any one time, thereby limiting the usage of system resources during times of peak processing. For example, you might have a system requiring a large number of connections but with a limited amount of memory to serve those connections. Adjusting this parameter can be useful in such an environment, where a period of high simultaneous activity could cause excessive operating system paging.

A value of -1 indicates that the limit is *max_coordagents*.

Recommendation: In most cases the default value for this parameter will be acceptable. In cases where the high concurrency of applications is causing problems, you can use benchmark testing to tune this parameter to optimize the performance of the database.

mon_heap_sz - Database system monitor heap size

This parameter determines the amount of the memory, in pages, to allocate for database system monitor data. Memory is allocated from the monitor heap when you perform database monitoring activities such as taking a snapshot, turning on a monitor switch, resetting a monitor, or activating an event monitor.

With Version 9.5, this database configuration parameter has a default value of *AUTOMATIC*, meaning that the monitor heap can increase as needed until the *instance_memory* limit is reached.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable online

Default [range]

Automatic [0 - 60 000]

Unit of measure

Pages (4 KB)

When allocated

When the database manager is started with the *db2start* command

When freed

When the database manager is stopped with the *db2stop* command

A value of zero prevents the database manager from collecting database system monitor data.

Recommendation: The amount of memory required for monitoring activity depends on the number of monitoring applications (applications taking snapshots or event monitors), which switches are set, and the level of database activity.

If the configured memory in this heap runs out and no more unreserved memory is available in the instance shared memory region, one of the following will occur:

- When the first application connects to the database for which this event monitor is defined, an error message is written to the administration notification log.
- If an event monitor being started dynamically using the SET EVENT MONITOR statement fails, an error code is returned to your application.
- If a monitor command or API subroutine fails, an error code is returned to your application.

nodetype - Machine node type

This parameter provides information about the DB2 products which you have installed on your machine and, as a result, information about the type of database manager configuration.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Informational

The following are the possible values returned by this parameter and the products associated with that node type:

- **Database server with local and remote clients** – a DB2 server product, supporting local and remote Data Server Runtime Clients, and capable of accessing other remote database servers.
- **Client** – a Data Server Runtime Client capable of accessing remote database servers.
- **Database server with local clients** – a DB2 relational database management system, supporting local Data Server Runtime Clients and capable of accessing other, remote database servers.
- **Partitioned database server with local and remote clients** – a DB2 server product, supporting local and remote Data Server Runtime Clients, and capable of accessing other remote database servers, and capable of parallelism.

notifylevel - Notify level

This parameter specifies the type of administration notification messages that are written to the administration notification log.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

3 [0 — 4]

On Linux and UNIX platforms, the administration notification log is a text file called *instance.nfy*. On Windows, all administration notification messages are written to the Event Log. The errors can be written by DB2, the Health Monitor, the Capture and Apply programs, and user applications.

Valid values for this parameter are:

- **0** — No administration notification messages captured. (This setting is not recommended.)
- **1** — Fatal or unrecoverable errors. Only fatal and unrecoverable errors are logged. To recover from some of these conditions, you might need assistance from DB2 service.
- **2** — Immediate action required. Conditions are logged that require immediate attention from the system administrator or the database administrator. If the condition is not resolved, it could lead to a fatal error. Notification of very significant, non-error activities (for example, recovery) might also be logged at this level. This level will capture Health Monitor alarms.
- **3** — Important information, no immediate action required. Conditions are logged that are non-threatening and do not require immediate action but might indicate a non-optimal system. This level will capture Health Monitor alarms, Health Monitor warnings, and Health Monitor attentions.
- **4** — Informational messages.

The administration notification log includes messages having values up to and including the value of *notifylevel*. For example, setting *notifylevel* to 3 will cause the administration notification log to include messages applicable to levels 1, 2, and 3.

For a user application to be able to write to the notification file or Windows Event Log, it must call the `db2AdminMsgWrite` API.

Recommendation: You might want to increase the value of this parameter to gather additional problem determination data to help resolve a problem. Note that you must set *notifylevel* to a value of 2 or higher for the Health Monitor to send any notifications to the contacts defined in its configuration.

num_initagents - Initial number of agents in pool

This parameter determines the initial number of idle agents that are created in the agent pool at DB2START time.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable online

Default [range]

0 [0–64 000]

The database manager always starts the *num_initagents* idle agents as part of the *db2start* command, except if the value of this parameter is greater than *num_poolagents* during start up and *num_poolagents* is not set to AUTOMATIC. In this case, the database manager only starts the *num_poolagents* idle agents since there is no reason to start more idle agents than can be pooled.

num_initfenced - Initial number of fenced processes

This parameter indicates the initial number of nonthreaded, idle *db2fmp* processes that are created in the *db2fmp* pool at *DB2START* time.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable online

Default [range]

0 [0–64 000]

Setting this parameter will reduce the initial startup time for running non-threadsafe C and Cobol routines. This parameter is ignored if *keepfenced* is not specified.

It is much more important to set *fenced_pool* to an appropriate size for your system than to start up a number of *db2fmp* processes at *DB2START* time.

In previous versions, this parameter was known as *num_initdaris*.

num_poolagents - Agent pool size

This parameter sets the maximum size of the idle agent pool.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable online

Default

100, Automatic [-1, 0–64 000]

This configuration parameter is set to AUTOMATIC with a value of 100 as the default. A setting of -1 is still supported, and translates into a value of 100. When this parameter is set to AUTOMATIC, the database manager automatically manages the number of idle agents to pool. Typically, this means that once an agent completes its work, it is not terminated, but becomes idle for a period of time. Depending on the workload and type of agent, it might be terminated after a certain amount of time.

When using AUTOMATIC, you can still specify a value for the `num_poolagents` configuration parameter. Additional idle agents will always be pooled when the current number of pooled idle agents is less than or equal to the value that you specified.

Examples:

num_poolagents is set to 100 and AUTOMATIC

As an agent becomes free, it is added to the idle agent pool, where at some point the database manager evaluates whether it should be terminated or not. At the time when the database manager considers terminating the agent, if the total number of idle agents pooled is greater than 100, this agent will be terminated. If there are less than 100 idle agents, the idle agent will remain awaiting work. Using the AUTOMATIC setting allows additional idle agents beyond 100 to be pooled, which might be useful during periods of heavier system activity when the frequency of work can fluctuate on a larger scale. For cases where there are likely to be less than 100 idle agents at any given time, agents are guaranteed to be pooled. Periods of light system activity can benefit from this by incurring a less start up cost for new work.

num_poolagents is configured dynamically

If the parameter value is increased to a value greater than the number of pooled agents, the effects are immediate. As new agents become idle, they are pooled. If the parameter value is decreased, the database manager does not immediately reduce the number of agents in the pool. Rather, the pool size remains as it is, and agents are terminated as they are used and become idle again—gradually reducing the number of agents in the pool to the new limit.

Recommendation: For most environments the default of 0 and AUTOMATIC will be sufficient. If you have a specific workload where you feel too many agents are being created and terminated, you can consider increasing the value of *num_poolagents* while leaving the parameter set to AUTOMATIC.

numdb - Maximum number of concurrently active databases including host and System i databases

This parameter specifies the number of local databases that can be concurrently active (that is, have applications connected to them), or the maximum number of different database aliases that can be cataloged on a DB2 Connect server.

Configuration type

Database manager

Applies to

- Database server with local and remote clients

- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

UNIX 8 [1 — 256]

Windows Database server with local and remote clients
8 [1 — 256]

Windows Database server with local clients
3 [1 — 256]

Unit of measure

Counter

Each database takes up storage, and an active database uses a new shared memory segment.

Recommendation: It is generally best to set this value to the actual number of databases that are already defined to the database manager, and to add about 10% to this value to allow for growth.

Changing the *numdb* parameter can impact the total amount of memory allocated. As a result, frequent updates to this parameter are not recommended. When updating this parameter, you should consider the other configuration parameters that can allocate memory for a database or an application connected to that database.

query_heap_sz - Query heap size

This parameter is deprecated in Version 9.5, but is still being used by pre-Version 9.5 data servers and clients. Any value specified for this configuration parameter will be ignored by the DB2 Version 9.5 database manager.

Note: The following information applies only to pre-Version 9.5 data servers and clients.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

1 000 [2 - 524 288]

Unit of measure

Pages (4 KB)

When allocated

When an application (either local or remote) connects to the database

When freed

When the application disconnects from the database, or detaches from the instance

This parameter specifies the **maximum** amount of memory that can be allocated for the query heap, ensuring that an application does not consume unnecessarily large amounts of virtual memory within an agent.

A query heap is used to store each query in the agent's private memory. The information for each query consists of the input and output SQLDA, the statement text, the SQLCA, the package name, creator, section number, and consistency token.

The query heap is also used for the memory allocated for blocking cursors. This memory consists of a cursor control block and a fully resolved output SQLDA.

The initial query heap allocated will be the same size as the application support layer heap, as specified by the *aslheapsz* parameter. The query heap size must be greater than or equal to two (2), and must be greater than or equal to the *aslheapsz* parameter. If this query heap is not large enough to handle a given request, it will be reallocated to the size required by the request (not exceeding *query_heap_sz*). If this new query heap is more than 1.5 times larger than *aslheapsz*, the query heap will be reallocated to the size of *aslheapsz* when the query ends.

Recommendation: In most cases the default value will be sufficient. As a minimum, you should set *query_heap_sz* to a value at least five times larger than *aslheapsz*. This will allow for queries larger than *aslheapsz* and provide additional memory for three or four blocking cursors to be open at a given time.

If you have very large LOBs, you might need to increase the value of this parameter so the query heap will be large enough to accommodate those LOBs.

release - Configuration file release level

This parameter specifies the release level of the configuration file.

Configuration type

Database manager, Database

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Informational

resync_interval - Transaction resync interval

This parameter specifies the time interval in seconds for which a transaction manager (TM), resource manager (RM) or sync point manager (SPM) should retry the recovery of any outstanding indoubt transactions found in the TM, the RM, or the SPM. This parameter is applicable when you have transactions running in a distributed unit of work (DUOW) environment. This parameter also applies to recovery of federated database systems.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

180 [1 - 60 000]

Unit of measure

Seconds

Recommendation: If, in your environment, indoubt transactions will not interfere with other transactions against your database, you might want to increase the value of this parameter. If you are using a DB2 Connect gateway to access DRDA2 application servers, you should consider the effect indoubt transactions might have at the application servers even though there will be no interference with local data access. If there are no indoubt transactions, the performance impact will be minimal.

rqrioblk - Client I/O block size

This parameter specifies the size of the communication buffer between remote applications and their database agents on the database server. It is also used to determine the I/O block size at the Data Server Runtime Client when a blocking cursor is opened.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

32 767 [4 096 - 65 535]

Unit of measure

Bytes

When allocated

- When a remote client application issues a connection request for a server database
- When a blocking cursor is opened, additional blocks are opened at the client

When freed

- When the remote application disconnects from the server database
- When the blocking cursor is closed

When a Data Server Runtime Client requests a connection to a remote database, this communication buffer is allocated on the client. On the database server, a

communication buffer of 32 767 bytes is initially allocated, until a connection is established and the server can determine the value of *rqrioblk* at the client. Once the server knows this value, it will reallocate its communication buffer if the client's buffer is not 32 767 bytes.

The memory for blocked cursors is allocated out of the application's private address space, so you should determine the optimal amount of private memory to allocate for each application program. If the Data Server Runtime Client cannot allocate space for a blocking cursor out of an application's private memory, a non-blocking cursor will be opened.

Recommendation: For non-blocking cursors, a reason for increasing the value of this parameter would be if the data (for example, large object data) to be transmitted by a single query statement is so large that the default value is insufficient.

You should also consider the effect of this parameter on the number and potential size of blocking cursors. Large row blocks might yield better performance if the number or size of rows being transferred is large (for example, if the amount of data is greater than 4 096 bytes). However, there is a trade-off in that larger record blocks increase the size of the working set memory for each connection.

Larger record blocks might also cause more fetch requests than are actually required by the application. You can control the number of fetch requests using the OPTIMIZE FOR clause on the SELECT statement in your application.

sheapthres - Sort heap threshold

This parameter is an instance-wide soft limit on the total amount of memory that can be consumed by private sorts at any given time. When the total private sort memory consumption for an instance reaches this limit, the memory allocated for additional incoming private sort requests is considerably reduced.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients
- OLAP functions

Parameter type

Configurable online

Propagation class

Immediate

Default [range]

UNIX 32-bit platforms

0 [0 - 2 097 152]

Windows 32-bit platforms

0 [0 - 2 097 152]

64-bit platforms

0 [0 - 2 147 483 647]

Unit of measure

Pages (4 KB)

Examples of operations that use the sort heap include: sorts, hash joins, dynamic bitmaps (used for index ANDing and Star Joins), and table in-memory operations.

Explicit definition of the threshold prevents the database manager from using excessive amounts of memory for large numbers of sorts.

There is no reason to increase the value of this parameter when moving from a non-partitioned to a partitioned database environment. Once you have tuned the database and database manager configuration parameters on a single database partition environment, the same values will in most cases work well in a partitioned database environment. The only way to set this parameter to different values on different nodes or database partitions is to create more than one DB2 instance. This will require managing different DB2 databases over different database partition groups. Such an arrangement defeats the purpose of many of the advantages of a partitioned database environment.

When the instance-level *sheapthres* is set to 0, then the tracking of sort memory consumption is done at the database level only and memory allocation for sorts is constrained by the value of the database-level *sheapthres_shr* configuration parameter.

Automatic tuning of *sheapthres_shr* is allowed only when the database manager configuration parameter *sheapthres* is set to 0.

This parameter will not be dynamically updatable if any of the following are true:

- The starting value for *sheapthres* is 0 and the target value is a value different from 0.
- The starting value for *sheapthres* is a value different from 0 and the target value is 0.

Recommendation: Ideally, you should set this parameter to a reasonable multiple of the largest *sortheap* parameter you have in your database manager instance. This parameter should be **at least** two times the largest *sortheap* defined for any database within the instance.

If you are doing private sorts and your system is not memory constrained, an ideal value for this parameter can be calculated using the following steps:

1. Calculate the typical sort heap usage for each database:
$$\begin{aligned} & \text{(typical number of concurrent agents running against the database)} \\ & * \text{(sortheap, as defined for that database)} \end{aligned}$$
2. Calculate the sum of the above results, which provides the total sort heap that could be used under typical circumstances for all databases within the instance.

You should use benchmarking techniques to tune this parameter to find the proper balance between sort performance and memory usage.

You can use the database system monitor to track the sort activity, using the post threshold sorts (*post_threshold_sorts*) monitor element.

spm_log_file_sz - Sync point manager log file size

This parameter identifies the sync point manager (SPM) log file size in 4 KB pages.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

256 [4 - 1000]

Unit of measure

Pages (4 KB)

The log file is contained in the `spmlog` sub-directory under `sqllib` and is created the first time SPM is started.

Recommendation: The sync point manager log file size should be large enough to maintain performance, but small enough to prevent wasted space. The size required depends on the number of transactions using protected conversations, and how often COMMIT or ROLLBACK is issued.

To change the size of the SPM log file:

1. Determine that there are no indoubt transactions by using the LIST DRDA INDOUBT TRANSACTIONS command.
2. If there are none, stop the database manager.
3. Update the database manager configuration with a new SPM log file size.
4. Go to the `$HOME/sqllib` directory and issue `rm -fr spmlog` to delete the current SPM log. (Note: This shows the AIX command. Other systems might require a different remove or delete command.)
5. Start the database manager. A new SPM log of the specified size is created during the startup of the database manager.

spm_log_path - Sync point manager log file path

This parameter specifies the directory where the sync point manager (SPM) logs are written.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]`sqllib/spmlog` [any valid path or device]

By default, the logs are written to the `sqllib/spmlog` directory, which, in a high-volume transaction environment, can cause an I/O bottleneck. Use this

parameter to have the SPM log files placed on a faster disk than the current sqllib/spmlog directory. This allows for better concurrency among the SPM agents.

spm_max_resync - Sync point manager resync agent limit

This parameter identifies the number of agents that can simultaneously perform resync operations.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

20 [10 — 256]

spm_name - Sync point manager name

This parameter identifies the name of the sync point manager (SPM) instance to the database manager.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default

Derived from the TCP/IP hostname

srvcon_auth - Authentication type for incoming connections at the server

This parameter specifies how and where user authentication is to take place when handling incoming connections at the server; it is used to override the current authentication type.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

Null [CLIENT; SERVER; SERVER_ENCRYPT; KERBEROS;
KRB_SERVER_ENCRYPT; GSSPLUGIN; GSS_SERVER_ENCRYPT]

If a value is not specified, DB2 uses the value of the *authentication* database manager configuration parameter.

For a description of each authentication type, see “authentication - Authentication type” on page 465.

srvcon_gssplugin_list - List of GSS API plug-ins for incoming connections at the server

This parameter specifies the GSS API plug-in libraries that are supported by the database server. It handles incoming connections at the server when the *srvcon_auth* parameter is specified as KERBEROS, KRB_SERVER_ENCRYPT, GSSPLUGIN or GSS_SERVER_ENCRYPT, or when *srvcon_auth* is not specified, and authentication is specified as KERBEROS, KRB_SERVER_ENCRYPT, GSSPLUGIN or GSS_SERVER_ENCRYPT.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

Null [any valid string]

By default, the value is null. If the authentication type is GSSPLUGIN and this parameter is NULL, an error is returned. If the authentication type is KERBEROS and this parameter is NULL, the DB2-supplied kerberos module or library is used. This parameter is not used if another authentication type is used.

When the authentication type is KERBEROS and the value of this parameter is not NULL, the list must contain exactly one Kerberos plug-in, and that plug-in is used for authentication (all other GSS plug-ins in the list are ignored). If there is more than one Kerberos plug-in, an error is returned.

Each GSS API plug-in name must be separated by a comma (,) with no space either before or after the comma. Plug-in names should be listed in the order of preference.

srvcon_pw_plugin - Userid-password plug-in for incoming connections at the server

This parameter specifies the name of the default userid-password plug-in library to be used for server-side authentication. It handles incoming connections at the server when the *srvcon_auth* parameter is specified as CLIENT, SERVER, SERVER_ENCRYPT, or DATA_ENCRYPT or when *srvcon_auth* is not specified, and *authentication* is specified as CLIENT, SERVER, SERVER_ENCRYPT, or DATA_ENCRYPT.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

Null [any valid string]

By default, the value is null and the DB2-supplied userid-password plug-in library is used. The plug-in will be used for all group lookups. For non-root installations, if the DB2 userid and password plug-in library is used, the db2rfe command must be run before using your DB2 product.

srv_plugin_mode - Server plug-in mode

This parameter specifies whether plug-ins are to run in fenced mode or unfenced mode. Unfenced mode is the only supported mode.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

UNFENCED

start_stop_time - Start and stop timeout

This parameter specifies the time, in minutes, within which all database partition servers must respond to a DB2START or a DB2STOP command. It is also used as the timeout value during an ADD DBPARTITIONNUM operation.

Configuration type

Database manager

Applies to

Database server with local and remote clients

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

10 [1 - 1 440]

Unit of measure

Minutes

Database partition servers that do not respond to a DB2START command within the specified time send a message to the db2start error log in the log subdirectory of the sql1ib subdirectory of the home directory for the instance. You should issue a DB2STOP on these nodes before restarting them.

Database partition servers that do not respond to a DB2STOP command within the specified time send a message to the db2stop error log in the log subdirectory of the sql1ib subdirectory of the home directory for the instance. You can either issue db2stop for each database partition server that does not respond, or for all of them. (Those that are already stopped will return stating that they are stopped.)

If a db2start or db2stop operation in a multi-partition database is not completed within the value specified by the *start_stop_time* database manager configuration parameter, the database partitions that have timed out will be killed internally. Environments with many database partitions with a low value for *start_stop_time* might experience this behavior. To resolve this behavior, increase the value of *start_stop_time*.

When adding a new database partition using one of the DB2START, START DATABASE MANAGER, or ADD DBPARTITIONNUM commands, the add database partition operation must determine whether or not each database in the instance is enabled for automatic storage. This is done by communicating with the catalog partition for each database. If automatic storage is enabled, the storage path definitions are retrieved as part of that communication. Likewise, if system temporary table spaces are to be created with the database partitions, the operation might have to communicate with another database partition server to retrieve the table space definitions for the database partitions that reside on that server. These factors should be considered when determining the value of the *start_stop_time* parameter.

svcename - TCP/IP service name

This parameter contains the name of the TCP/IP port which a database server will use to await communications from remote client nodes. This name must be the reserved for use by the database manager.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default

Null

In order to accept connection requests from a Data Server Runtime Client using TCP/IP, the database server must be listening on a port designated to that server. The system administrator for the database server must reserve a port (number *n*) and define its associated TCP/IP service name in the services file at the server.

The database server port (number *n*) and its TCP/IP service name need to be defined in the services file on the database client.

On Linux and UNIX systems, the services file is located in: `/etc/services`

The `svcname` parameter should be set to the service name associated with the main connection port so that when the database server is started, it can determine on which port to listen for incoming connection requests.

sysadm_group - System administration authority group name

This parameter defines the group name with SYSADM authority for the database manager instance.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default

Null

System administration (SYSADM) authority is the highest level of authority within the database manager and controls all database objects.

SYSADM authority is determined by the security facilities used in a specific operating environment.

- For the Windows operating system, this parameter can be set to any local group, and is defined in the Windows security database. Group names must adhere to the length limits specified in SQL and XML limits. If "NULL" is specified for this parameter, all members of the Administrators group have SYSADM authority.
- For Linux and UNIX systems, if "NULL" is specified as the value of this parameter, the SYSADM group defaults to the primary group of the instance owner.

If the value is not "NULL", the SYSADM group can be any valid UNIX group name.

To restore the parameter to its default (NULL) value, use `UPDATE DBM CFG USING SYSADM_GROUP NULL`. You must specify the keyword "NULL" in uppercase.

sysctrl_group - System control authority group name

This parameter defines the group name with system control (SYSCTRL) authority. SYSCTRL has privileges allowing operations affecting system resources, but does not allow direct access to data.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client

- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default

Null

Group names on all platforms are accepted as long as they adhere to the length limits specified in SQL and XML limits.

Attention: This parameter must be NULL for Windows clients when system security is used (that is, **authentication** is CLIENT, SERVER, or any other valid authentication). This is because the Windows operating systems do not store group information, thereby providing no way of determining if a user is a member of a designated SYSCTRL group. When a group name is specified, no user can be a member of it.

To restore the parameter to its default (NULL) value, use UPDATE DBM CFG USING SYSCTRL_GROUP NULL. You must specify the keyword NULL in uppercase.

sysmaint_group - System maintenance authority group name

This parameter defines the group name with system maintenance (SYSMAINT) authority.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default

Null

SYSMAINT has privileges to perform maintenance operations on all databases associated with an instance without having direct access to data.

Group names on all platforms are accepted as long as they adhere to the length limits specified in SQL and XML limits.

Attention: This parameter must be NULL for Windows clients when system security is used (that is, **authentication** is CLIENT, SERVER, or any other valid authentication). This is because the Windows operating systems do not store group information, thereby providing no way of determining if a user is a member of a designated SYSMAINT group. When a group name is specified, no user can be a member of it.

To restore the parameter to its default (NULL) value, use UPDATE DBM CFG USING SYSMON_GROUP NULL. You must specify the keyword NULL in uppercase.

sysmon_group - System monitor authority group name

This parameter defines the group name with system monitor (SYSMON) authority.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default

Null

Users having SYSMON authority at the instance level have the ability to take database system monitor snapshots of a database manager instance or its databases. SYSMON authority includes the ability to use the following commands:

- GET DATABASE MANAGER MONITOR SWITCHES
- GET MONITOR SWITCHES
- GET SNAPSHOT
- LIST ACTIVE DATABASES
- LIST APPLICATIONS
- LIST DCS APPLICATIONS
- RESET MONITOR
- UPDATE MONITOR SWITCHES

Users with SYSADM, SYSCTRL, or SYSMANT authority automatically have the ability to take database system monitor snapshots and to use these commands.

Group names on all platforms are accepted as long as they adhere to the length limits specified in SQL and XML limits.

To restore the parameter to its default (NULL) value, use UPDATE DBM CFG USING SYSMON_GROUP NULL. You must specify the keyword NULL in uppercase.

tm_database - Transaction manager database name

This parameter identifies the name of the transaction manager (TM) database for each DB2 instance.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client

- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

1ST_CONN [any valid database name]

A TM database can be:

- A local DB2 database
- A remote DB2 database that does not reside on a host or AS/400 system
- A DB2 for OS/390 Version 5 database if accessed via TCP/IP and the sync point manager (SPM) is not used.

The TM database is a database that is used as a logger and coordinator, and is used to perform recovery for indoubt transactions.

You can set this parameter to **1ST_CONN**, which will set the TM database to be the first database to which a user connects.

Recommendation: For simplified administration and operation, you might want to create a few databases over a number of instances and use these databases exclusively as TM databases.

tp_mon_name - Transaction processor monitor name

This parameter identifies the name of the transaction processing (TP) monitor product being used.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default

No default

Valid values

- CICS®
- MQ
- ENCINA
- CB
- SF
- TUXEDO
- TOPEND
- blank or some other value (for UNIX and Windows; no other possible values for Solaris or SINIX)

- If applications are run in a WebSphere Enterprise Server Edition CICS environment, this parameter should be set to "CICS"
- If applications are run in a WebSphere Enterprise Server Edition Encina[®] environment, this parameter should be set to "ENCINA"
- If applications are run in a WebSphere Enterprise Server Edition Component Broker environment, this parameter should be set to "CB"
- If applications are run in an IBM MQSeries[®] environment, this parameter should be set to "MQ"
- If applications are run in a BEA Tuxedo environment, this parameter should be set to "TUXEDO"
- If applications are run in an IBM San Francisco environment, this parameter should be set to "SF".

IBM WebSphere EJB and Microsoft Transaction Server users do not need to configure any value for this parameter.

If none of the above products are being used, this parameter should not be configured but left blank.

In previous versions of IBM DB2 on Windows, this parameter contained the path and name of the DLL which contained the XA Transaction Manager's functions *ax_reg* and *ax_unreg*. This format is still supported. If the value of this parameter does not match any of the above TP Monitor names, it will be assumed that the value is a library name which contains the *ax_reg* and *ax_unreg* functions. This is true for UNIX and Windows environments.

TXSeries[®] CICS and Encina Users: In previous versions of this product on Windows it was required to configure this parameter as "libEncServer:C" or "libEncServer:E". While this is still supported, it is no longer required. Configuring the parameter as "CICS" or "ENCINA" is sufficient.

MQSeries Users: In previous versions of this product on Windows it was required to configure this parameter as "mqmax". While this is still supported, it is no longer required. Configuring the parameter as "MQ" is sufficient.

Component Broker Users: In previous versions of this product on Windows it was required to configure this parameter as "somtrx1i". While this is still supported, it is no longer required. Configuring the parameter as "CB" is sufficient.

San Francisco Users: In previous versions of this product on Windows it was required to configure this parameter as "ibmsfDB2". While this is still supported, it is no longer required. Configuring the parameter as "SF" is sufficient.

The maximum length of the string that can be specified for this parameter is 19 characters.

It is also possible to configure this information in IBM DB2 Version 9.1's XA OPEN string. If multiple Transaction Processing Monitors are using a single DB2 instance, then it will be required to use this capability.

trust_allclnts - Trust all clients

This parameter and *trust_clntauth* are used to determine where users are validated to the database environment.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

YES [NO, YES, DRDAONLY]

This parameter is only active when the *authentication* parameter is set to CLIENT.

By accepting the default of “YES” for this parameter, all clients are treated as trusted clients. This means that the server assumes that a level of security is available at the client and the possibility that users can be validated at the client.

This parameter can only be changed to “NO” if the *authentication* parameter is set to CLIENT. If this parameter is set to “NO”, the untrusted clients must provide a userid and password combination when they connect to the server. Untrusted clients are operating system platforms that do not have a security subsystem for authenticating users.

Setting this parameter to “DRDAONLY” protects against all clients except clients from DB2 for OS/390 and z/OS, DB2 for VM and VSE, and DB2 for OS/400®. Only these clients can be trusted to perform client-side authentication. All other clients must provide a user ID and password to be authenticated by the server.

When *trust_allclnts* is set to “DRDAONLY”, the *trust_clntauth* parameter is used to determine where the clients are authenticated. If *trust_clntauth* is set to “CLIENT”, authentication occurs at the client. If *trust_clntauth* is set to “SERVER”, authentication occurs at the client if no password is provided, and at the server if a password is provided.

trust_clntauth - Trusted clients authentication

This parameter specifies whether a trusted client is authenticated at the server or the client when the client provides a userid and password combination for a connection. This parameter (and *trust_allclnts*) is only active if the *authentication* parameter is set to CLIENT. If a user ID and password are not provided, the client is assumed to have validated the user, and no further validation is performed at the server.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]
CLIENT [CLIENT, SERVER]

If this parameter is set to CLIENT (the default), the trusted client can connect without providing a user ID and password combination, and the assumption is that the operating system has already authenticated the user. If it is set to SERVER, the user ID and password will be validated at the server.

The numeric value for CLIENT is 0. The numeric value for SERVER is 1.

util_impact_lim - Instance impact policy

This parameter allows the database administrator (DBA) to limit the performance degradation of a throttled utility on the workload.

Configuration type
Database manager

Applies to

- Database server with local clients
- Database server with local and remote clients
- Partitioned database server with local and remote clients

Parameter type
Configurable Online

Propagation class
Immediate

Default [range]
10 [1 - 100]

Unit of measure
Percentage of allowable impact on workload

If the performance degradation is limited, the DBA can then run online utilities during critical production periods, and be guaranteed that the performance impact on production work will be within acceptable limits.

For example, a DBA specifying a *util_impact_lim* (impact policy) value of 10 can expect that a throttled backup invocation will not impact the workload by more than 10 percent.

If *util_impact_lim* is 100, no utility invocations will be throttled. In this case, the utilities can have an arbitrary (and undesirable) impact on the workload. If *util_impact_lim* is set to a value that is less than 100, it is possible to invoke utilities in throttled mode. To run in throttled mode, a utility must also be invoked with a non-zero priority.

Recommendation: Most users will benefit from setting *util_impact_lim* to a low value (for example, between 1 and 10).

A throttled utility will usually take longer to complete than an unthrottled utility. If you find that a utility is running for an excessively long time, increase the value of *util_impact_lim*, or disable throttling altogether by setting *util_impact_lim* to 100.

Database configuration parameters

alt_collate - Alternate collating sequence

This parameter specifies the collating sequence that is to be used for Unicode tables in a non-Unicode database.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

Null [IDENTITY_16BIT]

Until this parameter is set, Unicode tables and routines cannot be created in a non-Unicode database. Once set, this parameter cannot be changed or reset.

This parameter cannot be set for Unicode databases.

app_ctl_heap_sz - Application control heap size

This parameter is deprecated in Version 9.5, but is still being used by pre-Version 9.5 data servers and clients. Any value specified for this configuration parameter will be ignored by the DB2 Version 9.5 database manager. In Version 9.5, it has been replaced by the *appl_memory* configuration parameter.

Note: The following information applies only to pre-Version 9.5 data servers and clients.

For partitioned databases, and for non-partitioned databases with intra-parallelism enabled (*intra_parallel=ON*), this parameter specifies the average size of the shared memory area allocated for an application. For non-partitioned databases where intra-parallelism is disabled (*intra_parallel=OFF*), this is the maximum private memory that will be allocated for the heap. There is one application control heap per connection per database partition.

Configuration type

Database

Parameter type

Configurable

Default [range]

Database server with local and remote clients

- 128 [1 - 64 000] when *INTRA_PARALLEL* is not enabled
- 512 [1 - 64 000] when *INTRA_PARALLEL* is enabled

Database server with local clients

- 64 [1 - 64 000] (for non-UNIX platforms) when *INTRA_PARALLEL* is not enabled

- 512 [1 - 64 000] (for non-UNIX platforms) when INTRA_PARALLEL is enabled
- 128 [1 - 64 000] (for Linux and UNIX platforms) when INTRA_PARALLEL is not enabled
- 512 [1 - 64 000] (for Linux and UNIX platforms) when INTRA_PARALLEL is enabled

Partitioned database server with local and remote clients

512 [1 - 64 000]

Unit of measure

Pages (4 KB)

When allocated

When an application starts

When freed

When an application completes

The application control heap is required primarily for sharing information between agents working on behalf of the same request. Usage of this heap is minimal for non-partitioned databases when running queries with a degree of parallelism equal to 1.

This heap is also used to store descriptor information for declared temporary tables. The descriptor information for all declared temporary tables that have not been explicitly dropped is kept in this heap's memory and cannot be dropped until the declared temporary table is dropped.

Recommendation: Initially, start with the default value. You might have to set the value higher if you are running complex applications, if you have a system that contains a large number of database partitions, or if you use declared temporary tables. The amount of memory needed increases with the number of concurrently active declared temporary tables. A declared temporary table with many columns has a larger table descriptor size than a table with few columns, so having a large number of columns in an application's declared temporary tables also increases the demand on the application control heap.

appgroup_mem_sz - Maximum size of application group memory set

This parameter is deprecated in Version 9.5, but is still being used by pre-Version 9.5 data servers and clients. Any value specified for this configuration parameter will be ignored by the DB2 Version 9.5 database manager. In Version 9.5, it has been replaced by the *appl_memory* configuration parameter.

Note: The following information applies only to pre-Version 9.5 data servers and clients.

This parameter determines the size of the application group shared memory segment.

Configuration type

Database

Parameter type

Configurable

Default [range]

UNIX Database server with local clients (other than 32-bit HP-UX)
20 000 [1 - 1 000 000]

32-bit HP-UX

- Database server with local clients
- Database server with local and remote clients
- Partitioned database server with local and remote clients

10 000 [1 - 1 000 000]

Windows Database server with local clients

10 000 [1 - 1 000 000]

Database server with local and remote clients (other than 32-bit HP-UX)

30 000 [1 - 1 000 000]

Partitioned database server with local and remote clients (other than 32-bit HP-UX)

40 000 [1 - 1 000 000]

Unit of measure

Pages (4 KB)

Information that needs to be shared between agents working on the same application is stored in the application group shared memory segment.

In a partitioned database, or in a non-partitioned database with intra-partition parallelism enabled or concentrator enabled, multiple applications share one application group. One application group shared memory segment is allocated for the application group. Within the application group shared memory segment, each application will have its own application control heap, and all applications will share one application group shared heap.

The number of applications in one application group is calculated by:

$\text{appgroup_mem_sz} / \text{app_ctl_heap_sz}$

The application group shared heap size is calculated by:

$\text{appgroup_mem_sz} * \text{groupheap_ratio} / 100$

The size of each application control heap is calculated by:

$\text{app_ctl_heap_sz} * (100 - \text{groupheap_ratio}) / 100$

Recommendation: Retain the default value of this parameter unless you are experiencing performance problems.

appl_memory - Application Memory configuration parameter

This parameter allows DBAs and ISVs to control the maximum amount of application memory that is allocated by DB2 database agents to service application requests. By default, its value is set to AUTOMATIC, meaning that all application memory requests will be allowed as long as the total amount of memory allocated by the database partition is within the *instance_memory* limits.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Database server with local clients

- Partitioned database server with local and remote clients

Parameter type

Configurable online

Default [range]

Automatic [128 - 4 294 967 295]

Unit of measure

Pages (4 KB)

When allocated

During database activation

When freed

During database deactivation

Note: When *appl_memory* is set to AUTOMATIC, the initial application memory allocation at database activation time is minimal, and increases (or decreases) as needed. The change is applied in memory and the value of *appl_memory* does not change on disk as shown by `db2 get db cfg show detail`. On next activation, the value will be recalculated. If *appl_memory* is set to a specific value, then the requested amount of memory is allocated initially during database activation, and the application memory size does not change. If the initial amount of application memory cannot be allocated from the operating system, or exceeds the *instance_memory* limit, database activation fails with an SQL1084C error (Shared memory segments cannot be allocated).

applheapsz - Application heap size

In previous releases, the *applheapsz* database configuration parameter referred to the amount of application memory each individual database agent working for that application could consume. With Version 9.5, *applheapsz* refers to the total amount of application memory that can be consumed by the entire application. For DPF, Concentrator, or SMP configurations, this means that the *applheapsz* value used in previous releases may need to be increased under similar workloads, unless the AUTOMATIC setting is used.

With Version 9.5, this database configuration parameter has a default value of AUTOMATIC, meaning that it increases as needed until either the *appl_memory* limit is reached, or the *instance_memory* limit is reached.

Configuration type

Database

Parameter type

Configurable online

Default [range]

Automatic [16 - 60 000]

Unit of measure

Pages (4 KB)

When allocated

When an application associates with, or connects to, a database.

When freed

When the application disassociates or disconnects from the database.

Note: This parameter defines the maximum size of the application heap. One application heap is allocated per database application when the application first

connects with the database. The heap is shared by all database agents working for that application. (In previous releases, each database agent allocated its own application heap.) Memory is allocated from the application heap as needed to process the application, up to the limit specified by this parameter. When set to *AUTOMATIC*, the application heap is allowed to grow as needed up to either the *appl_memory* limit for the database, or the *instance_memory* limit for the database partition. The entire application heap is freed when the application disconnects with the database.

The online changed value takes effect at an application connection boundary, that is, after it is changed dynamically, currently connected applications still use the old value, but all newly connected applications will use the new value.

archretrydelay - Archive retry delay on error

This parameter specifies the number of seconds to wait after a failed archive attempt before trying to archive the log file again.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Default [range]

20 [0 - 65 535]

Subsequent retries will only take affect if the value of the *numarchretry* database configuration parameter is at least 1.

auto_del_rec_obj - Automated deletion of recovery objects configuration parameter

This parameter specifies whether database log files, backup images, and load copy images should be deleted when their associated recovery history file entry is pruned.

Configuration type

Database

Parameter type

Configurable online

Propagation class

Immediate

Default [range]

OFF [ON; OFF]

You can prune the entries in the recovery history file using the *PRUNE HISTORY* command or the *db2Prune* API. You can also configure the IBM Data Server database manager to automatically prune the recovery history file after each full database backup. If you set the *auto_del_rec_obj* database configuration parameter to *ON*, then the database manager will also delete the corresponding physical log

files, backup images, and load copy images when it prunes the history file. The database manager can only delete recovery objects such as database logs, backup images, and load copy images when your storage media is disk, or if you are using a storage manager, such as the Tivoli Storage Manager.

auto_maint - Automatic maintenance

This parameter is the parent of all the other automatic maintenance database configuration parameters (*auto_db_backup*, *auto_tbl_maint*, *auto_runstats*, *auto_stats_prof*, *auto_stmt_stats*, *auto_prof_upd*, and *auto_reorg*).

Configuration type

Database

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

ON [ON; OFF]

When this parameter is disabled, all of its child parameters are also disabled, but their settings, as recorded in the database configuration file, do not change. When this parent parameter is enabled, recorded values for its child parameters take effect. In this way, automatic maintenance can be enabled or disabled globally.

By default, this parameter is set to ON.

You can enable or disable individual automatic maintenance features independently by setting the following parameters:

auto_db_backup

This automated maintenance parameter enables or disables automatic backup operations for a database. A backup policy (a defined set of rules or guidelines) can be used to specify the automated behavior. The objective of the backup policy is to ensure that the database is being backed up regularly. The backup policy for a database is created automatically when the DB2 Health Monitor first runs. By default, this parameter is set to OFF. To be enabled, this parameter must be set to ON, and its parent parameter must also be enabled.

auto_tbl_maint

This parameter is the parent of all table maintenance parameters (*auto_runstats*, *auto_stats_prof*, *auto_prof_upd*, and *auto_reorg*). When this parameter is disabled, all of its child parameters are also disabled, but their settings, as recorded in the database configuration file, do not change. When this parent parameter is enabled, recorded values for its child parameters take effect. In this way, table maintenance can be enabled or disabled globally.

By default, this parameter is set to ON.

auto_runstats

This automated table maintenance parameter enables or disables automatic table runstats operations for a database. A runstats policy (a defined set of rules or guidelines) can be used to specify the automated behavior. Statistics collected by the runstats utility are used by the optimizer to determine the most efficient plan for accessing the physical data. To be enabled, this parameter must be set to On, and its parent parameters must also be enabled.

By default, this parameter is set to ON.

auto_stats_prof

When enabled, this automated table maintenance parameter turns on statistical profile generation, designed to improve applications whose workloads include complex queries, many predicates, joins, and grouping operations over several tables. To be enabled, this parameter must be set to ON, and its parent parameters must also be enabled.

By default, this parameter is set to OFF.

auto_stmt_stats

This parameter enables and disables the collection of real-time statistics. It is a child of the *auto_runstats* configuration parameter. This feature is enabled only if the parent, *auto_runstats* configuration parameter, is also enabled. For example, to enable *auto_stmt_stats*, set *auto_maint*, *auto_tbl_maint*, and *auto_runstats* to ON. To preserve the child value, the *auto_runstats* configuration parameter can be ON while the *auto_maint* configuration parameter is OFF. The corresponding Auto Runstats feature will still be OFF.

Assuming that both Auto Runstats and Auto Reorg are enabled, the settings are as follows:

Automatic maintenance	(AUTO_MAINT) = ON
Automatic database backup	(AUTO_DB_BACKUP) = OFF
Automatic table maintenance	(AUTO_TBL_MAINT) = ON
Automatic runstats	(AUTO_RUNSTATS) = ON
Automatic statement statistics	(AUTO_STMT_STATS) = OFF
Automatic statistics profiling	(AUTO_STATS_PROF) = OFF
Automatic profile updates	(AUTO_PROF_UPD) = OFF
Automatic reorganization	(AUTO_REORG) = ON

You can disable both Auto Runstats and Auto Reorg features temporarily by setting *auto_tbl_maint* to OFF. Both features can be enabled later by setting *auto_tbl_maint* back to ON. You do not need to issue db2stop or db2start commands to have the changes take effect.

By default, this parameter is set to OFF.

auto_prof_upd

When enabled, this automated table maintenance parameter (a child of *auto_stats_prof*) specifies that the runstats profile is to be updated with recommendations. When this parameter is disabled, recommendations are stored in the *opt_feedback_ranking* table, which you can inspect when manually updating the runstats profile. To be enabled, this parameter must be set to ON, and its parent parameters must also be enabled.

By default, this parameter is set to OFF.

auto_reorg

This automated table maintenance parameter enables or disables automatic table and index reorganization for a database. A reorganization policy (a

defined set of rules or guidelines) can be used to specify the automated behavior. To be enabled, this parameter must be set to ON, and its parent parameters must also be enabled.

By default, this parameter is set to OFF.

autorestart - Auto restart enable

This parameter determines whether the database manager can, in the event of an abnormal termination of the database, automatically call the restart database utility when an application connects to a database.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

On [On; Off]

The restart database utility performs a *Crash recovery* if the database terminated abnormally (because, for example, of a power failure or a system software failure) while applications were connected to it. It applies any committed transactions that were in the database buffer pool but were not written to disk at the time of the failure. It also backs out any uncommitted transactions that might have been written to disk.

If *autorestart* is not enabled, then an application that attempts to connect to a database which needs to have crash recovery performed (needs to be restarted) will receive a SQL1015N error. In this case, the application can call the restart database utility, or you can restart the database by selecting the restart operation of the recovery tool.

avg_appls - Average number of active applications

This parameter is used by the query optimizer to help estimate how much buffer pool will be available at run-time for the access plan chosen.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Statement boundary

Default [range]

Automatic [1 – maxappls]

Unit of measure

Counter

Recommendation: When running DB2 in a multi-user environment, particularly with complex queries and a large buffer pool, you might want the query optimizer to know that multiple query users are using your system so that the optimizer should be more conservative in assumptions of buffer pool availability.

When setting this parameter, you should estimate the number of complex query applications that typically use the database. This estimate should exclude all light OLTP applications. If you have trouble estimating this number, you can multiply the following:

- An average number of all applications running against your database. The database system monitor can provide information about the number of applications at any given time and using a sampling technique, you can calculate an average over a period of time. The information from the database system monitor includes both OLTP and non-OLTP applications.
- Your estimate of the percentage of complex query applications.

As with adjusting other configuration parameters that affect the optimizer, you should adjust this parameter in small increments. This allows you to minimize path selection differences.

You should consider rebinding applications (using the REBIND PACKAGE command) after changing this parameter.

backup_pending - Backup pending indicator

This parameter indicates whether you need to do a full backup of the database before accessing it.

Configuration type

Database

Parameter type

Informational

This parameter is only on if the database configuration is changed so that the database moves from being nonrecoverable to recoverable (that is, initially both the *logretain* and *userexit* parameters were set to NO, then either one or both of these parameters is set to YES, and the update to the database configuration is accepted).

blk_log_dsk_ful - Block on log disk full

This parameter can be set to prevent disk full errors from being generated when DB2 cannot create a new log file in the active log path.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

No [Yes; No]

Instead of generating a disk full error, DB2 will attempt to create the log file every five minutes until it succeeds. After each attempt, DB2 writes a message to the Administration Notification log. The only way that you can confirm that your application is hanging because of a log disk full condition is to monitor the Administration Notification log. Until the log file is successfully created, any user application that attempts to update table data will not be able to commit transactions. Read-only queries might not be directly affected; however, if a query

needs to access data that is locked by an update request, or a data page that is fixed in the buffer pool by the updating application, read-only queries will also appear to hang.

Setting *blk_log_dsk_ful* to *yes* causes applications to hang when DB2 encounters a log disk full error, thus allowing you to resolve the error and allowing the transaction to complete. You can resolve a disk full situation by moving old log files to another file system or by enlarging the file system, so that hanging applications can complete.

If *blk_log_dsk_ful* is set to *no*, then a transaction that receives a log disk full error will fail and will be rolled back. In some situations, the database will come down if a transaction causes a log disk full error.

catalogcache_sz - Catalog cache size

This parameter specifies the maximum space in pages that the catalog cache can use from the database heap.

Configuration type

Database

Parameter type

Configurable online

Propagation class

Immediate

Default [range]

-1 [MAXAPPLS*5]

Unit of measure

Pages (4 KB)

When allocated

When the database is initialized

When freed

When the database is shut down

This parameter is allocated out of the database shared memory, and is used to cache system catalog information. In a partitioned database system, there is one catalog cache for each database partition.

Caching catalog information at individual database partitions allows the database manager to reduce its internal overhead by eliminating the need to access the system catalogs (or the catalog node in a partitioned database environment) to obtain information that has previously been retrieved. The use of the catalog cache can help improve the overall performance of:

- Binding packages and compiling SQL and XQuery statements
- Operations that involve checking database-level privileges, routine privileges, global variable privileges and role authorizations
- Applications that are connected to non-catalog nodes in a partitioned database environment

By taking the default (-1) in a server or partitioned database environment, the value used to calculate the page allocation is five times the value specified for the

maxappls configuration parameter. The exception to this occurs if five times *maxappls* is less than 8. In this situation, the default value of -1 will set *catalogcache_sz* to 8.

Recommendation: Start with the default value and tune it by using the database system monitor. When tuning this parameter, you should consider whether the extra memory being reserved for the catalog cache might be more effective if it was allocated for another purpose, such as the buffer pool or package cache.

Tuning this parameter is particularly important if a workload involves many SQL or XQuery compilations for a brief period of time, with few or no compilations thereafter. If the cache is too large, memory might be wasted holding copies of information that will no longer be used.

In an partitioned database environment, consider if the *catalogcache_sz* at the catalog node needs to be set larger since catalog information that is required at non-catalog nodes will always first be cached at the catalog node.

The *cat_cache_lookups* (catalog cache lookups), *cat_cache_inserts* (catalog cache inserts), *cat_cache_overflows* (catalog cache overflows), and *cat_cache_size_top* (catalog cache high water mark) monitor elements can help you determine whether you should adjust this configuration parameter.

Note: The catalog cache exists on all nodes in a partitioned database environment. Since there is a local database configuration file for each node, each node's *catalogcache_sz* value defines the size of the local catalog cache. In order to provide efficient caching and avoid overflow scenarios, you need to explicitly set the *catalogcache_sz* value at each node and consider the feasibility of possibly setting the *catalogcache_sz* on non-catalog nodes to be smaller than that of the catalog node; keep in mind that information that is required to be cached at non-catalog nodes will be retrieved from the catalog node's cache. Hence, a catalog cache at a non-catalog node is like a subset of the information in the catalog cache at the catalog node.

In general, more cache space is required if a unit of work contains several dynamic SQL or XQuery statements or if you are binding packages that contain a large number of static SQL or XQuery statements.

chngpgs_thresh - Changed pages threshold

This parameter specifies the level (percentage) of changed pages at which the asynchronous page cleaners will be started, if they are not currently active.

Configuration type
Database

Parameter type
Configurable

Default [range]
60 [5 - 99]

Unit of measure
Percentage

Asynchronous page cleaners will write changed pages from the buffer pool (or the buffer pools) to disk before the space in the buffer pool is required by a database agent. As a result, database agents should not have to wait for changed pages to

be written out so that they might use the space in the buffer pool. This improves overall performance of the database applications.

When the page cleaners are started, they will build a list of the pages to write to disk. Once they have completed writing those pages to disk, they will become inactive again and wait for the next trigger to start.

When the `DB2_USE_ALTERNATE_PAGE_CLEANING` registry variable is set (that is, the alternate method of page cleaning is used), the `chnpggs_thresh` parameter has no effect, and the database manager automatically determines how many dirty pages to maintain in the buffer pool.

Recommendation: For databases with a heavy update transaction workload, you can generally ensure that there are enough clean pages in the buffer pool by setting the parameter value to be equal-to or less-than the default value. A percentage larger than the default can help performance if your database has a small number of very large tables.

codepage - Code page for the database

This parameter shows the code page that was used to create the database. The `codepage` parameter is derived based on the `codeset` parameter.

Configuration type

Database

Parameter type

Informational

codeset - Codeset for the database

This parameter shows the codeset that was used to create the database. Codeset is used by the database manager to determine `codepage` parameter values.

Configuration type

Database

Parameter type

Informational

collate_info - Collating information

This parameter determines the database's collating sequence. For a language-aware collation, the first 256 bytes contain the string representation of the collation name (for example, "SYSTEM_819_US").

This parameter can only be displayed using the `db2CfgGet` API. It **cannot** be displayed through the command line processor or the Control Center.

Configuration type

Database

Parameter type

Informational

This parameter provides 260 bytes of database collating information. The first 256 bytes specify the database collating sequence, where byte "n" contains the sort weight of the code point whose underlying decimal representation is "n" in the code page of the database.

The last 4 bytes contain internal information about the type of the collating sequence. The last four bytes of the parameter is an integer. The integer is sensitive to the endian order of the platform. The possible values are:

- **0** – The sequence contains non-unique weights
- **1** – The sequence contains all unique weights
- **2** – The sequence is the identity sequence, for which strings are compared byte for byte.
- **3** – The sequence is NLSCHAR, used for sorting characters in a TIS620-1 (code page 874) Thai database.
- **4** – The sequence is IDENTITY_16BIT, which implements the “CESU-8 Compatibility Encoding Scheme for UTF-16: 8-bit” algorithm as specified in the Unicode Technical Report #26 available at the Unicode Technical Consortium Web site at <http://www.unicode.org>
- **X'8001'** – The sequence is UCA400_NO, which implements the Unicode Collation[®] Algorithm (UCA) based on the Unicode Standard version 4.00, with normalization implicitly set to ON.
- **X'8002'** – The sequence is UCA400_LTH, which implements the Unicode Collation Algorithm (UCA) based on the Unicode Standard version 4.00, and sorts all Thai characters as per the Royal Thai Dictionary order.
- **X'8003'** – The sequence is UCA400_LSK, which implements the Unicode Collation Algorithm (UCA) based on the Unicode Standard version 4.00, and sorts all Slovakian characters properly.

Note: For a language-aware collation, the first 256 bytes contain the string representation of the collation name.

If you use this internal type information, you need to consider byte reversal when retrieving information for a database on a different platform.

You can specify the collating sequence at database creation time.

country/region - Database territory code

This parameter shows the *territory* code used to create the database.

Configuration type
Database

Parameter type
Informational

database_consistent - Database is consistent

This parameter indicates whether the database is in a consistent state.

Configuration type
Database

Parameter type
Informational

YES indicates that all transactions have been committed or rolled back so that the data is consistent. If the system “crashes” while the database is consistent, you do not need to take any special action to make the database usable.

NO indicates that a transaction is pending or some other task is pending on the database and the data is not consistent at this point. If the system “crashes” while

the database is not consistent, you will need to restart the database using the `RESTART DATABASE` command to make the database usable.

database_level - Database release level

This parameter indicates the release level of the database manager which can use the database.

Configuration type

Database

Parameter type

Informational

In the case of an incomplete or failed migration, this parameter will reflect the release level of the unmigrated database and might differ from the *release* parameter (the release level of the database configuration file). Otherwise the value of *database_level* will be identical to value of the *release* parameter.

database_memory - Database shared memory size

This parameter specifies the amount of memory that is reserved for the database shared memory region. If this amount is less than the amount calculated from the individual memory parameters (for example, locklist, utility heap, bufferpools, and so on), the larger amount will be used.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable online

Default [range]

Automatic [Computed, 0 - 4 294 967 295]

Unit of measure

Pages (4 KB)

When allocated

When the database is activated

When freed

When the database is deactivated

Setting this parameter to `AUTOMATIC` enables self-tuning. When enabled, the memory tuner determines the overall memory requirements for the database and increases or decreases the amount of memory allocated for database shared memory depending on the current database requirements. For example, if the current database requirements are high, and there is sufficient free memory on the system, more memory will be consumed by database shared memory. Once the database memory requirements drop, or the amount of free memory on the system drops too low, some database shared memory is released.

The memory tuner will always leave a minimum amount of memory free based on the calculated benefit to providing additional memory to the instance. If there is a great benefit to providing an instance with more memory, then the memory tuner

will maintain a lower amount of free memory. If the benefit is lower, then more free memory will be maintained. This allows databases to cooperate in the distribution of system memory.

Because the memory tuner trades memory resources between different memory consumers, there must be at least two memory consumers enabled for self-tuning to be active.

Automatic tuning of this configuration parameter will only occur when self-tuning memory is enabled for the database (the **self_tuning_mem** configuration parameter is set to ON).

To simplify the management of this parameter, the COMPUTED setting instructs the database manager to calculate the amount of memory needed, and to allocate it at database activation time. The database manager will also allocate some additional memory to satisfy peak memory requirements for any heap in the database shared memory region whenever a heap exceeds its configured size. Other operations, such as dynamic configuration updates, also have access to this additional memory. The db2pd command, with the **-memsets** option, can be used to monitor the amount of unused memory left in the database shared memory region.

Recommendation: This value will usually remain at AUTOMATIC. For environments that do not support the AUTOMATIC setting, this should be set to COMPUTED. For example, the additional memory can be used for creating new buffer pools, or for increasing the size of existing buffer pools.

Note: In Version 9.5, when you set **database_memory** configuration parameter to AUTOMATIC, the initial database shared memory allocation is the configured size of all heaps and buffer pools defined for the database, and the memory increases as needed. If **database_memory** is set to a specific value, then that requested amount of memory is allocated initially, during database activation. If the initial amount of memory cannot be allocated from the operating system, or exceeds the **instance_memory** limit, database activation fails with an SQL1084C error (Shared memory segments cannot be allocated).

As of DB2 Version 9.5 Fix Pack 2, if you set **database_memory** to AUTOMATIC on Solaris Operating Environment, the database manager uses pageable memory for the database shared memory. This use of smaller memory pages might result in some performance degradation. In Solaris operating systems on UltraSPARC, the database manager attempts to use 64 KB memory pages if they are available. If 64 KB memory pages are not available, the database manager will use 8 KB memory pages. In Solaris operating systems on Sun x64 systems, the database manager will use 4 KB memory pages. To maintain the use of large pages in shared memory on Solaris, set **database_memory** to COMPUTED or a numeric value.

Controlling DB2 Memory consumption:

When **instance_memory** is set to AUTOMATIC, a fixed upper bound on total memory consumption for the instance is set at instance startup (db2start). Actual memory consumption by the database manager varies depending on the workload. When self-tuning memory manager is enabled to perform **database_memory** tuning (by default for new databases), during run-time, self-tuning memory manager dynamically updates the size of performance-critical heaps within the database shared memory set according to the free physical memory on the system, while ensuring that

there is sufficient free **instance_memory** available for functional memory requirements. For more information, see the **instance_memory** configuration parameter.

Limitation on some Linux¹ kernels:

Due to operating system limitations on some Linux kernels, self-tuning memory manager currently does not allow setting **database_memory** to AUTOMATIC. However, this setting is now allowed on these kernels only when **instance_memory** is set to a specific value, and not AUTOMATIC. If **database_memory** is set to AUTOMATIC, and **instance_memory** is later set back to AUTOMATIC, the **database_memory** configuration parameter is automatically updated to COMPUTED during the next database activation. If some databases are already active, self-tuning memory manager stops tuning the overall **database_memory** sizes.

¹On Linux, this parameter supports the AUTOMATIC setting on RHEL5 and on SUSE 10 SP1 and newer. All other validated Linux distributions will return to COMPUTED if the kernel does not support this feature.

db_mem_thresh - Database memory threshold

This parameter represents the maximum percentage of committed, but currently unused, database shared memory that the database manager will allow before starting to release committed pages of memory back to the operating system.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

10 [0–100]

Unit of measure

Percentage

This database configuration parameter relates to how the database manager handles excess unused database shared memory. Typically, as pages of memory are touched by a process, they are committed, meaning that a page of memory has been allocated by the operating system and occupies space either in physical memory or in a page file on disk. Depending on the database workload, there might be peak database shared memory requirements at a certain times of day. Once the operating system has enough committed memory to meet those peak requirements, that memory remains committed, even after peak memory requirements have subsided.

Acceptable values are whole numbers in the range of 0 (immediately release any unused database shared memory) to 100 (never release any unused database shared memory). The default is 10 (release unused memory only when more than 10% of database shared memory is currently unused), which should be suitable for most workloads.

This configuration parameter can be updated dynamically. Care should be taken when updating this parameter, as setting the value too low could cause excessive memory thrashing on the box (memory pages constantly being committed and

then released), and setting the value too high might prevent the database manager from returning any database shared memory back to the operating system for other processes to use.

This configuration parameter will be ignored (meaning that unused database shared memory pages will remain committed) if the database shared memory region is pinned through the `DB2_PINNED_BP` registry variable, configured for large pages through the `DB2_LARGE_PAGE_MEM` registry variable, or if releasing of memory is explicitly disabled through the `DB2MEMDISCLAIM` registry variable.

Some versions of Linux do not support releasing subranges of a shared memory segment back to the operating system. On such platforms, this parameter will be ignored.

dbheap - Database heap

This parameter determines the maximum memory used by the database heap.

With Version 9.5, this database configuration parameter has a default value of `AUTOMATIC`, meaning that the database heap can increase as needed until either the `database_memory` limit is reached, or the `instance_memory` limit is reached.

Configuration type

Database

Parameter type

Configurable online

Propagation class

Immediate

Default [range]

Automatic [32 - 524 288]

Unit of measure

Pages (4 KB)

When allocated

When the database is activated

When freed

When the database is deactivated

There is one database heap per database, and the database manager uses it on behalf of all applications connected to the database. It contains control block information for tables, indexes, table spaces, and buffer pools. It also contains space for the log buffer (`logbufsz`) and temporary memory used by utilities. Therefore, the size of the heap will be dependent on a large number of variables. The control block information is kept in the heap until all applications disconnect from the database.

The minimum amount the database manager needs to get started is allocated at the first connection. The data area is expanded as needed until either the configured upper limit is reached, or, if set to `AUTOMATIC`, until all `database_memory` or `instance_memory`, or memory for both, is exhausted.

The following formula can be used as a rough guideline when deciding on a value to assign to the `dbheap` configuration parameter:

10K per tablespace + 4K per table + (1K + 4*extents used),
per range clustered table (RCT)

The *dbheap* value that you configure represents only a portion of the database heap that is allocated. The database heap is the main memory area used to satisfy global database memory requirements. It is sized to include basic allocations needed for the startup of a database in addition to the *dbheap* value. Tools which report memory usage such as Memory Tracker, Snapshot Monitor, and db2pd report the statistics of this larger database heap. There is no separate tracking of the allocations that are represented by the *dbheap* configuration parameter. Therefore, it is normal for the statistics on database heap memory usage reported from these tools to exceed the configured value for the *dbheap* parameter.

You can use the database system monitor to track the highest amount of memory that was used for the database heap, using the *db_heap_top* (maximum database heap allocated) element.

Note:

- Workload Management (WLM) work class sets and work action sets are stored in the database heap. However, a very small part of the memory is consumed for this.
- Trusted contexts, Workload Management, and Audit policy information is cached in memory for fast processing. This memory is allocated from the database heap. Therefore, user-defined trusted contexts, workload management, and audit policy objects would impose more memory requirements on the database heap. In this case, it is suggested that you set your database heap configuration to AUTOMATIC so that the database manager automatically manages the database heap size.

decflt_rounding - Decimal floating point rounding configuration parameter

This parameter allows you to specify the rounding mode for decimal floating point (DECFLOAT). The rounding mode affects decimal floating-point operations in the server, and in LOAD.

Configuration type

Database

Parameter type

Configurable

See “Consequences of changing decflt_rounding” on page 535 below.

Default [range]

ROUND_HALF_EVEN [ROUND_CEILING, ROUND_FLOOR, ROUND_HALF_UP, ROUND_DOWN]

DB2 supports five IEEE-compliant decimal floating point rounding modes. The rounding mode specifies how to round the result of a calculation when the result exceeds the precision. The definitions for all the rounding modes are as follows:

ROUND_CEILING

Round towards +infinity. If all of the discarded digits are zero or if the sign is negative the result is unchanged. Otherwise, the result coefficient should be incremented by 1 (rounded up).

ROUND_FLOOR

Round towards -infinity. If all of the discarded digits are zero or if the sign is positive the result is unchanged. Otherwise, the sign is negative and the result coefficient should be incremented by 1.

ROUND_HALF_UP

Round to nearest; if equidistant, round up 1. If the discarded digits represent greater than or equal to half (0.5) of the value of a 1 in the next left position then the result coefficient should be incremented by 1 (rounded up). Otherwise, the discarded digits (0.5 or less) are ignored.

ROUND_HALF_EVEN

Round to nearest; if equidistant, round so that the final digit is even. If the discarded digits represent greater than half (0.5) the value of a one in the next left position, then the result coefficient should be increment by 1 (rounded up). If they represent less than half, then the result coefficient is not adjusted, that is, the discarded digits are ignored. Otherwise, if they represent exactly half, the result coefficient is unaltered if its rightmost digit is even, or incremented by 1 (rounded up) if its rightmost digit is odd, to make an even digit. This rounding mode is the default rounding mode as per IEEE decimal floating point specification and is the default rounding mode in DB2 products.

ROUND_DOWN

Round towards 0 (truncation). The discarded digits are ignored.

Table 72 shows the result of rounding of 12.341, 12.345, 12.349, 12.355, and -12.345, each to 4 digits, under different rounding modes:

Table 72. Decimal floating point rounding modes

Rounding mode	12.341	12.345	12.349	12.355	-12.345
ROUND_DOWN	12.34	12.34	12.34	12.35	-12.34
ROUND_HALF_UP	12.34	12.35	12.35	12.36	-12.35
ROUND_HALF_EVEN	12.34	12.34	12.35	12.36	-12.34
ROUND_FLOOR	12.34	12.34	12.34	12.35	-12.35
ROUND_CEILING	12.35	12.35	12.35	12.36	-12.34

Consequences of changing decflt_rounding

- Previously constructed materialized query tables (MQTs) could contain results that differ from what would be produced with the new rounding mode. To correct this problem, refresh potentially impacted MQTs.
- The results of a trigger may be affected by the new rounding mode. Changing it has no effect on data that has already been written.
- Constraints that allowed data to be inserted into a table, if reevaluated, might reject that same data. Similarly constraints that did not allow data to be inserted into a table, if reevaluated, might accept that same data. Use the SET INTEGRITY statement to check for and correct such problems. The value of a generated column whose calculation is dependent on decflt_rounding could be different for two identical rows except for the generated column value, if one row was inserted before the change to decflt_rounding and the other was inserted after.
- The rounding mode is not compiled into sections. Therefore, static SQL does not need to be recompiled after changing *decflt_rounding*.

Note: The value of this configuration parameter is not changed dynamically but will become effective only after all applications disconnect from the database. If the database is activated, it must be deactivated.

dft_degree - Default degree

This parameter specifies the default value for the CURRENT DEGREE special register and the DEGREE bind option.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Connection

Default [range]

1 [-1(ANY), 1 - 32 767]

The default value is 1.

A value of 1 means no intra-partition parallelism. A value of -1 (or ANY) means the optimizer determines the degree of intra-partition parallelism based on the number of processors and the type of query.

The degree of intra-partition parallelism for an SQL statement is specified at statement compilation time using the CURRENT DEGREE special register or the DEGREE bind option. The maximum runtime degree of intra-partition parallelism for an active application is specified using the SET RUNTIME DEGREE command. The Maximum Query Degree of Parallelism (*max_querydegree*) configuration parameter specifies the maximum query degree of intra-partition parallelism for all SQL queries.

The actual runtime degree used is the lowest of:

- *max_querydegree* configuration parameter
- application runtime degree
- SQL statement compilation degree

dft_extent_sz - Default extent size of table spaces

This parameter sets the default extent size of table spaces.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

32 [2 – 256]

Unit of measure

Pages

When a table space is created, EXTENTSIZE n can be optionally specified, where n is the extent size. If you do not specify the extent size on the CREATE TABLESPACE statement, the database manager uses the value given by this parameter.

Recommendation: In many cases, you will want to explicitly specify the extent size when you create the table space. Before choosing a value for this parameter, you should understand how you would explicitly choose an extent size for the CREATE TABLESPACE statement.

dft_loadrec_ses - Default number of load recovery sessions

This parameter specifies the default number of sessions that will be used during the recovery of a table load.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

1 [1 - 30 000]

Unit of measure

Counter

The value should be set to an optimal number of I/O sessions to be used to retrieve a load copy. The retrieval of a load copy is an operation similar to restore. You can override this parameter through entries in the copy location file specified by the environment variable DB2LOADREC.

The default number of buffers used for load retrieval is two more than the value of this parameter. You can also override the number of buffers in the copy location file.

This parameter is applicable only if roll forward recovery is enabled.

dft_mttb_types - Default maintained table types for optimization

This parameter specifies the default value for the CURRENT MAINTAINED TABLE TYPES FOR OPTIMIZATION special register. The value of this register determines what types of refresh deferred materialized query tables will be used during query optimization.

Configuration type

Database

Parameter type

Configurable

Default [range]

SYSTEM [ALL, NONE, FEDERATED_TOOL, SYSTEM, USER, or a list of values]

You can specify a list of values separated by commas; for example, 'USER,FEDERATED_TOOL'. ALL or NONE cannot be listed with other values, and you cannot specify the same value more than once. For use with the db2CfgSet and db2CfgGet APIs, the acceptable parameter values are: 8 (ALL), 4 (NONE), 16 (FEDERATED_TOOL), 1 (SYSTEM) and 2 (USER). Multiple values can be specified

together using bitwise OR; for example, 18 would be the equivalent of `USER,FEDERATED_TOOL`. As before, the values 4 and 8 cannot be used with other values.

dft_prefetch_sz - Default prefetch size

This parameter sets the default prefetch size of table spaces.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

Automatic [0 - 32 767]

Unit of measure

Pages

When a table space is created, `PREFETCHSIZE n` can optionally be specified, where *n* represents the number of pages the database manager will read if prefetching is being performed. If you do not specify the prefetch size on invocation of the `CREATE TABLESPACE` statement, the database manager uses the current value of the `dft_prefetch_sz` parameter.

If a table space is created with `AUTOMATIC DFT_PREFETCH_SZ`, the prefetch size of the table space will become `AUTOMATIC`, which means that DB2 will automatically calculate and update the prefetch size of the table space, using the following equation:

$$\text{prefetch size} = (\# \text{ containers}) * (\# \text{ physical spindles}) * \text{extent size}$$

where the number of physical spindles defaults to 1 and can be specified through the DB2 registry variable `DB2_PARALLEL_IO`. This calculation is performed:

- At database start-up time
- When a table space is first created with `AUTOMATIC` prefetch size
- When the number of containers for a table space changes through execution of an `ALTER TABLESPACE` statement
- When the prefetch size for a table space is updated to be `AUTOMATIC` through execution of an `ALTER TABLESPACE` statement

The `AUTOMATIC` state of the prefetch size can be turned on or off as soon as the prefetch size is updated manually through invocation of the `ALTER TABLESPACE` statement.

Recommendation: Using system monitoring tools, you can determine if your CPU is idle while the system is waiting for I/O. Increasing the value of this parameter can help if the table spaces being used do not have a prefetch size defined for them.

This parameter provides the default for the entire database, and it might not be suitable for all table spaces within the database. For example, a value of 32 might be suitable for a table space with an extent size of 32 pages, but not suitable for a table space with an extent size of 25 pages. Ideally, you should explicitly set the prefetch size for each table space.

To help minimize I/O for table spaces defined with the default extent size (*dft_extent_sz*), you should set this parameter as a factor or whole multiple of the value of the *dft_extent_sz* parameter. For example, if the *dft_extent_sz* parameter is 32, you could set *dft_prefetch_sz* to 16 (a fraction of 32) or to 64 (a whole multiple of 32). If the prefetch size is a multiple of the extent size, the database manager might perform I/O in parallel, if the following conditions are true:

- The extents being prefetched are on different physical devices
- Multiple I/O servers are configured (*num_ioservers*).

dft_queryopt - Default query optimization class

The query optimization class is used to direct the optimizer to use different degrees of optimization when compiling SQL and XQuery queries. This parameter provides additional flexibility by setting the default query optimization class used when neither the SET CURRENT QUERY OPTIMIZATION statement nor the QUERYOPT option on the bind command are used.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Connection

Default [range]

5 [0 — 9]

Unit of measure

Query Optimization Class (see below)

The query optimization classes currently defined are:

- 0 - minimal query optimization.
- 1 - roughly comparable to DB2 Version 1.
- 2 - slight optimization.
- 3 - moderate query optimization.
- 5 - significant query optimization with heuristics to limit the effort expended on selecting an access plan. This is the default.
- 7 - significant query optimization.
- 9 - maximal query optimization

dft_refresh_age - Default refresh age

This parameter represents the maximum time duration since a REFRESH TABLE statement has been processed on a specific REFRESH DEFERRED materialized query table. After this time limit is exceeded, the materialized query table is not used to satisfy queries until the materialized query table is refreshed.

Configuration type

Database

Parameter type

Configurable

Default [range]

0 [0, 999999999999999 (ANY)]

This parameter has the default value used for the REFRESH AGE if the CURRENT REFRESH AGE special register is not specified. This parameter specifies a time stamp duration value with a data type of DECIMAL(20,6). If the CURRENT REFRESH AGE has a value of 9999999999999999 (ANY), and the QUERY OPTIMIZATION class has a value of two, or five or more, REFRESH DEFERRED materialized query tables are considered to optimize the processing of a dynamic query.

dft_sqlmathwarn - Continue upon arithmetic exceptions

This parameter sets the default value that determines the handling of arithmetic errors and retrieval conversion errors as errors or warnings during SQL statement compilation.

Configuration type

Database

Parameter type

Configurable

Default [range]

No [No, Yes]

For static SQL statements, the value of this parameter is associated with the package at bind time. For dynamic SQL DML statements, the value of this parameter is used when the statement is prepared.

Attention: If you change the *dft_sqlmathwarn* value for a database, the behavior of check constraints, triggers, and views that include arithmetic expressions might change. This might, in turn, have an impact on the data integrity of the database. You should only change the setting of *dft_sqlmathwarn* for a database after carefully evaluating how the new arithmetic exception handling behavior might impact check constraints, triggers, and views. Once changed, subsequent changes require the same careful evaluation.

As an example, consider the following check constraint, which includes a division arithmetic operation:

```
A/B > 0
```

When *dft_sqlmathwarn* is “No” and an INSERT with B=0 is attempted, the division by zero is processed as an arithmetic error. The insert operation fails because DB2 cannot check the constraint. If *dft_sqlmathwarn* is changed to “Yes”, the division by zero is processed as an arithmetic warning with a NULL result. The NULL result causes the predicate to evaluate to UNKNOWN and the insert operation succeeds. If *dft_sqlmathwarn* is changed back to “No”, an attempt to insert the same row will fail, because the division by zero error prevents DB2 from evaluating the constraint. The row inserted with B=0 when *dft_sqlmathwarn* was “Yes” remains in the table and can be selected. Updates to the row that cause the constraint to be evaluated will fail, while updates to the row that do not require constraint re-evaluation will succeed.

Before changing *dft_sqlmathwarn* from “No” to “Yes”, you should consider rewriting the constraint to explicitly handle nulls from arithmetic expressions. For example:

```
( A/B > 0 ) AND ( CASE  
                    WHEN A IS NULL THEN 1  
                    WHEN B IS NULL THEN 1
```



```

        WHEN A/B IS NULL THEN 0
        ELSE 1
        END
    = 1 )

```

can be used if both A and B are nullable. And, if A or B is not-nullable, the corresponding IS NULL WHEN-clause can be removed.

Before changing *dft_sqlmathwarn* from “Yes” to “No”, you should first check for data that might become inconsistent by using, for example, predicates such as the following:

```
WHERE A IS NOT NULL AND B IS NOT NULL AND A/B IS NULL
```

When inconsistent rows are isolated, you should take appropriate action to correct the inconsistency before changing *dft_sqlmathwarn*. You can also manually re-check constraints with arithmetic expressions after the change. To do this, first place the affected tables in a check pending state (with the OFF clause of the SET CONSTRAINTS statement), then request that the tables be checked (with the IMMEDIATE CHECKED clause of the SET CONSTRAINTS statement). Inconsistent data will be indicated by an arithmetic error, which prevents the constraint from being evaluated.

Recommendation: Use the default setting of no, unless you specifically require queries to be processed that include arithmetic exceptions. Then specify the value of yes. This situation can occur if you are processing SQL statements that, on other database managers, provide results regardless of the arithmetic exceptions that occur.

discover_db - Discover database

This parameter is used to prevent information about a database from being returned to a client when a discovery request is received at the server.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

Enable [Disable, Enable]

The default for this parameter is that discovery is enabled for this database.

By changing this parameter value to “Disable”, it is possible to hide databases with sensitive data from the discovery process. This can be done in addition to other database security controls on the database.

dlchktime - Time interval for checking deadlock

This parameter defines the frequency at which the database manager checks for deadlocks among all the applications connected to a database.

Configuration type

Database

Parameter type
Configurable online

Propagation class
Immediate

Default [range]
10 000 (10 seconds) [1 000 - 600 000]

Unit of measure
Milliseconds

A deadlock occurs when two or more applications connected to the same database wait indefinitely for a resource. The waiting is never resolved because each application is holding a resource that the other needs to continue.

Note:

1. In a partitioned database environment, this parameter applies to the catalog node only.
2. In a partitioned database environment, a deadlock is not flagged until after the second iteration.

Recommendation: Increasing this parameter decreases the frequency of checking for deadlocks, thereby increasing the time that application programs must wait for the deadlock to be resolved.

Decreasing this parameter increases the frequency of checking for deadlocks, thereby decreasing the time that application programs must wait for the deadlock to be resolved but increasing the time that the database manager takes to check for deadlocks. If the deadlock interval is too small, it can decrease runtime performance, because the database manager is frequently performing deadlock detection. If this parameter is set lower to improve concurrency, you should ensure that *maxlocks* and *locklist* are set appropriately to avoid unnecessary lock escalation, which can result in more lock contention and as a result, more deadlock situations.

dyn_query_mgmt - Dynamic SQL and XQuery query management

This parameter determines whether Query Patroller will capture information about submitted queries.

Configuration type
Database

Parameter type
Configurable Online

Default [range]
0 (DISABLE) [1(ENABLE), 0 (DISABLE)]

This parameter is relevant where DB2 Query Patroller is installed. If this parameter is set to "ENABLE", Query Patroller captures information about the query, such as the submitter ID and the estimated cost of execution, as calculated by the optimizer. These values are used to determine whether the query should be managed by Query Patroller, based on user- and system-level thresholds.

If this parameter is set to "DISABLE", Query Patroller does not capture any information about submitted queries, and no query management takes place.

enable_xmlchar - Enable conversion to XML configuration parameter

This parameter determines whether XMLPARSE operations can be performed on non-BIT DATA CHAR (or CHAR-type) expressions in an SQL statement.

Configuration type

Database

Parameter type

Configurable

Default [range]

Yes [Yes; No]

When pureXML™ features are used in a non-Unicode database, the XMLPARSE function can cause character substitutions to occur as SQL string data is converted from the client code page into the database code page, and then into Unicode for internal storage. Setting *enable_xmlchar* to NO blocks the usage of character data types during XML parsing, and any attempts to insert character types into a non-Unicode database will generate an error. The BLOB data type and FOR BIT DATA data types are still allowed when *enable_xmlchar* is set to NO, as code page conversion does not occur when these data types are used to pass XML data into a database.

By default, *enable_xmlchar* is set to YES so that parsing of character data types is allowed. In this case, you should ensure that any XML data to be inserted contains only code points that are part of the database code page, in order to avoid substitution characters being introduced during insertion of the XML data.

Note: The client needs to disconnect and reconnect to the agent for this change to be reflected.

failarchpath - Failover log archive path

This parameter specifies a path to which DB2 will try to archive log files if the log files cannot be archived to either the primary or the secondary (if set) archive destinations because of a media problem affecting those destinations. This specified path must reference a disk.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Default [range]

Null []

If there are log files in the path specified by the current value of *failarchpath*, any updates to *failarchpath* will not take effect immediately. Instead, the update will take effect when all applications disconnect.

groupheap_ratio - Percent of memory for application group heap

This parameter is deprecated in Version 9.5, but is still being used by pre-Version 9.5 data servers and clients. Any value specified for this configuration parameter will be ignored by the DB2 Version 9.5 database manager. In Version 9.5, it has been replaced by the *appl_memory* configuration parameter..

Note: The following information applies only to pre-Version 9.5 data servers and clients.

This parameter specifies the percentage of memory in the application control shared memory set devoted to the application group shared heap.

Configuration type

Database

Parameter type

Configurable

Default [range]

70 [1 – 99]

Unit of measure

Percentage

This parameter does not have any effect on a non-partitioned database with concentrator OFF and intra-partition parallelism disabled.

Recommendation: Retain the default value of this parameter unless you are experiencing performance problems.

hadr_db_role - HADR database role

This parameter indicates the current role of a database, whether the database is online or offline.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Database server with local clients

Parameter type

Informational

Valid values are: STANDARD, PRIMARY, or STANDBY.

Note: When a database is active, the HADR role of the database can also be determined using the GET SNAPSHOT FOR DATABASE command.

hadr_local_host - HADR local host name

This parameter specifies the local host for high availability disaster recovery (HADR) TCP communication.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Database server with local clients

Parameter type

Configurable

Default

Null

Either a host name or an IP address can be used. If a host name is specified and it maps to multiple IP addresses, an error is returned, and HADR will not start up. If the host name maps to multiple IP addresses (even if you specify the same host name on primary and standby), primary and standby can end up mapping this host name to different IP addresses, because some DNS servers return IP address lists in non-deterministic order.

A host name is in the form: myserver.ibm.com. An IP address is in the form: "12.34.56.78".

hadr_local_svc - HADR local service name

This parameter specifies the TCP service name or port number for which the local high availability disaster recovery (HADR) process accepts connections.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Database server with local clients

Parameter type

Configurable

Default

Null

The value for `hadr_local_svc` on the Primary or Standby database systems cannot be the same as the value of `svcname` or `svcname +1` on their respective nodes.

hadr_peer_window - HADR peer window configuration parameter

When you set `hadr_peer_window` to a non-zero time value, then a HADR primary-standby database pair continues to behave as though still in peer state, for the configured amount of time, if the primary database loses connection with the standby database. This helps ensure data consistency.

Configuration type

Database

Parameter type

Configurable

Default [range]

0 [0 - 4 294 967 295]

Unit of measure

Seconds

Usage notes:

- The value will need to be the same on both primary and standby databases.
- A recommended minimum value is 120 seconds.
- The **hadr_peer_window** value is ignored when the **hadr_syncmode** value is set to ASYNC. That is, the value is treated as if it were set to zero (0), since it is not meaningful in ASYNC mode.
- To avoid impacting the availability of the primary database when the standby database is intentionally shut down, for example, for maintenance, the peer window is not invoked if the standby database is explicitly deactivated while the HADR pair is in peer state.
- The TAKEOVER HADR command with the **PEER WINDOW ONLY** option will launch a takeover operation only if the HADR standby is presently inside the defined peer window.
- The takeover operation with the **PEER WINDOW ONLY** option may behave incorrectly if the primary database clock and the standby database clock are not synchronized to within 5 seconds of each other. That is, the operation may succeed when it should fail, or fail when it should succeed. You should use a time synchronization service (for example, NTP) to keep the clocks synchronized to the same source.
- On the standby databases, the peer window end time is based on the last heartbeat message received from the primary database rather than disconnection. Therefore, the standby database's remaining time in S-DisconnectedPeer state before transition to S-RemoteCatchupPending ranges from (**hadr_peer_window** - **hadr_timeout**) seconds to (**hadr_peer_window**) seconds, depending on when and how the disconnection occurred.

hadr_remote_host - HADR remote host name

This parameter specifies the TCP/IP host name or IP address of the remote high availability disaster recovery (HADR) database server.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Database server with local clients

Parameter type

Configurable

Default

Null

Similar to *hadr_local_host*, this parameter must map to only one IP address.

hadr_remote_inst - HADR instance name of the remote server

This parameter specifies the instance name of the remote server. Administration tools, such as the DB2 Control Center, use this parameter to contact the remote server. High availability disaster recovery (HADR) also checks whether a remote database requesting a connection belongs to the declared remote instance.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Database server with local clients

Parameter type
Configurable

Default
Null

hadr_remote_svc - HADR remote service name

This parameter specifies the TCP service name or port number that will be used by the remote high availability disaster recovery (HADR) database server.

Configuration type
Database

Applies to

- Database server with local and remote clients
- Database server with local clients

Parameter type
Configurable

Default
Null

hadr_syncmode - HADR synchronization mode for log write in peer state

This parameter specifies the synchronization mode, which determines how primary log writes are synchronized with the standby when the systems are in peer state.

Configuration type
Database

Applies to

- Database server with local and remote clients
- Database server with local clients

Parameter type
Configurable

Default [range]
NEARSYNC [ASYNC; SYNC]

Valid values for this parameter are:

SYNC This mode provides the greatest protection against transaction loss, but at a higher cost of transaction response time.

In this mode, log writes are considered successful only when logs have been written to log files on the primary database and when the primary database has received acknowledgement from the standby database that the logs have also been written to log files on the standby database. The log data is guaranteed to be stored at both sites.

NEARSYNC
This mode provides somewhat less protection against transaction loss, in exchange for a shorter transaction response time than that of SYNC mode.

In this mode, log writes are considered successful only when the log records have been written to the log files on the primary database and when the primary database has received acknowledgement from the standby system that the logs have also been written to main memory on the standby system. Loss of data occurs only if both sites fail simultaneously and if the target site has not transferred to nonvolatile storage all of the log data that it has received.

ASync

This mode has the highest probability of transaction loss in the event of primary failure, in exchange for the shortest transaction response time among the three modes.

In this mode, log writes are considered successful only when the log records have been written to the log files on the primary database and have been delivered to the TCP layer of the primary system's host machine. Because the primary system does not wait for acknowledgement from the standby system, transactions might be considered committed when they are still on their way to the standby.

hadr_timeout - HADR timeout value

This parameter specifies the time (in seconds) that the high availability disaster recovery (HADR) process waits before considering a communication attempt to have failed.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Database server with local clients

Parameter type

Configurable

Default [range]

120 [1 - 4 294 967 295]

indexrec - Index re-creation time

This parameter indicates when the database manager will attempt to rebuild invalid indexes, and whether or not any index build will be redone during DB2 rollforward or HADR log replay on the standby database.

Configuration type

Database and Database Manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

UNIX Database Manager

restart [restart; restart_no_redo; access; access_no_redo]

Windows Database Manager

restart [restart; restart_no_redo; access; access_no_redo]

Database

Use system setting [system; restart; restart_no_redo; access; access_no_redo]

There are five possible settings for this parameter:

SYSTEM

use system setting specified in the database manager configuration file to decide when invalid indexes will be rebuilt, and whether any index build log records are to be redone during DB2 rollforward or HADR log replay. (Note: This setting is only valid for database configurations.)

ACCESS

Invalid indexes are rebuilt when the index is first accessed. Any fully logged index builds are redone during DB2 rollforward or HADR log replay. When HADR is started and an HADR takeover occurs, any invalid indexes are rebuilt after takeover when the underlying table is first accessed.

ACCESS_NO_REDO

Invalid indexes will be rebuilt when the underlying table is first accessed. Any fully logged index build will not be redone during DB2 rollforward or HADR log replay and those indexes will be left invalid. When HADR is started and an HADR takeover takes place, any invalid indexes will be rebuilt after takeover when the underlying table is first accessed.

RESTART

The default value for *indexrec*. Invalid indexes will be rebuilt when a RESTART DATABASE command is either explicitly or implicitly issued. Any fully logged index build will be redone during DB2 rollforward or HADR log replay. When HADR is started and an HADR takeover takes place, any invalid indexes will be rebuilt at the end of takeover.

Note that a RESTART DATABASE command is implicitly issued if the *autorestart* parameter is enabled.

RESTART_NO_REDO

Invalid indexes will be rebuilt when a RESTART DATABASE command is either explicitly or implicitly issued. (A RESTART DATABASE command is implicitly issued if the *autorestart* parameter is enabled.) Any fully logged index build will not be redone during DB2 rollforward or HADR log replay and instead those indexes will be rebuilt when rollforward completes or when HADR takeover takes place.

Indexes can become invalid when fatal disk problems occur. If this happens to the data itself, the data could be lost. However, if this happens to an index, the index can be recovered by re-creating it. If an index is rebuilt while users are connected to the database, two problems could occur:

- An unexpected degradation in response time might occur as the index file is re-created. Users accessing the table and using this particular index would wait while the index was being rebuilt.

- Unexpected locks might be held after index re-creation, especially if the user transaction that caused the index to be re-created never performed a COMMIT or ROLLBACK.

Recommendation: The best choice for this option on a high-user server and if restart time is not a concern, would be to have the index rebuilt at DATABASE RESTART time as part of the process of bringing the database back online after a crash.

Setting this parameter to “ACCESS” or to “ACCESS_NO_REDO” will result in a degradation of the performance of the database manager while the index is being re-created. Any user accessing that specific index or table would have to wait until the index is recreated.

If this parameter is set to “RESTART”, the time taken to restart the database will be longer due to index re-creation, but normal processing would not be impacted once the database has been brought back online.

Note: At database recovery time, all SQL procedure executables on the file system that belong to the database being recovered are removed. If *indexrec* is set to RESTART, all SQL procedure executables are extracted from the database catalog and put back on the file system at the next connection to the database. If *indexrec* is not set to RESTART, an SQL executable is extracted to the file system only on first execution of that SQL procedure.

The difference between the RESTART and the RESTART_NO_REDO values, or between the ACCESS and the ACCESS_NO_REDO values, is only significant when full logging is activated for index build operations, such as CREATE INDEX and REORG INDEX operations, or for an index rebuild. You can activate logging by enabling the *logindexbuild* database configuration parameter or by enabling LOG INDEX BUILD when altering a table. By setting *indexrec* to either RESTART or ACCESS, operations involving a logged index build can be rolled forward without leaving the index object in an invalid state, which would require the index to be rebuilt at a later time.

jdk_64_path - 64-Bit Software Developer’s Kit for Java installation path DAS

This parameter specifies the directory under which the 64-Bit Software Developer’s Kit (SDK) for Java, to be used for running DB2 administration server functions, is installed.

Configuration type

DB2 Administration Server

Applies to

DB2 Administration Server

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

Null [any valid path]

Note: This is different from the *jdk_path* configuration parameter, which specifies a 32-bit SDK for Java.

Environment variables used by the Java interpreter are computed from the value of this parameter. This parameter is only used on those platforms that support both 32- and 64-bit instances.

Those platforms are:

- 64-bit kernels of AIX, HP-UX, and Solaris operating systems
- 64-bit Windows on X64 and IPF
- 64-bit Linux kernel on AMD64 and Intel® EM64T systems (x64), POWER, and zSeries.

On all other platforms, only **jdk_path** is used.

Because there is no default value for this parameter, you should specify a value when you install the SDK for Java.

This parameter can only be updated from a Version 8 command line processor (CLP).

locklist - Maximum storage for lock list

This parameter indicates the amount of storage that is allocated to the lock list. There is one lock list per database and it contains the locks held by all applications concurrently connected to the database.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

UNIX Automatic [4 - 524 288]

Windows Database server with local and remote clients

Automatic [4 - 524 288]

Windows 64-bit Database server with local clients

Automatic [4 - 524 288]

Windows 32-bit Database server with local clients

Automatic [4 - 524 288]

Unit of measure

Pages (4 KB)

When allocated

When the first application connects to the database

When freed

When last application disconnects from the database

Locking is the mechanism that the database manager uses to control concurrent access to data in the database by multiple applications. Both rows and tables can be locked. The database manager can also acquire locks for internal use.

When this parameter is set to AUTOMATIC, it is enabled for self tuning. This allows the memory tuner to dynamically size the memory area controlled by this parameter as the workload requirements change. Because the memory tuner trades

memory resources between different memory consumers, there must be at least two memory consumers enabled for self tuning in order for self tuning to be active

The value of *locklist* is tuned together with the *maxlocks* parameter, therefore disabling self tuning of the *locklist* parameter automatically disables self tuning of the *maxlocks* parameter. Enabling self tuning of the *locklist* parameter automatically enables self tuning of the *maxlocks* parameter.

Automatic tuning of this configuration parameter will only occur when self tuning memory is enabled for the database (the *self_tuning_mem* configuration parameter is set to "ON.")

On 32-bit platforms, each lock requires 48 or 96 bytes of the lock list, depending on whether other locks are held on the object:

- 96 bytes are required to hold a lock on an object that has no other locks held on it
- 48 bytes are required to record a lock on an object that has an existing lock held on it.

On 64-bit platforms (except HP-UX/PA-RISC), each lock requires 64 or 128 bytes of the lock list, depending on whether other locks are held on the object:

- 128 bytes are required to hold a lock on an object that has no other locks held on it
- 64 bytes are required to record a lock on an object that has an existing lock held on it.

On 64-bit HP-UX/PA-RISC, each lock requires 80 or 160 bytes of the lock list, depending on whether or not other locks are held on the object.

When the percentage of the lock list used by one application reaches *maxlocks*, the database manager will perform lock escalation, from row to table, for the locks held by the application. Although the escalation process itself does not take much time, locking entire tables (versus individual rows) decreases concurrency, and overall database performance might decrease for subsequent accesses against the affected tables. Suggestions of how to control the size of the lock list are:

- Perform frequent COMMITs to release locks.
- When performing many updates, lock the entire table before updating (using the SQL LOCK TABLE statement). This will use only one lock, keeps others from interfering with the updates, but does reduce concurrency of the data.
You can also use the LOCKSIZE option of the ALTER TABLE statement to control how locking is done for a specific table.
Use of the Repeatable Read isolation level might result in an automatic table lock.
- Use the Cursor Stability isolation level when possible to decrease the number of share locks held. If application integrity requirements are not compromised use Uncommitted Read instead of Cursor Stability to further decrease the amount of locking.
- Set *locklist* to AUTOMATIC. The lock list will increase synchronously to avoid lock escalation or a lock list full situation.

Once the lock list is full, performance can degrade since lock escalation will generate more table locks and fewer row locks, thus reducing concurrency on shared objects in the database. Additionally there might be more deadlocks between applications (since they are all waiting on a limited number of table

locks), which will result in transactions being rolled back. Your application will receive an SQLCODE of -912 when the maximum number of lock requests has been reached for the database.

Recommendation: If lock escalations are causing performance concerns you might need to increase the value of this parameter or the *maxlocks* parameter. You can use the database system monitor to determine if lock escalations are occurring. Refer to the *lock_escals* (*lock escalations*) monitor element.

The following steps might help in determining the number of pages required for your lock list:

1. Calculate a lower bound for the size of your lock list, using *one* of the following calculations, depending on your environment:

a.

$$(512 * x * \text{maxappls}) / 4096$$

b. with Concentrator enabled:

$$(512 * x * \text{max_coordagents}) / 4096$$

c. in a partitioned database with Concentrator enabled:

$$(512 * x * \text{max_coordagents} * \text{number of database partitions}) / 4096$$

where 512 is an estimate of the average number of locks per application and *x* is the number of bytes required for each lock against an object that has an existing lock (40 bytes on 32-bit platforms, 64 bytes on 64-bit platforms).

2. Calculate an upper bound for the size of your lock list:

$$(512 * y * \text{maxappls}) / 4096$$

where *y* is the number of bytes required for the first lock against an object (80 bytes on 32-bit platforms, 128 bytes on 64-bit platforms).

3. Estimate the amount of concurrency you will have against your data and based on your expectations, choose an initial value for *locklist* that falls between the upper and lower bounds that you have calculated.
4. Using the database system monitor, as described below, tune the value of this parameter.

If *maxappls* or *max_coordagents* are set to AUTOMATIC in your applicable scenario, you should also set *locklist* to AUTOMATIC.

You can use the database system monitor to determine the maximum number of locks held by a given transaction. Refer to the *locks_held_top* (*maximum number of locks held*) monitor element.

This information can help you validate or adjust the estimated number of locks per application. In order to perform this validation, you will have to sample several applications, noting that the monitor information is provided at a transaction level, not an application level.

You might also want to increase *locklist* if *maxappls* is increased, or if the applications being run perform infrequent commits.

You should consider rebinding applications (using the REBIND command) after changing this parameter.

locktimeout - Lock timeout

This parameter specifies the number of seconds that an application will wait to obtain a lock, helping avoid global deadlocks for applications.

Configuration type

Database

Parameter type

Configurable

Default [range]

-1 [-1; 0 - 32 767]

Unit of measure

Seconds

If you set this parameter to 0, locks are not waited for. In this situation, if no lock is available at the time of the request, the application immediately receives a -911.

If you set this parameter to -1, lock timeout detection is turned off. In this situation a lock will be waited for (if one is not available at the time of the request) until either of the following:

- The lock is granted
- A deadlock occurs.

Recommendation: In a transaction processing (OLTP) environment, you can use an initial starting value of 30 seconds. In a query-only environment you could start with a higher value. In both cases, you should use benchmarking techniques to tune this parameter.

The value should be set to quickly detect waits that are occurring because of an abnormal situation, such as a transaction that is stalled (possibly as a result of a user leaving their workstation). You should set it high enough so valid lock requests do not time out because of peak workloads, during which time, there is more waiting for locks.

You can use the database system monitor to help you track the number of times an application (connection) experienced a lock timeout or that a database detected a timeout situation for all applications that were connected.

High values of the *lock_timeout* (number of lock timeouts) monitor element can be caused by:

- Too low a value for this configuration parameter.
- An application (transaction) that is holding locks for an extended period. You can use the database system monitor to further investigate these applications.
- A concurrency problem, that could be caused by lock escalations (from row-level to a table-level lock).

log_retain_status - Log retain status indicator

If set (when the *logretain* parameter setting is equal to Recovery), this parameter indicates that log files are being retained for use in roll-forward recovery.

Configuration type

Database

Parameter type

Informational

logarchmeth1 - Primary log archive method

This parameter specifies the media type of the primary destination for archived logs.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Default [range]

OFF [LOGRETAIN, USEREXIT, DISK, TSM, VENDOR]

OFF Specifies that the log archiving method is not to be used. If both *logarchmeth1* and *logarchmeth2* are set to OFF, the database is considered to be using circular logging and will not be rollforward recoverable. This is the default.

LOGRETAIN

This value can only be used for *logarchmeth1* and is equivalent to setting the *logretain* configuration parameter to RECOVERY. If you specify this value, the *logretain* configuration parameters will automatically be updated.

USEREXIT

This value is only valid for *logarchmeth1* and is equivalent to setting the *userexit* configuration parameter to ON. If specify this value, the *userexit* configuration parameter will be automatically updated.

DISK This value must be followed by a colon(:) and then a fully qualified existing path name where the log files will be archived. For example, if you set *logarchmeth1* to DISK:/u/dbuser/archived_logs the archive log files will be placed in a directory called /u/dbuser/archived_logs.

Note: If you are archiving to tape, you can use the db2tapemgr utility to store and retrieve log files.

TSM If specified without any additional configuration parameters, this value indicates that log files should be archived on the local TSM server using the default management class. If followed by a colon(:) and a TSM management class, the log files will be archived using the specified management class.

When archiving logs using TSM, before using the management class specified by the database configuration parameter, TSM first attempts to bind the object to the management class specified in the INCLUDE-EXCLUDE list found in the TSM client options file. If a match is not found, the default TSM management class specified on the TSM server will be used. TSM will then rebind the object to the management class specified by the database configuration parameter.

Thus, the default management class, as well as the management class specified by the database configuration parameter, must contain an archive copy group, or the archive operation will fail.

VENDOR

Specifies that a vendor library will be used to archive the log files. This value must be followed by a colon(:) and the name of the library. The APIs provided in the library must use the backup and restore APIs for vendor products.

Note:

1. If either *logarchmeth1* or *logarchmeth2* is set to a value other than OFF, the database is configured for rollforward recovery.
2. If you update the *userexit* or *logretain* configuration parameters *logarchmeth1* will automatically be updated and vice versa. However, if you are using either *userexit* or *logretain*, *logarchmeth2* must be set to OFF.

logarchmeth2 - Secondary log archive method

This parameter specifies the media type of the secondary destination for archived logs.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Default [range]

OFF [LOGRETAIN, USEREXIT, DISK, TSM, VENDOR]

OFF Specifies that the log archiving method is not to be used. If both *logarchmeth1* and *logarchmeth2* are set to OFF, the database is considered to be using circular logging and will not be rollforward recoverable. This is the default.

LOGRETAIN

This value can only be used for *logarchmeth1* and is equivalent to setting the *logretain* configuration parameter to RECOVERY. If you specify this value, the *logretain* configuration parameters will automatically be updated.

USEREXIT

This value is only valid for *logarchmeth1* and is equivalent to setting the *userexit* configuration parameter to ON. If specify this value, the *userexit* configuration parameter will be automatically updated.

DISK This value must be followed by a colon(:) and then a fully qualified existing path name where the log files will be archived. For example, if you set *logarchmeth1* to DISK:/u/dbuser/archived_logs the archive log files will be placed in a directory called /u/dbuser/archived_logs.

Note: If you are archiving to tape, you can use the `db2tapemgr` utility to store and retrieve log files.

TSM If specified without any additional configuration parameters, this value indicates that log files should be archived on the local TSM server using the default management class. If followed by a colon(:) and a TSM management class, the log files will be archived using the specified management class.

VENDOR

Specifies that a vendor library will be used to archive the log files. This value must be followed by a colon(:) and the name of the library. The APIs provided in the library must use the backup and restore APIs for vendor products.

Note:

1. If either `logarchmeth1` or `logarchmeth2` is set to a value other than OFF, the database is configured for rollforward recovery.
2. If you update the `userexit` or `logretain` configuration parameters `logarchmeth1` will automatically be updated and vice versa. However, if you are using either `userexit` or `logretain`, `logarchmeth2` must be set to OFF.

If this path is specified, log files will be archived to both this destination and the destination specified by the `logarchmeth1` database configuration parameter.

logarchopt1 - Primary log archive options

This parameter specifies the options field for the primary destination for archived logs (if required).

Configuration type

Database

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Default [range]

Null [not applicable]

logarchopt2 - Secondary log archive options

This parameter specifies the options field for the secondary destination for archived logs (if required).

Configuration type

Database

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type
Configurable Online

Default [range]
Null [not applicable]

logbufsz - Log buffer size

This parameter allows you to specify the amount of the database heap (defined by the *dbheap* parameter) to use as a buffer for log records before writing these records to disk.

Configuration type
Database

Parameter type
Configurable

Default [range]

32-bit platforms
8 [4 - 4 096]

64-bit platforms
8 [4 - 65 535]

Starting from Version 9.5 Fix Pack 3, the range has been increased to 8 [4 - 131 070]

Unit of measure
Pages (4 KB)

Log records are written to disk when one of the following occurs:

- A transaction commits or a group of transactions commit, as defined by the *mincommit* configuration parameter
- The log buffer is full
- As a result of some other internal database manager event.

This parameter must also be less than or equal to the *dbheap* parameter. Buffering the log records will result in more efficient logging file I/O because the log records will be written to disk less frequently and more log records will be written at each time.

Recommendation: Increase the size of this buffer area if there is considerable read activity on a dedicated log disk, or there is high disk utilization. When increasing the value of this parameter, you should also consider the *dbheap* parameter since the log buffer area uses space controlled by the *dbheap* parameter.

You can use the database system monitor to determine how much of the log buffer space is used for a particular transaction (or unit of work). Refer to the *log_space_used* (unit of work log space used) monitor element.

logfilsiz - Size of log files

This parameter defines the size of each primary and secondary log file. The size of these log files limits the number of log records that can be written to them before they become full and a new log file is required.

Configuration type
Database

Parameter type
Configurable

Default [range]

UNIX 1000 [4 - 524 286]

Starting from Version 9.5 Fix Pack 3, the range has been increased to 1000 [4 - 1 048 572]

Windows

1000 [4 - 524 286]

Starting from Version 9.5 Fix Pack 3, the range has been increased to 1000 [4 - 1 048 572]

Unit of measure
Pages (4 KB)

The use of primary and secondary log files as well as the action taken when a log file becomes full are dependent on the type of logging that is being performed:

- Circular logging

A primary log file can be reused when the changes recorded in it have been committed. If the log file size is small and applications have processed a large number of changes to the database without committing the changes, a primary log file can quickly become full. If all primary log files become full, the database manager will allocate secondary log files to hold the new log records.

- Log retention logging

When a primary log file is full, the log is archived and a new primary log file is allocated.

Recommendation: You must balance the size of the log files with the number of primary log files:

- The value of the *logfilsiz* should be increased if the database has a large number of update, delete, or insert transactions running against it which will cause the log file to become full very quickly.

Note: The upper limit of log file size, combined with the upper limit of the number of log files (*logprimary* + *logsecond*), gives an upper limit of 512 GB of active log space. Starting from Version 9.5 Fix Pack 3, there is now 1024 GB of active log space available.

A log file that is too small can affect system performance because of the overhead of archiving old log files, allocating new log files, and waiting for a usable log file.

- The value of the *logfilsiz* should be reduced if disk space is scarce, since primary logs are preallocated at this size.

A log file that is too large can reduce your flexibility when managing archived log files and copies of log files, since some media might not be able to hold an entire log file.

If you are using log retention, the current active log file is closed and truncated when the last application disconnects from a database. When the next connection to the database occurs, the next log file is used. Therefore, if you understand the logging requirements of your concurrent applications, you might be able to determine a log file size that will not allocate excessive amounts of wasted space.

loghead - First active log file

This parameter contains the name of the log file that is currently active.

Configuration type

Database

Parameter type

Informational

logindexbuild - Log index pages created

This parameter specifies whether index creation, recreation, or reorganization operations are to be logged so that indexes can be reconstructed during DB2 rollforward operations or high availability disaster recovery (HADR) log replay procedures.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Default [range]

Off [On; Off]

logpath - Location of log files

This parameter contains the current path being used for logging purposes.

Configuration type

Database

Parameter type

Informational

You cannot change this parameter directly as it is set by the database manager after a change to the *newlogpath* parameter becomes effective.

When a database is created, the recovery log file for it is created in a subdirectory of the directory containing the database. The default is a subdirectory named SQLOGDIR under the directory created for the database.

logprimary - Number of primary log files

This parameter allows you to specify the number of primary log files to be preallocated. The primary log files establish a fixed amount of storage allocated to the recovery log files.

Configuration type

Database

Parameter type

Configurable

Default [range]
3 [2 - 256]

Unit of measure
Counter

When allocated

- The database is created
- A log is moved to a different location (which occurs when the *logpath* parameter is updated)
- When the database is next started following an increase following an increase in the value of this parameter (*logprimary*), provided that the database is not started as an HADR standby database
- A log file has been archived and a new log file is allocated (the *logretain* or *userexit* parameter must be enabled)
- If the *logfilsiz* parameter has been changed, the log files are re-sized during the next database startup, provided that it is not started as an HADR standby database

When freed

Not freed unless this parameter decreases. If decreased, unneeded log files are deleted during the next connection to the database.

Under circular logging, the primary logs are used repeatedly in sequence. That is, when a log is full, the next primary log in the sequence is used if it is available. A log is considered available if all units of work with log records in it have been committed or rolled-back. If the next primary log in sequence is not available, then a secondary log is allocated and used. Additional secondary logs are allocated and used until the next primary log in the sequence becomes available or the limit imposed by the *logsecond* parameter is reached. These secondary log files are dynamically deallocated as they are no longer needed by the database manager.

The number of primary and secondary log files must comply with the following:

- If *logsecond* has a value of -1, $logprimary \leq 256$.
- If *logsecond* does not have a value of -1, $(logprimary + logsecond) \leq 256$.

Recommendation: The value chosen for this parameter depends on a number of factors, including the type of logging being used, the size of the log files, and the type of processing environment (for example, length of transactions and frequency of commits).

Increasing this value will increase the disk requirements for the logs because the primary log files are preallocated during the very first connection to the database.

If you find that secondary log files are frequently being allocated, you might be able to improve system performance by increasing the log file size (*logfilsiz*) or by increasing the number of primary log files.

For databases that are not frequently accessed, in order to save disk storage, set the parameter to 2. For databases enabled for roll-forward recovery, set the parameter larger to avoid the overhead of allocating new logs almost immediately.

You can use the database system monitor to help you size the primary log files. Observation of the following monitor values over a period of time will aid in better tuning decisions, as average values might be more representative of your ongoing requirements.

- *sec_log_used_top* (maximum secondary log space used)
- *tot_log_used_top* (maximum total log space used)
- *sec_logs_allocated* (secondary logs allocated currently)

logretain - Log retain enable

This parameter is deprecated in Version 9.5, but is still being used by pre-Version 9.5 data servers and clients. Any value specified for this configuration parameter will be ignored by the DB2 Version 9.5 database manager.

Note: The following information applies only to pre-Version 9.5 data servers and clients.

This parameter determines whether active log files are retained and available for roll-forward recovery.

Configuration type

Database

Parameter type

Configurable

Default [range]

No [Recovery; No]

The values are as follows:

- No, to indicate that logs are not retained.
- Recovery, to indicate that the logs are retained, and can be used for forward recovery.

If **logretain** is set to Recovery or **userexit** is set to Yes, the active log files will be retained and become online archive log files for use in roll-forward recovery. This is called log retention logging.

After **logretain** is set to Recovery or **userexit** is set to Yes (or both), you must make a full backup of the database. This state is indicated by the **backup_pending** flag parameter.

Note:

Both **logarchmeth1** or **logretain** will enable rollforward recovery. However, only one method should be enabled for a database at one time.

If using **logarchmeth1**, do not set the **logretain** and **userexit** configuration parameters. If the **logretain** configuration parameter is set to recover, the value for **logarchmeth1** will automatically be set to **logretain**.

It is recommended that **logarchmeth1** (and **logarchmeth2**) be used rather than **logretain** and **userexit** to activate archive logging and rollforward recovery. The **logretain** and **userexit** options have been kept to support users who have not yet migrated to **logarchmeth1**.

logsecond - Number of secondary log files

This parameter specifies the number of secondary log files that are created and used for recovery log files (only as needed).

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

2 [-1; 0 – 254]

Unit of measure

Counter

When allocatedAs needed when *logprimary* is insufficient (see detail below)**When freed**

Over time as the database manager determines they will no longer be required.

When the primary log files become full, the secondary log files (of size *logfilsiz*) are allocated one at a time as needed, up to a maximum number as controlled by this parameter. An error code will be returned to the application, and the database will be shut down, if more secondary log files are required than are allowed by this parameter.

If you set *logsecond* to -1, the database is configured with infinite active log space. There is no limit on the size or the number of in-flight transactions running on the database. If you set *logsecond* to -1, you still use the *logprimary* and *logfilsiz* configuration parameters to specify how many log files the database manager should keep in the active log path. If the database manager needs to read log data from a log file, but the file is not in the active log path, the database manager retrieves the log file from the archive to the active log path. (The database manager retrieves the files to the overflow log path, if you have configured one.) Once the log file is retrieved, the database manager will cache this file in the active log path so that other reads of log data from the same file will be fast. The database manager will manage the retrieval, caching, and removal of these log files as required.

If your log path is a raw device, you must configure the *overflowlogpath* configuration parameter in order to set *logsecond* to -1.

By setting *logsecond* to -1, you will have no limit on the size of the unit of work or the number of concurrent units of work. However, rollback (both at the savepoint level and at the unit of work level) could be very slow due to the need to retrieve log files from the archive. Crash recovery could also be very slow for the same reason. The database manager writes a message to the administration notification log to warn you that the current set of active units of work has exceeded the primary log files. This is an indication that rollback or crash recovery could be extremely slow.

To set *logsecond* to -1, the *logarchmeth1* configuration parameter must be set to a value other than OFF or LOGRETAIN.

Recommendation: Use secondary log files for databases that have periodic needs for large amounts of log space. For example, an application that is run once a

month might require log space beyond that provided by the primary log files. Since secondary log files do not require permanent file space they are advantageous in this type of situation.

max_log - Maximum log per transaction

This parameter specifies if there is a limit to the percentage of log space that a transaction can consume, and what that limit is.

Configuration type

Database

Parameter type

Configurable online

Propagation class

Immediate

Default [range]

0 [0 — 100]

Unit of measure

Percentage

If the value is not 0, this parameter indicates the percentage of primary log space that can be consumed by one transaction.

If the value is set to 0, there is no limit regarding how much space (as a percentage of total primary log space) one single transaction can consume. This was the behavior of transactions prior to Version 8.

maxappls - Maximum number of active applications

This parameter specifies the maximum number of concurrent applications that can be connected (both local and remote) to a database. Since each application that attaches to a database causes some private memory to be allocated, allowing a larger number of concurrent applications will potentially use more memory.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

Automatic [1 - 60 000]

Unit of measure

Counter

Setting *maxappls* to *automatic* has the effect of allowing any number of connected applications. The database manager will dynamically allocate the resources it needs to support new applications.

If you do not want to set this parameter to *automatic*, the value of this parameter must be equal to or greater than the sum of the connected applications, plus the number of these same applications that might be concurrently in the process of completing a two-phase commit or rollback. Then add to this sum the anticipated number of indoubt transactions that might exist at any one time.

When an application attempts to connect to a database, but *maxappls* has already been reached, an error is returned to the application indicating that the maximum number of applications have been connected to the database.

In a partitioned database environment, this is the maximum number of applications that can be concurrently active against a database partition. This parameter limits the number of active applications against the database partition on a database partition server, regardless of whether the server is the coordinator node for the application or not. The catalog node in a partitioned database environment requires a higher value for *maxappls* than is the case for other types of environments because, in the partitioned database environment, every application requires a connection to the catalog node.

Recommendation: Increasing the value of this parameter without lowering the *maxlocks* parameter or increasing the *locklist* parameter could cause you to reach the database limit on locks (*locklist*) rather than the application limit and as a result cause pervasive lock escalation problems.

To a certain extent, the maximum number of applications is also governed by *max_coordagents*. An application can only connect to the database, if there is an available connection (*maxappls*) as well as an available coordinating agent (*max_coordagents*).

maxfilop - Maximum database files open per application

This parameter specifies the maximum number of file handles that can be open for each database.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Transaction boundary

Default [range]

AIX, Sun, HP, and Linux 64-bit

61 440 [64 - 61 440]

Linux 32-bit

30 720 [64 - 30 720]

Windows 32-bit

32 768 [64 - 32 768]

Windows 64-bit

65 335 [64 - 65 335]

Unit of measure

Counter

If opening a file causes this value to be exceeded, some files in use by this database are closed. If *maxfilop* is too small, the overhead of opening and closing files will become excessive and might degrade performance.

Both SMS table spaces and DMS table space file containers are treated as files in the database manager's interaction with the operating system, and file handles are required. More files are generally used by SMS table spaces compared to the

number of containers used for a DMS file table space. Therefore, if you are using SMS table spaces, you will need a larger value for this parameter compared to what you would require for DMS file table spaces.

You can also use this parameter to ensure that the overall total of file handles used by the database manager does not exceed the operating system limit by limiting the number of handles per database to a specific number; the actual number will vary depending on the number of databases running concurrently.

maxlocks - Maximum percent of lock list before escalation

This parameter defines a percentage of the lock list held by an application that must be filled before the database manager performs lock escalation.

Configuration type

Database

Parameter type

Configurable online

Propagation class

Immediate

Default [range]

Automatic [1 - 100]

Unit of measure

Percentage

Lock escalation is the process of replacing row locks with table locks, reducing the number of locks in the list. When the number of locks held by any one application reaches this percentage of the total lock list size, lock escalation will occur for the locks held by that application. Lock escalation also occurs if the lock list runs out of space.

The database manager determines which locks to escalate by looking through the lock list for the application and finding the table with the most row locks. If after replacing these with a single table lock, the **maxlocks** value is no longer exceeded, lock escalation will stop. If not, it will continue until the percentage of the lock list held is below the value of **maxlocks**. The **maxlocks** parameter multiplied by the **maxappls** parameter cannot be less than 100.

When this parameter is set to AUTOMATIC, it is enabled for self tuning. This allows the memory tuner to dynamically size the memory area controlled by this parameter as the workload requirements change. Because the memory tuner trades memory resources between different memory consumers, there must be at least two memory consumers enabled for self tuning in order for self tuning to be active.

The value of **locklist** is tuned together with the **maxlocks** parameter, therefore disabling self tuning of the **locklist** parameter automatically disables self tuning of the **maxlocks** parameter. Enabling self tuning of the **locklist** parameter automatically enables self tuning of the **maxlocks** parameter.

Automatic tuning of this configuration parameter will only occur when self tuning memory is enabled for the database (the **self_tuning_mem** configuration parameter is set to ON).

On 32-bit platforms, each lock requires 48 or 96 bytes of the lock list, depending on whether other locks are held on the object:

- 96 bytes are required to hold a lock on an object that has no other locks held on it
- 48 bytes are required to record a lock on an object that has an existing lock held on it.

On 64-bit platforms (except HP-UX/PA-RISC), each lock requires 64 or 128 bytes of the lock list, depending on whether other locks are held on the object:

- 128 bytes are required to hold a lock on an object that has no other locks held on it.
- 64 bytes are required to record a lock on an object that has an existing lock held on it.

On 64-bit HP-UX/PA-RISC, each lock requires 80 or 160 bytes of the lock list, depending on whether or not other locks are held on the object.

Recommendation: The following formula allows you to set **maxlocks** to allow an application to hold twice the average number of locks:

$$\text{maxlocks} = 2 * 100 / \text{maxapps}$$

Where 2 is used to achieve twice the average and 100 represents the largest percentage value allowed. If you have only a few applications that run concurrently, you could use the following formula as an alternative to the first formula:

$$\text{maxlocks} = 2 * 100 / (\text{average number of applications running concurrently})$$

One of the considerations when setting **maxlocks** is to use it in conjunction with the size of the lock list (**locklist**). The actual limit of the number of locks held by an application before lock escalation occurs is:

- $\text{maxlocks} * \text{locklist} * 4096 / (100 * 48)$ on a 32-bit system
- $\text{maxlocks} * \text{locklist} * 4096 / (100 * 80)$ on a 64-bit system HP-UX/PA-RISC environment
- $\text{maxlocks} * \text{locklist} * 4096 / (100 * 64)$ on other 64-bit systems

Where 4096 is the number of bytes in a page, 100 is the largest percentage value allowed for **maxlocks**, and 48 is the number of bytes per lock on a 32-bit system, 80 is the number of bytes per lock on a HP-UX/PA-RISC 64-bit system, and 64 is the number of bytes per lock on other 64-bit systems. If you know that one of your applications requires 1000 locks, and you do not want lock escalation to occur, then you should choose values for **maxlocks** and **locklist** in this formula so that the result is greater than 1000. (Using 10 for **maxlocks** and 100 for **locklist**, this formula results in greater than the 1000 locks needed.)

If **maxlocks** is set too low, lock escalation happens when there is still enough lock space for other concurrent applications. If **maxlocks** is set too high, a few applications can consume most of the lock space, and other applications will have to perform lock escalation. The need for lock escalation in this case results in poor concurrency.

You can use the database system monitor to help you track and tune this configuration parameter.

min_dec_div_3 - Decimal division scale to 3

This parameter is provided as a quick way to enable a change to computation of the scale for decimal division in SQL.

Configuration type

Database

Parameter type

Configurable

Default [range]

No [Yes, No]

The *min_dec_div_3* database configuration parameter changes the resulting scale of a decimal arithmetic operation involving division. It can be set to "Yes" or "No". The default value for *min_dec_div_3* is "No". If the value is "No", the scale is calculated as $31-p+s-s'$. If set to "Yes", the scale is calculated as $\text{MAX}(3, 31-p+s-s')$. This causes the result of decimal division to always have a scale of at least 3. Precision is always 31.

Changing this database configuration parameter might cause changes to applications for existing databases. This can occur when the resulting scale for decimal division would be impacted by changing this database configuration parameter. Listed below are some possible scenarios that might impact applications. These scenarios should be considered before changing the *min_dec_div_3* on a database server with existing databases.

- If the resulting scale of one of the view columns is changed, a view that is defined in an environment with one setting could fail with SQLCODE -344 when referenced after the database configuration parameter is changed. The message SQL0344N refers to recursive common table expressions, however, if the object name (first token) is a view, then you will need to drop the view and create it again to avoid this error.
- A static package will not change behavior until the package is rebound, either implicitly or explicitly. For example, after changing the value from NO to YES, the additional scale digits might not be included in the results until rebind occurs. For any changed static packages, an explicit REBIND command can be used to force a rebind.
- A check constraint involving decimal division might restrict some values that were previously accepted. Such rows now violate the constraint but will not be detected until one of the columns involved in the check constraint row is updated or the SET INTEGRITY statement with the IMMEDIATE CHECKED option is processed. To force checking of such a constraint, perform an ALTER TABLE statement in order to drop the check constraint and then perform an ALTER TABLE statement to add the constraint again.

Note: *min_dec_div_3* also has the following limitations:

1. The command GET DB CFG FOR DBNAME will not display the *min_dec_div_3* setting. The best way to determine the current setting is to observe the side-effect of a decimal division result. For example, consider the following statement:

```
VALUES (DEC(1,31,0)/DEC(1,31,5))
```

If this statement returns sqlcode SQL0419N, the database does not have *min_dec_div_3* support, or it is set to "No". If the statement returns 1.000, *min_dec_div_3* is set to "Yes".

2. *min_dec_div_3* does not appear in the list of configuration keywords when you run the following command: ? UPDATE DB CFG

mincommit - Number of commits to group

This parameter allows you to delay the writing of log records to disk until a minimum number of commits have been performed, helping reduce the database manager overhead associated with writing log records.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

1 [1 – 25]

Unit of measure

Counter

This delay will improve performance when you have multiple applications running against a database and many commits are requested by the applications within a very short time frame.

This grouping of commits will only occur when the value of this parameter is greater than one and when the number of applications connected to the database is greater than or equal to the value of this parameter. When commit grouping is being performed, application commit requests could be held until either one second has elapsed or the number of commit requests equals the value of this parameter.

This parameter should be incremented by small amounts only; for example one (1). You should also use multi-user tests to verify that increasing the value of this parameter provides the expected results.

Changes to the value specified for this parameter take effect immediately; you do not have to wait until all applications disconnect from the database.

Recommendation: Increase this parameter from its default value if multiple read/write applications typically request concurrent database commits. This will result in more efficient logging file I/O as it will occur less frequently and write more log records each time it does occur.

You could also sample the number of transactions per second and adjust this parameter to accommodate the peak number of transactions per second (or some large percentage of it). Accommodating peak activity would minimize the overhead of writing log records during transaction intensive periods.

If you increase *mincommit*, you might also need to increase the *logbufsz* parameter to avoid having a full log buffer force a write during these transaction intensive periods. In this case, the *logbufsz* should be equal to:

$\text{mincommit} * (\text{log space used, on average, by a transaction})$

You can use the database system monitor to help you tune this parameter in the following ways:

- Calculating the peak number of transactions per second:
Taking monitor samples throughout a typical day, you can determine your transaction intensive periods. You can calculate the total transactions by adding the following monitor elements:
 - *commit_sql_stmts* (commit statements attempted)
 - *rollback_sql_stmts* (rollback statements attempted)
 Using this information and the available timestamps, you can calculate the number of transactions per second.
- Calculating the log space used per transaction:
Using sampling techniques over a period of time and a number of transactions, you can calculate an average of the log space used with the following monitor element:
 - *log_space_used* (unit of work log space used)

mirrorlogpath - Mirror log path

This parameter allows you to specify a string of up to 242 bytes for the mirror log path. The string must point to a path name, and it must be a fully qualified path name, not a relative path name.

Configuration type

Database

Parameter type

Configurable

Default [range]

Null [any valid path or device]

Note: In a single or multi-partition DB2 ESE environment, the node number is automatically appended to the path. This is done to maintain the uniqueness of the path in multiple logical node configurations.

If *mirrorlogpath* is configured, DB2 will create active log files in both the log path and the mirror log path. All log data will be written to both paths. The mirror log path has a duplicated set of active log files, such that if there is a disk error or human error that destroys active log files on one of the paths, the database can still function.

If the mirror log path is changed, there might be log files in the old mirror log path. These log files might not have been archived, so you might need to archive these log files manually. Also, if you are running replication on this database, replication might still need the log files from before the log path change. If the database is configured with the User Exit Enable (*userexit*) database configuration parameter set to Yes, and if all the log files have been archived either by DB2 automatically or by yourself manually, then DB2 will be able to retrieve the log files to complete the replication process. Otherwise, you can copy the files from the old mirror log path to the new mirror log path.

If *logpath* or *newlogpath* specifies a raw device as the location where the log files are stored, mirror logging, as indicated by *mirrorlogpath*, is not allowed. If *logpath* or *newlogpath* specifies a file path as the location where the log files are stored, mirror logging is allowed and *mirrorlogpath* must also specify a file path.

Recommendation: Just like the log files, the mirror log files should be on a physical disk that does not have high I/O.

It is strongly recommended that this path be on a separate device than the primary log path.

You can use the database system monitor to track the number of I/Os related to database logging.

The following data elements return the amount of I/O activity related to database logging. You can use an operating system monitor tool to collect information about other disk I/O activity, then compare the two types of I/O activity.

- *log_reads* (number of log pages read)
- *log_writes* (number of log pages written).

multipage_alloc - Multipage file allocation enabled

Multipage file allocation is used to improve insert performance. It applies to SMS table spaces only. If enabled, all SMS table spaces are affected: there is no selection possible for individual SMS table spaces.

Configuration type

Database

Parameter type

Informational

The default for the parameter is Yes: multipage file allocation is enabled.

Following database creation, this parameter cannot be set to No. Multipage file allocation cannot be disabled once it has been enabled. The `db2empfa` tool can be used to enable multipage file allocation for a database that currently has it disabled.

newlogpath - Change the database log path

This parameter allows you to specify a string of up to 242 bytes to change the location where the log files are stored.

Configuration type

Database

Parameter type

Configurable

Default [range]

Null [any valid path or device]

The string can point to either a path name or to a raw device. Note that as of DB2 Version 9, the use of raw devices for database logging is deprecated. As an alternative to using raw logs, you can use either direct input/output (DIO) or concurrent input/output (CIO).

If the string points to a path name, it must be a fully qualified path name, not a relative path name.

In a single or multi-partition DB2 ESE environment, the node number is automatically appended to the path. This is done to maintain the uniqueness of the path in multiple logical node configurations.

If you want to use replication, and your log path is a raw device, the *overflowlogpath* configuration parameter must be configured.

To specify a device, specify a string that the operating system identifies as a device. For example:

- On Windows, \\.\d: or \\.\PhysicalDisk5

Note: You must have Windows Version 4.0 with Service Pack 3 or later installed to be able to write logs to a device.

- On Linux and UNIX platforms, /dev/rdblog8

Note: You can only specify a device on AIX, Windows 2000, Windows, Solaris, HP-UX, and Linux platforms.

The new setting does not become the value of *logpath* until both of the following occur:

- The database is in a consistent state, as indicated by the *database_consistent* parameter.
- All applications are disconnected from the database

When the first new connection is made to the database, the database manager will move the logs to the new location specified by *logpath*.

There might be log files in the old log path. These log files might not have been archived. You might need to archive these log files manually. Also, if you are running replication on this database, replication might still need the log files from before the log path change. If the database is configured with the User Exit Enable (*userexit*) database configuration parameter set to Yes, and if all the log files have been archived either by DB2 automatically or by yourself manually, then DB2 will be able to retrieve the log files to complete the replication process. Otherwise, you can copy the files from the old log path to the new log path.

If *logpath* or *newlogpath* specifies a raw device as the location where the log files are stored, mirror logging, as indicated by *mirrorlogpath*, is not allowed. If *logpath* or *newlogpath* specifies a file path as the location where the log files are stored, mirror logging is allowed and *mirrorlogpath* must also specify a file path.

Recommendation: Ideally, the log files will be on a physical disk which does **not** have high I/O. For instance, avoid putting the logs on the same disk as the operating system or high volume databases. This will allow for efficient logging activity with a minimum of overhead such as waiting for I/O.

You can use the database system monitor to track the number of I/Os related to database logging.

The monitor elements *log_reads* (number of log pages read) and *log_writes* (number of log pages written) return the amount of I/O activity related to database logging. You can use an operating system monitor tool to collect information about other disk I/O activity, then compare the two types of I/O activity.

Do not use a network or local file system that is shared as the log path for both the primary and standby databases in a DB2 High Availability Disaster Recovery (HADR) database pair. The primary and standby databases each have copies of the transaction logs – the primary database ships logs to the standby database. If the log path for both the primary and standby databases points to the same physical location, then the primary and standby database would use the same physical files for their respective copies of the logs. The database manager returns an error if the database manager detects a shared log path.

num_db_backups - Number of database backups

This parameter specifies the number of database backups to retain for a database.

Configuration type

Database

Parameter type

Configurable online

Propagation class

Transaction boundary

Default [range]

12 [1 - 32 767]

After the specified number of backups is reached, old backups are marked as expired in the recovery history file. Recovery history file entries for the table space backups and load copy backups that are related to the expired database backup are also marked as expired. When a backup is marked as expired, the physical backups can be removed from where they are stored (for example, disk, tape, TSM). The next database backup will prune the expired entries from the recovery history file.

The *rec_his_retentn* configuration parameter should be set to a value compatible with the value of *num_db_backups*. For example, if *num_db_backup* is set to a large value, the value for *rec_his_retentn* should be large enough to support that number of backups.

num_freqvalues - Number of frequent values retained

This parameter allows you to specify the number of “most frequent values” that will be collected when the WITH DISTRIBUTION option is specified on the RUNSTATS command.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

10 [0 - 32 767]

Unit of measure

Counter

Increasing the value of this parameter increases the amount of statistics heap (*stat_heap_sz*) used when collecting statistics.

The “most frequent value” statistics help the optimizer understand the distribution of data values within a column. A higher value results in more information being available to the query optimizer but requires additional catalog space. When 0 is specified, no frequent-value statistics are retained, even if you request that distribution statistics be collected.

You can also specify the number of frequent values retained as part of the RUNSTATS command at the table or the column level. by using the

NUM_FREQVALUES option. If none is specified, the *num_freqvalues* configuration parameter value is used. Changing the number of frequent values retained through the RUNSTATS command is easier than making the change using the *num_freqvalues* database configuration parameter.

Updating this parameter can help the optimizer obtain better selectivity estimates for some predicates (=, <, >) over data that is non-uniformly distributed. More accurate selectivity calculations might result in the choice of more efficient access plans.

After changing the value of this parameter, you need to:

- Run the RUNSTATS command again to collect statistics with the changed number of frequent values
- Rebind any packages containing static SQL or XQuery statements.

When using RUNSTATS, you have the ability to limit the number of frequent values collected at both the table level and the column level. This allows you to optimize on space occupied in the catalogs by reducing the distribution statistics for columns where they could not be exploited and yet still using the information for critical columns.

Recommendation: In order to update this parameter you should determine the degree of non-uniformity in the most important columns (in the most important tables) that typically have selection predicates. This can be done using an SQL SELECT statement that provides an ordered ranking of the number of occurrences of each value in a column. You should not consider uniformly distributed, unique, long, or LOB columns. A reasonable practical value for this parameter lies in the range of 10 to 100.

Note that the process of collecting frequent value statistics requires significant CPU and memory (*stat_heap_sz*) resources.

num_iocleaners - Number of asynchronous page cleaners

This parameter allows you to specify the number of asynchronous page cleaners for a database.

Configuration type

Database

Parameter type

Configurable

Default [range]

Automatic [0 – 255]

Unit of measure

Counter

These page cleaners write changed pages from the buffer pool to disk before the space in the buffer pool is required by a database agent. As a result, database agents should not have to wait for changed pages to be written out so that they might use the space in the buffer pool. This improves overall performance of the database applications.

If you set the parameter to zero (0), no page cleaners are started and as a result, the database agents will perform all of the page writes from the buffer pool to disk. This parameter can have a significant performance impact on a database

stored across many physical storage devices, since in this case there is a greater chance that one of the devices will be idle. If no page cleaners are configured, your applications might encounter periodic log full conditions.

If this parameter is set to `AUTOMATIC`, the number of page cleaners started will be based on the number of CPUs configured on the current machine, as well as the number of local logical database partitions in a partitioned database environment. There will always be at least one page cleaner started when this parameter is set to `AUTOMATIC`.

The number of page cleaners to start when this parameter is set to `AUTOMATIC` will be calculated using the following formula:

$$\text{number of page cleaners} = \max(\text{ceil}(\# \text{ CPUs} / \# \text{ local logical DPs}) - 1, 1)$$

This formula ensures that the number of page cleaners is distributed almost evenly across your logical database partitions, and that there are no more page cleaners than there are CPUs.

If the applications for a database primarily consist of transactions that update data, an increase in the number of cleaners will speed up performance. Increasing the page cleaners will also decrease recovery time from soft failures, such as power outages, because the contents of the database on disk will be more up-to-date at any given time.

Recommendation: Consider the following factors when setting the value for this parameter:

- Application type
 - If it is a query-only database that will not have updates, set this parameter to be zero (0). The exception would be if the query work load results in many TEMP tables being created (you can determine this by using the explain utility).
 - If transactions are run against the database, set this parameter to be between one and the number of physical storage devices used for the database.
- Workload
 - Environments with high update transaction rates might require more page cleaners to be configured.
- Buffer pool sizes
 - Environments with large buffer pools might also require more page cleaners to be configured.

You can use the database system monitor to help you tune this configuration parameter using information from the event monitor about write activity from a buffer pool:

- The parameter can be reduced if both of the following conditions are true:
 - *pool_data_writes* is approximately equal to *pool_async_data_writes*
 - *pool_index_writes* is approximately equal to *pool_async_index_writes*.
- The parameter should be increased if either of the following conditions are true:
 - *pool_data_writes* is much greater than *pool_async_data_writes*
 - *pool_index_writes* is much greater than *pool_async_index_writes*.

num_ioservers - Number of I/O servers

This parameter specifies the number of I/O servers for a database. No more than this number of I/Os for prefetching and utilities can be in progress for a database at any time.

Configuration type

Database

Parameter type

Configurable

Default [range]

Automatic [1 – 255]

Unit of measure

Counter

When allocated

When an application connects to a database

When freed

When an application disconnects from a database

I/O servers, also called prefetchers, are used on behalf of the database agents to perform prefetch I/O and asynchronous I/O by utilities such as backup and restore. An I/O server waits while an I/O operation that it initiated is in progress. Non-prefetch I/Os are scheduled directly from the database agents and as a result are not constrained by *num_ioservers*.

If this parameter is set to AUTOMATIC, the number of prefetchers started will be based on the parallelism settings of the table spaces in the current database partition. (Parallelism settings are controlled by the DB2_PARALLEL_IO environment variable.) For each DMS table space, the value of this parallelism setting will be multiplied by the maximum number of containers in the table space stripe set. For each SMS table space, the value of this parallelism setting will be multiplied by the number of containers in the table space. The largest result over all table spaces in the current database partition will be used as the number of prefetchers to start. There will always be at least three prefetchers started when this parameter is set to AUTOMATIC.

When this parameter is set to AUTOMATIC, the number of prefetchers to start will be calculated at database activation time based on the following formula:

```
number of prefetchers = max( max over all table spaces
( parallelism setting * [SMS: # containers;
  DMS: max # containers in stripe set] ), 3 )
```

Recommendation: In order to fully exploit all the I/O devices in the system, a good value to use is generally one or two more than the number of physical devices on which the database resides. It is better to configure additional I/O servers, since there is minimal overhead associated with each I/O server and any unused I/O servers will remain idle.

num_log_span - Number log span

This parameter specifies whether there is a limit to how many log files one transaction can span, and what that limit is.

Configuration type

Database

Parameter type
Configurable online

Propagation class
Immediate

Default [range]
0 [0 - 65 535]

Unit of measure
Counter

If the value is not 0, this parameter indicates the number of active log files that one active transaction is allowed to span.

If the value is set to 0, there is no limit to how many log files one single transaction can span. This was the behavior of transactions prior to Version 8.

num_quantiles - Number of quantiles for columns

This parameter controls the number of quantiles that will be collected when the WITH DISTRIBUTION option is specified on the RUNSTATS command.

Configuration type
Database

Parameter type
Configurable Online

Propagation class
Immediate

Default [range]
20 [0 - 32 767]

Unit of measure
Counter

Increasing the value of this parameter increases the amount of statistics heap (*stat_heap_sz*) used when collecting statistics.

The “quantile” statistics help the optimizer understand the distribution of data values within a column. A higher value results in more information being available to the query optimizer but requires additional catalog space. When 0 or 1 is specified, no quantile statistics are retained, even if you request that distribution statistics be collected.

You can also specify the number of quantiles collected as part of the RUNSTATS command at the table or the column level, by using the NUM_QUANTILES option. If none is specified, the *num_quantiles* configuration parameter value is used. Changing the number of quantiles that will be collected through the RUNSTATS command is easier than making the change using the *num_quantiles* database configuration parameter.

Updating this parameter can help obtain better selectivity estimates for range predicates over data that is non-uniformly distributed. Among other optimizer decisions, this information has a strong influence on whether an index scan or a table scan will be chosen. (It is more efficient to use a table scan to access a range of values that occur frequently and it is more efficient to use an index scan for a range of values that occur infrequently.)

After changing the value of this parameter, you need to:

- Run the RUNSTATS command again to collect statistics with the changed number of frequent values
- Rebind any packages containing static SQL or XQuery statements.

When using RUNSTATS, you have the ability to limit the number of quantiles collected at both the table level and the column level. This allows you to optimize on space occupied in the catalogs by reducing the distribution statistics for columns where they could not be exploited and yet still using the information for critical columns.

Recommendation: This default value for this parameter guarantees a maximum estimation error of approximately 2.5% for any single-sided range predicate (>, >=, <, or <=), and a maximum error of 5% for any BETWEEN predicate. A simple way to approximate the number of quantiles is:

- Determine the maximum error that is tolerable in estimating the number of rows of any range query, as a percentage, P.
- The number of quantiles should be approximately 100/P if most of your predicates are BETWEEN predicates, and 50/P if most of your predicates are other types of range predicates (<, <=, >, or >=).

For example, 25 quantiles should result in a maximum estimate error of 4% for BETWEEN predicates and of 2% for ">" predicates. A reasonable practical value for this parameter lies in the range of 10 to 50.

numarchretry - Number of retries on error

This parameter specifies the number of times that DB2 is to try archiving a log file to the primary or the secondary archive directory before trying to archive log files to the failover directory.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Default [range]

5 [0 - 65 535]

This parameter is only used if the *failarchpath* database configuration parameter is set. If *numarchretry* is not set, DB2 will continuously retry archiving to the primary or the secondary log path.

numsegs - Default number of SMS containers

This parameter is deprecated in Version 9.5, but is still being used by pre-Version 9.5 data servers and clients. Any value specified for this configuration parameter will be ignored by the DB2 Version 9.5 database manager.

Note: The following information applies only to pre-Version 9.5 data servers and clients.

Configuration type

Database

Parameter type

Informational

Unit of measure

Counter

This parameter indicates the number of containers that will be created within the default table spaces. It also shows the information used when you created your database, whether it was specified explicitly or implicitly on the CREATE DATABASE command.

This parameter only applies to SMS table spaces; the CREATE TABLESPACE statement **does not** use it in any way.

overflowlogpath - Overflow log path

This parameter specifies a location for DB2 to find log files needed for a rollforward operation, as well as where to store active log files retrieved from the archive. It also gives a location for finding and storing log files needed for using db2ReadLog API.

Configuration type

Database

Parameter type

Configurable online

Propagation class

Immediate

Default [range]

NULL [any valid path]

This parameter can be used for several functions, depending on your logging requirements.

- This parameter allows you to specify a location for DB2 to find log files that are needed for a rollforward operation. It is similar to the OVERFLOW LOG PATH option on the ROLLFORWARD command. Instead of always specifying OVERFLOW LOG PATH on every ROLLFORWARD command, you can set this configuration parameter once. However, if both are used, the OVERFLOW LOG PATH option will overwrite the *overflowlogpath* configuration parameter, for that particular rollforward operation.
- If *logsecond* is set to -1, *overflowlogpath* allows you to specify a directory for DB2 to store active log files retrieved from the archive. (Active log files have to be retrieved for rollback operations if they are no longer in the active log path). Without *overflowlogpath*, DB2 will retrieve the log files into the active log path. Using *overflowlogpath* allows you to provide additional resource for DB2 to store the retrieved log files. The benefit includes spreading the I/O cost to different disks, and allowing more log files to be stored in the active log path.
- If you need to use the db2ReadLog API (prior to DB2 V8, db2ReadLog was called sqlurlog) for replication, for example, *overflowlogpath* allows you to specify a location for DB2 to search for log files that are needed for this API. If the log file is not found (in either the active log path or the overflow log path) and the

database is configured with *userexit* enabled, DB2 will retrieve the log file. *overflowlogpath* also allows you to specify a directory for DB2 to store the log files retrieved. The benefit comes from reducing the I/O cost on the active log path and allowing more log files to be stored in the active log path.

- If you have configured a raw device for the active log path, *overflowlogpath* must be configured if you want to set *logsecond* to -1, or if you want to use the db2ReadLog API.

To set *overflowlogpath*, specify a string of up to 242 bytes. The string must point to a path name, and it must be a fully qualified path name, not a relative path name. The path name must be a directory, not a raw device.

Note: In a single or multi-partition DB2 ESE environment, the node number is automatically appended to the path. This is done to maintain the uniqueness of the path in multiple logical node configurations.

pagesize - Database default page size

This parameter contains the value that was used as the default page size when the database was created. Possible values are: 4 096, 8 192, 16 384 and 32 768. When a buffer pool or table space is created in that database, the same default page size applies.

Configuration type
Database

Parameter type
Informational

pckcachesz - Package cache size

This parameter is allocated out of the database shared memory, and is used for caching of sections for static and dynamic SQL and XQuery statements on a database.

Configuration type
Database

Parameter type
Configurable online

Propagation class
Immediate

Default [range]

32-bit platforms
Automatic [-1, 32 - 128 000]

64-bit platforms
Automatic [-1, 32 - 524 288]

Unit of measure
Pages (4 KB)

When allocated
When the database is initialized

When freed
When the database is shut down

In a partitioned database system, there is one package cache for each database partition.

Caching packages allows the database manager to reduce its internal overhead by eliminating the need to access the system catalogs when reloading a package; or, in the case of dynamic SQL or XQuery statements, eliminating the need for compilation. Sections are kept in the package cache until one of the following occurs:

- The database is shut down
- The package or dynamic SQL or XQuery statement is invalidated
- The cache runs out of space.

This caching of the section for a static or dynamic SQL or XQuery statement can improve performance, especially when the same statement is used multiple times by applications connected to a database. This is particularly important in a transaction processing environment.

When this parameter is set to `AUTOMATIC`, it is enabled for self tuning. When `self_tuning_mem` is set to `ON`, the memory tuner will dynamically size the memory area controlled by `pckcachesz` as the workload requirements change. Because the memory tuner trades memory resources between different memory consumers, there must be at least two memory consumers enabled for self tuning in order for self tuning to be active.

Automatic tuning of this configuration parameter will only occur when self tuning memory is enabled for the database (the `self_tuning_mem` configuration parameter is set to "ON.")

When this parameter is set to `-1`, the value used to calculate the page allocation is eight times the value specified for the `maxappls` configuration parameter. The exception to this occurs if eight times `maxappls` is less than 32. In this situation, the default value of `-1` will set `pckcachesz` to 32.

Recommendation: When tuning this parameter, you should consider whether the extra memory being reserved for the package cache might be more effective if it was allocated for another purpose, such as the buffer pool or catalog cache. For this reason, you should use benchmarking techniques when tuning this parameter.

Tuning this parameter is particularly important when several sections are used initially and then only a few are run repeatedly. If the cache is too large, memory is wasted holding copies of the initial sections.

The following monitor elements can help you determine whether you should adjust this configuration parameter:

- `pkg_cache_lookups` (package cache lookups)
- `pkg_cache_inserts` (package cache inserts)
- `pkg_cache_size_top` (package cache high water mark)
- `pkg_cache_num_overflows` (package cache overflows)

Note: The package cache is a working cache, so you cannot set this parameter to zero. There must be sufficient memory allocated in this cache to hold all sections of the SQL or XQuery statements currently being executed. If there is more space

allocated than currently needed, then sections are cached. These sections can simply be executed the next time they are needed without having to load or compile them.

The limit specified by the *pckcachesz* parameter is a soft limit. This limit can be exceeded, if required, if memory is still available in the database shared set. You can use the *pkg_cache_size_top* monitor element to determine the largest that the package cache has grown, and the *pkg_cache_num_overflows* monitor element to determine how many times the limit specified by the *pckcachesz* parameter has been exceeded.

priv_mem_thresh - Private memory threshold

This parameter is deprecated in Version 9.5, but is still being used by pre-Version 9.5 data servers and clients. Any value specified for this configuration parameter will be ignored by the DB2 Version 9.5 database manager.

Note: The following information applies only to pre-Version 9.5 data servers and clients.

Configuration type

Database manager

Applies to

- Database server with local and remote clients
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable

Default [range]

20 000 [-1; 32 - 112 000]

Unit of measure

Pages (4 KB)

This parameter is used to determine the amount of unused agent private memory that will be kept allocated, ready to be used by new agents that are started. It does not apply to Linux and UNIX platforms.

A value of -1 will cause this parameter to use the value of the *min_priv_mem* parameter.

Recommendation: When setting this parameter, you should consider the client connection/disconnection patterns as well as the memory requirements of other processes on the same machine.

If there is only a brief period during which many clients are concurrently connected to the database, a high threshold will prevent unused memory from being decommitted and made available to other processes. This case results in poor memory management which can affect other processes which require memory.

If the number of concurrent clients is more uniform and there are frequent fluctuations in this number, a high threshold will help to ensure memory is available for the client processes and reduce the overhead to allocate and deallocate memory.

rec_his_retentn - Recovery history retention period

This parameter specifies the number of days that historical information on backups will be retained.

Configuration type

Database

Parameter type

Configurable

Default [range]

366 [-1; 0 - 30 000]

Unit of measure

Days

If the recovery history file is not needed to keep track of backups, restores, and loads, this parameter can be set to a small number.

If value of this parameter is -1, the number of entries indicating full database backups (and any table space backups that are associated with the database backup) will correspond with the value specified by the *num_db_backups* parameter. Other entries in the recovery history file can only be pruned by explicitly using the available commands or APIs.

No matter how small the retention period, the most recent full database backup plus its restore set will always be kept, unless you use the PRUNE utility with the FORCE option.

restore_pending - Restore pending

This parameter states whether a RESTORE PENDING status exists in the database.

Configuration type

Database

Parameter type

Informational

restrict_access - Database has restricted access configuration parameter

This parameter indicates whether the database was created using the restrictive set of default actions. In other words, if it was created with the RESTRICTIVE clause in the CREATE DATABASE command.

Configuration type

Database

Parameter type

Informational

YES The RESTRICTIVE clause was used in the CREATE DATABASE command when this database was created.

NO The RESTRICTIVE clause was not used in the CREATE DATABASE command when this database was created.

rollfwd_pending - Roll forward pending indicator

This parameter informs you whether or not a roll-forward recovery is required, and where it is required.

Configuration type

Database

Parameter type

Informational

This parameter can indicate one of the following states:

- **DATABASE**, meaning that a roll-forward recovery procedure is required for this database
- **TABLESPACE**, meaning that one or more table spaces need to be rolled forward
- **NO**, meaning that the database is usable and no roll-forward recovery is required.

The recovery (using ROLLFORWARD DATABASE) must complete before you can access the database or table space.

self_tuning_mem- Self-tuning memory

This parameter determines whether the memory tuner will dynamically distribute available memory resources as required between memory consumers that are enabled for self-tuning.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

Single-database partition environments

ON [ON; OFF]

Multi-database partition environments

OFF [ON; OFF]

In a database that is migrated from an earlier version, *self_tuning_mem* will be set to OFF.

Because memory is being traded between memory consumers, there must be at least two memory consumers enabled for self-tuning in order for the memory tuner to be active. When *self_tuning_mem* is set to ON, but there are less than two memory consumers enabled for self-tuning, the memory tuner is inactive. (The exception to this is the sort heap memory area, which can be tuned regardless of whether other memory consumers are enabled for self-tuning or not.) When *database_memory* is set to a numeric value, it is considered enabled for self-tuning.

This parameter is ON by default in single database partition environments. In multi-database partition environments, it is OFF by default.

The memory consumers that can be enabled for self-tuning include:

- Buffer pools (controlled by the size parameter of the ALTER BUFFERPOOL and CREATE BUFFERPOOL statements)
- Package cache (controlled by the *pckcachesz* configuration parameter)
- Lock List (controlled by the *locklist* and *maxlocks* configuration parameters)
- Sort heap (controlled by the *sheapthres_shr* and *sortheap* configuration parameters)
- Database shared memory (controlled by the *database_memory* configuration parameter)

To view the current setting for this parameter, use the GET DATABASE CONFIGURATION command specifying the SHOW DETAIL parameter. The possible settings returned for this parameter are:

```
Self Tuning Memory      (SELF_TUNING_MEM) = OFF
Self Tuning Memory      (SELF_TUNING_MEM) = ON (Active)
Self Tuning Memory      (SELF_TUNING_MEM) = ON (Inactive)
Self Tuning Memory      (SELF_TUNING_MEM) = ON
```

The following values indicate:

- ON (Active) - the memory tuner is actively tuning the memory on the system
- ON (Inactive) - that although the parameter is set ON, self-tuning is not occurring because there are less than two memory consumers enabled for self-tuning
- ON without (Active) or (Inactive) - from a query without the SHOW DETAIL option, or without a database connection.

In partitioned environments, the *self_tuning_mem* configuration parameter will only show ON (Active) for the database partition on which the tuner is running. On all other nodes *self_tuning_mem* will show ON (Inactive). As a result, to determine if the memory tuner is active in a partitioned database, you must check the *self_tuning_mem* parameter on all database partitions.

If you have migrated to DB2 Version 9 from an earlier version of DB2 and you plan to use the self-tuning memory feature, you should configure the following health indicators to disable threshold or state checking:

- Shared Sort Memory Utilization - db.sort_shrmem_util
- Percentage of sorts that overflowed - db.spilled_sorts
- Long Term Shared Sort Memory Utilization - db.max_sort_shrmem_util
- Lock List Utilization - db.locklist_util
- Lock Escalation Rate - db.lock_escal_rate
- Package Cache Hit Ratio - db.pkgcache_hitratio

One of the objectives of the self-tuning memory feature is to avoid having memory allocated to a memory consumer when it is not immediately required. Therefore, utilization of the memory allocated to a memory consumer might approach 100% before more memory is allocated. By disabling these health indicators, you will avoid unnecessary alerts triggered by the high rate of memory utilization by a memory consumer.

Instances created in DB2 Version 9 will have these health indicators disabled by default.

seqdetect - Sequential detection flag

This parameter controls whether the database manager is allowed to detect sequential page reading during I/O activity.

Configuration type

Database

Parameter type

Configurable online

Propagation class

Immediate

Default [range]

Yes [Yes; No]

The database manager can monitor I/O, and if sequential page reading is occurring the database manager can activate I/O prefetching. This type of sequential prefetch is known as *sequential detection*.

If this parameter is set to No, prefetching takes place only if the database manager knows it will be useful, for example table sorts, table scans, or list prefetch.

Recommendation: In most cases, you should use the default value for this parameter. Try turning sequential detection off, only if other tuning efforts were unable to correct serious query performance problems.

sheapthres_shr - Sort heap threshold for shared sorts

This parameter represents a soft limit on the total amount of database shared memory that can be used by sort memory consumers at any one time.

Configuration type

Database

Applies to

OLAP functions

Parameter type

Configurable online

Propagation class

Immediate

Default [range]**32-bit platforms**

Automatic [250 - 524 288]

64-bit platforms

Automatic [250 - 2 147 483 647]

Unit of measure

Pages (4 KB)

There are other sort memory consumers in addition to sort, like hash join, index ANDing, block index ANDing, merge join, and in-memory tables. When the total amount of shared memory for shared sort memory consumers approaches the *sheapthres_shr* limit, a memory throttling mechanism is activated and the future shared sort memory consumer requests might be granted less memory than requested, but will always be granted more than the minimum they need for finishing the task. Once the *sheapthres_shr* limit is exceeded, all requests of shared sort memory from sort memory consumers will be granted the minimum amount of memory required to finish the task. When the total amount of shared memory for active shared sort memory consumers reaches this limit, subsequent sorts could fail (SQL0955C).

When the value of the database manager configuration parameter *sheapthres* is 0, all sort memory consumers for the database will use the database shared memory with *sheapthres_shr* instead of private sort memory.

When *sheapthres_shr* is set to AUTOMATIC, it is enabled for self tuning. This allows the memory tuner to dynamically size the memory area controlled by this parameter as the workload requirements change. Because the memory tuner trades memory resources between different memory consumers, there must be at least two memory consumers enabled for self tuning in order for self tuning to be active. Memory consumers include SHEAPTHRES_SHR, PCKCACHESZ, BUFFER POOL (each buffer pool counts as one), LOCKLIST, and DATABASE_MEMORY.

Automatic tuning of *sheapthres_shr* is allowed only when the database manager configuration parameter *sheapthres* is set to 0.

The value of *sortheap* is tuned together with the *sheapthres_shr* parameter therefore disabling self tuning of the *sortheap* parameter automatically disables self tuning of the *sheapthres_shr* parameter. Enabling self tuning of the *sheapthres_shr* parameter automatically enables self tuning of the *sortheap* parameter.

Automatic tuning of this configuration parameter will only occur when self tuning memory is enabled for the database (the *self_tuning_mem* configuration parameter is set to "ON.")

When the value of this parameter is updated online, only new requests of shared-sort memory made after the update will use the new value. It is recommended that you reduce the value of *sortheap* before reducing the value of *sheapthres_shr* and to increase the value of *sheapthres_shr* before increasing the value of *sortheap*.

When the database manager configuration parameter *sheapthres* is greater than 0, *sheapthres_shr* is only meaningful in two cases:

- if the *intra_parallel* database manager configuration parameter is set to *yes*, because when *intra_parallel* is set to *no*, there will be no shared sorts.
- if the Concentrator is on (that is, when *max_connections* is greater than *max_coordagents*), because sorts that use a cursor declared with the WITH HOLD option will be allocated from shared memory.

softmax - Recovery range and soft checkpoint interval

This parameter determines the frequency of soft checkpoints and the recovery range, which help out in the crash recovery process.

Configuration Type

Database

Parameter Type

Configurable

Default [range]

100 [1 – 100 * *logprimary*]

Unit of Measure

Percentage of the size of one primary log file

This parameter is used to:

- Influence the number of logs that need to be recovered following a crash (such as a power failure). For example, if the default value is used, the database

manager will try to keep the number of logs that need to be recovered to 1. If you specify 300 as the value of this parameter, the database manager will try to keep the number of logs that need to be recovered to 3.

To influence the number of logs required for crash recovery, the database manager uses this parameter to trigger the page cleaners to ensure that pages older than the specified recovery window are already written to disk.

- Determine the frequency of soft checkpoints.

At the time of a database failure resulting from an event such as a power failure, there might have been changes to the database which:

- Have not been committed, but updated the data in the buffer pool
- Have been committed, but have not been written from the buffer pool to the disk
- Have been committed and written from the buffer pool to the disk.

When a database is restarted, the log files will be used to perform a crash recovery of the database which ensures that the database is left in a consistent state (that is, all committed transactions are applied to the database and all uncommitted transactions are not applied to the database).

To determine which records from the log file need to be applied to the database, the database manager uses information recorded in a log control file. (The database manager actually maintains two copies of the log control file, `SQLLOGCTL.LFH.1` and `SQLLOGCTL.LFH.2`, so that if one copy is damaged, the database manager can still use the other copy.) These log control files are periodically written to disk, and, depending on the frequency of this event, the database manager might be applying log records of committed transactions or applying log records that describe changes that have already been written from the buffer pool to disk. These log records have no impact on the database, but applying them introduces some overhead into the database restart process.

The log control files are always written to disk when a log file is full, and during soft checkpoints. You can use this configuration parameter to trigger additional soft checkpoints.

The timing of soft checkpoints is based on the difference between the “current state” and the “recorded state”, given as a percentage of the `logfilesiz`. The “recorded state” is determined by the oldest valid log record indicated in the log control files on disk, while the “current state” is determined by the log control information in memory. (The oldest valid log record is the first log record that the recovery process would read.) The soft checkpoint will be taken if the value calculated by the following formula is greater than or equal to the value of this parameter:

$$(\text{ (space between recorded and current states) } / \text{ logfilesiz }) * 100$$

Recommendation: You might want to increase or reduce the value of this parameter, depending on whether your acceptable recovery window is greater than or less than one log file. Lowering the value of this parameter will cause the database manager both to trigger the page cleaners more often and to take more frequent soft checkpoints. These actions can reduce both the number of log records that need to be processed and the number of redundant log records that are processed during crash recovery.

Note however, that more page cleaner triggers and more frequent soft checkpoints increase the overhead associated with database logging, which can impact the

performance of the database manager. Also, more frequent soft checkpoints might not reduce the time required to restart a database, if you have:

- Very long transactions with few commit points.
- A very large buffer pool and the pages containing the committed transactions are not written back to disk very frequently. (Note that the use of asynchronous page cleaners can help avoid this situation.)

In both of these cases, the log control information kept in memory does not change frequently and there is no advantage in writing the log control information to disk, unless it has changed.

sortheap - Sort heap size

This parameter defines the maximum number of private memory pages to be used for private sorts, or the maximum number of shared memory pages to be used for shared sorts.

Configuration type

Database

Applies to

OLAP functions

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

32-bit platforms

Automatic [16 - 524 288]

64-bit platforms

Automatic [16 - 4 194 303]

Unit of measure

Pages (4 KB)

When allocated

As needed to perform sorts

When freed

When sorting is complete

If the sort is a private sort, then this parameter affects agent private memory. If the sort is a shared sort, then this parameter affects the database shared memory. Each sort has a separate sort heap that is allocated as needed, by the database manager. This sort heap is the area where data is sorted. If directed by the optimizer, a smaller sort heap than the one specified by this parameter is allocated using information provided by the optimizer.

When this parameter is set to AUTOMATIC, it is enabled for self tuning. This allows the memory tuner to dynamically size the memory area controlled by this parameter as the workload requirements change.

The value of *sortheap* is tuned together with the *sheapthres_shr* parameter, therefore disabling self tuning of the *sortheap* parameter can not be done without disabling self tuning of the *sheapthres_shr* parameter. Enabling self tuning of the *sheapthres_shr*

parameter automatically enables self tuning of the *sortheap* parameter. The *sortheap* parameter can, however, be enabled for self tuning without the *sheapthres_shr* parameter being AUTOMATIC.

Automatic tuning of *sortheap* is allowed only when the database manager configuration parameter *sheapthres* is set to 0.

Automatic tuning of this configuration parameter will only occur when self tuning memory is enabled for the database (the *self_tuning_mem* configuration parameter is set to "ON.")

Recommendation: When working with the sort heap, you should consider the following:

- Appropriate indexes can minimize the use of the sort heap.
- Hash join buffers, block index ANDing, merge join, table in memory and dynamic bitmaps (used for index ANDing and Star Joins) use sort heap memory. Increase the size of this parameter when these techniques are used.
- Increase the size of this parameter when frequent large sorts are required.
- When increasing the value of this parameter, you should examine whether the *sheapthres* and *sheapthres_shr* parameters in the database manager configuration file also need to be adjusted.
- The sort heap size is used by the optimizer in determining access paths. You should consider rebinding applications (using the REBIND command) after changing this parameter.

When the *sortheap* value is updated, the database manager will immediately start using this new value for any current or new sorts.

stat_heap_sz - Statistics heap size

This parameter indicates the *maximum* size of the heap used in collecting statistics using the RUNSTATS command.

With Version 9.5, this database configuration parameter has a default value of AUTOMATIC, meaning that it increases as needed until either the *appl_memory* limit is reached, or the *instance_memory* limit is reached.

Configuration type

Database

Parameter type

Configurable online

Default [range]

Automatic [1 096 - 524 288]

Unit of measure

Pages (4 KB)

When allocated

When the RUNSTATS utility is started

When freed

When the RUNSTATS utility is completed

Recommendation: The default setting of AUTOMATIC is recommended.

stmtheap - Statement heap size

This parameter specifies the size of the statement heap, which is used as a work space for the SQL or XQuery compiler during compilation of an SQL or XQuery statement.

With Version 9.5, this database configuration parameter has a default value of *AUTOMATIC*, meaning that it increases as needed until either the *appl_memory* limit is reached, or the *instance_memory* limit is reached.

Configuration type

Database

Parameter type

Configurable Online

Propagation class

Statement boundary

Default [range]

For both 32-bit and 64-bit platforms

Automatic [128 - 524 288]

Unit of measure

Pages (4 KB)

When allocated

For each statement during precompiling or binding

When freed

When precompiling or binding of each statement is complete

This area does not stay permanently allocated, but is allocated and released for every SQL or XQuery statement handled. Note that for dynamic SQL or XQuery statements, this work area will be used during execution of your program; whereas, for static SQL or XQuery statements, it is used during the bind process but not during program execution.

Recommendation: In most cases the default *AUTOMATIC* setting for this parameter is acceptable. When set to *AUTOMATIC*, there is an internal limit on the total amount of memory allocated during the dynamic programming join enumeration phase of compilation. If this limit is exceeded, the statement is compiled using greedy join enumeration, and is only limited by the amount of remaining *appl_memory* or *instance_memory*, or both. If your application is receiving *SQL0437W* warnings, and the runtime performance for your query is not acceptable, you might want to consider setting a sufficiently large manual *stmtheap* value to ensure that dynamic join enumeration is always used.

Note: Dynamic join enumeration occurs only at optimization classes 3 and higher (5 is the default).

territory - Database territory

This parameter shows the territory used to create the database. *territory* is used by the database manager when processing data that is territory sensitive.

Configuration type

Database

Parameter type

Informational

trackmod - Track modified pages enable

This parameter specifies whether the database manager will track database modifications so that the backup utility can detect which subsets of the database pages must be examined by an incremental backup and potentially included in the backup image.

Configuration type

Database

Parameter type

Configurable

Default [range]

No [Yes, No]

After setting this parameter to "Yes", you must take a full database backup in order to have a baseline against which incremental backups can be taken. Also, if this parameter is enabled and if a table space is created, then a backup must be taken which contains that table space. This backup could be either a database backup or a table space backup. Following the backup, incremental backups will be permitted to contain this table space.

tsm_mgmtclass - Tivoli Storage Manager management class

The Tivoli Storage Manager management class determines how the TSM server should manage the backup versions of the objects being backed up.

Configuration type

Database

Parameter type

Configurable

Default [range]

Null [any string]

The default is that there is no DB2-specified management class.

When performing any TSM backup, before using the management class specified by the database configuration parameter, TSM first attempts to bind the backup object to the management class specified in the INCLUDE-EXCLUDE list found in the TSM client options file. If a match is not found, the default TSM management class specified on the TSM server will be used. TSM will then rebind the backup object to the management class specified by the database configuration parameter.

Thus, the default management class, as well as the management class specified by the database configuration parameter, must contain a backup copy group, or the backup operation will fail.

tsm_nodename - Tivoli Storage Manager node name

This parameter is used to override the default setting for the node name associated with the Tivoli Storage Manager (TSM) product.

Configuration type

Database

Parameter type

Configurable online

Propagation class
Statement boundary

Default [range]
Null [any string]

The node name is needed to allow you to restore a database that was backed up to TSM from another node.

The default is that you can only restore a database from TSM on the same node from which you did the backup. It is possible for the *tsm_nodename* to be overridden during a backup done through DB2 (for example, with the BACKUP DATABASE command).

tsm_owner - Tivoli Storage Manager owner name

This parameter is used to override the default setting for the owner associated with the Tivoli Storage Manager (TSM) product.

Configuration type
Database

Parameter type
Configurable online

Propagation class
Statement boundary

Default [range]
Null [any string]

The owner name is needed to allow you to restore a database that was backed up to TSM from another node. It is possible for the *tsm_owner* to be overridden during a backup done through DB2 (for example, with the BACKUP DATABASE command).

Note: The owner name is case sensitive.

The default is that you can only restore a database from TSM on the same node from which you did the backup.

tsm_password - Tivoli Storage Manager password

This parameter is used to override the default setting for the password associated with the Tivoli Storage Manager (TSM) product.

Configuration type
Database

Parameter type
Configurable online

Propagation class
Statement boundary

Default [range]
Null [any string]

The password is needed to allow you to restore a database that was backed up to TSM from another node.

Note: If the *tsm_nodename* is overridden during a backup done with DB2 (for example, with the BACKUP DATABASE command), the *tsm_password* might also have to be set.

The default is that you can only restore a database from TSM on the same node from which you did the backup. It is possible for the *tsm_nodename* to be overridden during a backup done with DB2.

user_exit_status - User exit status indicator

If set to On, this parameter indicates that the database manager is enabled for roll-forward recovery and that the user exit program will be used to archive and retrieve log files when called by the database manager.

Configuration type

Database

Parameter type

Informational

userexit - User exit enable

This parameter is deprecated in Version 9.5, but is still being used by pre-Version 9.5 data servers and clients. Any value specified for this configuration parameter will be ignored by the DB2 Version 9.5 database manager.

Note: The following information applies only to pre-Version 9.5 data servers and clients.

If this parameter is enabled, log retention logging is performed regardless of how the *logretain* parameter is set. This parameter also indicates that a user exit program should be used to archive and retrieve the log files.

Configuration type

Database

Parameter type

Configurable

Default [range]

Off [On; Off]

Log files are archived when the log file is full. They are retrieved when the ROLLFORWARD utility needs to use them to restore a database.

After *logretain*, or *userexit*, or both of these parameters are enabled, you must make a full backup of the database. This state is indicated by the *backup_pending* flag parameter.

If both of these parameters are de-selected, roll-forward recovery becomes unavailable for the database because logs will no longer be retained. In this case, the database manager deletes all log files in the *logpath* directory (including online archive log files), allocates new active log files, and reverts to circular logging.

util_heap_sz - Utility heap size

This parameter indicates the maximum amount of memory that can be used simultaneously by the BACKUP, RESTORE, and LOAD (including load recovery) utilities.

Configuration type

Database

Parameter type

Configurable online

Propagation class

Immediate

Default [range]

5000 [16 - 524 288]

Unit of measure

Pages (4 KB)

When allocated

As required by the database manager utilities

When freed

When the utility no longer needs the memory

Recommendation: Use the default value unless your utilities run out of space, in which case you should increase this value. If memory on your system is constrained, you might want to lower the value of this parameter to limit the memory used by the database utilities. If the parameter is set too low, you might not be able to run utilities concurrently. You should update this parameter dynamically as needed. For a small number of utilities, set this parameter to a small value. For a large number of utilities, or for memory intensive utilities, you should set this parameter to a larger value.

vendoropt - Vendor options

This parameter specifies additional parameters that DB2 might need to use to communicate with storage systems during backup, restore, or load copy operations.

Configuration type

Database

Applies to

- Database server with local and remote clients
- Client
- Database server with local clients
- Partitioned database server with local and remote clients

Parameter type

Configurable Online

Default [range]

Null []

Restriction

You cannot use the **vendoropt** configuration parameter to specify vendor-specific options for snapshot backup or restore operations. You must use the **OPTIONS** parameter of the backup or restore utilities instead.

wlm_collect_int - Workload management collection interval configuration parameter

This parameter specifies a collect and reset interval, in minutes, for workload management (WLM) statistics.

Every x *wlm_collect_int* minutes, (where x is the value of the *wlm_collect_int* parameter) all workload management statistics are collected and sent to any active statistics event monitor; then the statistics are reset. If an active event monitor exists, depending on how it was created, the statistics are written either to file or to a table. If it does not exist, the statistics are only reset and not collected.

The collect and reset process is initiated from the catalog partition. The *wlm_collect_int* parameter must be specified on the catalog partition. It is not used on other partitions.

Configuration type

Database

Parameter type

Configurable online

Default [range]

0 [0 (no collection performed), 5 - 32 767]

The workload management statistics collected by a statistics event monitor can be used to monitor both short term and long term system behavior. A small interval can be used to obtain both short term and long term system behavior because the results can be merged together to obtain long term behavior. However, having to manually merge the results from different intervals complicates the analysis. If it's not required, a small interval unnecessarily increases the overhead. Therefore, reduce the interval to capture shorter term behavior, and increase the interval to reduce overhead when only analysis of long term behavior is sufficient.

The interval needs to be customized per database, not for each SQL request, or command invocation, or application. There are no other configuration parameters that need to be considered.

Note: All WLM statistics table functions return statistics that have been accumulated since the last time the statistics were reset. The statistics will be reset regularly on the interval specified by this configuration parameter.

DB2 Administration Server (DAS) configuration parameters

authentication - Authentication type DAS

Determines how and where authentication of a user takes place.

Configuration type

DB2 Administration Server

Applies to

DB2 Administration Server

Parameter type

Configurable

Default [range]

SERVER_ENCRYPT [SERVER_ENCRYPT; KERBEROS_ENCRYPT]

If **authentication** is SERVER_ENCRYPT, then the user ID and password are sent from the client to the server so authentication can take place on the server. User IDs and passwords sent over the network are encrypted.

A value of `KERBEROS_ENCRYPT` means that authentication is performed at a Kerberos server using the Kerberos security protocol for authentication.

Note: The `KERBEROS_ENCRYPT` authentication type is only supported on servers running Windows.

This parameter can only be updated from a Version 9 command line processor (CLP).

contact_host - Location of contact list

This parameter specifies the location where the contact information used for notification by the Scheduler and the Health Monitor is stored.

Configuration type

DB2 Administration Server

Applies to

DB2 Administration Server

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

Null [any valid DB2 administration server TCP/IP hostname]

The location is defined to be a DB2 administration server's TCP/IP hostname. Allowing `contact_host` to be located on a remote DAS provides support for sharing a contact list across multiple DB2 administration servers. If `contact_host` is not specified, the DAS assumes the contact information is local.

This parameter can only be updated from a Version 8 command line processor (CLP).

das_codepage - DAS code page

This parameter indicates the code page used by the DB2 administration server.

Configuration type

DB2 Administration Server

Applies to

DB2 Administration Server

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

Null [any valid DB2 code page]

If the parameter is null, then the default code page of the system is used. This parameter should be compatible with the locale of the local DB2 instances. Otherwise, the DB2 administration server cannot communicate with the DB2 instances.

This parameter can only be updated from a Version 8 command line processor (CLP).

das_territory - DAS territory

This parameter shows the territory used by the DB2 administration server.

Configuration type

DB2 Administration Server

Applies to

DB2 Administration Server

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

Null [any valid DB2 territory]

If the parameter is null, then the default territory of the system is used.

This parameter can only be updated from a Version 8 command line processor (CLP).

dasadm_group - DAS administration authority group name

This parameter defines the group name with DAS Administration (DASADM) authority for the DAS.

Configuration type

DB2 Administration Server

Applies to

DB2 Administration Server

Parameter type

Configurable

Default [range]

Null [any valid group name]

DASADM authority is the highest level of authority within the DAS.

DASADM authority is determined by the security facilities used in a specific operating environment.

- For the Windows operating systems, this parameter can be set to any local group that is defined in the Windows security database. Group names are accepted as long as they are 30 bytes or less in length. If "NULL" is specified for this parameter, all members of the Administrators group have DASADM authority.
- For Linux and UNIX systems, if "NULL" is specified as the value of this parameter, the DASADM group defaults to the primary group of the instance owner.

If the value is not "NULL", the DASADM group can be any valid UNIX group name.

This parameter can only be updated from a Version 8 command line processor (CLP).

db2system - Name of the DB2 server system

This parameter specifies the name that is used by your users and database administrators to identify the DB2 server system.

Configuration type

DB2 Administration Server

Applies to

DB2 Administration Server

Parameter type

Configurable Online

Default [range]

TCP/IP host name [any valid system name]

If possible, this name should be unique within your network.

This name is displayed in the system level of the Control Center's object tree to aid administrators in the identification of server systems that can be administered from the Control Center.

When using the 'Search the Network' function of the Configuration Assistant, DB2 discovery returns this name and it is displayed at the system level in the resulting object tree. This name aids users in identifying the system that contains the database they want to access. A value for *db2system* is set at installation time as follows:

- On Windows, the setup program sets it equal to the computer name specified for the Windows system.
- On UNIX systems, it is set equal to the UNIX system's TCP/IP hostname.

discover - DAS discovery mode

This parameter determines the type of discovery mode that is started when the DB2 Administration Server starts.

Configuration type

DB2 Administration Server

Applies to

DB2 Administration Server

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

SEARCH [DISABLE; KNOWN; SEARCH]

- If discover = SEARCH, the administration server handles SEARCH discovery requests from clients. SEARCH provides a superset of the functionality provided by KNOWN discovery. When discover = SEARCH, the administration server will handle both SEARCH and KNOWN discovery requests from clients.
- If discover = KNOWN, the administration server handles only KNOWN discovery requests from clients.
- If discover = DISABLE, then the administration server will not handle any type of discovery request. The information for this server system is essentially hidden from clients.

The default discovery mode is SEARCH.

This parameter can only be updated from a Version 8 command line processor (CLP).

exec_exp_task - Execute expired tasks

This parameter specifies whether or not the Scheduler will execute tasks that have been scheduled in the past, but have not yet been executed.

Configuration type

DB2 Administration Server

Applies to

DB2 Administration Server

Parameter type

Configurable

Default [range]

No [Yes; No]

The Scheduler only detects expired tasks when it starts up. For example, if you have a job scheduled to run every Saturday, and the Scheduler is turned off on Friday and then restarted on Monday, the job scheduled for Saturday is now a job that is scheduled in the past. If *exec_exp_task* is set to Yes, your Saturday job will run when the Scheduler is restarted.

This parameter can only be updated from a Version 8 command line processor (CLP).

jdk_path - Software Developer's Kit for Java installation path DAS

This parameter specifies the directory under which the Software Developer's Kit (SDK) for Java, to be used for running DB2 administration server functions, is installed.

Configuration type

DB2 Administration Server

Applies to

DB2 Administration Server

Parameter type

Configurable Online

Propagation class

Immediate

Default [range]

Default Java install path [any valid path]

Environment variables used by the Java interpreter are computed from the value of this parameter.

On Windows operating systems, Java files (if needed) are placed under the sqllib directory (in java\jdk) during DB2 installation. The *jdk_path* configuration parameter is then set to sqllib\java\jdk. Java is never actually installed by DB2 on Windows platforms; the files are merely placed under the sqllib directory, and this is done regardless of whether or not Java is already installed.

This parameter can only be updated from a Version 8 command line processor (CLP).

sched_enable - Scheduler mode

This parameter indicates whether or not the Scheduler is started by the administration server.

Configuration type
DB2 Administration Server

Applies to
DB2 Administration Server

Parameter type
Configurable

Default [range]
Off [On; Off]

The Scheduler allows tools such as the Task Center to schedule and execute tasks at the administration server.

This parameter can only be updated from a Version 8 command line processor (CLP).

sched_userid - Scheduler user ID

This parameter specifies the user ID used by the Scheduler to connect to the tools catalog database. This parameter is only relevant if the tools catalog database is remote to the DB2 administration server.

Configuration type
DB2 Administration Server

Applies to
DB2 Administration Server

Parameter type
Informational

Default [range]
Null [any valid user ID]

The userid and password used by the Scheduler to connect to the remote tools catalog database are specified using the db2admin command.

smtp_server - SMTP server

When the Scheduler is on, this parameter identifies the SMTP server that the Scheduler will use to send e-mail and pager notifications.

Configuration type
DB2 Administration Server

Applies to
DB2 Administration Server

Parameter type
Configurable Online

Propagation class
Immediate

Default [range]

Null [any valid SMTP server TCP/IP hostname]

This parameter is used by the Scheduler and the Health Monitor.

This parameter can only be updated from a Version 8 command line processor (CLP).

toolscat_db - Tools catalog database

This parameter indicates the tools catalog database used by the Scheduler.

Configuration type

DB2 Administration Server

Applies to

DB2 Administration Server

Parameter type

Configurable

Default [range]

Null [any valid database alias]

This database must be in the database directory of the instance specified by *toolscat_inst*.

This parameter can only be updated from a Version 8 command line processor (CLP).

toolscat_inst - Tools catalog database instance

This parameter indicates the instance name that is used by the Scheduler, along with *toolscat_db* and *toolscat_schema*, to identify the tools catalog database.

Configuration type

DB2 Administration Server

Applies to

DB2 Administration Server

Parameter type

Configurable

Default [range]

Null [any valid instance]

The tools catalog database contains task information created by the Task Center and the Control Center. The tools catalog database must be listed in the database directory of the instance specified by this configuration parameter. The database can be local or remote. If the tools catalog database is local, the instance must be configured for TCP/IP. If the database is remote, the database partition cataloged in the database directory must be a TCP/IP node.

This parameter can only be updated from a Version 8 command line processor (CLP).

toolscat_schema - Tools catalog database schema

This parameter indicates the schema of the tools catalog database used by the Scheduler.

Configuration type
DB2 Administration Server

Applies to
DB2 Administration Server

Parameter type
Configurable

Default [range]
Null [any valid schema]

The schema is used to uniquely identify a set of tools catalog tables and views within the database.

This parameter can only be updated from a Version 8 command line processor (CLP).

Part 5. Appendixes

Appendix A. Overview of the DB2 technical information

DB2 technical information is available through the following tools and methods:

- *DB2 Information Center*
 - Topics (Task, concept and reference topics)
 - Help for DB2 tools
 - Sample programs
 - Tutorials
- DB2 books
 - PDF files (downloadable)
 - PDF files (from the DB2 PDF DVD)
 - printed books
- Command line help
 - Command help
 - Message help

Note: The *DB2 Information Center* topics are updated more frequently than either the PDF or the hard-copy books. To get the most current information, install the documentation updates as they become available, or refer to the *DB2 Information Center* at ibm.com[®].

You can access additional DB2 technical information such as technotes, white papers, and IBM Redbooks[®] publications online at [ibm.com](http://www.ibm.com). Access the DB2 Information Management software library site at <http://www.ibm.com/software/data/sw-library/>.

Documentation feedback

We value your feedback on the DB2 documentation. If you have suggestions for how to improve the DB2 documentation, send an email to db2docs@ca.ibm.com. The DB2 documentation team reads all of your feedback, but cannot respond to you directly. Provide specific examples wherever possible so that we can better understand your concerns. If you are providing feedback on a specific topic or help file, include the topic title and URL.

Do not use this email address to contact DB2 Customer Support. If you have a DB2 technical issue that the documentation does not resolve, contact your local IBM service center for assistance.

If you would like to help IBM make the IBM Information Management products easier to use, take the Consumability Survey: <http://www.ibm.com/software/data/info/consumability-survey/>.

DB2 technical library in hardcopy or PDF format

The following tables describe the DB2 library available from the IBM Publications Center at www.ibm.com/shop/publications/order. English DB2 Version 9.5 manuals in PDF format and translated versions can be downloaded from www.ibm.com/support/docview.wss?rs=71&uid=swg2700947.

Although the tables identify books available in print, the books might not be available in your country or region.

The form number increases each time a manual is updated. Ensure that you are reading the most recent version of the manuals, as listed below.

Note: The *DB2 Information Center* is updated more frequently than either the PDF or the hard-copy books.

Table 73. DB2 technical information

Name	Form Number	Available in print	Last updated
<i>Administrative API Reference</i>	SC23-5842-02	Yes	April, 2009
<i>Administrative Routines and Views</i>	SC23-5843-02	No	April, 2009
<i>Call Level Interface Guide and Reference, Volume 1</i>	SC23-5844-02	Yes	April, 2009
<i>Call Level Interface Guide and Reference, Volume 2</i>	SC23-5845-02	Yes	April, 2009
<i>Command Reference</i>	SC23-5846-02	Yes	April, 2009
<i>Data Movement Utilities Guide and Reference</i>	SC23-5847-02	Yes	April, 2009
<i>Data Recovery and High Availability Guide and Reference</i>	SC23-5848-02	Yes	April, 2009
<i>Data Servers, Databases, and Database Objects Guide</i>	SC23-5849-02	Yes	April, 2009
<i>Database Security Guide</i>	SC23-5850-02	Yes	April, 2009
<i>Developing ADO.NET and OLE DB Applications</i>	SC23-5851-02	Yes	April, 2009
<i>Developing Embedded SQL Applications</i>	SC23-5852-02	Yes	April, 2009
<i>Developing Java Applications</i>	SC23-5853-02	Yes	April, 2009
<i>Developing Perl and PHP Applications</i>	SC23-5854-02	No	April, 2009
<i>Developing User-defined Routines (SQL and External)</i>	SC23-5855-02	Yes	April, 2009
<i>Getting Started with Database Application Development</i>	GC23-5856-02	Yes	April, 2009
<i>Getting Started with DB2 installation and administration on Linux and Windows</i>	GC23-5857-02	Yes	April, 2009
<i>Internationalization Guide</i>	SC23-5858-02	Yes	April, 2009

Table 73. DB2 technical information (continued)

Name	Form Number	Available in print	Last updated
<i>Message Reference, Volume 1</i>	GI11-7855-01	No	April, 2009
<i>Message Reference, Volume 2</i>	GI11-7856-01	No	April, 2009
<i>Migration Guide</i>	GC23-5859-02	Yes	April, 2009
<i>Net Search Extender Administration and User's Guide</i>	SC23-8509-02	Yes	April, 2009
<i>Partitioning and Clustering Guide</i>	SC23-5860-02	Yes	April, 2009
<i>Query Patroller Administration and User's Guide</i>	SC23-8507-01	Yes	April, 2009
<i>Quick Beginnings for IBM Data Server Clients</i>	GC23-5863-02	No	April, 2009
<i>Quick Beginnings for DB2 Servers</i>	GC23-5864-02	Yes	April, 2009
<i>Spatial Extender and Geodetic Data Management Feature User's Guide and Reference</i>	SC23-8508-02	Yes	April, 2009
<i>SQL Reference, Volume 1</i>	SC23-5861-02	Yes	April, 2009
<i>SQL Reference, Volume 2</i>	SC23-5862-02	Yes	April, 2009
<i>System Monitor Guide and Reference</i>	SC23-5865-02	Yes	April, 2009
<i>Text Search Guide</i>	SC23-5866-01	Yes	April, 2009
<i>Troubleshooting Guide</i>	GI11-7857-02	No	April, 2009
<i>Tuning Database Performance</i>	SC23-5867-02	Yes	April, 2009
<i>Visual Explain Tutorial</i>	SC23-5868-00	No	
<i>What's New</i>	SC23-5869-02	Yes	April, 2009
<i>Workload Manager Guide and Reference</i>	SC23-5870-02	Yes	April, 2009
<i>pureXML Guide</i>	SC23-5871-02	Yes	April, 2009
<i>XQuery Reference</i>	SC23-5872-02	No	April, 2009

Table 74. DB2 Connect-specific technical information

Name	Form Number	Available in print	Last updated
<i>Quick Beginnings for DB2 Connect Personal Edition</i>	GC23-5839-02	Yes	April, 2009
<i>Quick Beginnings for DB2 Connect Servers</i>	GC23-5840-02	Yes	April, 2009
<i>DB2 Connect User's Guide</i>	SC23-5841-02	Yes	April, 2009

Table 75. Information Integration technical information

Name	Form Number	Available in print	Last updated
Information Integration: Administration Guide for Federated Systems	SC19-1020-01	Yes	March, 2008
Information Integration: ASNCLP Program Reference for Replication and Event Publishing	SC19-1018-02	Yes	March, 2008
Information Integration: Configuration Guide for Federated Data Sources	SC19-1034-01	No	
Information Integration: SQL Replication Guide and Reference	SC19-1030-01	Yes	March, 2008
Information Integration: Introduction to Replication and Event Publishing	SC19-1028-01	Yes	March, 2008

Ordering printed DB2 books

If you require printed DB2 books, you can buy them online in many but not all countries or regions. You can always order printed DB2 books from your local IBM representative. Keep in mind that some softcopy books on the *DB2 PDF Documentation DVD* are unavailable in print. For example, neither volume of the *DB2 Message Reference* is available as a printed book.

Printed versions of many of the DB2 books available on the DB2 PDF Documentation DVD can be ordered for a fee from IBM. Depending on where you are placing your order from, you may be able to order books online, from the IBM Publications Center. If online ordering is not available in your country or region, you can always order printed DB2 books from your local IBM representative. Note that not all books on the DB2 PDF Documentation DVD are available in print.

Note: The most up-to-date and complete DB2 documentation is maintained in the DB2 Information Center at <http://publib.boulder.ibm.com/infocenter/db2luw/v9r5>.

To order printed DB2 books:

- To find out whether you can order printed DB2 books online in your country or region, check the IBM Publications Center at <http://www.ibm.com/shop/publications/order>. You must select a country, region, or language to access publication ordering information and then follow the ordering instructions for your location.
- To order printed DB2 books from your local IBM representative:
 1. Locate the contact information for your local representative from one of the following Web sites:
 - The IBM directory of world wide contacts at www.ibm.com/planetwide
 - The IBM Publications Web site at <http://www.ibm.com/shop/publications/order>. You will need to select your country, region, or

- language to the access appropriate publications home page for your location. From this page, follow the "About this site" link.
2. When you call, specify that you want to order a DB2 publication.
 3. Provide your representative with the titles and form numbers of the books that you want to order. For titles and form numbers, see "DB2 technical library in hardcopy or PDF format" on page 607.

Displaying SQL state help from the command line processor

DB2 returns an SQLSTATE value for conditions that could be the result of an SQL statement. SQLSTATE help explains the meanings of SQL states and SQL state class codes.

To invoke SQL state help, open the command line processor and enter:

```
? sqlstate or ? class code
```

where *sqlstate* represents a valid five-digit SQL state and *class code* represents the first two digits of the SQL state.

For example, ? 08003 displays help for the 08003 SQL state, and ? 08 displays help for the 08 class code.

Accessing different versions of the DB2 Information Center

For DB2 Version 9.5 topics, the DB2 Information Center URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/>

For DB2 Version 9 topics, the DB2 Information Center URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9/>

For DB2 Version 8 topics, go to the Version 8 Information Center URL at: <http://publib.boulder.ibm.com/infocenter/db2luw/v8/>

Displaying topics in your preferred language in the DB2 Information Center

The DB2 Information Center attempts to display topics in the language specified in your browser preferences. If a topic has not been translated into your preferred language, the DB2 Information Center displays the topic in English.

- To display topics in your preferred language in the Internet Explorer browser:
 1. In Internet Explorer, click the **Tools** —> **Internet Options** —> **Languages...** button. The Language Preferences window opens.
 2. Ensure your preferred language is specified as the first entry in the list of languages.
 - To add a new language to the list, click the **Add...** button.

Note: Adding a language does not guarantee that the computer has the fonts required to display the topics in the preferred language.

- To move a language to the top of the list, select the language and click the **Move Up** button until the language is first in the list of languages.
- 3. Clear the browser cache and then refresh the page to display the DB2 Information Center in your preferred language.
- To display topics in your preferred language in a Firefox or Mozilla browser:

1. Select the button in the **Languages** section of the **Tools** —> **Options** —> **Advanced** dialog. The Languages panel is displayed in the Preferences window.
2. Ensure your preferred language is specified as the first entry in the list of languages.
 - To add a new language to the list, click the **Add...** button to select a language from the Add Languages window.
 - To move a language to the top of the list, select the language and click the **Move Up** button until the language is first in the list of languages.
3. Clear the browser cache and then refresh the page to display the DB2 Information Center in your preferred language.

On some browser and operating system combinations, you might have to also change the regional settings of your operating system to the locale and language of your choice.

Updating the DB2 Information Center installed on your computer or intranet server

If you have installed the DB2 Information Center locally, you can obtain and install documentation updates from IBM.

Updating your locally-installed *DB2 Information Center* requires that you:

1. Stop the *DB2 Information Center* on your computer, and restart the Information Center in stand-alone mode. Running the Information Center in stand-alone mode prevents other users on your network from accessing the Information Center, and allows you to apply updates. Non-Administrative and Non-Root *DB2 Information Centers* always run in stand-alone mode. .
2. Use the update feature to see what updates are available. If there are updates that you would like to install, you can use the update feature to obtain and install them.

Note: If your environment requires installing the *DB2 Information Center* updates on a machine that is not connected to the internet, you have to mirror the update site to a local file system using a machine that is connected to the internet and has the *DB2 Information Center* installed. If many users on your network will be installing the documentation updates, you can reduce the time required for individuals to perform the updates by also mirroring the update site locally and creating a proxy for the update site. If update packages are available, use the update feature to get the packages. However, the update feature is only available in stand-alone mode.

3. Stop the stand-alone Information Center, and restart the *DB2 Information Center* on your computer.

Note: On Windows Vista, the commands listed below must be run as an administrator. To launch a command prompt or graphical tool with full administrator privileges, right-click on the shortcut and then select **Run as administrator**.

To update the *DB2 Information Center* installed on your computer or intranet server:

1. Stop the *DB2 Information Center*.
 - On Windows, click **Start** → **Control Panel** → **Administrative Tools** → **Services**. Then right-click on **DB2 Information Center** service and select **Stop**.

- On Linux, enter the following command:
`/etc/init.d/db2icdv95 stop`
2. Start the Information Center in stand-alone mode.
 - On Windows:
 - a. Open a command window.
 - b. Navigate to the path where the Information Center is installed. By default, the *DB2 Information Center* is installed in the `Program_files\IBM\DB2 Information Center\Version 9.5` directory, where `Program_files` represents the location of the Program Files directory.
 - c. Navigate from the installation directory to the `doc\bin` directory.
 - d. Run the `help_start.bat` file:
`help_start.bat`
 - On Linux:
 - a. Navigate to the path where the Information Center is installed. By default, the *DB2 Information Center* is installed in the `/opt/ibm/db2ic/V9.5` directory.
 - b. Navigate from the installation directory to the `doc/bin` directory.
 - c. Run the `help_start` script:
`help_start`

The systems default Web browser launches to display the stand-alone Information Center.

3. Click the **Update** button (🔧). On the right hand panel of the Information Center, click **Find Updates**. A list of updates for existing documentation displays.
4. To initiate the installation process, check the selections you want to install, then click **Install Updates**.
5. After the installation process has completed, click **Finish**.
6. Stop the stand-alone Information Center:
 - On Windows, navigate to the installation directory's `doc\bin` directory, and run the `help_end.bat` file:
`help_end.bat`

Note: The `help_end` batch file contains the commands required to safely terminate the processes that were started with the `help_start` batch file. Do not use `Ctrl-C` or any other method to terminate `help_start.bat`.

 - On Linux, navigate to the installation directory's `doc/bin` directory, and run the `help_end` script:
`help_end`

Note: The `help_end` script contains the commands required to safely terminate the processes that were started with the `help_start` script. Do not use any other method to terminate the `help_start` script.
7. Restart the *DB2 Information Center*.
 - On Windows, click **Start** → **Control Panel** → **Administrative Tools** → **Services**. Then right-click on **DB2 Information Center** service and select **Start**.
 - On Linux, enter the following command:
`/etc/init.d/db2icdv95 start`

The updated *DB2 Information Center* displays the new and updated topics.

DB2 tutorials

The DB2 tutorials help you learn about various aspects of DB2 products. Lessons provide step-by-step instructions.

Before you begin

You can view the XHTML version of the tutorial from the Information Center at <http://publib.boulder.ibm.com/infocenter/db2help/>.

Some lessons use sample data or code. See the tutorial for a description of any prerequisites for its specific tasks.

DB2 tutorials

To view the tutorial, click on the title.

“pureXML” in *pureXML Guide*

Set up a DB2 database to store XML data and to perform basic operations with the native XML data store.

“Visual Explain” in *Visual Explain Tutorial*

Analyze, optimize, and tune SQL statements for better performance using Visual Explain.

DB2 troubleshooting information

A wide variety of troubleshooting and problem determination information is available to assist you in using DB2 database products.

DB2 documentation

Troubleshooting information can be found in the DB2 Troubleshooting Guide or the Database fundamentals section of the DB2 Information Center. There you will find information on how to isolate and identify problems using DB2 diagnostic tools and utilities, solutions to some of the most common problems, and other advice on how to solve problems you might encounter with your DB2 database products.

DB2 Technical Support Web site

Refer to the DB2 Technical Support Web site if you are experiencing problems and want help finding possible causes and solutions. The Technical Support site has links to the latest DB2 publications, TechNotes, Authorized Program Analysis Reports (APARs or bug fixes), fix packs, and other resources. You can search through this knowledge base to find possible solutions to your problems.

Access the DB2 Technical Support Web site at http://www.ibm.com/software/data/db2/support/db2_9/

Terms and Conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal use: You may reproduce these Publications for your personal, non commercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these Publications, or any portion thereof, without the express consent of IBM.

Commercial use: You may reproduce, distribute and display these Publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these Publications, or reproduce, distribute or display these Publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the Publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Appendix B. 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/region or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country/region 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.

This document may provide links or references to non-IBM Web sites and resources. IBM makes no representations, warranties, or other commitments whatsoever about any non-IBM Web sites or third-party resources that may be referenced, accessible from, or linked from this document. A link to a non-IBM Web site does not mean that IBM endorses the content or use of such Web site or

its owner. In addition, IBM is not a party to or responsible for any transactions you may enter into with third parties, even if you learn of such parties (or use a link to such parties) from an IBM site. Accordingly, you acknowledge and agree that IBM is not responsible for the availability of such external sites or resources, and is not responsible or liable for any content, services, products, or other materials on or available from those sites or resources. Any software provided by third parties is subject to the terms and conditions of the license that accompanies that software.

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

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

IBM Canada Limited
Office of the Lab Director
8200 Warden Avenue
Markham, Ontario
L6G 1C7
CANADA

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

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

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

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

This information may contain 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 may contain 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.

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

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

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

The following terms are trademarks or registered trademarks of other companies

- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Intel trademark information
- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

Index

A

- active directory
 - configuring DB2 95
 - DB2 objects 96
 - extending the directory schema 97
 - Lightweight Directory Access Protocol (LDAP) 77
 - security 96
 - support 95
- ADC (automatic dictionary creation) 47
- ADMIN_COPY_SCHEMA procedure
 - example of schema copy 218
- AFTER triggers 303
- agent pool size configuration parameter 498
- agent process
 - applheapsz configuration parameter 520
 - aslheapsz configuration parameter 463
 - maximum number of agents 493
 - maximum number of concurrent agents 494
 - priority of agents configuration parameter 462
- agent_stack_sz database manager configuration parameter 460
- agentpri database manager configuration parameter 462
- agents
 - configuration 35
 - configuration parameters affecting number of 456
 - overview 36
- AIX
 - large page support 4
 - system commands
 - vmo 4
 - vmtune 4
- alert summaries
 - DB2 Health Monitor 124
- aliases
 - chaining process 261
 - creating 130
 - dropping 139
 - overview 261
- alt_collate configuration parameter 517
- ALTER COLUMN clause
 - in table columns 259
- alter materialized query table properties 257
- ALTER TABLESPACE statement
 - examples 193
- ALTER triggers
 - description 302
- app_ctl_heap_sz database configuration parameter 517
- append mode tables
 - comparison with other table types 225
- appgroup_mem_sz database manager configuration parameter 518
- appl_memory configuration parameter
 - interaction between memory parameters 25
- appl_memory database configuration parameter 519
- application control heap size configuration parameter 517
- application performance
 - comparison of sequences and identity columns 324
 - sequences 323
- application process
 - connection states 135
- application programs
 - controlling sequences 322
- application support layer heap size configuration parameter 463
- application-directed distributed unit of work 134
- applications
 - control heap, setting 517
 - maximum number of coordinating agents at node 491
 - requesters 131
- archretrydelay configuration parameter 521
- aslheapsz configuration parameter 463
- ATTACH command 72
- attributes
 - Netscape LDAP 91
- audit_buf_sz configuration parameter 464
- authentication
 - trust all clients configuration parameter 515
 - trusted clients authentication configuration parameter 515
- authentication configuration parameter 465
- authentication DAS configuration parameter 596
- AUTHID identifier
 - restrictions 347
- authorities
 - defining group names
 - system administration authority group name configuration parameter 510
 - system control authority group name configuration parameter 510
 - system maintenance authority group name configuration parameter 511
- auto restart enable configuration parameter 524
- auto_del_rec_obj database configuration parameter 521
- auto_maint configuration parameter 522
- Autoconfigure API 56
- AUTOCONFIGURE command 56
 - sample output 56
- automatic
 - prefetch size adjustment
 - after adding or dropping containers 175
 - table space re-sizing 40, 171
- automatic dictionary creation (ADC) 47
- automatic features 17
 - enabled by default 17
- automatic maintenance
 - about 19
 - windows 19
- automatic memory tuning 28
- automatic statistics collection 54
 - description 17
- automatic storage
 - about 38
 - description 17
 - for databases 44, 121
 - restrictions 47, 124
 - table spaces 38, 163
- automatic storage path
 - adding 46, 124
- automatic storage table spaces
 - altering 206
- autonomic computing
 - about 17

avg_appls configuration parameter 524

B

backup

track modified pages 592

backup_pending configuration parameter 525

base tables

comparison with other table types 225

BEFORE DELETE triggers

description 302

BEFORE triggers 303

comparison with check constraints 276

description 302

bi-directional indexes 290

binding

changing configuration parameters 439, 440

database utilities 129

blank data type 235

blk_log_dsk_ful configuration parameter

details 525

block-structured devices 189

books

printed

ordering 610

buffer pools

about 143

creating 146

designing 143

dropping 148

effect on query optimization 456

memory (protection) 145

modifying 147

built-in functions 245

bypass federated authentication configuration parameter 479

C

caching

file system for table spaces 176

call level interface (CLI)

binding to a database 129

capacity

expansion 3

catalog cache size configuration parameter 526

CATALOG DATABASE command

example 128

catalog views

description 336

catalog_noauth configuration parameter 466

catalogcache_sz database configuration parameter 526

change the database log path configuration parameter 571

character serial devices 189

check constraints 267

comparison with BEFORE triggers 276

description 235

designing 275

check option

in views

examples 337

chnpgps_thresh configuration parameter 527

CIO/DIO

using as the default 178

client I/O block size configuration parameter 502

client interface copy

default 7

client support

client I/O block size configuration parameter 502

TCP/IP service name configuration parameter 509

clnt_krb_plugin configuration parameter 467

clnt_pw_plugin configuration parameter 467

cluster managers

cluster manager name configuration parameter 468

cluster_mgr configuration parameter 468

clustered indexes 290

clustering indexes

guidelines and considerations 292

code pages

database configuration parameter 528

codepage database configuration parameter 528

codeset database configuration parameter 528

collate_info database configuration parameter 528

column data

constraining 234

column data type

specifying 227

column properties

changing 258

columns

altering 259

definition

modifying 259

implicitly hidden 253

implicitly hidden 246

ordering 242

comm_bandwidth database manager configuration parameter

description 468

effect on query optimization 456

command line processor (CLP)

binding to a database 129

commit

number of commits to group (mincommit) 569

communications

connection elapse time 469

compression

data rows 49, 243

compression dictionary creation

automated 17

concurrency control

maximum number of active applications 564

concurrent transactions 131

configuration

agent and process model 35

changing database parameters 440

file system caching 179

memory heaps 32

Configuration Advisor

about 55

defining the scope of configuration parameters 55

description 17

generating recommended values 56

sample output 56

configuration file release level configuration parameter 501

configuration files

description 439

location 439

configuration parameters

affecting number of agents 456

affecting query optimization 456

agent_stack_sz 460

agentpri 462

alt_collate 517

app_ctl_heap_sz 517

configuration parameters (*continued*)

appgroup_mem_sz 518
 appl_memory 519
 applheapsz 520
 archretrydelay 521
 aslheapsz 463
 audit_buf_sz 464
 authentication 465
 authentication (DAS) 596
 auto_del_rec_obj 521
 auto_maint 522
 autorestart 524
 avg_appls 524
 backup_pending 525
 blk_log_dsk_ful 525
 catalog_noauth 466
 catalogcache_sz 526
 chngpgs_thresh 527
 clnt_krb_plugin 467
 clnt_pw_plugin 467
 cluster_mgr 468
 codepage 528
 codeset 528
 collate_info 528
 comm_bandwidth 468
 conn_elapse 469
 contact_host 597
 cpuspeed 470
 das_codepage 597
 das_territory 598
 dasadm_group 598
 database
 changing 439
 database_consistent 529
 database_level 530
 database_memory 530
 db_mem_thresh 532
 db2system 599
 dbheap 533
 decflt_rounding 534
 description 439
 dft_account_str 470
 dft_degree 536
 dft_extent_sz 536
 dft_loadrec_ses 537
 dft_monswitches 471
 dft_mttb_types 537
 dft_prefetch_sz 538
 dft_queryopt 539
 dft_refresh_age 539
 dft_sqlmathwarn 540
 dftdbpath 472
 diaglevel 473
 diagpath 474
 dir_cache 475
 discover 476
 discover (DAS) 599
 discover_db 541
 discover_inst 477
 dlchktime 541
 dyn_query_mgmt 542
 enable_xmlchar 543
 exec_exp_task 600
 failarchpath 543
 fcm_num_buffers 477
 fcm_num_channels 478
 fed_noauth 479

configuration parameters (*continued*)

federated 479
 federated_async 480
 fenced_pool 480
 group_plugin 482
 groupheap_ratio 544
 hadr_db_role 544
 hadr_local_host 544
 hadr_local_svc 545
 hadr_peer_window 545
 hadr_remote_host 546
 hadr_remote_inst 546
 hadr_remote_svc 547
 hadr_syncmode 547
 hadr_timeout 548
 health_mon 482
 indexrec 483, 548
 instance_memory 485
 interaction between memory parameters 25
 intra_parallel 487
 java_heap_sz 487
 jdk_64_path 550
 jdk_path 488
 jdk_path (DAS) 600
 keepfenced 489
 local_gssplugin 490
 locklist 551
 locktimeout 554
 log_retain_status 554
 logarchmeth1 555
 logarchmeth2 556
 logarchopt1 557
 logarchopt2 557
 logbufsz 558
 logfilsiz 558
 loghead 560
 logindexbuild 560
 logpath 560
 logprimary 560
 logretain 562
 logsecond 563
 max_connections 490
 restrictions 458
 max_connretries 491
 max_coordagents 491
 restrictions 458
 max_querydegree 492
 max_time_diff 493
 maxagents 493
 maxappls 564
 maxcagents 494
 maxfilop 565
 maxlocks 566
 maxlog 564
 min_dec_div_3 568
 mincommit 569
 mirrorlogpath 570
 mon_heap_sz 495
 multipage_alloc 571
 newlogpath 571
 nodetype 496
 notifylevel 496
 num_db_backups 573
 num_freqvalues 573
 num_initagents 497
 num_initfenced 498
 num_iocleaners 574

configuration parameters (*continued*)

- num_ioservers 576
- num_poolagents 498
- num_quantiles 577
- numarchretry 578
- numdb 499
- numlogspan 576
- numsegs 579
- overflowlogpath 579
- pagesize 580
- pckcachesz 580
- priv_mem_thresh 582
- query_heap_sz 500
- rec_his_retentn 583
- release 501
- restore_pending 583
- restrict_access 583
- resync_interval 501
- rollfwd_pending 584
- rqioblk 502
- sched_enable 601
- sched_userid 601
- self_tuning_mem 584
- seqdetect 586
- sheaphres 503
- sheaphres_shr 586
- smtp_server 601
- softmax 587
- sorheap 589
- spm_log_file_sz 505
- spm_log_path 505
- spm_max_resync 506
- spm_name 506
- srv_plugin_mode 508
- srvcon_auth 506
- srvcon_gssplugin_list 507
- srvcon_pw_plugin 508
- start_stop_time 508
- stat_heap_sz 590
- stmtheap 591
- summary
 - database 443
 - database manager 443
 - section heading descriptions 443
- svcname 509
- sysadm_group 510
- sysctrl_group 510
- sysmaint_group 511
- sysmon_group 512
- territory 591
- tm_database 512
- toolscat_db 602
- toolscat_inst 602
- toolscat_schema 603
- tp_mon_name 513
- trackmod 592
- trust_allclnts 515
- trust_clntauth 515
- tsm_mgmtclass 592
- tsm_nodename 592
- tsm_owner 593
- tsm_password 593
- user_exit_status 594
- userexit 594
- using the Configuration Advisor to define the scope 55
- util_heap_sz 595
- util_impact_lim 516

configuration parameters (*continued*)

- vendoropt 595
- wlm_collect_int configuration parameter 596
- configuring
 - LDAP user for applications 102
- conn_elapse configuration parameter 469
- connection elapse time configuration parameter 469
- connection states
 - application processes 135
 - description 136
- connections
 - elapse time 469
- constraints
 - (table) check 269
 - comparison with BEFORE triggers 276
 - creating 284
 - defining
 - foreign keys 277
 - referential constraints 277
 - description 267
 - designing 274
 - check constraints 275
 - dropping 286
 - informational 269, 274, 282
 - interaction with foreign keys 280
 - modifying 284
 - NOT NULL 268
 - primary key 269
 - referential 269
 - table check 269
 - unique 269
 - unique (key) 268
 - viewing definitions for a table 285
- contact_host configuration parameter 597
- containers
 - DMS table spaces 193
 - dropping containers from 203
 - modifying containers in 202
 - rebalancing and dropping 194
 - reduction of containers in 203
- constraints
 - types 267
- Coordinated Universal Time 493
- copy schema
 - operation, restarting 219
- cpuspeed configuration parameter
 - described 470
 - effect on query optimization 456
- CREATE DATABASE command
 - example of 120
- CREATE TABLE statement
 - defining referential constraints 277
- CREATE TABLESPACE statement
 - adjusting system temporary table spaces page sizes 166
- creating
 - LDAP users 101
- CURRENT SCHEMA special register 215

D

- DAS discovery mode configuration parameter 599
- das_codepage configuration parameter 597
- das_territory configuration parameter 598
- dasadm_group configuration parameter 598
- data
 - representation 138

- data access optimization
 - overview 19
- Data Definition Language (DDL)
 - description 109
 - statements
 - description 109
- data defragmentation
 - overview 19
- data organization schemes
 - table partitioning 253
- data partitions
 - creating 255
- data row compression
 - enabling 49, 243
- data servers
 - capacity management 3
 - overview 3
- data types
 - default values 235
- database configuration file
 - changing 116
 - creating 113
- database configuration parameters
 - recommended values 56
- database directories
 - structure 110
- database heap configuration parameter 533
- database manager
 - binding utilities 129
 - limits 353
 - machine node type configuration parameter 496
 - multiple instances 12
 - start timeout 508
 - stop timeout 508
- database manager configuration parameters
 - recommended values 56
 - summary 443
- database objects
 - naming rules
 - NLS 350
 - overview 348
 - Unicode 351
 - overview 223
 - statement dependencies when modifying objects 281
- database partitions
 - cataloging
 - node directory 113
 - node directory 113
 - overview 141
- database recovery log
 - allocating during database creation 117
- database system monitor
 - default database system monitor switches configuration parameter 471
- database territory code configuration parameter 529
- database_consistent configuration parameter 529
- database_level configuration parameter 530
- database_memory configuration parameter
 - interaction between memory parameters 25
- database_memory database configuration parameter
 - description 530
 - self-tuning 20
- database-managed space (DMS)
 - containers
 - dropping 194
 - rebalancing 194
 - reducing size 203
- database-managed space (DMS) (*continued*)
 - description 156
 - devices 169
 - table space containers 194
 - table space maps 159
 - table spaces
 - altering 193
 - compared to SMS table spaces 167
 - containers (dropping) 203
 - containers (reducing) 203
 - creating 189
 - workloads 168
- databases
 - aliases
 - creating 130
 - appl_memory configuration parameter 519
 - automatic storage 44, 121
 - autorestart configuration parameter 524
 - backup_pending configuration parameter 525
 - backups
 - automated 17, 19
 - cataloging
 - overview 128
 - codepage configuration parameter 528
 - codeset configuration parameter 528
 - collating information 528
 - configuration parameter summary 443
 - configuring across multiple partitions 36
 - designing
 - overview 109
 - distributed 109
 - dropping
 - DROP DATABASE command 138
 - estimating size requirements 118
 - maximum number of concurrently active databases 499
 - package dependencies 281
 - partitioned 109
 - relational 109
 - release level configuration parameter 501
 - restoring 125
 - territory code configuration parameter 529
 - territory configuration parameter 591
- DATE data type
 - default value 235
- db_mem_thresh configuration parameter 532
- DB2 administration server (DAS)
 - configuration parameters
 - authentication 596
 - contact_host 597
 - das_codepage 597
 - das_territory 598
 - dasadm_group 598
 - db2system 599
 - exec_exp_task 600
 - jdk_64_path 550
 - jdk_path 600
 - sched_enable 601
 - sched_userid 601
 - smtp_server 601
 - toolscat_db 602
 - toolscat_inst 602
 - toolscat_schema 603
 - multiple DB2 copies setup 10
 - ownership rules 369
- DB2 copies
 - default IBM database client interface copy 7

DB2 copies (*continued*)

- multiple on same computer
 - DB2 administration server (DAS) setting 10
 - default instance setting 11
- updating 13

DB2 Information Center

- languages 611
- updating 612
- versions 611
- viewing in different languages 611

DB2 servers

- post-migration tasks
 - adjusting system temporary table space page sizes 166

DB2_ALLOCATION_SIZE registry variable

- description 404

DB2_ALTERNATE_GROUP_LOOKUP environment variable 385

DB2_ANTIJOIN variable 399

DB2_APM_PERFORMANCE variable 404

DB2_ASYNC_IO_MAXFILOP registry variable

- description 404

DB2_ATS_ENABLE registry variable 421

DB2_AVOID_PREFETCH variable 404

DB2_CAPTURE_LOCKTIMEOUT registry variable

- description 377

DB2_CLP_EDITOR variable 385

DB2_CLPHISTSIZE variable 385

DB2_CLPPROMPT registry variable 397

DB2_COLLECT_TS_REC_INFO registry variable 377

DB2_COMMIT_ON_EXIT registry variable 421

DB2_CONNRETRIES_INTERVAL registry variable

- description 377

DB2_COPY_NAME environment variable 385

DB2_CREATE_DB_ON_PATH registry variable 421

DB2_DIAGPATH variable

- description 385

DB2_DISABLE_FLUSH_LOG registry variable 421

DB2_DISPATCHER_PEEKTIMEOUT registry variable 421

DB2_DJ_INI variable 421

DB2_DOCHOST variable 421

DB2_DOCPORT variable 421

DB2_ENABLE_AUTOCONFIG_DEFAULT variable 421

DB2_ENABLE_LDAP variable

- description 421

DB2_EVALUNCOMMITTED registry variable

- description 404

DB2_EVMON_EVENT_LIST_SIZE registry variable

- description 421

DB2_EVMON_STMT_FILTER registry variable 421

DB2_EXTENDED_IN2JOIN variable 404

DB2_EXTENDED_IO_FEATURES variable

- description 404

DB2_EXTENDED_OPTIMIZATION variable 404

DB2_EXTSECURITY registry variable 421

DB2_FALLBACK variable 421

DB2_FMP_COMM_HEAPSZ variable 421

DB2_FORCE_APP_ON_MAX_LOG registry variable 377

DB2_FORCE-NLS_CACHE registry variable

- description 394

DB2_GRP_LOOKUP variable 421

DB2_HADR_BUF_SIZE variable 421

DB2_HADR_NO_IP_CHECK variable 421

DB2_HADR_SORCVBUF registry variable 421

DB2_HADR_SOSNDBUF registry variable 421

DB2_HASH_JOIN registry variable

- description 404

DB2_INLIST_TO_NLJN registry variable 399

DB2_IO_PRIORITY_SETTING registry variable 404

DB2_KEEPTABLELOCK registry variable 404

DB2_LARGE_PAGE_MEM registry variable

- description 404

DB2_LIC_STAT_SIZE registry variable 377

DB2_LIKE_VARCHAR registry variable 399

DB2_LOAD_COPY_NO_OVERRIDE variable 421

DB2_MAP_XML_AS_CLOB_FOR_DLC registry variable

- description 421

DB2_MAX_CLIENT_CONNRETRIES registry variable

- description 377

DB2_MAX_INACT_STMTS variable 404

DB2_MAX_LOB_BLOCK_SIZE variable 421

DB2_MAX_NON_TABLE_LOCKS variable 404

DB2_MDC_ROLLOUT registry variable

- description 404

DB2_MEM_TUNING_RANGE variable 404

DB2_MEMORY_PROTECT registry variable

- description 421

DB2_MINIMIZE_LISTPREFETCH registry variable 399

DB2_MMAP_READ variable 404

DB2_MMAP_WRITE variable 404

DB2_NEW_CORR_SQ_FF variable 399

DB2_NO_FORK_CHECK registry variable

- description 404

DB2_NUM_CKPW_DAEMONS registry variable 421

DB2_NUM_FAILOVER_NODES registry variable 398

DB2_OPT_MAX_TEMP_SIZE variable 399

DB2_OPTSTATS_LOG registry variable

- description 421

DB2_OVERRIDE_BPF variable 404

DB2_PARALLEL_IO registry variable

- description 207, 385

DB2_PARTITIONEDLOAD_DEFAULT registry variable

- description 398

DB2_PINNED_BP registry variable

- description 404

DB2_REDUCED_OPTIMIZATION registry variable 399

DB2_RESOLVE_CALL_CONFLICT registry variable 421

DB2_RESOURCE_POLICY registry variable

- description 404

DB2_SELECTIVITY registry variable 399

DB2_SELUDI_COMM_BUFFER registry variable 404

DB2_SERVER_CONTIMEOUT registry variable 421

DB2_SERVER_ENCALG registry variable 421

DB2_SET_MAX_CONTAINER_SIZE registry variable

- description 404

DB2_SKIPDELETED registry variable 404

DB2_SKIPINSERTED registry variable 404

DB2_SMS_TRUNC_TMPTABLE_THRESH variable 404

DB2_SORT_AFTER_TQ variable 404

DB2_SQLROUTINE_PREPOPTS registry variable 399

DB2_SYSTEM_MONITOR_SETTINGS registry variable

- description 377

DB2_THREAD_SUSPENSION registry variable

- description 421

DB2_TRUNCATE_REUSESTORAGE registry variable 421

DB2_TRUSTED_BINDIN registry variable

- description 404

DB2_UPDDBCFG_SINGLE_DBPARTITION variable

- description 385

DB2_USE_ALTERNATE_PAGE_CLEANING registry variable

- description 404

DB2_USE_DB2JCT2_JROUTINE variable

- description 421

DB2_USE_PAGE_CONTAINER_TAG variable

- description 385

DB2_USE_PAGE_CONTAINER_TAG variable (continued)
 performance impact 207

DB2_UTIL_MSGPATH registry variable 421

DB2_VENDOR_INI registry variable
 description 421

DB2_VIEW_REOPT_VALUES registry variable 377

DB2_WORKLOAD aggregate registry variable
 description 371, 385

DB2_XBSA_LIBRARY registry variable 421

DB2ACCOUNT registry variable
 description 377

DB2ADMINSERVER variable 421

DB2ASSUMEUPDATE registry variable 404

DB2AUTH
 registry variables 421

DB2BIDI registry variable
 description 377

DB2BPVARS registry variable
 description 404

DB2BQTIME variable 397

DB2BQTRY variable 397

DB2CHECKCLIENTINTERVAL variable 394

DB2CHGPWD_EEE registry variable 398

DB2CHKPTR variable 404

DB2CHKSQlda variable 404

DB2CLIINIPATH variable
 description 421

DB2CODEPAGE registry variable
 description 377

DB2COMM variable 394

DB2CONNECT_DISCONNECT_ON_INTERRUPT
 variable 421

DB2CONNECT_IN_APP_PROCESS environment variable 385

DB2CONSOLECP registry variable 377

DB2COUNTRY registry variable
 description 377

DB2DBDFT registry variable 377

DB2DBMSADDR registry variable 377

DB2DEFPREP registry variable
 description 421

DB2DISCOVERYTIME registry variable 377

DB2DMNBCKCTLr registry variable
 description 421

DB2DOMAINLIST variable
 description 385

db2envar.bat command
 using to switch DB2 copies 11

DB2ENVLIST environment variable 385

DB2FCMCOMM variable 394

DB2FODC registry variable
 description 377

DB2GRAPHICUNICODESERVER registry variable
 description 377

db2icrt command
 creating instances 68

db2idrop command
 dropping instances 75

DB2INCLUDE registry variable 377

DB2INSTANCE
 setting 11

DB2INSTANCE environment variable
 defining default instance 12
 description 385

DB2INSTDEF
 setting 11

DB2INSTDEF registry variable 377

DB2INSTOWNER registry variable 377

DB2INSTPROF registry variable
 description 385
 location 439

DB2IQTIME variable 397

db2iupdt command
 updating instance configuration
 Linux 69
 UNIX 69
 Windows 70

DB2LDAP_BASEDN variable
 description 421

DB2LDAP_CLIENT_PROVIDER registry variable
 description 421
 IBM LDAP client 88

DB2LDAP_KEEP_CONNECTION registry variable
 description 421

DB2LDAP_SEARCH_SCOPE variable
 description 421

DB2LDAPCACHE variable 421

DB2LDAPHOST variable
 description 421

DB2LDAPSecurityConfig environment variable 385

db2ldcfg command
 configuring LDAP user 102

DB2LIBPATH environment variable 385

DB2LOADREC registry variable
 description 421

DB2LOCALE registry variable
 description 377

DB2LOCK_TO_RB variable 421

DB2LOGINRESTRICTIONS variable 385

DB2MAXFSCRSEARCH variable 404

DB2MEMDISCLAIM registry variable 404

DB2MEMMAXFREE registry variable
 description 404

db2move command
 COPY schema errors 219
 schema copying examples 218

DB2NODE environment variable
 description 385

db2nodes.cfg file
 creating 114
 overview 66

DB2NOEXITLIST registry variable
 description 421

DB2NTMEMSIZE variable 404

DB2NTNOCACHE registry variable
 description 404
 NO FILE SYSTEM CACHING clause comparison 176

DB2NTPRICLASS registry variable
 description 404

DB2NTWORKSET variable 404

DB2OPTIONS environment variable
 description 385

DB2PATH environment variable 385

DB2PORTRANGE registry variable 398

DB2PRIORITIES registry variable
 description 404

DB2PROCESSORS environment variable 385

DB2RCMD_LEGACY_MODE environment variable 385

DB2REMOTEPREG variable 421

DB2ROUTINE_DEBUG registry variable 421

DB2RQTIME variable 397

DB2RSHCMD registry variable 394

DB2RSHTIMEOUT registry variable 394

DB2SATELLITEID variable 421

- db2SelectDB2Copy API
 - using to switch DB2 copies 11
- db2set command
 - managing registry and environment variables 363, 365
- DB2SORCVBUF variable
 - description 394
- DB2SORT variable 421
- DB2SOSNDBUF variable
 - description 394
- db2system configuration parameter 599
- DB2SYSTEM environment variable 385
- DB2TCP_CLIENT_CONTIMEOUT registry variable 394
- DB2TCP_CLIENT_RCVTIMEOUT registry variable
 - description 394
- DB2TCPCONNMGERS registry variable 394
- DB2TERRITORY registry variable
 - description 377
- DBCS (double-byte character set)
 - See double-byte character set (DBCS) 350
- dbheap database configuration parameter
 - overview 533
- DDL (Data Definition Language)
 - See Data Definition Language (DDL) 109
- deadlocks
 - checking for 541
 - dlchktime configuration parameter 541
- decflt_rounding database configuration parameter 534
- decimal division scale to 3 configuration parameter 568
- DECLARE GLOBAL TEMPORARY TABLE statement
 - overview 254
- default database path configuration parameter 472
- default number of SMS containers configuration
 - parameter 579
- deferred index cleanup
 - monitoring 60
- deletable views
 - description 339
- delete rule
 - description 269
- delimited identifiers
 - naming rules 349
- dependent rows
 - overview 269
- dependent tables
 - overview 269
- descendent row
 - overview 269
- descendent table
 - overview 269
- DETACH command
 - detaching from instances 72
- dft_account_str configuration parameter 470
- dft_degree configuration parameter
 - description 536
 - effect on query optimization 456
- dft_extent_sz configuration parameter 536
- dft_loadrec_ses configuration parameter 537
- dft_mon_bufpool configuration parameter 471
- dft_mon_lock configuration parameter 471
- dft_mon_sort configuration parameter 471
- dft_mon_stmt configuration parameter 471
- dft_mon_table configuration parameter 471
- dft_mon_timestamp configuration parameter 471
- dft_mon_uow configuration parameter 471
- dft_monswitches configuration parameter 471
- dft_mttb_types configuration parameter 537
- dft_prefetch_sz configuration parameter 538
- dft_queryopt configuration parameter 539
- dft_refresh_age configuration parameter 539
- dft_sqlmathwarn configuration parameter 540
- dftdbpath configuration parameter 472
- diaglevel configuration parameter
 - description 473
- diagpath configuration parameter 474
- dictionaries
 - automated creation 17, 47
- dir_cache configuration parameter 475
- directories
 - instance 66
 - local database
 - description 113
 - viewing 138
 - node
 - cataloguing database partition 113
 - viewing 113
 - system database
 - description 114
 - viewing 138
- directory cache support configuration parameter
 - description 475
- directory schema
 - extending
 - IBM Tivoli Directory Server 90
 - Sun One Directory Server 93
- discover server instance configuration parameter 477
- discover_db configuration parameter 541
- discover_inst configuration parameter 477
- discovery feature
 - discovery mode configuration parameter 476
- discovery mode configuration parameter 476
- distinct user-defined data type 235
- distributed relational databases
 - connecting to 131
 - remote units of work 131
- dlchktime configuration parameter 541
- DMS (database-managed space)
 - See database-managed space (DMS) 156
- documentation
 - overview 607
 - PDF 607
 - printed 607
 - terms and conditions of use 614
- double-byte character set (DBCS)
 - naming rules 350
- dyn_query_mgmt configuration parameter
 - description 542

E

- enable_xmlchar database configuration parameter
 - description 543
- environment variables
 - declaring 365
 - Linux 369
 - overview 372
 - profile registry 363
 - setting
 - Linux 369
 - UNIX 369
 - Windows 367
 - UNIX
 - setting 369
 - Windows 367
- exec_exp_task configuration parameter 600

- expressions
 - NEXT VALUE 321
 - PREVIOUS VALUE 321
- extent sizes
 - table spaces 182

F

- failarchpath configuration parameter 543
- FCM (Fast Communications Manager)
 - channels 478
 - monitor elements
 - fcm_num_channels 478
- fcm_num_buffers configuration parameter 477
- fcm_num_channel configuration parameter 478
- fed_noauth configuration parameter 479
- federated configuration parameter 479
- federated databases
 - system support configuration parameter 479
- federated_async database manager configuration parameter 480
- fenced_pool database manager configuration parameter 480
- fenced-mode processes
 - running vendor library functions 37
- FILE SYSTEM CACHING clause 176
- file systems
 - caching for table spaces 176, 179
- first active log file configuration parameter 560
- first-fit order 238
- foreign key constraints
 - referential constraints 277
 - referential integrity rules 269
 - rules for defining 277
- foreign key constraint 267
- foreign keys
 - composite 277
 - constraint name 277
 - constraints 269
 - interaction with referential constraints 280
 - referential integrity implications 281
 - rules for defining 277

G

- generated columns
 - defining 232
 - examples 232
 - modifying 258
- global (user-defined) temporary tables
 - creating 254
- global level profile registry 363
- global temporary tables
 - comparison with other table types 225
- group_plugin configuration parameter 482
- groupheap_ratio database manager configuration parameter 544
- groups
 - naming rules 350

H

- hadr_db_role configuration parameter 544
- hadr_local_host configuration parameter 544
- hadr_local_svc configuration parameter 545
- hadr_peer_window database configuration parameter
 - description 545

- hadr_remote_host configuration parameter 546
- hadr_remote_inst configuration parameter 546
- hadr_remote_svc configuration parameter 547
- hadr_syncmode configuration parameter 547
- hadr_timeout configuration parameter 548
- health monitor
 - description 17
 - health_mon configuration parameter 482
- health_mon configuration parameter 482
- heaps
 - configuring 32
- help
 - configuring language 611
 - SQL statements 611

I

- I/O
 - parallelism
 - using RAID devices 207
 - table space considerations 183
- IBM database client interface copies
 - default 7
- IBM eNetwork Directory
 - object classes and attributes 78
- IBM SecureWay Directory Server
 - extending the directory schema for 90
- identifiers
 - length limits 353
- identity columns
 - comparing with sequences 326
 - comparison with sequences 324
 - defining on a new table 233
 - example 233
 - modifying 258
- IMPLICIT_SCHEMA authority 211
- indexes
 - asynchronous cleanup 58, 60
 - bi-directional 290
 - cleaning up 58, 60
 - clustered 290
 - creating 298
 - deferred cleanup 60
 - description 289
 - Design Advisor 294
 - designing 292
 - dropping 300
 - guidelines and considerations 292
 - improving performance 290
 - modifying 298
 - non-clustered 290
 - non-unique 290
 - overview 289
 - rebuilding 299
 - renaming 298
 - space requirements 294
 - tools for designing 294
 - unique 290
- indexrec configuration parameter 483, 548
- informational constraints 267
 - description 269, 274
 - designing 282
- initial number of agents in pool configuration parameter 497
- initial number of fenced processes configuration parameter 498
- insert rule
 - for referential integrity 269

- insert rule (*continued*)
 - referential constraints 269
- insertable views
 - using 340
- instance level profile registry 363
- instance profile registry 363
- instance_memory configuration parameter
 - interaction between memory parameters 25
- instances
 - auto-starting 71
 - creating 63
 - creating additional 68
 - default 63, 65
 - default, setting 11
 - designing 64
 - directory 66
 - instance_memory configuration parameter 485
 - modifying 69
 - multiple 12
 - multiple (Linux, UNIX) 66
 - multiple (Windows) 13, 67
 - overview 12
 - removing 75
 - running concurrently 15, 73
 - running multiple (Windows) 14
 - setting the current 370
 - starting (Linux, UNIX) 71
 - starting (Windows) 72
 - stopping (Linux, UNIX) 73
 - stopping (Windows) 74
 - updating the configuration
 - UNIX 69
 - Windows 70
 - working with 71
- INSTEAD OF triggers 304
 - description 302
- intra_parallel database manager configuration parameter 487

J

- java_heap_sz database manager configuration parameter 487
- jdk_64_path configuration parameter 550
- jdk_path configuration parameter
 - description 488
- jdk_path DAS configuration parameter 600

K

- keepfenced configuration parameter
 - description 489
- keys
 - foreign constraints 269
 - parent constraints 269

L

- label-based access control (LBAC)
 - limits 353
- large objects (LOBs)
 - caching behavior 169
- large page support
 - AIX 64-bit environment 4
- larger RID
 - adjusting system temporary table space page sizes 166
- LBAC (label-based access control)
 - limits 353

- LBAC (label-based access control) (*continued*)
 - optimistic locking 246
 - security labels
 - component name length 353
 - name length 353
 - security policies
 - name length 353
- LDAP (Lightweight Directory Access Protocol)
 - attaching remotely 104
 - cataloging a node entry 100
 - DB2 Connect 89
 - deregistering
 - databases 101
 - servers 101
 - description 77
 - directory service 119
 - disabling 102
 - enabling 98
 - extending directory schema 88
 - object classes and attributes 78
 - refreshing entries 104
 - registering
 - databases 100
 - DB2 servers 99
 - host databases 89
 - rerouting clients 103
 - searching
 - directory domains 105
 - directory partitions 105
 - security 77
 - setting registry variables 102
 - supporting 88
 - updating protocol information 103
 - user creation 101
 - Windows 2003 active directory 97
- library functions
 - running in fenced-mode processes 37
- Lightweight Directory Access Protocol (LDAP)
 - attaching remotely 104
 - cataloging node entries 100
 - creating a user 101
 - DB2 Connect 89
 - deregistering
 - databases 101
 - servers 101
 - description 77
 - directory service 119
 - disabling 102
 - enabling 98
 - extending directory schema 88
 - object classes and attributes 78
 - refreshing entries 104
 - registering
 - databases 100
 - DB2 servers 99
 - host databases 89
 - searching
 - directory domains 105
 - directory partitions 105
 - security 77
 - setting registry variables 102
 - supporting 88
 - updating protocol information 103
 - Windows 2003 active directory 97
- limits
 - identifier length 353
 - SQL 353

- LOBs (large objects)
 - caching behavior 169
- local database directory
 - description 113
 - viewing 138
- local_gssplugin configuration parameter 490
- locklist configuration parameter
 - description 551
 - query optimization 456
- locks
 - maximum percent of lock list before escalation 566
 - maximum storage for lock list 551
 - optimistic locking 245
 - time interval for checking deadlock configuration parameter 541
- locktimeout configuration parameter 554
- log files
 - space requirements 118
- log_retain_status configuration parameter 554
- logarchmeth1 configuration parameter 555
- logarchmeth2 configuration parameter 556
- logarchopt1 configuration parameter 557
- logarchopt2 configuration parameter 557
- LOGBUFFSZ configuration parameter 558
- logfilesiz configuration parameter 558
- loghead configuration parameter 560
- logindexbuild configuration parameter 560
- logpath configuration parameter 560
- logprimary configuration parameter 560
- logretain database configuration parameter 562
- logs
 - block on log disk full configuration parameter 525
 - first active log file configuration parameter 560
 - location of log files configuration parameter 560
 - log buffer size configuration parameter 558
 - log retain enable configuration parameter 562
 - log retain status indicator configuration parameter 554
 - mirror log path configuration parameter 570
 - newlogpath configuration parameter 571
 - number of primary log files configuration parameter 560
 - number of secondary log files configuration parameter 563
 - overflow log path configuration parameter 579
 - raw devices 186
 - recovery range and soft checkpoint interval configuration parameter 587
 - size of log files configuration parameter 558
 - user exit enable configuration parameter 594
- logsecond configuration parameter 563
- LONG data
 - caching behavior 169
- long fields
 - caching behavior 169

M

- maintenance
 - automatic 19
 - windows 19
- maintenance windows
 - automatic 19
- materialized query tables (MQTs)
 - altering properties 257
 - dropping 262
 - refreshing data 257
- max_connections database manager configuration parameter
 - restrictions 458
- max_connretries configuration parameter 491
- max_coordagents database manager configuration parameter 491
 - restrictions 458
- max_logicagents configuration parameter 490
- max_querydegree configuration parameter 492
- max_time_diff configuration parameter 493
- maxagents database manager configuration parameter 493
- maxappls configuration parameter 564
 - effect on memory use 21
- maxcagents database manager configuration parameter 494
- maxcoordagents configuration parameter 21
- MAXDARI configuration parameter
 - renamed to fenced_pool configuration parameter 480
- maxfilop database configuration parameter 565
- maximum database files open per application configuration parameter 565
- maximum Java interpreter heap size configuration parameter 487
- maximum log per transaction configuration parameter 564
- maximum number of active applications configuration parameter 564
- maximum number of agents configuration parameter 493
- maximum number of concurrent agents configuration parameter 494
- maximum number of concurrently active databases configuration parameter 499
- maximum number of coordinating agents configuration parameter 491
- maximum number of fenced processes configuration parameter 480
- maximum percent of lock list before escalation configuration parameter 566
- maximum query degree of parallelism configuration parameter 492
 - effect on query optimization 456
- maximum size of application group memory set configuration parameter 518
- maximum storage for lock list configuration parameter 551
- maximum time difference among nodes configuration parameter 493
- maxlocks configuration parameter 566
- maxlog configuration parameter 564
- memory
 - applheapsz configuration parameter 520
 - application memory configuration parameter 519
 - aslheapsz configuration parameter 463
 - configuration 32
 - self-tuning memory 20
 - dbheap configuration parameter 533
 - instance memory configuration parameter 485
 - interaction between memory parameters 25
 - organization of use 21
 - package cache size configuration parameter 580
 - sort heap size configuration parameter 589
 - sort heap threshold configuration parameter 503
 - statement heap size configuration parameter 591
 - when allocated 21
- memory configuration
 - overview 36
- memory tuner
 - partitioned database environments 31
- min_dec_div_3 configuration parameter 568
- mincommit configuration parameter 569
- mirror log path configuration parameter 570
- mirrorlogpath configuration parameter 570
- mon_heap_sz database manager configuration parameter 495

- MQs (materialized query tables)
 - altering properties 257
 - refreshing data 257
- multidimensional clustering (MDC) tables
 - comparison with other table types 225
 - deferred index cleanup 60
- multipage_alloc configuration parameter 571
- multiple DB2 copies
 - default IBM database client interface copy 7
 - overview 7
 - running instances concurrently 15, 73
 - setting the default instance 11
- multiple instances 66
 - overview 12
 - Windows 13, 67

N

- naming rules
 - DB2 objects 348
 - delimited identifiers and object names 349
 - national languages 350
 - restrictions 347
 - schema name restrictions 215
 - Unicode 351
 - users, user IDs and groups 350
- nested views
 - definition 339
- Netscape
 - LDAP directory support 91
- newlogpath configuration parameter 571
- NEXT VALUE expression
 - sequences 321
 - using identity columns 326
- NO FILE SYSTEM CACHING clause 176
- NO_SORT_MGJOIN 399
- NO_SORT_NLJOIN 399
- node configuration files
 - creating 114
- node connection retries configuration parameter 491
- node directories
 - cataloguing database partition 113
 - description 113
 - viewing 113
- node level profile registry 363
- nodes
 - connection elapse time 469
 - coordinating agents 491
 - maximum time difference among 493
- nodetype configuration parameter 496
- non-buffered I/O
 - enabling/disabling 176
- non-clustered indexes 290
- non-unique indexes 290
- NOT NULL constraints
 - overview 268
- NOT NULL constraints
 - types 267
- notices 617
- notify level configuration parameter
 - overview 496
- NULL data type 235
- num_db_backups configuration parameter
 - overview 573
- num_freqvalues configuration parameter 573
- num_initfenced database manager configuration parameter 498

- num_iocleaners configuration parameter 574
- num_ioservers configuration parameter 576
- num_poolagents database manager configuration parameter 498
- num_quantiles configuration parameter 577
- numarchretry configuration parameter 578
- number log span configuration parameter 576
- number of commits to group configuration parameter 569
- number of database backups configuration parameter 573
- NUMDB
 - configuration parameter 499
 - effect on memory use 21
- numinitagents configuration parameter 497
- numlogspan configuration parameter 576
- numsegs database configuration parameter
 - overview 579

O

- object names
 - rules 349
- offline maintenance 19
- OLTP (online transaction processing)
 - table space design 168
- online maintenance 19
- online transaction processing (OLTP)
 - table space design 168
- optimistic locking
 - about 245
 - conditions 252
 - enabling 253
 - implicitly hidden columns 253
 - implicitly hidden columns 246
 - LBAC considerations 246
 - overview 244
 - planning enablement 252
 - restrictions 246
 - row change token 253
 - scenario A
 - optimistic locking enabled 263
 - scenario C
 - using implicitly hidden columns 265
 - time-based update detection 248, 253
 - usage scenarios 263
 - using RID() functions 253
- ordering DB2 books 610
- overflowlogpath configuration parameter 579

P

- packages
 - inoperative 281
- page sizes
 - tables 238
- pages
 - sizes
 - database default 580
 - table spaces 183
- pagesize configuration parameter 580
- parallelism
 - configuration parameters
 - dft_degree 536
 - intra_parallel 487
 - max_querydegree 492

- parallelism (*continued*)
 - I/O
 - Redundant Array of Independent Disks (RAID) devices 207
- parent key
 - overview 269
- parent row
 - overview 269
- parent tables
 - overview 269
- partitioned database environments
 - self-tuning memory 29, 31
- partitioned tables
 - comparison with other table types 225
 - creating 255
- pckcachesz configuration parameter 580
- performance
 - improving with indexes 290
 - managing sequences 322
 - table spaces 207
- Performance Configuration wizard
 - renamed to Configuration Advisor 116
- pool size for agents
 - controlling 498
- post-migration tasks
 - DB2 servers
 - adjusting system temporary table space page sizes 166
- prefetch size
 - automatic adjustment 175
- PREVIOUS VALUE expression
 - identity columns 326
 - overview 321
- primary key constraint
 - overview 267
- primary keys
 - description 235, 269
 - designing 275
- priv_mem_thresh database manager configuration parameter
 - description 582
- problem determination
 - information available 614
 - tutorials 614
- process model
 - configuration simplification 35
 - overview 36
- processors
 - adding 3
- profiles
 - registry 363
- properties
 - columns
 - changing 258
- protocols
 - TCP/IP service name configuration parameter 509

Q

- queries
 - statement heap size configuration parameter 591
- query optimization
 - configuration parameters 456
- query workload
 - table space design 168
- query_heap_sz database manager configuration parameter 500

R

- RAID (Redundant Array of Independent Disks) devices
 - optimizing table space performance 207
- range-clustered tables
 - comparison with other table types 225
- raw devices 189
- raw I/O
 - setting up on Linux 187
 - specifying 186
- raw logs 186
- read-only views
 - using 341
- rebalancing
 - across containers 193
- rec_his_retentn configuration parameter 583
- recovery
 - auto restart enable configuration parameter 524
 - backup pending indicator configuration parameter 525
 - default number of load recovery sessions configuration parameter 537
 - index re-creation time configuration parameter 483, 548
 - log retain status indicator configuration parameter 554
 - number of database backups configuration parameter 573
 - restore pending configuration parameter 583
 - roll forward pending indicator configuration parameter 584
 - summary tables
 - inoperative 260
 - user exit status indicator configuration parameter 594
 - views
 - inoperative 343
- recovery history file
 - retention period configuration parameter 583
- recovery log
 - allocating during database creation 117
- recovery range and soft checkpoint interval configuration parameter 587
- Redundant Array of Independent Disks (RAID)
 - optimizing performance 207
- referential constraints
 - defining 277
 - description 269
 - interaction with foreign keys 280
 - PRIMARY KEY clause, CREATE/ALTER TABLE statements 277
 - REFERENCES clause, CREATE/ALTER TABLE statements 277
- referential integrity
 - constraints 269
 - delete rule 269
 - description 235
 - insert rule 269
 - update rule 269
- registry variables
 - aggregate 371
 - DB2_ALLOCATION_SIZE 404
 - DB2_ALTERNATE_GROUP_LOOKUP 385
 - DB2_ANTIJOIN 399
 - DB2_APM_PERFORMANCE 404
 - DB2_ASYNC_IO_MAXFILOP 404
 - DB2_ATS_ENABLE 421
 - DB2_AVOID_PREFETCH 404
 - DB2_CAPTURE_LOCKTIMEOUT 377
 - DB2_CLP_EDITOR 385
 - DB2_CLPHISTSIZE 385
 - DB2_CLPPROMPT 397
 - DB2_COLLECT_TS_REC_INFO 377

registry variables (continued)

DB2_COMMIT_ON_EXIT 421
 DB2_CONNRETRIES_INTERVAL 377
 DB2_COPY_NAME 385
 DB2_CREATE_DB_ON_PATH 421
 DB2_DIAGPATH 385
 DB2_DISABLE_FLUSH_LOG 421
 DB2_DISPATCHER_PEEKTIMEOUT 421
 DB2_DJ_INI 421
 DB2_DOCHOST 421
 DB2_DOCPORT 421
 DB2_ENABLE_AUTOCONFIG_DEFAULT 421
 DB2_ENABLE_LDAP 421
 DB2_EVALUNCOMMITTED 404
 DB2_EVMON_EVENT_LIST_SIZE 421
 DB2_EVMON_STMT_FILTER 421
 DB2_EXTENDED_IO_FEATURES 404
 DB2_EXTENDED_OPTIMIZATION 404
 DB2_EXTSECURITY 421
 DB2_FALLBACK 421
 DB2_FCM_SETTINGS 398
 DB2_FMP_COMM_HEAPSZ 421
 DB2_FORCE_APP_ON_MAX_LOG 377
 DB2_FORCE-NLS_CACHE 394
 DB2_GRP_LOOKUP 421
 DB2_HADR_BUF_SIZE 421
 DB2_HADR_NO_IP_CHECK 421
 DB2_HADR_SORCVBUF 421
 DB2_HADR_SOSNDBUF 421
 DB2_HASH_JOIN 404
 DB2_INLIST_TO_NLJN 399
 DB2_IO_PRIORITY_SETTING 404
 DB2_KEEP_AS_AND_DMS_CONTAINERS_OPEN 404
 DB2_KEEPTABLELOCK 404
 DB2_LARGE_PAGE_MEM 404
 DB2_LIC_STAT_SIZE 377
 DB2_LIKE_VARCHAR 399
 DB2_LOAD_COPY_NO_OVERRIDE 421
 DB2_MAP_XML_AS_CLOB_FOR_DLC 421
 DB2_MAX_CLIENT_CONNRETRIES 377
 DB2_MAX_INACT_STMTS 404
 DB2_MAX_LOB_BLOCK_SIZE 421
 DB2_MAX_NON_TABLE_LOCKS 404
 DB2_MDC_ROLLOUT 404
 DB2_MEM_TUNING_RANGE 404
 DB2_MEMORY_PROTECT 421
 DB2_MINIMIZE_LISTPREFETCH 399
 DB2_MMAP_READ 404
 DB2_MMAP_WRITE 404
 DB2_NEW_CORR_SQ_FF 399
 DB2_NO_FORK_CHECK 404
 DB2_NUM_CKPW_DAEMONS 421
 DB2_NUM_FAILOVER_NODES 398
 DB2_OBJECT_TABLE_ENTRIES 404
 DB2_OPT_MAX_TEMP_SIZE 399
 DB2_OPTSTATS_LOG 421
 DB2_OVERRIDE_BPF 404
 DB2_PARALLEL_IO 385
 DB2_PARTITIONEDLOAD_DEFAULT 398
 DB2_PINNED_BP 404
 DB2_REDUCED_OPTIMIZATION 399
 DB2_RESOLVE_CALL_CONFLICT 421
 DB2_RESOURCE_POLICY 404
 DB2_SELECTIVITY 399
 DB2_SELUDI_COMM_BUFFER 404
 DB2_SERVER_CONTIMEOUT 421
 DB2_SERVER_ENCALG 421

registry variables (continued)

DB2_SET_MAX_CONTAINER_SIZE 404
 DB2_SKIPDELETED 404
 DB2_SKIPINSERTED 404
 DB2_SMS_TRUNC_TMPTABLE_THRESH 404
 DB2_SORT_AFTER_TQ 404
 DB2_SQLROUTINE_PREPOPTS 399
 DB2_SYSTEM_MONITOR_SETTINGS 377
 DB2_THREAD_SUSPENSION 421
 DB2_TRUNCATE_REUSESTORAGE 421
 DB2_TRUSTED_BINDIN 404
 DB2_UPDDBCFCG_SINGLE_DBPARTITION 385
 DB2_USE_ALTERNATE_PAGE_CLEANSING 404
 DB2_USE_DB2JCCT2_JROUTINE 421
 DB2_USE_PAGE_CONTAINER_TAG 385
 DB2_UTIL_MSGPATH 421
 DB2_VENDOR_INI 421
 DB2_VIEW_REOPT_VALUES 377
 DB2_WORKLOAD 385
 DB2_XBSA_LIBRARY 421
 DB2ACCOUNT 377
 DB2ADMINSERVER 421
 DB2ASSUMEUPDATE 404
 DB2BIDI 377
 DB2BPVARS 404
 DB2BQTIME 397
 DB2BQTRY 397
 DB2CHECKCLIENTINTERVAL 394
 DB2CHGPWD_ESE 398
 DB2CHKPTR 404
 DB2CHKSQLDA 404
 DB2CLIINIPATH 421
 DB2CODEPAGE 377
 DB2COMM 394
 DB2CONNECT_DISCONNECT_ON_INTERRUPT 421
 DB2CONNECT_IN_APP_PROCESS 385
 DB2CONSOLECP 377
 DB2COUNTRY 377
 DB2DBDFT 377
 DB2DBMSADDR 377
 DB2DEFPREP 421
 DB2DISCOVERYTIME 377
 DB2DMNBCKCTLR 421
 DB2DOMAINLIST 385
 DB2ENVLIST 385
 DB2FCMCOMM 394
 DB2FODC 377
 DB2GRAPHICUNICODESERVER 377
 DB2INCLUDE 377
 DB2INSTANCE 385
 DB2INSTDEF 377
 DB2INSTOWNER 377
 DB2INSTPROF 385
 DB2IQTIME 397
 DB2LDAP_BASEDN 421
 DB2LDAP_CLIENT_PROVIDER 421
 DB2LDAP_KEEP_CONNECTION 421
 DB2LDAP_SEARCH_SCOPE 421
 DB2LDAPCACHE 421
 DB2LDAPHOST 421
 DB2LDAPSecurityConfig 385
 DB2LIBPATH 385
 DB2LOADREC 421
 DB2LOCALE 377
 DB2LOCK_TO_RB 421
 DB2LOGINRESTRICTIONS 385
 DB2MAXFSCRSEARCH 404

- registry variables (*continued*)
 - DB2MEMDISCLAIM 404
 - DB2MEMMAXFREE 404
 - DB2NODE 385
 - DB2NOEXITLIST 421
 - DB2NTMEMSIZE 404
 - DB2NTNOCACHE 404
 - DB2NTPRICLASS 404
 - DB2NTWORKSET 404
 - DB2OPTIONS 385
 - DB2PATH 385
 - DB2PORTRANGE 398
 - DB2PRIORITIES 404
 - DB2PROCESSORS 385
 - DB2RCMD_LEGACY_MODE 385
 - DB2REMOTEPREG 421
 - DB2ROUTINE_DEBUG 421
 - DB2RQTIME 397
 - DB2RSHCMD 394
 - DB2RSHTIMEOUT 394
 - DB2SATELLITEID 421
 - DB2SLOGON 377
 - DB2SORCVBUF 394
 - DB2SORT 421
 - DB2SOSNDBUF 394
 - DB2SYSTEM 385
 - DB2TCP_CLIENT_CONTIMEOUT 394
 - DB2TCP_CLIENT_RCVTIMEOUT 394
 - DB2TCPCONNMGERS 394
 - DB2TERRITORY 377
 - declaring 365
 - environment variables 363
 - NO_SORT_MGJOIN keyword 399
 - NO_SORT_NLJOIN keyword 399
 - overview 372
 - registry variable
 - DB2_LOGGER_NON_BUFFERED_IO 404
 - registry variables
 - DB2_HADR_PEER_WAIT_LIMIT 421
- regular tables
 - comparison with other table types 225
- release configuration parameter 501
- remote units of work
 - distributed relational databases 131
- reorganization utility
 - binding to a database 129
- rerouting clients
 - LDAP 103
- restore database
 - implications 125
- restore_pending configuration parameter 583
- restrict_access configuration parameter 583
- restrictions
 - automatic storage 47, 124
 - naming rules 347
- RESTRICTIVE option
 - CREATE DATABASE
 - database configuration parameter 583
- result tables
 - comparison with other table types 225
- resync_interval configuration parameter 501
- RID_BIT() and RID()
 - built-in functions 250
- RID_BIT() and RID() built-in function 250
- RID_BIT() built-in function 248
- rollforward utility
 - roll forward pending indicator 584
- rollfwd_pending configuration parameter 584
- rollout
 - deferred cleanup 60
- ROW CHANGE TIMESTAMP column 248
- row change timestamps 249
- row compression
 - enabling 49, 243
 - update logs 242
- row identifier (RID_BIT or RID) 245
- row identifier (RID_BIT or RID) built-in function 245
- rows
 - change tokens 248
 - dependent 269
 - descendent 269
 - parent 269
 - self-referencing 269
- rqrioblk configuration parameter 502
- RUNSTATS command
 - automatic statistics collection 50, 54

S

- scenarios
 - time-based update detection 266
- sched_enable configuration parameter 601
- sched_userid configuration parameter 601
- schemas
 - copying 216
 - creating 216
 - db2move COPY errors 219
 - description 211, 215
 - designing 212
 - dropping 222
 - naming restrictions and recommendations 215
 - restarting failed copy schema operation 219
 - troubleshooting tips 216
- scope
 - adding 259
- security
 - plug-ins
 - configuration parameters 465, 467, 506, 508
- security labels (LBAC)
 - component name length 353
 - name length 353
 - policies
 - name length 353
- self tuning memory
 - description 17
 - disabling 27
 - enabling 27, 584
 - non-uniform environments 31
 - limitations 24
 - monitoring 28
 - overview 20
 - partitioned database environments 29, 31
- self_tuning_mem
 - configuration parameter 584
- self-referencing row 269
- self-referencing table 269
- seqdetect configuration parameter 586
- sequence expressions
 - description 326
- sequences
 - application performance 323
 - comparing with identify columns 326
 - comparison with identity columns 324
 - creating 325

- sequences (*continued*)
 - designing 321
 - dropping 328
 - examples 329
 - generating 321
 - managing behavior 322
 - modifying 327
 - recovering databases that use sequences 325
 - using 326
 - values 330
 - viewing 328
- sequential values
 - generating 326
- set integrity pending state 269
- shared file handle table 37
- sheaphres configuration parameter 503
- sheaphres_shr configuration parameter 586
- size limits
 - identifier length 353
 - SQL 353
- size requirements
 - estimating 118
- SMS (system managed space)
 - device considerations 169
 - table spaces
 - compared to DMS table spaces 167
 - creating 189
 - description 154
 - workload considerations 168
- SMS directories
 - in non-automatic storage databases 110
- SMS table spaces
 - altering 193
- smtp_server configuration parameter 601
- softmax configuration parameter 587
- sorheap database configuration parameter
 - description 589
 - effect on query optimization 456
- sorting
 - sort heap size configuration parameter 589
 - sort heap threshold configuration parameter 503
 - sort heap threshold for shared sorts 586
- source tables
 - creating 255
- space compression
 - tables 240
- spm_log_file_sz configuration parameter 505
- spm_log_path configuration parameter 505
- spm_max_resync configuration parameter 506
- spm_name configuration parameter 506
- SQL (Structured Query Language)
 - limits 353
- SQL PL statements
 - supported in trigger-actions 311
- SQL statements
 - displaying help 611
 - inoperative 281
 - optimization
 - configuration parameters 456
 - statement heap size configuration parameter 591
- SQLDBCON configuration file 439, 440
- SQLDBCON database configuration file 113
- SQLDBCONF configuration file 439, 440
- SQLDBCONF database configuration file 113
- srv_plugin_mode configuration parameter 508
- srvcon_auth configuration parameter 506
- srvcon_gssplugin_list configuration parameter 507
- srvcon_pw_plugin configuration parameter 508
- staging tables
 - creating 256
 - dropping 262
- start and stop timeout configuration parameter 508
- start_stop_time configuration parameter 508
- stat_heap_sz database configuration parameter 590
- statement heap size configuration parameter 591
- statistics
 - automatic collection 50, 54
- statistics profiling
 - about 19
- STMM (Self Tuning Memory Manager)
 - enabling 27
 - limitations 24
 - monitoring 28
- STMM (Self-Tuning Memory Manager)
 - enabling 584
- stmheap database configuration parameter 591
 - effect on query optimization 456
- storage
 - automatic
 - table spaces 38, 163
 - automatic for databases 44, 121
- storage paths
 - automatic
 - adding 46, 124
 - monitoring 124
- string
 - data types
 - zero-length 235
- stripe sets 159
 - DMS table spaces 193
- summary tables
 - comparison with other table types 225
 - recovering inoperative 260
- Sun One Directory Server
 - extending directory schema for 93
- svcname configuration parameter 509
- SWITCH ONLINE clause 207
- switching DB2 copies 11
- sysadm_group configuration parameter 510
- SYSCAT.INDEXES view
 - viewing constraint definitions for a table 285
- SYSCATSPACE table spaces 185
- sysctrl_group configuration parameter 510
- sysmaint_group configuration parameter 511
- sysmon_group configuration parameter 512
- system catalog views
 - description 336
- system clock
 - change considerations 249
- system database directory
 - description 114
 - viewing 138
- system managed space (SMS)
 - table spaces
 - description 154
- system temporary table spaces 189
 - page sizes
 - post-migration tasks for DB2 servers 166

T

- table check constraints 269
- table partitioning
 - data organization schemes 253

- table space maps 159
- table spaces
 - adding
 - containers 193
 - altering 193
 - automatic storage 206
 - DMS containers 193
 - SMS containers 193
 - automatic re-sizing 40, 171
 - automatic storage 38, 163
 - choosing extent sizes 182
 - containers
 - extending 202
 - file example 189
 - creating 189
 - database managed space (DMS) 156
 - description 151
 - designing 152
 - device container example 189
 - disk I/O considerations 183
 - dropping 208
 - initial 185
 - mapping to tables 170
 - maps 159
 - page sizes 183
 - performance 207
 - renaming 206
 - resizing container 202
 - switching states 207
 - system managed space (SMS) 154
 - system temporary 189
 - temporary
 - recommendations 165
 - types 154
 - SMS or DMS 167
 - user temporary 189
 - without file system caching 176, 179
 - workload considerations 168
- tables
 - adding columns 258
 - aliases 261
 - append mode 225
 - base 225
 - check constraints
 - overview 235
 - types 269
 - creating
 - overview 253
 - data type definitions 235
 - default columns 235
 - defining
 - referential constraints 277
 - dependent 269
 - descendent 269
 - description 225
 - design concepts 227
 - designing 227
 - dropping 261
 - dropping columns 258
 - estimating size requirements 118
 - examples 263
 - generated columns 232
 - global temporary 225
 - identity columns 233
 - mapping to table spaces 170
 - mismatch 255
 - modifying 257
 - tables (continued)
 - modifying DEFAULT clause column definitions 258
 - multidimensional clustering 225
 - page sizes 238
 - parent 269
 - partitioned 225
 - primary keys 235
 - range-clustered 225
 - referential integrity 235
 - refreshing 257
 - regular 225
 - renaming tables 260
 - result 225
 - scenarios 263
 - self-referencing 269
 - shared file handle 37
 - source 255
 - space compression 240
 - space requirements 236
 - summary 225
 - target 255
 - temporary 225
 - typed 225
 - Unicode table and data considerations 236
 - unique constraints 235
 - user 238
 - viewing definitions 261
 - TCP/IP service name configuration parameter 509
 - temporary table spaces
 - recommendations 165
 - temporary tables
 - comparison with other table types 225
 - global (user-defined) 254
 - TEMPSPACE1 table space 185
 - terms and conditions
 - use of publications 614
 - territory configuration parameter 591
 - time
 - deadlock configuration parameter, interval for
 - checking 541
 - difference among nodes, maximum 493
 - time-based update detection 248
 - scenarios 266
 - TIMESTAMP data type
 - overview 235
 - timestamps
 - row changes 249
 - Tivoli Storage Manager (TSM)
 - management class configuration parameter 592
 - node name configuration parameter 592
 - owner name configuration parameter 593
 - password configuration parameter 593
 - tm_database configuration parameter 512
 - toolscat_db configuration parameter 602
 - toolscat_inst configuration parameter 602
 - toolscat_schema configuration parameter 603
 - tp_mon_name configuration parameter 513
 - track modified pages enable configuration parameter 592
 - trackmod configuration parameter 592
 - transaction processing monitors
 - transaction processing monitor name configuration
 - parameter 513
 - transition tables
 - for triggers, referencing old and new table result sets 313
 - transition variables
 - using triggers, accessing old and new column values 312

- triggered-actions
 - coding 311
 - conditions
 - WHEN clause 311
 - supported SQL PL statements 311
- triggers
 - accessing old and new column values 312
 - activation time 308
 - AFTER
 - example 303
 - AFTER clause 308
 - BEFORE
 - examples 303
 - BEFORE clause 308
 - cascading 301
 - coding triggered-actions 311
 - comparison with check constraints 276
 - conditions 311
 - constraints, interaction 278, 316
 - creating 314
 - description 301
 - designing 305
 - dropping 316
 - examples
 - defining actions 318
 - defining business rules 319
 - preventing operations on tables 320
 - granularity rules 307
 - INSTEAD OF
 - example 304
 - INSTEAD OF clause 308
 - interactions 278, 316
 - maximum name length 353
 - modifying 316
 - referencing old and new table result sets 313
 - triggering events 307
 - types 302
- troubleshooting
 - online information 614
 - tutorials 614
- trust_allclnts configuration parameter 515
- trust_clntauth configuration parameter 515
- tsm_mgmtclass configuration parameter 592
- tsm_nodename configuration parameter 592
- tsm_owner configuration parameter 593
- tsm_password configuration parameter 593
- tuning partition
 - determining 31
- tutorials
 - problem determination 614
 - troubleshooting 614
 - Visual Explain 614
- typed tables
 - comparison with other table types 225
- typed views
 - description 335
 - modifying 343

U

- Unicode
 - table and data considerations 236
- Unicode (UCS-2)
 - identifiers 351
 - naming rules 351
- unique constraints 268
 - definition 269

- unique constraints (*continued*)
 - description 235
 - designing 274
- unique constraint 267
- unique indexes 290
- unique keys
 - description 269
 - generating using sequences 321
- UNIQUERULE column
 - viewing constraint definitions for a table 285
- units of work (UOW)
 - application-directed distributed 134
 - semantics 137
- updatable views
 - using 341
- update rule
 - for referential integrity 269
 - referential constraints 269
- updates
 - DB2 Information Center 612
- user data
 - directories 474
- user exit enable configuration parameter 594
- user exit status indicator configuration parameter 594
- user IDs
 - naming rules 350
- user table page limits 238
- user temporary table spaces
 - creating 189
- user_exit_status configuration parameter 594
- user-defined (global) temporary tables
 - creating 254
- user-defined functions
 - used with views 342
- userexit database configuration parameter 594
- USERSPACE1 table space 185
- util_heap_sz configuration parameter 595
- util_impact_lim configuration parameter
 - described 516
- utility operations
 - constraint implications 281
- utility throttling
 - about 58
 - description 17

V

- values
 - sequence 330
- VARCHAR data type
 - in table columns 259
- vendor code
 - in fenced vendor processes 37
- vendoropt configuration parameter 595
- view
 - aliases 261
- views
 - creating 341
 - definition of nested views 339
 - deletable
 - using 339
 - description 335
 - designing 336
 - dropping 344
 - inoperative 343
 - insertable 340
 - modifying 343

- views (*continued*)
 - overview 335
 - read-only
 - using 341
 - recovering inoperative 343
 - updatable 341
 - using user-defined functions 342
 - with check option
 - examples 337
- Vista
 - user data directories 474
- Visual Explain
 - tutorial 614
- vmo AIX system command 4
- vmtune AIX system command 4

W

- Windows operating systems
 - active directory
 - DB2 object creation 96
 - LDAP object classes and attributes 78
 - extending the directory schema
 - Windows 2003 97
- wizards
 - Configuration Advisor 116
- wlm_collect_int database configuration parameter
 - description 596

X

- XQuery statements
 - inoperative 281
 - optimization
 - configuration parameters 456
 - statement heap size configuration parameter 591



Printed in USA

SC23-5849-02



Spine information:

DB2 Version 9.5 for Linux, UNIX, and Windows

Data Servers, Databases, and Database Objects Guide

