




Administration Guide: Planning



Administration Guide: Planning

Before using this information and the product it supports, be sure to read the general information under *Notices*.

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, 2006. All rights reserved.**

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

Contents

About this book.	vii
Who should use this book	viii
How this book is structured	viii

Part 1. Database concepts. 1

Chapter 1. Basic relational database concepts 3

About databases	3
Database objects	3
Configuration parameters.	12
Environment variables and the profile registry.	14
Business rules for data.	17
Data security	20
Authentication	20
Authorization	21
Units of work	22
High availability disaster recovery (HADR) feature overview	23
Developing a backup and recovery strategy	25

Chapter 2. Automatic maintenance. 29

About automatic maintenance	29
Automatic features enabled by default	30
Automatic database backup	31
Automatic reorganization.	32
Automatic statistics collection by table	33
Automatic statistics profiling using automatic statistics collection	34
Storage used by automatic statistics collection and profiling	35
Maintenance windows.	35
Offline maintenance	36
Online maintenance	36

Chapter 3. Parallel database systems 37

Parallelism	37
Input/output parallelism	37
Query parallelism	37
Utility parallelism	40
Partitioned database environments	41
Database partition and processor environments	42
Single database partition on a single processor	42
Single database partition with multiple processors	43
Multiple database partition configurations	44
Summary of parallelism best suited to each hardware environment	48

Part 2. Database design 51

Chapter 4. Logical database design 53

What to record in a database	53
--	----

Database relationships.	54
One-to-many and many-to-one relationships	54
Many-to-many relationships	55
One-to-one relationships	55
Ensure that equal values represent the same entity	56
Column definitions	56
Primary keys	58
Identifying candidate key columns	59
Identity columns	60
Normalization	61
First normal form	62
Second normal form	62
Third normal form	63
Fourth normal form	64
Constraints	65
Unique constraints	66
Referential constraints	66
Table check constraints	69
Informational constraints	69
Triggers	70
Additional database design considerations	71

Chapter 5. Physical database design 73

Database directories and files	73
Space requirements for database objects	75
Space requirements for system catalog tables	76
Space requirements for user table data	77
Space requirements for long field data	78
Space requirements for large object data.	79
Space requirements for indexes.	80
Space requirements for log files.	82
Space requirements for temporary tables	83
XML storage object overview	84
Guidelines for storage requirements for XML documents.	84
Database partition groups	85
Database partition group design	87
Distribution maps	88
Distribution keys	89
Table collocation.	91
Database partition compatibility	91
Data partitions	92
Table partitioning	93
Table partitioning keys	96
Data organization schemes	99
Partitioned tables	104
Data organization schemes in DB2 and Informix databases.	105
Replicated materialized query tables.	111
Table space design.	112
SYSTOOLSPACE and SYSTOOLSTMPSPACE table spaces	115
System managed space	117
SMS table spaces	119

Database managed space	120
DMS table spaces	123
DMS device considerations	124
Table space maps	125
How containers are added and extended in DMS table spaces	129
Rebalancing	129
Without rebalancing (using stripe sets)	135
How containers are dropped and reduced in DMS table spaces	137
Comparison of SMS and DMS table spaces	140
Table space disk I/O	141
Workload considerations in table space design	143
Extent size	144
Relationship between table spaces and buffer pools	145
Relationship between table spaces and database partition groups	146
Storage management view	146
Stored procedures for the storage management tool	147
Storage management view tables	148
Thresholds	160
Temporary table space design	161
Temporary tables in SMS table spaces	162
Catalog table space design	163
Optimizing table space performance when data is on RAID devices	164
Considerations when choosing table spaces for your tables	166
DB2 table types.	167
Range-clustered tables	168
Range-clustered tables and out-of-range record key values	171
Range-clustered table locks	172
Multidimensional clustering tables	172
Comparison of regular and MDC tables	173
Block indexes	175
Working with an MDC table	177
Block indexes and query performance	180
Maintaining clustering automatically during INSERT operations	183
Block maps	185
Deletion from an MDC table	187
Updating an MDC table.	187
Load considerations for MDC tables.	188
Logging considerations for MDC tables.	188
Block index considerations for MDC tables	188
Designing multidimensional clustering (MDC) tables	189
Multidimensional clustering (MDC) table creation, placement, and use	197

Chapter 6. Designing partitioned databases 203

Updating a single database in a transaction	203
Using multiple databases in a single transaction	204
Updating a single database in a multi-database transaction	204
Updating multiple databases in a transaction	205
DB2 transaction manager	206
DB2 Database transaction manager configuration	207

Updating a database from a host or iSeries client	210
Two-phase commit	210
Error recovery during two-phase commit	213
Error recovery if autorestart=off	214

Chapter 7. Designing for XA-compliant transaction managers. 215

X/Open distributed transaction processing model	215
Application program (AP)	216
Transaction manager (TM)	217
Resource managers (RM)	218
Resource manager setup.	219
Database connection considerations	219
xa_open string formats	221
Updating host or iSeries database servers with an XA-compliant transaction manager	227
Resolving indoubt transactions manually	227
Indoubt transaction management APIs	230
Security considerations for XA transaction managers.	231
Configuration considerations for XA transaction managers.	232
XA function supported by DB2 Database for Linux, UNIX, and Windows	233
XA switch usage and location	233
Using the DB2 Database for Linux, UNIX, and Windows XA switch	234
XA interface problem determination.	235
XA transaction manager configuration	236
Configuring IBM WebSphere Application Server	236
Configuring IBM TXSeries CICS	236
Configuring IBM TXSeries Encina	236
Configuring BEA Tuxedo	238

Part 3. Appendixes 241

Appendix A. Incompatibilities between releases. 243

Deprecated and discontinued features	243
Version 9 incompatibilities with previous releases and changed behaviors	260
Version 8 incompatibilities with previous releases	285

Appendix B. National language support (NLS) 313

National language versions.	313
Supported territory codes and code pages.	313
Availability of Asian fonts (Linux)	334
Simplified Chinese locale coding set.	335
Displaying Indic characters in the DB2 GUI tools	336
Enabling and disabling euro symbol support.	336
Character-conversion guidelines	338
Conversion table files for euro-enabled code pages	339
Conversion tables for code pages 923 and 924	343
Choosing a language for your database.	344
Locale setting for the DB2 Administration Server	345
Enabling bidirectional support.	345
Bidirectional-specific CCSIDs	347

Bidirectional support with DB2 Connect	349
Collating sequences	351
Collating Thai characters	352
Date and time formats by territory code	353
Unicode character encoding	355
UCS-2	355
UTF-8	356
UTF-16	356
Unicode implementation in DB2 Database for Linux, UNIX, and Windows	357
AIX, UNIX, and Linux distributions and code pages	358
Code Page/CCSID Numbers	359
Thai and Unicode collation algorithm differences	360
Unicode handling of data types	360
Creating a Unicode database	362
Converting non-Unicode databases to Unicode	362
Unicode literals.	364
String comparisons in a Unicode database.	364
Installing the previous tables for converting between code page 1394 and Unicode	366
Alternative Unicode conversion table for the coded character set identifier (CCSID) 943	366
Replacing the Unicode conversion tables for coded character set identifier (CCSID) 943 with Microsoft conversion tables	368
Alternative Unicode conversion table for the coded character set identifier (CCSID) 954	369
Replacing the Unicode conversion table for coded character set identifier (CCSID) 954 with the Microsoft conversion table	370
Alternative Unicode conversion table for the coded character set identifier (CCSID) 5026	371
Replacing the Unicode conversion table for coded character set identifier (CCSID) 5026 with the Microsoft conversion table	371

Alternative Unicode conversion table for the coded character set identifier (CCSID) 5035	372
Replacing the Unicode conversion table for coded character set identifier (CCSID) 5035 with the Microsoft conversion table	373
Alternative Unicode conversion table for the coded character set identifier (CCSID) 5039	374
Replacing the Unicode conversion table for coded character set identifier (CCSID) 5039 with the Microsoft conversion table	375

Appendix C. DB2 Database technical information 377

Overview of the DB2 technical information	377
Documentation feedback	377
DB2 technical library in hardcopy or PDF format	378
Ordering printed DB2 books	380
Displaying SQL state help from the command line processor	381
Accessing different versions of the DB2 Information Center	382
Displaying topics in your preferred language in the DB2 Information Center	382
Updating the DB2 Information Center installed on your computer or intranet server	383
DB2 tutorials	385
DB2 troubleshooting information	385
Terms and Conditions	386

Appendix D. Notices 387

Trademarks	389
----------------------	-----

Index 391

Contacting IBM 399

About this book

The Administration Guide: Planning provides information necessary to use and administer the DB2[®] relational database management system (RDBMS) products, and includes information about database planning and design.

Many of the tasks described in this book can be performed using different interfaces:

- The **Command Line Processor**, which allows you to access and manipulate databases from a command-line interface. From this interface, you can also execute SQL statements and DB2 utility functions. Most examples in this book illustrate the use of this interface. For more information about using the command line processor, see the *Command Reference*.
- The **application programming interface**, which allows you to execute DB2 utility functions within an application program. For more information about using the application programming interface, see the *Administrative API Reference*.
- The **Control Center**, which allows you to use a graphical user interface to manage and administer your data and database components. You can invoke the Control Center using the db2cc command on a Linux[™] or Windows[®] command line, or using the Start menu on Windows platforms. The Control Center presents your database components as a hierarchy of objects in an object tree. This Control Center tree includes your systems, instances, databases, tables, views, triggers, and indexes. From the tree you can perform actions on your database objects, such as creating new tables, reorganizing data, configuring and tuning databases, and backing up and restoring table spaces. In many cases, wizards and launchpads are available to help you perform these tasks more quickly and easily.

The Control Center is available in three views:

- **Basic**. This view provides you with the core DB2 database functions. From this view you can work with all the databases to which you have been granted access, including their related objects such as tables and stored procedures. It provides you with the essentials for working with your data.
- **Advanced**. This view provides you with all of the objects and actions available in the Control Center. Use this view if you are working in an enterprise environment and you want to connect to DB2 Version 9.1 for z/OS[®] (DB2 for z/OS) or IMS.
- **Custom**. This view provides you with the ability to tailor the Control Center to your needs. You select the objects and actions that you want to appear in your view.

For help on using the Control Center, select **Getting started** from the **Help** pull-down on the Control Center window.

There are other tools that you can use to perform administration tasks. They include:

- The **Command Editor** which replaces the Command Center and is used to generate, edit, run, and manipulate SQL statements; IMS and DB2 commands; work with the resulting output; and to view a graphical representation of the access plan for explained SQL statements.

- The Development Center which provides support for native SQL Persistent Storage Module (PSM) stored procedures; for Java™ stored procedures for iSeries™ Version 5 Release 3 and later; user-defined functions (UDFs); and structured types.
- The Health Center which provides a tool to assist DBAs in the resolution of performance and resource allocation problems.
- The Tools Settings which you can use to change the settings for the Control Center, and the Health Center.
- The Memory Visualizer which helps database administrators monitor the memory-related performance of an instance and all of its databases organized in a hierarchical tree.
- The Indoubt Transaction Manager window which is used to display indoubt transactions. That is, the transactions that are waiting to be committed, rolled back, or forgotten for a selected database and one or more selected partitions.
- The Information Catalog Manager which is used to provide a graphical representation of data relationships and object definitions when working in a warehouse environment.
- The Journal which you can use to schedule jobs that are to run unattended.
- The Data Warehouse Center which manages warehouse objects.

Who should use this book

This book is intended primarily for database administrators, system administrators, security administrators and system operators who need to plan and design databases that can 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

The major subject areas discussed in the chapters of this book are as follows:

Database Concepts

- Chapter 1, “Basic relational database concepts,” presents an overview of database objects and database concepts.
- Chapter 3, “Parallel database systems,” provides an introduction to the types of parallelism available with DB2 databases.

Database Design

- Chapter 4, “Logical database design,” discusses the concepts and guidelines for logical database design.
- Chapter 5, “Physical database design,” discusses the guidelines for physical database design, including space requirements and table space design.
- Chapter 6, “Designing partitioned databases,” discusses how you can access multiple databases in a single transaction.
- Chapter 7, “Designing for XA-compliant transaction managers,” discusses how you can use your databases in a distributed transaction processing environment.

Appendixes

- Appendix A, “Incompatibilities between releases,” presents the incompatibilities introduced by Version 8 and Version 9, as well as planned future incompatibilities.

- Appendix B, “National language support (NLS),” introduces DB2 National Language Support, including information about territories, languages, and code pages.

Part 1. Database concepts

Chapter 1. Basic relational database concepts

About databases

A relational database presents data as a collection of tables. A table consists of a defined set of columns and any number of rows. The data in each table is logically related, and relationships can be defined between tables. Data can be viewed and manipulated based on mathematical principles and operations called relations (such as, INSERT, SELECT, and UPDATE).

A database is self-describing in that it contains, in addition to data, a description of its own structure. It includes a set of system catalog tables, which describe the logical and physical structure of the data; a configuration file, which contains the parameter values associated with the database; and a recovery log, which records ongoing transactions and transactions that can be archived.

Databases can be local or remote. A local database is physically located on the workstation in use, while a database on another machine is considered remote.

You can:

- Create a database.
- Add a database to the Control Center.
- Drop a database from the Control Center.
- Back up a database.
- Restore a database.
- Configure a database.
- Catalog a database.
- Uncatalog a database.
- Connect to a database.
- Monitor a database with the event monitor.
- Work with partitioned databases.
- Work with federated systems.

For z/OS and OS/390® systems, the default database, DSNDB04, is predefined in the DB2 installation process. This database has a default buffer pool (BP0), and a default DB2 storage group (SYSDEFLT).

Related concepts:

- “Tables” in *SQL Reference, Volume 1*

Database objects

Systems:

DB2 databases are organized around a hierarchy of database objects. The highest-level object in the hierarchy is a system. A system represents an installation of DB2. The Control Center maintains a list of systems that it knows about and records the information needed to communicate with each system (such as its network address, operating system, and communication protocol). The Control Center supports both DB2 and IMS™ systems.

A system can have one or more DB2 instances, each of which can manage one or more databases. The databases may be partitioned with their table spaces residing in database partition groups. The table spaces in turn store table data.

You can:

- Add a system to the Control Center.
- Attach to a system.
- Remove a system from the Control Center.

Instances:

An *instance* (sometimes called a *database manager*) is DB2 code that manages data. It controls what can be done to the data, and manages system resources assigned to it. Each instance is a complete environment. It contains all the database partitions defined for a given parallel database system. An instance has its own databases (which other instances cannot access), and all its database partitions share the same system directories. It also has security separate from other instances on the same computer (system).

Databases:

A *relational database* presents data as a collection of tables. A table consists of a defined number of columns and any number of rows. Each database includes a set of system catalog tables that describe the logical and physical structure of the data, a configuration file containing the parameter values allocated for the database, and a recovery log with ongoing transactions and transactions to be archived.

Database partitions:

A *database partition* consists of its own data, indexes, configuration files, and transaction key. It is sometimes referred to as a node or database node. Tables can be located in one or more database partitions. When a table's data is distributed across multiple partitions, some of its rows are stored in one partition, and other rows are stored in other partitions. Data retrieval and update requests are decomposed automatically into sub-requests, and run in parallel among the applicable database partitions. The fact that a database can be split across database partitions is transparent to users.

Database partition groups:

A *database partition group* is a set of one or more database partitions. When you want to create tables for the database, you first create the database partition group where the table spaces will be stored, then you create the table space where the tables will be stored.

In earlier versions of DB2 Universal Database™ (UDB), *database partition groups* were known as *nodegroups*.

Table spaces:

A database is organized into parts called *table spaces*. A table space is a place to store tables. When creating a table, you can decide to have certain objects such as indexes and large object (LOB) data kept separately from the rest of the table data. A table space can also be spread over one or more physical storage devices. The

following diagram shows some of the flexibility you have in spreading data over table spaces:

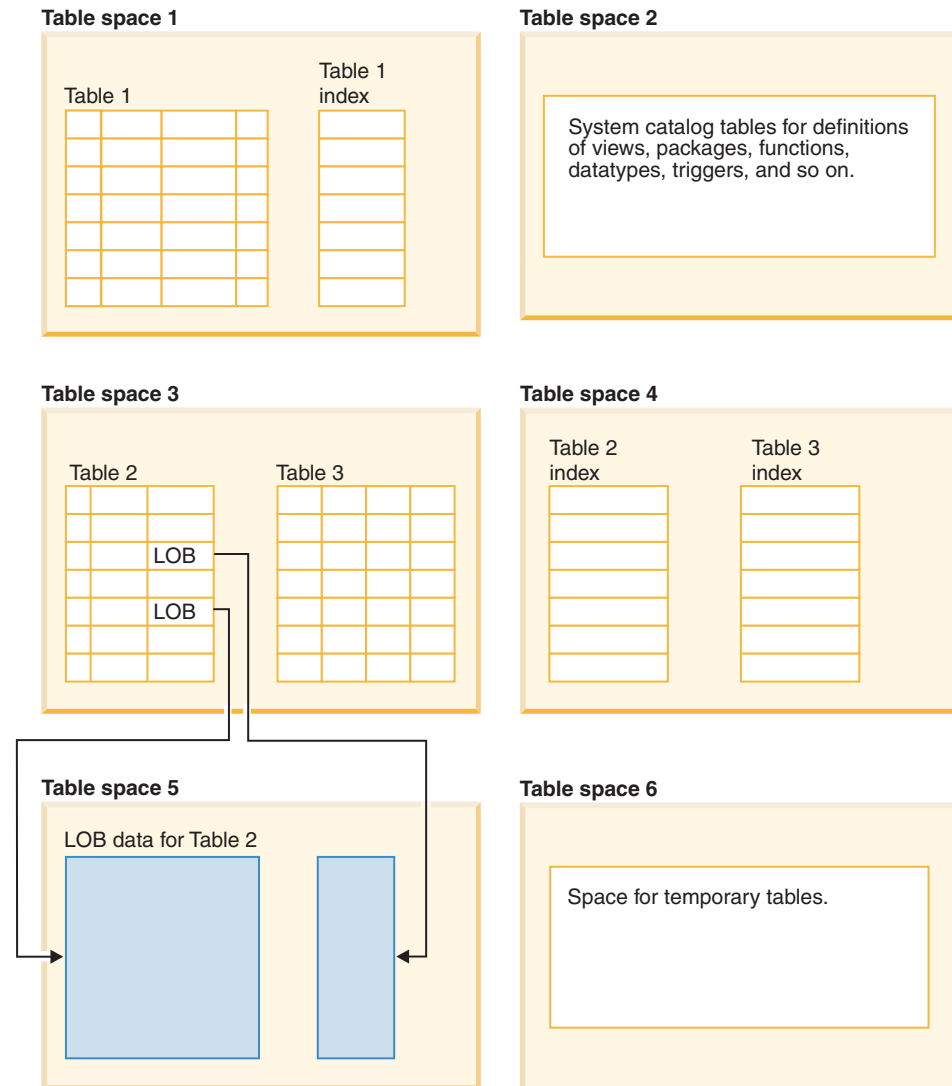


Figure 1. Table space flexibility

Table spaces reside in database partition groups. Table space definitions and attributes are recorded in the database system catalog.

Containers are assigned to table spaces. A *container* is an allocation of physical storage (such as a file or a device).

A table space can be either system managed space (SMS), or database managed space (DMS). For an SMS table space, each container is a directory in the file space of the operating system, and the operating system's file manager controls the storage space. For a DMS table space, each container is either a fixed size pre-allocated file, or a physical device such as a disk, and the database manager controls the storage space.

Figure 2 illustrates the relationship between tables, table spaces, and the two types of space. It also shows that tables, indexes, and long data are stored in table spaces.

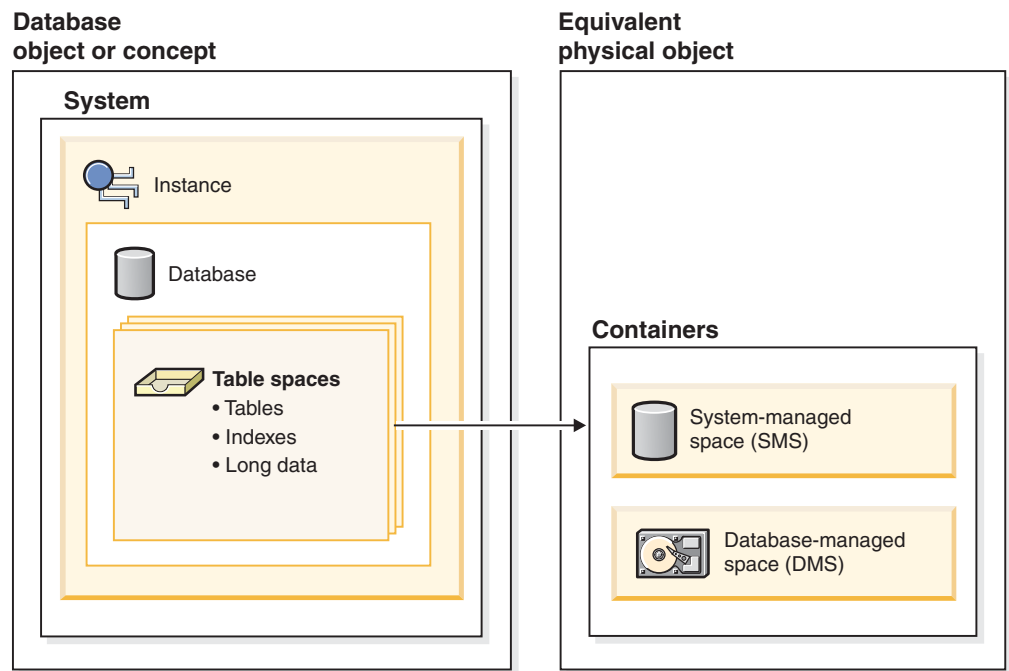


Figure 2. Table spaces and container types that hold data

Figure 3 on page 7 shows the three table space types: *regular*, *temporary*, and *large*.

Tables containing user data exist in regular table spaces. The default user table space is called `USERSPACE1`. The system catalog tables exist in a regular table space. The default system catalog table space is called `SYSCATSPACE`.

Tables containing long field data or large object data, such as multimedia objects, exist in large table spaces or in regular table spaces. The base column data for these columns is stored in a regular table space, while the long field or large object data can be stored in the same regular table space or in a specified large table space.

Indexes can be stored in regular table spaces or large table spaces.

Temporary table spaces are classified as either system or user. *System temporary table spaces* are used to store internal temporary data required during SQL operations such as sorting, reorganizing tables, creating indexes, and joining tables. These operations require extra space to process the result set. Although you can create any number of system temporary table spaces, it is recommended that you create only one, using the page size that the majority of your tables use. The default system temporary table space is called `TEMPSPACE1`. Any user and application may use system temporary table spaces. *User temporary table spaces* are used to store declared global temporary tables that store application temporary data. The tables used within user temporary table spaces are created using the `DECLARE GLOBAL TEMPORARY TABLE` statement. User temporary table spaces are *not* created by default at database creation time. Access to user temporary table spaces

is controlled. Remember to grant the appropriate USE privileges on the user temporary tables with the GRANT statement.

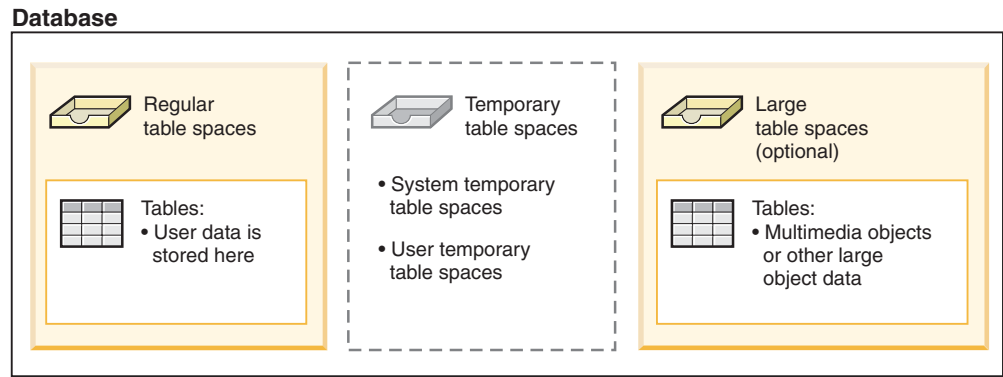


Figure 3. Types of table spaces

Tables:

A relational database presents data as a collection of tables. A *table* consists of data logically arranged in columns and rows. All database and table data is assigned to table spaces. The data in the table is logically related, and relationships can be defined between tables. Data can be viewed and manipulated based on mathematical principles and operations called *relations*.

Table data is accessed through Structured Query Language (SQL), a standardized language for defining and manipulating data in a relational database. A *query* is used in applications or by users to retrieve data from a database. The query uses SQL to create a statement in the form of

```
SELECT <data_name> FROM <table_name>
```

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 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 4 on page 8 shows the relationship between tables and views.

Database

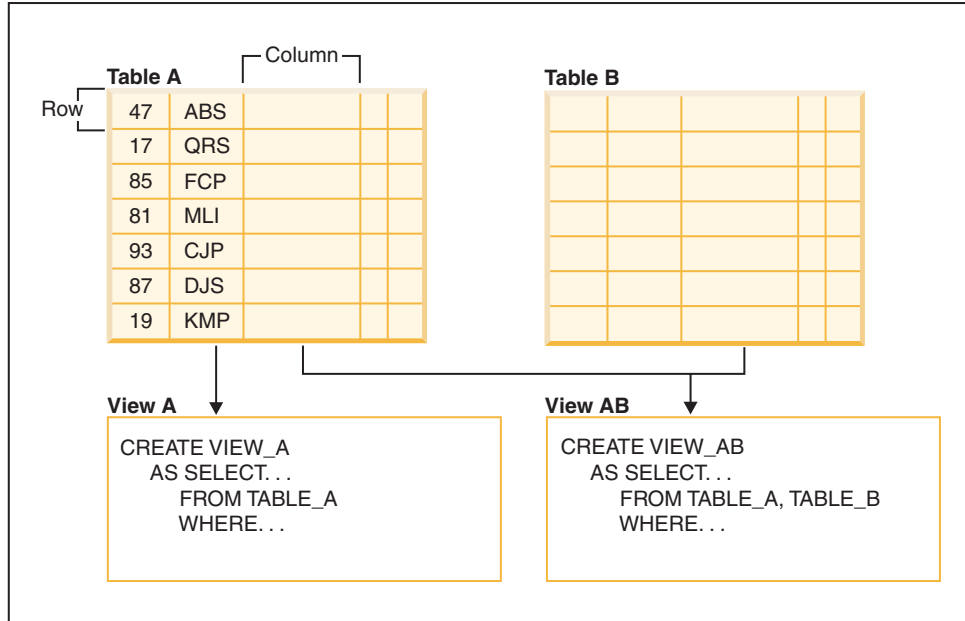


Figure 4. Relationship between tables and views

Indexes:

As data is added to a table, unless other actions have been carried out on the table or the data being added, it 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 field or column from within a row of data may be used as a value that can identify the entire row. One or more columns may be needed to identify the row. This identifying column or columns is known as a *key*. A column may be used in more than one key.

An index is ordered by the values within a key.

keys may be unique or non-unique. Each table should have at least one unique key; but may 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.

An *index* is a set of one or more keys, each key pointing to a row in a table. For example, table A in Figure 5 on page 9 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.

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.

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 5 shows the relationship between an index and a table.

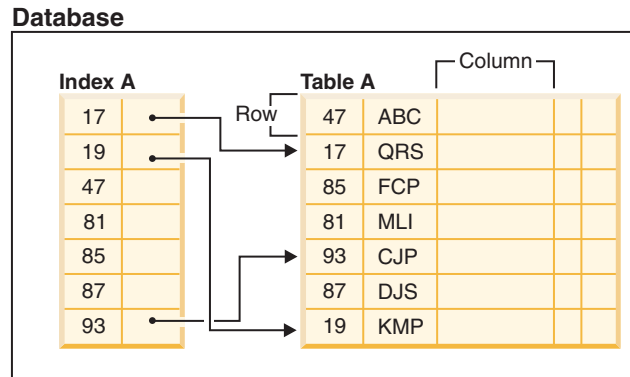


Figure 5. Relationship between an index and a table

Figure 6 illustrates the relationships among some database objects. It also shows that tables, indexes, and long data are stored in table spaces.

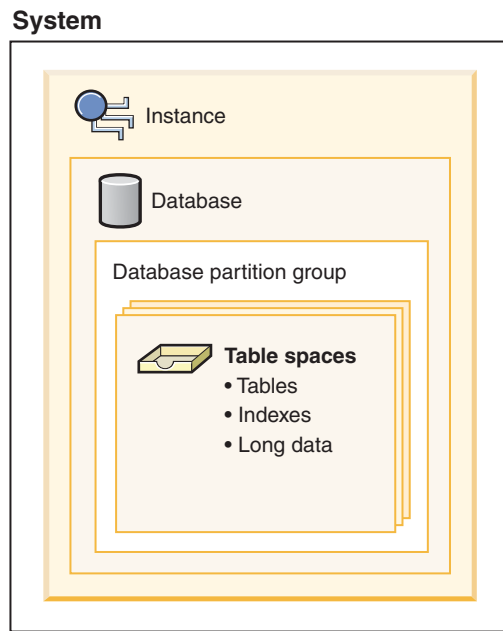


Figure 6. Relationships among selected database objects

Schemas:

A *schema* is an identifier, such as a user ID, that helps group tables and other database objects. A schema can be owned by an individual, and the owner can control access to the data and the objects within it.

A schema is also an object in the database. It may be created automatically when the first object in a schema is created. Such an object can be anything that can be

qualified by a schema name, such as a table, index, view, package, distinct type, function, or trigger. You must have `IMPLICIT_SCHEMA` authority if the schema is to be created automatically, or you can create the schema explicitly.

A schema name is used as the first part of a two-part object name. When an object is created, you can assign it to a specific schema. If you do not specify a schema, it is assigned to the default schema, which is usually the user ID of the person who created the object. The second part of the name is the name of the object. For example, a user named Smith might have a table named `SMITH.PAYROLL`.

System catalog tables:

Each database includes a set of *system catalog tables*, which describe the logical and physical structure of the data. DB2 creates and maintains an extensive set of system catalog tables for each database. These tables contain information about the definitions of database objects such as user tables, views, and indexes, as well as security information about the authority that users have on these objects. They are created when the database is created, and are updated during the course of normal operation. You cannot explicitly create or drop them, but you can query and view their contents using the catalog views.

Containers:

A *container* is a physical storage device. It can be identified by a directory name, a device name, or a file name.

A container is assigned to a table space. A single table space can span many containers, but each container can belong to only one table space.

Figure 7 on page 11 illustrates the relationship between tables and a table space within a database, and the associated containers and disks.

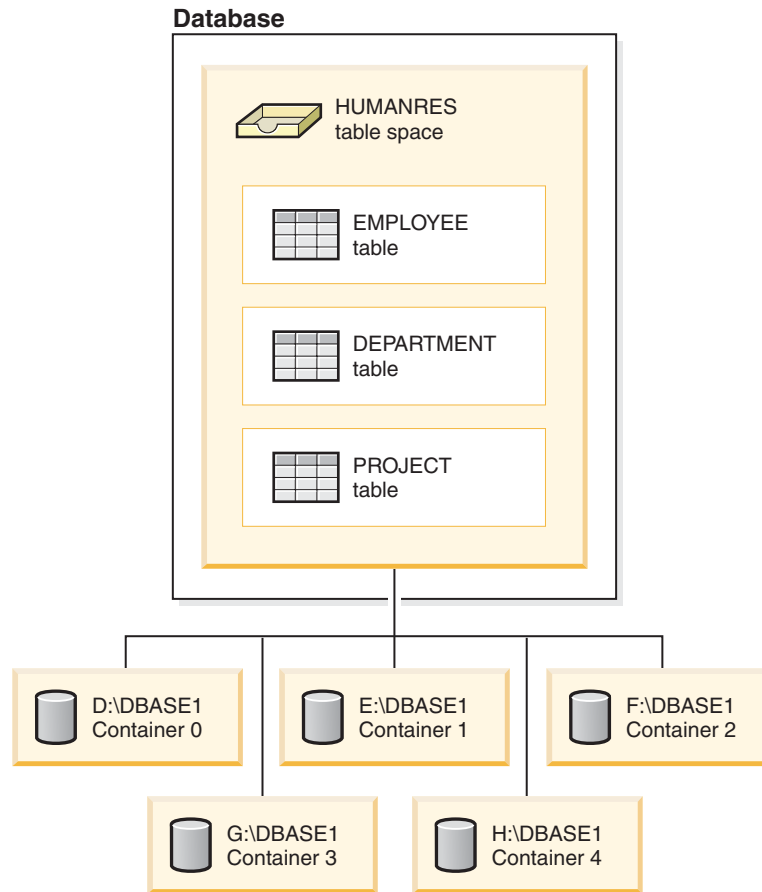


Figure 7. Relationship between a table space and its containers

The EMPLOYEE, DEPARTMENT, and PROJECT tables are in the HUMANRES table space which spans containers 0, 1, 2, 3, and 4. This example shows each container existing on a separate disk.

Data for any table will be stored on all containers in a table space in a round-robin fashion. This balances the data across the containers that belong to a given table space. The number of pages that the database manager writes to one container before using a different one is called the *extent size*.

Buffer pools:

A *buffer pool* is the amount of main memory allocated to cache table and index data pages as they are being read from disk, or being modified. The purpose of the buffer pool is to improve system performance. Data can be accessed much faster from memory than from disk; therefore, the fewer times the database manager needs to read from or write to a disk (I/O), the better the performance. (You can create more than one buffer pool, although for most situations only one is required.)

The configuration of the buffer pool is the single most important tuning area, because you can reduce the delay caused by slow I/O.

Figure 8 on page 12 illustrates the relationship between a buffer pool and containers.

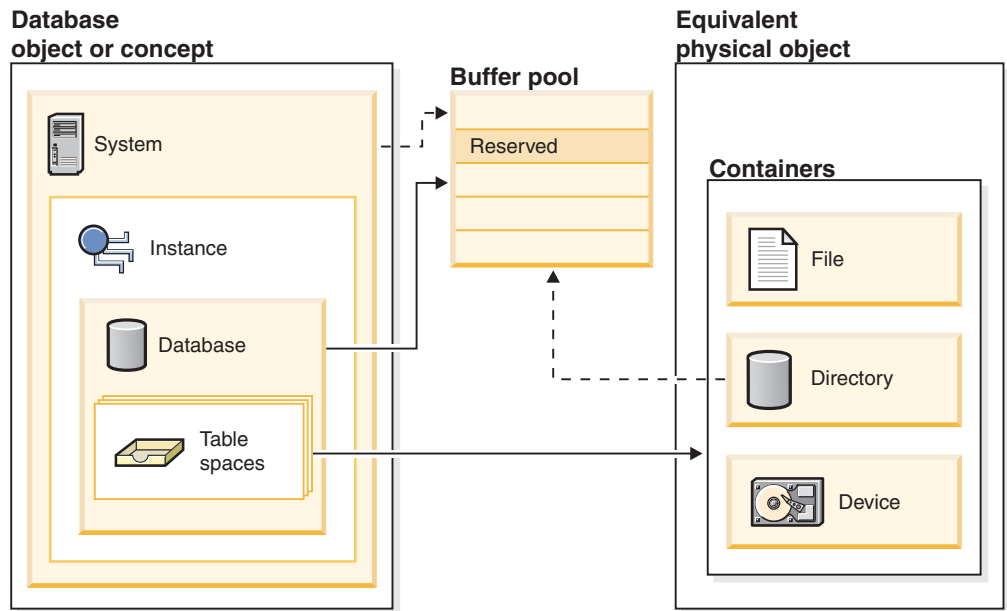


Figure 8. Relationship between the buffer pool and containers

Related concepts:

- “Indexes” in *SQL Reference, Volume 1*
- “Relational databases” in *SQL Reference, Volume 1*
- “Schemas” in *SQL Reference, Volume 1*
- “Table spaces and other storage structures” in *SQL Reference, Volume 1*
- “Tables” in *SQL Reference, Volume 1*
- “Views” in *SQL Reference, Volume 1*

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.

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 UNIX[®]-based 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 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.

Parameters for an individual database are stored in a configuration file named `SQLDBCON`. 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.

In a partitioned database environment, a separate `SQLDBCON` file exists for each database partition. The values in the `SQLDBCON` file may be the same or different at each database partition, but the recommendation is that the database configuration parameter values be the same on all database partitions.

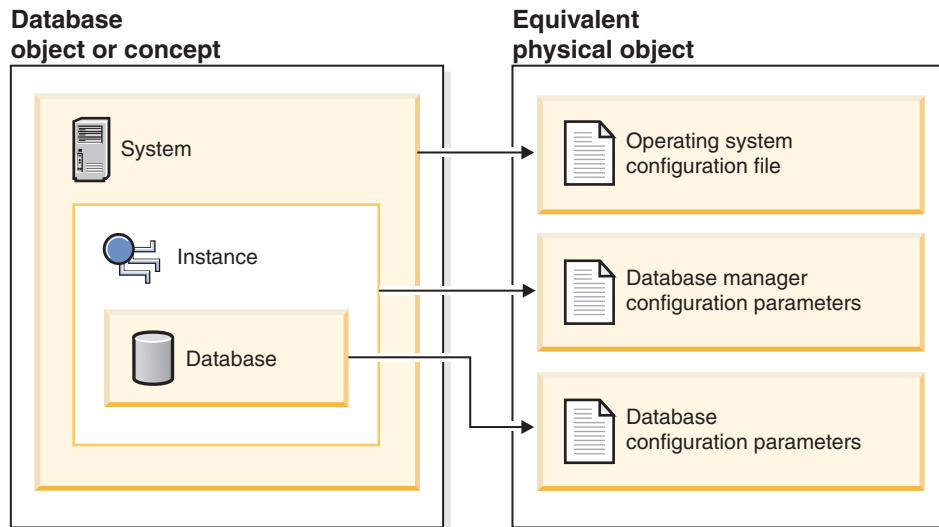


Figure 9. Relationship between database objects and configuration files

Related concepts:

- “Configuration parameters that affect query optimization” in *Performance Guide*

Related tasks:

- “Configuring DB2 with configuration parameters” in *Performance Guide*

Environment variables and the profile registry

Environment and registry variables control your database environment.

You can use the Configuration Assistant (**db2ca**) to configure configuration parameters and registry variables.

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. Windows users do not need SYSADM authority to update registry variables. Use the **db2set** command to update registry variables without restarting; this information is stored immediately in the profile registries. The DB2 registry applies the updated information to DB2 server instances and DB2 applications started after the changes are made.

When updating the registry, changes do not affect the currently running DB2 applications or users. Applications started following the update use the new values.

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 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 F00=BAR
```

or

```
db2set F00=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.

Related concepts:

- “DB2 registry and environment variables” in *Performance Guide*

Related tasks:

- “Declaring, showing, changing, resetting, and deleting registry and environment variables” in *Administration Guide: Implementation*

Business rules for data

Within any business, data must often adhere to certain restrictions or rules. For example, an employee number must be unique. DB2 Database for Linux, UNIX, and Windows provides *constraints* as a way to enforce such rules. Triggers are also used to enforce business rules on your data.

DB2 V9.1 provides the following types of constraints:

- NOT NULL constraint
- Unique constraint
- Primary key constraint
- Foreign key constraint
- Check constraint
- Informational constraint

NOT NULL constraint

NOT NULL constraints prevent null values from being entered into a column.

unique constraint

Unique constraints ensure that the values in a set of columns are unique and not null for all rows in the table. For example, a typical unique constraint in a DEPARTMENT table might be that the department number is unique and not null.

The following figure shows that a duplicate record is prevented from being added to a table when a unique constraint exists for the table.

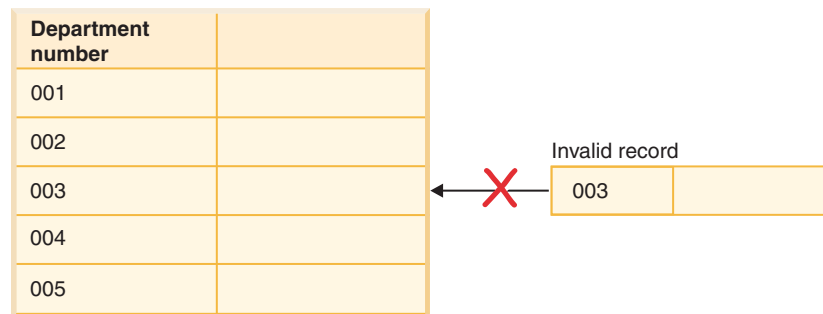


Figure 10. Unique constraints prevent duplicate data

The database manager enforces the constraint during insert and update operations, ensuring data integrity.

primary key constraint

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. They are also beneficial, because they order the data when data is exported or reorganized.

In the following tables, DEPTNO and EMPNO are the primary keys for the DEPARTMENT and EMPLOYEE tables.

Table 1. DEPARTMENT Table

DEPTNO (Primary Key)	DEPTNAME	MGRNO
A00	Spiffy Computer Service Division	000010
B01	Planning	000020
C01	Information Center	000030
D11	Manufacturing Systems	000060

Table 2. EMPLOYEE Table

EMPNO (Primary Key)	FIRSTNAME	LASTNAME	WORKDEPT (Foreign Key)	PHONENO
000010	Christine	Haas	A00	3978
000030	Sally	Kwan	C01	4738
000060	Irving	Stern	D11	6423
000120	Sean	O'Connell	A00	2167
000140	Heather	Nicholls	C01	1793
000170	Masatoshi	Yoshimura	D11	2890

foreign key constraint

Foreign key constraints (also known as 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.

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.

The following figure 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.

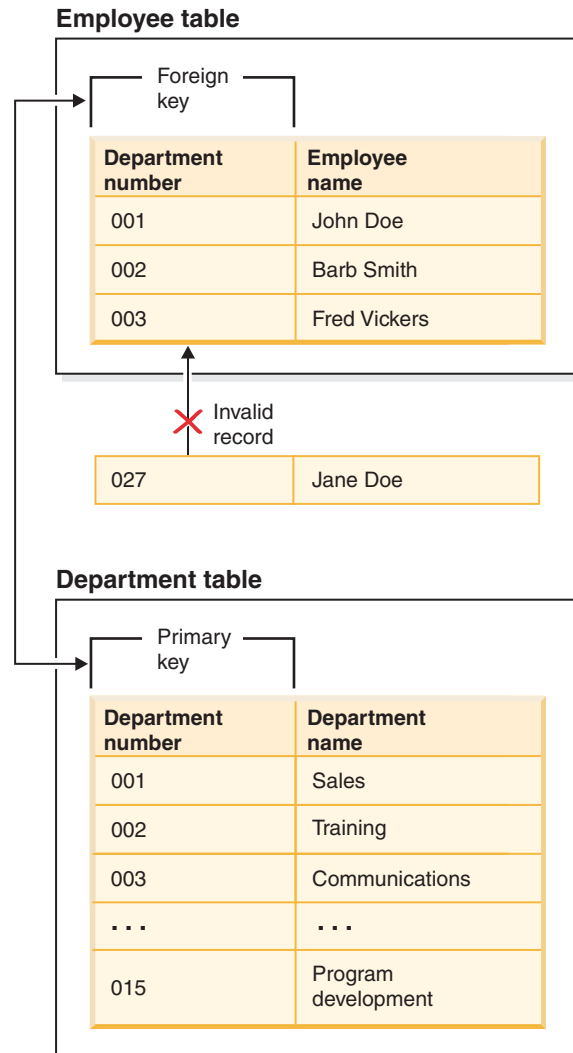


Figure 11. Foreign and primary key constraints

check constraint

A check constraint is a database rule that specifies the values allowed in one or more columns of every row of a table.

For example, in an EMPLOYEE table, you can define the Type of Job column to be "Sales", "Manager", or "Clerk". With this constraint, any record with a different value in the Type of Job column is not valid, and would be rejected, enforcing rules about the type of data allowed in the table.

informational constraint

An informational constraint is a rule that can be used by the SQL compiler but is not enforced by the database manager. The purpose of the constraint is not to have additional verification of data by the database manager, rather it is to improve query performance.

Informational constraints are defined using the CREATE TABLE or ALTER TABLE statements. You add referential integrity or check constraints but 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.

In addition to using constraints to enforce business rules on your data, you can also use *triggers* in your database. Triggers are more complex and potentially more powerful than constraints. They define a set of actions that are executed in conjunction with, or triggered by, an INSERT, UPDATE, or DELETE clause on a specified base table. 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 be used in a banking application to raise an alert if a withdrawal from an account did not fit a customer's standard withdrawal patterns.

Related concepts:

- "Constraints" on page 65
- "Triggers" on page 70

Data security

Two security levels control access to DB2 Database for Linux, UNIX, and Windows data and functions. Access to DB2 is managed by facilities specific to the operating environment (authentication), whereas access within DB2 is managed by the database manager (authorization).

Authentication is the process by which a system verifies a user's identity. User authentication is completed by a security facility outside DB2, often part of the operating system or a separate product.

Once a user is authenticated, the databasemanager determines if that user is allowed to access DB2 data or resources. Authorization is the process whereby DB2 obtains information about the authenticated user, indicating which database operations the user can perform, and which data objects the user can access. An authorization ID designates the authorized user's access. Authorization can be broken down into two categories: privileges and authorities.

Privileges enable a user to create or access database resources. Authorities provide a way both to group privileges, and to control maintenance and utility operations for instances, databases, and database objects.

Related concepts:

- "About databases" on page 3

Authentication

Authentication of a user is completed using a security facility outside of DB2 Database for Linux, UNIX, and Windows. The security facility can be part of the operating system, a separate product or, in certain cases, may not exist at all. On UNIX based systems, the security facility is in the operating system itself.

The security facility requires two items to authenticate a user: a user ID and a password. The user ID identifies the user to the security facility. By supplying the correct password, information known only to the user and the security facility, the user's identity (corresponding to the user ID) is verified.

Once authenticated:

- The user must be identified to DB2 using an SQL authorization name or *authid*. This name can be the same as the user ID, or a mapped value. For example, on

UNIX operating systems, a DB2 *authid* is derived by transforming to uppercase letters a UNIX user ID that follows DB2 naming conventions.

- A list of groups to which the user belongs is obtained. Group membership may be used when authorizing the user. Groups are security facility entities that must also map to DB2 authorization names. This mapping is done in a method similar to that used for user IDs.

DB2 V9.1 uses the security facility to authenticate users in one of two ways:

- DB2 uses a successful security system login as evidence of identity, and allows:
 - Use of local commands to access local data
 - Use of remote connections when the server trusts the client authentication.
- DB2 accepts a user ID and password combination. It uses successful validation of this pair by the security facility as evidence of identity and allows:
 - Use of remote connections where the server requires proof of authentication
 - Use of operations where the user wants to run a command under an identity other than the identity used for login.

DB2 on AIX® can log failed password attempts with the operating system, and detect when a client has exceeded the number of allowable login tries, as specified by the LOGINRETRIES parameter.

Related concepts:

- “Authentication methods for your server” in *Administration Guide: Implementation*
- “Authorization” on page 21
- “Authorization, privileges, and object ownership” in *Administration Guide: Implementation*

Authorization

Authorization is the process whereby DB2 obtains information about an authenticated DB2 user, indicating the database operations that user may perform, and what data objects may be accessed. With each user request, there may be more than one authorization check, depending on the objects and operations involved.

Authorization is performed using DB2 facilities. DB2 tables and configuration files are used to record the permissions associated with each authorization name. When an authenticated user tries to access data, the authorization name of the user, and those of groups to which the user belongs, are compared with the recorded permissions. Based on this comparison, DB2 decides whether to allow the requested access.

There are three types of permissions recorded by DB2 Database for Linux, UNIX, and Windows: privileges, authority levels, and LBAC credentials.

A *privilege* defines a single permission for an authorization name, enabling a user to create or access database resources. Privileges are stored in the database catalogs.

Authority levels provide a method of grouping privileges and control over higher-level database manager maintenance and utility operations. Database-specific authorities are stored in the database catalogs; system authorities

are associated with group membership, and the group names that are associated with the authority levels are stored in the database manager configuration file for a given instance.

LBAC credentials are LBAC security labels and LBAC rule exemptions that allow access to data protected by label-based access control (LBAC). LBAC credentials are stored in the database catalogs.

Groups provide a convenient means of performing authorization for a collection of users without having to grant or revoke privileges for each user individually. Unless otherwise specified, group authorization names can be used anywhere that authorization names are used for authorization purposes. In general, group membership is considered for dynamic SQL and non-database object authorizations (such as instance level commands and utilities), but is not considered for static SQL. The exception to this general case occurs when privileges are granted to PUBLIC: these are considered when static SQL is processed. Specific cases where group membership does not apply are noted throughout the DB2 documentation, where applicable.

Related concepts:

- “Authorization and privileges” in *SQL Reference, Volume 1*
- “Authorization, privileges, and object ownership” in *Administration Guide: Implementation*
- “Label-based access control (LBAC) overview” in *Administration Guide: Implementation*

Units of work

A transaction is commonly referred to in DB2 Database for Linux, UNIX, and Windows as a *unit of work*. A unit of work is a recoverable sequence of operations within an application process. It is used by the database manager to ensure that a database is in a consistent state. Any reading from or writing to the database is done within a unit of work.

For example, a bank transaction might involve the transfer of funds from a savings account to a checking account. After the application subtracts an amount from the savings account, the two accounts are inconsistent, and remain so until the amount is added to the checking account. When *both* steps are completed, a point of consistency is reached. The changes can be committed and made available to other applications.

A unit of work is started implicitly when the first SQL statement is issued against the database. All subsequent reads and writes by the same application are considered part of the same unit of work. The application must end the unit of work by issuing either a COMMIT or a ROLLBACK statement. The COMMIT statement makes permanent all changes made within a unit of work. The ROLLBACK statement removes these changes from the database. If the application ends normally without either of these statements being explicitly issued, the unit of work is automatically committed. If it ends abnormally in the middle of a unit of work, the unit of work is automatically rolled back. Once issued, a COMMIT or a ROLLBACK cannot be stopped. With some multi-threaded applications, or some operating systems (such as Windows), if the application ends normally without either of these statements being explicitly issued, the unit of work is automatically rolled back. It is recommended that your applications always explicitly commit or roll back complete units of work. If part of a unit of work does not complete

successfully, the updates are rolled back, leaving the participating tables as they were before the transaction began. This ensures that requests are neither lost nor duplicated.

There is no physical representation of a unit of work because it is a series of instructions (SQL statements).

Related reference:

- “COMMIT statement” in *SQL Reference, Volume 2*
- “ROLLBACK statement” in *SQL Reference, Volume 2*

High availability disaster recovery (HADR) feature overview

DB2 Database for Linux, UNIX, and Windows high availability disaster recovery (HADR) is a database replication feature that provides a high availability solution for both partial and complete site failures. HADR protects against data loss by replicating data changes from a source database, called the *primary*, to a target database, called the *standby*.

A partial site failure can be caused by a hardware, network, or software (DB2 database or operating system) failure. Without HADR, the database management system (DBMS) server or the machine where the database resides has to be rebooted or restarted. This process could take several minutes to complete. With HADR, the standby database can take over the primary database role in a matter of seconds.

A complete site failure can occur when a disaster, such as a fire, causes the entire site to be destroyed. Since HADR uses TCP/IP for communication between the primary and standby databases, the two databases can be situated in different locations. For example, your primary database might be located at your head office in one city, while your standby database is located at your sales office in another city. If a disaster occurs at the primary site, data availability is maintained by having the remote standby database take over as the primary database with full DB2 functionality. After a takeover operation occurs, you can bring the original primary database back up and return it to its status of primary database; this is known as *failback*.

With HADR, you can choose the level of protection you want from potential loss of data by specifying one of three synchronization modes: synchronous (SYNC), near synchronous (NEARSYNC), and asynchronous (ASYNC). These modes indicate how data changes are propagated between the two systems. The synchronization mode selected will determine how close to being a replica the standby database will be when compared to the primary database. For example, using synchronous mode, HADR can guarantee that any transaction committed on the primary is also committed on the standby.

Synchronization allows you to have failover and failback between the two systems.

Data changes are recorded in database log records which are shipped from the primary system to the standby system. HADR is tightly-coupled with DB2 logging and recovery.

HADR requires that both systems have the same hardware, operating system, and DB2 software. (There may be some minor differences during times when the systems are being upgraded.)

The HADR standby database is established either by restoring it from a backup of the primary database, or by initializing it from a split-mirror copy of the primary database. Once HADR is started, the standby database will retrieve log records from the primary database and replay them against its own copy of the database. The log records are applied to the standby database until the standby database “catches up” to the in-memory log set of the primary database. At this point, the HADR pairing transitions to PEER state where the primary database sends new log pages to the standby database as well as writing the pages to its local disk. The log pages are replayed on the standby database as they arrive. Through continuous log replay, the standby database is maintained as a time-delayed replica of the primary database.

When a failure occurs on the primary database, you can then easily fail over to the standby database. Once you have failed over to the standby database, it becomes the new primary database. Because the standby database server is already online, failover can be completed very quickly. This keeps your time without database activity to a minimum.

HADR can also be used to maintain database availability across certain hardware or software release upgrades. You can upgrade your hardware, operating system, or DB2 FixPak level on the standby while the primary is available to applications. You can then transfer the applications to the upgraded system while the original primary is upgraded.

The performance of the new primary database immediately after failover may not be exactly the same as on the old primary database before the failure. The new primary database needs some time to populate the statement cache, the buffer pool, and other memory locations used by the database manager. Although the replaying of the log data from the old primary partly places data in the buffer pool and system catalog caches, it is not complete because it is only based on write activity. Frequently accessed index pages, catalog information for tables that is queried but not updated, statement caches, and access plans will all be missing from the caches. However, the whole process is faster than if you were starting up a new DB2 database.

Once the failed former primary server is repaired, it can be reintegrated as a standby database if the two copies of the database can be made consistent. After reintegration, a failback operation can be performed so that the original primary database is once again the primary database.

The HADR feature is available only on DB2 Enterprise Server Edition (ESE). It is disabled in other editions such as Personal Edition, and ESE with the database partitioning feature (DPF).

HADR takes place at the database level, not at the instance level. This means that a single instance could include the primary database (A), the standby database (B), and a standard (non-HADR) database (C). However, an instance cannot contain both the primary and standby for a single database because HADR requires that each copy of the database has the same database name.

Related concepts:

- “High availability” in *Data Recovery and High Availability Guide and Reference*

Developing a backup and recovery strategy

A database can become unusable because of hardware or software failure, or both. You might, at one time or another, encounter storage problems, power interruptions, or application failures, and each failure scenario requires a different recovery action. Protect your data against the possibility of loss by having a well rehearsed recovery strategy in place. Some of the questions that you should answer when developing your recovery strategy are:

- Will the database be recoverable?
- How much time can be spent recovering the database?
- How much time will pass between backup operations?
- How much storage space can be allocated for backup copies and archived logs?
- Will table space level backups be sufficient, or will full database backups be necessary?
- Should I configure a standby system, either manually or through high availability disaster recovery (HADR)?

A database recovery strategy should ensure that all information is available when it is required for database recovery. It should include a regular schedule for taking database backups and, in the case of partitioned database environments, include backups when the system is scaled (when database partition servers or nodes are added or dropped). Your overall strategy should also include procedures for recovering command scripts, applications, user-defined functions (UDFs), stored procedure code in operating system libraries, and load copies.

Different recovery methods are discussed in the sections that follow, and you will discover which recovery method is best suited to your business environment.

The concept of a database *backup* is the same as any other data backup: taking a copy of the data and then storing it on a different medium in case of failure or damage to the original. The simplest case of a backup involves shutting down the database to ensure that no further transactions occur, and then simply backing it up. You can then recreate the database if it becomes damaged or corrupted in some way.

The recreation of the database is called *recovery*. *Version recovery* is the restoration of a previous version of the database, using an image that was created during a backup operation. *Rollforward recovery* is the reapplication of transactions recorded in the database log files after a database or a table space backup image has been restored.

Crash recovery is the automatic recovery of the database if a failure occurs before all of the changes that are part of one or more units of work (transactions) are completed and committed. This is done by rolling back incomplete transactions and completing committed transactions that were still in memory when the crash occurred.

Recovery log files and the recovery history file are created automatically when a database is created (Figure 12 on page 26). These log files are important if you need to recover data that is lost or damaged.

Each database includes *recovery logs*, which are used to recover from application or system errors. In combination with the database backups, they are used to recover the consistency of the database right up to the point in time when the error occurred.

The *recovery history file* contains a summary of the backup information that can be used to determine recovery options, if all or part of the database must be recovered to a given point in time. It is used to track recovery-related events such as backup and restore operations, among others. This file is located in the database directory.

The *table space change history file*, which is also located in the database directory, contains information that can be used to determine which log files are required for the recovery of a particular table space.

You cannot directly modify the recovery history file or the table space change history file; however, you can delete entries from the files using the PRUNE HISTORY command. You can also use the *rec_his_retentn* database configuration parameter to specify the number of days that these history files will be retained.

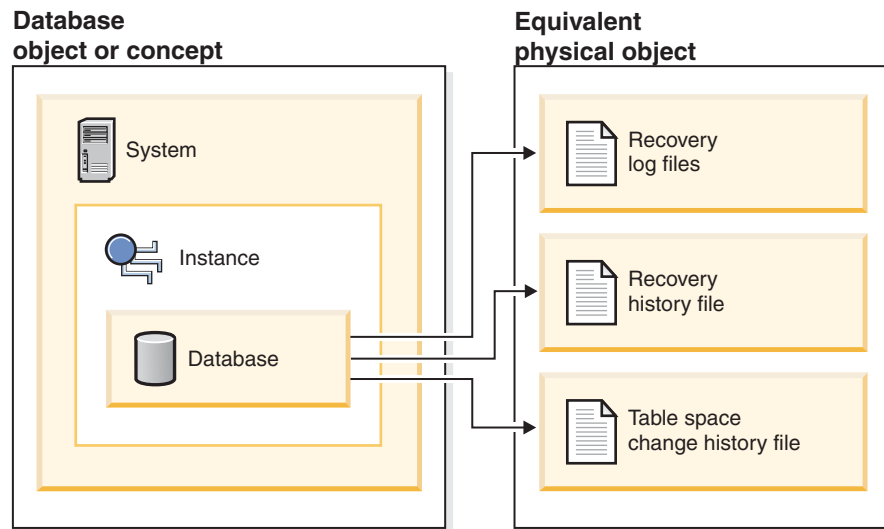


Figure 12. Database recovery files

Data that is easily recreated can be stored in a non-recoverable database. This includes data from an outside source that is used for read-only applications, and tables that are not often updated, for which the small amount of logging does not justify the added complexity of managing log files and rolling forward after a restore operation. If both the *logarchmeth1* and *logarchmeth2* database configuration parameters are set to "OFF" then the database is *Non-recoverable*. This means that the only logs that are kept are those required for crash recovery. These logs are known as *active logs*, and they contain current transaction data. Version recovery using *offline* backups is the primary means of recovery for a non-recoverable database. (An offline backup means that no other application can use the database when the backup operation is in progress.) Such a database can only be restored offline. It is restored to the state it was in when the backup image was taken and rollforward recovery is not supported.

Data that *cannot* be easily recreated should be stored in a recoverable database. This includes data whose source is destroyed after the data is loaded, data that is

manually entered into tables, and data that is modified by application programs or users after it is loaded into the database. *Recoverable databases* have the *logarchmeth1* or *logarchmeth2* database configuration parameters set to a value other than "OFF". Active logs are still available for crash recovery, but you also have the *archived logs*, which contain committed transaction data. Such a database can only be restored offline. It is restored to the state it was in when the backup image was taken. However, with rollforward recovery, you can roll the database *forward* (that is, past the time when the backup image was taken) by using the active and archived logs to either a specific point in time, or to the end of the active logs.

Recoverable database backup operations can be performed either offline or *online* (online meaning that other applications can connect to the database during the backup operation). Online table space restore and rollforward operations are supported only if the database is recoverable. If the database is non-recoverable, database restore and rollforward operations must be performed offline. During an online backup operation, rollforward recovery ensures that *all* table changes are captured and reapplied if that backup is restored.

If you have a recoverable database, you can back up, restore, and roll individual table spaces forward, rather than the entire database. When you back up a table space online, it is still available for use, and simultaneous updates are recorded in the logs. When you perform an online restore or rollforward operation on a table space, the table space itself is not available for use until the operation completes, but users are not prevented from accessing tables in other table spaces.

Automated backup operations:

Since it can be time-consuming to determine whether and when to run maintenance activities such as backup operations, you can use the Configure Automatic Maintenance wizard to do this for you. With automatic maintenance, you specify your maintenance objectives, including when automatic maintenance can run. DB2 then uses these objectives to determine if the maintenance activities need to be done and then runs only the required maintenance activities during the next available maintenance window (a user-defined time period for the running of automatic maintenance activities).

Note: You can still perform manual backup operations when automatic maintenance is configured. DB2 will only perform automatic backup operations if they are required.

Related concepts:

- "Crash recovery" in *Data Recovery and High Availability Guide and Reference*
- "High availability disaster recovery overview" in *Data Recovery and High Availability Guide and Reference*
- "Rollforward recovery" in *Data Recovery and High Availability Guide and Reference*
- "Version recovery" in *Data Recovery and High Availability Guide and Reference*

Related reference:

- "logarchmeth1 - Primary log archive method configuration parameter" in *Performance Guide*
- "rec_his_retentn - Recovery history retention period configuration parameter" in *Performance Guide*

Chapter 2. Automatic maintenance

About automatic maintenance

The DB2 product 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 to ensure that they are optimized for performance and recoverability. These activities include:

- Backup of the database. DB2 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 backup provides users with a solution to help ensure their database is being backed up both properly and regularly, without either having to worry about when to back up, or having any knowledge of the backup command.
- Data defragmentation (table or index reorganization). This maintenance activity can increase the efficiency with which the DB2 database manager accesses your tables. Automatic reorganization manages offline table and index reorganization without users having to worry about when and how to reorganize their data.
- Data access optimization (running statistics). The DB2 database manager updates the system catalog statistics on the data in a table, the data in a table's 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, and incorrectly specified statistics and by generating statistical profiles based on query feedback.

For users, it can be time-consuming to determine whether and when to run maintenance activities. Automatic maintenance takes the burden off of users. With automatic maintenance, you can specify your maintenance objectives, including when automatic maintenance can run. The DB2 database manager uses the objectives you have specified to determine if the maintenance activities need to be done and runs only the required maintenance activities during the next available maintenance window (a user-defined time period for the running of automatic maintenance activities).

Enablement of the automatic maintenance features is controlled using the automatic maintenance database configuration parameters. These are a hierarchical set of switches to allow for simplicity and flexibility in managing the enablement of these features. You can automate database maintenance activities to run only when they are needed using the Configure Automatic Maintenance wizard. The DB2 database manager uses the objectives you have specified using the Configure Automatic Maintenance wizard to determine whether the maintenance activities need to be done. Then the DB2 database manager runs only the required maintenance activities during the next available maintenance window. The maintenance window is a time period specified by you for the running of automatic maintenance activities.

Related concepts:

- “Maintenance windows” on page 35
- “Offline maintenance” on page 36
- “Online maintenance” on page 36

Automatic features enabled by default

DB2 includes several automatic features that are enabled by default when you create a database. These features are designed to assist you in managing your database system. This means that your system is capable of self-diagnosis and can anticipate problems before they happen by monitoring real-time data against historical problem data. In some cases, these tools can be configured to automatically make changes to your system so that service disruptions are avoided.

The following automatic features are enabled by default:

Automatic statistics collection

Automatic statistics collection helps to improve database performance by collecting up-to-date table statistics. DB2 determines which statistics are required by your workload and which statistics need to be updated. Then, the RUNSTATS utility is automatically invoked in the background to ensure the correct statistics are collected and maintained. Statistics are first collected on the tables that need it the most. The DB2 optimizer can then choose an access plan based on accurate statistics. You can disable automatic statistics collection after a database is created by setting the database configuration parameter AUTO_RUNSTATS 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 DB2 will place your table space data. Then, DB2 will manage the container and space allocation for the table spaces as they are created and populated.

Configuration Advisor

When you create a database, this tool is automatically invoked 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 will have better performance than a database created with the DB2 default values. It also means that you will spend less time tuning your system after the database has been created. The Configuration Advisor can be invoked at any time (even after your databases are populated) to recommend and optionally apply a set of configuration parameters to optimize DB2 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, DB2 will let you know about it and will also advise 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.

Self tuning memory (single partition databases only)

This feature simplifies the task of memory configuration by automatically adjusting the values for several memory configuration parameters based on the memory requirements of the system's workload. The memory tuner dynamically distributes available memory resources among several memory consumers including sort, the package cache, the lock list, and buffer pools. The memory tuner is responsive to significant changes in workload characteristics and iteratively adjusts the values of the memory configuration parameters and the sizes of the buffer pools to optimize performance. You can disable self tuning memory after a database is created by setting the database configuration parameter `SELF_TUNING_MEM` to `OFF`.

Utility throttling

This feature regulates the performance impact of maintenance utilities, so that they can run concurrently during production periods. While the impact policy for throttled utilities is defined by default, you must set the impact priority when a utility is invoked for it to run throttled. The throttling system will ensure the throttled utilities are run as aggressively as possible without violating the impact policy. Currently, you can throttle statistics collection, backup operations, and rebalancing operations.

Related concepts:

- "Introduction to the health monitor" in *System Monitor Guide and Reference*
- "Automatic storage databases" in *Administration Guide: Implementation*
- "Quick-start tips for performance tuning" in *Performance Guide*
- "Self tuning memory" in *Performance Guide*
- "Automatic statistics collection" in *Performance Guide*

Related reference:

- "db2AutoConfig API - Access the Configuration Advisor" in *Administrative API Reference*
- "SET UTIL_IMPACT_PRIORITY command" in *Command Reference*
- "util_impact_lim - Instance impact policy configuration parameter" in *Performance Guide*
- "auto_maint - Automatic maintenance configuration parameter" in *Performance Guide*

Automatic database backup

A database may become unusable due to a wide variety of hardware or software failures. Automatic database backup simplifies database backup management tasks for the DBA by always ensuring that a recent full backup of the database is performed as needed. It determines the need to perform a backup operation based on one or more of the following measures:

- You have never completed a full database backup
- The time elapsed since the last full backup is more than a specified number of hours
- The transaction log space consumed since the last backup is more than a specified number of 4 KB pages (in archive logging mode only).

Protect your data by planning and implementing a disaster recovery strategy for your system. If suitable to your needs, you may incorporate the automatic database backup feature as part of your backup and recovery strategy.

If the database is enabled for roll-forward recovery (archive logging), then automatic database backup can be enabled for either online or offline backup. Otherwise, only offline backup is available. Automatic database backup supports disk, tape, Tivoli® Storage Manager (TSM), and vendor DLL media types.

Through the Configure Automatic Maintenance wizard in the Control Center or Health Center, you can configure:

- The requested time or number of log pages between backups
- The backup media
- Whether it will be an online or offline backup.

If backup to disk is selected, the automatic backup feature will regularly delete backup images from the directory specified in the Configure Automatic Maintenance wizard. Only the most recent backup image is guaranteed to be available at any given time. It is recommended that this directory be kept exclusively for the automatic backup feature and not be used to store other backup images.

The automatic database backup feature can be enabled or disabled by using the **auto_db_backup** and **auto_maint** database configuration parameters. In a partitioned database environment, the automatic database backup runs on each database partition if the database configuration parameters are enabled on that database partition.

Related concepts:

- “Developing a backup and recovery strategy” on page 25
- “Automatic statistics collection” in *Performance Guide*
- “Catalog statistics” in *Performance Guide*
- “Table and index management for MDC tables” in *Performance Guide*
- “Table and index management for standard tables” in *Performance Guide*
- “Table reorganization” in *Performance Guide*
- “Health monitor” in *System Monitor Guide and Reference*

Related reference:

- “auto_maint - Automatic maintenance configuration parameter” in *Performance Guide*

Automatic reorganization

After many changes to table data, logically sequential data may be on non-sequential physical pages so the database manager has to perform additional read operations to access data.

Among other information, the statistical information collected by RUNSTATS shows the data distribution within a table. In particular, analysis of these statistics can indicate when and what kind of reorganization is necessary. Automatic reorganization determines the need for reorganization on tables by using the REORGCHK formulas. It periodically evaluates tables that have had their statistics updated to see if reorganization is required. If so, it internally schedules a classic

reorganization for the table. This requires that your applications function without write access to the tables being reorganized.

The automatic reorganization feature can be enabled or disabled by using the **auto_reorg**, **auto_tbl_maint**, and **auto_maint** database configuration parameters.

In a partitioned database environment, the determination to carry out automatic reorganization and the initiation of automatic reorganization, is done on the catalog partition. The database configuration parameters need to be enabled on the catalog partition only. The reorganization runs on all of the database partitions on which the target tables reside.

If you are unsure about when and how to reorganize your tables and indexes, you can incorporate automatic reorganization as part of your overall database maintenance plan.

Tables considered for automatic reorganization are configurable by you using the Automatic Maintenance wizard from the Control Center or Health Center.

Related concepts:

- “Table reorganization” in *Performance Guide*

Related tasks:

- “Choosing a table reorganization method” in *Performance Guide*
- “Determining when to reorganize tables” in *Performance Guide*
- “Enabling automatic table and index reorganization” in *Performance Guide*

Automatic statistics collection by table

When the SQL compiler optimizes SQL query plans, its decisions are heavily influenced by statistical information about the size of the database tables and indexes. The optimizer also uses information about the distribution of data in specific columns of tables and indexes if these columns are used to select rows or join tables. The optimizer uses this information to estimate the costs of alternative access plans for each query. Having out-of-date or incomplete statistics for a table or index could lead the optimizer to select a plan that is much more inefficient than other alternatives, slowing down query execution. In addition, deciding which statistics to collect for a given workload is complex, and keeping the statistics up-to-date through running the RUNSTATS utility can be time-consuming.

Once enabled, automatic statistics collection works in the background by determining the minimum set of statistics that give optimal performance improvement. The decision to collect or update statistics is taken by observing and learning how often tables are modified and how much the table statistics have changed. The automatic statistics collection algorithm learns over time how fast the statistics change on a per table basis and internally schedules RUNSTATS execution accordingly.

Normal database maintenance activities such as when you might use the RUNSTATS utility, the REORG utility, or altering or dropping a table, are not affected by the enablement of this feature.

If you are unsure about how often to collect statistics for the tables in your database, you may incorporate the automatic statistics collection feature as part of your overall database maintenance plan.

The automatic statistics collection feature can be enabled or disabled by using the **auto_runstats**, **auto_tbl_maint**, and **auto_maint** database configuration parameters. Or, you can use the Configure Automatic Maintenance wizard from the Control Center or Health Center to enable automatic statistics collection.

In a partitioned database environment, the determination to carry out automatic statistics collection and the initiation of automatic statistics collection, is done on the catalog partition. The **auto_runstats** configuration parameter needs to be enabled on the catalog partition only. The actual statistics collection is done by RUNSTATS and is collected as follows:

1. If the catalog partition has table data, then collect statistics on the catalog partition. RUNSTATS always collects statistics on the database partition where it is initiated if that database partition contains table data.
2. Otherwise, collection of statistics is done on the first database partition in the database partition list.

Related concepts:

- “Catalog statistics” in *Performance Guide*

Related tasks:

- “Collecting catalog statistics” in *Performance Guide*

Automatic statistics profiling using automatic statistics collection

Missing or outdated statistics can make the optimizer pick a slower query plan. It is important to note that not all statistics are important for a given workload. For example, statistics on columns not appearing in any query predicate are unlikely to have any impact. Sometimes statistics on several columns (column group statistics) are needed in order to adjust for correlations between these columns.

Automatic statistics profiling analyzes optimizer behavior by only considering columns that were used in previous queries and also knowing columns or column combinations where estimation errors occurred. In order to detect errors and recommend or change a statistical profile, the statistical profile generator mines information collected when the query is compiled as well as information accumulated when the query ran. This approach is reactive as the action is taken after the query has been seen and eventually after a plan has been chosen and run.

Automatic statistics profiling advises on how to collect statistics using the RUNSTATS utility by detecting outdated, missing, and incorrectly specified statistics and generating statistical profiles based on query feedback.

If suitable to your needs, you may incorporate the automatic statistics profiling feature as part of your overall database maintenance plan.

Automatic statistics profiling interacts with automatic statistics collection and advises on when to collect statistics.

The automatic statistics profiling feature can be enabled or disabled by using the **auto_stats_prof**, **auto_tbl_maint** and **auto_maint** database configuration

parameters. If the **auto_prof_upd** database configuration parameter is also enabled, then the statistical profiles generated are used to update the RUNSTATS user profiles. Automatic statistics profiling is not available for partitioned database environments or when symmetric multi-processor (SMP) parallelism, also called intrapartition parallelism, is enabled.

Related concepts:

- “Catalog statistics” in *Performance Guide*

Related tasks:

- “Collecting catalog statistics” in *Performance Guide*

Storage used by automatic statistics collection and profiling

The automatic statistics collection and reorganization features store working data in tables in your database. These tables are created in the SYSTOOLSPACE table space. The SYSTOOLSPACE table space is created automatically with default options when the database is activated. Storage requirements for these tables are proportional to the number of tables in the database and should be calculated as approximately 1 KB per table. If this is a significant size for your database, you may want to drop and re-create the table space yourself and allocate storage appropriately. The automatic maintenance and health monitor tables in the table space are automatically re-created. Any history captured in those tables is lost when the table space is dropped.

Maintenance windows

A maintenance window is a user-defined time period for the running of automatic maintenance activities. This is different than a task schedule. When a maintenance window occurs, each automatic maintenance activity is not necessarily run. Instead, the DB2 database manager evaluates the system based on the need for each maintenance activity to be run. If the maintenance requirements are not met, then the maintenance activity is run. If the database is already well maintained, the maintenance activity is not run.

You may need to think about when you would like the automatic maintenance activities to be run. The automatic maintenance activities (backup, statistics collection, statistics profiling, and reorganization) consume resources on your system and may affect the performance of your database when they are run. Automatic reorganization and offline database backup also restrict access to the tables and database when these utilities are run. It is therefore necessary to provide appropriate periods of time when these maintenance activities can be internally scheduled to be run by the DB2 database manager. These can be specified as offline and online maintenance time periods using the automatic maintenance wizard from the Control Center or Health Center.

Offline database backups and table and index reorganization are run in the offline maintenance time period. These features run to completion even if they go beyond the time period specified. The internal scheduling mechanism learns over time and estimates job completion times. If the offline time period is too small for a particular database backup or reorganization activity, the scheduler will not start the job the next time around and relies on the health monitor to provide notification of the need to increase the offline maintenance time period.

Automatic statistics collection and profiling as well as online database backups are run in the online maintenance time period. To minimize the impact on the system, they are throttled by the adaptive utility throttling mechanism. The internal scheduling mechanism uses the online maintenance time period to start the online jobs. These features run to completion even if they go beyond the time period specified.

Related concepts:

- “About automatic maintenance” on page 29
- “Offline maintenance” on page 36
- “Online maintenance” on page 36

Offline maintenance

Offline maintenance activities are maintenance activities that can occur only when there is some interruption of user access to the database. 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 will be forced off.
- During an offline data defragmentation (table or index reorganization), applications can access the data in tables in the database but cannot make updates.

Note: Data access optimization maintenance activities (running statistics) can only be performed online.

Related concepts:

- “About automatic maintenance” on page 29
- “Maintenance windows” on page 35
- “Online maintenance” on page 36

Online maintenance

Online maintenance activities are maintenance activities that can occur while users are connected to the database. When online maintenance activities run, any currently connected applications are allowed to remain connected, and new connections can be established.

Related concepts:

- “About automatic maintenance” on page 29
- “Maintenance windows” on page 35
- “Offline maintenance” on page 36

Chapter 3. Parallel database systems

This chapter discusses different ways of dividing and retrieving data to improve the speed of data access and the resulting response times for applications. Information about data distribution, table partitioning, parallelism and the use of single and multiple processors is included.

Parallelism

Components of a task, such as a database query, can be run in parallel to dramatically enhance performance. The nature of the task, the database configuration, and the hardware environment, all determine how the DB2 database product will perform a task in parallel. These considerations are interrelated, and should be considered together when you work on the physical and logical design of a database. The following types of parallelism are supported by the DB2 database system:

- I/O
- Query
- Utility

Input/output parallelism

When there are multiple containers for a table space, the database manager can exploit *parallel I/O*. Parallel I/O refers to the process of writing to, or reading from, two or more I/O devices simultaneously; it can result in significant improvements in throughput.

Query parallelism

There are two types of query parallelism: interquery parallelism and intraquery parallelism.

Interquery parallelism refers to the ability of the database to accept queries from multiple applications at the same time. Each query runs independently of the others, but DB2 runs all of them at the same time. DB2 database products have always supported this type of parallelism.

Intraquery parallelism refers to the simultaneous processing of parts of a single query, using either *intrapartition parallelism*, *interpartition parallelism*, or both.

Intrapartition parallelism

Intrapartition parallelism refers to the ability to break up a query into multiple parts. Some DB2 utilities also perform this type of parallelism.

Intrapartition parallelism subdivides what is usually considered a single database operation such as index creation, database loading, or SQL queries into multiple parts, many or all of which can be run in parallel *within a single database partition*.

Figure 13 on page 38 shows a query that is broken into four pieces that can be run in parallel, with the results returned more quickly than if the query were run in serial fashion. The pieces are copies of each other. To utilize intrapartition parallelism, you must configure the database appropriately. You can choose the

degree of parallelism or let the system do it for you. The degree of parallelism represents the number of pieces of a query running in parallel.

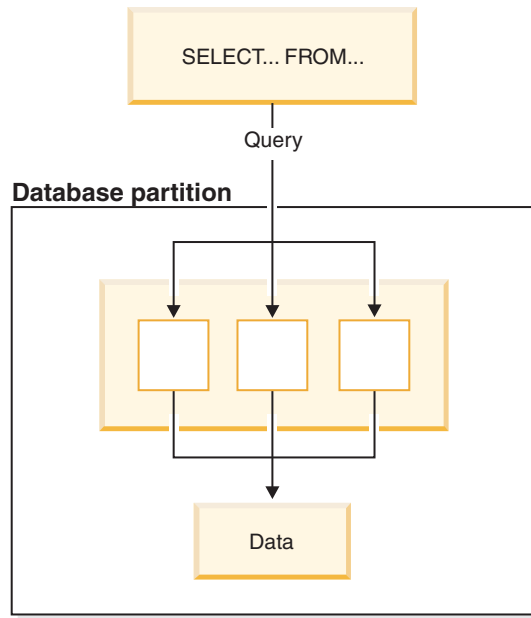


Figure 13. Intrapartition parallelism

Interpartition parallelism

Interpartition parallelism refers to the ability to break up a query into multiple parts across multiple partitions of a partitioned database, on one machine or multiple machines. The query is run in parallel. Some DB2 utilities also perform this type of parallelism.

Interpartition parallelism subdivides what is usually considered a single database operation such as index creation, database loading, or SQL queries into multiple parts, many or all of which can be run in parallel *across multiple partitions of a partitioned database on one machine or on multiple machines*.

Figure 14 on page 39 shows a query that is broken into four pieces that can be run in parallel, with the results returned more quickly than if the query were run in serial fashion on a single database partition.

The degree of parallelism is largely determined by the number of database partitions you create and how you define your database partition groups.

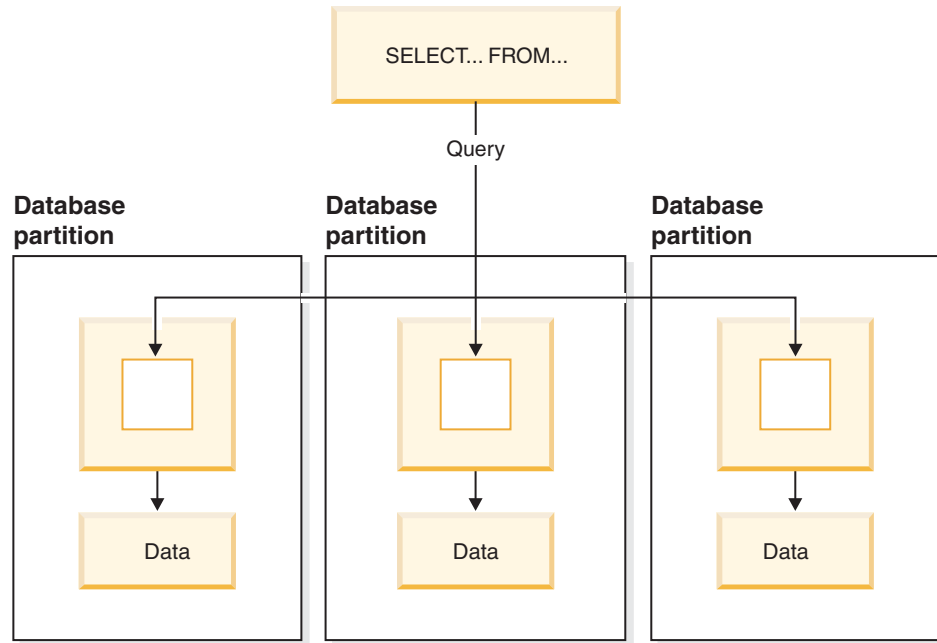


Figure 14. Interpartition parallelism

Simultaneous intrapartition and interpartition parallelism

You can use intrapartition parallelism and interpartition parallelism at the same time. This combination provides two dimensions of parallelism, resulting in an even more dramatic increase in the speed at which queries are processed.

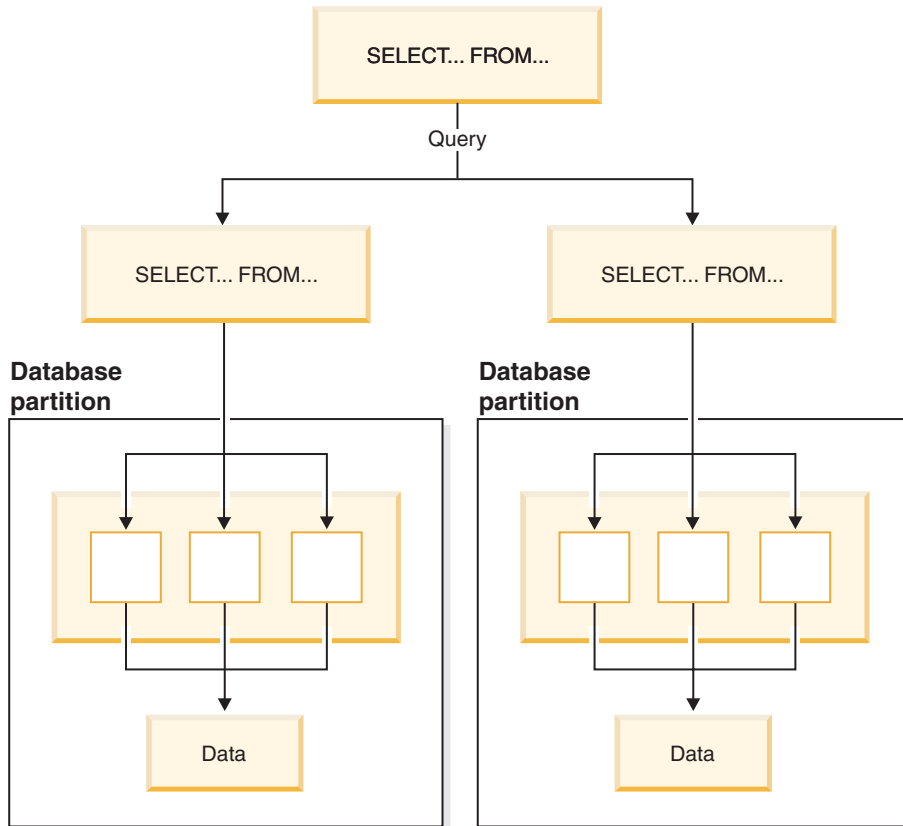


Figure 15. Simultaneous interpartition and intrapartition parallelism

Utility parallelism

DB2 utilities can take advantage of intrapartition parallelism. They can also take advantage of interpartition parallelism; where multiple database partitions exist, the utilities execute in each of the database partitions in parallel.

The load utility can take advantage of intrapartition parallelism and I/O parallelism. Loading data is a CPU-intensive task. The load utility takes advantage of multiple processors for tasks such as parsing and formatting data. It can also use parallel I/O servers to write the data to containers in parallel.

In a partitioned database environment, the LOAD command takes advantage of intrapartition, interpartition, and I/O parallelism by parallel invocations at each database partition where the table resides.

During index creation, the scanning and subsequent sorting of the data occurs in parallel. The DB2 system exploits both I/O parallelism and intrapartition parallelism when creating an index. This helps to speed up index creation when a CREATE INDEX statement is issued, during restart (if an index is marked invalid), and during the reorganization of data.

Backing up and restoring data are heavily I/O-bound tasks. The DB2 system exploits both I/O parallelism and intrapartition parallelism when performing

backup and restore operations. Backup exploits I/O parallelism by reading from multiple table space containers in parallel, and asynchronously writing to multiple backup media in parallel.

Related concepts:

- “Database partition and processor environments” on page 42

Partitioned database environments

DB2 Database for Linux, UNIX, and Windows extends the database manager to the parallel, multi-partition environment. 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.

A *single-partition database* is a database having only one database partition. All data in the database is stored in that single database partition. In this case database partition groups, while present, provide no additional capability.

A *multi-partition database* is a database with two or more database partitions. Tables can be located in one or more database partitions. When a table is in a database partition group consisting of multiple database partitions, some of its rows are stored in one database partition, and other rows are stored in other database partitions.

Usually, a single database partition exists on each physical machine, and the processors on each system are used by the database manager at each database partition to manage its part of the total data in the database.

Because data is distributed across database partitions, you can use the power of multiple processors on multiple physical machines to satisfy requests for information. Data retrieval and update requests are decomposed automatically into sub-requests, and executed in parallel among the applicable database partitions. The fact that databases are split across database partitions is transparent to users issuing SQL statements.

User interaction occurs through one database partition, known as the *coordinator partition* for that user. The coordinator partition runs on the same database partition as the application, or in the case of a remote application, the database partition to which that application is connected. Any database partition can be used as a coordinator partition.

DB2 allows you to store data across several database partitions in the database. This means that the data is physically stored across more than one database partition, and yet can be accessed as though it were located in the same place. Applications and users accessing data in a multi-partition database do not need to be aware of the physical location of the data.

The data, while physically split, is used and managed as a logical whole. Users can choose how to distribute their data by declaring distribution keys. Users can also determine across which and over how many database partitions their data is distributed by selecting the table space and the associated database partition group in which the data should be stored. Suggestions for distribution and replication can be done using the DB2 Design Advisor. In addition, an updatable distribution map

is used with a hashing algorithm to specify the mapping of distribution key values to database partitions, which determines the placement and retrieval of each row of data. As a result, you can spread the workload across a multi-partition database for large tables, while allowing smaller tables to be stored on one or more database partitions. Each database partition has local indexes on the data it stores, resulting in increased performance for local data access.

You are not restricted to having all tables divided across all database partitions in the database. DB2 supports *partial declustering*, which means that you can divide tables and their table spaces across a subset of database partitions in the system.

An alternative to consider when you want tables to be positioned on each database partition, is to use materialized query tables and then replicate those tables. You can create a materialized query table containing the information that you need, and then replicate it to each database partition.

Database partition and processor environments

This section provides an overview of the following hardware environments:

- Single database partition on a single processor (uniprocessor)
- Single database partition with multiple processors (SMP)
- Multiple database partition configurations
 - Database partitions with one processor (MPP)
 - Database partitions with multiple processors (cluster of SMPs)
 - Logical database partitions

Capacity and scalability are discussed for each environment. *Capacity* refers to the number of users and applications able to access the database. This is in large part determined by memory, agents, locks, I/O, and storage management. *Scalability* refers to the ability of a database to grow and continue to exhibit the same operating characteristics and response times.

Single database partition on a single processor

This environment is made up of memory and disk, but contains only a single CPU (see Figure 16 on page 43). It is referred to by many different names, including stand-alone database, client/server database, serial database, uniprocessor system, and single node or non-parallel environment.

The database in this environment serves the needs of a department or small office, where the data and system resources (including a single processor or CPU) are managed by a single database manager.




Figure 16. Single database partition on a single processor

Capacity and scalability

In this environment you can add more disks. Having one or more I/O servers for each disk allows for more than one I/O operation to take place at the same time.

A single-processor system is restricted by the amount of disk space the processor can handle. As workload increases, a single CPU may not be able to process user requests any faster, regardless of other components, such as memory or disk, that you may add. If you have reached maximum capacity or scalability, you can consider moving to a single database partition system with multiple processors.

Single database partition with multiple processors

This environment is typically made up of several equally powerful processors within the same machine (see Figure 17 on page 44), and is called a *symmetric multiprocessor (SMP)* system. Resources, such as disk space and memory, are *shared*.

With multiple processors available, different database operations can be completed more quickly. DB2 database systems can also divide the work of a single query among available processors to improve processing speed. Other database operations, such as loading data, backing up and restoring table spaces, and creating indexes on existing data, can take advantage of multiple processors.

Symmetric multiprocessor (SMP) environment




Figure 17. Single partition database symmetric multiprocessor environment

Capacity and scalability

In this environment you can add more processors. However, since the different processors may attempt to access the same data, limitations with this environment can appear as your business operations grow. With shared memory and shared disks, you are effectively sharing all of the database data.

You can increase the I/O capacity of the database partition associated with your processor by increasing the number of disks. You can establish I/O servers to specifically deal with I/O requests. Having one or more I/O servers for each disk allows for more than one I/O operation to take place at the same time.

If you have reached maximum capacity or scalability, you can consider moving to a system with multiple database partitions.

Multiple database partition configurations

You can divide a database into multiple database partitions, each on its own machine. Multiple machines with multiple database partitions can be grouped together. This section describes the following database partition configurations:

- Database partitions on systems with one processor
- Database partitions on systems with multiple processors
- Logical database partitions

Database partitions with one processor

In this environment, there are many database partitions. Each database partition resides on its own machine, and has its own processor, memory, and disks (Figure 18 on page 45). All the machines are connected by a communications facility. This environment is referred to by many different names, including: cluster, cluster of uniprocessors, massively parallel processing (MPP) environment, and shared-nothing configuration. The latter name accurately reflects the arrangement of resources in this environment. Unlike an SMP environment, an MPP

environment has no shared memory or disks. The MPP environment removes the limitations introduced through the sharing of memory and disks.

A partitioned database environment allows a database to remain a logical whole, despite being physically divided across more than one database partition. The fact that data is distributed remains transparent to most users. Work can be divided among the database managers; each database manager in each database partition works against its own part of the database.




Figure 18. Massively parallel processing (MPP) environment

Capacity and scalability: In this environment you can add more database partitions (nodes) to your configuration. On some platforms, for example the RS/6000® SP™, the maximum number is 512 nodes. However, there may be practical limits on managing a high number of machines and instances.

If you have reached maximum capacity or scalability, you can consider moving to a system where each database partition has multiple processors.

Database partitions with multiple processors

An alternative to a configuration in which each database partition has a single processor, is a configuration in which each database partition has multiple processors. This is known as an *SMP cluster* (Figure 19 on page 46).

This configuration combines the advantages of SMP and MPP parallelism. This means that a query can be performed in a single database partition across multiple processors. It also means that a query can be performed in parallel across multiple database partitions.

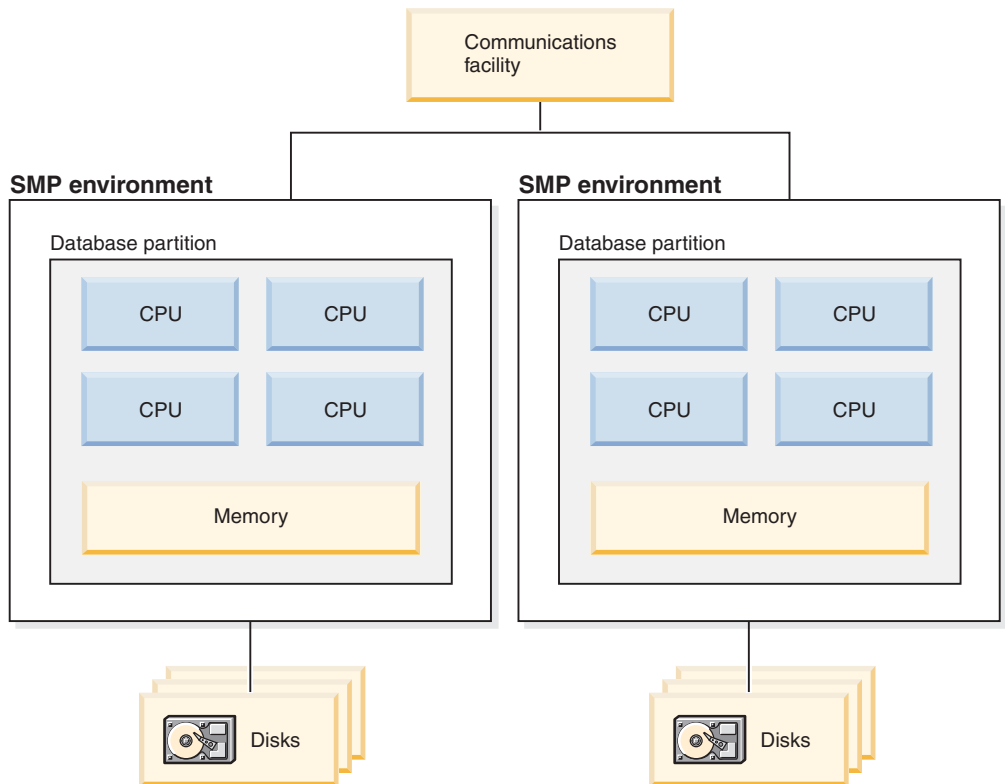


Figure 19. Several symmetric multiprocessor (SMP) environments in a cluster

Capacity and scalability: In this environment you can add more database partitions, and you can add more processors to existing database partitions.

Logical database partitions

A logical database partition differs from a physical partition in that it is not given control of an entire machine. Although the machine has shared resources, database partitions do not share the resources. Processors are shared but disks and memory are not.

Logical database partitions provide scalability. Multiple database managers running on multiple logical partitions may make fuller use of available resources than a single database manager could. Figure 20 on page 47 illustrates the fact that you may gain more scalability on an SMP machine by adding more database partitions; this is particularly true for machines with many processors. By distributing the database, you can administer and recover each database partition separately.

Big SMP environment




Figure 20. Partitioned database with symmetric multiprocessor environment

Figure 21 on page 48 illustrates that you can multiply the configuration shown in Figure 20 to increase processing power.

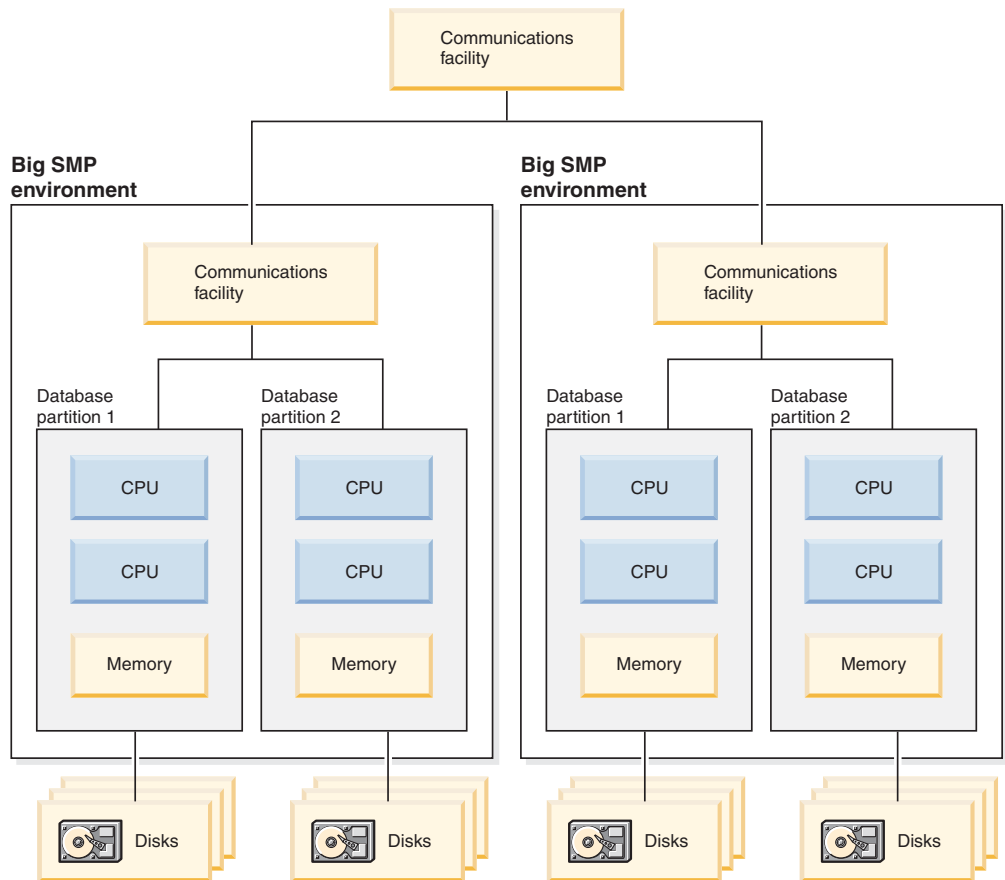


Figure 21. Partitioned database with symmetric multiprocessor environments clustered together

Note: The ability to have two or more database partitions coexist on the same machine (regardless of the number of processors) allows greater flexibility in designing high availability configurations and failover strategies. Upon machine failure, a database partition can be automatically moved and restarted on a second machine that already contains another database partition of the same database.

Summary of parallelism best suited to each hardware environment

The following table summarizes the types of parallelism best suited to take advantage of the various hardware environments.

Table 3. Types of Parallelism Possible in Each Hardware Environment

Hardware Environment	I/O Parallelism	Intra-Query Parallelism	
		Intra-Partition Parallelism	Inter-Partition Parallelism
Single Database Partition, Single Processor	Yes	No(1)	No
Single Database Partition, Multiple Processors (SMP)	Yes	Yes	No
Multiple Database Partitions, One Processor (MPP)	Yes	No(1)	Yes

Table 3. Types of Parallelism Possible in Each Hardware Environment (continued)

Hardware Environment	I/O Parallelism	Intra-Query Parallelism	
		Intra-Partition Parallelism	Inter-Partition Parallelism
Multiple Database Partitions, Multiple Processors (cluster of SMPs)	Yes	Yes	Yes
Logical Database Partitions	Yes	Yes	Yes
<p>Note: (1) There may be an advantage to setting the degree of parallelism (using one of the configuration parameters) to some value greater than one, even on a single processor system, especially if the queries you execute are not fully utilizing the CPU (for example, if they are I/O bound).</p>			

Related concepts:

- “Parallelism” on page 37

Part 2. Database design

Chapter 4. Logical database design

When designing a database, you want to create an accurate representation of your environment that will serve as a basis for expansion. In addition, your database design should maintain the consistency and integrity of your data. You can achieve this by creating a design that will reduce redundancy and eliminate anomalies that can occur when updating your database. The topics in this chapter will discuss the elements of logical database design.

What to record in a database

The first step in developing a database design is to identify the types of data to be stored in database tables. A database includes information about the *entities* in an organization or business, and their relationships to each other. In a relational database, *entities* are represented as *tables*.

An *entity* is a person, object, or concept about which you want to store information. Some of the entities described in the sample tables are employees, departments, and projects.

In the sample employee table, the entity "employee" has *attributes*, or properties, such as employee number, name, work department, and job description. Those properties appear as the *columns* EMPNO, FIRSTNME, LASTNAME, WORKDEPT, and JOB.

An *occurrence* of the entity "employee" consists of the values in all of the columns for one employee. Each employee has a unique employee number (EMPNO) that can be used to identify an occurrence of the entity "employee". Each row in a table represents an occurrence of an entity or relationship. For example, in the following table the values in the first row describe an employee named Haas.

Table 4. Occurrences of Employee Entities and their Attributes

EMPNO	FIRSTNME	LASTNAME	WORKDEPT	JOB
000010	Christine	Haas	A00	President
000020	Michael	Thompson	B01	Manager
000120	Sean	O'Connell	A00	Clerk
000130	Dolores	Quintana	C01	Analyst
000030	Sally	Kwan	C01	Manager
000140	Heather	Nicholls	C01	Analyst
000170	Masatoshi	Yoshimura	D11	Designer

There is a growing need to support non-traditional database applications such as multimedia. You may want to consider attributes to support multimedia objects such as documents, video or mixed media, image, and voice.

Within a table, each column of a row is related in some way to all the other columns of that row. Some of the relationships expressed in the sample tables are:

- Employees are assigned to departments

- Dolores Quintana is assigned to Department C01
- Employees perform a job
 - Dolores works as an Analyst
- Employees manage departments
 - Sally manages department C01.

"Employee" and "department" are entities; Sally Kwan is part of an occurrence of "employee," and C01 is part of an occurrence of "department". The same relationship applies to the same columns in every row of a table. For example, one row of a table expresses the relationship that Sally Kwan manages Department C01; another, the relationship that Sean O'Connell is a clerk in Department A00.

The information contained within a table depends on the relationships to be expressed, the amount of flexibility needed, and the data retrieval speed desired.

In addition to identifying the entity relationships within your enterprise, you also need to identify other types of information, such as the business rules that apply to that data.

Related concepts:

- "Column definitions" on page 56
- "Database relationships" on page 54

Database relationships

Several types of relationships can be defined in a database. Consider the possible relationships between employees and departments.

One-to-many and many-to-one relationships

An employee can work in only one department; this relationship is *single-valued* for employees. On the other hand, one department can have many employees; this relationship is *multi-valued* for departments. The relationship between employees (single-valued) and departments (multi-valued) is a *one-to-many* relationship.

To define tables for each one-to-many and each many-to-one relationship:

1. Group all the relationships for which the "many" side of the relationship is the same entity.
2. Define a single table for all the relationships in the group.

In the following example, the "many" side of the first and second relationships is "employees" so an employee table, EMPLOYEE, is defined.

Table 5. Many-to-One Relationships

Entity	Relationship	Entity
Employees	are assigned to	departments
Employees	work at	jobs
Departments	report to	(administrative) departments

In the third relationship, "departments" is on the "many" side, so a department table, DEPARTMENT, is defined.

The following tables show these different relationships.

The EMPLOYEE table:

EMPNO	WORKDEPT	JOB
000010	A00	President
000020	B01	Manager
000120	A00	Clerk
000130	C01	Analyst
000030	C01	Manager
000140	C01	Analyst
000170	D11	Designer

The DEPARTMENT table:

DEPTNO	ADMRDEPT
C01	A00
D01	A00
D11	D01

Many-to-many relationships

A relationship that is multi-valued in both directions is a many-to-many relationship. An employee can work on more than one project, and a project can have more than one employee. The questions "What does Dolores Quintana work on?", and "Who works on project IF1000?" both yield multiple answers. A many-to-many relationship can be expressed in a table with a column for each entity ("employees" and "projects"), as shown in the following example.

The following table shows how a many-to-many relationship (an employee can work on many projects, and a project can have many employees working on it) is represented.

The employee activity (EMP_ACT) table:

EMPNO	PROJNO
000030	IF1000
000030	IF2000
000130	IF1000
000140	IF2000
000250	AD3112

One-to-one relationships

One-to-one relationships are single-valued in both directions. A manager manages one department; a department has only one manager. The questions, "Who is the manager of Department C01?", and "What department does Sally Kwan manage?" both have single answers. The relationship can be assigned to either the DEPARTMENT table or the EMPLOYEE table. Because all departments have

managers, but not all employees are managers, it is most logical to add the manager to the DEPARTMENT table, as shown in the following example.

The following table shows the representation of a one-to-one relationship.

The DEPARTMENT table:

DEPTNO	MGRNO
A00	000010
B01	000020
D11	000060

Ensure that equal values represent the same entity

You can have more than one table describing the attributes of the same set of entities. For example, the EMPLOYEE table shows the number of the department to which an employee is assigned, and the DEPARTMENT table shows which manager is assigned to each department number. To retrieve both sets of attributes simultaneously, you can join the two tables on the matching columns, as shown in the following example. The values in WORKDEPT and DEPTNO represent the same entity, and represent a *join path* between the DEPARTMENT and EMPLOYEE tables.

The DEPARTMENT table:

DEPTNO	DEPTNAME	MGRNO	ADMRDEPT
D21	Administration Support	000070	D01

The EMPLOYEE table:

EMPNO	FIRSTNAME	LASTNAME	WORKDEPT	JOB
000250	Daniel	Smith	D21	Clerk

When you retrieve information about an entity from more than one table, ensure that equal values represent the same entity. The connecting columns can have different names (like WORKDEPT and DEPTNO in the previous example), or they can have the same name (like the columns called DEPTNO in the department and project tables).

Related concepts:

- “Column definitions” on page 56
- “What to record in a database” on page 53

Column definitions

Within a relational table, each row of data in the table is a collection of related data values. There are characteristics to each piece of data in each row. Columns are used to identify and classify each piece of data.

Each column in a table must have a name that is unique for that table.

The *data type* and *length* specify the type of data and the maximum length that are valid for the column. Data types may be chosen from those provided by the database manager or you may choose to create your own user-defined types.

Examples of data type categories are: numeric, character string, double-byte (or graphic) character string, date-time, and binary string.

Large object (LOB) data types support multi-media objects such as documents, video, image and voice. These objects are implemented using the following data types:

- A *binary large object* (BLOB) string. Examples of BLOBs are photographs of employees, voice, and video.
- A *character large object* (CLOB) string, where the sequence of characters can be either single- or multi-byte characters, or a combination of both. An example of a CLOB is an employee's resume.
- A *double-byte character large object* (DBCLOB) string, where the sequence of characters is double-byte characters. An example of a DBCLOB is a Japanese resume.

A *user-defined type* (UDT), is a data type that is derived from an existing type. You may need to define types that are derived from and share characteristics with existing data types, but that are nevertheless considered to be separate from them and incompatible with them.

A *structured type* is a user-defined type whose structure is defined in the database. It contains a sequence of named *attributes*, each of which has a data type. A structured type may be defined as a *subtype* of another structured type, called its *supertype*. A subtype inherits all the attributes of its supertype and may have additional attributes defined. The set of structured types that are related to a common supertype is called a *type hierarchy*, and the supertype that does not have any supertype is called the *root type* of the type hierarchy.

A structured type may be used as the type of a table or a view. The names and data types of the attributes of the structured types, together with the object identifier, become the names and data types of the columns of this *typed table* or *typed view*. Rows of the typed table or typed view can be thought of as a representation of instances of the structured type.

A structured type cannot be used as the data type of a column of a table or a view. There is also no support for retrieving a whole structured type instance into a host variable in an application program.

A *reference type* is a companion type to the structured type. Similar to a distinct type, a reference type is a scalar type that shares a common representation with one of the built-in data types. This same representation is shared for all types in the type hierarchy. The reference type representation is defined when the root type of a type hierarchy is created. When using a reference type, a structured type is specified as a parameter of the type. This parameter is called the *target type* of the reference.

The target of a reference is always a row in a typed table or view. When a reference type is used, it may have a *scope* defined. The scope identifies a table (called the *target table*) or view (called the *target view*) that contains the target row of a reference value. The target table or view must have the same type as the target

type of the reference type. An instance of a scoped reference type uniquely identifies a row in a typed table or typed view, called its *target row*.

A *user-defined function* (UDF) can be used for a number of reasons, including invoking routines that allow comparison or conversion between user-defined types. UDFs extend and add to the support provided by built-in SQL functions, and can be used wherever a built-in function can be used. There are two types of UDFs:

- An external function, which is written in a programming language
- A sourced function, which will be used to invoke other UDFs

For example, two numeric data types are European Shoe Size and American Shoe Size. Both types represent shoe size, but they are incompatible, because the measurement base is different and cannot be compared. A user-defined function can be invoked to convert one shoe size to another.

Some columns cannot have meaningful values in all rows because:

- A column value is not applicable to the row.
For example, a column containing an employee's middle initial is not applicable to an employee who has no middle initial.
- A value is applicable, but is not yet known.
For example, the MGRNO column might not contain a valid manager number because the previous manager of the department has been transferred, and a new manager has not been appointed yet.

In both situations, you can choose between allowing a NULL value (a special value indicating that the column value is unknown or not applicable), or allowing a non-NULL default value to be assigned by the database manager or by the application.

Primary keys

A *key* is a set of columns that can be used to identify or access a particular row or rows. The key is identified in the description of a table, index, or referential constraint. The same column can be part of more than one key.

A *unique key* is a key that is constrained so that no two of its values are equal. The columns of a unique key cannot contain NULL values. For example, an employee number column can be defined as a unique key, because each value in the column identifies only one employee. No two employees can have the same employee number.

The mechanism used to enforce the uniqueness of the key is called a *unique index*. The unique index of a table is a column, or an ordered collection of columns, for which each value identifies (functionally determines) a unique row. A unique index can contain NULL values.

The *primary key* is one of the unique keys defined on a table, but is selected to be the key of first importance. There can be only one primary key on a table.

A *primary index* is automatically created for the primary key. The primary index is used by the database manager for efficient access to table rows, and it allows the database manager to enforce the uniqueness of the primary key. (You can also define indexes on non-primary key columns to efficiently access data when processing queries.)

If a table does not have a "natural" unique key, or if arrival sequence is the method used to distinguish unique rows, using a time stamp as part of the key can be helpful.

Primary keys for some of the sample tables are:

Table	Key Column
Employee table	EMPNO
Department table	DEPTNO
Project table	PROJNO

The following example shows part of the PROJECT table, including its primary key column.

Table 6. A Primary Key on the PROJECT Table

PROJNO (Primary Key)	PROJNAME	DEPTNO
MA2100	Weld Line Automation	D01
MA2110	Weld Line Programming	D11

If every column in a table contains duplicate values, you cannot define a primary key with only one column. A key with more than one column is a *composite key*. The combination of column values should define a unique entity. If a composite key cannot be defined easily, you may consider creating a new column that has unique values.

The following example shows a primary key containing more than one column (a composite key):

Table 7. A Composite Primary Key on the EMP_ACT Table

EMPNO (Primary Key)	PROJNO (Primary Key)	ACTNO (Primary Key)	EMPTIME	EMSTDATE (Primary Key)
000250	AD3112	60	1.0	1982-01-01
000250	AD3112	60	.5	1982-02-01
000250	AD3112	70	.5	1982-02-01

Identifying candidate key columns

To identify candidate keys, select the smallest number of columns that define a unique entity. There may be more than one candidate key. In Table 8, there appear to be many candidate keys. The EMPNO, the PHONENO, and the LASTNAME columns each uniquely identify the employee.

Table 8. EMPLOYEE Table

EMPNO (Primary Key)	FIRSTNAME	LASTNAME	WORKDEPT (Foreign Key)	PHONENO
000010	Christine	Haas	A00	3978
000030	Sally	Kwan	C01	4738
000060	Irving	Stern	D11	6423
000120	Sean	O'Connell	A00	2167
000140	Heather	Nicholls	C01	1793

Table 8. EMPLOYEE Table (continued)

EMPNO (Primary Key)	FIRSTNAME	LASTNAME	WORKDEPT (Foreign Key)	PHONENO
000170	Masatoshi	Yoshimura	D11	2890

The criteria for selecting a primary key from a pool of candidate keys should be persistence, uniqueness, and stability:

- Persistence means that a primary key value for each row always exists.
- Uniqueness means that the key value for each row is different from all the others.
- Stability means that primary key values never change.

Of the three candidate keys in the example, only EMPNO satisfies all of these criteria. An employee may not have a phone number when joining a company. Last names can change, and, although they may be unique at one point, are not guaranteed to be so. The employee number column is the best choice for the primary key. An employee is assigned a unique number only once, and that number is generally not updated as long as the employee remains with the company. Since each employee must have a number, values in the employee number column are persistent.

Related concepts:

- “Identity columns” on page 60

Identity columns

An *identity column* provides a way for DB2 Database for Linux, UNIX, and Windows to automatically generate a unique numeric value for each row in a table. A table can have a single column that is defined with the identity attribute. Examples of an identity column include order number, employee number, stock number, and incident number.

Values for an identity column can be generated always or generated by default.

- An identity column that is defined as *generated always* prevents the overriding of values in an SQL statement. Its values are always generated by DB2 database manager; applications are not allowed to provide an explicit value. There is no guarantee on the uniqueness of values found within generated always columns. To guarantee uniqueness of values in the column, a unique index should be defined on the column.
- An identity column that is defined as *generated by default* gives applications a way to explicitly provide a value for the identity column. If a value is not given, DB2 generates one, but cannot guarantee the uniqueness of the value in this case. Generated by default is meant to be used for data propagation, in which the contents of an existing table are copied, or for the unloading and reloading of a table.

Identity columns are ideally suited to the task of generating unique primary key values. Applications can use identity columns to avoid the concurrency and performance problems that can result when an application generates its own unique counter outside of the database. For example, one common application-level implementation is to maintain a 1-row table containing a counter. Each transaction locks this table, increments the number, and then commits; that is,

only one transaction at a time can increment the counter. In contrast, if the counter is maintained through an identity column, much higher levels of concurrency can be achieved because the counter is not locked by transactions. One uncommitted transaction that has incremented the counter will not prevent subsequent transactions from also incrementing the counter.

The counter for the identity column is incremented (or decremented) independently of the transaction. If a given transaction increments an identity counter two times, that transaction may see a gap in the two numbers that are generated because there may be other transactions concurrently incrementing the same identity counter (that is, inserting rows into the same table). If an application must have a consecutive range of numbers, that application should take an exclusive lock on the table that has the identity column. This decision must be weighed against the resulting loss of concurrency. Furthermore, it is possible that a given identity column can appear to have generated gaps in the number, because a transaction that generated a value for the identity column has rolled back, or the database that has cached a range of values has been deactivated before all of the cached values were assigned.

The sequential numbers that are generated by the identity column have the following additional properties:

- The values can be of any exact numeric data type with a scale of zero; that is, SMALLINT, INTEGER, BIGINT, or DECIMAL with a scale of zero. (Single- and double-precision floating-point are considered to be approximate numeric data types.)
- Consecutive values can differ by any specified integer increment. The default increment is 1.
- The counter value for the identity column is recoverable. If a failure occurs, the counter value is reconstructed from the logs, thereby guaranteeing that unique values continue to be generated.
- Identity column values can be cached to give better performance.

Related concepts:

- “Primary keys” on page 58

Normalization

Normalization helps eliminate redundancies and inconsistencies in table data. It is the process of reducing tables to a set of columns where all the non-key columns depend on the primary key column. If this is not the case, the data can become inconsistent during updates.

This section briefly reviews the rules for first, second, third, and fourth normal form. The fifth normal form of a table, which is covered in many books on database design, is not described here.

Form Description

First At each row and column position in the table, there exists *one* value, never a set of values.

Second Each column that is not part of the key is dependent upon the key.

Third Each non-key column is independent of other non-key columns, and is dependent only upon the key.

Fourth No row contains two or more independent multi-valued facts about an entity.

First normal form

A table is in *first normal form* if there is only one value, never a set of values, in each cell. A table that is in first normal form does not necessarily satisfy the criteria for higher normal forms.

For example, the following table violates first normal form because the WAREHOUSE column contains several values for each occurrence of PART.

Table 9. Table Violating First Normal Form

PART (Primary Key)	WAREHOUSE
P0010	Warehouse A, Warehouse B, Warehouse C
P0020	Warehouse B, Warehouse D

The following example shows the same table in first normal form.

Table 10. Table Conforming to First Normal Form

PART (Primary Key)	WAREHOUSE (Primary Key)	QUANTITY
P0010	Warehouse A	400
P0010	Warehouse B	543
P0010	Warehouse C	329
P0020	Warehouse B	200
P0020	Warehouse D	278

Second normal form

A table is in *second normal form* if each column that is not part of the key is dependent upon the *entire* key.

Second normal form is violated when a non-key column is dependent upon *part* of a composite key, as in the following example:

Table 11. Table Violating Second Normal Form

PART (Primary Key)	WAREHOUSE (Primary Key)	QUANTITY	WAREHOUSE_ADDRESS
P0010	Warehouse A	400	1608 New Field Road
P0010	Warehouse B	543	4141 Greenway Drive
P0010	Warehouse C	329	171 Pine Lane
P0020	Warehouse B	200	4141 Greenway Drive
P0020	Warehouse D	278	800 Massey Street

The primary key is a composite key, consisting of the PART and the WAREHOUSE columns together. Because the WAREHOUSE_ADDRESS column depends only on the value of WAREHOUSE, the table violates the rule for second normal form.

The problems with this design are:

- The warehouse address is repeated in every record for a part stored in that warehouse.
- If the address of a warehouse changes, every row referring to a part stored in that warehouse must be updated.
- Because of this redundancy, the data might become inconsistent, with different records showing different addresses for the same warehouse.
- If at some time there are no parts stored in a warehouse, there might not be a row in which to record the warehouse address.

The solution is to split the table into the following two tables:

Table 12. PART_STOCK Table Conforming to Second Normal Form

PART (Primary Key)	WAREHOUSE (Primary Key)	QUANTITY
P0010	Warehouse A	400
P0010	Warehouse B	543
P0010	Warehouse C	329
P0020	Warehouse B	200
P0020	Warehouse D	278

Table 13. WAREHOUSE Table Conforms to Second Normal Form

WAREHOUSE (Primary Key)	WAREHOUSE_ADDRESS
Warehouse A	1608 New Field Road
Warehouse B	4141 Greenway Drive
Warehouse C	171 Pine Lane
Warehouse D	800 Massey Street

There is a performance consideration in having the two tables in second normal form. Applications that produce reports on the location of parts must join both tables to retrieve the relevant information.

Third normal form

A table is in *third normal form* if each non-key column is independent of other non-key columns, and is dependent only on the key.

The first table in the following example contains the columns EMPNO and WORKDEPT. Suppose a column DEPTNAME is added (see Table 15 on page 64). The new column depends on WORKDEPT, but the primary key is EMPNO. The table now violates third normal form. Changing DEPTNAME for a single employee, John Parker, does not change the department name for other employees in that department. There are now two different department names used for department number E11. The inconsistency that results is shown in the updated version of the table.

Table 14. Unnormalized EMPLOYEE_DEPARTMENT Table Before Update

EMPNO (Primary Key)	FIRSTNAME	LASTNAME	WORKDEPT	DEPTNAME
000290	John	Parker	E11	Operations
000320	Ramlal	Mehta	E21	Software Support
000310	Maude	Setright	E11	Operations

Table 15. Unnormalized EMPLOYEE_DEPARTMENT Table After Update. Information in the table has become inconsistent.

EMPNO (Primary Key)	FIRSTNAME	LASTNAME	WORKDEPT	DEPTNAME
000290	John	Parker	E11	Installation Mgmt
000320	Ramlal	Mehta	E21	Software Support
000310	Maude	Setright	E11	Operations

The table can be normalized by creating a new table, with columns for WORKDEPT and DEPTNAME. An update like changing a department name is now much easier; only the new table needs to be updated.

An SQL query that returns the department name along with the employee name is more complex to write, because it requires joining the two tables. It will probably also take longer to run than a query on a single table. Additional storage space is required, because the WORKDEPT column must appear in both tables.

The following tables are defined as a result of normalization:

Table 16. EMPLOYEE Table After Normalizing the EMPLOYEE_DEPARTMENT Table

EMPNO (Primary Key)	FIRSTNAME	LASTNAME	WORKDEPT
000290	John	Parker	E11
000320	Ramlal	Mehta	E21
000310	Maude	Setright	E11

Table 17. DEPARTMENT Table After Normalizing the EMPLOYEE_DEPARTMENT Table

DEPTNO (Primary Key)	DEPTNAME
E11	Operations
E21	Software Support

Fourth normal form

A table is in fourth normal form if no row contains two or more independent multi-valued facts about an entity.

Consider these entities: employees, skills, and languages. An employee can have several skills and know several languages. There are two relationships, one between employees and skills, and one between employees and languages. A table is not in fourth normal form if it represents both relationships, as in the following

example:

Table 18. Table Violating Fourth Normal Form

EMPNO (Primary Key)	SKILL (Primary Key)	LANGUAGE (Primary Key)
000130	Data Modelling	English
000130	Database Design	English
000130	Application Design	English
000130	Data Modelling	Spanish
000130	Database Design	Spanish
000130	Application Design	Spanish

Instead, the relationships should be represented in two tables:

Table 19. EMPLOYEE_SKILL Table Conforming to Fourth Normal Form

EMPNO (Primary Key)	SKILL (Primary Key)
000130	Data Modelling
000130	Database Design
000130	Application Design

Table 20. EMPLOYEE_LANGUAGE Table Conforming to Fourth Normal Form

EMPNO (Primary Key)	LANGUAGE (Primary Key)
000130	English
000130	Spanish

If, however, the attributes are interdependent (that is, the employee applies certain languages only to certain skills), the table should *not* be split.

A good strategy when designing a database is to arrange all data in tables that are in fourth normal form, and then to decide whether the results give you an acceptable level of performance. If they do not, you can rearrange the data in tables that are in third normal form, and then reassess performance.

Constraints

A *constraint* is a rule that the database manager enforces.

There are four types of constraints:

- A *unique 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 *referential 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* 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 a rule that can be used by the SQL compiler, but that is not enforced by the database manager.

Referential and table check constraints can be turned on or off. It is generally a good idea, for example, to turn off the enforcement of a constraint when large amounts of data are loaded into a database.

Unique constraints

A *unique constraint* is the rule that the values of a key are valid only if they are unique within a table. Unique constraints are optional and can be defined in the CREATE TABLE or ALTER TABLE statement using the PRIMARY KEY clause or the UNIQUE clause. 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.

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 that 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.

Referential constraints

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.

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 are optional and 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, ADD CONSTRAINT, and SET INTEGRITY statements.

Referential constraints with a delete or an update rule of RESTRICT are enforced before all other referential constraints. Referential constraints with a delete or an update rule of NO ACTION behave like RESTRICT in most cases.

Note that referential constraints, check constraints, and triggers can be combined.

Referential integrity rules involve the following concepts and terminology:

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 *self-referencing constraint*.

Self-referencing row

A row that is a parent of itself.

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.

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.

The delete rule of a referential constraint applies when a row of the parent table is deleted. More precisely, the rule applies 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.

Each referential constraint in which a table is a parent has its own delete rule, and all applicable delete rules are used to determine the result of a delete operation. Thus, a row cannot be deleted if it has dependents in a referential constraint with a delete rule of RESTRICT or NO ACTION, or the deletion cascades to any of its descendants that are dependents in a referential constraint with the delete rule of RESTRICT or NO ACTION.

The deletion of a row from parent table P involves other tables and can affect rows of these tables:

- If table D is a dependent of P and the delete rule is RESTRICT or NO ACTION, then D is involved in the operation but is not affected by the operation.
- If D is a dependent of P and the delete rule is SET NULL, then D is involved in the operation, and rows of D can be updated during the operation.
- If D is a dependent of P and the delete rule is CASCADE, then D is involved in the operation and rows of D can be deleted during the operation.

If rows of D are deleted, then the delete operation on P is said to be propagated to D. If D is also a parent table, then the actions described in this list apply, in turn, to the dependents of D.

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.
- 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.

Table check constraints

A *table check constraint* is a rule that specifies the values allowed in one or more columns of every row in a table. A constraint is optional, and can be defined using the CREATE TABLE or the ALTER TABLE statement. Specifying table check constraints is done through a restricted form of a search condition. One of the restrictions is that a column name in a table check constraint on table T must identify a column of table T.

A table can have an arbitrary number of table check constraints. A table check constraint is enforced by applying its search condition to each row that is inserted or updated. An error occurs if the result of the search condition is false for any row.

When one or more table check constraints is defined in the ALTER TABLE statement for a table with existing data, the existing data is checked against the new condition before the ALTER TABLE statement completes. The SET INTEGRITY statement can be used to put the table in *set integrity pending* state, which allows the ALTER TABLE statement to proceed without checking the data.

Informational constraints

An *informational constraint* is a rule that can be used by the SQL compiler to improve the access path 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.

Use the CREATE TABLE or ALTER TABLE statement to define a referential or table check constraint, specifying constraint attributes that determine whether or not the database manager is to enforce the constraint and whether or not the constraint is to be used for query optimization.

Related reference:

- “Interaction of triggers and constraints” in *SQL Reference, Volume 1*
- “SET INTEGRITY statement” in *SQL Reference, Volume 2*

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 SQL statements that are executed whenever the trigger is activated. The SQL 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 SQLstates. If the trigger activation time is after the trigger event, triggered actions can include statements that select, insert, update, delete, or signal SQLstates.

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, and separate transition tables can be defined for old and new values.

Multiple triggers can be specified for a combination of table, event, or 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 SQL 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, temporary tables are used to avoid access conflicts, and this can have a noticeable impact on performance, particularly in partitioned database environments.

Related concepts:

- “Triggers in application development” in *Developing SQL and External Routines*

Related tasks:

- “Creating triggers” in *Administration Guide: Implementation*

Related reference:

- “Interaction of triggers and constraints” in *SQL Reference, Volume 1*

Additional database design considerations

When designing a database, it is important to consider which tables users should be able to access. Access to tables is granted or revoked through authorizations. The highest level of authority is system administration authority (SYSADM). A user with SYSADM authority can assign other authorizations, including database administrator authority (DBADM).

For *audit* purposes, you may have to record every update made to your data for a specified period. For example, you may want to update an audit table each time an employee’s salary is changed. Updates to this table could be made automatically if an appropriate trigger is defined. Audit activities can also be carried out through the DB2 Database for Linux, UNIX, and Windows audit facility.

For performance reasons, you may only want to access a selected amount of data, while maintaining the base data as *history*. You should include within your design, the requirements for maintaining this historical data, such as the number of months or years of data that is required to be available before it can be purged.

You may also want to make use of *summary* information. For example, you may have a table that has all of your employee information in it. However, you would

like to have this information divided into separate tables by division or department. In this case, a materialized query table for each division or department based on the data in the original table would be helpful.

Security implications should also be identified within your design. For example, you may decide to support user access to certain types of data through security tables. You can define access levels to various types of data, and who can access this data. Confidential data, such as employee and payroll data, would have stringent security restrictions.

You can create tables that have a *structured type* associated with them. With such typed tables, you can establish a hierarchical structure with a defined relationship between those tables called a *type hierarchy*. The type hierarchy is made up of a single root type, supertypes, and subtypes.

A *reference type* representation is defined when the root type of a type hierarchy is created. The target of a reference is always a row in a typed table or view.

When working in a High Availability Disaster Recovery (HADR) environment, there are several recommendations:

- The two instances of the HADR environment should be identical in hardware and software. This means:
 - The host computers for the HADR primary and standby databases should be identical.
 - The operating system on the primary and standby systems should have the same version including patches. (During an upgrade this may not be possible. However, the period when the instances are out of step should be kept as short as possible to limit any difficulties arising from the differences. Those differences could include the loss of support for new features during a failover as well as any issues affecting normal non-failover operations.)
 - A TCP/IP interface must be available between the two HADR instances.
 - A high speed, high capacity network is recommended.
- There are DB2 database requirements that should be considered.
 - The database release used by the primary and the standby should be identical.
 - The primary and standby DB2 software must have the same bit size. (That is, both should be at 32-bit or at 64-bit.)
 - The table spaces and their corresponding containers should be identical on the primary and standby databases. Those characteristics that must be symmetrical for the table spaces include (but are not limited to): the table space type (DMS or SMS), table space size, the container paths, the container sizes, and the container file type (raw device or file system). Relative container paths may be used, in which case the relative path must be the same on each instance; they may map to the same or different absolute paths.
 - The primary and standby databases need not have the same database path (as declared when using the CREATE DATABASE command).
 - Buffer pool operations on the primary are replayed on the standby, which suggests the importance of the primary and standby databases having the same amount of memory.

Chapter 5. Physical database design

After you have completed your logical database design, there are a number of issues that you should consider about the physical environment in which your database and tables reside. These issues include understanding the files that are created to support and manage your database, understanding how much space is required to store your data, determining how to use the table spaces that are required to store your data, and determining the structure of the tables used to hold the data. This chapter will discuss these issues.

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 recommended that you explicitly state where you would like the database created.

In the directory you specify 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_directory>/<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. Each file has a duplicate copy to provide a backup.
- The files SQLSPCS.1 and SQLSPCS.2 contain table space information. Each file has a duplicate copy to provide a backup.
- The files SQLSGF.1 and SQLSGF.2 contain storage path information associated with the database's automatic storage. Each file has a duplicate copy to provide a backup.
- The SQLDBCON file contains database configuration information. Do not edit this file. To change configuration parameters, use either the Control Center or the command-line statements UPDATE DATABASE CONFIGURATION and RESET DATABASE CONFIGURATION.

- 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 DB2TSCHNG.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 and SQLOGMIR.LFH, contain information about the active logs.
Recovery processing uses information from this file 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

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)

Related concepts:

- “Comparison of SMS and DMS table spaces” on page 140
- “Database managed space” on page 120
- “DMS device considerations” on page 124
- “DMS table spaces” on page 123
- “SMS table spaces” on page 119
- “Understanding the recovery history file” in *Data Recovery and High Availability Guide and Reference*

Related reference:

- “CREATE DATABASE command” in *Command Reference*

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.

From the Control Center, you can access a number of utilities that are designed to assist you in determining the size requirements of various database objects:

- You can select an object and then use the “Estimate Size” utility. This utility can tell you the current size of an existing object, such as a table. You can then change the object, and the utility will calculate new estimated values for the object. The utility will help you approximate storage requirements, taking future growth into account. It provides possible size ranges for the object: both the smallest size, based on current values, and the largest possible size.
- You can determine the relationships between objects by using the “Show Related” window.
- You can select any database object on the instance and request “Generate DDL”. This function uses the **db2look** utility to generate data definition statements for the database.

In each of these cases, either the “Show SQL” or the “Show Command” button is available to you. You can save the resulting SQL statements or commands in script files to be used later. All of these utilities have online help to assist you.

Keep these utilities in mind as you plan your physical database.

When estimating the size of a database, the contribution of the following must be considered:

- System Catalog Tables
- User Table Data
- Long Field Data
- Large Object (LOB) Data
- Index Space
- Log File Space
- Temporary Work Space

Space requirements related to the following are not discussed:

- 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

Related concepts:

- “Space requirements for indexes” on page 80
- “Space requirements for large object data” on page 79
- “Space requirements for log files” on page 82
- “Space requirements for long field data” on page 78
- “Space requirements for system catalog tables” on page 76
- “Space requirements for temporary tables” on page 83
- “Space requirements for user table data” on page 77

Related reference:

- “db2look - DB2 statistics and DDL extraction tool command” in *Command Reference*

Space requirements for 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.

Related concepts:

- “Space requirements for database objects” on page 75
- “System catalog tables” in *Administration Guide: Implementation*

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 invoked, 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, DB2 Database for Linux, UNIX, and Windows will attempt to physically cluster the data according to the key order of that clustering index. When a row is inserted into the table, DB2 will first look up its key value in the clustering index. If the key value is found, DB2 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, DB2 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, DB2 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. An exception occurs when the row size is smaller than the page size divided by 255. In such a case, there is wasted space on each page. (There is a maximum of only 255 rows per page.) To reduce this wasted space, a smaller page size may be appropriate.

You cannot restore a backup 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.

Related concepts:

- "Space requirements for database objects" on page 75

Space requirements for long field data

Long field 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.

Related concepts:

- “Space requirements for database objects” on page 75

Space requirements for large object 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.

Related concepts:

- “Space requirements for database objects” on page 75

Related reference:

- “Large objects (LOBs)” in *SQL Reference, Volume 1*

Space requirements for indexes

For each index, the space needed can be estimated as:

$$(\text{average index key size} + 9) * \text{number of rows} * 2$$

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 VARCHARIC columns, use an average of the current data size, plus two bytes. Do not use the maximum declared size.)
- The factor of “2” is for overhead, such as non-leaf pages and free space.

Notes:

1. For every column that allows NULLs, 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”.

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 INSPECT command. To convert type-1 indexes to type-2 indexes, use the REORG INDEXES 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:

$$(\text{average index key size} + 9) * \text{number of rows} * 3.2$$

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.

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
- $K = \text{averageKeySize}$

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 .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.

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).
- $M = U / (9 + \text{minimumKeySize}$ for non-leaf pages)
- $K = \text{averageKeySize}$ 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 $(100 - \text{pctfree})/100$, 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:

$$T * \text{pagesize}$$

Related concepts:

- "Indexes" in *SQL Reference, Volume 1*
- "Index cleanup and maintenance" in *Performance Guide*
- "Space requirements for database objects" on page 75

Space requirements for log files

You will require 32 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 will provide a sufficient amount of disk space.

If the database is enabled for roll-forward recovery, special log space requirements should be taken into consideration:

- With the *logretain* configuration parameter enabled, 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 *userexit* configuration parameter enabled, 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 for LOGRETAIN to enable archive logging. DB2 Database for Linux, UNIX, and Windows 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.

Related concepts:

- “Space requirements for database objects” on page 75
- “Log mirroring” in *Data Recovery and High Availability Guide and Reference*
- “Understanding recovery logs” in *Data Recovery and High Availability Guide and Reference*

Related reference:

- “mirrorlogpath - Mirror log path configuration parameter” in *Performance Guide*
- “logprimary - Number of primary log files configuration parameter” in *Performance Guide*
- “logsecond - Number of secondary log files configuration parameter” in *Performance Guide*
- “logfilsiz - Size of log files configuration parameter” in *Performance Guide*

Space requirements for temporary tables

Some SQL 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 is used.

You can use the database system monitor and the query table space 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.

Related concepts:

- “Space requirements for database objects” on page 75

Related reference:

- “sqlbmtsqr API - Get the query data for all table spaces” in *Administrative API Reference*
- “Query compiler variables” in *Performance Guide*

XML storage object overview

DB2 tables can store well formed XML documents in XML columns, alongside columns that contain relational data. In much the same way that LONG VARCHAR and LOB data are stored apart from the other contents of a table, DB2 stores XML data contained in table columns of the type XML in auxiliary XML storage objects. When stored in system managed space, the files associated with XML storage objects have the file type extension .xda.

XML storage objects are separate from, but dependent upon their parent table objects. For each XML value stored in a row of an XML table column, DB2 maintains a record, called an XML data specifier (XDS), which specifies where to retrieve the XML data stored on disk from the associated XML storage object.

You can store XML documents of up to 2 gigabytes in size in a database. Because XML data can be quite large, you may want to monitor the buffering activity for XML data separately from the buffering activity for other data. A number of monitor elements are available to help you gauge the buffer pool activity for XML storage objects.

Related concepts:

- “Preference of database managed table spaces for native XML data store performance” in *Performance Guide*
- “Guidelines for storage requirements for XML documents” on page 84
- “Native XML data store overview” in *XML Guide*
- “XML data specifier” in *Data Movement Utilities Guide and Reference*

Related reference:

- “Buffer pool activity monitor elements” in *System Monitor Guide and Reference*

Guidelines for storage requirements for XML documents

The amount of space that an XML document occupies in a DB2 database is determined by the initial size of the document in raw form and by a number of other properties. The following list includes the most important of these properties:

Document structure

XML documents that contain complex markup tagging require a larger amount of storage space than documents with simple markup. For example, an XML document that has many nested elements, each containing a small amount of text or having short attribute values, occupies more storage space than an XML document composed primarily of textual content.

Node names

The length of element names, attribute names, namespace prefixes and similar, non-content data also affect storage size. Any information unit of this type that exceeds 4 bytes in raw form is compressed for storage, resulting in comparatively greater storage efficiency for longer node names.

Ratio of attributes to elements

Typically, the more attributes that are used per element, the lower the amount of storage space that is required for the XML document.

Document codepage

XML documents with encoding that uses more than one byte per character occupy a larger amount storage space than documents using a single-byte character set.

Document validation

XML documents are annotated after having been validated against an XML schema. The addition of type information after validation results in an increased storage requirement.

Related concepts:

- “XML storage object overview” on page 84

Database partition groups

A database partition group is a set of one or more database partitions defined as belongin to a database. When you want to create tables for the database, you first create the database partition group where the table spaces will be stored, then you create the table space where the tables will be stored.

You can define named subsets of one or more database partitions in a database. Each subset you define is known as a *database partition group*. Each subset that contains more than one database partition is known as a *multipartition database partition group*. Multipartition database partition groups can only be defined with database partitions that belong to the same instance. A database partition group can contain as few as one database partition, or span all of the database partitions in the database.

Figure 22 on page 86 shows an example of a database with five database partitions in which:

- A database partition group spans all but one of the database partitions (Database Partition Group 1).
- A database partition group contains one database partition (Database Partition Group 2).
- A database partition group contains two database partitions. (Database Partition Group 3).
- The database partition within Database Partition Group 2 is shared (and overlaps) with Database Partition Group 1.
- There is a single database partition within Database Partition Group 3 that is shared (and overlaps) with Database Partition Group 1.

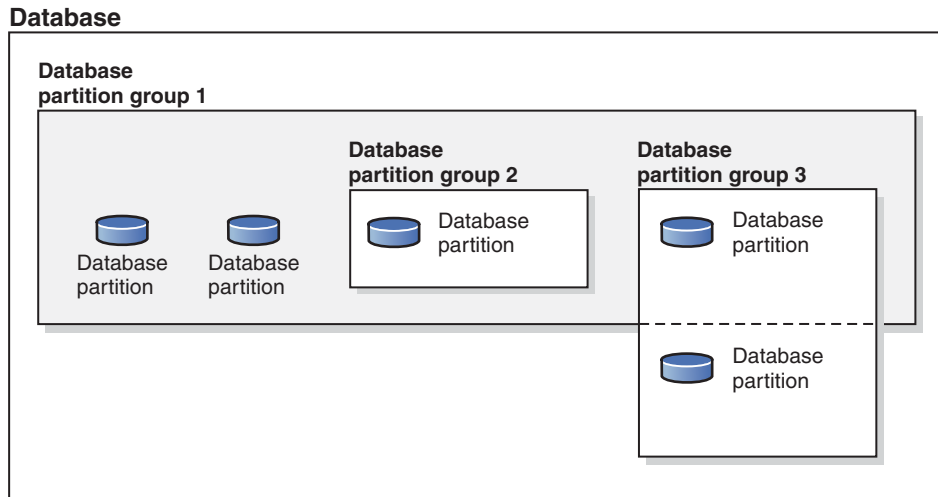


Figure 22. Database partition groups in a database

You create a new database partition group using the `CREATE DATABASE PARTITION GROUP` statement. You can modify it using the `ALTER DATABASE PARTITION GROUP` statement. Data is divided across all the database partitions in a database partition group, and you can add or drop one or more database partitions from a database partition group. If you are using a multipartition database partition group, you must look at several database partition group design considerations.

Each database partition that is part of the database system configuration must already be defined in a *database partition configuration file* called `db2nodes.cfg`. A database partition group can contain as few as one database partition, or as many as the entire set of database partitions defined for the database system.

When a database partition group is created or modified, a *distribution map* is associated with it. A distribution map, in conjunction with a *distribution key* and a hashing algorithm, is used by the database manager to determine which database partition in the database partition group will store a given row of data.

In a non-partitioned database, no distribution key or distribution map is required. A database partition is a part of the database, complete with user data, indexes, configuration files, and transaction logs. Default database partition groups that were created when the database was created are used by the database manager. `IBMCATGROUP` is the default database partition group for the table space containing the system catalogs. `IBMTEMPGROUP` is the default database partition group for system temporary table spaces. `IBMDEFAULTGROUP` is the default database partition group for the table spaces containing the user defined tables that you may choose to put there. A user temporary table space for a declared temporary table can be created in `IBMDEFAULTGROUP` or any user-created database partition group but not in `IBMTEMPGROUP`.

When working with database partition groups you can:

- Create a database partition group.
- Change the comment associated with a database partition group.
- Add database partitions to a database partition group.
- Drop database partitions from a database partition group.
- Redistribute table data within a database partition group.

Related concepts:

- “Database partition group design” on page 87
- “Distribution keys” on page 89
- “Distribution maps” on page 88

Related reference:

- “ALTER DATABASE PARTITION GROUP statement” in *SQL Reference, Volume 2*
- “CREATE DATABASE PARTITION GROUP statement” in *SQL Reference, Volume 2*

Database partition group design

There are no database partition group design considerations if you are using a single-partition database.

The DB2 Design Advisor is a tool that can be used to recommend database partition groups. The DB2 Design Advisor can be accessed from the Control Center and using `db2advis` from the command line processor.

If you are using a multiple partition database partition group, consider the following design points:

- In a multipartition database partition group, you can only create a unique index if it is a superset of the distribution key.
- Depending on the number of database partitions in the database, you may have one or more single-partition database partition groups, and one or more multipartition database partition groups present.
- Each database partition must be assigned a unique number. The same database partition may be found in one or more database partition groups.
- To ensure fast recovery of the database partition containing system catalog tables, avoid placing user tables on the same database partition. This is accomplished by placing user tables in database partition groups that do not include the database partition in the `IBMCATGROUP` database partition group.

You should place small tables in single-partition database partition groups, except when you want to take advantage of *collocation* with a larger table. Collocation is the placement of rows from different tables that contain related data in the same database partition. Collocated tables allow DB2 Database for Linux, UNIX, and Windows to utilize more efficient join strategies. Collocated tables can reside in a single-partition database partition group. Tables are considered collocated if they reside in a multipartition database partition group, have the same number of columns in the distribution key, and if the data types of the corresponding columns are compatible. Rows in collocated tables with the same distribution key value are placed on the same database partition. Tables can be in separate table spaces in the same database partition group, and still be considered collocated.

You should avoid extending medium-sized tables across too many database partitions. For example, a 100 MB table may perform better on a 16-partition database partition group than on a 32-partition database partition group.

You can use database partition groups to separate online transaction processing (OLTP) tables from decision support (DSS) tables, to ensure that the performance of OLTP transactions is not adversely affected.

Related concepts:

- “Database partition groups” on page 85
- “Database partition compatibility” on page 91
- “Distribution keys” on page 89
- “Distribution maps” on page 88
- “Replicated materialized query tables” on page 111
- “Table collocation” on page 91

Related reference:

- “db2advis - DB2 design advisor command” in *Command Reference*

Distribution maps

In a partitioned database environment, the database manager must know where to find the data it needs. The database manager uses a map, called a *distribution map*, to find the data.

A distribution map is an internally generated array containing either 4 096 entries for multiple-partition database partition groups, or a single entry for single-partition database partition groups. For a single-partition database partition group, the distribution map has only one entry containing the number of the database partition where all the rows of a database table are stored. For multiple-partition database partition groups, the numbers of the database partition group are specified in a way such that each database partition is used one after the other to ensure an even distribution across the entire map. Just as a city map is organized into sections using a grid, the database manager uses a *distribution key* to determine the location (the database partition) where the data is stored.

For example, assume that you have a database created on four database partitions (numbered 0–3). The distribution map for the IBMDEFAULTGROUP database partition group of this database would be:

```
0 1 2 3 0 1 2 ...
```

If a database partition group had been created in the database using database partitions 1 and 2, the distribution map for that database partition group would be:

```
1 2 1 2 1 2 1 ...
```

If the distribution key for a table to be loaded into the database is an integer with possible values between 1 and 500 000, the distribution key is hashed to a number between 0 and 4 095. That number is used as an index into the distribution map to select the database partition for that row.

Figure 23 on page 89 shows how the row with the distribution key value (c1, c2, c3) is mapped to number 2, which, in turn, references database partition n5.

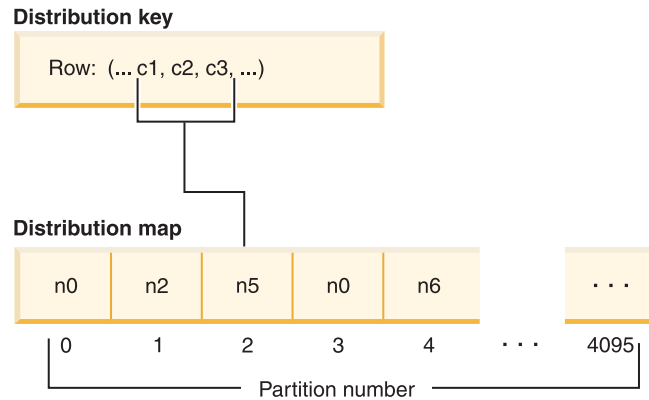


Figure 23. Data distribution using a distribution map

A distribution map is a flexible way of controlling where data is stored in a multi-partition database. If you need to change the data distribution across the database partitions in your database, you can use the data redistribution utility. This utility allows you to rebalance or introduce skew into the data distribution.

You can use the `sqlugtpi API - Get table distribution information`) to obtain a copy of a distribution map that you can view.

Related concepts:

- “Database partition group design” on page 87
- “Database partition groups” on page 85
- “Distribution keys” on page 89

Related reference:

- “`sqlugtpi API - Get table distribution information`” in *Administrative API Reference*

Distribution keys

A *distribution key* is a column (or group of columns) that is used to determine the database partition in which a particular row of data is stored. A distribution key is defined on a table using the CREATE TABLE statement. If a distribution key is not defined for a table in a table space that is divided across more than one database partition in a database partition group, one is created by default from the first column of the primary key. If no primary key is specified, the default distribution key is the first non-long field column defined on that table. (*Long* includes all long data types and all large object (LOB) data types). If you are creating a table in a table space associated with a single-partition database partition group, and you want to have a distribution key, you must define the distribution key explicitly. One is not created by default.

If no columns satisfy the requirement for a default distribution key, the table is created without one. Tables without a distribution key are only allowed in single-partition database partition groups. You can add or drop distribution keys later, using the ALTER TABLE statement. Altering the distribution key can only be done to a table whose table space is associated with a single-partition database partition group.

Choosing a good distribution key is important. You should take into consideration:

- How tables are to be accessed
- The nature of the query workload
- The join strategies employed by the database system.

If collocation is not a major consideration, a good distribution key for a table is one that spreads the data evenly across all database partitions in the database partition group. The distribution key for each table in a table space that is associated with a database partition group determines if the tables are collocated. Tables are considered collocated when:

- The tables are placed in table spaces that are in the same database partition group
- The distribution keys in each table have the same number of columns
- The data types of the corresponding columns are partition-compatible.

These characteristics ensure that rows of collocated tables with the same distribution key values are located on the same database partition.

An inappropriate distribution key can cause uneven data distribution. Columns with unevenly distributed data, and columns with a small number of distinct values should not be chosen as distribution keys. The number of distinct values must be great enough to ensure an even distribution of rows across all database partitions in the database partition group. The cost of applying the distribution algorithm is proportional to the size of the distribution key. The distribution key cannot be more than 16 columns, but fewer columns result in better performance. Unnecessary columns should not be included in the distribution key.

The following points should be considered when defining distribution keys:

- Creation of a multiple-partition table that contains only long data types (LONG VARCHAR, LONG VARGRAPHIC, BLOB, CLOB, or DBCLOB) is not supported.
- The distribution key definition cannot be altered.
- The distribution key should include the most frequently joined columns.
- The distribution key should be made up of columns that often participate in a GROUP BY clause.
- Any unique key or primary key must contain all of the distribution key columns.
- In an online transaction processing (OLTP) environment, all columns in the distribution key should participate in the transaction by using equal (=) predicates with constants or host variables. For example, assume you have an employee number, *emp_no*, that is often used in transactions such as:

```
UPDATE emp_table SET ... WHERE
emp_no = host-variable
```

In this case, the EMP_NO column would make a good single column distribution key for EMP_TABLE.

Database partitioning is the method by which the placement of each row in the table is determined. The method works as follows:

1. A hashing algorithm is applied to the value of the distribution key, and generates a number between zero (0) and 4095.
2. The distribution map is created when a database partition group is created. Each of the numbers is sequentially repeated in a round-robin fashion to fill the distribution map.

3. The number is used as an index into the distribution map. The number at that location in the distribution map is the number of the database partition where the row is stored.

Related concepts:

- “Database partition group design” on page 87
- “Database partition groups” on page 85
- “Distribution maps” on page 88
- “The Design Advisor” in *Performance Guide*

Related reference:

- “ALTER TABLE statement” in *SQL Reference, Volume 2*

Table collocation

You may discover that two or more tables frequently contribute data in response to certain queries. In this case, you will want related data from such tables to be located as close together as possible. In an environment where the database is physically divided among two or more database partitions, there must be a way to keep the related pieces of the divided tables as close together as possible. The ability to do this is called *table collocation*.

Tables are collocated when they are stored in the same database partition group, and when their distribution keys are compatible. Placing both tables in the same database partition group ensures a common distribution map. The tables may be in different table spaces, but the table spaces must be associated with the same database partition group. The data types of the corresponding columns in each distribution key must be *partition-compatible*.

DB2 Database for Linux, UNIX, and Windows can recognize, when accessing more than one table for a join or a subquery, that the data to be joined is located at the same database partition. When this happens, DB2 can perform the join or subquery at the database partition where the data is stored, instead of having to move data between database partitions. This ability has significant performance advantages.

Related concepts:

- “Database partition group design” on page 87
- “Database partition groups” on page 85
- “Database partition compatibility” on page 91
- “Distribution keys” on page 89

Database partition compatibility

The base data types of corresponding columns of distribution keys are compared and can be declared *partition-compatible*. Partition-compatible data types have the property that two variables, one of each type, with the same value, are mapped to the same number by the same partitioning algorithm.

Partition-compatibility has the following characteristics:

- A base data type is compatible with another of the same base data type.
- Internal formats are used for DATE, TIME, and TIMESTAMP data types. They are not compatible with each other, and none are compatible with CHAR.

- Partition-compatibility is not affected by columns with NOT NULL or FOR BIT DATA definitions.
- NULL values of compatible data types are treated identically; those of non-compatible data types may not be.
- Base data types of a user-defined type are used to analyze partition-compatibility.
- Decimals of the same value in the distribution key are treated identically, even if their scale and precision differ.
- Trailing blanks in character strings (CHAR, VARCHAR, GRAPHIC, or VARGRAPHIC) are ignored by the hashing algorithm.
- BIGINT, SMALLINT, and INTEGER are compatible data types.
- REAL and FLOAT are compatible data types.
- CHAR and VARCHAR of different lengths are compatible data types.
- GRAPHIC and VARGRAPHIC are compatible data types.
- Partition-compatibility does not apply to LONG VARCHAR, LONG VARGRAPHIC, CLOB, DBCLOB, and BLOB data types, because they are not supported as distribution keys.

Related concepts:

- “Database partition group design” on page 87
- “Database partition groups” on page 85
- “Distribution keys” on page 89

Data partitions

A data partition is a set of table rows, stored separately from other sets of rows, and grouped by the specifications provided in the PARTITION BY clause of the CREATE TABLE statement. If a table is created using the PARTITION BY clause, then the table is partitioned.

A partitioned table uses a data organization scheme in which table data is divided across multiple storage objects, called data partitions or ranges, according to values in one or more table partitioning key columns of the table. Data from a given table is partitioned into multiple storage objects based on the specifications provided in the PARTITION BY clause of the CREATE TABLE statement. These storage objects can be in different table spaces, in the same table space, or a combination of both. All the table spaces specified must have the same: pagesize, extensize, storage mechanism (DMS, SMS), and type (REGULAR or LARGE) and all the table spaces must be in the same database partition group.

A partitioned table simplifies the rolling in and rolling out of table data and a partitioned table can contain vastly more data than an ordinary table. You can create a partitioned table with a maximum of 32767 data partitions. Data partitions can be added to, attached to, and detached from a partitioned table, and you can store multiple data partition ranges from a table in one table space.

The ranges specified for each data partition can be generated automatically or manually when creating a table.

Data partitions are referred to in various ways throughout the DB2 library. The following list represents the most common references:

- DATAPARTITIONNAME is the permanent name assigned to a data partition for a given table at create time. This column value is stored in the SYSCAT.DATAPARTITIONS catalog view. This name is not preserved on an attach or detach operation.
- DATAPARTITIONID is the permanent identifier assigned to a data partition for a given table at create time. It is used to uniquely identify a particular data partition in a given table. This identifier is not preserved on an attach or detach operation. This value is system generated and may appear in output from various utilities.
- SEQNO indicates the order of a particular data partition range with regards to other data partition ranges in the table, with detached data partitions sorting after all visible and attached data partitions.

Related concepts:

- “Data organization schemes in DB2 and Informix databases” on page 105
- “Optimization strategies for partitioned tables” in *Performance Guide*
- “Partitioned tables” on page 104
- “Table partitioning” on page 93

Related tasks:

- “Adding data partitions to partitioned tables” in *Administration Guide: Implementation*
- “Altering partitioned tables” in *Administration Guide: Implementation*
- “Creating partitioned tables” in *Administration Guide: Implementation*
- “Dropping a data partition” in *Administration Guide: Implementation*
- “Approaches to migrating existing tables and views to partitioned tables” in *Administration Guide: Implementation*
- “Attaching a data partition” in *Administration Guide: Implementation*
- “Detaching a data partition” in *Administration Guide: Implementation*
- “Rotating data in a partitioned table” in *Administration Guide: Implementation*
- “Approaches to defining ranges on partitioned tables” in *Administration Guide: Implementation*

Related reference:

- “Examples of rolling in and rolling out partitioned table data” in *Administration Guide: Implementation*
- “CREATE TABLE statement” in *SQL Reference, Volume 2*
- “Guidelines and restrictions on altering partitioned tables with attached or detached data partitions” in *Administration Guide: Implementation*

Table partitioning

Table partitioning is a data organization scheme in which table data is divided across multiple storage objects called data partitions or ranges according to values in one or more table columns. Each data partition is stored separately. These storage objects can be in different table spaces, in the same table space, or a combination of both.

Storage objects behave much like individual tables, making it easy to accomplish fast roll-in by incorporating an existing table into a partitioned table using the ALTER TABLE ...ATTACH statement. Likewise, easy roll-out is accomplished with

the ALTER TABLE ...DETACH statement. Query processing can also take advantage of the separation of the data to avoid scanning irrelevant data, resulting in better query performance for many data warehouse style queries.

Table data is partitioned as specified in the PARTITION BY clause of the CREATE TABLE statement. The columns used in this definition are referred to as the table partitioning key columns.

This organization scheme can be used in isolation or in combination with other organization schemes. By combining the DISTRIBUTE BY and PARTITION BY clauses of the CREATE TABLE statement, data can be spread across database partitions spanning multiple table spaces. The DB2 organization schemes include:

- DISTRIBUTE BY HASH
- PARTITION BY RANGE
- ORGANIZE BY DIMENSIONS

Table partitioning functionality is available with DB2 Version 9.1 Enterprise Server Edition for Linux, UNIX, and Windows.

Benefits of table partitioning:

If any of the following circumstances apply to you and your organization, consider the numerous benefits of table partitioning:

- You have a data warehouse that would benefit from easier roll-in and roll-out of table data
- You have a data warehouse that includes large tables
- You are considering a migration to a DB2 V9.1 database from a previous release or a competitive database product
- You need to use Hierarchical Storage Management (HSM) solutions more effectively

Table partitioning offers easy roll-in and roll-out of table data, easier administration, flexible index placement and better query processing.

Efficient roll-in and roll-out

Table partitioning allows for the efficient roll-in and roll-out of table data. You can achieve this by using the ATTACH PARTITION and DETACH PARTITION clauses of the ALTER TABLE statement. Rolling in partitioned table data allows a new range to be easily incorporated into a partitioned table as an additional data partition. Rolling out partitioned table data allows you to easily separate ranges of data from a partitioned table for subsequent purging or archiving.

Easier administration of large tables

Table level administration is more flexible because you can perform administrative tasks on individual data partitions. These tasks include: detaching and reattaching of a data partition, backing up and restoring individual data partitions, and reorganizing individual indexes. Time consuming maintenance operations can be shortened by breaking them down into a series of smaller operations. For example, backup operations can work data partition by data partition when the data partitions are placed in separate table spaces. Thus, it is possible to backup one data partition of a partitioned table at a time.

Flexible index placement

Indexes can now be placed in different table spaces allowing for more granular control of index placement. Some benefits of this new design include:

- Improved performance of drop index and online index create.
- The ability to use different values for any of the table space characteristics between each index on the table (for example, different page sizes for each index may be appropriate to ensure better space utilization).
- Reduced IO contention providing more efficient concurrent access to the index data for the table.
- When individual indexes are dropped space will immediately become available to the system without the need for an index reorganization.
- If you choose to perform index reorganization, an individual index can be reorganized.

Both DMS and SMS table spaces support the use of indexes in a different location than the table.

Improved performance for business intelligence style queries

Query processing is enhanced to automatically eliminate data partitions based on predicates of the query. This functionality, known as Data Partition Elimination, benefits many decision support queries.

The following example creates a table *customer* where rows with *l_shipdate* >= '01/01/2006' and *l_shipdate* <= '03/31/2006' are stored in table space *ts1*, rows with *l_shipdate* >= '04/01/2006' and *l_shipdate* <= '06/30/2006' are in table space *ts2*, etc.

```
CREATE TABLE customer (l_shipdate, l_name CHAR(30))
IN ts1, ts2, ts3, ts4, ts5
PARTITION BY RANGE(l_shipdate) (STARTING FROM ('01/01/2006')
ENDING AT ('12/31/2006') EVERY (3 MONTHS))
```

Related concepts:

- “Data organization schemes” on page 99
- “Data partitions” on page 92
- “Partitioned materialized query table behavior” in *Administration Guide: Implementation*
- “Partitioned tables” on page 104
- “Understanding index behavior on partitioned tables” in *Performance Guide*
- “Optimization strategies for partitioned tables” in *Performance Guide*
- “Understanding clustering index behavior on partitioned tables” in *Performance Guide*

Related tasks:

- “Altering partitioned tables” in *Administration Guide: Implementation*
- “Adding data partitions to partitioned tables” in *Administration Guide: Implementation*
- “Approaches to defining ranges on partitioned tables” in *Administration Guide: Implementation*
- “Approaches to migrating existing tables and views to partitioned tables” in *Administration Guide: Implementation*
- “Attaching a data partition” in *Administration Guide: Implementation*

- “Creating partitioned tables” in *Administration Guide: Implementation*
- “Detaching a data partition” in *Administration Guide: Implementation*
- “Dropping a data partition” in *Administration Guide: Implementation*
- “Rotating data in a partitioned table” in *Administration Guide: Implementation*

Related reference:

- “Examples of rolling in and rolling out partitioned table data” in *Administration Guide: Implementation*
- “CREATE TABLE statement” in *SQL Reference, Volume 2*
- “Command Line Processor (CLP) samples” in *Samples Topics*

Table partitioning keys

A table partitioning key is an ordered set of one or more columns in a table. The values in the table partitioning key columns are used to determine in which data partition each table row belongs.

To define the table partitioning key on a table use the CREATE TABLE statement with the PARTITION BY clause.

Choosing an effective table partitioning key column is essential to taking full advantage of the benefits of table partitioning. The following guidelines can help you to choose the most effective table partitioning key columns for your partitioned table:

- Define ranges to match the data roll-in size. It is most common to partition data on a date or time column.
- Define range granularity to match data roll-out. It is most common to use month or quarter.
- Partition on a column that provides advantages in partition elimination.

Supported data types:

The following data types (including synonyms) are supported for use as a table partitioning key column:

SMALLINT	INTEGER
INT	BIGINT
FLOAT	REAL
DOUBLE	DECIMAL
DEC	NUMERIC
NUM	CHARACTER
CHAR	VARCHAR
DATE	TIME
GRAPHIC	VARGRAPHIC
CHARACTER VARYING	TIMESTAMP
CHAR VARYING	CHARACTER FOR BIT DATA
CHAR FOR BIT DATA	VARCHAR FOR BIT DATA
CHARACTER VARYING FOR BIT DATA	CHAR VARYING FOR BIT DATA
User defined types (distinct)	

Unsupported data types:

The following data types can appear in a partitioned table, but are not supported for use as a table partitioning key column:

- User defined types (structured)
- LONG VARCHAR
- LONG VARCHAR FOR BIT DATA
- BLOB
- BINARY LARGE OBJECT
- CLOB
- CHARACTER LARGE OBJECT
- DBCLOB
- LONG VARGRAPHIC
- REF
- Varying length string for C
- Varying length string for Pascal

The following data types are not supported in a partitioned table:

- XML
- DATALINK

If you choose to automatically generate data partitions using the EVERY clause of the CREATE TABLE statement, only one column can be used as the table partitioning key. If you choose to manually generate data partitions by specifying each range in the PARTITION BY clause of the CREATE TABLE statement, multiple columns can be used as the table partitioning key, as shown in the following example:

```
CREATE TABLE sales (year INT, month INT)
  IN tbsp1, tbsp2, tbsp3, tbsp4, tbsp5, tbsp6, tbsp7, tbsp8
  PARTITION BY RANGE(year, month)
  (STARTING FROM (2001, 1) ENDING (2001,3) IN tbsp1,
   ENDING (2001,6) IN tbsp2, ENDING (2001,9)
   IN tbsp3, ENDING (2001,12) IN tbsp4,
   ENDING (2002,3) IN tbsp5, ENDING (2002,6)
   IN tbsp6, ENDING (2002,9) IN tbsp7,
   ENDING (2002,12) IN tbsp8)
```

This results in eight data partitions, one for each quarter in year 2001 and 2002.

Note:

1. When multiple columns are used as the table partitioning key, they are treated as a composite key (which are similar to composite keys in an index), in the sense that trailing columns are dependent on the leading columns. Each starting or ending value (all of the columns, together) must be specified in 512 characters or less. This limit corresponds to the size of the LOWVALUE and HIGHVALUE columns of the SYSCAT.DATAPARTITIONS catalog view. A starting or ending value specified with more than 512 characters will result in error SQL0636N, reason code 9.
2. Table partitioning is multicolumn not multidimension. In table partitioning, all columns used are part of a single dimension.

Generated columns:

Generated columns can be used as table partitioning keys. This example creates a table with twelve data partitions, one for each month. All rows for January of any year will be placed in the first data partition, rows for February in the second, and so on.

Example 1

```
CREATE TABLE monthly_sales (sales_date date,  
sales_month int GENERATED ALWAYS AS (month(sales_date)))  
PARTITION BY RANGE (sales_month)  
(STARTING FROM 1 ENDING AT 12 EVERY 1);
```

Note:

1. You cannot alter or drop the expression of a generated column that is used in the table partitioning key. Adding a generated column expression on a column that is used in the table partitioning key is not permitted. Attempting to add, drop or alter a generated column expression for a column used in the table partitioning key results in error (SQL0270N rc=52).
2. Data partition elimination will not be used for range predicates if the generated column is not monotonic, or the optimizer can not detect that it is monotonic. In the presence of non-monotonic expressions, data partition elimination can only take place for equality or IN predicates. For a detailed discussion and examples of monotonicity see Multidimensional clustering (MDC) table creation, placement, and use.

Related concepts:

- “Table partitioning” on page 93
- “Data partitions” on page 92
- “Optimization strategies for partitioned tables” in *Performance Guide*
- “Partitioned tables” on page 104

Related tasks:

- “Approaches to defining ranges on partitioned tables” in *Administration Guide: Implementation*
- “Adding data partitions to partitioned tables” in *Administration Guide: Implementation*
- “Creating partitioned tables” in *Administration Guide: Implementation*
- “Dropping a data partition” in *Administration Guide: Implementation*
- “Approaches to migrating existing tables and views to partitioned tables” in *Administration Guide: Implementation*
- “Attaching a data partition” in *Administration Guide: Implementation*
- “Detaching a data partition” in *Administration Guide: Implementation*
- “Rotating data in a partitioned table” in *Administration Guide: Implementation*

Related reference:

- “SYSCAT.DATAPARTITIONS catalog view” in *SQL Reference, Volume 1*
- “Examples of rolling in and rolling out partitioned table data” in *Administration Guide: Implementation*
- “CREATE TABLE statement” in *SQL Reference, Volume 2*
- “DESCRIBE command” in *Command Reference*

- “DESCRIBE statement” in *SQL Reference, Volume 2*

Data organization schemes

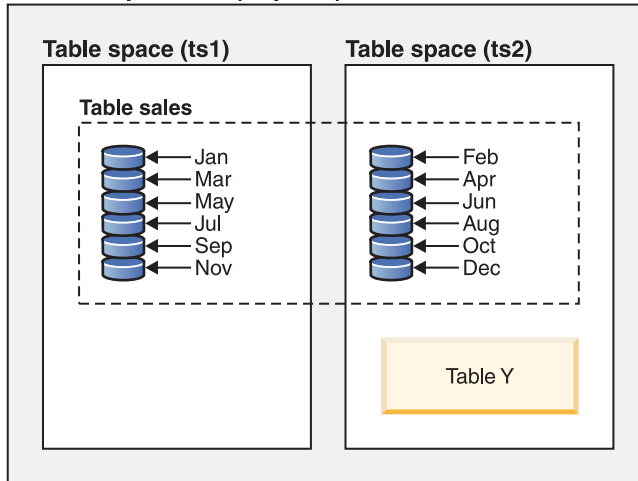
With the introduction of table partitioning, a DB2 database offers a three level data organization scheme. Each clause of the CREATE TABLE statement includes an algorithm to indicate how the data should be organized. The following three clauses demonstrate the levels of data organization that can be used together in any combination:

- DISTRIBUTE BY to spread data evenly across database partitions (to enable intra-query parallelism and to balance the load across each database partition)(database partitioning)
- PARTITION BY to group rows with similar values of a single dimension in the same data partition (table partitioning)
- ORGANIZE BY to group rows with similar values on multiple dimensions in the same table extent (multidimensional clustering)

This syntax allows consistency between the clauses as well as allowing for future algorithms of data organization. Each of these clauses can be used in isolation or in combination with one another. By combining the DISTRIBUTE BY and PARTITION BY clauses of the CREATE TABLE statement data can be spread across database partitions spanning multiple table spaces. This approach allows for similar behavior to the Informix[®] Dynamic Server and Informix Extended Parallel Server hybrid functionality.

In a single table, you can combined the clauses used in each data organization scheme to create more sophisticated partitioning schemes. For example, DB2 Database Partitioning Feature (DPF) is not only compatible, but also complementary to table partitioning.

Database partition (dbpart1)



Legend


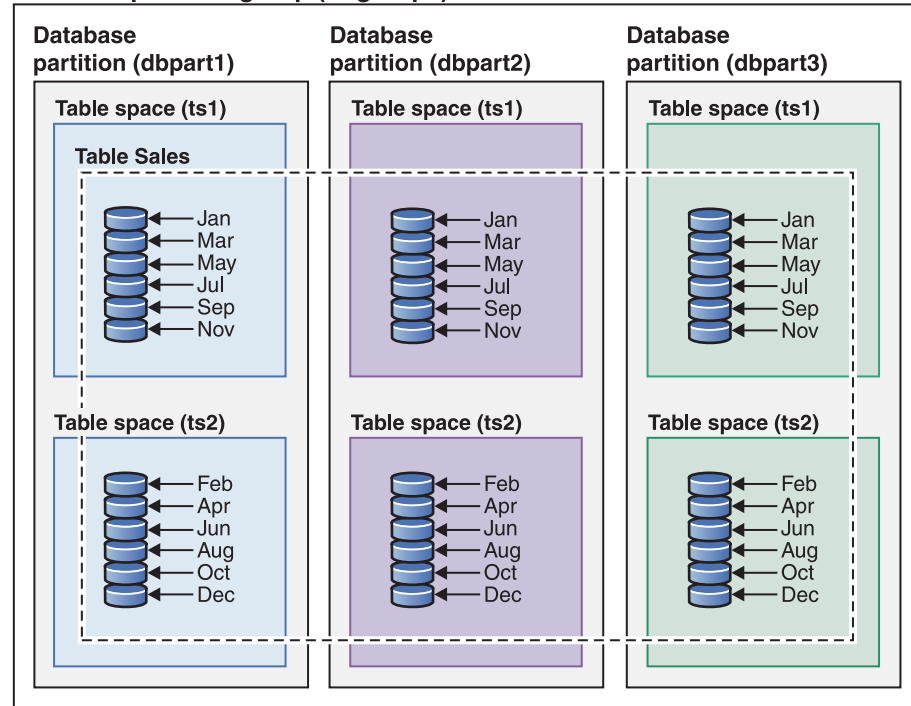


Figure 24. Demonstrating the table partitioning organization scheme where a table representing monthly sales data is partitioned into multiple data partitions. The table also spans two table spaces (ts1 and ts2).

Database partition group (dbgroup1)



Legend




Figure 25. Demonstrating the complementary organization schemes of database partitioning and table partitioning. A table representing monthly sales data is partitioned into multiple data partitions, spanning two table spaces (ts1 and ts2) that are distributed across multiple database partitions (dbpart1, dbpart2, dbpart3) of a database partition group (dbgroup1).

The salient distinction between multidimensional clustering (MDC) and table partitioning is multi-dimension versus single dimension. MDC is suitable to cubes (that is, tables with multiple dimensions), while table partitioning works well if there is a single dimension which is central to the database design, such as a DATE column. MDC and table partitioning are complementary when both of these conditions are met. This is demonstrated in Figure 26 on page 102.

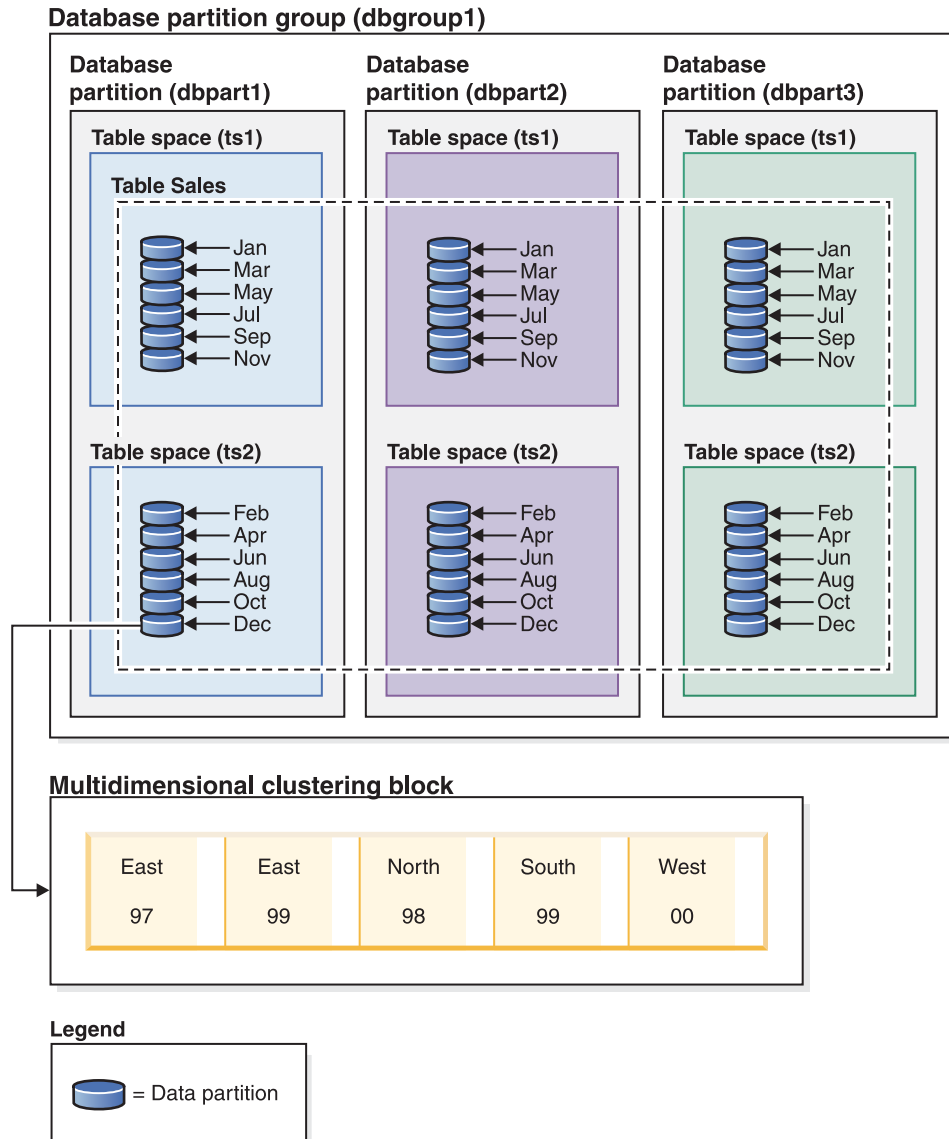


Figure 26. A representation of the database partitioning, table partitioning and multi-dimensional organization schemes where data from table SALES is not only distributed across multiple database partitions, partitioned across table spaces ts1 and ts2, but also groups rows with similar values on both the date and region dimensions.

There is another data organization scheme which cannot be used in conjunction with any of those listed above. This scheme is ORGANIZE BY KEY SEQUENCE. It is used to insert each record into a row that was reserved for that record at the time of table creation (Range-clustered table).

Data organization terminology:

Database partitioning

A data organization scheme in which table data is divided across multiple database partitions based on the hash values in one or more distribution key columns of the table, and based on the use of a distribution map of the database partitions. Data from a given table is distributed based on the specifications provided in the DISTRIBUTE BY HASH clause of the CREATE TABLE statement.

Database partition

A portion of a database on a database partition server consisting of its own user data, indexes, configuration file, and transaction logs. Database partitions can be logical or physical.

Table partitioning

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 based on the specifications provided in the PARTITION BY RANGE clause of the CREATE TABLE statement. These storage objects can be in different table spaces.

Data partition

A set of table rows, stored separately from other sets of rows, grouped by the specifications provided in the PARTITION BY RANGE clause of the CREATE TABLE statement.

Multidimensional clustering (MDC)

A table whose data is physically organized into blocks along one or more dimensions, or clustering keys, specified in the ORGANIZE BY DIMENSIONS clause.

Benefits of each data organization scheme:

Understanding the benefits of each data organization scheme can help you to determine the best approach when planning, designing, or reassessing your database system requirements. Table 21 provides a high-level view of common customer requirements and shows how the various data organization schemes can help you to meet those requirements.

Table 21. Using table partitioning with the Database Partitioning Feature

Issue	Recommended scheme	Explanation
Data roll-out	Table partitioning	Uses detach to roll-out large amounts of data with minimal disruption
Parallel query execution (query performance)	Database Partitioning Feature	Provides query parallelism for improved query performance
Data partition elimination (query performance)	Table partitioning	Provides data partition elimination for improved query performance
Maximization of query performance	Both	Maximum query performance when used together: query parallelism and data partition elimination are complementary
Heavy administrator workload	Database Partitioning Feature	Execute many tasks for each database partition

Table 22. Using table partitioning with MDC tables

Issue	Recommended scheme	Explanation
Data availability during roll-out	Table partitioning	Use the DETACH PARTITION clause to roll out large amounts of data with minimal disruption.

Table 22. Using table partitioning with MDC tables (continued)

Issue	Recommended scheme	Explanation
Query performance	Both	MDC is best for querying multiple dimensions. Table partitioning helps through data partition elimination.
Minimal reorganization	MDC	MDC maintains clustering, which reduces the need to reorganize.

Note: Table partitioning is now recommended over UNION ALL views.

Related concepts:

- “Data organization schemes in DB2 and Informix databases” on page 105
- “Data partitions” on page 92
- “Database partition group design” on page 87
- “Designing multidimensional clustering (MDC) tables” on page 189
- “Multidimensional clustering (MDC) table creation, placement, and use” on page 197
- “Multidimensional clustering tables” on page 172
- “Range-clustered tables” on page 168
- “Table partitioning” on page 93
- “Database database partition group impact on query optimization” in *Performance Guide*
- “Optimization strategies for partitioned tables” in *Performance Guide*
- “Examples of range-clustered tables” in *Administration Guide: Implementation*
- “Guidelines for using range-clustered tables” in *Administration Guide: Implementation*
- “Database partitioning across multiple database partitions” in *SQL Reference, Volume 1*

Related tasks:

- “Creating partitioned tables” in *Administration Guide: Implementation*
- “Creating a table in a partitioned database environment” in *Administration Guide: Implementation*
- “Creating a table in multiple table spaces” in *Administration Guide: Implementation*
- “Attaching a data partition” in *Administration Guide: Implementation*
- “Detaching a data partition” in *Administration Guide: Implementation*
- “Rotating data in a partitioned table” in *Administration Guide: Implementation*

Related reference:

- “CREATE TABLE statement” in *SQL Reference, Volume 2*

Partitioned tables

Partitioned tables use a data organization scheme in which table data is divided across multiple storage objects, called data partitions or ranges, according to values in one or more table partitioning key columns of the table. Data from a given table is partitioned into multiple storage objects based on the specifications provided in

the PARTITION BY clause of the CREATE TABLE statement. These storage objects can be in different table spaces, in the same table space, or a combination of both. Table partitioning functionality is available with DB2 Version 9.1 Enterprise Server Edition for Linux, UNIX, and Windows.

Table partitioning offers easy roll-in and roll-out of table data, easier administration, flexible index placement and better query processing.

Partitioned hierarchical or temporary tables, range-clustered tables, and partitioned views are not supported.

The following column types are not supported for use in partitioned tables:

- XML
- DATALINK

Related concepts:

- “Table partitioning” on page 93
- “Table partitioning keys” on page 96
- “Data partitions” on page 92
- “Data organization schemes in DB2 and Informix databases” on page 105

Related tasks:

- “Adding data partitions to partitioned tables” in *Administration Guide: Implementation*
- “Altering partitioned tables” in *Administration Guide: Implementation*
- “Altering a table” in *Administration Guide: Implementation*
- “Creating partitioned tables” in *Administration Guide: Implementation*
- “Dropping a data partition” in *Administration Guide: Implementation*
- “Attaching a data partition” in *Administration Guide: Implementation*
- “Detaching a data partition” in *Administration Guide: Implementation*
- “Rotating data in a partitioned table” in *Administration Guide: Implementation*
- “Approaches to defining ranges on partitioned tables” in *Administration Guide: Implementation*

Related reference:

- “Examples of rolling in and rolling out partitioned table data” in *Administration Guide: Implementation*
- “CREATE TABLE statement” in *SQL Reference, Volume 2*
- “Command Line Processor (CLP) samples” in *Samples Topics*

Data organization schemes in DB2 and Informix databases

Table partitioning is a data organization scheme in which table data is divided across multiple storage objects called data partitions or ranges according to values in one or more table columns. Each data partition is stored separately. These storage objects can be in different table spaces, in the same table space, or a combination of both. Table data is partitioned as specified in the PARTITION BY clause of the CREATE TABLE statement. The columns used in this definition are referred to as the table partitioning key columns. DB2 table partitioning maps to the data fragmentation approach to data organization offered by Informix Dynamic Server and Informix Extended Parallel Server.

The Informix approach:

Informix supports several data organization schemes, which are called *fragmentation* in the Informix products. One of the more commonly used types of fragmentation is FRAGMENT BY EXPRESSION. This type of fragmentation works much like a CASE statement, where there is an expression associated with each fragment of the table. These expressions are checked in order to determine where to place a row.

An Informix and DB2 database system comparison:

DB2 database provides a rich set of complementary features that map directly to the Informix data organization schemes, making it relatively easy for customers to convert from the Informix syntax to the DB2 syntax. The DB2 database manager handles complicated Informix schemes using a combination of generated columns and the PARTITION BY RANGE clause of the CREATE TABLE statement. Table 23 compares data organizations schemes used in Informix and DB2 database products.

Table 23. A mapping of all Informix and DB2 data organization schemes

Data organization scheme	Informix syntax	DB2 Version 9.1 syntax
<ul style="list-style-type: none">• Informix: expression-based• DB2: table partitioning	FRAGMENT BY EXPRESSION	PARTITION BY RANGE
<ul style="list-style-type: none">• Informix: round-robin• DB2: default	FRAGMENT BY ROUND ROBIN	No syntax: DB2 database manager automatically spreads data among containers
<ul style="list-style-type: none">• Informix: range distribution• DB2: table partitioning	FRAGMENT BY RANGE	PARTITION BY RANGE
<ul style="list-style-type: none">• Informix: system defined-hash• DB2: database partitioning	FRAGMENT BY HASH	DISTRIBUTE BY HASH
<ul style="list-style-type: none">• Informix: HYBRID• DB2: database partitioning with table partitioning	FRAGMENT BY HYBRID	DISTRIBUTE BY HASH, PARTITION BY RANGE
<ul style="list-style-type: none">• Informix: n/a• DB2: Multidimensional clustering	n/a	ORGANIZE BY DIMENSION

Examples:

The following examples provide details on how to accomplish DB2 database equivalent outcomes for any Informix fragment by expression scheme.

Example 1: The following basic create table statement shows Informix fragmentation and the equivalent table partitioning syntax for a DB2 database system:

Informix syntax:

```
CREATE TABLE demo(a INT) FRAGMENT BY EXPRESSION
  a = 1 IN db1,
  a = 2 IN db2,
  a = 3 IN db3;
```

DB2 syntax:

```
CREATE TABLE demo(a INT) PARTITION BY RANGE(a)
(STARTING(1) IN db1,
STARTING(2) IN db2,
STARTING(3) ENDING(3) IN db3);
```

Informix XPS supports a two-level fragmentation scheme known as hybrid where data is spread across co-servers with one expression and within the co-server with a second expression. This allows all co-servers to be active on a query (that is, there is data on all co-servers) as well as allowing the query to take advantage of data partition elimination.

The DB2 database system achieves the equivalent organization scheme to the Informix hybrid using a combination of the DISTRIBUTE BY and PARTITION BY clauses of the CREATE TABLE statement.

Example 2: The following example shows the syntax for the combined clauses:

Informix syntax

```
CREATE TABLE demo(a INT, b INT) FRAGMENT BY HYBRID HASH(a)
EXPRESSION b = 1 IN dbs11,
b = 2 IN dbs12;
```

DB2 syntax

```
CREATE TABLE demo(a INT, b INT) IN dbs11, dbs12
DISTRIBUTE BY HASH(a),
PARTITION BY RANGE(b) (STARTING 1 ENDING 2 EVERY 1);
```

In addition, you can use multidimensional clustering to gain an extra level of data organization:

```
CREATE TABLE demo(a INT, b INT, c INT) IN dbs11, dbs12
DISTRIBUTE BY HASH(a),
PARTITION BY RANGE(b) (STARTING 1 ENDING 2 EVERY 1)
ORGANIZE BY DIMENSIONS(c);
```

Thus, all rows with the same value of column **a** are in the same database partition. All rows with the same value of column **b** are in the same table space. For a given value of **a** and **b**, all rows with the same value **c** are clustered together on disk. This approach is ideal for OLAP-type drill-down operations, because only one or several extents (blocks) in a single table space on a single database partition must be scanned to satisfy this type of query.

Table partitioning applied to common application problems:

The following sections discuss how to apply the various features of DB2 table partitioning to common application problems. In each section, particular attention is given to best practices for mapping various Informix fragmentation schemes into equivalent DB2 table partitioning schemes.

Considerations for creating simple data partition ranges:

One of the most common applications of table partitioning is to partition a large fact table based on a date key. If you need to create uniformly sized ranges of dates, consider using the automatically generated form of the CREATE TABLE syntax.

Examples:

Example 1: The following example shows the automatically generated form of the syntax:

```
CREATE TABLE orders
(
  l_orderkey DECIMAL(10,0) NOT NULL,
  l_partkey INTEGER,
  l_suppkey INTEGER,
  l_linenumber INTEGER,
  l_quantity DECIMAL(12,2),
  l_extendedprice DECIMAL(12,2),
  l_discount DECIMAL(12,2),
  l_tax DECIMAL(12,2),
  l_returnflag CHAR(1),
  l_linestatus CHAR(1),
  l_shipdate DATE,
  l_commitdate DATE,
  l_receiptdate DATE,
  l_shipinstruct CHAR(25),
  l_shipmode CHAR(10),
  l_comment VARCHAR(44)
  PARTITION BY RANGE(l_shipdate)
  (STARTING '1/1/1992' ENDING '12/31/1993' EVERY 1 MONTH);
```

This creates 24 ranges, one for each month in 1992-1993. Attempting to insert a row with `l_shipdate` outside of that range results in an error.

Example 2: Compare the preceding example to the following Informix syntax:

```
create table orders
(
  l_orderkey decimal(10,0) not null,
  l_partkey integer,
  l_suppkey integer,
  l_linenumber integer,
  l_quantity decimal(12,2),
  l_extendedprice decimal(12,2),
  l_discount decimal(12,2),
  l_tax decimal(12,2),
  l_returnflag char(1),
  l_linestatus char(1),
  l_shipdate date,
  l_commitdate date,
  l_receiptdate date,
  l_shipinstruct char(25),
  l_shipmode char(10),
  l_comment varchar(44)
) fragment by expression
l_shipdate < '1992-02-01' in ldfs1,
l_shipdate >= '1992-02-01' and l_shipdate < '1992-03-01' in ldfs2,
l_shipdate >= '1992-03-01' and l_shipdate < '1992-04-01' in ldfs3,
l_shipdate >= '1992-04-01' and l_shipdate < '1992-05-01' in ldfs4,
l_shipdate >= '1992-05-01' and l_shipdate < '1992-06-01' in ldfs5,
l_shipdate >= '1992-06-01' and l_shipdate < '1992-07-01' in ldfs6,
l_shipdate >= '1992-07-01' and l_shipdate < '1992-08-01' in ldfs7,
l_shipdate >= '1992-08-01' and l_shipdate < '1992-09-01' in ldfs8,
l_shipdate >= '1992-09-01' and l_shipdate < '1992-10-01' in ldfs9,
l_shipdate >= '1992-10-01' and l_shipdate < '1992-11-01' in ldfs10,
l_shipdate >= '1992-11-01' and l_shipdate < '1992-12-01' in ldfs11,
l_shipdate >= '1992-12-01' and l_shipdate < '1993-01-01' in ldfs12,
l_shipdate >= '1993-01-01' and l_shipdate < '1993-02-01' in ldfs13,
l_shipdate >= '1993-02-01' and l_shipdate < '1993-03-01' in ldfs14,
l_shipdate >= '1993-03-01' and l_shipdate < '1993-04-01' in ldfs15,
l_shipdate >= '1993-04-01' and l_shipdate < '1993-05-01' in ldfs16,
l_shipdate >= '1993-05-01' and l_shipdate < '1993-06-01' in ldfs17,
l_shipdate >= '1993-06-01' and l_shipdate < '1993-07-01' in ldfs18,
l_shipdate >= '1993-07-01' and l_shipdate < '1993-08-01' in ldfs19,
```



```

l_shipdate >= '1993-08-01' and l_shipdate < '1993-09-01' in ldfs20,
l_shipdate >= '1993-09-01' and l_shipdate < '1993-10-01' in ldfs21,
l_shipdate >= '1993-10-01' and l_shipdate < '1993-11-01' in ldfs22,
l_shipdate >= '1993-11-01' and l_shipdate < '1993-12-01' in ldfs23,
l_shipdate >= '1993-12-01' and l_shipdate < '1994-01-01' in ldfs24,
l_shipdate >= '1994-01-01' in ldfs25;

```

Notice that the Informix syntax provides an open ended range at the top and bottom to catch dates that are not in the expected range. The DB2 syntax can be modified to match the Informix syntax by adding ranges that make use of MINVALUE and MAXVALUE.

Example 3: The following example modifies Example 1 to mirror the Informix syntax::

```

CREATE TABLE orders
(
  l_orderkey DECIMAL(10,0) NOT NULL,
  l_partkey INTEGER,
  l_suppkey INTEGER,
  l_linenumber INTEGER,
  l_quantity DECIMAL(12,2),
  l_extendedprice DECIMAL(12,2),
  l_discount DECIMAL(12,2),
  l_tax DECIMAL(12,2),
  l_returnflag CHAR(1),
  l_linestatus CHAR(1),
  l_shipdate DATE,
  l_commitdate DATE,
  l_receiptdate DATE,
  l_shipinstruct CHAR(25),
  l_shipmode CHAR(10),
  l_comment VARCHAR(44)
) PARTITION BY RANGE(l_shipdate)
(STARTING MINVALUE,
 STARTING '1/1/1992' ENDING '12/31/1993' EVERY 1 MONTH,
 ENDING MAXVALUE);

```

This technique allows any date to be inserted into the table.

Partition by expression using generated columns:

Although DB2 database does not directly support partitioning by expression, partitioning on a generated column is supported, making it possible to achieve the same result.

Consider the following usage guidelines before deciding whether to use this approach:

- The generated column is a real column that occupies physical disk space. Tables that make use of a generated column can be slightly larger.
- Altering the generated column expression for the column on which a partitioned table is partitioned is not supported. Attempting to do so will result in the message SQL0190. Adding a new data partition to a table that uses generated columns in the manner described in the next section generally requires you to alter the expression that defines the generated column. Altering the expression that defines a generated column is not currently supported.
- There are limitations on when you can apply data partition elimination when a table uses generated columns.

Examples:

Example 1: The following uses the Informix syntax, where it is appropriate to use generated columns. In this example, the column to be partitioned on holds Canadian provinces and territories. Because the list of provinces is unlikely to change, the generated column expression is unlikely to change.

```
CREATE TABLE customer (
  cust_id INT,
  cust_prov CHAR(2))
FRAGMENT BY EXPRESSION
  cust_prov = "AB" IN dbspace_ab
  cust_prov = "BC" IN dbspace_bc
  cust_prov = "MB" IN dbspace_mb
  ...
  cust_prov = "YT" IN dbspace_yt
REMAINDER IN dbspace_remainder;
```

Example 2: In this example, the DB2 table is partitioned using a generated column:

```
CREATE TABLE customer (
  cust_id INT,
  cust_prov CHAR(2),
  cust_prov_gen GENERATED ALWAYS AS (CASE
    WHEN cust_prov = 'AB' THEN 1
    WHEN cust_prov = 'BC' THEN 2
    WHEN cust_prov = 'MB' THEN 3
    ...
    WHEN cust_prov = 'YT' THEN 13
  ELSE 14 END))
IN tbspace_ab, tbspace_bc, tbspace_mb, .... tbspace_remainder
PARTITION BY RANGE (cust_prov_gen)
(STARTING 1 ENDING 14 EVERY 1);
```

Here the expressions within the case statement match the corresponding expressions in the FRAGMENT BY EXPRESSION clause. The case statement maps each original expression to a number, which is stored in the generated column (cust_prov_gen in this example). This column is a real column stored on disk, so the table could occupy slightly more space than would be necessary if DB2 supported partition by expression directly. This example uses the short form of the syntax. Therefore, the table spaces in which to place the data partitions must be listed in the IN clause of the CREATE TABLE statement. Using the long form of the syntax requires a separate IN clause for each data partition.

Note: This technique can be applied to any FRAGMENT BY EXPRESSION clause.

Related concepts:

- “Data partitions” on page 92
- “Partitioned database environments” on page 41
- “Partitioned tables” on page 104
- “Table partitioning” on page 93
- “Optimization strategies for partitioned tables” in *Performance Guide*
- “Understanding clustering index behavior on partitioned tables” in *Performance Guide*
- “Understanding index behavior on partitioned tables” in *Performance Guide*
- “Attributes of detached data partitions” in *Administration Guide: Implementation*
- “Database partitioning across multiple database partitions” in *SQL Reference, Volume 1*
- “Large object behavior in partitioned tables” in *SQL Reference, Volume 1*

- “Partitioned materialized query table behavior” in *Administration Guide: Implementation*

Related tasks:

- “Adding data partitions to partitioned tables” in *Administration Guide: Implementation*
- “Altering partitioned tables” in *Administration Guide: Implementation*
- “Creating partitioned tables” in *Administration Guide: Implementation*
- “Dropping a data partition” in *Administration Guide: Implementation*
- “Enabling database partitioning in a database” in *Administration Guide: Implementation*
- “Approaches to migrating existing tables and views to partitioned tables” in *Administration Guide: Implementation*
- “Attaching a data partition” in *Administration Guide: Implementation*
- “Detaching a data partition” in *Administration Guide: Implementation*
- “Rotating data in a partitioned table” in *Administration Guide: Implementation*
- “Approaches to defining ranges on partitioned tables” in *Administration Guide: Implementation*

Related reference:

- “Examples of rolling in and rolling out partitioned table data” in *Administration Guide: Implementation*
- “CREATE TABLE statement” in *SQL Reference, Volume 2*
- “Guidelines and restrictions on altering partitioned tables with attached or detached data partitions” in *Administration Guide: Implementation*

Replicated materialized query tables

A *materialized query table* is a table that is defined by a query that is also used to determine the data in the table. Materialized query tables can be used to improve the performance of queries. If DB2 Database for Linux, UNIX, and Windows determines that a portion of a query could be resolved using a materialized query table, the query may be rewritten by the database manager to use the materialized query table.

In a partitioned database environment, you can replicate materialized query tables and use them to improve query performance. A *replicated materialized query table* is based on a table that may have been created in a single-partition database partition group, but that you want replicated across all of the database partitions in another database partition group. To create the replicated materialized query table, invoke the CREATE TABLE statement with the REPLICATED keyword.

By using replicated materialized query tables, you can obtain collocation between tables that are not typically collocated. Replicated materialized query tables are particularly useful for joins in which you have a large fact table and small dimension tables. To minimize the extra storage required, as well as the impact of having to update every replica, tables that are to be replicated should be small and updated infrequently.

Note: You should also consider replicating larger tables that are updated infrequently: the one-time cost of replication is offset by the performance benefits that can be obtained through collocation.

By specifying a suitable predicate in the subselect clause used to define the replicated table, you can replicate selected columns, selected rows, or both.

Related concepts:

- “Database partition group design” on page 87
- “The Design Advisor” in *Performance Guide*

Related tasks:

- “Creating a materialized query table” in *Administration Guide: Implementation*

Related reference:

- “CREATE TABLE statement” in *SQL Reference, Volume 2*

Table space design

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). For improved performance, each container should use a different disk. Figure 27 illustrates the relationship between tables and table spaces within a database, and the containers associated with that database.

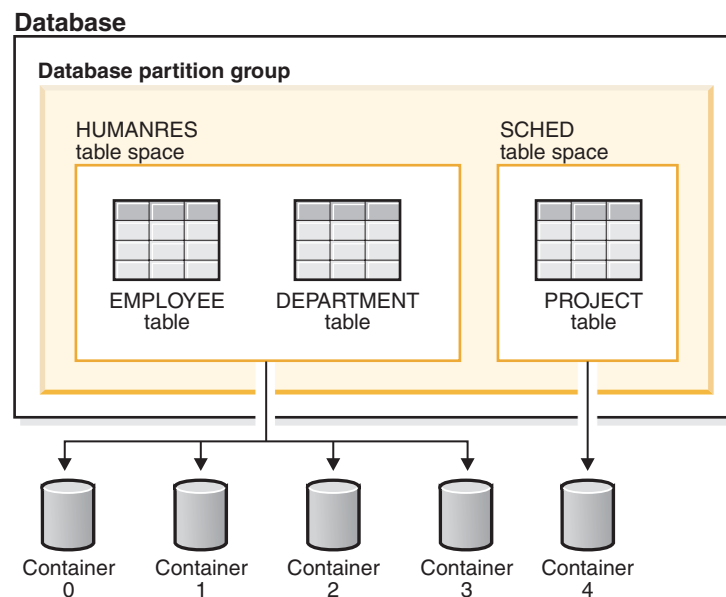


Figure 27. Table spaces and tables in a database

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 28 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.

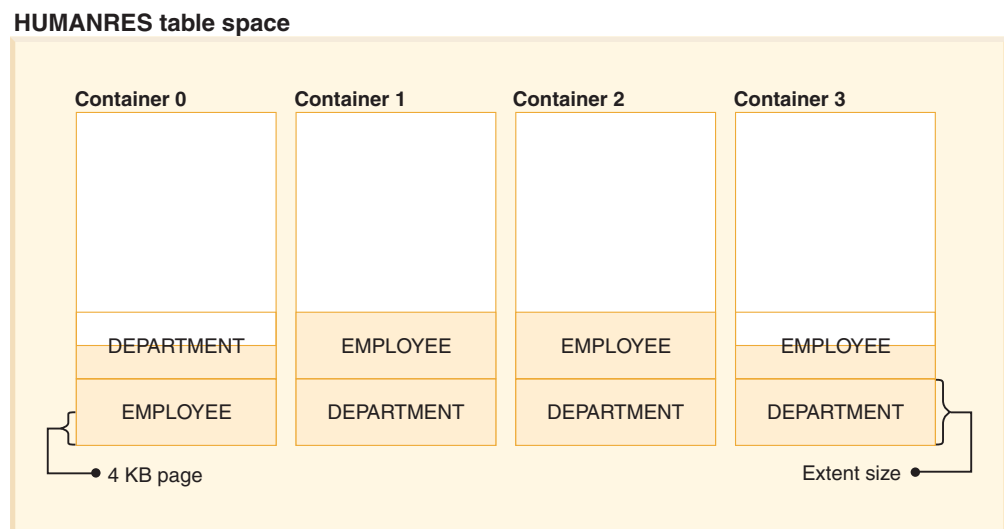


Figure 28. Containers and extents in a table space

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 base 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.

In a partitioned database environment, the catalog node will contain all three default table spaces, and the other database partitions will each contain only TEMPSPACE1 and USERSPACE1.

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.

Related concepts:

- "Catalog table space design" on page 163
- "Comparison of SMS and DMS table spaces" on page 140
- "Database managed space" on page 120
- "Extent size" on page 144
- "Relationship between table spaces and buffer pools" on page 145

- “Relationship between table spaces and database partition groups” on page 146
- “System managed space” on page 117
- “SYSTOOLSPACE and SYSTOOLSTMPSPACE table spaces” on page 115
- “Temporary table space design” on page 161
- “Workload considerations in table space design” on page 143
- “Table space disk I/O” on page 141
- “Table spaces and other storage structures” in *SQL Reference, Volume 1*

Related tasks:

- “Optimizing table space performance when data is on RAID devices” on page 164
- “Creating a table space” in *Administration Guide: Implementation*

Related reference:

- “CREATE TABLE statement” in *SQL Reference, Volume 2*
- “CREATE TABLESPACE statement” in *SQL Reference, Volume 2*

SYSTOOLSPACE and SYSTOOLSTMPSPACE table spaces

The SYSTOOLSPACE table space is a user data table space used by the DB2 administration tools and some SQL administrative routines for storing historical data and configuration information. The following tools and SQL administrative routines use the SYSTOOLSPACE table space:

- Design advisor
- Alter table notebook
- Configure Automatic Maintenance wizard
- Storage management tool
- db2look command
- Automatic statistics collection (including the Statistics Collection Required health indicator)
- Automatic reorganization (including the Reorganization Required health indicator)
- GET_DBSIZE_INFO stored procedure
- ADMIN_COPY_SCHEMA stored procedure
- ADMIN_DROP_SCHEMA stored procedure
- SYSINSTALLOBJECTS stored procedure
- ALTOBJ stored procedure

The SYSTOOLSPACE table space is created the first time any of the above are used (except for DB2LOOK, ALTOBJ, ADMIN_COPY_SCHEMA and ADMIN_DROP_SCHEMA).

The SYSTOOLSTMPSPACE table space is a user temporary table space used by the REORGCHK_TB_STATS, REORGCHK_IX_STATS and the ADMIN_CMD stored procedures for storing temporary data. The SYSTOOLSTMPSPACE table space will be created the first time any of these stored procedures is invoked (except for ADMIN_CMD).

Notes:

1. If the DB2 registry variable DB2_WORKLOAD is set to SAP, neither the SYSTOOLSPACE nor the SYSTOOLSTMPSPACE will be created automatically.
2. The Reorganization Required and Statistics Collection Required health indicators and the Health Monitor are enabled by default on all new databases. These two health indicators are evaluated by the Health Monitor approximately every two hours. This means that the SYSTOOLSPACE and SYSTOOLSTMPSPACE table spaces are created automatically for new databases after they have been active for two hours unless the health monitor or these health indicators are explicitly disabled.
3. The automatic statistics collection feature is enabled by default on all new databases. This feature is evaluated approximately every two hours. This means that the SYSTOOLSPACE and SYSTOOLSTMPSPACE table spaces are created automatically for new databases after they have been active for two hours unless the automatic statistic collection feature is explicitly disabled.

If the default definition for either table space is not preferred, you can create the table spaces manually (or drop and recreate them if they have already been created automatically). The table space definitions may vary (for example, you can use a DMS or SMS table space, or you can enable or disable automatic storage), however the table spaces must be created in the IBMCATGROUP database partition group. If you attempt to create them in any other database partition group, error SQL1258N will be returned.

Following is an example of how to create the SYSTOOLSPACE and SYSTOOLSTMPSPACE table spaces manually. This example uses the same definitions that are used when table spaces are created automatically:

If the database is using automatic storage:

```
CREATE TABLESPACE SYSTOOLSPACE IN IBMCATGROUP
  MANAGED BY AUTOMATIC STORAGE
  EXTENTSIZE 4

CREATE USER TEMPORARY TABLESPACE SYSTOOLSTMPSPACE IN IBMCATGROUP
  MANAGED BY AUTOMATIC STORAGE
  EXTENTSIZE 4
```

If the database is not using automatic storage:

```
CREATE TABLESPACE SYSTOOLSPACE IN IBMCATGROUP
  MANAGED BY DATABASE USING ( FILE 'SYSTOOLSPACE' 32 M )
  AUTORESIZE YES
  EXTENTSIZE 4

CREATE USER TEMPORARY TABLESPACE SYSTOOLSTMPSPACE IN IBMCATGROUP
  MANAGED BY SYSTEM USING ( 'SYSTOOLSTMPSPACE' )
  EXTENTSIZE 4
```

By default, the use of SYSTOOLSTMPSPACE will be granted to the PUBLIC group as long as the database is not created using restricted access.

Related concepts:

- “The Design Advisor” in *Performance Guide*
- “Automatic reorganization” on page 32

Related tasks:

- “Using automatic statistics collection” in *Performance Guide*

- “Altering a table” in *Administration Guide: Implementation*

Related reference:

- “ADMIN_COPY_SCHEMA procedure – Copy a specific schema and its objects” in *Administrative SQL Routines and Views*
- “ADMIN_DROP_SCHEMA procedure – Drop a specific schema and its objects” in *Administrative SQL Routines and Views*
- “ALTOBJ procedure” in *Administrative SQL Routines and Views*
- “db2look - DB2 statistics and DDL extraction tool command” in *Command Reference*
- “Health indicators summary” in *System Monitor Guide and Reference*
- “Storage management view” on page 146
- “GET_DBSIZE_INFO procedure” in *Administrative SQL Routines and Views*
- “SYSINSTALLOBJECTS procedure” in *Administrative SQL Routines and Views*

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. The user decides on the location of the files, DB2 Database for Linux, UNIX, and Windows 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.

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 (SQL00001.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 SQL00001.DAT in the next container. This process continues until all of the containers contain SQL00001.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 is also used for index (SQLnnnnn.INX), long field (SQLnnnnn.LF), LOB (SQLnnnnn.LB and SQLnnnnn.LBA) and XML (SQLnnnnn.XDA) files.

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 (SQL00001.DAT in the above example) and factoring into account the number of containers. Containers are numbered sequentially, starting at 0.

Related concepts:

- “Comparison of SMS and DMS table spaces” on page 140
- “Database managed space” on page 120
- “Table space design” on page 112

Related reference:

- “db2empfa - Enable multipage file allocation command” in *Command Reference*
- “multipage_alloc - Multipage file allocation enabled configuration parameter” in *Performance Guide*

SMS table spaces

System Managed Space (SMS) table spaces store data in operating system files. 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.

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. Prior to version 8.2, the default setting of the configuration parameter was NO which caused one page to be allocated at a time. This default could be changed with the db2empfa tool which allows you to enable multipage file allocation. When you run db2empfa, the multipage_alloc database configuration parameter is set to YES.

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.

Related concepts:

- “Comparison of SMS and DMS table spaces” on page 140
- “Table space design” on page 112

Related tasks:

- “Adding a container to an SMS table space on a database partition” in *Administration Guide: Implementation*

Related reference:

- “db2empfa - Enable multipage file allocation command” in *Command Reference*
- “multipage_alloc - Multipage file allocation enabled configuration parameter” in *Performance Guide*

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 DB2 Database for Linux, UNIX, and Windows. The database administrator decides which devices and files to use, and DB2 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.
- The maximum size of a regular table space is 512 GB for 32 KB pages. The maximum size of a large table space is 16TB. See SQL and XQuery 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. No operational difference exists between a file and a device; however, a file can be less efficient because of the run-time overhead associated with the file system. Files are useful when:
 - Devices are not directly supported
 - A device is not available
 - Maximum performance is not required
 - You do not want to set up devices.
- 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 is one exception to this general statement regarding contiguous placement of pages in storage. 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 (SMP).

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 Space Map Pages (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 29 shows how extents are allocated in a logical table space address map.




Figure 29. Logical table space address map

Related concepts:

- “Comparison of SMS and DMS table spaces” on page 140
- “How containers are added and extended in DMS table spaces” on page 129
- “System managed space” on page 117
- “Table space design” on page 112
- “Table space maps” on page 125

DMS table spaces

With database-managed space (DMS) table spaces, the database manager controls the storage space. A list of devices or files is selected to belong to a table space when the DMS table space is defined. The space on those devices or files is managed by the DB2 database manager. As with SMS table spaces and containers, DMS table spaces and the database manager use striping by extent to ensure an even distribution of data across all containers.

DMS table spaces differ from SMS table spaces in that for DMS table spaces, space is allocated when the table space is created and not allocated when needed.

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 may, or may not, 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.

Note: Like SMS table spaces, DMS file containers can take advantage of file-system prefetching and caching. However, DMS table spaces cannot.

There is one exception to this general statement regarding contiguous placement of pages in storage. 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.

Unlike SMS table spaces, the containers that make up a DMS table space do not need to be close to being equal in their capacity. However, it is recommended that the containers are equal, or close to being equal, in their capacity. Also, if any container is full, any available free space from other containers can be used in a DMS table space.

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 other 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 (SMP).

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 Space Map Pages (SMP) for the table space. SMP extents are spread at regular intervals throughout the table space. Each SMP extent is simply 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.

Related concepts:

- “Comparison of SMS and DMS table spaces” on page 140
- “Database directories and files” on page 73
- “Database managed space” on page 120
- “DMS device considerations” on page 124
- “SMS table spaces” on page 119
- “Table space design” on page 112

Related tasks:

- “Adding a container to a DMS table space” in *Administration Guide: Implementation*

Related reference:

- “CREATE TABLESPACE statement” in *SQL Reference, Volume 2*

DMS device considerations

If you use Database Managed Storage (DMS) device containers for table spaces, consider the following factors for effective administration:

- **File system caching**

File system caching is performed as follows:

- For DMS file containers (and all SMS containers), the operating system might cache pages in the file system cache
- For DMS device container table spaces, the operating system does not cache pages in the file system cache.

Note: On Windows, the registry variable DB2NTNOCACHE specifies whether or not DB2 will open 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 sortheap can be increased.

- **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 storage (SMS) or database managed storage (DMS) file containers, file system caching above can eliminate I/O that would otherwise have been required.

Table spaces using database managed storage (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.

Related concepts:

- “DMS table spaces” on page 123
- “SMS table spaces” on page 119
- “Database directories and files” on page 73

Table space maps

A table space map is DB2 V9.1’s internal representation of a DMS table space that describes the logical to physical conversion of page locations in a table space. The following information describes why a table space map is useful, and where the information in a table space map comes from.

In a DB2 Database for Linux, UNIX, and Windows 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.

The following figure shows the logical address map for a DMS table space.

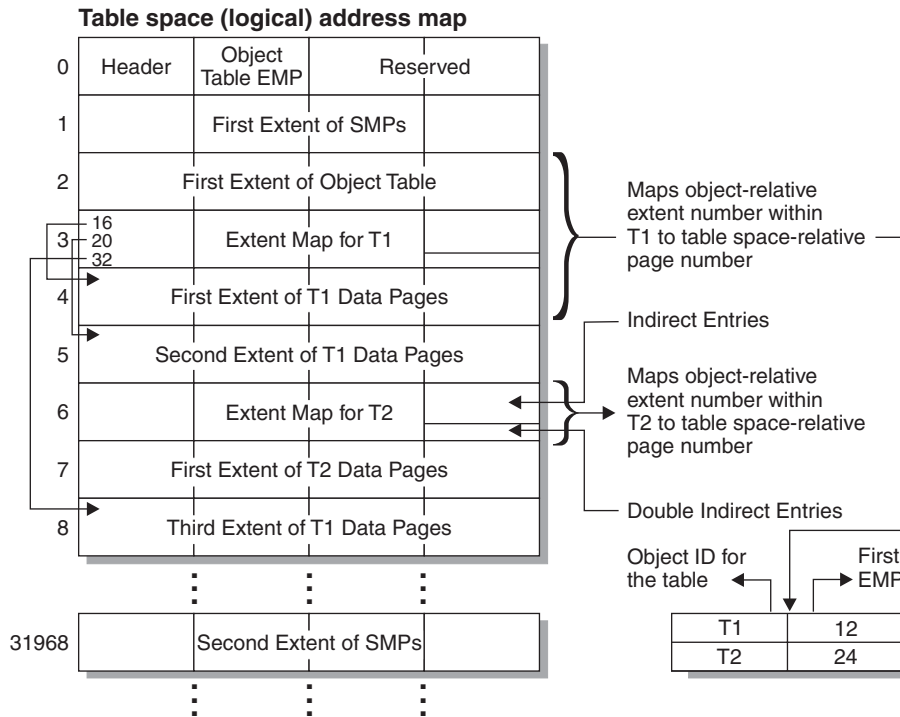


Figure 30. DMS table spaces

Within the table space address map there are two types of map pages: extent map pages (EMP) and space map pages (SMP).

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-related extent number to a table space-related page number where the object extent is located. Direct EMP entries directly map object-related addresses to table space-related 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 31.

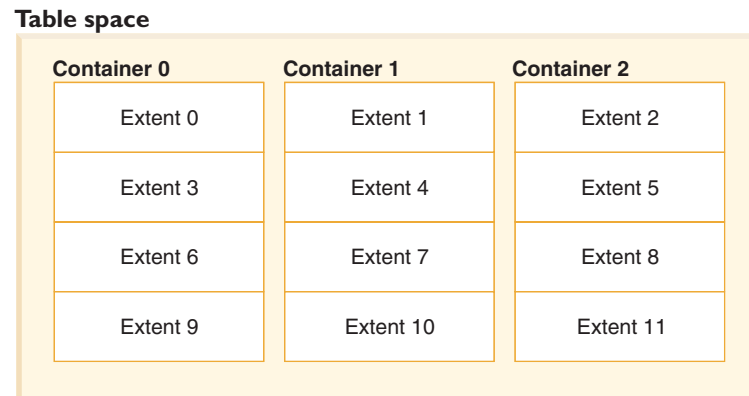


Figure 31. Table space with three containers and 12 extents

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 Number	Stripe Set	Stripe Offset	Max Extent	Max Page	Start Stripe	End Stripe	Adj.	Containers
[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 32, in which each vertical line corresponds to a container, each horizontal line is called a *stripe*, and each cell number corresponds to an extent.

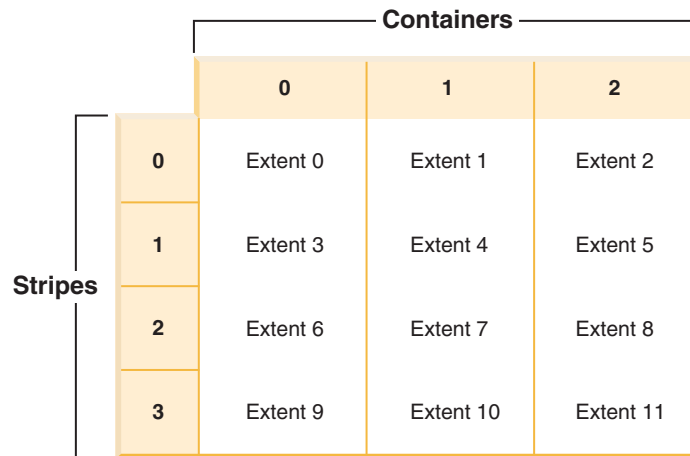


Figure 32. Table space with three containers and 12 extents, with stripes highlighted

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 33.

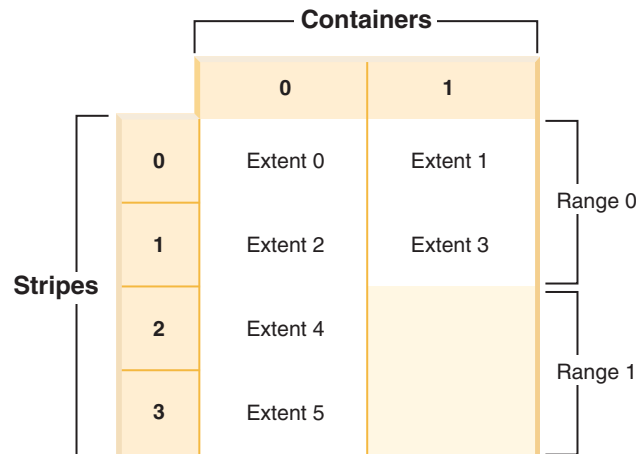


Figure 33. Table space with two containers, with ranges highlighted

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	Max Extent	Max Page	Start Stripe	End Stripe	Adj.	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.

Related concepts:

- “Snapshot monitor” in *System Monitor Guide and Reference*
- “Database managed space” on page 120
- “How containers are added and extended in DMS table spaces” on page 129
- “How containers are dropped and reduced in DMS table spaces” on page 137

Related reference:

- “GET SNAPSHOT command” in *Command Reference*

How containers are added and extended in DMS table spaces

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” on page 130.)

The ALTER TABLESPACE statement lets you add a container to an existing table space or extend a container to increase its storage capacity.

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” on page 131). 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” on page 133.) 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 34 on page 131.

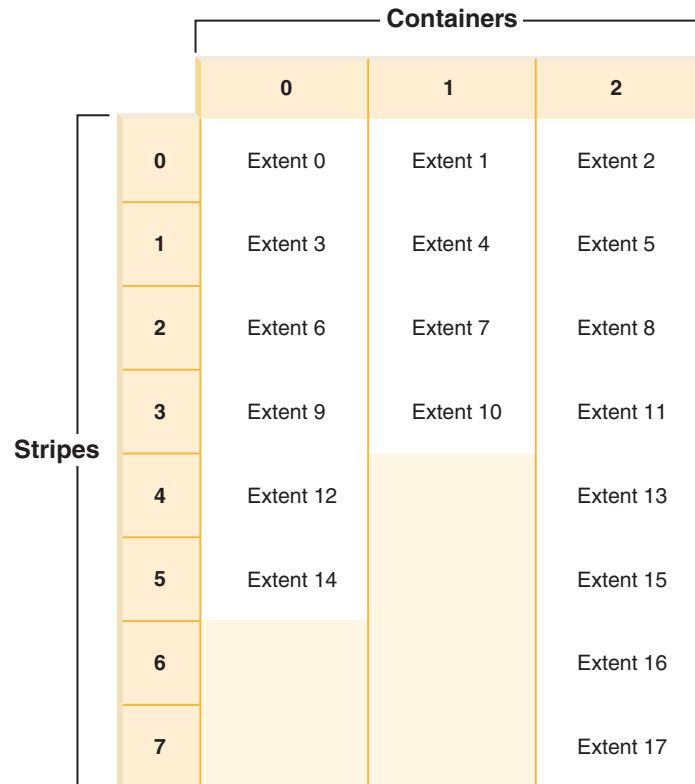


Figure 34. Table space with three containers and 18 extents

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 35 on page 132.

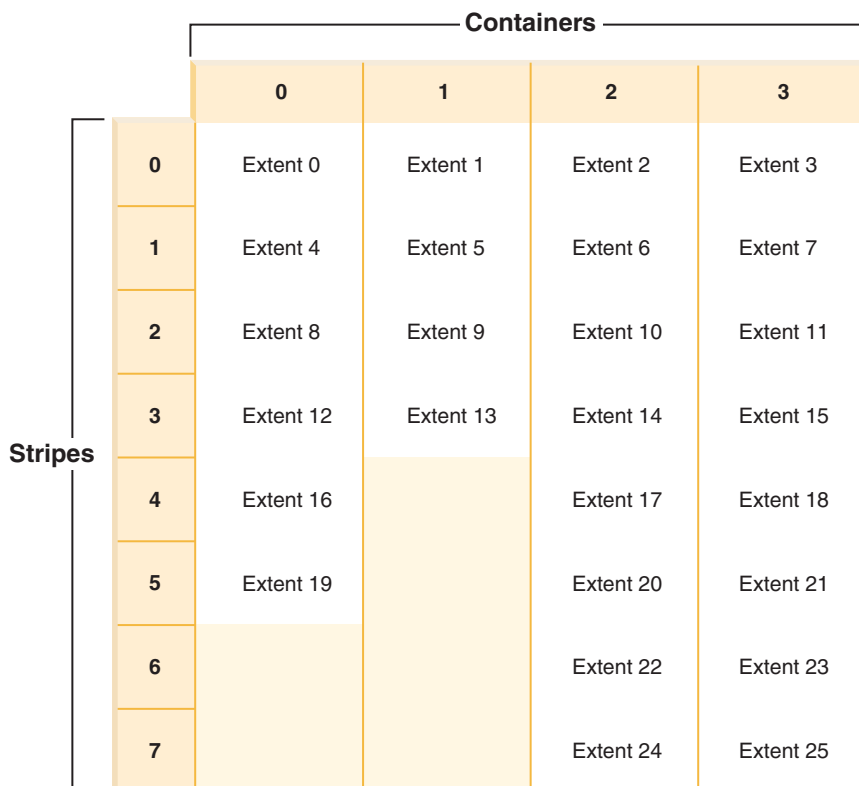


Figure 35. Table space with four containers and 26 extents

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 36.

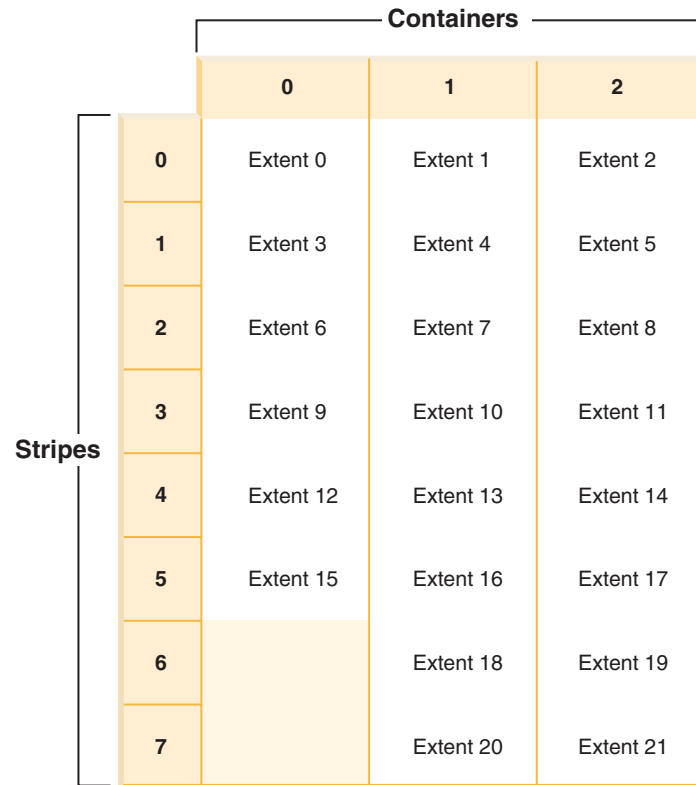


Figure 36. Table space with three containers and 22 extents

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” on page 130. 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 37 on page 134.)

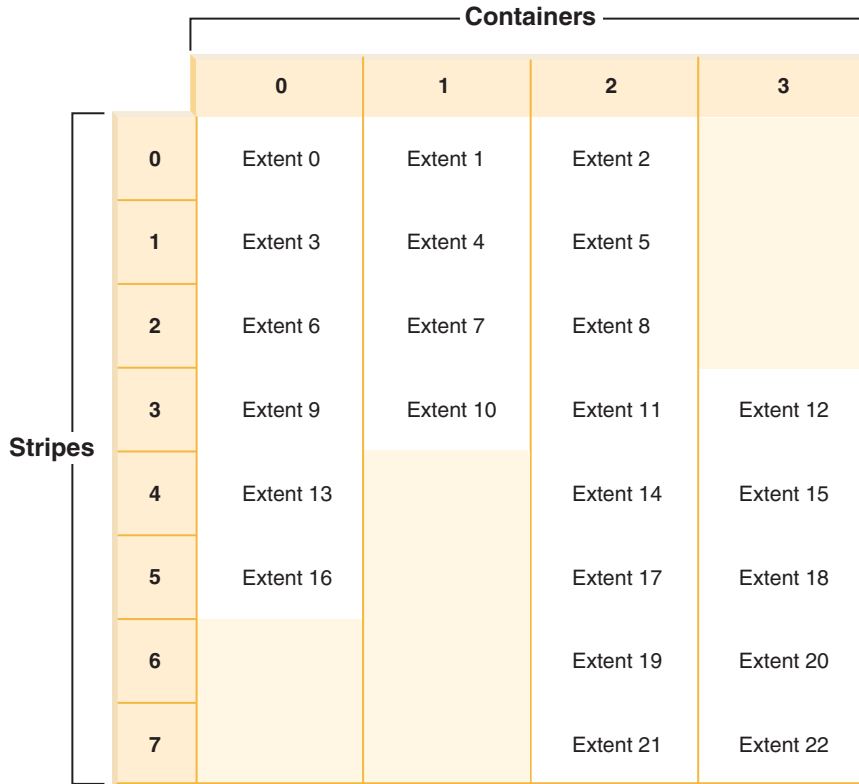


Figure 37. Table space with four containers and 23 extents

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` option on the `ALTER TABLESPACE` statement. To add containers and rebalance the data, use the `ADD` option 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` option 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` option 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 38.

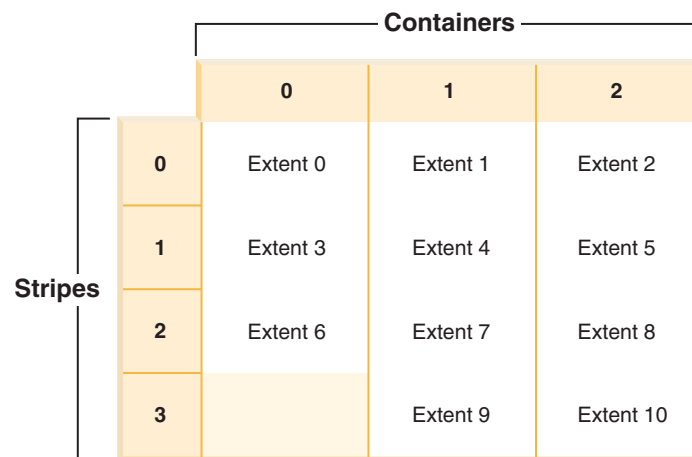


Figure 38. Table space with three containers and 11 extents

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` option, 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 39.

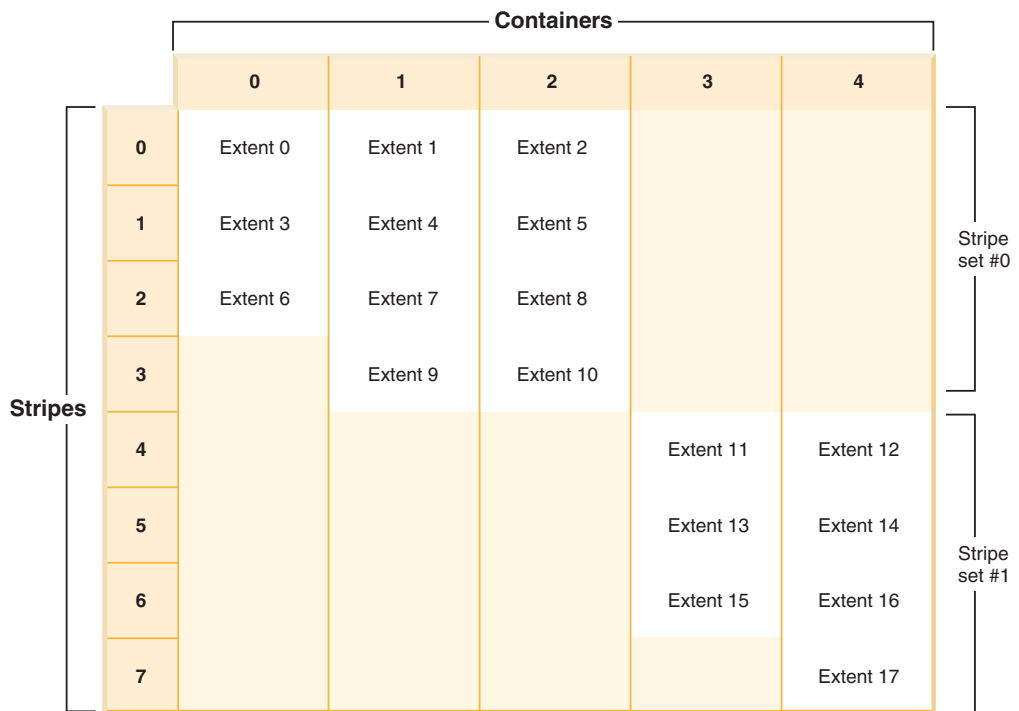


Figure 39. Table space with two stripe sets

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` option 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.

DB2 Database for Linux, UNIX, and Windows 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.

Related concepts:

- “Table space maps” on page 125

Related tasks:

- “Adding a container to a DMS table space” in *Administration Guide: Implementation*
- “Modifying containers in a DMS table space” in *Administration Guide: Implementation*

Related reference:

- “Table space activity monitor elements” in *System Monitor Guide and Reference*
- “ALTER TABLESPACE statement” in *SQL Reference, Volume 2*
- “GET SNAPSHOT command” in *Command Reference*

How containers are dropped and reduced in DMS table spaces

With a DMS table space, you can drop a container from the table space or reduce the size of a container. You use the ALTER TABLESPACE statement to accomplish this.

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 40.

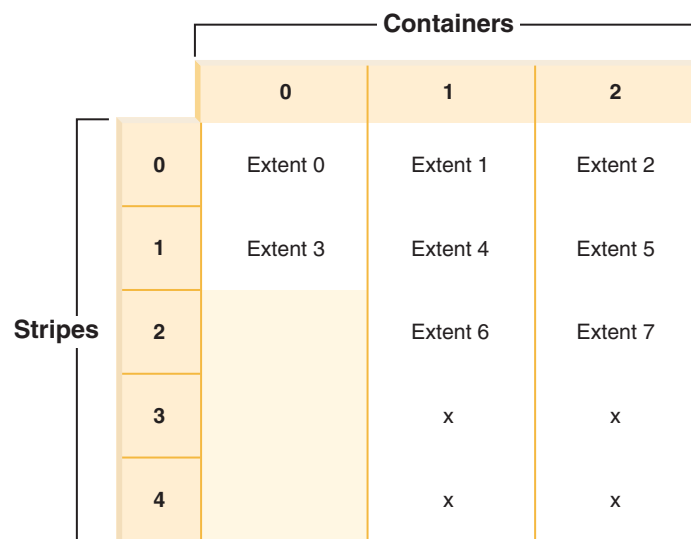


Figure 40. Table space with 12 extents, including four extents with no data

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 41.




Figure 41. 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.

Related concepts:

- “Table space maps” on page 125

Related tasks:

- “Modifying containers in a DMS table space” in *Administration Guide: Implementation*

Related reference:

- “ALTER TABLESPACE statement” in *SQL Reference, Volume 2*
- “GET SNAPSHOT command” in *Command Reference*
- “Table space activity monitor elements” in *System Monitor Guide and Reference*

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 distributed 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 and 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, you could have a table with 64 GB of regular table data, 64 GB of index data and 2 TB of long data. If you are using 8 KB pages, the table data and the index data can be as much as 128 GB. If you are using 16 KB pages, it can be as much as 256 GB. If you are using 32 KB pages, the table data and the index data can be as much as 512 GB.

- Indexes created on distributed 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.
- If all table data is in a single table space, a table space can be dropped and redefined with less overhead than dropping and redefining a table.
- In general, a well-tuned set of DMS table spaces will outperform SMS table spaces.

Notes:

1. On the Solaris operating system, DMS table spaces with raw devices are strongly recommended for performance-critical workloads.
2. For performance-sensitive applications, particularly those involving a large number of insert operations, it is recommended 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 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.

Related concepts:

- “Database managed space” on page 120
- “System managed space” on page 117
- “Table space design” on page 112

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. Following are concepts that you should understand before considering further the issues surrounding 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, DB2 Database for Linux, UNIX, and Windows 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. (The default value for all table spaces in the database is set by the *dft_prefetch_sz* database configuration parameter.) The PREFETCHSIZE parameter tells DB2 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, DB2 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 DB2 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 DB2 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.

Related concepts:

- “Prefetching data into the buffer pool” in *Performance Guide*
- “Table space design” on page 112

Related reference:

- “ALTER TABLESPACE statement” in *SQL Reference, Volume 2*
- “CREATE TABLESPACE statement” in *SQL Reference, Volume 2*

Workload considerations in table space design

The primary type of workload being managed by DB2 Database for Linux, UNIX, and Windows 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. 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 `chngpgs_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. This allows DB2 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. There is, however, an exception to this. If your row size is smaller than:

$$\text{pagesize} / 255$$

there will be wasted space on each page (there is a maximum of 255 rows per 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.
- 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.

Related concepts:

- “Database managed space” on page 120
- “System managed space” on page 117

Related reference:

- “ALTER TABLESPACE statement” in *SQL Reference, Volume 2*
- “CREATE TABLESPACE statement” in *SQL Reference, Volume 2*
- “SQL and XQuery limits” in *SQL Reference, Volume 1*
- “chnpggs_thresh - Changed pages threshold configuration parameter” in *Performance Guide*

Extent size

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 object. This is where long field data, if your table has one or more LONG columns, is stored.
- Two LOB 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 tables.

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.) So, 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.

Otherwise, 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.

Related concepts:

- “Table space design” on page 112

Related reference:

- “CREATE TABLESPACE statement” in *SQL Reference, Volume 2*
- “db2empfa - Enable multipage file allocation command” in *Command Reference*
- “multipage_alloc - Multipage file allocation enabled configuration parameter” in *Performance Guide*

Relationship between table spaces and buffer pools

Each table space is associated with a specific buffer pool. The default buffer pool is IBMDEFAULTBP. If another buffer pool is to be associated with a table space, the buffer pool must exist (it is defined with the CREATE BUFFERPOOL statement), it must have the same page size, and the association is defined when the table space is created (using the CREATE TABLESPACE statement). The association between the table space and the buffer pool can be changed 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.

Note: If you have determined that a page size of 8 KB, 16 KB, or 32 KB is required by your database, each table space with one of these page sizes must be mapped to a buffer pool with the same page size.

The storage required for all the buffer pools must be available to the database manager when the database is started. If DB2 Database for Linux, UNIX, and Windows is unable to obtain the required storage, the database manager will start up with default buffer pools (one each of 4 KB, 8 KB, 16 KB, and 32 KB page sizes), and issue a warning.

In a partitioned database environment, you can create a buffer pool of the same size for all database partitions in the database. You can also create buffer pools of different sizes on different database partitions.

Related concepts:

- “Table spaces and other storage structures” in *SQL Reference, Volume 1*

Related reference:

- “ALTER BUFFERPOOL statement” in *SQL Reference, Volume 2*
- “ALTER TABLESPACE statement” in *SQL Reference, Volume 2*
- “CREATE BUFFERPOOL statement” in *SQL Reference, Volume 2*
- “CREATE TABLESPACE statement” in *SQL Reference, Volume 2*

Relationship between 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.

Related concepts:

- “Table spaces and other storage structures” in *SQL Reference, Volume 1*
- “Database partition groups” on page 85
- “Table space design” on page 112

Related reference:

- “CREATE DATABASE PARTITION GROUP statement” in *SQL Reference, Volume 2*
- “CREATE TABLESPACE statement” in *SQL Reference, Volume 2*

Storage management view

Use the Storage Management view to monitor the storage state of a partitioned database. The Storage Management view is the graphical interface to the Storage Management tool. In the Storage Management view, you can take storage snapshots for a database, a database partition group, or a table space. When a table space snapshot is taken, statistical information is collected from the system catalogs and database monitor for tables, indexes, and containers defined under the scope of the given table space. When a database or database partition group snapshot is taken, statistical information is collected for all the table spaces defined in the given database or database partition group. When a database snapshot is taken,

statistical information is collected for all the database partition groups within the database. Different types of storage snapshots can be used to help you monitor different aspects of storage:

- Space usage can be monitored through snapshots of table spaces.
- On partitioned databases only: Data skew (database distribution) can be monitored best through snapshots of database partition groups.
- Cluster ratio of indexes can be captured through both database partition group snapshots and table space snapshots. The cluster ratio of indexes is presented through the detail view of the index folder.

The Storage Management view also enables you to set thresholds for data skew, space usage, and index cluster ratio. If a target object exceeds a specified threshold, the icons beside the object and its parent object in the Storage Management view are marked with a warning flag or an alarm flag.

Note: You can only set data skew thresholds for partitioned databases.

Use the Storage Management launchpad to guide you through the tasks necessary to set up the Storage Management tool. The Storage Management tool provides you with the ability to manage the storage of a specific database or database partition over the long term. It also allows you to capture data distribution snapshots and to view storage history. Three stored procedure functions are automatically created for the storage management tool when the database is created: `SYSPROC.CREATE_STORAGEEMGMT_TABLES`, `SYSPROC.DROP_STORAGEEMGMT_TABLES`, and `SYSPROC.CAPTURE_STORAGEEMGMT_INFO`. Their respective packages are bound on demand.

Note: You can open the Storage Management Setup launchpad from a database, database partition group, or table space object in the Control Center. The launchpad will lead you through the one-time-only setup process for using the Storage Management tool. After you have captured a snapshot for the selected object or its parent object using the Storage Management Setup launchpad, you will be able to open the Storage Management view.

Related reference:

- “`CAPTURE_STORAGEEMGMT_INFO` procedure – Retrieve storage-related information for a given root object” in *Administrative SQL Routines and Views*
- “`CREATE_STORAGEEMGMT_TABLES` procedure – Create storage management tables” in *Administrative SQL Routines and Views*
- “`DROP_STORAGEEMGMT_TABLES` procedure – Drop all storage management tables” in *Administrative SQL Routines and Views*
- “Storage management view tables” on page 148

Stored procedures for the storage management tool

The following table shows the stored procedure functions that are created for the storage management tool. The stored procedures are automatically created when the database is created. Also, their respective packages are bound on demand.

Table 24. Stored procedures for the storage management tool

Fully qualified name	Parameters	Functionality
SYSPROC.CREATE_STORAGEMGMT_TABLES	in_tbspace VARCHAR(128) input - table space name	Creates all storage management tables under a fixed "DB2TOOLS" schema, in the table space specified by input.
SYSPROC.DROP_STORAGEMGMT_TABLES	dropSpec SMALLINT input - 0 / 1	Attempt to drop all storage management tables. When dropSpec=0, the process will stop when any error is encountered; when dropSpec=1, the process will continue ignoring any error it has encountered.
SYSPROC.CAPTURE_STORAGEMGMT_INFO	in_rootType SMALLINT input all valid values are given in STMG_OBJECT_TYPE table in_rootSchema VARCHAR(128) input - schema name of the storage snapshot root object in_rootName VARCHAR(128) input- name of the root object	Attempt to collect for system catalog and snapshot the storage-related information for the given root object, as well as its the storage objects defined within its scope. All the storage objects are specified in STMG_OBJECT_TYPE table.

Related reference:

- "Storage management view" on page 146

Storage management view tables

STMG_OBJECT_TYPE table:

The STMG_OBJECT_TYPE table contains one row for each supported storage type that can be monitored.

The STMG_OBJECT_TYPE must be specified as the first parameter to the capture_storagegmt_info() stored procedure. For example:
sysproc.capture_storagegmt_info(<stmg_object_type>, <object_schema>, <object_name>)

The first parameter, stmg_object_type, is defined by the entries in this table.

Table 25. STMG_OBJECT_TYPE table

Column name	Data type	Nullable	Description
OBJ_TYPE	INTEGER	N	Integer value corresponds to a type of storage object 0 - Database 1 - Database Partition Group 2 - Table Space 3 - Table Space Container 4 - Table 5 - Index

Table 25. STMG_OBJECT_TYPE table (continued)

Column name	Data type	Nullable	Description
TYPE_NAME	VARCHAR	N	Descriptive name of the storage object type STMG_DATABASE STMG_DBPGROUP STMG_TABLESPACE STMG_CONTAINER STMG_TABLE STMG_INDEX

STMG_THRESHOLD_REGISTRY table:

The STMG_THRESHOLD_REGISTRY table contains one row for each storage threshold type. The enabled thresholds are used by the analysis process when a storage snapshot is taken. If a threshold type is enabled, the threshold analysis will be performed on the data being monitored and threshold exceeded columns will be updated with the appropriate values for the specified threshold type.

Example::

To disable threshold analysis for table space usage:

```
db2 UPDATE SYSTOOLS.STMG_THRESHOLD_REGISTRY SET ENABLED = 'N'
WHERE STMG_TH_TYPE = 1
```

Table 26. STMG_THRESHOLD_REGISTRY table

Column name	Data type	Nullable	Description
STMG_TH_TYPE	INTEGER	N	Integer value corresponds to a storage threshold type 1 = STMG SPACE USAGE THRESHOLD 2 = STMG DATA SKEW THRESHOLD 3 = STMG CLUSTER RATIO THRESHOLD
ENABLED	CHARACTER	N	Y = the threshold is enabled N = the threshold is not enabled and therefore will not be compared against during storage analysis
STMG_TH_NAME	VARCHAR	Y	Descriptive name of the storage threshold STMG CLUSTER RATIO THRESHOLD STMG SPACE USAGE THRESHOLD STMG DATA SKEW THRESHOLD

STMG_CURR_THRESHOLD table:

The STMG_CURR_THRESHOLD table contains one row for each threshold type which is explicitly set for a storage object. When a new storage snapshot is taken, and threshold analysis is enabled for the objects being captured (see the Table 26), the values in this table are used to determine the warning and alarm thresholds that are set for each type of threshold being monitored. If an object under analysis does not have thresholds explicitly set in this table, the thresholds for the parent

object for that object type are used. By default, this table contains three rows, one for each threshold type. The thresholds in these three rows are set for the database object, the parent of all other objects in the database. All objects included in the storage snapshot analysis will automatically inherit these thresholds from the database object unless a threshold is set explicitly on a child object such as a table space or table.

Example::

To set the space usage warning and alarm thresholds for all objects in the database to 90 and 95:

```
db2 UPDATE SYSTOOLS.STMG_CURR_THRESHOLD SET WARNING_THRESHOLD = 90,
      ALARM_THRESHOLD = 95
      WHERE STMG_TH_TYPE = 1 AND OBJ_TYPE = 0
```

Table 27. STMG_CURR_THRESHOLD table

Column name	Data type	Nullable	Description
STMG_TH_TYPE	INTEGER	N	Integer value corresponds to a storage threshold type. See Table 26 on page 149 for a definition of threshold types.
OBJ_TYPE	INTEGER	N	Integer value corresponds to a type of storage object. See Table 25 on page 148 for a definition of threshold types.
OBJ_NAME	VARCHAR	N	The name of the storage object.
OBJ_SCHEMA	VARCHAR	N	The schema of the storage object. "- " is used when schema is not applicable for the object
WARNING_THRESHOLD	SMALLINT	Y	The value of the warning threshold set for the storage object.
ALARM_THRESHOLD	SMALLINT	Y	The value of the alarm threshold set for the storage object.

STMG_ROOT_OBJECT table:

The STMG_ROOT_OBJECT table contains one row for the root object of each storage snapshot. Complete storage snapshots can be deleted by deleting entries from this table.

Examples::

1. Delete all storage management snapshots:

```
db2 DELETE FROM SYSTOOLS.STMG_ROOT_OBJECT
```

2. Delete all table space snapshots:

```
db2 DELETE FROM SYSTOOLS.STMG_ROOT_OBJECT WHERE OBJ_TYPE = 2
```

Table 28. STMG_ROOT_OBJECT table

Column name	Data type	Nullable	Description
STMG_TIMESTAMP	TIMESTAMP	N	The timestamp of the storage snapshot. It indicates when the data capturing process started.
OBJ_TYPE	INTEGER	N	Integer value corresponds to a type of storage object. See Table 25 on page 148 for a definition of threshold types.
ROOT_ID	VARCHAR	N	The ID of the root object.

STMG_OBJECT table:

The STMG_OBJECT table contains one row for each storage object that is analyzed by the storage snapshots taken so far.

Note: Within a column, “(PK)” indicates a primary key.

Table 29. STMG_OBJECT table

Column name	Data type	Nullable	Description
STMG_TIMESTAMP (PK)	TIMESTAMP	N	The timestamp of the storage snapshot. It indicates the time the data capturing process started.
OBJ_ID (PK)	VARCHAR	N	The unique identifier for each storage object under a given storage snapshot timestamp.
ROOT_ID	CHARACTER	N	The ID of the root object.
OBJ_TYPE	INTEGER	N	Integer value corresponds to a type of storage object. See Table 25 on page 148 for a definition of threshold types.
OBJ_SCHEMA	VARCHAR	N	The schema of the storage object. “-” is used when schema is not applicable for the object
OBJ_NAME	VARCHAR	N	The name of the storage object.
DBPG_NAME	VARCHAR	Y	The name of the database partition group the object residing in. Null if not applicable.
TS_NAME	VARCHAR	Y	The name of the table space the object residing in. Null if not applicable.

STMG_HIST_THRESHOLD table:

The STMG_HIST_THRESHOLD table contains one row for each threshold used for the analyzing the storage objects at the time the storage snapshots are taken. This is basically a snapshot of what was in the SYSTOOLS.STMG_CURR_THRESHOLD table at the time of the snapshot.

Note: Within a column, “(PK)” indicates a primary key.

Table 30. STMG_HIST_THRESHOLD table

Column name	Data type	Nullable	Description
STMG_TIMESTAMP (PK)	TIMESTAMP	N	The timestamp of the storage snapshot. It indicates the time the data capturing process started.
STMG_TH_TYPE (PK)	INTEGER	N	Integer value corresponds to a storage threshold type. See Table 26 on page 149 for a definition of threshold types.
OBJ_ID (PK)	VARCHAR	N	The unique identifier for each storage object under a given storage snapshot timestamp.

Table 30. STMG_HIST_THRESHOLD table (continued)

Column name	Data type	Nullable	Description
WARNING_THRESHOLD	SMALLINT	Y	The value of the warning threshold set for the storage object at the time the storage snapshot was taken.
ALARM_THRESHOLD	SMALLINT	Y	The value of the alarm threshold set for the storage object at the time the storage snapshot was taken

STMG_DATABASE table:

The STMG_DATABASE table contains one row for each detailed entry of database storage snapshots.

Table 31. STMG_DATABASE table

Column name	Data type	Nullable	Description
STMG_TIMESTAMP (PK)	TIMESTAMP	N	The timestamp of the storage snapshot. It indicates when the data capturing process started.
OBJ_ID (PK)	VARCHAR	N	The unique identifier for each storage object under a given storage snapshot timestamp.
COMPLETE_TIMESTAMP	TIMESTAMP	Y	The timestamp of when the data capturing process has completed for the database, identified by OBJ_ID column.
REMARKS	VARCHAR	Y	User-specified remarks.

STMG_DBPGROUP table:

The STMG_DBPGROUP table contains one row for each detailed entry of database partition group storage snapshots.

Note: Within a column, "(PK)" indicates a primary key.

Table 32. STMG_DBPGROUP table

Column name	Data type	Nullable	Description
STMG_TIMESTAMP (PK)	TIMESTAMP	N	The timestamp of the storage snapshot. It indicates when the data capturing process started.
OBJ_ID (PK)	VARCHAR	N	The unique identifier for each storage object under a given storage snapshot timestamp.
COMPLETE_TIMESTAMP	TIMESTAMP	Y	The timestamp of when the data capturing process has completed for the database partition group, identified by OBJ_ID column.
PARTITON_COUNT	SMALLINT	Y	The number of database partitions included in the database partition group.

Table 32. STMG_DBPGROUP table (continued)

Column name	Data type	Nullable	Description
TARGET_LEVEL	BIGINT	Y	The average data size, in bytes, over all the database partitions contained by the database partition group. It is the target level of even data distribution.
DATA_SKEW	SMALLINT	Y	A percentage of the maximum data size deviation from the TARGET_LEVEL among all the database partitions. This value is used during data capture and analysis process to be compared against the data distribution skew set for the database partition group in the Table 27 on page 150.
TOTAL_SIZE	BIGINT	Y	The total size, in bytes, over all the database partitions contained by the database partition group. It is the sum of the total size (number of pages multiplied by page size) of all table spaces defined under the database partition group. For DMS table spaces, the total size is the allocated size; for SMS table spaces, it is the size of the currently used by the table space.
DATA_SIZE	BIGINT	Y	The data size, in bytes, over all the database partitions contained by the database partition group. It is the sum of the data size (number of data pages multiplied by page size) of all table spaces defined under the database partition group.
PERCENT_USED	SMALLINT	Y	A percentage value of data size over total size. This value is compared against the space usage threshold during the data capture and analysis process. In the case of SMS table spaces, the space usage threshold for the table space or its parent database partition group should be set to 100 to avoid unnecessary alarms.
REMARKS	VARCHAR	Y	User-specified remarks.

STMG_DBPARTITION table:

The STMG_DBPARTITION table contains one row for each detailed entry of database partition storage snapshots. This is meant to be used along with the STMG_DBPGROUP table.

Note: Within a column, “(PK)” indicates a primary key.

Table 33. STMG_DBPARTITION table

Column name	Data type	Nullable	Description
STMG_TIMESTAMP (PK)	TIMESTAMP	N	The timestamp of the storage snapshot. It indicates when the data capturing process started.
OBJ_ID (PK)	VARCHAR	N	The unique identifier for each storage object under a given storage snapshot timestamp.
PARTITION_NUM (PK)	INTEGER	Y	The database partition number.
COMPLETE_TIMESTAMP	TIMESTAMP	Y	The timestamp of when the data capturing process has completed for the database partition, identified by OBJ_ID column.
DBPG_NAME	CHARACTER	Y	The name of database partition group.
IN_USE	CHARACTER	Y	Status of the database partition at the time of the storage snapshot. Same as IN_USE column in SYSCAT.DBPARTITIONGROUPDEF.
HOST_NAME	VARCHAR	Y	The host name of the database partition.
HOST_SYSTEM_SIZE	BIGINT	Y	NOT AVAILABLE.
EST_DATA_SIZE	BIGINT	Y	The estimated data size on the database partition, within the database partition group scope. This value is calculated as the sum of the data size for that portion of the table found on the given partition.

STMG_TABLESPACE table:

The STMG_TABLESPACE table contains one row for each detailed entry of table space storage snapshots.

Note: Within a column, “(PK)” indicates a primary key.

Table 34. STMG_TABLESPACE table

Column name	Data type	Nullable	Description
STMG_TIMESTAMP (PK)	TIMESTAMP	N	The timestamp of the storage snapshot. It indicates when the data capturing process started.
OBJ_ID (PK)	VARCHAR	N	The unique identifier for each storage object under a given storage snapshot timestamp.
COMPLETE_TIMESTAMP	TIMESTAMP	Y	The timestamp of when the data capturing process has completed for the table space, identified by OBJ_ID column.
TYPE	CHARACTER	Y	As defined in SYSCAT.TABLESPACES.
DATATYPE	CHARACTER	Y	As defined in SYSCAT.TABLESPACES.
TOTAL_SIZE	BIGINT	Y	As defined in SYSCAT.TABLESPACES.

Table 34. STMG_TABLESPACE table (continued)

Column name	Data type	Nullable	Description
PERCENT_USED	SMALLINT	Y	As defined in SYSCAT.TABLESPACES. This is used during data capture and analysis process to be compared against the space usage threshold in the STMG_CURR_THRESHOLD table.
DATA_SIZE	BIGINT	Y	DATA_PAGE * PAGE_SIZE.
DATA_PAGE	BIGINT	Y	USED_PAGES as defined in SYSPROC.SNAPSHOT_TBS_CFG table UDF.
EXTENT_SIZE	INTEGER	Y	As defined in SYSCAT.TABLESPACES.
PREFETCH_SIZE	INTEGER	Y	As defined in SYSCAT.TABLESPACES.
OVERHEAD	DOUBLE	Y	As defined in SYSCAT.TABLESPACES.
TRANSFER_RATE	DOUBLE	Y	As defined in SYSCAT.TABLESPACES.
BUFFERPOOL_ID	INTEGER	Y	As defined in SYSCAT.TABLESPACES.
PAGE_SIZE	INTEGER	Y	As defined in SYSCAT.TABLESPACES.

STMG_CONTAINER table:

The STMG_CONTAINER table contains one row for each detailed entry of container storage snapshots.

Note: Within a column, “(PK)” indicates a primary key.

Table 35. STMG_CONTAINER table

Column name	Data type	Nullable	Description
STMG_TIMESTAMP (PK)	TIMESTAMP	N	The timestamp of the storage snapshot. It indicates when the data capturing process started.
OBJ_ID (PK)	VARCHAR	N	The unique identifier for each storage object under a given storage snapshot timestamp.
COMPLETE_TIMESTAMP	TIMESTAMP	Y	The timestamp of when the data capturing process has completed for the container, identified by OBJ_ID column.
TABLESPACE_ID	INTEGER	Y	tablespace_id - Table Space Identification monitor element
CONTAINER_ID	INTEGER	Y	container_id - Container Identification monitor element
PARTITION_NUM	INTEGER	Y	node_number - Node Number monitor element
CONTAINER_TYPE	CHARACTER	Y	container_type - Container Type monitor element
TOTAL_PAGES	BIGINT	Y	container_total_pages - Total Pages in Container monitor element
USABLE_PAGES	BIGINT	Y	container_usable_pages - Usable Pages in Container monitor element

Table 35. *STMG_CONTAINER* table (continued)

Column name	Data type	Nullable	Description
ACCESSIBLE	BIGINT	Y	container_accessible - Accessibility of Container monitor element
STRIPE_SET	BIGINT	Y	container_stripe_set - Stripe Set monitor element
FILESYSTEM_NODENAME	BIGINT	Y	The node name of the file system in which the container is defined.
FILESYSTEM_ID	BIGINT	Y	The unique file system identifier.
FILESYSTEM_MOUNT_POINT	VARCHAR	Y	The file system mount point.
FILESYSTEM_TYPE_NAME	VARCHAR	Y	File system type. For example, jfs, jfs2, ext2, or ntfs.
FILESYSTEM_DEVICE_TYPE	BIGINT	Y	File system device type.
FILESYSTEM_TOTAL_SIZE	BIGINT	Y	The total file system size in bytes.
FILESYSTEM_FREE_SIZE	BIGINT	Y	The total file system free size in bytes.
REMARKS	VARCHAR	Y	User-specified remarks.

STMG_TABLE table:

The *STMG_TABLE* table contains one or more rows for each table included in the specified snapshot type. A database snapshot would insert entries for each table in the database. A table space snapshot would insert one or more rows for each table in the specified table space, a table snapshot would insert entries for the table specified in the snapshot command.

For non-partitioned tables, there would be exactly one row per table. For partitioned tables, there would one row per table space that the table resides in. For example, if a partitioned table was spread over 5 table spaces, there would be 5 rows in the *STMG_TABLE* for that table. Each row would contain information specific to a table space with one exception: Information that relates to table totals for partitioned tables are a summation of values taken from all the table spaces; each row would show the same value where a table total is kept.

Note: Within a column, “(PK)” indicates a primary key.

Table 36. *STMG_TABLE* table

Column name	Data type	Nullable	Description
STMG_TIMESTAMP (PK)	TIMESTAMP	N	The timestamp of the storage snapshot. It indicates when the data capturing process started.
OBJ_ID (PK)	VARCHAR	N	The unique identifier for each storage object under a given storage snapshot timestamp.
COMPLETE_TIMESTAMP	TIMESTAMP	Y	The timestamp of when the data capturing process has completed for the table, identified by OBJ_ID column.
DBPG_NAME	VARCHAR	Y	The name of the database partition group in which the table resides.
TOTAL_ROW_COUNT	BIGINT	Y	Total row count of the table.

Table 36. STMG_TABLE table (continued)

Column name	Data type	Nullable	Description
AVG_ROW_COUNT	BIGINT	Y	The average row count from all portions of the table.
TARGET_LEVEL	BIGINT	Y	The average data size on each database partition, in bytes.
DATA_SKEW	SMALLINT	Y	The maximum percentage of the ROW_COUNT value deviated from the TARGET_LEVEL, over all portions of the table, for the given table. This is used during data capture and analysis process to be compared against the data skew threshold in the STMG_CURR_THRESHOLD table.
AVG_ROW_LENGTH	BIGINT	Y	The average row length of the table. If this statistic has been collected, it will be the sum of the average column length of all the columns in this table; when there is no statistical data, this value is calculated by adding the fixed columns' length with the percentage of the variable columns' length.
COLCOUNT	INTEGER	Y	As defined in SYSCAT.TABLES.
ESTIMATED_SIZE	BIGINT	Y	As defined in SYSCAT.TABLES.
NPAGES	INTEGER	Y	As defined in SYSCAT.TABLES.
FPAGES	INTEGER	Y	As defined in SYSCAT.TABLES.
OVERFLOW	INTEGER	Y	As defined in SYSCAT.TABLES.
MAIN_TBSPACE	VARCHAR	Y	As defined in SYSCAT.TABLES.
INDEX_TBSPACE	VARCHAR	Y	As defined in SYSCAT.TABLES.
LONG_TBSPACE	VARCHAR	Y	As defined in SYSCAT.TABLES.
REMARKS	VARCHAR	Y	User-specified remarks.
TABLE_PARTITIONED	CHAR(1)	N	Specifies whether the table is divided into one or more data partitions. Has value "Y" if table is partitioned and "N" otherwise.

STMG_TBPARTITION table:

The STMG_TBPARTITION table contains one row for each detailed entry of table partition storage snapshots.

Note: Within a column, "(PK)" indicates a primary key.

Table 37. STMG_TBPARTITION table

Column name	Data type	Nullable	Description
STMG_TIMESTAMP (PK)	TIMESTAMP	N	The timestamp of the storage snapshot. It indicates when the data capturing process started.
OBJ_ID (PK)	VARCHAR	N	The unique identifier for each storage object under a given storage snapshot timestamp.

Table 37. *STMG_TBPARTITION* table (continued)

Column name	Data type	Nullable	Description
PARTITION_NUM (PK)	INTEGER	N	The partition number of the database partition where the table partition resides.
COMPLETE_TIMESTAMP	TIMESTAMP	Y	The timestamp of when the data capturing process has completed for the table partition, identified by OBJ_ID column.
DBPG_NAME	VARCHAR	Y	The name of the database partition group where the table resides.
ROWCOUNT	BIGINT	Y	The number of rows in this table partition.
REMARKS	VARCHAR	Y	User-specified remarks.

STMG_INDEX table:

The *STMG_INDEX* table contains one row for each detailed entry of index storage snapshots.

Note: Within a column, “(PK)” indicates a primary key.

Table 38. *STMG_INDEX* table

Column name	Data type	Nullable	Description
STMG_TIMESTAMP (PK)	TIMESTAMP	N	The timestamp of the storage snapshot. It indicates when the data capturing process started.
OBJ_ID (PK)	VARCHAR	N	The unique identifier for each storage object under a given storage snapshot timestamp.
COMPLETE_TIMESTAMP	TIMESTAMP	Y	The timestamp of when the data capturing process has completed for the index, identified by OBJ_ID column.
DBPG_NAME	VARCHAR	Y	The name of the database partition group in which the index resides.
TB_SCHEMA	VARCHAR	Y	As TABNAME defined in SYSCAT.INDEXES.
TB_NAME	VARCHAR	Y	As TABSCHEMA defined in SYSCAT.INDEXES.
COLCOUNT	INTEGER	Y	As defined in SYSCAT.INDEXES.
ESTIMATED_SIZE	BIGINT	Y	As defined in SYSCAT.INDEXES.
NLEAF	INTEGER	Y	As defined in SYSCAT.INDEXES.
NLEVELS	SMALLINT	Y	As defined in SYSCAT.INDEXES.
FIRSTKEYCARD	BIGINT	Y	As defined in SYSCAT.INDEXES.
FIRST2KEYCARD	BIGINT	Y	As defined in SYSCAT.INDEXES.
FIRST3KEYCARD	BIGINT	Y	As defined in SYSCAT.INDEXES.
FIRST4KEYCARD	BIGINT	Y	As defined in SYSCAT.INDEXES.
FULLKEYCARD	BIGINT	Y	As defined in SYSCAT.INDEXES.

Table 38. STMG_INDEX table (continued)

Column name	Data type	Nullable	Description
CLUSTERRATIO	SMALLINT	Y	As defined in SYSCAT.INDEXES, this is used during data capture and analysis process to compare against the threshold set for the given index.
CLUSTERFACTOR	BIGINT	Y	As defined in SYSCAT.INDEXES.
SEQUENTIAL_PAGES	INTEGER	Y	As defined in SYSCAT.INDEXES.
DENSITY	INTEGER	Y	As defined in SYSCAT.INDEXES.
REMARKS	VARCHAR	Y	User-specified remarks.

STMG_OBJ_HISTORICAL_THRESHOLDS view:

The STMG_OBJ_HISTORICAL_THRESHOLDS view contains one row for each captured snapshot object. This view can be used to determine the thresholds that were set for a given object at the time of the snapshot. It can also be used to determine easily which objects have exceeded their thresholds for data skew, cluster ratio, and space usage.

Note: Within a column, "(PK)" indicates a primary key.

Table 39. STMG_OBJ_HISTORICAL_THRESHOLDS view

Column name	Data type	Nullable	Description
STMG_TIMESTAMP (PK)	TIMESTAMP	N	The timestamp of the storage snapshot. It indicates when the data capturing process started.
OBJ_ID (PK)	VARCHAR	N	The unique identifier for each storage object under a given storage snapshot timestamp.
OBJ_NAME (PK)	VARCHAR	N	The name of the storage object.
OBJ_SCHEMA (PK)	VARCHAR	N	The schema of the storage object. "- " is used when schema is not applicable for the object.
DBPG_NAME	VARCHAR	Y	The name of the database partition group where the object resides. Null if not applicable.
TS_NAME	VARCHAR	Y	The name of the table space in which the object resides. Null if not applicable.
SPACE_WARNING_THRESHOLD	SMALLINT	Y	The space usage warning threshold. Null if not applicable.
SPACE_ALARM_THRESHOLD	SMALLINT	Y	The space usage alarm threshold. Null if not applicable.
SPACE_THRESHOLD_EXCEEDED	SMALLINT	Y	The space usage threshold exceeded value. 1 if exceeded; 0 otherwise. Null if not applicable.
SKEW_WARNING_THRESHOLD	SMALLINT	Y	The data skew warning threshold. Null if not applicable.
SKEW_ALARM_THRESHOLD	SMALLINT	Y	The data skew alarm threshold. Null if not applicable.

Table 39. *STMG_OBJ_HISTORICAL_THRESHOLDS* view (continued)

Column name	Data type	Nullable	Description
SKEW_THRESHOLD_EXCEEDED	SMALLINT	Y	The data skew threshold exceeded value. 1 if exceeded; 0 otherwise. Null if not applicable.
CLUSTER_WARNING_THRESHOLD	SMALLINT	Y	The cluster ratio warning threshold. Null if not applicable.
CLUSTER_ALARM_THRESHOLD	SMALLINT	Y	The cluster ratio alarm threshold. Null if not applicable.
CLUSTER_THRESHOLD_EXCEEDED	SMALLINT	Y	The cluster ratio threshold exceeded value. 1 if exceeded; 0 otherwise. Null if not applicable.

Related reference:

- “Storage management view” on page 146

Thresholds

For storage management, thresholds are used to monitor the storage usage of your database. In the Storage Management view, you can set warning and alarm thresholds for the Storage Management tool to compare against the real time readings of the system. If an object’s storage state exceeds the safe levels, or thresholds, that you have set for it, an alert flag will be shown beside the object in the Storage Management view.

When a database is created, there are default thresholds set for the database object. All of its children (objects within the database scope) inherit the default threshold. However, you can overwrite the default thresholds by providing specific values for any of the objects. Once a threshold is set for an object, all the objects defined under its scope will inherit its threshold setting, unless otherwise specified.

The Storage Management view monitors three types of thresholds: space usage, data skew, and cluster ratio.

- Space usage measures the percentage of available storage space that is used by an object. Space usage is monitored through table spaces. The space usage of an object is represented as a percentage of total storage space, with a value of 0 to 100.
- Data skew measures the distribution of data by measuring an object’s deviation from the average data level, as a percentage. The data skew is monitored through tables and database partition groups. When a data skew threshold is exceeded, the Redistribute Data wizard can be used to even out data distribution differences among the database partitions in the database partition group. The disk skew of an object is represented as a percentage showing the object’s deviation from the average data level, ranging from -100 to 100. A negative integer indicates that the data level is less than the average data level; a positive integer indicates that the data level is greater than the average.
- Cluster ratio measures the degree to which the rows in a table are arranged in the same order specified by a given index. A higher cluster ratio indicates that the data rows are stored in the same physical sequence as the index. A low cluster ratio indicates the index and data rows are stored in a different physical sequence. Cluster ratio is represented as a percentage, with a value of 0 to 100.

In the Health Center, the criteria for the health indicators that measure a continuous range of values is defined in terms of thresholds. Thresholds define boundaries or zones and are configured as single-bounded with either increasing or decreasing values. There are three boundaries or zones: normal, warning, and alarm. If the value of a health indicator falls into the warning zone, a warning alert is issued. Similarly, if the indicator value falls into the alarm zone, an alarm alert is generated.

Related reference:

- “Storage management view” on page 146

Temporary table space design

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 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 spaces is *not* created by default at database creation time.

It is recommended that you define a single SMS temporary table space with a page size equal to the page size used in the majority of your regular 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 is inserted, or a batch of sequential rows is fetched. Therefore, a larger page size typically results in better performance, because fewer logical or physical page I/O requests are required to read a given amount of data. This is not always the case when the average temporary table row size is smaller than the page size divided by 255. A maximum of 255 rows can exist on any page, regardless of the page size. For example, a query that requires a temporary table with 15-byte rows would be better served by a 4 KB temporary table space page size, because 255 such rows can all be contained within a 4 KB page. An 8 KB (or larger) page size would result in at least 4 KB (or more) bytes of wasted space on each temporary table page, and would not reduce the number of required I/O requests.
- If more than fifty percent of the regular table spaces in your database use the same page size, it can be advantageous to define your temporary table spaces with the same page size. The reason for this is that this arrangement enables your temporary table space to share the same buffer pool space with most or all of your regular table spaces. This, in turn, simplifies buffer pool tuning.
- 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 target table space. Of course, this type of reorganization requires that there be extra space in the target table space for the reorganization process.

- If you are reliant on system temporary tables in SMS system temporary table spaces because of your work environment, you may want to consider using the registry variable `DB2_SMS_TRUNC_TMPTABLE_THRESH`. In the past when system temporary tables were no longer needed, they were truncated to a file size of zero. The need for a new system temporary table would have a performance cost associated with it. Using this registry variable allows for leaving non-zero system temporary tables on the system to avoid the performance cost of repeated creations and truncations of system temporary tables.
- In general, when temporary table spaces of differing 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.
- SMS is almost always a better choice than DMS for temporary table spaces because:
 - There is more overhead in the creation of a temporary table when using DMS versus SMS.
 - Disk space is allocated on demand in SMS, whereas it must be pre-allocated in DMS. Pre-allocation can be difficult: Temporary table spaces hold transient data that can have a very large peak storage requirement, and a much smaller average storage requirement. With DMS, the peak storage requirement must be pre-allocated, whereas with SMS, the extra disk space can be used for other purposes during off-peak hours.
 - The database manager attempts to keep temporary table pages in memory, rather than writing them out to disk. As a result, the performance advantages of DMS are less significant.

Related concepts:

- “System managed space” on page 117
- “Table space design” on page 112
- “Temporary tables in SMS table spaces” on page 162

Related reference:

- “REORG INDEXES/TABLE command” in *Command Reference*

Temporary tables in SMS table spaces

Temporary tables in SMS table spaces are not deleted by default once they are no longer needed. Instead, files associated with temporary tables are truncated to a length of zero. In cases where temporary tables are used repeatedly, this avoids some of the performance cost of deleting and recreating temporary tables.

This reuse of temporary tables benefits users whose workload involves dealing with many small temporary tables on smaller systems such as Windows where the file system calls are relatively expensive; and users whose disk storage is distributed, requiring network messages to complete file system operations.

By default, files that hold temporary tables are truncated to a zero length or to the extent size specified in the `DB2_SMS_TRUNC_TMPTABLE_THRESH` registry variable once they are no longer needed. You can set the number of extents to be used by specifying a value for the `DB2_SMS_TRUNC_TMPTABLE_THRESH` registry variable. You should increase the value associated with this registry variable if your workload repeatedly uses large SMS temporary tables and you can afford to leave space allocated between uses.

You can turn off this feature by specifying a value of 0 for the `DB2_SMS_TRUNC_TMPTABLE_THRESH` registry variable. You might want to do this if your system has restrictive space limitations and you are experiencing repeated out of disk errors for SMS temporary table spaces.

The first connection to the database deletes any previously allocated files. If you want to clear out existing temporary tables, you should drop all database connections and reconnect, or deactivate the database and reactivate it. If you want to ensure that space for temporary tables stays allocated, use the `ACTIVATE DATABASE` command to start the database. This will avoid the repeated cost of startup on the first connect to the database.

Related concepts:

- “Temporary table space design” on page 161

Catalog table space design

An SMS table space is recommended for database catalogs, for the following reasons:

- The database catalog consists of many tables of varying sizes. When using a DMS table space, a minimum of two extents are allocated for each table object. Depending on the extent size chosen, a significant amount of allocated and unused space may result. When using a DMS table space, a small extent size (two to four pages) should be chosen; otherwise, an SMS table space should be used.
- There are large object (LOB) columns in the catalog tables. LOB data is not kept in the buffer pool with other data, but is read from disk each time it is needed. Reading LOBs from disk reduces performance. Since a file system usually has its own cache, using an SMS table space, or a DMS table space built on file containers, makes avoidance of I/O possible if the LOB has previously been referenced.

Given these considerations, an SMS table space is a somewhat better choice for the catalogs.

Another factor to consider is whether you will need to enlarge the catalog table space in the future. While some platforms have support for enlarging the underlying storage for SMS containers, and while you can use redirected restore to enlarge an SMS table space, the use of a DMS table space facilitates the addition of new containers.

Note: When creating a database, three table spaces are defined, including the SYSCATSPACE table space for the system catalog tables. The page size that becomes the default for all table spaces is set when the database is created. If a page size greater than 4096 (or 4 KB) is chosen, the page size for the catalog tables is restricted to a row size that it would have if the catalog table space had a page size of 4 KB. The default database page size is stored as an informational database configuration parameter called *pagesize*.

Related concepts:

- “Database managed space” on page 120
- “System managed space” on page 117
- “Table space design” on page 112
- “System catalog tables” in *Administration Guide: Implementation*

Optimizing table space performance when data is on RAID devices

To optimize performance when data is placed on Redundant Array of Independent Disks (RAID) devices, you should do the following for each table space that uses a RAID device:

- Define a single container for the table space (using the RAID device).
- Make the EXTENTSIZE of the table space equal to, or a multiple of, the RAID stripe size.
- Ensure that the PREFETCHSIZE of the table space is:
 - the RAID stripe size multiplied by the number of RAID parallel devices (or a whole multiple of this product), and
 - a multiple of the EXTENTSIZE.
- Use the DB2_PARALLEL_IO registry variable to enable parallel I/O for the table space.

DB2_PARALLEL_IO:

When reading data from, or writing data to table space containers, DB2 Database for Linux, UNIX, and Windows may use parallel I/O if the number of containers in the database is greater than 1. However, there are situations when it would be beneficial to have parallel I/O enabled for single container table spaces. For example, if the container is created on a single RAID device that is composed of more than one physical disk, you may want to issue parallel read and write calls.

To force parallel I/O for a table space that has a single container, you can use the DB2_PARALLEL_IO registry variable. This variable can be set to "*" (asterisk), meaning every table space, or it can be set to a list of table space IDs separated by commas. For example:

```
db2set DB2_PARALLEL_IO=*      {turn parallel I/O on for all table spaces}
db2set DB2_PARALLEL_IO=1,2,4,8 {turn parallel I/O on for table spaces 1, 2,
                               4, and 8}
```

After setting the registry variable, DB2 must be stopped (**db2stop**), and then restarted (**db2start**), for the changes to take effect.

DB2_PARALLEL_IO also affects table spaces with more than one container defined. If you do not set the registry variable, the I/O parallelism is equal to the number of containers in the table space. If you set the registry variable, the I/O

parallelism is equal to the result of prefetch size divided by extent size. You might want to set the registry variable if the individual containers in the table space are striped across multiple physical disks.

For example, a table space has two containers and the prefetch size is four times the extent size. If the registry variable is not set, a prefetch request for this table space will be broken into two requests (each request will be for two extents). Provided that the prefetchers are available to do work, two prefetchers can be working on these requests in parallel. In the case where 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.

In this example, if each of the two containers had a single disk dedicated to it, setting the registry variable for this table space might result in contention on those disks since 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.

DB2_USE_PAGE_CONTAINER_TAG:

By default, DB2 uses the first extent of each DMS container (file or device) to store a container tag. The container tag is DB2's metadata for the container. In earlier versions of the DB2 database system, the first page was used for the container tag, instead of the first extent, and as a result less space in the container was used to store the tag. (In earlier versions of the DB2 database system, the `DB2_STRIPED_CONTAINERS` registry variable was used to create table spaces with an extent sized tag. However, because this is now the default behavior, this registry variable no longer has any affect.)

When the `DB2_USE_PAGE_CONTAINER_TAG` registry variable is set to ON, any new DMS containers created will be created with a one-page tag, instead of a one-extent tag (the default). There will be no impact to existing containers that were created before the registry variable was set.

Setting this registry variable to ON is not recommended unless you have very tight space constraints, or you require behavior consistent with pre-Version 8 databases.

Setting this registry variable to ON can have a negative impact on I/O performance if RAID devices are used for table space containers. When using RAID devices for table space containers, it is suggested that the table space be created with an extent size that is equal to, or a multiple of, the RAID stripe size. However, if this registry variable is set to ON, a one-page container tag will be used and the extents will not line up with the RAID stripes. As a result it may be necessary during an I/O request to access more physical disks than would be optimal. Users are thus strongly advised against setting this registry variable.

Procedure:

To create containers with one-page container tags, set this registry variable to ON, and then stop and restart the instance:

```
db2set DB2_USE_PAGE_CONTAINER_TAG=ON
db2stop
db2start
```

To stop creating containers with one-page container tags, reset this registry variable, and then stop and restart the instance.

```
db2set DB2_USE_PAGE_CONTAINER_TAG=  
db2stop  
db2start
```

The Control Center, the LIST TABLESPACE CONTAINERS command, and the GET SNAPSHOT FOR TABLESPACES command do not show whether a container has been created with a page or extent sized tag. They use the label “file” or “device,” depending on how the container was created. To verify whether a container was created with a page- or extent-size tag, you can use the /DTSF option of DB2DART to dump table space and container information, and then look at the type field for the container in question. The query container APIs (sqlbftcq and sqlbtcq), can be used to create a simple application that will display the type.

Related concepts:

- “Table space design” on page 112

Related reference:

- “System environment variables” in *Performance Guide*

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.

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 of allocating space one page at a time, 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 LOB and long field 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 tables 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.

Related concepts:

- “Comparison of SMS and DMS table spaces” on page 140
- “Database managed space” on page 120
- “Database partition groups” on page 85
- “System managed space” on page 117

DB2 table types

DB2 Database for Linux, UNIX, and Windows provides the following types of tables:

- Regular tables, which are implemented as a heap
- Append mode tables, which are regular tables that are optimized primarily for INSERTs
- Multidimensional clustering (MDC) tables, which are implemented as tables that are physically clustered on more than one key, or dimension, at the same time
- Range-clustered tables (RCT), which are implemented as sequential clusters of data that provide fast, direct access
- Partitioned tables, which are implemented as tables with data divided across multiple data partitions according to values in the table partitioning key columns for the table.

Each type of table has characteristics that make it useful when working in a particular business environment. For each table that you use, consider which table types would best suit your needs.

Regular tables with indexes are the “general purpose” table choice.

Regular tables are placed into append mode through an ALTER TABLE statement. Append mode tables are suitable where you need to add new data and retrieve existing data such as where you are dealing with customer accounts in a banking environment. There you record each change to the account through debits, credits, and transfers. You also have customers who want to review the history of changes to that account.

Multidimensional clustering 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.

Range-clustered 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.

Partitioned tables allow easier roll-in and roll-out of table data, easier administration, flexible index placement, and better query processing than regular tables.

Related concepts:

- “Multidimensional clustering tables” on page 172
- “Range-clustered tables” on page 168

Related tasks:

- “Creating a materialized query table” in *Administration Guide: Implementation*
- “Creating and populating a table” in *Administration Guide: Implementation*

Range-clustered tables

A range-clustered table (RCT) is a table layout scheme where each record in the table has a predetermined record ID (RID) which is an internal identifier used to locate a record in a table.

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.

An algorithm is used to equate the value of the key for the record with the location of a specific row within a table. The basic algorithm is fairly simple. In its most basic form (using a single column instead of two or more columns to make up the key), the algorithm maps a sequence number to a logical row number. The algorithm also uses the record’s key to determine the logical page number and slot number. This process provides exceptionally fast access to records; that is, to specific rows in the table.

The algorithm does not involve hashing because hashing does not preserve key-value ordering. Preserving key-value ordering is essential because it eliminates the need to reorganize the table data over time.

Each record key in the table should have the following characteristics:

- Unique
- Not null
- An integer (SMALLINT, INTEGER, or BIGINT)
- Monotonically increasing

- Within a predetermined set of ranges based on each column in the key
The ALLOW OVERFLOW option is used when creating the table to allow key values to exceed the defined range. The DISALLOW OVERFLOW option is used when creating the table where key values will not exceed the defined range. In this case, if a record is inserted out of the boundary indicated by the range, an SQL error message is returned.

Applications where tightly clustered (dense) sequence key ranges are likely are excellent candidates for range-clustered tables. When using this type of key to create a range-clustered table, the key is used to generate the logical location of a row in a table. This process avoids the need for a separate index.

Advantages associated with a range-clustered table structure include the following factors:

- Direct access
Access is through a range-clustered table key-to-RID mapping function.
- Less maintenance
A secondary structure such as a B+ tree does not need to be updated for every INSERT, UPDATE, or DELETE.
- Less logging
There is less logging done for range-clustered tables when compared to a similarly sized regular table and associated B+ tree index.
- Less buffer pool memory required
There is no additional memory required to store a secondary structure.
- Order properties of B+ tree tables
The ordering of the records is the same as what was achieved by B+ tree tables without requiring extra levels or B+ tree next-key locking schemes. With RCT, the code path length is reduced compared to regular B+ tree indexes. To obtain this advantage, however, the range-clustered table must be created with DISALLOW OVERFLOW and the data must be dense, not sparse.
- One less index
Mapping each key to a location on disk means that the table can be created with one less index than would have been necessary otherwise. With range-clustered tables, the application requirements for accessing the data in the table might make a second, separate index unnecessary. You may still choose to create regular indexes, especially if the application requires it.

Indexes are used to perform the following functions:

- Locate a record based on a key from the record
- Apply start and stop key scans
- Distribute data vertically
By using an RCT, the only property of an index that is not accounted for is vertical distribution of data.

When deciding to use range-clustered tables, consider the following characteristics which differentiate them from regular tables:

- Range-clustered tables have no free-space control records (FSCR).
- Space is preallocated.
Space for the table is preallocated and reserved for use by the table even when records for the table are not filled in. At table creation time, there are no records

in the table; however, the entire range of pages is preallocated. Preallocation is based on the record size and the maximum number of records to be stored.

- If variable length fields such as VARCHAR are used in each record, the maximum length of the field is used and the overall record size is a fixed length. The overall fixed length of each record is used with the maximum number of records to determine the space required.
 - This can result in additional space being allocated that cannot be effectively utilized.
 - If key values are sparse, there is unused space and poor range scan performance.
 - Range scans must visit all possible records within a range even if the rows containing those key values have not yet been inserted into the database.
- No schema modifications permitted.

If a schema modification is required on a range-clustered table, the table must be recreated to include the new schema name for the table and all the data from the old table. In particular:

- Altering a key range is not supported.
This is important since if a table's ranges need to be altered, a new table with the desired ranges must be created and the new table populated with the data from the old table.
- Duplicate key values are not allowed.
 - Key values outside the defined range are not allowed.
This is true for range-clustered tables defined to DISALLOW OVERFLOW only.
 - NULL values are explicitly disallowed.
 - Range-cluster index is not materialized
An index with RCT key properties is indicated in the system catalogs and can be selected by the optimizer, but the index is not materialized on disk. With a regular table, space also needs to be given to each index associated with a table. With a RCT, no space is required for the RCT index. The optimizer uses the information in the system catalogs that refers to this RCT index to ensure that the correct access method for the table can be chosen.
 - Creating a primary or a unique key on the same definition as the range-clustered table index is not permitted since it would be redundant.
 - Range-clustered tables retain the original key value ordering, a feature that guarantees the clustering of rows within a table.

In addition to those considerations, there are some incompatibilities that either limit places where range-clustered tables can be used, or other utilities that do not work with these tables. The limitations on range-clustered tables include:

- Range clustered tables on partitioned tables are not supported.
If you attempt to create a partitioned table with range clustering, the error message SQL0270 rc=87 is returned.
- Declared global temporary tables (DGTT) are not supported.
These temp tables are not allowed to use the range cluster property.
- Automatic summary tables (AST) are not supported.
These tables are not allowed to use the range cluster property.
- Load utility is not supported.
Rows must be inserted one at a time through an import operation or a parallel inserting application.
- REORG TABLE utility is not supported.

Range-clustered tables that are defined to `DISALLOW OVERFLOW` will not need to be reorganized. Those range-clustered tables defined to `ALLOW OVERFLOW` are still not permitted to have the data in this overflow region reorganized.

- Range-clustered tables on one logical machine only.
On the Enterprise Server Edition (ESE) with the Database Partitioning Feature (DPF), a range-clustered table cannot exist in a database containing more than one database partition.
- The design advisor will not recommend range-clustered tables.
- Range-clustered tables are, by definition, already clustered.
This means that the following clustering schemes are incompatible with range-clustered tables:
 - Multi-dimensional clustered (MDC) table
 - Clustering indexes
- Value and default compression are not supported.
- Reverse scans on the range-clustered table are not supported.
- The `REPLACE` option on the `IMPORT` command is not supported.
- The `WITH EMPTY TABLE` option on the `ALTER TABLE ... ACTIVATE NOT LOGGED INITIALLY` statement is not supported.

Related concepts:

- “Range-clustered tables and out-of-range record key values” on page 171
- “Examples of range-clustered tables” in *Administration Guide: Implementation*

Related reference:

- “Restrictions on native XML data store” in *XML Guide*

Range-clustered tables and out-of-range record key values

You control the behavior of a range-clustered table (RCT) that allows overflow records by using the `CREATE TABLE` statement and the `ALLOW OVERFLOW` option. In this way, you ensure that all of the pages required by the table within the defined range are allocated immediately.

Once created, any records with keys that fall into the defined range work the same way, regardless of whether the table is created with the overflow option allowed or disallowed. The difference occurs when there is a record with a key that falls outside of the defined range. In this case, when the table allows overflow records, the record is placed in the overflow area, which is dynamically allocated. As more records are added from outside the defined range, they are placed into the growing overflow area. Actions against the table that involve this overflow area will require longer processing time because the overflow area must be accessed as part of the action. The larger the overflow area, the longer it will take to access the overflow area. After prolonged use of the overflow area, consider reducing its size by exporting the data from the table to a new range-clustered table that you have defined using new, extended ranges.

There might be times when you do not want records placed into a range-clustered table to have record key values falling outside of an allowed or defined range. For this type of RCT to exist, you must use the `DISALLOW OVERFLOW` option on the

CREATE TABLE statement. Once you have created this type of RCT, you might have to accept error messages if a record key value falls outside of the allowed or defined range.

Related reference:

- “CREATE TABLE statement” in *SQL Reference, Volume 2*

Range-clustered table locks

Within normal processing, locking of records takes place to ensure that only one application or user has access to a record or group of records at any given time. With range-clustered tables, instead of key and next-key locking, “discrete locking” is used. This method locks all records that are effected by, or might be effected by, the operation requested by the application or user. The number of locks that are obtained depends on the isolation level.

Qualifying rows in range-clustered tables that are currently empty but have been preallocated are locked. This avoids the need for next-key locking. As a result, fewer locks are required for a dense, range-clustered table.

Related concepts:

- “Locks and concurrency control” in *Performance Guide*

Multidimensional clustering tables

Multidimensional clustering (MDC) provides an elegant method for clustering data in tables along multiple dimensions in a flexible, continuous, and automatic way. MDC can significantly improve query performance. In addition, MDC can significantly reduce the overhead of data maintenance, such as reorganization and index maintenance operations during insert, update, and delete operations. MDC is primarily intended for data warehousing and large database environments, but it can also be used in online transaction processing (OLTP) environments.

Related concepts:

- “Indexes” in *SQL Reference, Volume 1*
- “Block index considerations for MDC tables” on page 188
- “Block indexes” on page 175
- “Optimization strategies for MDC tables” in *Performance Guide*
- “Table and index management for MDC tables” in *Performance Guide*
- “Block indexes and query performance” on page 180
- “Block maps” on page 185
- “Comparison of regular and MDC tables” on page 173
- “Deletion from an MDC table” on page 187
- “Designing multidimensional clustering (MDC) tables” on page 189
- “Load considerations for MDC tables” on page 188
- “Logging considerations for MDC tables” on page 188
- “Maintaining clustering automatically during INSERT operations” on page 183
- “Multidimensional clustering (MDC) table creation, placement, and use” on page 197
- “Updating an MDC table” on page 187

- “Working with an MDC table” on page 177
- “Multidimensional clustering considerations when loading data” in *Data Movement Utilities Guide and Reference*

Related reference:

- “Lock modes for table and RID index scans of MDC tables” in *Performance Guide*
- “Locking for block index scans for MDC tables” in *Performance Guide*

Comparison of regular and MDC tables

Regular tables have indexes that are record-based. Any clustering of the indexes is restricted to a single dimension. Prior to Version 8, DB2 Universal Database supported only single-dimensional clustering of data, through clustering indexes. Using a clustering index, DB2 attempts to maintain the physical order of data on pages in the key order of the index when records are inserted and updated in the table. Clustering indexes greatly improve the performance of range queries that have predicates containing the key (or keys) of the clustering index. Performance is improved with a good clustering index because only a portion of the table needs to be accessed, and more efficient prefetching can be performed.

Data clustering using a clustering index has some drawbacks, however. First, because space is filled up on data pages over time, clustering is not guaranteed. An insert operation will attempt to add a record to a page nearby to those having the same or similar clustering key values, but if no space can be found in the ideal location, it will be inserted elsewhere in the table. Therefore, periodic table reorganizations may be necessary to re-cluster the table and to setup pages with additional free space to accommodate future clustered insert requests.

Second, only one index can be designated as the “clustering” index, and all other indexes will be unclustered, because the data can only be physically clustered along one dimension. This limitation is related to the fact that the clustering index is record-based, as all indexes have been prior to Version 8.1.

Third, because record-based indexes contain a pointer for every single record in the table, they can be very large in size.

Clustering index




Figure 42. A regular table with a clustering index

The table in Figure 42 has two record-based indexes defined on it:

- A clustering index on “Region”
- Another index on “Year”

The “Region” index is a clustering index which means that as keys are scanned in the index, the corresponding records should be found for the most part on the same or neighboring pages in the table. In contrast, the “Year” index is unclustered which means that as keys are scanned in that index, the corresponding records will likely be found on random pages throughout the table. Scans on the clustering index will exhibit better I/O performance and will benefit more from sequential prefetching, the more clustered the data is to that index.

MDC introduces indexes that are block-based. “Block indexes” point to blocks or groups of records instead of to individual records. By physically organizing data in an MDC table into blocks according to clustering values, and then accessing these blocks using block indexes, MDC is able not only to address all of the drawbacks of clustering indexes, but to provide significant additional performance benefits.

First, MDC enables a table to be physically clustered on more than one key, or dimension, simultaneously. With MDC, the benefits of single-dimensional clustering are therefore extended to multiple dimensions, or clustering keys. Query performance is improved where there is clustering of one or more specified dimensions of a table. Not only will these queries access only those pages having records with the correct dimension values, these qualifying pages will be grouped into blocks, or extents.

Second, although a table with a clustering index can become unclustered over time, an MDC table is able to maintain and guarantee its clustering over all dimensions automatically and continuously. This eliminates the need to reorganize MDC tables to restore the physical order of the data.

Third, in MDC the clustering indexes are block-based. These indexes are drastically smaller than regular record-based indexes, so take up much less disk space and are faster to scan.

Block indexes

The MDC table shown in Figure 43 is physically organized such that records having the same “Region” and “Year” values are grouped together into separate blocks, or extents. An extent is a set of contiguous pages on disk, so these groups of records are clustered on physically contiguous data pages. Each table page belongs to exactly one block, and all blocks are of equal size (that is, an equal number of pages). The size of a block is equal to the extent size of the table space, so that block boundaries line up with extent boundaries. In this case, two block indexes are created, one for the “Region” dimension, and another for the “Year” dimension. These block indexes contain pointers only to the blocks in the table. A scan of the “Region” block index for all records having “Region” equal to “East” will find two blocks that qualify. All records, and only those records, having “Region” equal to “East” will be found in these two blocks, and will be clustered on those two sets of contiguous pages or extents. At the same time, and completely independently, a scan of the “Year” index for records between 1999 and 2000 will find three blocks that qualify. A data scan of each of these three blocks will return all records and only those records that are between 1999 and 2000, and will find these records clustered on the sequential pages within each of the blocks.

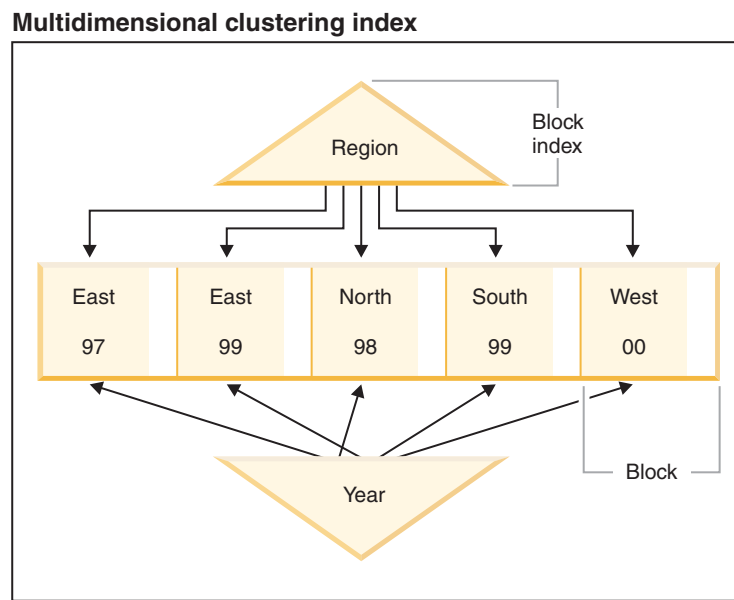


Figure 43. A multidimensional clustering table

In addition to these clustering improvements, MDC tables provide the following benefits:

- Probes and scans of block indexes are much faster due to their incredibly small size in relation to record-based indexes
- Block indexes and the corresponding organization of data allows for fine-grained “database partition elimination”, or selective table access
- Queries that utilize the block indexes benefit from the reduced index size, optimized prefetching of blocks, and guaranteed clustering of the corresponding data
- Reduced locking and predicate evaluation is possible for some queries

- Block indexes have much less overhead associated with them for logging and maintenance because they only need to be updated when adding the first record to a block, or removing the last record from a block
- Data rolled in can reuse the contiguous space left by data previously rolled out.

Note: An MDC table defined with even just a single dimension can benefit from these MDC attributes, and can be a viable alternative to a regular table with a clustering index. This decision should be based on many factors, including the queries that make up the workload, and the nature and distribution of the data in the table. Refer to “Considerations when choosing dimensions” and “MDC Advisor Feature on the DB2 Advisor”.

When you create a table, you can specify one or more keys as dimensions along which to cluster the data. Each of these MDC dimensions can consist of one or more columns similar to regular index keys. A *dimension block index* will be automatically created for each of the dimensions specified, and it will be used by the optimizer to quickly and efficiently access data along each dimension. A *composite block index* will also automatically be created, containing all columns across all dimensions, and will be used to maintain the clustering of data over insert and update activity. A composite block index will only be created if a single dimension does not already contain all the dimension key columns. The composite block index may also be selected by the optimizer to efficiently access data that satisfies values from a subset, or from all, of the column dimensions.

Note: The usefulness of this index during query processing depends on the order of its key parts. The key part order is determined by the order of the columns encountered by the parser when parsing the dimensions specified in the ORGANIZE BY clause of the CREATE TABLE statement. Refer to section “Block index considerations for MDC tables” for more information.

Block indexes are structurally the same as regular indexes, except that they point to blocks instead of records. Block indexes are smaller than regular indexes by a factor of the block size multiplied by the average number of records on a page. The number of pages in a block is equal to the extent size of the table space, which can range from 2 to 256 pages. The page size can be 4 KB, 8 KB, 16 KB, or 32 KB.




Figure 44. How row indexes differ from block indexes

As seen in Figure 44, in a block index there is a single index entry for each block compared to a single entry for each row. As a result, a block index provides a significant reduction in disk usage and significantly faster data access.

In an MDC table, every unique combination of dimension values form a logical *cell*, which may be physically made up of one or more blocks of pages. The logical cell will only have enough blocks associated with it to store the records having the dimension values of that logical cell. If there are no records in the table having the dimension values of a particular logical cell, no blocks will be allocated for that logical cell. The set of blocks that contain data having a particular dimension key value is called a *slice*.

Related concepts:

- “Block index considerations for MDC tables” on page 188
- “Block indexes and query performance” on page 180
- “Block maps” on page 185

Related reference:

- “Restrictions on native XML data store” in *XML Guide*

Working with an MDC table

As an example of how to work with an MDC table, we will imagine an MDC table called “Sales” that records sales data for a national retailer. The table is clustered along the dimensions “YearAndMonth” and “Region”. Records in the table are stored in blocks, which contain enough consecutive pages on disk to fill an extent. In Figure 45 on page 178, a block is represented by a rectangle, and is numbered according to the logical order of allocated extents in the table. The grid in the diagram represents the logical database partitioning of these blocks, and each square represents a logical cell. A column or row in the grid represents a slice for a particular dimension. For example, all records containing the value ‘South-central’ in the “Region” column are found in the blocks contained in the slice defined by

the 'South-central' column in the grid. In fact, each block in this slice also only contains records having 'South-central' in the "Region" field. Thus, a block is contained in this slice or column of the grid if and only if it contains records having 'South-central' in the "Region" field.

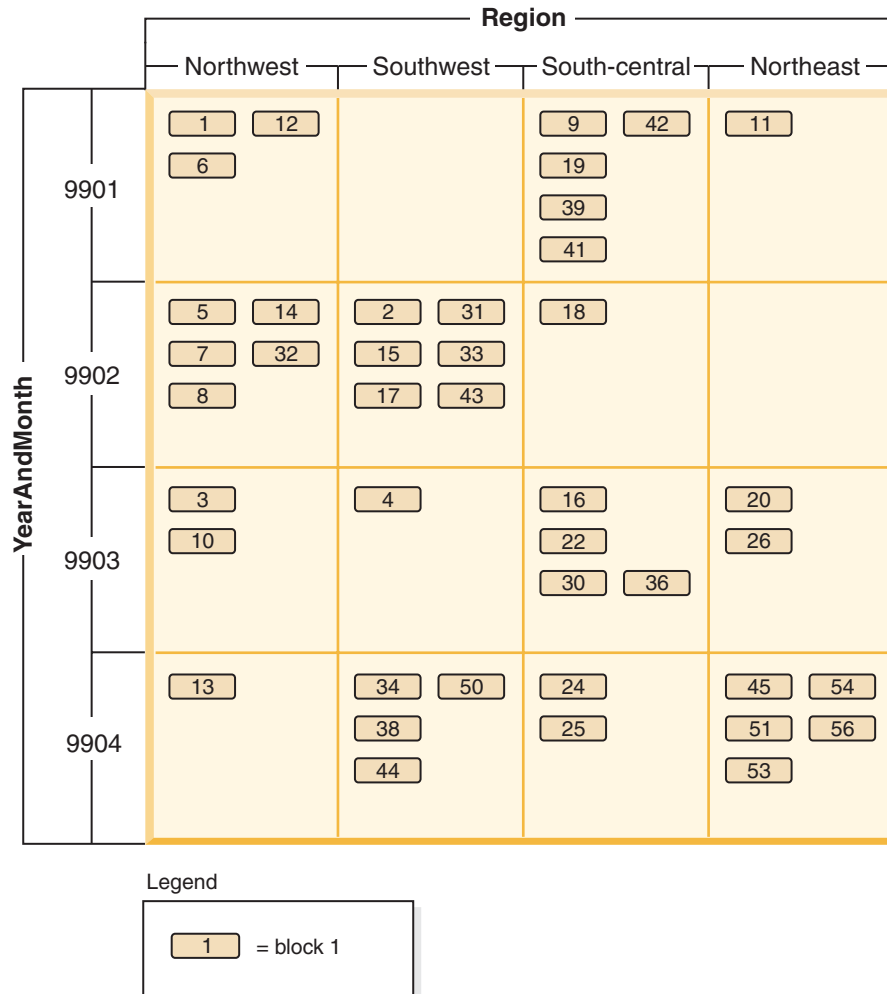


Figure 45. Multidimensional table with dimensions of 'Region' and 'YearAndMonth' that is called Sales

To determine which blocks comprise a slice, or equivalently, which blocks contain all records having a particular dimension key value, a dimension block index is automatically created for each dimension when the table is created.

In Figure 46 on page 179, a dimension block index is created on the "YearAndMonth" dimension, and another on the "Region" dimension. Each dimension block index is structured in the same manner as a traditional RID index, except that at the leaf level the keys point to a block identifier (BID) instead of a record identifier (RID). A RID identifies the location of a record in the table by a physical page number and a slot number — the slot on the page where the record is found. A BID represents a block by the physical page number of the first page of that extent, and a dummy slot (0). Because all pages in the block are physically consecutive starting from that one, and we know the size of the block, all records in the block can be found using this BID.

A slice, or the set of blocks containing pages with all records having a particular key value in a dimension, will be represented in the associated dimension block index by a BID list for that key value.

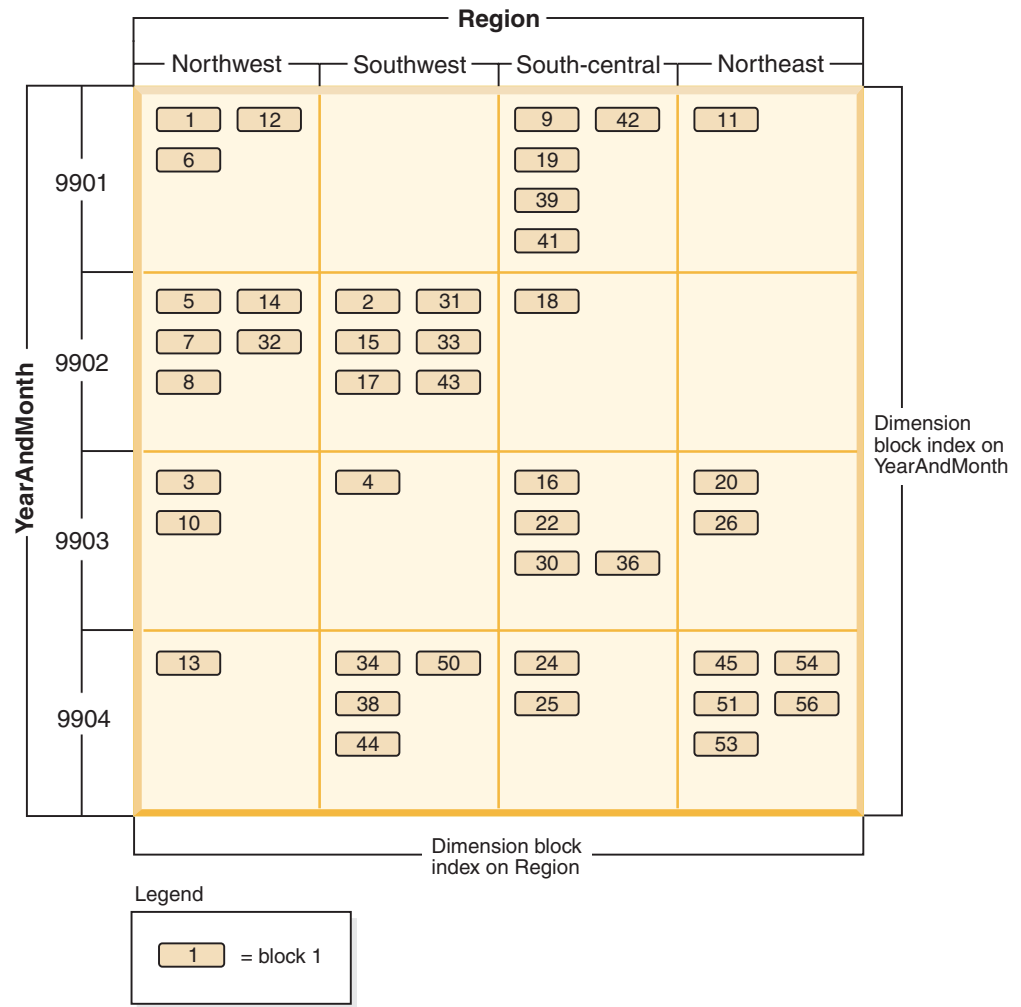


Figure 46. Sales table with dimensions of 'Region' and 'YearAndMonth' showing dimension block indexes

Figure 47 shows how a key from the dimension block index on "Region" would appear. The key is made up of a key value, namely 'South-central', and a list of BIDs. Each BID contains a block location. In Figure 47, the block numbers listed are the same that are found in the 'South-central' slice found in the grid for the Sales table (see Figure 45 on page 178).




Figure 47. Key from the dimension block index on 'Region'

Similarly, to find the list of blocks containing all records having '9902' for the "YearAndMonth" dimension, look up this value in the "YearAndMonth" dimension block index, shown in Figure 48.




Figure 48. Key from the dimension block index on 'YearAndMonth'

Related concepts:

- "Multidimensional clustering (MDC) table creation, placement, and use" on page 197
- "Multidimensional clustering tables" on page 172

Block indexes and query performance

Scans on any of the block indexes of an MDC table provide clustered data access, because each BID corresponds to a set of sequential pages in the table that is guaranteed to contain data having the specified dimension value. Moreover, dimensions or slices can be accessed independently from each other through their block indexes without compromising the cluster factor of any other dimension or slice. This provides the multidimensionality of multidimensional clustering.

Queries that take advantage of block index access can benefit from a number of factors that improve performance. First, the block index is so much smaller than a regular index, the block index scan is very efficient. Second, prefetching of the data pages does not rely on sequential detection when block indexes are used. DB2 looks ahead in the index, prefetching the data pages of the blocks into memory using big-block I/O, and ensuring that the scan does not incur the I/O when the data pages are accessed in the table. Third, the data in the table is clustered on sequential pages, optimizing I/O and localizing the result set to a selected portion of the table. Fourth, if a block-based buffer pool is used with its block size being the extent size, then MDC blocks will be prefetched from sequential pages on disk into sequential pages in memory, further increasing the effect of clustering on performance. Finally, the records from each block are retrieved using a mini-relational scan of its data pages, which is often a faster method of scanning data than through RID-based retrieval.

Queries use can use block indexes to narrow down a portion of the table having a particular dimension value or range of values. This provides a fine-grained form of "database partition elimination", that is, block elimination. This can translate into better concurrency for the table, because other queries, loads, inserts, updates and deletes may access other blocks in the table without interacting with this query's data set.

If the Sales table is clustered on three dimensions, the individual dimension block indexes can also be used to find the set of blocks containing records which satisfy a query on a subset of all of the dimensions of the table. If the table has dimensions of "YearAndMonth", "Region" and "Product", this can be thought of as a logical cube, as illustrated in Figure 49 on page 181.

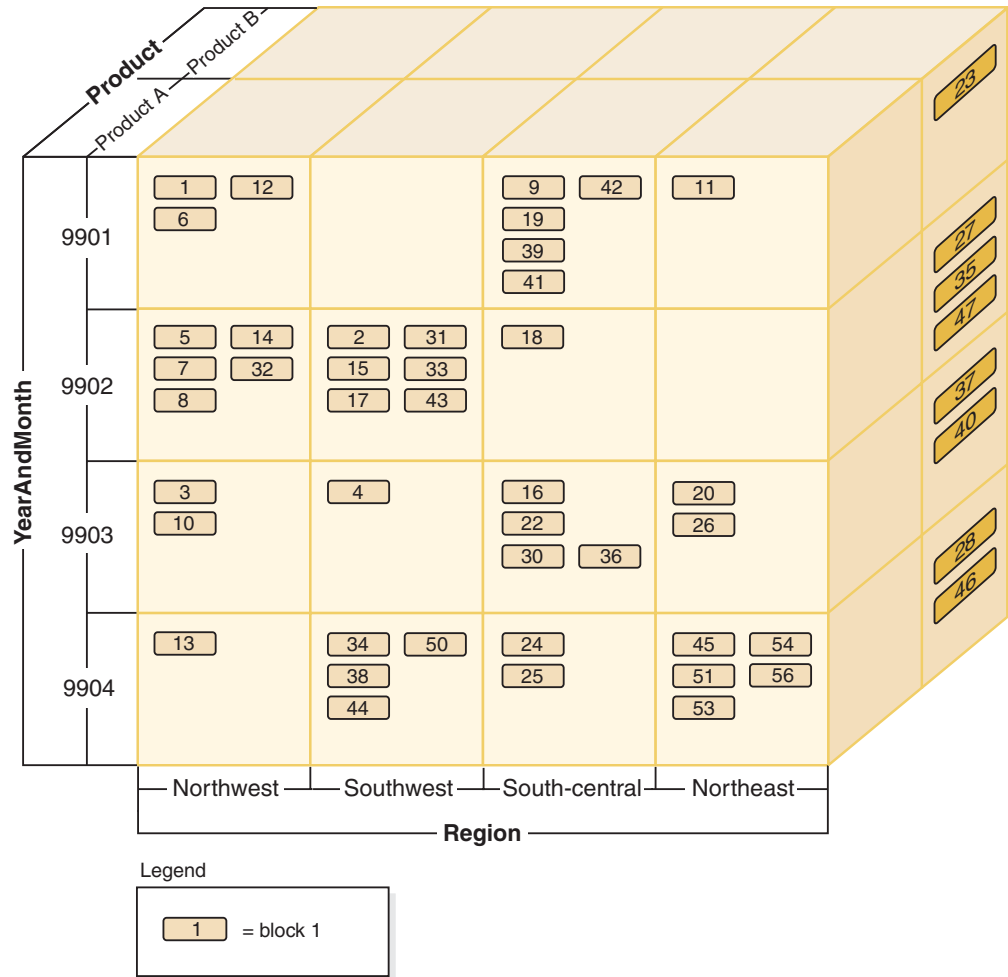


Figure 49. Multidimensional table with dimensions of 'Region', 'YearAndMonth', and 'Product'

Four block indexes will be created for the MDC table shown in Figure 49: one for each of the individual dimensions, "YearAndMonth", "Region", and "Product"; and another with all of these dimension columns as its key. To retrieve all records having a "Product" equal to "ProductA" and "Region" equal to "Northeast", DB2 would first search for the ProductA key from the "Product" dimension block index. (See Figure 50.) DB2 then determines the blocks containing all records having "Region" equal to "Northeast", by looking up the "Northeast" key in the "Region" dimension block index. (See Figure 51.)

Product A	1	2	3	...	11	...	20	22	24	25	26	30	...	56
-----------	---	---	---	-----	----	-----	----	----	----	----	----	----	-----	----

Figure 50. Key from dimension block index on 'Product'

Northeast	11	20	23	26	27	28	35	37	40	45	46	47	51	53	54	56
-----------	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----

Figure 51. Key from dimension block index on 'Region'

Block index scans can be combined through the use of the logical AND and logical OR operators and the resulting list of blocks to scan also provides clustered data access.

Using the example above, in order to find the set of blocks containing all records having both dimension values, you have to find the intersection of the two slices. This is done by using the logical AND operation on the BID lists from the two block index keys. The common BID values are 11, 20, 26, 45, 54, 51, 53, and 56.

The following example illustrates how using the logical OR operation with block indexes to satisfy a query having predicates that involve two dimensions. Figure 52 assumes an MDC table where the two dimensions are "Color" and "Nation". The goal is to retrieve all those records in the MDC table that meet the conditions of having "Color" of "blue" or having a "Nation" name "USA".

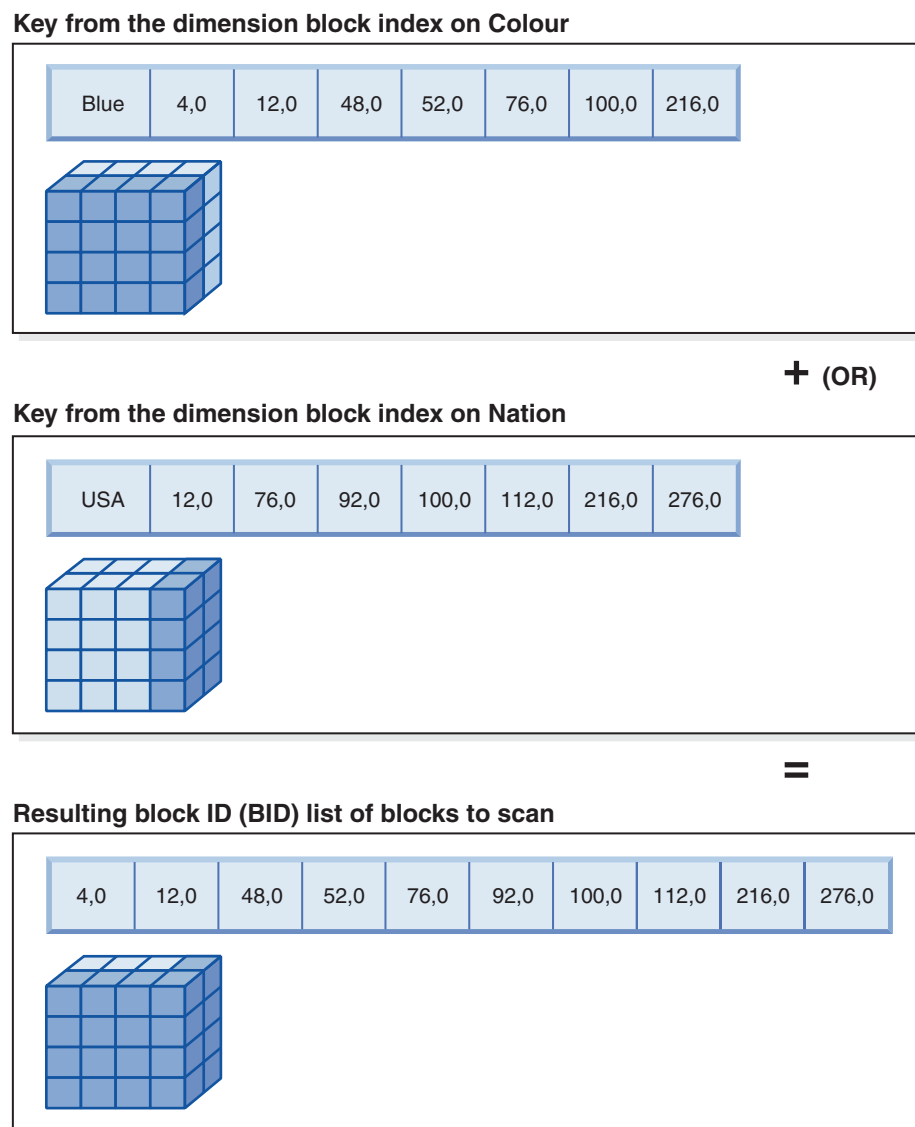


Figure 52. How the logical OR operation can be used with block indexes

This diagram shows how the result of two separate block index scans are combined to determine the range of values that meet the predicate restrictions.

Based on the predicates from the SELECT statement, two separate dimension block index scans are done; one for the blue slice, and another for the USA slice. A logical OR operation is done in memory in order to find the union of the two slices, and determine the combined set of blocks found in both slices (including the removal of duplicate blocks).

Once DB2 has list of blocks to scan, DB2 can do a mini-relational scan of each block. Prefetching of the blocks can be done, and will involve just one I/O per block, as each block is stored as an extent on disk and can be read into the buffer pool as a unit. If predicates need to be applied to the data, dimension predicates need only be applied to one record in the block, because all records in the block are guaranteed to have the same dimension key values. If other predicates are present, DB2 only needs to check these on the remaining records in the block.

MDC tables also support regular RID-based indexes. Both RID and block indexes can be combined using a logical AND operation, or a logical OR operation, with the index. Block indexes provide the optimizer with additional access plans to choose from, and do not prevent the use of traditional access plans (RID scans, joins, table scans, and others). Block index plans will be costed by the optimizer along with all other possible access plans for a particular query, and the most inexpensive plan will be chosen.

The DB2 Design Advisor can help to recommend RID-based indexes on MDC tables, or to recommend MDC dimensions for a table.

Related concepts:

- “Block index considerations for MDC tables” on page 188
- “Block indexes” on page 175
- “Block maps” on page 185

Maintaining clustering automatically during INSERT operations

Automatic maintenance of data clustering in an MDC table is ensured using the composite block index. It is used to dynamically manage and maintain the physical clustering of data along the dimensions of the table over the course of INSERT operations. A key is found in this composite block index only for each of those logical cells of the table that contain records. This block index is therefore used during an INSERT to quickly and efficiently determine if a logical cell exists in the table, and only if so, determine exactly which blocks contain records having that cell’s particular set of dimension values.

When an insert occurs:

- The composite block index is probed for the logical cell corresponding to the dimension values of the record to be inserted.
- If the key of the logical cell is found in the index, its list of block ID (BIDs) gives the complete list of blocks in the table having the dimension values of the logical cell. (See Figure 53 on page 184.) This limits the numbers of extents of the table to search for space to insert the record.
- If the key of the logical cell is not found in the index; or, if the extents containing these values are full, a new block is assigned to the logical cell. If possible, the reuse of an empty block in the table occurs first before extending the table by another new extent of pages (a new block).

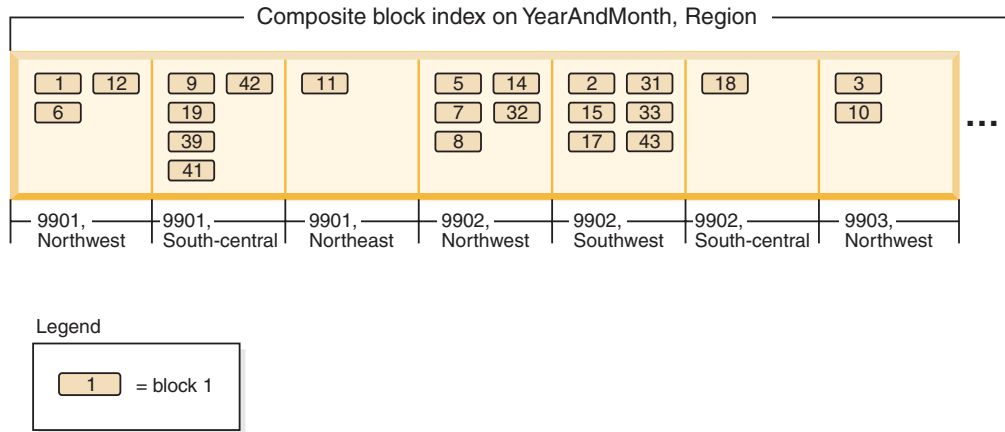


Figure 53. Composite block index on 'YearAndMonth', 'Region'

Data records having particular dimension values are guaranteed to be found in a set of blocks that contain only and all the records having those values. Blocks are made up of consecutive pages on disk. As a result, access to these records is sequential, providing clustering. This clustering is automatically maintained over time by ensuring that records are only inserted into blocks from cells with the record's dimension values. When existing blocks in a logical cell are full, an empty block is reused or a new block is allocated and added to the set of blocks for that logical cell. When a block is emptied of data records, the block ID (BID) is removed from the block indexes. This disassociates the block from any logical cell values so that it can be reused by another logical cell in the future. Thus, cells and their associated block index entries are dynamically added and removed from the table as needed to accommodate only the data that exists in the table. The composite block index is used to manage this, because it maps logical cell values to the blocks containing records having those values.

Because clustering is automatically maintained in this way, reorganization of an MDC table is never needed to re-cluster data. However, reorganization can still be used to reclaim space. For example, if cells have many sparse blocks where data could fit on fewer blocks, or if the table has many pointer-overflow pairs, a reorganization of the table would compact records belonging to each logical cell into the minimum number of blocks needed, as well as remove pointer-overflow pairs.

The following example illustrates how the composite block index can be used for query processing. If you want to find all records in the Sales table having "Region" of 'Northwest' and "YearAndMonth" of '9903', DB2 would look up the key value 9903, Northwest in the composite block index, as shown in Figure 54 on page 185. The key is made up a key value, namely '9903, Northwest', and a list of BIDs. You can see that the only BIDs listed are 3 and 10, and indeed there are only two blocks in the Sales table containing records having these two particular values.




Figure 54. Key from composite block index on 'YearAndMonth', 'Region'

To illustrate the use of the composite block index during insert, take the example of inserting another record with dimension values 9903 and Northwest. DB2 would look up this key value in the composite block index and find BIDs for blocks 3 and 10. These blocks contain all records and the only records having these dimension key values. If there is space available, DB2 inserts the new record into one of these blocks. If there is no space on any pages in these blocks, DB2 allocates a new block for the table, or uses a previously emptied block in the table. Note that, in this example, block 48 is currently not in use by the table. DB2 inserts the record into the block and associates this block to the current logical cell by adding the BID of the block to the composite block index and to each dimension block index. See Figure 55 for an illustration of the keys of the dimension block indexes after the addition of Block 48.




Figure 55. Keys from the dimension block indexes after addition of Block 48

Related concepts:

- "Block maps" on page 185
- "Block indexes" on page 175

Block maps

When a block is emptied, it is disassociated from its current logical cell values by removing its BID from the block indexes. The block can then be reused by another logical cell. This reduces the need to extend the table with new blocks. When a new block is needed, previously emptied blocks need to be found quickly without having to search the table for them.

The block map is a new structure used to facilitate locating empty blocks in the MDC table. The block map is stored as a separate object:

- In SMS, as a separate .BKM file
- In DMS, as a new object descriptor in the object table.

The block map is an array containing an entry for each block of the table. Each entry comprises a set of status bits for a block. The status bits include:

- In use. The block is assigned to a logical cell.
- Load. The block is recently loaded; not yet visible by scans.
- Constraint. The block is recently loaded; constraint checking is still to be done.
- Refresh. The block is recently loaded; materialized query views still need to be refreshed.




Figure 56. How a block map works

In Figure 56, the left side shows the block map array with different entries for each block in the table. The right side shows how each extent of the table is being used: some are free, most are in use, and records are only found in blocks marked in use in the block map. For simplicity, only one of the two dimension block indexes is shown in the diagram.

Notes:

1. There are pointers in the block index only to blocks which are marked IN USE in the block map.
2. The first block is reserved. This block contains system records for the table.

Free blocks are found easily for use in a cell, by scanning the block map for FREE blocks, that is, those without any bits set.

Table scans also use the block map to access only extents currently containing data. Any extents not in use do not need to be included in the table scan at all. To illustrate, a table scan in this example (Figure 56) would start from the third extent

(extent 2) in the table, skipping the first reserved extent and the following empty extent, scan blocks 2, 3 and 4 in the table, skip the next extent (not touching any of that extent's data pages), and then continue scanning from there.

Related concepts:

- “Block index considerations for MDC tables” on page 188
- “Block indexes” on page 175
- “Block indexes and query performance” on page 180

Deletion from an MDC table

When a record is deleted in an MDC table, if it is not the last record in the block, the DB2 database system merely deletes the record and removes its RID from any record-based indexes defined on the table. When a delete removes the last record in a block, however, DB2 frees the block by changing its IN_USE status bit and removing the block's BID from all block indexes. Again, if there are record-based indexes as well, the RID is removed from them.

Note: Therefore, block index entries need only be removed once per entire block and only if the block is completely emptied, instead of once per deleted row in a record-based index.

Related concepts:

- “Multidimensional clustering tables” on page 172

Updating an MDC table

In an MDC table, updates of non-dimension values are done in place just as they are done with regular tables. If the update affects a variable length column and the record no longer fits on the page, another page with sufficient space is found. The search for this new page begins within the same block. If there is no space in that block, the algorithm to insert a new record is used to find a page in the logical cell with enough space. There is no need to update the block indexes, unless no space is found in the cell and a new block needs to be added to the cell.

Updates of dimension values are treated as a delete of the current record followed by an insert of the changed record, because the record is changing the logical cell to which it belongs. If the deletion of the current record causes a block to be emptied, the block index needs to be updated. Similarly, if the insert of the new record requires it to be inserted into a new block, the block index needs to be updated.

Block indexes only need to be updated when inserting the first record into a block or when deleting the last record from a block. Index overhead associated with block indexes for maintenance and logging is therefore much less than the index overhead associated with regular indexes. For every block index that would have otherwise been a regular index, the maintenance and logging overhead is greatly reduced.

MDC tables are treated like any existing table; that is, triggers, referential integrity, views, and materialized query tables can all be defined upon them.

Related concepts:

- “Multidimensional clustering tables” on page 172

Load considerations for MDC tables

If you roll data in to your data warehouse on a regular basis, you can use MDC tables to your advantage. In MDC tables, load will first reuse previously emptied blocks in the table before extending the table and adding new blocks for the remaining data. After you have deleted a set of data, for example, all the data for a month, you can use the load utility to roll in the next month of data and it can reuse the blocks that have been emptied after the (committed) deletion.

When loading data into MDC tables, the input data can be either sorted or unsorted. If unsorted, consider doing the following:

- Increase the *util_heap* configuration parameter.
Increasing the utility heap size will affect all load operations in the database (as well as backup and restore operations).
- Increase the value given with the DATA BUFFER clause of the LOAD command.
Increasing this value will affect a single load request. The utility heap size must be large enough to accommodate the possibility of multiple concurrent load requests.
- Ensure the page size used for the buffer pool is the same as the largest page size for the temporary table space.

Load begins at a block boundary, so it is best used for data belonging to new cells or for the initial populating of a table.

Related concepts:

- “Multidimensional clustering tables” on page 172

Logging considerations for MDC tables

In cases where columns previously or otherwise indexed by RID indexes are now dimensions and so are indexed with block indexes, index maintenance and logging are significantly reduced. Only when the last record in an entire block is deleted does DB2 need to remove the BID from the block indexes and log this index operation. Similarly, only when a record is inserted to a new block (if it is the first record of a logical cell or an insert to a logical cell of currently full blocks) does DB2 need to insert a BID in the block indexes and log that operation. Because blocks can be between 2 and 256 pages of records, this block index maintenance and logging will be relatively small. Inserts and deletes to the table and to RID indexes will still be logged.

Related concepts:

- “Multidimensional clustering tables” on page 172

Block index considerations for MDC tables

When you define dimensions for an MDC table, dimension block indexes are created. In addition, a composite block index may also be created when multiple dimensions are defined. If you have defined only one dimension for your MDC table, however, DB2 will create only one block index, which will serve both as the dimension block index and as the composite block index. Similarly, if you create an

MDC table that has dimensions on column A, and on (column A, column B), DB2 will create a dimension block index on column A and a dimension block index on column A, column B. Because a composite block index is a block index of all the dimensions in the table, the dimension block index on column A, column B will also serve as the composite block index.

The composite block index is also used in query processing to access data in the table having specific dimension values. Note that the order of key parts in the composite block index may affect its use or applicability for query processing. The order of its key parts is determined by the order of columns found in the entire ORGANIZE BY DIMENSIONS clause used when creating the MDC table. For example, if a table is created using the statement

```
CREATE TABLE t1 (c1 int, c2 int, c3 int, c4 int)
  ORGANIZE BY DIMENSIONS (c1, c4, (c3,c1), c2)
```

then the composite block index will be created on columns (c1,c4,c3,c2). Note that although c1 is specified twice in the dimensions clause, it is used only once as a key part for the composite block index, and in the order in which it is first found. The order of key parts in the composite block index makes no difference for insert processing, but may do so for query processing. Therefore, if it is more desirable to have the composite block index with column order (c1,c2,c3,c4), then the table should be created using the statement

```
CREATE TABLE t1 (c1 int, c2 int, c3 int, c4 int)
  ORGANIZE BY DIMENSIONS (c1, c2, (c3,c1), c4)
```

Related concepts:

- “Block indexes” on page 175
- “Block indexes and query performance” on page 180
- “Block maps” on page 185

Designing multidimensional clustering (MDC) tables

Once you have decided to work with multidimensional clustering tables, the dimensions that you choose will depend not only on the type of queries that will use the tables and benefit from block-level clustering, but even more importantly on the amount and distribution of your actual data. Aspects of designing MDC tables and some guidance regarding the selection of appropriate dimensions and block sizes can be seen using the related concepts links.

Queries that will benefit from MDC:

The first consideration when choosing clustering dimensions for your table is the determination of which queries will benefit from clustering at a block level. Typically, there will be several candidates when choosing dimensions based on the queries that make up the work to be done on the data. The ranking of these candidates is important. Columns, especially those with low cardinalities, that are involved in equality or range predicate queries will show the greatest benefit from, and should be considered as candidates for, clustering dimensions. You will also want to consider creating dimensions for foreign keys in an MDC fact table involved in star joins with dimension tables. You should keep in mind the performance benefits of automatic and continuous clustering on more than one dimension, and of clustering at an extent or block level.

There are many queries that can take advantage of multidimensional clustering. Examples of such queries follow. In some of these examples, assume that there is an MDC table t1 with dimensions c1, c2, and c3. In the other examples, assume that there is an MDC table mdctable with dimensions color and nation.

Example 1:

```
SELECT .... FROM t1 WHERE c3 < 5000
```

This query involves a range predicate on a single dimension, so it can be internally rewritten to access the table using the dimension block index on c3. The index is scanned for block identifiers (BIDs) of keys having values less than 5000, and a mini-relational scan is applied to the resulting set of blocks to retrieve the actual records.

Example 2:


```
SELECT .... FROM t1 WHERE c2 IN (1,2037)
```

This query involves an IN predicate on a single dimension, and can trigger block index based scans. This query can be internally rewritten to access the table using the dimension block index on c2. The index is scanned for BIDs of keys having values of 1 and 2037, and a mini-relational scan is applied to the resulting set of blocks to retrieve the actual records.

Example 3:


```
SELECT * FROM MDCTABLE WHERE COLOR='BLUE' AND NATION='USA'
```

Key from the dimension block index on Colour



+ (AND)

Key from the dimension block index on Nation



=

Resulting block ID (BID) list of blocks to scan




Figure 57. A query request that uses a logical AND operation with two block indexes

To carry out this query request, the following is done (and is shown in Figure 57):

- A dimension block index lookup is done: one for the Blue slice and another for the USA slice.
- A block logical AND operation is carried out to determine the intersection of the two slices. That is, the logical AND operation determines only those blocks that are found in both slices.
- A mini-relation scan of the resulting blocks in the table is carried out.

Example 4:

```
SELECT ... FROM t1
WHERE c2 > 100 AND c1 = '16/03/1999' AND c3 > 1000 AND c3 < 5000
```

This query involves range predicates on c2 and c3 and an equality predicate on c1, along with a logical AND operation. This can be internally rewritten to access the table on each of the dimension block indexes:

- A scan of the c2 block index is done to find BIDs of keys having values greater than 100

- A scan of the c3 block index is done to find BIDs of keys having values between 1000 and 5000
- A scan of the c1 block index is done to find BIDs of keys having the value '16/03/1999'.

A logical AND operation is then done on the resulting BIDs from each block scan, to find their intersection, and a mini-relational scan is applied to the resulting set of blocks to find the actual records.

Example 5:

```
SELECT * FROM MDCTABLE WHERE COLOR='BLUE' OR NATION='USA'
```

To carry out this query request, the following is done:

- A dimension block index lookup is done: one for each slice.
- A logical OR operation is done to find the union of the two slices.
- A mini-relation scan of the resulting blocks in the table is carried out.

Example 6:

```
SELECT .... FROM t1 WHERE c1 < 5000 OR c2 IN (1,2,3)
```

This query involves a range predicate on the c1 dimension and a IN predicate on the c2 dimension, as well as a logical OR operation. This can be internally rewritten to access the table on the dimension block indexes c1 and c2. A scan of the c1 dimension block index is done to find values less than 5000 and another scan of the c2 dimension block index is done to find values 1, 2, and 3. A logical OR operation is done on the resulting BIDs from each block index scan, then a mini-relational scan is applied to the resulting set of blocks to find the actual records.

Example 7:

```
SELECT .... FROM t1 WHERE c1 = 15 AND c4 < 12
```

This query involves an equality predicate on the c1 dimension and another range predicate on a column that is not a dimension, along with a logical AND operation. This can be internally rewritten to access the dimension block index on c1, to get the list of blocks from the slice of the table having value 15 for c1. If there is a RID index on c4, an index scan can be done to retrieve the RIDs of records having c4 less than 12, and then the resulting list of blocks undergoes a logical AND operation with this list of records. This intersection eliminates RIDs not found in the blocks having c1 of 15, and only those listed RIDs found in the blocks that qualify are retrieved from the table.


If there is no RID index on c4, then the block index can be scanned for the list of qualifying blocks, and during the mini-relational scan of each block, the predicate c4 < 12 can be applied to each record found.

Example 8:

Given a scenario where there are dimensions for color, year, nation and a row ID (RID) index on the part number, the following query is possible.


```
SELECT * FROM MDCTABLE WHERE COLOR='BLUE' AND PARTNO < 1000
```

Key from the dimension block index on Colour



+ (AND)

Row IDs (RID) from RID index on Partno



=

Resulting row IDs to fetch




Figure 58. A query request that uses a logical AND operation on a block index and a row ID (RID) index


To carry out this query request, the following is done (and is shown in Figure 58):

- A dimension block index lookup and a RID index lookup are done.
- A logical AND operation is used with the blocks and RIDs to determine the intersection of the slice and those rows meeting the predicate condition.
- The result is only those RIDs that also belong to the qualifying blocks.

Example 9:

```
SELECT * FROM MDCTABLE WHERE COLOR='BLUE' OR PARTNO < 1000
```

Key from the dimension block index on Colour



+ (OR)

Row IDs (RID) from RID index on Partno



=

Resulting blocks and RIDs to fetch



Figure 59. How block index and row ID using a logical OR operation works

To carry out this query request, the following is done (and is shown in Figure 59):

- A dimension block index lookup and a RID index lookup are done.
- A logical OR operation is used with the blocks and RIDs to determine the union of the slice and those rows meeting the predicate condition.
- The result is all of the rows in the qualifying blocks, plus additional RIDs that fall outside the qualifying blocks that meet the predicate condition. A mini-relational scan of each of the blocks is performed to retrieve their records, and the additional records outside these blocks are retrieved individually.

Example 10:

```
SELECT ... FROM t1 WHERE c1 < 5 OR c4 = 100
```

This query involves a range predicate on dimension c1 and an equality predicate on a non-dimension column c4, as well as a logical OR operation. If there is a RID

index on the c4 column, this may be internally rewritten to do a logical OR operation using the dimension block index on c1 and the RID index on c4. If there is no index on c4, a table scan may be chosen instead, since all records must be checked. The logical OR operation would use a block index scan on c1 for values less than 4, as well as a RID index scan on c4 for values of 100. A mini-relational scan is performed on each block that qualifies, because all records within those blocks will qualify, and any additional RIDs for records outside of those blocks are retrieved as well.

Example 11:

```
SELECT .... FROM t1,d1,d2,d3
WHERE t1.c1 = d1.c1 and d1.region = 'NY'
      AND t2.c2 = d2.c3 and d2.year='1994'
      AND t3.c3 = d3.c3 and d3.product='basketball'
```

This query involves a star join. In this example, t1 is the fact table and it has foreign keys c1, c2, and c3, corresponding to the primary keys of d1, d2, and d3, the dimension tables. The dimension tables do not have to be MDC tables. Region, year, and product are columns of the respective dimension tables which can be indexed using regular or block indexes (if the dimension tables are MDC tables). When accessing the fact table on c1, c2, and c3 values, block index scans of the dimension block indexes on these columns can be done, followed by a logical AND operation using the resulting BIDs. When there is a list of blocks, a mini-relational scan can be done on each block to get the records.

Density of cells:

The choices made for the appropriate dimensions and for the extent size are of **critical** importance to MDC design. These factors determine the table's expected cell density. They are important because an extent is allocated for every existing cell, regardless of the number of records in the cell. The right choices will take advantage of block-based indexing and multidimensional clustering, resulting in performance gains. The goal is to have densely-filled blocks to get the most benefit from multidimensional clustering, and to get optimal space utilization.

Thus, a very important consideration when designing a multidimensional table is the expected density of cells in the table, based on present and anticipated data. You can choose a set of dimensions, based on query performance, that cause the potential number of cells in the table to be very large, based on the number of possible values for each of the dimensions. The number of possible cells in the table is equal to the Cartesian product of the cardinalities of each of the dimensions. For example, if you cluster the table on dimensions Day, Region and Product and the data covers 5 years, you might have 1821 days * 12 regions * 5 products = 109 260 different possible cells in the table. Any cell that contains only a few records will still require an entire block of pages allocated to it, in order to store the records for that cell. If the block size is large, this table could end up being much larger than it really needs to be.

There are several design factors that can contribute to optimal cell density:

- Varying the number of dimensions.
- Varying the granularity of one or more dimensions.
- Varying the block (extent) size and page size of the table space.

Carry out the following steps to achieve the best design possible:

1. Identify candidate dimensions.

Determine which queries will benefit from block-level clustering. Examine the potential workload for columns which have some or all of the following characteristics:

- Range and equality of any IN-list predicates
- Roll-in or roll-out of data
- Group-by and order-by clauses
- Join clauses (especially in star schema environments).

2. Estimate the number of cells.

Identify how many potential cells are possible in a table organized along a set of candidate dimensions. Determine the number of unique combinations of the dimension values that occur in the data. If the table exists, an exact number can be determined for the current data by simply selecting the number of distinct values in each of the columns that will be dimensions for the table.

Alternatively, an approximation can be determined if you only have the statistics for a table, by multiplying the column cardinalities for the dimension candidates.

Note: If your table is in a partitioned database environment, and the distribution key is not related to any of the dimensions considered, you will have to determine an average amount of data per cell by taking all of the data and dividing by the number of database partitions.

3. Estimate the space occupancy or density.

On average, consider that each cell has one partially-filled block where only a few rows are stored. There will be more partially-filled blocks as the number of rows per cell becomes smaller. Also, note that on average (assuming little or no data skew), the number of records per cell can be found by dividing the number of records in the table by the number of cells. However, if your table is in a partitioned database environment, you need to consider how many records there are per cell on each database partition, as blocks are allocated for data on a database partition basis. When estimating the space occupancy and density in a partitioned database environment, you need to consider the number of records per cell on average on each database partition, not across the entire table. See the section called “Multidimensional clustering (MDC) table creation, placement, and use” for more information.

There are several ways to improve the density:

- Reduce the block size so that partially-filled blocks take up less space.

Reduce the size of each block by making the extent size appropriately small. Each cell that has a partially-filled block, or that contains only one block with few records on it, wastes less space. The trade-off, however, is that for those cells having many records, more blocks are needed to contain them. This increases the number of block identifiers (BIDs) for these cells in the block indexes, making these indexes larger and potentially resulting in more inserts and deletes to these indexes as blocks are more quickly emptied and filled. It also results in more small groupings of clustered data in the table for these more populated cell values, versus a smaller number of larger groupings of clustered data.

- Reduce the number of cells by reducing the number of dimensions, or by increasing the granularity of the cells with a generated column.

You can roll up one or more dimensions to a coarser granularity in order to give it a lower cardinality. For example, you can continue to cluster the data in the previous example on Region and Product, but replace the dimension of Day with a dimension of YearAndMonth. This gives cardinalities of 60 (12 months times 5 years), 12, and 5 for YearAndMonth, Region, and Product,

with a possible number of cells of 3600. Each cell then holds a greater range of values and is less likely to contain only a few records.

You should also take into account predicates commonly used on the columns involved, such as whether many are on Month of Date, or Quarter, or Day. This affects the desirability of changing the granularity of the dimension. If, for example, most predicates are on particular days and you have clustered the table on Month, DB2 Database for Linux, UNIX, and Windows can use the block index on YearAndMonth to quickly narrow down which months contain the days desired and access only those associated blocks. When scanning the blocks, however, the Day predicate must be applied to determine which days qualify. However, if you cluster on Day (and Day has high cardinality), the block index on Day can be used to determine which blocks to scan, and the Day predicate only has to be reapplied to the first record of each cell that qualifies. In this case, it may be better to consider rolling up one of the other dimensions to increase the density of cells, as in rolling up the Region column, which contains 12 different values, to Regions West, North, South and East, using a user-defined function.

Related concepts:

- “The Design Advisor” in *Performance Guide*
- “Multidimensional clustering (MDC) table creation, placement, and use” on page 197
- “Multidimensional clustering tables” on page 172

Multidimensional clustering (MDC) table creation, placement, and use

There are many factors that should be considered when creating MDC tables. The following sections discuss how your decisions on how to create, place, and use your MDC tables could be influenced by your current database environment (for example, whether you have a partitioned database or not), and by your choices of dimensions for your MDC table. Also discussed is the DB2 Design Advisor, and how it can be used to provide advice on some of these issues.

Moving data from an existing table to a multidimensional clustering (MDC) table:

To improve query performance and reduce the overhead of data maintenance operations in a data warehouse or large database environment, you can move data from regular tables into multidimensional clustering (MDC) tables. To move data from an existing table to an MDC table: export your data, drop the original table (optional), create a multidimensional clustering (MDC) table (using the CREATE TABLE statement with the ORGANIZE BY DIMENSIONS clause), and load the MDC table with your data.

An ALTER TABLE procedure called SYSPROC.ALTOBJ can be used to carry out the translation of data from an existing table to an MDC table. The procedure is called from the DB2 Design Advisor. The time required to translate the data between the tables can be significant and depends on the size of the table and the amount of data that needs to be translated.

The ALTOBJ procedure does the following when altering a table:

- Drop all dependent objects of the table
- Rename the table
- Create the table using the new definition

- Recreate all dependent objects of the table
- Transform existing data in the table into the data required in the new table. That is, the selecting of data from the old table and loading that data into the new one where column functions may be used to transform from a old data type to a new data type.

Multidimensional clustering (MDC) tables in SMS table spaces:

If you plan to store MDC tables in an SMS table space, we strongly recommend that you use multipage file allocation.

Note: Multipage file allocation is the default for newly created databases in Version 8.2 and later.

The reason for this recommendation is that MDC tables are always extended by whole extents, and it is important that all the pages in these extents are physically consecutive. Therefore, there are no space advantage to disabling multipage file allocation; and furthermore, enabling it will significantly increase the chances that the pages in each extent are physically consecutive.

MDC Advisor feature on the DB2 Design Advisor:

The DB2 Design Advisor (db2advis), formerly known as the Index Advisor, has an MDC feature. This feature recommends clustering dimensions for use in an MDC table, including coarsifications on base columns in order to improve workload performance. The term *coarsification* refers to a mathematic expression to reduce the cardinality (the number of distinct values) of a clustering dimension. A common example of a coarsification is the date where coarsification could be by date, week of the date, month of the date, or quarter of the year.

A requirement to use the MDC feature of the DB2 Design Advisor is the existence of at least several extents of data within the database. The DB2 Design Advisor uses the data to model data density and cardinality.

If the database does not have data in the tables, the DB2 Design Advisor will not recommend MDC, even if the database contains empty tables but has a mocked up set of statistics to imply a populated database.

The recommendation includes identifying potential generated columns that define coarsification of dimensions. The recommendation does not include possible block sizes. The extent size of the table space is used when making recommendations for MDC tables. The assumption is that the recommended MDC table will be created in the same table space as the existing table, and will therefore have the same extent size. The recommendations for MDC dimensions would change depending on the extent size of the table space since the extent size impacts the number of records that can fit into a block or cell. This directly affects the density of the cells.

Only single-column dimensions, and not composite-column dimensions, are considered, although single or multiple dimensions may be recommended for the table. The MDC feature will recommend coarsifications for most supported data types with the goal of reducing the cardinality of cells in the resulting MDC solution. The data type exceptions include: CHAR, VARCHAR, GRAPHIC, and VARGRAPH data types. All supported data types are cast to INTEGER and are coarsified through a generated expression.

The goal of the MDC feature of the DB2 Design Advisor is to select MDC solutions that result in improved performance. A secondary goal is to keep the storage expansion of the database constrained to a modest level. A statistical method is used to determine the maximum storage expansion on each table.

The analysis operation within the advisor includes not only the benefits of block index access but also the impact of MDC on insert, update, and delete operations against dimensions of the table. These actions on the table have the potential to cause records to be moved between cells. The analysis operation also models the potential performance impact of any table expansion resulting from the organization of data along particular MDC dimensions.

The MDC feature is enabled using the `-m <advise type>` flag on the `db2advvis` utility. The “C” advise type is used to indicate multidimensional clustering tables. The advise types are: “I” for index, “M” for materialized query tables, “C” for MDC, and “P” for partitioned database environment. The advise types can be used in combination with each other.

Note: The DB2 Design Advisor will not explore tables that are less than 12 extents in size.

The advisor will analyze both MQTs and regular base tables when coming up with recommendations.

The output from the MDC feature includes:

- Generated column expressions for each table for coarsified dimensions that appear in the MDC solution.
- An ORGANIZE BY clause recommended for each table.

The recommendations are reported both to stdout and to the ADVISE tables that are part of the explain facility.

Multidimensional clustering (MDC) tables and a partitioned database environment:

Multidimensional clustering can be used in conjunction with a partitioned database environment. In fact, MDC can complement a partitioned database environment. A partitioned database environment is used to distribute data from a table across multiple physical or logical nodes in order to:

- Take advantage of multiple machines to increase processing requests in parallel.
- Increase the physical size of the table beyond a single database partition’s limits.
- Improve the scalability of the database.

The reason for distributing a table is independent of whether the table is an MDC table or a regular table. For example, the rules for the selection of columns to make up the distribution key are the same. The distribution key for an MDC table can involve any column, whether those columns make up part of a dimension of the table or not.

If the distribution key is identical to a dimension from the table, then each database partition will contain a different portion of the table. For instance, if our example MDC table is distributed by color across two database partitions, then the Color column will be used to divide the data. As a result, the Red and Blue slices may be found on one database partition and the Yellow slice on the other. If the distribution key is not identical to the dimensions from the table, then each

database partition will have a subset of data from each slice. When choosing dimensions and estimating cell occupancy (see the section called “Density of cells”), note that on average the total amount of data per cell is determined by taking all of the data and dividing by the number of database partitions.

Multidimensional clustering (MDC) tables with multiple dimensions:

If you know that certain predicates will be heavily used in queries, you can cluster the table on the columns involved, using the ORGANIZE BY DIMENSIONS clause.

Example 1:

```
CREATE TABLE T1 (c1 DATE, c2 INT, c3 INT, c4 DOUBLE)
  ORGANIZE BY DIMENSIONS (c1, c3, c4)
```

The table in Example 1 is clustered on the values within three native columns forming a logical cube (that is, having three dimensions). The table can now be logically sliced up during query processing on one or more of these dimensions such that only the blocks in the appropriate slices or cells will be processed by the relational operators involved. Note that the size of a block (the number of pages) will be the extent size of the table.

Multidimensional clustering (MDC) tables with dimensions based on more than one column:

Each dimension can be made up of one or more columns. As an example, you can create a table that is clustered on a dimension containing two columns.

Example 2:

```
CREATE TABLE T1 (c1 DATE, c2 INT, c3 INT, c4 DOUBLE)
  ORGANIZE BY DIMENSIONS (c1, (c3, c4))
```

In Example 2, the table will be clustered on two dimensions, c1 and (c3,c4). Thus, in query processing, the table can be logically sliced up on either the c1 dimension, or on the composite (c3, c4) dimension. The table will have the same number of blocks as the table in Example 1, but one less dimension block index. In Example 1, there will be three dimension block indexes, one for each of the columns c1, c3, and c4. In Example 2, there will be two dimension block indexes, one on the column c1 and the other on the columns c3 and c4. The main differences between these two approaches is that, in Example 1, queries involving just c4 can use the dimension block index on c4 to quickly and directly access blocks of relevant data. In Example 2, c4 is a second key part in a dimension block index, so queries involving just c4 involve more processing. However, in Example 2 DB2 Database for Linux, UNIX, and Windows will have one less block index to maintain and store.

The DB2 Design Advisor does not make recommendations for dimensions containing more than one column.

Multidimensional clustering (MDC) tables with column expressions as dimensions:

Column expressions can also be used for clustering dimensions. The ability to cluster on column expressions is useful for rolling up dimensions to a coarser granularity, such as rolling up an address to a geographic location or region, or rolling up a date to a week, month, or year. In order to implement the rolling up of dimensions in this way, you can use generated columns. This type of column

definition will allow the creation of columns using expressions that can represent dimensions. In Example 3, the statement creates a table clustered on one base column and two column expressions.

Example 3:

```
CREATE TABLE T1(c1 DATE, c2 INT, c3 INT, c4 DOUBLE,  
  c5 DOUBLE GENERATED ALWAYS AS (c3 + c4),  
  c6 INT GENERATED ALWAYS AS (MONTH(C1))  
  ORGANIZE BY DIMENSIONS (c2, c5, c6)
```

In Example 3, column c5 is an expression based on columns c3 and c4, while column c6 rolls up column c1 to a coarser granularity in time. This statement will cluster the table based on the values in columns c2, c5, and c6.

Range queries on a generated column dimension require monotonic column functions:

Expressions must be monotonic to derive range predicates for dimensions on generated columns. If you create a dimension on a generated column, queries on the base column will be able to take advantage of the block index on the generated column to improve performance, with one exception. For range queries on the base column (date, for example) to use a range scan on the dimension block index, the expression used to generate the column in the CREATE TABLE statement must be monotonic. Although a column expression can include any valid expression (including user-defined functions (UDFs)), if the expression is non-monotonic, only equality or IN predicates are able to use the block index to satisfy the query when these predicates are on the base column.

As an example, assume that we create an MDC table with dimensions on the generated column month, where `month = INTEGER (date)/100`. For queries on the dimension (month), block index scans can be done. For queries on the base column (date), block index scans can also be done to narrow down which blocks to scan, and then apply the predicates on date to the rows in those blocks only.

The compiler generates additional predicates to be used in the block index scan. For example, with the query:

```
SELECT * FROM MDCTABLE WHERE DATE > "1999/03/03" AND DATE < "2000/01/15"
```

the compiler generates the additional predicates: “month >= 199903” and “month < 200001” which can be used as predicates for a dimension block index scan. When scanning the resulting blocks, the original predicates are applied to the rows in the blocks.

A non-monotonic expression will only allow equality predicates to be applied to that dimension. A good example of a non-monotonic function is MONTH() as seen in the definition of column c6 in Example 3. If the c1 column is a date, timestamp, or valid string representation of a date or timestamp, then the function returns an integer value in the range of 1 to 12. Even though the output of the function is deterministic, it actually produces output similar to a step function (that is, a cyclic pattern):

```
MONTH(date('99/01/05')) = 1  
MONTH(date('99/02/08')) = 2  
MONTH(date('99/03/24')) = 3  
MONTH(date('99/04/30')) = 4  
...
```

```
MONTH(date('99/12/09')) = 12
MONTH(date('00/01/18')) = 1
MONTH(date('00/02/24')) = 2
...
```

Although date in this example is continually increasing, MONTH(date) is not. More specifically, it is not guaranteed that whenever date1 is larger than date2, MONTH(date1) is greater than or equal to MONTH(date2). It is this condition that is required for monotonicity. This non-monotonicity is allowed, but it limits the dimension in that a range predicate on the base column cannot generate a range predicate on the dimension. However, a range predicate on the expression is fine, for example, where month(c1) between 4 and 6. This can use the index on the dimension in the usual way, with a starting key of 4 and a stop key of 6.

To make this function monotonic, you would have to include the year as the high order part of the month. DB2 V9.1 provides an extension to the INTEGER built-in function to help in defining a monotonic expression on date. INTEGER(date) returns an integer representation of the date, which then can be divided to find an integer representation of the year and month. For example, INTEGER(date('2000/05/24')) returns 20000524, and therefore INTEGER(date('2000/05/24'))/100 = 200005. The function INTEGER(date)/100 is monotonic.

Similarly, the built-in functions DECIMAL and BIGINT also have extensions so that you can derive monotonic functions. DECIMAL(timestamp) returns a decimal representation of a timestamp, and this can be used in monotonic expressions to derive increasing values for month, day, hour, minute, and so on. BIGINT(date) returns a big integer representation of the date, similar to INTEGER(date).

DB2 will determine the monotonicity of an expression, where possible, when creating the generated column for the table, or when creating a dimension from an expression in the dimensions clause. Certain functions can be recognized as monotonicity-preserving, such as DATENUM(), DAYS(), YEAR(). Also, various mathematical expressions such as division, multiplication, or addition of a column and a constant are monotonicity-preserving. Where DB2 determines that an expression is not monotonicity-preserving, or if it cannot determine this, the dimension will only support the use of equality predicates on its base column.

Related concepts:

- “Multidimensional clustering considerations when loading data” in *Data Movement Utilities Guide and Reference*
- “Designing multidimensional clustering (MDC) tables” on page 189
- “Extent size” on page 144
- “Multidimensional clustering tables” on page 172

Related tasks:

- “Defining dimensions on a table” in *Administration Guide: Implementation*

Related reference:

- “CREATE TABLE statement” in *SQL Reference, Volume 2*
- “db2empfa - Enable multipage file allocation command” in *Command Reference*

Chapter 6. Designing partitioned databases

This chapter discusses issues related to managing transactions when working with partitioned databases. Updating a single database is discussed first. This is followed by a discussion of more complex considerations associated with transactions that use multiple databases. The concept of transaction managers is also introduced, as well as considerations regarding updating a database from a host or iSeries. This chapter also discusses two-phase commit when managing a multi-site update and error recovery when working with transactions using two-phase commit.

Updating a single database in a transaction

The simplest form of transaction is to read from and write to only one database within a single unit of work. This type of database access is called a *remote unit of work*.




Figure 60. Using a single database in a transaction

Figure 60 shows a database client running a funds transfer application that accesses a database containing checking and savings account tables, as well as a banking fee schedule. The application must:

- Accept the amount to transfer from the user interface
- Subtract the amount from the savings account, and determine the new balance
- Read the fee schedule to determine the transaction fee for a savings account with the given balance
- Subtract the transaction fee from the savings account
- Add the amount of the transfer to the checking account
- Commit the transaction (unit of work).

Procedure:

To set up such an application, you must do the following as part of the preparation to carry out the transaction within the environment:

1. Create the tables for the savings account, checking account and banking fee schedule in the same database

2. If physically remote, set up the database server to use the appropriate communications protocol
3. If physically remote, catalog the node and the database to identify the database on the database server
4. Precompile your application program to specify a type 1 connection; that is, specify CONNECT 1 (the default) on the PRECOMPILE PROGRAM command.

Related concepts:

- “Units of work” on page 22

Related tasks:

- “Updating a single database in a multi-database transaction” on page 204
- “Updating multiple databases in a transaction” on page 205

Related reference:

- “PRECOMPILE command” in *Command Reference*

Using multiple databases in a single transaction

When using multiple databases in a single transaction, the requirements for setting up and administering your environment are different depending on the number of databases that are being updated in the transaction. The following topics discuss these requirements.

Updating a single database in a multi-database transaction

If your data is distributed across multiple databases, you may wish to update one database while reading from one or more other databases. This type of access can be performed within a single unit of work (transaction).

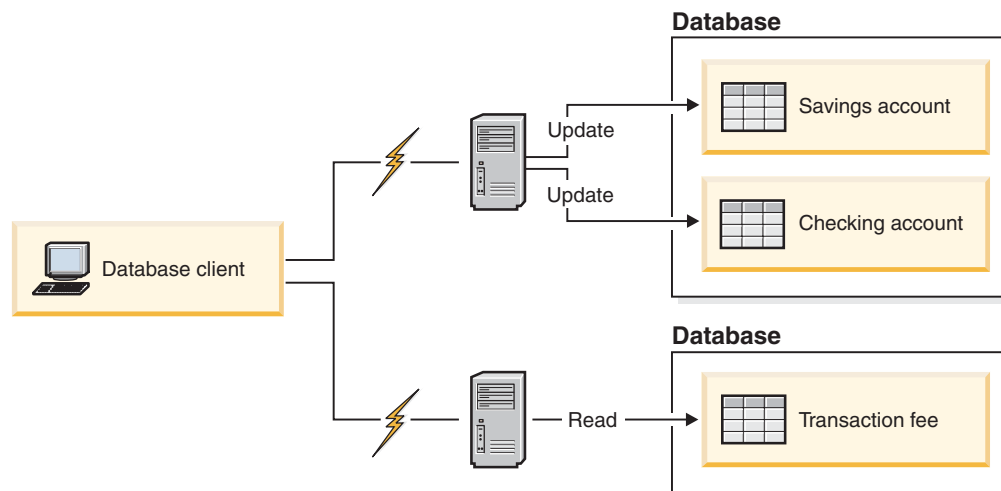


Figure 61. Using multiple databases in a single transaction

Figure 61 shows a database client running a funds transfer application that accesses two database servers: one containing the checking and savings accounts, and another containing the banking or transaction fee payment table.

Procedure:

To set up a funds transfer application for this environment, you must:

1. Create the necessary tables in the appropriate databases
2. If physically remote, set up the database servers to use the appropriate communications protocols
3. If physically remote, catalog the nodes and the databases to identify the databases on the database servers
4. Precompile your application program to specify a type 2 connection (that is, specify CONNECT 2 on the PRECOMPILE PROGRAM command), and one-phase commit (that is, specify SYNCPOINT ONEPHASE on the PRECOMPILE PROGRAM command).

If databases are located on a host or iSeries database server, you require DB2 Connect™ for connectivity to these servers.

Related concepts:

- “Units of work” on page 22

Related tasks:

- “Updating a single database in a transaction” on page 203
- “Updating multiple databases in a transaction” on page 205

Related reference:

- “PRECOMPILE command” in *Command Reference*

Updating multiple databases in a transaction

If your data is distributed across multiple databases, you may want to read and update several databases in a single transaction. This type of database access is called a *multisite update*.




Figure 62. Updating multiple databases in a single transaction

Figure 62 on page 205 shows a database client running a funds transfer application that accesses three database servers: one containing the checking account, another containing the savings account, and the third containing the banking fee schedule.

Procedure:

To set up a funds transfer application for this environment, you have two options:

1. With the DB2 transaction manager (TM):
 - a. Create the necessary tables in the appropriate databases
 - b. If physically remote, set up the database servers to use the appropriate communications protocols
 - c. If physically remote, catalog the nodes and the databases to identify the databases on the database servers
 - d. Precompile your application program to specify a type 2 connection (that is, specify `CONNECT 2` on the `PRECOMPILE PROGRAM` command), and two-phase commit (that is, specify `SYNCPOINT TWOPHASE` on the `PRECOMPILE PROGRAM` command)
 - e. Configure the DB2 transaction manager (TM).
2. Using an XA-compliant transaction manager:
 - a. Create the necessary tables in the appropriate databases
 - b. If physically remote, set up the database servers to use the appropriate communications protocols
 - c. If physically remote, catalog the nodes and the databases to identify the databases on the database servers
 - d. Precompile your application program to specify a type 2 connection (that is, specify `CONNECT 2` on the `PRECOMPILE PROGRAM` command), and one-phase commit (that is, specify `SYNCPOINT ONEPHASE` on the `PRECOMPILE PROGRAM` command)
 - e. Configure the XA-compliant transaction manager to use the DB2 databases.

Related concepts:

- “DB2 transaction manager” on page 206
- “Units of work” on page 22

Related tasks:

- “Updating a single database in a multi-database transaction” on page 204
- “Updating a single database in a transaction” on page 203

Related reference:

- “PRECOMPILE command” in *Command Reference*

DB2 transaction manager

The DB2 Database for Linux, UNIX, and Windows transaction manager (TM) assigns identifiers to transactions, monitors their progress, and takes responsibility for transaction completion and failure. The DB2 database system and DB2 Connect provide a transaction manager. The DB2 TM stores transaction information in the designated TM database.

The database manager provides transaction manager functions that can be used to coordinate the updating of several databases within a single unit of work. The

database client automatically coordinates the unit of work, and uses a *transaction manager database* to register each transaction and track its completion status.

You can use the DB2 transaction manager with DB2 databases. If you have resources other than DB2 databases that you want to participate in a two-phase commit transaction, you can use an XA-compliant transaction manager.

Related concepts:

- “DB2 Database transaction manager configuration” on page 207
- “Two-phase commit” on page 210
- “Units of work” on page 22

DB2 Database transaction manager configuration

If you are using an XA-compliant transaction manager, such as IBM® WebSphere®, BEA Tuxedo, or Microsoft Transaction Server, you should follow the configuration instructions for that product.

When using DB2 Database for Linux, UNIX, and Windows to coordinate your transactions, you must fulfill certain configuration requirements. Configuration is straightforward if you use TCP/IP exclusively for communications and DB2 Database for Linux, UNIX, and Windows or DB2 Universal Database for iSeries V5, z/OS or OS/390 are the only database servers involved in your transactions.

DB2 Connect no longer supports SNA two phase commit access to host or iSeries servers.

DB2 Database for Linux, UNIX, and Windows and DB2 Universal Database for z/OS, OS/390, and iSeries V5 using TCP/IP Connectivity

If each of the following statements is true for your environment, the configuration steps for multisite update are straightforward.

- All communications with remote database servers (including DB2 UDB for z/OS, OS/390, and iSeries V5) use TCP/IP exclusively.
- DB2 Database for Linux, UNIX, and Windows or DB2 Universal Database for z/OS, OS/390 or iSeries V5 are the only database servers involved in the transaction.

The database that will be used as the transaction manager database is determined at the database client by the database manager configuration parameter *tm_database*. Consider the following factors when setting this configuration parameter:

- The transaction manager database can be:
 - A DB2 Universal Database for UNIX or Windows Version 8 database
 - A DB2 for z/OS and OS/390 Version 7 database or a DB2 for OS/390 Version 5 or 6 database
 - A DB2 for iSeries V5 database

DB2 for z/OS, OS/390, and iSeries V5 are the recommended database servers to use as the transaction manager database. z/OS, OS/390, and iSeries V5 systems are, generally, more secure than workstation servers, reducing the possibility of accidental power downs, reboots, and so on. Therefore the recovery logs, used in the event of resynchronization, are more secure.

- If a value of `1ST_CONN` is specified for the `tm_database` configuration parameter, the first database to which an application connects is used as the transaction manager database.
Care must be taken when using `1ST_CONN`. You should only use this configuration if it is easy to ensure that all involved databases are cataloged correctly; that is, if the database client initiating the transaction is in the same instance that contains the participating databases, including the transaction manager database.
- If using TCP/IP version 6. The IP address is created depending on the operating system configuration mode chosen.
- If using Auto Configuration mode. The MAC address is extracted from the IPv6 address and is used within the internal DB2 Coordinator's Unit of Work Identifier. No configuration changes are required.
- If using Manual Configuration mode. The internal DB2 Coordinator's Unit of Work Identifier is created using the last 6 bytes of the IPv6 address. To prevent collision, the user must ensure that the last 6 bytes of the IPv6 addresses within the network are unique.

Notes:

1. DB2 Coordinator is the DB2 client and configuration changes must be performed on the system where the DB2 client exists.
2. If your application attempts to disconnect from the database being used as the transaction manager database, you will receive a warning message, and the connection will be held until the unit of work is committed.

Configuration parameters for transaction managers

You should consider the following configuration parameters when you are setting up your environment to support transaction managers.

Database Manager Configuration Parameters

- `tm_database`
This parameter identifies the name of the Transaction Manager (TM) database for each DB2 instance.
- `spm_name`
This parameter identifies the name of the DB2 Connect sync point manager instance to the database manager. For resynchronization to be successful, the name must be unique across your network.
- `resync_interval`
This parameter identifies the time interval (in seconds) after which the DB2 Transaction Manager, the DB2 server database manager, and the DB2 Connect sync point manager or the DB2 sync point manager should retry the recovery of any outstanding indoubt transactions.
- `spm_log_file_sz`
This parameter specifies the size (in 4 KB pages) of the SPM log file.
- `spm_max_resync`
This parameter identifies the number of agents that can simultaneously perform resynchronization operations.
- `spm_log_path`
This parameter identifies the log path for the SPM log files.

Database Configuration Parameters

- `maxappls`

This parameter specifies the maximum permitted number of active applications. Its value must be equal to or greater than the sum of the connected applications, plus the number of these applications that may be concurrently in the process of completing a two-phase commit or rollback, plus the anticipated number of indoubt transactions that might exist at any one time.

- *autorestart*

This database configuration parameter specifies whether the RESTART DATABASE routine will be invoked automatically when needed. The default value is YES (that is, enabled).

A database containing indoubt transactions requires a restart database operation to start up. If *autorestart* is not enabled when the last connection to the database is dropped, the next connection will fail and require an explicit RESTART DATABASE invocation. This condition will exist until the indoubt transactions have been removed, either by the transaction manager's resynchronization operation, or through a heuristic operation initiated by the administrator. When the RESTART DATABASE command is issued, a message is returned if there are any indoubt transactions in the database. The administrator can then use the LIST INDOUBT TRANSACTIONS command and other Command Line Processor (CLP) commands to find get information about those indoubt transactions.

Related concepts:

- "DB2 transaction manager" on page 206

Related tasks:

- "Configuring BEA Tuxedo" on page 238
- "Configuring IBM TXSeries CICS" on page 236
- "Configuring IBM TXSeries Encina" on page 236
- "Configuring IBM WebSphere Application Server" on page 236

Related reference:

- "autorestart - Auto restart enable configuration parameter" in *Performance Guide*
- "maxappls - Maximum number of active applications configuration parameter" in *Performance Guide*
- "spm_log_path - Sync point manager log file path configuration parameter" in *Performance Guide*
- "spm_log_file_sz - Sync point manager log file size configuration parameter" in *Performance Guide*
- "spm_name - Sync point manager name configuration parameter" in *Performance Guide*
- "spm_max_resync - Sync point manager resync agent limit configuration parameter" in *Performance Guide*
- "tm_database - Transaction manager database name configuration parameter" in *Performance Guide*
- "resync_interval - Transaction resync interval configuration parameter" in *Performance Guide*

Updating a database from a host or iSeries client

Applications executing on host or iSeries clients can access data residing on DB2 Database for Linux, UNIX, and Windows database servers. TCP/IP is the only protocol used for this access. DB2 Database for Linux, UNIX, and Windows servers on all platforms no longer support SNA access from remote clients.

Previous to version 8, TCP/IP access from host or iSeries clients only supported one-phase commit access. DB2 Database for Linux, UNIX, and Windows now allows TCP/IP two-phase commit access from host or iSeries clients. There is no need to use the Syncpoint Manager (SPM) when using TCP/IP two-phase commit access.

The DB2 TCP/IP listener must be active on the server to be accessed by the host or iSeries client. You can check that the TCP/IP listener is active by using the `db2set` command to validate that the registry variable `DB2COMM` has a value of `"tcpip"`; and that the database manager configuration parameter `svcename` is set to the service name by using the `GET DBM CFG` command. If the listener is not active, it can be made active by using the `db2set` command and the `UPDATE DBM CFG` command.

Related reference:

- "spm_name - Sync point manager name configuration parameter" in *Performance Guide*
- "Communications variables" in *Performance Guide*

Two-phase commit

Figure 63 on page 211 illustrates the steps involved in a multisite update. Understanding how a transaction is managed will help you to resolve the problem if an error occurs during the two-phase commit process.




Figure 63. Updating multiple databases

- 0** The application is prepared for two-phase commit. This can be accomplished through precompilation options. This can also be accomplished through DB2 Database for Linux, UNIX, and Windows CLI (Call Level Interface) configuration.
- 1** When the database client wants to connect to the SAVINGS_DB database, it first internally connects to the transaction manager (TM) database. The TM database returns an acknowledgment to the database client. If the database manager configuration parameter *tm_database* is set to 1ST_CONN, SAVINGS_DB becomes the transaction manager database for the duration of this application instance.
- 2** The connection to the SAVINGS_DB database takes place and is acknowledged.

- 3** The database client begins the update to the SAVINGS_ACCOUNT table. This begins the unit of work. The TM database responds to the database client, providing a transaction ID for the unit of work. Note that the registration of a unit of work occurs when the first SQL statement in the unit of work is run, not during the establishment of a connection.
- 4** After receiving the transaction ID, the database client registers the unit of work with the database containing the SAVINGS_ACCOUNT table. A response is sent back to the client to indicate that the unit of work has been registered successfully.
- 5** SQL statements issued against the SAVINGS_DB database are handled in the normal manner. The response to each statement is returned in the SQLCA when working with SQL statements embedded in a program.
- 6** The transaction ID is registered at the FEE_DB database containing the TRANSACTION_FEE table, during the first access to that database within the unit of work.
- 7** Any SQL statements against the FEE_DB database are handled in the normal way.
- 8** Additional SQL statements can be run against the SAVINGS_DB database by setting the connection, as appropriate. Since the unit of work has already been registered with the SAVINGS_DB database **4**, the database client does not need to perform the registration step again.
- 9** Connecting to, and using the CHECKING_DB database follows the same rules described in **6** and **7**.
- 10** When the database client requests that the unit of work be committed, a *prepare* message is sent to all databases participating in the unit of work. Each database writes a "PREPARED" record to its log files, and replies to the database client.
- 11** After the database client receives a positive response from all of the databases, it sends a message to the transaction manager database, informing it that the unit of work is now ready to be committed (PREPARED). The transaction manager database writes a "PREPARED" record to its log file, and sends a reply to inform the client that the second phase of the commit process can be started.
- 12** During the second phase of the commit process, the database client sends a message to all participating databases to tell them to commit. Each database writes a "COMMITTED" record to its log file, and releases the locks that were held for this unit of work. When the database has completed committing the changes, it sends a reply to the client.
- 13** After the database client receives a positive response from all participating databases, it sends a message to the transaction manager database, informing it that the unit of work has been completed. The transaction manager database then writes a "COMMITTED" record to its log file, indicating that the unit of work is complete, and replies to the client, indicating that it has finished.

Related concepts:

- "DB2 transaction manager" on page 206
- "Units of work" on page 22

Error recovery during two-phase commit

Recovering from error conditions is a normal task associated with application programming, system administration, database administration and system operation. Distributing databases over several remote servers increases the potential for error resulting from network or communications failures. To ensure data integrity, the database manager provides the two-phase commit process. The following explains how the database manager handles errors during the two-phase commit process:

- **First Phase Error**

If a database communicates that it has failed to prepare to commit the unit of work, the database client will roll back the unit of work during the second phase of the commit process. A prepare message will *not* be sent to the transaction manager database in this case.

During the second phase, the client sends a rollback message to all participating databases that successfully prepared to commit during the first phase. Each database then writes an "ABORT" record to its log file, and releases the locks that were held for this unit of work.

- **Second Phase Error**

Error handling at this stage is dependent upon whether the second phase will commit or roll back the transaction. The second phase will only roll back the transaction if the first phase encountered an error.

If one of the participating databases fails to commit the unit of work (possibly due to a communications failure), the transaction manager database will retry the commit on the failed database. The application, however, will be informed that the commit was successful through the SQLCA. DB2 Database for Linux, UNIX, and Windows will ensure that the uncommitted transaction in the database server is committed. The database manager configuration parameter *resync_interval* is used to specify how long the transaction manager database should wait between attempts to commit the unit of work. All locks are held at the database server until the unit of work is committed.

If the transaction manager database fails, it will resynchronize the unit of work when it is restarted. The resynchronization process will attempt to complete all *indoubt transactions*; that is, those transactions that have finished the first phase, but have not completed the second phase of the commit process. The database manager associated with the transaction manager database performs the resynchronization by:

1. Connecting to the databases that indicated they were "PREPARED" to commit during the first phase of the commit process.
2. Attempting to commit the indoubt transactions at those databases. (If the indoubt transactions cannot be found, the database manager assumes that the database successfully committed the transactions during the second phase of the commit process.)
3. Committing the indoubt transactions in the transaction manager database, after all indoubt transactions have been committed in the participating databases.

If one of the participating databases fails and is restarted, the database manager for this database will query the transaction manager database for the status of this transaction, to determine whether the transaction should be rolled back. If the transaction is not found in the log, the database manager assumes that the transaction was rolled back, and will roll back the indoubt transaction in this database. Otherwise, the database waits for a commit request from the transaction manager database.

If the transaction was coordinated by a transaction processing monitor (XA-compliant transaction manager), the database will always depend on the TP monitor to initiate the resynchronization.

If, for some reason, you cannot wait for the transaction manager to automatically resolve indoubt transactions, there are actions you can take to manually resolve them. This manual process is sometimes referred to as "making a heuristic decision".

Error recovery if autorestart=off

If the *autorestart* database configuration parameter is set to OFF, and there are indoubt transactions in either the TM or RM databases, the RESTART DATABASE command is required to start the resynchronization process. When issuing the RESTART DATABASE command from the command line processor, use different sessions. If you restart a different database from the same session, the connection established by the previous invocation will be dropped, and must be restarted once again. Issue the TERMINATE command to drop the connection after no more indoubt transactions are returned by the LIST INDOUBT TRANSACTIONS command.

Related concepts:

- "Two-phase commit" on page 210

Related tasks:

- "Resolving indoubt transactions manually" on page 227

Related reference:

- "autorestart - Auto restart enable configuration parameter" in *Performance Guide*
- "LIST INDOUBT TRANSACTIONS command" in *Command Reference*
- "RESTART DATABASE command" in *Command Reference*
- "TERMINATE command" in *Command Reference*

Chapter 7. Designing for XA-compliant transaction managers

You may want to use your databases with an XA-compliant transaction manager if you have resources other than DB2 databases that you want to participate in a two-phase commit transaction. If your transactions only access DB2 databases, you should use the DB2 transaction manager, described in "Updating multiple databases in a transaction" on page 205.

The topics in this chapter will assist you in using the database manager with an XA-compliant transaction manager, such as IBM WebSphere or BEA Tuxedo.

If you are looking for information about Microsoft Transaction Server, see the *Call Level Interface Guide and Reference, Volume 1*.

If you are using an XA-compliant transaction manager, or are implementing one, more information is available from our technical support web site:

<http://www.ibm.com/software/data/db2/udb/winoss2unix/support>

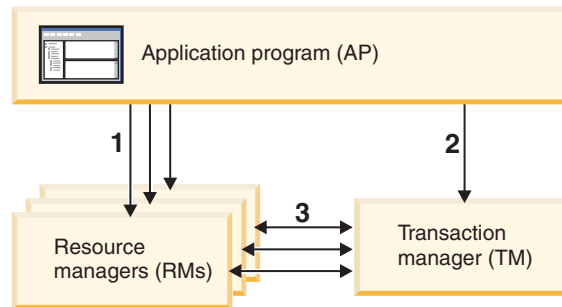
Once there, choose "DB2", then search the web site using the keyword "XA" for the latest available information on XA-compliant transaction managers.

X/Open distributed transaction processing model

The X/Open Distributed Transaction Processing (DTP) model includes three interrelated components:

- Application program (AP)
- Transaction manager (TM)
- Resources managers (RM)

Figure 64 on page 216 illustrates this model, and shows the relationship among these components.



Legend

- 1 - AP uses resources from a set of RMs
- 2 - AP defines transaction boundaries through TM interfaces
- 3 - TM and RMs exchange transaction information

Figure 64. X/Open distributed transaction processing (DTP) model

Application program (AP)

The application program (AP) defines transaction boundaries, and defines the application-specific actions that make up the transaction.

For example, a CICS® application program might want to access resource managers (RMs), such as a database and a CICS Transient Data Queue, and use programming logic to manipulate the data. Each access request is passed to the appropriate resource managers through function calls specific to that RM. In the case of DB2 products, these could be function calls generated by the DB2 database precompiler for each SQL statement, or database calls coded directly by the programmer using the APIs.

A transaction manager (TM) product usually includes a transaction processing (TP) monitor to run the user application. The TP monitor provides APIs to allow an application to start and end a transaction, and to perform application scheduling and load balancing among the many users who want to run the application. The application program in a distributed transaction processing (DTP) environment is really a combination of the user application and the TP monitor.

To facilitate an efficient online transaction processing (OLTP) environment, the TP monitor pre-allocates a number of server processes at startup, and then schedules and reuses them among the many user transactions. This conserves system resources, by allowing more concurrent users to be supported with a smaller number of server processes and their corresponding RM processes. Reusing these processes also avoids the overhead of starting up a process in the TM and RMs for each user transaction or program. (A program invokes one or more transactions.) This also means that the server processes are the actual "user processes" to the TM and the RMs. This has implications for security administration and application programming.

The following types of transactions are possible from a TP monitor:

- Non-XA transactions

These transactions involve RMs that are not defined to the TM, and are therefore not coordinated under the two-phase commit protocol of the TM. This might be necessary if the application needs to access an RM that does not support the XA interface. The TP monitor simply provides efficient scheduling of applications and load balancing. Since the TM does not explicitly "open" the RM for XA processing, the RM treats this application as any other application that runs in a non-DTP environment.

- Global transactions

These transactions involve RMs that are defined to the TM, and are under the TM's two-phase commit control. A global transaction is a unit of work that could involve one or more RMs. A *transaction branch* is the part of work between a TM and an RM that supports the global transaction. A global transaction could have multiple transaction branches when multiple RMs are accessed through one or more application processes that are coordinated by the TM.

Loosely coupled global transactions exist when each of a number of application processes accesses the RMs as if they are in a separate global transaction, but those applications are under the coordination of the TM. Each application process will have its own transaction branch within an RM. When a commit or rollback is requested by any one of the APs, TM, or RMs, the transaction branches are completed altogether. It is the application's responsibility to ensure that resource deadlock does not occur among the branches. (Note that the transaction coordination performed by the DB2 transaction manager for applications prepared with the SYNCPOINT(TWOPHASE) option is roughly equivalent to these loosely coupled global transactions.

Tightly coupled global transactions exist when multiple application processes take turns to do work under the same transaction branch in an RM. To the RM, the two application processes are a single entity. The RM must ensure that resource deadlock does not occur within the transaction branch.

Transaction manager (TM)

The transaction manager (TM) assigns identifiers to transactions, monitors their progress, and takes responsibility for transaction completion and failure. The transaction branch identifiers (known as XIDs) are assigned by the TM to identify both the global transaction, and the specific branch within an RM. This is the correlation token between the log in a TM and the log in an RM. The XID is needed for two-phase commit, or rollback, to perform the *resynchronization* operation (also known as a *resync*) on system startup, or to let the administrator perform a *heuristic* operation (also known as *manual intervention*), if necessary.

After a TP monitor is started, it asks the TM to open all the RMs that a set of application servers have defined. The TM passes **xa_open** calls to the RMs, so that they can be initialized for DTP processing. As part of this startup procedure, the TM performs a resync to recover all *indoubt transactions*. An indoubt transaction is a global transaction that was left in an uncertain state. This occurs when the TM (or at least one RM) becomes unavailable after successfully completing the first phase (that is, the prepare phase) of the two-phase commit protocol. The RM will not know whether to commit or roll back its branch of the transaction until the TM can reconcile its own log with the RM logs when they become available again. To perform the resync operation, the TM issues a **xa_recover** call one or more times to each of the RMs to identify all the indoubt transactions. The TM compares the replies with the information in its own log to determine whether it should inform the RMs to **xa_commit** or **xa_rollback** those transactions. If an RM has already

committed or rolled back its branch of an indoubt transaction through a heuristic operation by its administrator, the TM issues an **xa_forget** call to that RM to complete the resync operation.

When a user application requests a commit or a rollback, it must use the API provided by the TP monitor or TM, so that the TM can coordinate the commit and rollback among all the RMs involved. For example, when a CICS application issues the CICS SYNCPOINT request to commit a transaction, the CICS XA TM (implemented in the Encina Server) will in turn issue XA calls, such as **xa_end**, **xa_prepare**, **xa_commit**, or **xa_rollback** to request the RM to commit or roll back the transaction. The TM could choose to use one-phase instead of two-phase commit if only one RM is involved, or if an RM replies that its branch is read-only.

Resource managers (RM)

A resource manager (RM) provides access to shared resources, such as databases.

The DB2 system, as resource manager of a database, can participate in a *global transaction* that is being coordinated by an XA-compliant TM. As required by the XA interface, the database manager provides a *db2xa_switch* external C variable of type *xa_switch_t* to return the XA switch structure to the TM. This data structure contains the addresses of the various XA routines to be invoked by the TM, and the operating characteristics of the RM.

There are two methods by which the RM can register its participation in each global transaction: *static registration* and *dynamic registration*:

- Static registration requires the TM to issue (for every transaction) the **xa_start**, **xa_end**, and **xa_prepare** series of calls to all the RMs defined for the server application, regardless of whether a given RM is used by the transaction. This is inefficient if not every RM is involved in every transaction, and the degree of inefficiency is proportional to the number of defined RMs.
- Dynamic registration (used by DB2) is flexible and efficient. An RM registers with the TM using an **ax_reg** call only when the RM receives a request for its resource. Note that there is no performance disadvantage with this method, even when there is only one RM defined, or when every RM is used by every transaction, because the **ax_reg** and the **xa_start** calls have similar paths in the TM.

The XA interface provides two-way communication between a TM and an RM. It is a system-level interface between the two DTP software components, not an ordinary application program interface to which an application developer codes. However, application developers should be familiar with the programming restrictions that the DTP software components impose.

Although the XA interface is invariant, each XA-compliant TM may have product-specific ways of integrating an RM. For information about integrating your DB2 product as a resource manager with a specific transaction manager, see the appropriate TM product documentation.

Related concepts:

- “Security considerations for XA transaction managers” on page 231
- “X/Open XA Interface programming considerations” in *Developing SQL and External Routines*
- “XA function supported by DB2 Database for Linux, UNIX, and Windows” on page 233

Related tasks:

- “Updating multiple databases in a transaction” on page 205

Resource manager setup

Each database is defined as a separate resource manager (RM) to the transaction manager (TM), and the database must be identified with an `xa_open` string.

When setting up a database as a resource manager, you do not need the `xa_close` string. If provided, this string will be ignored by the database manager.

Database connection considerations

Automatic client reroute (ACR)

Whenever a server crashes, each client that is connected to that server gets a communication error which terminates the connection and concludes in an application error. In application environments where availability is important, the user will either have a redundant setup or will fail the server over to a standby node. In either case, the DB2 Database for Linux, UNIX, and Windows client code will attempt to re-establish the connection to either the original database (which may be running on a failover node where the IP address fails over as well), or to a new database on a different server. The application is then notified using an `SQLCODE` to indicate that the connection has been rerouted and that the specific transaction being run has been rolled back. At that point, the application can choose to rerun that transaction or continue on.

Data consistency between the failed primary database and the “failed to” standby database when using ACR is very dependent upon the state of the database logs in the database to which the connection has been rerouted. For the purposes of this discussion, we will call this database the “standby database” and the server on which this standby database resides the “standby server”. If the standby database is an exact copy of the failed primary database at the point in time of the failure then the data at the standby database will be consistent and there will be no data integrity issues. However, if the standby database is not an exact copy of the failed primary database then there may be data integrity issues resulting from inconsistent transaction outcomes for transactions which have been prepared by the XA Transaction Manager but yet to be committed. These are known as indoubt transactions. The Database Administrator and application developers who are using the ACR function must be aware of the risk of data integrity problems when using this capability.

The following sections describe the various DB2 Database for Linux, UNIX, and Windows environments and the risks of data integrity problems in each.

High availability disaster recovery (HADR):

DB2’s High Availability Disaster Recovery feature (HADR) can be used to control the level of log duplication between the primary and standby databases when the application regains connectivity after a primary database failure. The database configuration parameter which controls the level of log duplication is called `hadr_syncmode`. There are three possible values for this parameter:

- SYNC

This mode provides the greatest protection against transaction loss at the cost of longest transaction response time among the three modes. As the name of this

mode suggests, SYNC is used to synchronize the writing of the transaction log in the primary database and in the standby database. Synchronization is accomplished when the primary database has written its own log files and it has received acknowledgement from the standby database that the logs have also been written on the standby database.

If an XA Transaction Manager is being used to coordinate transactions involving DB2 resources, then it is strongly recommended that SYNC mode be used. SYNC mode will guarantee data integrity as well as transaction resynchronization integrity when a client is rerouted to the standby database since it is an exact replica of the primary database.

- NEARSYNC

This mode provides slightly less protection against transaction loss, in exchange for a shorter transaction response time when compared with SYNC mode. The primary database considers log write successful only when logs have been written to its own log files and it has received acknowledgement from the standby database that the logs have also been written to main memory on the standby database. If the standby database crashes before it can copy the logs from memory to disk, the logs are lost on the standby database in the short term.

Given the possibility that database logs are lost, and the situation where the standby database is not an exact replica of the primary database, it is possible that data integrity will be compromised. The compromise occurs if the given transaction was indoubt and then the primary database crashes. Assume the transaction outcome is COMMIT. When the XA TM issues the subsequent XA_COMMIT request, it will fail since the primary database has crashed. Since the XA_COMMIT request has failed, the XA TM will need to recover this transaction on this database by issuing an XA_RECOVER request. The standby database will respond by returning the list of all its transactions which are INDOUBT. If the standby database were to crash and restart before the “in memory,” database logs were written to disk, and before the XA_RECOVER request was issued by the XA TM, the standby database would have lost the log information about the transaction and could not return it in response to the XA_RECOVER request. The XA TM would then assume the database committed this transaction. But, what has really occurred is the data manipulation will have been lost and the appearance that the transaction was rolled back. This results in a data integrity issue since all other resources involved in this transaction were COMMITTED by the XA TM.

Using NEARSYNC is a good compromise between data integrity and transaction response time since the likelihood of both the primary and standby databases crashing should be low. However, a database administrator still needs to understand that there is a possibility of data integrity problems.

- ASYNC

This mode has the greatest chance of transaction loss in the event of primary failure, in exchange for the shortest transaction response time among the three modes. The primary database considers log write successful only when logs have been written to its own log files and the logs have been delivered to the TCP layer on the primary database’s host machine. The primary database does not wait for acknowledgement of any kind from the standby database. The logs may be still on their way to the standby database when the primary database considers relevant transactions committed.

If the same scenario as described in NEARSYNC occurs, the likelihood of loss of transaction information is higher than with NEARSYNC. Therefore, the likelihood of data integrity issues is higher than with NEARSYNC and, obviously, with SYNC.

DB2 ESE Partitioned Database Environments:

The use of ACR in partitioned database environments can also lead to data integrity issues. If the standby database is defined to be a different database partition of the same database, then recovery of indoubt transactions in scenarios as described in the High Availability Disaster Recovery NEARSYNC section above, may result in data integrity problems. This occurs because the database partitions do not share database transaction logs. Therefore the standby database (database partition B) will have no knowledge of indoubt transactions that exist at the primary database (database partition A).

DB2 ESE Non Partitioned Database Environments:

The use of ACR in non-partitioned database environments can also lead to data integrity issues. Assuming disk failover technology, such as IBM AIX High Availability Cluster Multiprocessor (HACMP™), Microsoft Cluster Service (MSCS), or HP's Service Guard, is not in use then the standby database will not have the database transaction logs that existed on the primary database when it failed. Therefore, the recovery of indoubt transactions in scenarios as described in the High Availability Disaster Recovery NEARSYNC section above, can result in data integrity problems.

Transactions accessing partitioned databases

In a partitioned database environment, user data may be distributed across database partitions. An application accessing the database connects and sends requests to one of the database partitions (the coordinator node). Different applications can connect to different database partitions, and the same application can choose different database partitions for different connections.

For transactions against a database in a partitioned database environment, all access must be through the *same* database partition. That is, the same database partition must be used from the start of the transaction until (and including) the time that the transaction is committed.

Any transaction against the partitioned database must be committed before disconnecting.

Related concepts:

- "X/Open distributed transaction processing model" on page 215
- "High availability disaster recovery overview" in *Data Recovery and High Availability Guide and Reference*

Related reference:

- "xa_open string formats" on page 221

xa_open string formats

xa_open string format for DB2 Database for Linux, UNIX, and Windows and DB2 Connect Version 8 FixPak 3 and later:

This is the format for the xa_open string:

```
parm_id1 = <parm value>,parm_id2 = <parm value>, ...
```

It does not matter in what order these parameters are specified. Valid values for *parm_id* are described below.

Note: Unless explicitly stated, these parameters are not case sensitive and have no default value.

AXLIB

Library that contains the TP monitor's **ax_reg** and **ax_unreg** functions. This value is used by DB2 to obtain the addresses of the required **ax_reg** and **ax_unreg** functions. It can be used to override assumed values based on the TPM parameter, or it can be used by TP monitors that do not appear on the list for TPM. On AIX, if the library is an archive library, the archive member should be specified in addition to the library name. For example:
AXLIB=/usr/mqm/lib/libmqmax_r.a(libmqmax_r.o). This parameter is optional.

CHAIN_END

xa_end chaining flag. Valid values are T, F, or no value. XA_END chaining is an optimization that can be used by DB2 to reduce network flows. If the TP monitor environment is such that it can be guaranteed that **xa_prepare** will be invoked within the same thread or process immediately following the call to **xa_end**, and if CHAIN_END is on, the **xa_end** flag will be chained with the **xa_prepare** command, thus eliminating one network flow. A value of T means that CHAIN_END is on; a value of F means that CHAIN_END is off; no specified value means that CHAIN_END is on. This parameter can be used to override the setting derived from a specified TPM value. If this parameter is not specified, the default value of F is used.

CREG

xa_start chaining flag. Valid values are T, or F, or no value. **xa_start** chaining is an optimization that is used by DB2 to reduce network flows. The parameter is only valid if the TP monitor is using static registration (see SREG). The TP monitor environment is such that it can guarantee that an SQL statement will be invoked immediately after the call to the XA API **xa_start**. If CREG is set to T, the SQL statement is chained to the **xa_start** request, thus eliminating one network flow. This parameter can be used to override the setting derived from a specified TPM value. If this parameter is not specified, the default value of F is used.

CT

Connect Timeout. Valid values are 0 - 32767. CT specifies the amount of time, in seconds, that an application will wait when attempting to establish a connection with the server. If a connection is not established in the amount of time specified, an error will be returned. Specifying a value of 0 means that the application will attempt to wait until a connection is established regardless of how long it takes. However, it is possible that the connection attempt will be terminated by the default TCP/IP timeout setting. If this parameter is not specified, the default value of 0 is used.

DB

Database alias. Database alias used by the application to access the database. This parameter must be specified.

HOLD_CURSOR

Specifies whether cursors are held across transaction commits. Valid values are T, F, or no value. TP monitors typically reuse threads or processes for multiple applications. To ensure that a newly loaded application does not inherit cursors opened by a previous application, cursors are closed after a commit. If HOLD_CURSORS is on, cursors with hold attributes are not closed, and will persist across transaction commit boundaries. When using this option, the

global transaction must be committed or rolled back from the same thread of control. If `HOLD_CURSOR` is off, the opening of any cursors with hold attributes will be rejected. A value of T means that `HOLD_CURSOR` is on; a value of F means that `HOLD_CURSOR` is off; no specified value means that `HOLD_CURSOR` is on. This parameter can be used to override the setting derived from a specified TPM value. If this parameter is not specified, the default value of F is used.

PWD

Password. A password that is associated with the user ID. Required if a user ID is specified. This parameter is case sensitive.

SREG

Static Registration. Valid values are T, or F, or no value. DB2 supports two methods of registering a global transaction. The first is Dynamic Registration, where DB2 calls the TP's `ax_reg` function to register the transaction (see AXLIB). The second method is Static Registration, where the TP calls the XA API `xa_start` to initiate a global transaction. Please note both dynamic and static registration are mutually exclusive. If this parameter is not specified, the default value of F is used.

SUSPEND_CURSOR

Specifies whether cursors are to be kept when a transaction thread of control is suspended. Valid values are T, F, or no value. TP monitors that suspend a transaction branch can reuse the suspended thread or process for other transactions. If `SUSPEND_CURSOR` is off, all cursors except cursors with hold attributes are closed. On resumption of the suspended transaction, the application must obtain the cursors again. If `SUSPEND_CURSOR` is on, any open cursors are not closed, and are available to the suspended transaction on resumption. A value of T means that `SUSPEND_CURSOR` is on; a value of F means that `SUSPEND_CURSOR` is off; no specified value means that `SUSPEND_CURSOR` is on. This parameter can be used to override the setting derived from a specified TPM value. If this parameter is not specified, the default value of F is used.

TOC

The entity ("Thread of Control") to which all DB2 XA Connections are bound. Valid values are T, or P, or not set. TOC is the entity where all DB2 XA Connections are bound. All DB2 XA Connections formed within an entity must be unique. That is, they cannot have two connections to the same database within the entity. The TOC has two parameters: T (OS Thread) and P (OS Process). When set to a value of T, all DB2 XA Connections formed under a particular OS Thread are unique to that thread only. Multiple threads cannot share DB2 XA Connections. Each OS thread has to form its own set of DB2 XA Connections. When set to a value of P, all DB2 XA Connections are unique to the OS Process and all XA Connections can be shared between OS threads. If this parameter is not specified, the default value of T is used.

TPM

Transaction processing monitor name. Name of the TP monitor being used. For supported values, see the next table. This parameter can be specified to allow multiple TP monitors to use a single DB2 instance. The specified value will override the value specified in the `tp_mon_name` database manager configuration parameter. This parameter is optional.

UID

User ID. Specifies the user ID that has authority to connect to the database. Required if a password is specified. This parameter is case sensitive.

UREGNM

User Registry Name. When an identity mapping service is being used, this parameter gives the name of the registry to which the user name given in the UID parameter belongs.

TCTX

Specifies whether or not the transaction should use a trusted connection. Valid values are TRUE or FALSE. If this parameter is set to TRUE it tells the transaction manager to try to open a trusted connection.

TPM and tp_mon_name values:

The xa_open string TPM parameter and the *tp_mon_name* database manager configuration parameter are used to indicate to DB2 which TP monitor is being used. The *tp_mon_name* value applies to the entire DB2 instance. The TPM parameter applies only to the specific XA resource manager. The TPM value overrides the *tp_mon_name* parameter. Valid values for the TPM and *tp_mon_name* parameters are as follows:

Table 40. Valid Values for TPM and *tp_mon_name*

TPM Value	TP Monitor Product	Internal Settings
CICS	IBM TxSeries CICS	AXLIB=libEncServer (for Windows) =/usr/lpp/encina/lib/libEncServer (for UNIX based systems) HOLD_CURSOR=T CHAIN_END=T SUSPEND_CURSOR=F TOC=T
ENCINA	IBM TxSeries Encina® Monitor	AXLIB=libEncServer (for Windows) =/usr/lpp/encina/lib/libEncServer (for UNIX based systems) HOLD_CURSOR=F CHAIN_END=T SUSPEND_CURSOR=F TOC=T
MQ	IBM MQSeries®	AXLIB=mqmax (for Windows) =/usr/mqm/lib/libmqmax_r.a (for AIX threaded applications) =/usr/mqm/lib/libmqmax.a (for AIX non-threaded applications) =/opt/mqm/lib/libmqmax.so (for Solaris) =/opt/mqm/lib/libmqmax_r.sl (for HP threaded applications) =/opt/mqm/lib/libmqmax.sl (for HP non-threaded applications) =/opt/mqm/lib/libmqmax_r.so (for Linux threaded applications) =/opt/mqm/lib/libmqmax.so (for Linux non-threaded applications) HOLD_CURSOR=F CHAIN_END=F SUSPEND_CURSOR=F TOC=P

Table 40. Valid Values for TPM and tp_mon_name (continued)

TPM Value	TP Monitor Product	Internal Settings
CB	IBM Component Broker	AXLIB=somtrx1i (for Windows) =libsomtrx1 (for UNIX based systems) HOLD_CURSOR=F CHAIN_END=T SUSPEND_CURSOR=F TOC=T
SF	IBM San Francisco	AXLIB=ibmsfDB2 HOLD_CURSOR=F CHAIN_END=T SUSPEND_CURSOR=F TOC=T
TUXEDO	BEA Tuxedo	AXLIB=libtux HOLD_CURSOR=F CHAIN_END=F SUSPEND_CURSOR=F TOC=T
MTS	Microsoft Transaction Server	It is not necessary to configure DB2 for MTS. MTS is automatically detected by DB2's ODBC driver.
JTA	Java Transaction API	It is not necessary to configure DB2 for Enterprise Java Servers (EJS) such as IBM WebSphere. DB2's JDBC driver automatically detects this environment. Therefore this TPM value is ignored.

xa_open string format for earlier versions:

Earlier versions of DB2 used the xa_open string format described here. This format is still supported for compatibility reasons. Applications should be migrated to the new format when possible.

Each database is defined as a separate resource manager (RM) to the transaction manager (TM), and the database must be identified with an xa_open string that has the following syntax:

```
"database_alias<,userid,password>"
```

The *database_alias* is required to specify the alias name of the database. The alias name is the same as the database name unless you have explicitly cataloged an alias name after you created the database. The user name and password are optional and, depending on the authentication method, are used to provide authentication information to the database.

Examples:

1. You are using IBM TxSeries CICS on Windows. The TxSeries documentation indicates that you need to configure *tp_mon_name* with a value of libEncServer:C. This is still an acceptable format; however, with DB2 Database for Linux, UNIX, and Windows or DB2 Connect Version 8 FixPak 3 and later, you have the option of:

- Specifying a *tp_mon_name* of CICS (recommended for this scenario):

```
db2 update dbm cfg using tp_mon_name CICS
```

For each database defined to CICS in the Region—> Resources—> Product—> XAD—> Resource manager initialization string, specify:

```
db=dbalias,uid=userid,pwd=password
```

- For each database defined to CICS in the Region—> Resources—> Product—> XAD—> Resource manager initialization string, specify:
`db=dbalias,uid=userid,pwd=password,tpm=cics`
2. You are using IBM MQSeries on Windows. The MQSeries documentation indicates that you need to configure *tp_mon_name* with a value of mqmax. This is still an acceptable format; however, with DB2 Database for Linux, UNIX, and Windows or DB2 Connect Version 8 FixPak 3 and later, you have the option of:
 - Specifying a *tp_mon_name* of MQ (recommended for this scenario):
`db2 update dbm cfg using tp_mon_name MQ`

For each database defined to CICS in the Region—> Resources—> Product—> XAD—> Resource manager initialization string, specify:
`uid=userid,db=dbalias,pwd=password`

 - For each database defined to CICS in the Region—> Resources—> Product—> XAD—> Resource manager initialization string, specify:
`uid=userid,db=dbalias,pwd=password,tpm=mq`
 3. You are using both IBM TxSeries CICS and IBM MQSeries on Windows. A single DB2 instance is being used. In this scenario, you would configure as follows:
 - a. For each database defined to CICS in the Region—> Resources—> Product—> XAD—> Resource manager initialization string, specify:
`pwd=password,uid=userid,tpm=cics,db=dbalias`
 - b. For each database defined as a resource in the queue manager properties, specify an XaOpenString as:
`db=dbalias,uid=userid,pwd=password,tpm=mq`
 4. You are developing your own XA-compliant transaction manager (XA TM) on Windows, and you want to tell DB2 that library "myaxlib" has the required functions **ax_reg** and **ax_unreg**. Library "myaxlib" is in a directory specified in the PATH statement. You have the option of:
 - Specifying a *tp_mon_name* of myaxlib:
`db2 update dbm cfg using tp_mon_name myaxlib`

and, for each database defined to the XA TM, specifying an xa_open string:
`db=dbalias,uid=userid,pwd=password`

 - For each database defined to the XA TM, specifying an xa_open string:
`db=dbalias,uid=userid,pwd=password,axlib=myaxlib`
 5. You are developing your own XA-compliant transaction manager (XA TM) on Windows, and you want to tell DB2 that library "myaxlib" has the required functions **ax_reg** and **ax_unreg**. Library "myaxlib" is in a directory specified in the PATH statement. You also want to enable XA END chaining. You have the option of:
 - For each database defined to the XA TM, specifying an xa_open string:
`db=dbalias,uid=userid,pwd=password,axlib=myaxlib,chain_end=T`
 - For each database defined to the XA TM, specifying an xa_open string:
`db=dbalias,uid=userid,pwd=password,axlib=myaxlib,chain_end`

Related concepts:

- "X/Open distributed transaction processing model" on page 215

Related reference:

- “tp_mon_name - Transaction processor monitor name configuration parameter” in *Performance Guide*

Updating host or iSeries database servers with an XA-compliant transaction manager

Host and iSeries database servers may be updatable depending upon the architecture of the XA Transaction Manager.

Procedure:

To support commit sequences from different processes, the DB2 Connect connection concentrator must be enabled. To enable the DB2 Connect connection concentrator, set the database manager configuration parameter *max_connections* to a value greater than *max_coordagents*. Note that the DB2 Connect connection concentrator requires a DB2 Universal Database (DB2 UDB) Version 7.1 client or later to support XA commit sequences from different processes.

You will also require DB2 Connect with the DB2 sync point manager (SPM) configured.

Related reference:

- “maxagents - Maximum number of agents configuration parameter” in *Performance Guide*
- “max_connections - Maximum number of client connections configuration parameter” in *Performance Guide*

Resolving indoubt transactions manually

An XA-compliant transaction manager (Transaction Processing Monitor) uses a two-phase commit process similar to that used by the DB2 transaction manager. The principal difference between the two environments is that the TP monitor provides the function of logging and controlling the transaction, instead of the DB2 transaction manager and the transaction manager database.

Errors similar to those that occur for the DB2 transaction manager can occur when using an XA-compliant transaction manager. Similar to the DB2 transaction manager, an XA-compliant transaction manager will attempt to resynchronize indoubt transactions.

If you cannot wait for the transaction manager to automatically resolve indoubt transactions, you can manually resolve them. This manual process is sometimes referred to as “making a heuristic decision”.

The LIST INDOUBT TRANSACTIONS command (using the WITH PROMPTING option), or the related set of APIs (*db2XaListIndTrans*, *sqlxphcm*, *sqlxhfrg*, *sqlxphrl*), allows you to query, commit, and roll back indoubt transactions. In addition, it also allows you to “forget” transactions that have been heuristically committed or rolled back, by removing the log records and releasing the log space.

Restrictions:

Manually resolve indoubt transactions by using these commands (or related APIs) with *extreme caution*, and only as a last resort. The best strategy is to wait for the

transaction manager to drive the resynchronization process. You could experience data integrity problems if you manually commit or roll back a transaction in one of the participating databases, and the opposite action is taken against another participating database. Recovering from data integrity problems requires you to understand the application logic, to identify the data that was changed or rolled back, and then to perform a point-in-time recovery of the database, or manually undo or reapply the changes.

If you cannot wait for the transaction manager to initiate the resynchronization process, and you must release the resources tied up by an indoubt transaction, heuristic operations are necessary. This situation could occur if the transaction manager will not be available for an extended period of time to perform the resynchronization, and the indoubt transaction is tying up resources that are urgently needed. An indoubt transaction ties up the resources that were associated with this transaction before the transaction manager or resource managers became unavailable. For the database manager, these resources include locks on tables and indexes, log space, and storage taken up by the transaction. Each indoubt transaction also decreases (by one) the maximum number of concurrent transactions that can be handled by the database. Moreover, an offline backup cannot be taken unless all indoubt transactions have been resolved.

The heuristic forget function is required in the following situations:

- When a heuristically committed or rolled back transaction causes a log full condition, indicated in output from the LIST INDOUBT TRANSACTIONS command
- When an offline backup is to be taken

The heuristic forget function releases the log space occupied by an indoubt transaction. The implication is that if a transaction manager eventually performs a resynchronization operation for this indoubt transaction, it could potentially make the wrong decision to commit or roll back other resource managers, because there is no log record for the transaction in this resource manager. In general a “missing” log record implies that the resource manager has rolled back the transaction.

Procedure:

To manually resolve indoubt transactions:

1. Connect to the database for which you require all transactions to be complete.
2. Display the indoubt transactions:
 - For DB2 database servers, use the LIST INDOUBT TRANSACTIONS WITH PROMPTING command. The *xid* represents the global transaction ID, and is identical to the *xid* used by the transaction manager and by other resource managers participating in the transaction.
 - For host or iSeries database servers, you may use one of the following:
 - You can obtain indoubt information directly from the host or iSeries server.
To obtain indoubt information directly from DB2 for z/OS and OS/390, invoke the DISPLAY THREAD TYPE(INDOUBT) command. Use the RECOVER command to make a heuristic decision. To obtain indoubt information directly from DB2 for iSeries, invoke the **wrkcmdfn** command.
 - You can obtain indoubt information from the DB2 Connect server used to access the host or iSeries database server.

To obtain indoubt information from the DB2 Connect server, first connect to the DB2 sync point manager by connecting to the DB2 instance represented by the value of the *spm_name* database manager configuration parameter. Then issue the LIST DRDA INDOUBT TRANSACTIONS WITH PROMPTING command to display indoubt transactions and to make heuristic decisions. Alternatively, you can call the *sqlcspqy* API from a client application to list DRDA® indoubt transactions.

3. For each indoubt transaction that has been listed or displayed, use the information shown about the application and the operating environment to determine the other participating resource managers.
4. Determine the actions to take with each indoubt transaction:
 - If the transaction manager is available, and the indoubt transaction in a resource manager was caused by the resource manager not being available in the second commit phase, or for an earlier resynchronization process, you should do the following:
 - a. Check the transaction manager's log to determine what action has been taken against the other resource managers.
 - b. Take the same action against the database; that is, use the LIST INDOUBT TRANSACTIONS WITH PROMPTING command, to either heuristically commit or heuristically roll back the transaction.
 - If the transaction manager is *not* available, use the status of the transaction in the other participating resource managers to determine what action you should take:
 - If at least one of the other resource managers has committed the transaction, heuristically commit the transaction in all the resource managers.
 - If at least one of the other resource managers has rolled back the transaction, heuristically roll back the transaction.
 - If the transaction is in the “prepared” (indoubt) state in all of the participating resource managers, heuristically roll back the transaction.
 - If one or more of the other resource managers is not available, heuristically roll back the transaction.

To obtain indoubt transaction information from DB2 on UNIX or Windows, connect to the database and issue the LIST INDOUBT TRANSACTIONS WITH PROMPTING command, or call the *db2XaListIndTrans* API from a client application.

Related concepts:

- “Indoubt transaction management APIs” on page 230
- “Two-phase commit” on page 210

Related reference:

- “LIST DRDA INDOUBT TRANSACTIONS command” in *Command Reference*
- “LIST INDOUBT TRANSACTIONS command” in *Command Reference*
- “db2XaListIndTrans API - List indoubt transactions” in *Administrative API Reference*
- “sqlcspqy API - List DRDA indoubt transactions” in *Administrative API Reference*
- “sqlxhfrg API - Forget transaction status” in *Administrative API Reference*
- “sqlxphcm API - Commit an indoubt transaction” in *Administrative API Reference*
- “sqlxphrl API - Roll back an indoubt transaction” in *Administrative API Reference*

Indoubt transaction management APIs

Related samples:

- “dbxamon.c -- Show and roll back indoubt transactions.”

Indoubt transaction management APIs

Databases can be used in a distributed transaction processing (DTP) environment.

A set of APIs is provided for tool writers to perform heuristic functions on indoubt transactions when the resource owner (such as the database administrator) cannot wait for the Transaction Manager (TM) to perform the *re-sync* action. This condition may occur if, for example, the communication line is broken, and an indoubt transaction is tying up needed resources. For the database manager, these resources include locks on tables and indexes, log space, and storage used by the transaction. Each indoubt transaction also decreases, by one, the maximum number of concurrent transactions that could be processed by the database manager.

The heuristic APIs have the capability to query, commit, and roll back indoubt transactions, and to cancel transactions that have been heuristically committed or rolled back, by removing the log records and releasing log pages.

Attention: The heuristic APIs should be used with caution and only as a last resort. The TM should drive the re-sync events. If the TM has an operator command to start the re-sync action, it should be used. If the user cannot wait for a TM-initiated re-sync, heuristic actions are necessary.

Although there is no set way to perform these actions, the following guidelines may be helpful:

- Use the `db2XaListIndTrans` function to display the indoubt transactions. They have a status = 'P' (prepared), and are not connected. The *gtrid* portion of an *xid* is the global transaction ID that is identical to that in other resource managers (RM) that participate in the global transaction.
- Use knowledge of the application and the operating environment to identify the other participating RMs.
- If the transaction manager is CICS, and the only RM is a CICS resource, perform a heuristic rollback.
- If the transaction manager is not CICS, use it to determine the status of the transaction that has the same *gtrid* as does the indoubt transaction.
- If at least one RM has committed or rolled back, perform a heuristic commit or a rollback.
- If they are all in the prepared state, perform a heuristic rollback.
- If at least one RM is not available, perform a heuristic rollback.

If the transaction manager is available, and the indoubt transaction is due to the RM not being available in the second phase, or in an earlier re-sync, the DBA should determine from the TM's log what action has been taken against the other RMs, and then do the same. The *gtrid* is the matching key between the TM and the RMs.

Do not execute `sqlxhfrg` unless a heuristically committed or rolled back transaction happens to cause a log full condition. The `forget` function releases the log space occupied by this indoubt transaction. If a transaction manager eventually performs a re-sync action for this indoubt transaction, the TM could make the wrong

decision to commit or to roll back other RMs, because no record was found in this RM. In general, a missing record implies that the RM has rolled back.

Related reference:

- “db2XaListIndTrans API - List indoubt transactions” in *Administrative API Reference*
- “sqlcspqy API - List DRDA indoubt transactions” in *Administrative API Reference*
- “sqlxhfrg API - Forget transaction status” in *Administrative API Reference*
- “sqlxphcm API - Commit an indoubt transaction” in *Administrative API Reference*
- “sqlxphrl API - Roll back an indoubt transaction” in *Administrative API Reference*

Security considerations for XA transaction managers

The TP monitor pre-allocates a set of server processes and runs the transactions from different users under the IDs of the server processes. To the database, each server process appears as a big application that has many units of work, all being run under the same ID associated with the server process.

For example, in an AIX environment using CICS, when a TXSeries[®] CICS region is started, it is associated with the AIX user name under which it is defined. All the CICS Application Server processes are also being run under this TXSeries CICS “master” ID, which is usually defined as “cics”. CICS users can invoke CICS transactions under their DCE login ID, and while in CICS, they can also change their ID using the CESN signon transaction. In either case, the end user’s ID is not available to the RM. Consequently, a CICS Application Process might be running transactions on behalf of many users, but they appear to the RM as a single program with many units of work from the same “cics” ID. Optionally, you can specify a user ID and password on the xa_open string, and that user ID will be used, instead of the “cics” ID, to connect to the database.

There is not much impact on static SQL statements, because the binder’s privileges, not the end user’s privileges, are used to access the database. This does mean, however, that the EXECUTE privilege of the database packages must be granted to the server ID, and not to the end user ID.

For dynamic statements, which have their access authentication done at run time, access privileges to the database objects must be granted to the server ID and not to the actual user of those objects. Instead of relying on the database to control the access of specific users, you must rely on the TP monitor system to determine which users can run which programs. The server ID must be granted all privileges that its SQL users require.

To determine who has accessed a database table or view, you can perform the following steps:

1. From the SYSCAT.PACKAGEDEP catalog view, obtain a list of all packages that depend on the table or view.
2. Determine the names of the server programs (for example, CICS programs) that correspond to these packages through the naming convention used in your installation.
3. Determine the client programs (for example, CICS transaction IDs) that could invoke these programs, and then use the TP monitor’s log (for example, the CICS log) to determine who has run these transactions or programs, and when.

Indoubt transaction management APIs

Related concepts:

- “X/Open distributed transaction processing model” on page 215

Configuration considerations for XA transaction managers

You should consider the following configuration parameters when you are setting up your TP monitor environment:

- *tp_mon_name*

This database manager configuration parameter identifies the name of the TP monitor product being used (“CICS”, or “ENCINA”, for example).

- *tpname*

This database manager configuration parameter identifies the name of the remote transaction program that the database client must use when issuing an allocate request to the database server, using the APPC communications protocol. The value is set in the configuration file at the server, and must be the same as the transaction processor (TP) name configured in the SNA transaction program.

- *tm_database*

Because DB2 Database for Linux, UNIX, and Windows does *not* coordinate transactions in the XA environment, this database manager configuration parameter is not used for XA-coordinated transactions.

- *maxappls*

This database configuration parameter specifies the maximum number of active applications allowed. The value of this parameter must be equal to or greater than the sum of the connected applications, plus the number of these applications that may be concurrently in the process of completing a two-phase commit or rollback. This sum should then be increased by the anticipated number of indoubt transactions that might exist at any one time.

For a TP monitor environment (for example, TXSeries CICS), you may need to increase the value of the *maxappls* parameter. This would help to ensure that all TP monitor processes can be accommodated.

- *autorestart*

This database configuration parameter specifies whether the RESTART DATABASE routine will be invoked automatically when needed. The default value is YES (that is, enabled).

A database containing indoubt transactions requires a restart database operation to start up. If *autorestart* is not enabled when the last connection to the database is dropped, the next connection will fail and require an explicit RESTART DATABASE invocation. This condition will exist until the indoubt transactions have been removed, either by the transaction manager’s resync operation, or through a heuristic operation initiated by the administrator. When the RESTART DATABASE command is issued, a message is returned if there are any indoubt transactions in the database. The administrator can then use the LIST INDOUBT TRANSACTIONS command and other command line processor commands to find get information about those indoubt transactions.

Related concepts:

- “X/Open distributed transaction processing model” on page 215

Related reference:

- “tpname - APPC transaction program name configuration parameter” in *Performance Guide*

- “autorestart - Auto restart enable configuration parameter” in *Performance Guide*
- “LIST INDOUBT TRANSACTIONS command” in *Command Reference*
- “maxappls - Maximum number of active applications configuration parameter” in *Performance Guide*
- “RESTART DATABASE command” in *Command Reference*
- “tm_database - Transaction manager database name configuration parameter” in *Performance Guide*
- “tp_mon_name - Transaction processor monitor name configuration parameter” in *Performance Guide*

XA function supported by DB2 Database for Linux, UNIX, and Windows

DB2 Database for Linux, UNIX, and Windows supports the XA91 specification defined in *X/Open CAE Specification Distributed Transaction Processing: The XA Specification*, with the following exceptions:

- Asynchronous services
The XA specification allows the interface to use asynchronous services, so that the result of a request can be checked at a later time. The database manager requires that the requests be invoked in synchronous mode.
- Registration
The XA interface allows two ways to register an RM: static registration and dynamic registration. DB2 supports both dynamic and static registration. DB2 provides two switches to control the type of registration used.
 - `db2xa_switch_std` for dynamic registration
 - `db2xa_switch_static_std` for static registration
- Association migration
DB2 V9.1 does not support transaction migration between threads of control.

XA switch usage and location

As required by the XA interface, the database manager provides a `db2xa_switch_std` and a `db2xa_switch_static_std` external C variable of type `xa_switch_t` to return the XA switch structure to the TM. Other than the addresses of various XA functions, the following fields are returned:

Field	Value
name	The product name of the database manager. For example, IBM DB2 Version 9.1 for AIX.
flags	For <code>db2xa_switch_std</code> TMREGISTER TMNOMIGRATE is set Explicitly states that DB2 V9.1 uses dynamic registration, and that the TM should not use association migration. Implicitly states that asynchronous operation is not supported. For <code>db2xa_switch_static_std</code> TMNOMIGRATE is set Explicitly states that DB2 V9.1 uses static registration, and that the TM should not use association migration. Implicitly states that asynchronous operation is not supported.
version	Must be zero.

Using the DB2 Database for Linux, UNIX, and Windows XA switch

The XA architecture requires that a Resource Manager (RM) provide a *switch* that gives the XA Transaction Manager (TM) access to the RM's *xa_* routines. An RM switch uses a structure called *xa_switch_t*. The switch contains the RM's name, non-NULL pointers to the RM's XA entry points, a flag, and a version number.

Linux and UNIX

The switch for DB2 Database for Linux, UNIX, and Windows can be obtained through either of the following two ways:

- Through one additional level of indirection. In a C program, this can be accomplished by defining the macro:

```
#define db2xa_switch_std (*db2xa_switch_std)
#define db2xa_switch_static_std (*db2xa_switch_std)
```

prior to using *db2xa_switch_std* or *db2xa_switch_static_std*.

- By calling **db2xacic_std** or **db2xacicst_std**

DB2 provides these APIs, which return the address of the *db2xa_switch_std* or *db2xa_switch_static_std* structure. This function is prototyped as:

```
struct xa_switch_t * SQL_API_FN db2xacic_std( )
struct xa_switch_t * SQL_API_FN db2xacicst_std( )
```

With either method, you must link your application with *libdb2*.

Windows

The pointer to the *xa_switch* structure, *db2xa_switch_std*, or *db2xa_switch_static_std* is exported as DLL data. This implies that a Windows application using this structure must reference it in one of three ways:

- Through one additional level of indirection. In a C program, this can be accomplished by defining the macro:

```
#define db2xa_switch_std (*db2xa_switch_std)
#define db2xa_switch_static_std (*db2xa_switch_std)
```

prior to using *db2xa_switch_std* or *db2xa_switch_static_std*.

- If using the Microsoft Visual C++ compiler, *db2xa_switch_std* or *db2xa_switch_static_std* can be defined as:

```
extern __declspec(dllimport) struct xa_switch_t db2xa_switch_std
extern __declspec(dllimport) struct xa_switch_t db2xa_switch_static_std
```

- By calling **db2xacic_std** or **db2xacicst_std**

DB2 provides this API, which returns the address of the *db2xa_switch_std* or *db2xa_switch_static_std* structure. This function is prototyped as:

```
struct xa_switch_t * SQL_API_FN db2xacic_std( )
struct xa_switch_t * SQL_API_FN db2xacicst_std( )
```

With any of these methods, you must link your application with *db2api.lib*.

Example C Code

The following code illustrates the different ways in which the *db2xa_switch_std* or *db2xa_switch_static_std* can be accessed via a C program on any DB2 V9.1 platform. Be sure to link your application with the appropriate library.

```
#include <stdio.h>
#include <xa.h>

struct xa_switch_t * SQL_API_FN db2xacic_std( );
```

```

#ifdef DECLSPEC_DEFN
extern __declspec(dllimport) struct xa_switch_t db2xa_switch_std;
#else
#define db2xa_switch_std (*db2xa_switch_std)
extern struct xa_switch_t db2xa_switch_std;
#endif

main( )
{
    struct xa_switch_t *foo;
    printf ( "%s \n", db2xa_switch_std.name );
    foo = db2xacic_std();
    printf ( "%s \n", foo->name );
    return ;
}

```

Related concepts:

- “X/Open distributed transaction processing model” on page 215

XA interface problem determination

When an error is detected during an XA request from the TM, the application program may not be able to get the error code from the TM. If your program abends, or gets a cryptic return code from the TP monitor or the TM, you should check the First Failure Service Log, which reports XA error information when diagnostic level 3 or greater is in effect.

You should also consult the console message, TM error file, or other product-specific information about the external transaction processing software that you are using.

The database manager writes all XA-specific errors to the First Failure Service Log with SQLCODE -998 (transaction or heuristic errors) and the appropriate reason codes. Following are some of the more common errors:

- Invalid syntax in the xa_open string.
- Failure to connect to the database specified in the open string as a result of one of the following:
 - The database has not been cataloged.
 - The database has not been started.
 - The server application’s user name or password is not authorized to connect to the database.
- Communications error.

Related concepts:

- “X/Open distributed transaction processing model” on page 215

Related reference:

- “xa_open string formats” on page 221

XA transaction manager configuration

Configuring IBM WebSphere Application Server

IBM WebSphere Application Server is a Java-based application server. It can use the DB2 Database for Linux, UNIX, and Windows XA support via the Java Transaction API (JTA) provided by the DB2 JDBC driver. Refer to IBM WebSphere documentation regarding how to use the Java Transaction API with WebSphere Application Server. WebSphere Application Server documentation can be viewed online at <http://www.ibm.com/software/webservers/appserv/infocenter.html>.

Configuring IBM TXSeries CICS

For information about how to configure IBM TXSeries CICS to use DB2 Database for Linux, UNIX, and Windows as a resource manager, refer to your *IBM TXSeries CICS Administration Guide*. TXSeries documentation can be viewed online at http://www.transarc.com/Library/documentation/websphere/WAS-EE/en_US/html/.

Host and iSeries database servers can participate in CICS-coordinated transactions.

Configuring IBM TXSeries Encina

Following are the various APIs and configuration parameters required for the integration of Encina Monitor and DB2 Database for Linux, UNIX, and Windows servers, or DB2 for z/OS and OS/390, DB2 for iSeries, or DB2 for VSE & VM when accessed through DB2 Connect. TXSeries documentation can be viewed online at http://www.transarc.com/Library/documentation/websphere/WAS-EE/en_US/html/.

Host and iSeries database servers can participate in Encina-coordinated transactions.

Configuring DB2 Database for Linux, UNIX, and Windows

To configure DB2 Database for Linux, UNIX, and Windows:

1. Each database name must be defined in the DB2 database directory. If the database is a remote database, a node directory entry must also be defined. You can perform the configuration using the Configuration Assistant, or the DB2 command line processor (CLP). For example:

```
DB2 CATALOG DATABASE inventdb AS inventdb AT NODE host1 AUTH SERVER
DB2 CATALOG TCP/IP NODE host1 REMOTE hostname1 SERVER svcname1
```

2. The DB2 client can optimize its internal processing for Encina if it knows that it is dealing with Encina. You can specify this by setting the *tp_mon_name* database manager configuration parameter to ENCINA. The default behavior is no special optimization. If *tp_mon_name* is set, the application must ensure that the thread that performs the unit of work also immediately commits the work after ending it. No other unit of work may be started. If this is *not* your environment, ensure that the *tp_mon_name* value is NONE (or, through the CLP, that the value is set to NULL). The parameter can be updated through the Control Center or the CLP. The CLP command is:

```
db2 update dbm cfg using tp_mon_name ENCINA
```


Configuring Encina for Each Resource Manager

To configure Encina for each resource manager (RM), an administrator must define the Open String, Close String, and Thread of Control Agreement for each DB2 database as a resource manager before the resource manager can be registered for transactions in an application. The configuration can be performed using the Enconcole full screen interface, or the Encina command line interface. For example:

```
monadmin create rm inventdb -open "db=inventdb,uid=user1,pwd=password1"
```

There is one resource manager configuration for each DB2 database, and each resource manager configuration must have an rm name ("logical RM name"). To simplify the situation, you should make it identical to the database name.

The xa_open string contains information that is required to establish a connection to the database. The content of the string is RM-specific. The xa_open string of DB2 contains the alias name of the database to be opened, and optionally, a user ID and password to be associated with the connection. Note that the database name defined here must also be cataloged into the regular database directory required for all database access.

The xa_close string is not used by DB2.

The Thread of Control Agreement determines if an application agent thread can handle more than one transaction at a time.

If you are accessing DB2 for z/OS and OS/390, DB2 for iSeries, or DB2 for VSE & VM, you must use the DB2 Syncpoint Manager.

Referencing a DB2 Database for Linux, UNIX, and Windows database from an Encina application

To reference a DB2 Database for Linux, UNIX, and Windows database from an Encina application:

1. Use the Encina Scheduling Policy API to specify how many application agents can be run from a single TP monitor application process. For example:

```
rc = mon_SetSchedulingPolicy (MON_EXCLUSIVE)
```

2. Use the Encina RM Registration API to provide the XA switch and the logical RM name to be used by Encina when referencing the RM in an application process. For example:

```
rc = mon_RegisterRmi ( &db2xa_switch, /* xa switch */
                    "inventdb",      /* logical RM name */
                    &rmiId );        /* internal RM ID */
```

The XA switch contains the addresses of the XA routines in the RM that the TM can call, and it also specifies the functionality that is provided by the RM. The XA switch of DB2 V9.1 is db2xa_switch, and it resides in the DB2 client library (db2app.dll on Windows operating systems and libdb2 on UNIX based systems).

The logical RM name is the one used by Encina, and is not the actual database name that is used by the SQL application that runs under Encina. The actual database name is specified in the xa_open string in the Encina RM Registration API. The logical RM name is set to be the same as the database name in this example.

The third parameter returns an internal identifier or handle that is used by the TM to reference this connection.

Related concepts:

Indoubt transaction management APIs

- “DB2 Connect and transaction processing monitors” in *DB2 Connect User’s Guide*

Related reference:

- “tp_mon_name - Transaction processor monitor name configuration parameter” in *Performance Guide*
- “xa_open string formats” on page 221

Configuring BEA Tuxedo

What follows is a description of the process to configure BEA Tuxedo for use with DB2 Database for Linux, UNIX, and Windows. There are some differences that are noted based on whether Tuxedo is working with a 64-bit instance of DB2 Database for Linux, UNIX, and Windows or a 32-bit instance of DB2 Database for Linux, UNIX, and Windows.

Note: There are new names for the XA switch data structures: *db2xa_switch_std* and *db2xa_switch_static_std*. There are also new names for the APIs: *db2xacic* and *db2xacicst*. The old switch data structure and API names can be used but only when working with a 32-bit instance of DB2 Database for Linux, UNIX, and Windows.

Procedure:

To configure Tuxedo to use DB2 Database for Linux, UNIX, and Windows as a resource manager, perform the following steps:

1. Install Tuxedo as specified in the documentation for that product. Ensure that you perform all basic Tuxedo configuration, including the log files and environment variables.
You also require a compiler and the DB2 Application Development Client. Install these if necessary.
2. At the Tuxedo server ID, set the DB2INSTANCE environment variable to reference the instance that contains the databases that you want Tuxedo to use. Set the PATH variable to include the DB2 program directories. Confirm that the Tuxedo server ID can connect to the DB2 databases.
3. Update the *tp_mon_name* database manager configuration parameter with the value TUXEDO.
4. Add a definition for DB2 V9.1 to the Tuxedo resource manager definition file. In the examples that follow, *UDB_XA* is the locally-defined Tuxedo resource manager name for DB2 V9.1, and *db2xa_switch_std* is the DB2-defined name for a structure of type *xa_switch_t*:
 - For AIX. In the file `${TUXDIR}/udataobj/RM`, add the definition:

```
# DB2 UDB
UDB_XA:db2xa_switch_std:-L${DB2DIR} /lib -ldb2
```

where `{TUXDIR}` is the directory where you installed Tuxedo, and `{DB2DIR}` is the DB2 instance directory.
 - For Windows. In the file `%TUXDIR%\udataobj\rm`, add the definition:

```
# DB2 UDB
UDB_XA;db2xa_switch_std;%DB2DIR%\lib\db2api.lib
```

where `%TUXDIR%` is the directory where you installed Tuxedo, and `%DB2DIR%` is the DB2 instance directory.
5. Build the Tuxedo transaction monitor server program for DB2:

- For AIX:

```
{TUXDIR}/bin/buildtms -r UDB_XA -o {TUXDIR}/bin/TMS_UBD
```

where {TUXDIR} is the directory where you installed Tuxedo.

- For Windows:

```
%TUXDIR%\bin\buildtms -r UDB_XA -o %TUXDIR%\bin\TMS_UBD
```

6. Build the application servers. In the examples that follow, the -r option specifies the resource manager name, the -f option (used one or more times) specifies the files that contain the application services, the -s option specifies the application service names for this server, and the -o option specifies the output server file name:

- For AIX:

```
{TUXDIR}/bin/buildserver -r UDB_XA -f svcfile.o -s SVC1,SVC2  
-o UDBserver
```

where {TUXDIR} is the directory where you installed Tuxedo.

- For Windows:

```
%TUXDIR%\bin\buildserver -r UDB_XA -f svcfile.o -s SVC1,SVC2  
-o UDBserver
```

where %TUXDIR% is the directory where you installed Tuxedo.

7. Set up the Tuxedo configuration file to reference the DB2 server. In the *GROUPS section of the UDBCONFIG file, add an entry similar to:

```
UDB_GRP  LMID=simp GRPNO=3  
TMSNAME=TMS_UBD TMSCOUNT=2  
OPENINFO="UDB_XA:db=sample,uid=db2_user,pwd=db2_user_pwd"
```

where the TMSNAME parameter specifies the transaction monitor server program that you built previously, and the OPENINFO parameter specifies the resource manager name. This is followed by the database name, and the DB2 database user ID and password, which are used for authentication.

The application servers that you built previously are referenced in the *SERVERS section of the Tuxedo configuration file.

8. If the application is accessing data residing on DB2 for z/OS and OS/390, DB2 for iSeries, or DB2 for VM&VSE, the DB2 Connect XA concentrator will be required.
9. Start Tuxedo:

```
tmbboot -y
```

After the command completes, Tuxedo messages should indicate that the servers are started. In addition, if you issue the DB2 command LIST APPLICATIONS ALL, you should see two connections (in this situation) specified by the TMSCOUNT parameter in the UDB group in the Tuxedo configuration file, UDBCONFIG.

Related concepts:

- “DB2 Connect and transaction processing monitors” in *DB2 Connect User’s Guide*

Related reference:

- “LIST APPLICATIONS command” in *Command Reference*
- “tp_mon_name - Transaction processor monitor name configuration parameter” in *Performance Guide*

Indoubt transaction management APIs

Part 3. Appendixes

Appendix A. Incompatibilities between releases

This section identifies the incompatibilities that exist between DB2 Version 9 and previous releases of DB2 Universal Database.

An *incompatibility* is a part of DB2 database that works differently than it did in a previous release. If used in an existing application, it will produce an unexpected result, require a change to the application, or reduce performance. In this context, "application" refers to:

- Application program code
- Third-party utilities
- Interactive SQL queries
- Command or API invocation.

Incompatibilities introduced in DB2 Universal Database Version 8 and DB2 Version 9 are described. They are grouped according to the following categories:

- System Catalog Information
- Application Programming
- SQL
- Database Security and Tuning
- Utilities and Tools
- Connectivity and Coexistence
- Messages
- Configuration Parameters.

Each incompatibility section includes a description of the incompatibility, the symptom or effect of the incompatibility, and possible resolutions. There is also an indicator at the beginning of each incompatibility description that identifies the operating system to which the incompatibility applies:

Windows

Microsoft Windows® platforms supported by DB2 databases

UNIX UNIX®-based platforms supported by DB2 databases

Deprecated and discontinued features

This section describes current and future deprecated and discontinued features. In addition, any planned incompatibilities that users of DB2 database systems should keep in mind when coding new applications, or when modifying existing applications are presented here. Knowing about these changes will facilitate your current application development and future planning to move to newer versions of DB2.

For example, in the reference material you will see the new parameters in the command or SQL statement syntax and description with a note that says that this new parameter is replacing another parameter. However, that other, older parameter will continue to be recognized for a period of time. The actual time for this continued support of the older parameter is not explicitly stated because it is

difficult to predict the future. The time of overlap between the old and the new parameter allows you time to plan for the changes to be applied to your applications.

There are also new functions or features for the product that are first discussed in the "What's New" document.

Some of these deprecated function or features, or the new functions and features, will have an impact for you if you are a current customer of our product. The impacts are outlined as part of the discussion on migration.

What follows is a list of those differences in the current release from the previous release of DB2.

System catalog information:

PK_COLNAMES and FK_COLNAMES in a future version of DB2:

Operating systems affected:

All supported operating systems are affected.

Change:

The SYSCAT.REFERENCES columns PK_COLNAMES and FK_COLNAMES will no longer be available.

Symptom:

When referenced, an error is returned because the columns no longer exist.

Explanation:

These columns are obsolete and have been replaced.

Resolution:

Change your tools or applications that reference the SYSCAT.REFERENCES columns PK_COLNAMES and FK_COLNAMES to use the SYSCAT.KEYCOLUSE view instead.

COLNAMES no longer available in a future version of DB2:

Operating systems affected:

All supported operating systems are affected.

Change:

The SYSCAT.INDEXES column COLNAMES will no longer be available.

This column will only contain valid information if the column names are less than or equal to 30 bytes, and if there are less than or equal to 16 columns in the index. Either a blank or a NULL value is returned if any column name exceeds 30 bytes, or if there are greater than 16 columns.

Symptom:

Column does not exist and an error is returned.

Explanation:

Tools or applications are coded to use the obsolete COLNAMES column.

Resolution:

Change the tool or application to use the SYSCAT.INDEXCOLUSE view instead.

Application programming:**iCheckPending parameter of the db2Load API is deprecated:****Operating systems affected:**

All supported operating systems are affected.

Change:

The iCheckPending parameter of the db2load API is deprecated. The replacement parameter is iSetIntegrityPending.

Explanation:

The iCheckPending parameter of the db2Load API is deprecated. It is an input parameter of the db2Load API to specify whether a table should be put into the check pending state.

Note: The set integrity pending state replaces the check pending state. They are equivalent states.

Resolution:

Use the iSetIntegrityPending parameter with the db2Load API. The values to use with this new parameter are: SQLU_SI_PENDING_CASCADE_IMMEDIATE or SQLU_SI_PENDING_CASCADE_DEFERRED.

User defined functions (UDFs) and procedures to be deprecated:**Operating systems affected:**

All supported operating systems are affected.

Change:

The following UDFs are deprecated: GET_DBM_CONFIG, SNAP_GET_CONTAINER, SNAP_GET_DB, SNAP_GET_DYN_SQL, SNAP_GET_STO_PATHS, SNAP_GET_TAB, SNAP_GET_TBSP, SNAP_GET_TBSP_PART, SNAPSHOT_AGENT, SNAPSHOT_APPL, SNAPSHOT_APPL_INFO, SNAPSHOT_BP, SNAPSHOT_CONTAINER, SNAPSHOT_DATABASE, SNAPSHOT_DBM, SNAPSHOT_DYN_SQL, SNAPSHOT_FCM, SNAPSHOT_FCMNODE, SNAPSHOT_LOCK, SNAPSHOT_LOCKWAIT, SNAPSHOT QUIESCERS, SNAPSHOT_RANGES,

SNAPSHOT_STATEMENT, SNAPSHOT_SUBSECT, SNAPSHOT_SWITCHES, SNAPSHOT_TABLE, SNAPSHOT_TBREORG, SNAPSHOT_TBS, SNAPSHOT_TBS_CFG, SQLCACHE_SNAPSHOT.

The following procedures are deprecated: GET_DB_CONFIG, SNAPSHOT_FILEW, and SYSINSTALLROUTINES.

Deprecating these routines means that there will be no further investment in the routines. The documentation for the routines is updated to indicate that the routine is deprecated, but is being maintained for compatibility. At some point in the future these UDFs and routines will be removed from the catalogs and documentation.

Explanation:

New routines based on the SQL Administrative API standards, are replacing old functions created before the standards were implemented.

Resolution:

New equivalent functions with similar names beginning with SNAP_GET_ are available. Different parameters and additional columns may be associated with the new functions. You should review the differences before using the replacement functions within your applications.

For additional information on the deprecated UDFs and procedures, and the new equivalent functions and views, refer to “Deprecated SQL administrative routines and their replacement routines or views”.

Use the new functions, routines and views; applications using the old functions and procedures should consider a plan to update the applications by moving to the new functions and routines. The old functions and procedures will continue to be supported for compatibility. However, this support will be removed in a future version or release of the product.

Default function entry points in external routine libraries is deprecated:

Operating systems affected:

Only the 32-bit AIX and Windowsoperating systems are affected.

Change:

In some future version or release, we will no longer support loading a library name and assuming the default entry point.

Within the AIX and Windowsoperating system environments, support for the default function entry points in external routine libraries is deprecated.

Explanation:

There is a risk of instance failure when only specifying the library name and using the default entry point when routines are run in trusted (not fenced) mode.

Resolution:

From this point forward, when creating stored procedures and functions, do not rely on the database manager to resolve and load the function specified by the default entry point. Instead, specify the complete entry point and library name when loading routine libraries. For new routines, specify the *!proc-id* (for a procedure) or the *!func-id* (for a function) value as part of the EXTERNAL NAME clause value. For existing routines, provide an explicit entry point value for routine definitions that specify the EXTERNAL NAME clause. This can be done using the ALTER FUNCTION statement.

Remove type-1 index support:

Operating systems affected:

All supported operating systems are affected.

Change:

The type-1 index support has been removed.

Explanation:

A new type of index was introduced in Version 8, called a type-2 index. With type-1 indexes, that is indexes created prior to Version 8, a key is physically removed from a leaf page as part of the deletion or update of a table row. With type-2 indexes, keys are marked as deleted when a row is deleted or updated, but they are not physically removed until after the deletion or update is committed. When support for re-creation of type-1 indexes is removed, you will not have to rebuild your indexes manually. Type-1 indexes will continue to function correctly. All actions that result in the re-creation of indexes will automatically convert type-1 indexes to type-2 indexes. In a future version, support for type-1 indexes will be removed.

Resolution:

Use the newer type-2 indexes. This can be done by converting the older indexes manually (by request during a REORG of the indexes). All new indexes use the new type-2 indexes.

DB2 JDBC type 2 driver is deprecated:

Operating systems affected:

All supported operating systems are affected.

Change:

The DB2 JDBC type 2 driver was deprecated in version 8.2, and remains deprecated in version 9.1. Support for the driver will be removed in a future release.

Resolution:

Use the IBM DB2 Driver for JDBC and SQLJ.

Remove Type 3 JDBC driver support:

Operating systems affected:

All supported operating systems are affected.

Change:

The Type 3 JDBC driver support has been removed.

Explanation:

The db2jd is not be shipped with the product.

Resolution:

Use the IBM DB2 Driver for JDBC and SQLJ.

Application libraries have changed:**Operating systems affected:**

All supported operating systems are affected.

Change:

The following changes have been made:

- db2app.dll was extended. It includes its original information, plus the information from the db2util.dll, db2abind.dll, and db2cli.dll libraries.
- db2api.dll was extended. It includes its original information, plus the information from the db2cli.dll library.

Explanation:

The library information is being consolidated.

Resolution:

Stubs for the db2util.dll, db2abind.dll, and db2cli.dll libraries are still available for backwards compatibility. These stubs will be removed in a future version or release of the product.

SQL:**Some SQL administrative routines have been replaced:****Operating systems affected:**

All supported operating systems are affected.

Change:

Some of the existing administrative routines have been replaced by newer, more comprehensive routines or views.

Explanation:

Expanding the support of SQL administrative routines in Version 9 required the replacement of some existing routines.

Resolution:

Applications that use Version 8 table functions should be modified to use the new functions or administrative views. The new table functions have the same base names as the original functions but are suffixed with “_Vxx” to identify the version of the product in which they were added. However, the administrative views will always be based on the most current version of the table functions, and therefore allow for more application portability.

For additional information on the new routines, refer to “Deprecated SQL administrative routines and their replacement routines or views”.

Partitioning key to distribution key terminology change:

Operating systems affected:

All supported operating systems are affected.

Change:

The term “partitioning key” is changed to “distribution key”. A distribution key is a column (or group of columns) that is used to determine the database partition in which a particular row of data is stored. A table partitioning key is an ordered set of one or more columns that is used to determine the data partition in which each table row belongs.

Explanation:

The introduction of table partitioning required that there be a redefinition of “partitioning key”.

Resolution:

The term “distribution key” is used in the documentation where it once was “partitioning key”.

PARTITIONING KEY clause changes on the ALTER TABLE statement:

Operating systems affected:

All supported operating systems are affected.

Change:

The ADD PARTITIONING KEY clause of the ALTER TABLE statement is being deprecated. This clause is being replaced by the ADD DISTRIBUTE BY HASH clause.

The DROP PARTITIONING KEY clause of the ALTER TABLE statement is being deprecated. This clause is being replaced by the DROP DISTRIBUTION clause.

Explanation:

The introduction of table partitioning required that there be a redefinition of “partitioning key” resulting in changed syntax on the ALTER TABLE statement.

Resolution:

The old PARTITIONING KEY clause in the syntax is supported for backwards compatibility. However, the clause not be supported in a future release of the product. Therefore, you should plan to convert applications that use this old syntax to the new syntax.

Database security and tuning:

Extended storage is no longer supported:

Operating systems affected:

All supported operating systems are affected.

Change:

Extended storage is no longer supported. With DB2 products moving to 64-bit environments, the need for extended storage is removed.

Following your migration to DB2 Version 9, values in the catalog views will change. For example, the ESTORE column within the SYSCAT.BUFFERPOOLS catalog view will always be “N”. Any data definition language (DDL) that may be run to attempt to change this value will be tolerated but will have no effect.

There are monitor elements that are still present, but are deprecated in Version 9. The four monitor elements are:

- *pool_data_to_estore*
- *pool_index_to_estore*
- *pool_data_from_estore*
- *pool_index_from_estore*

In a future release or version of DB2 products, using the monitor elements relating to extended storage, and the output generated by the contents of those elements when requesting a GET SNAPSHOT command, will no longer be available.

In addition, the configuration parameters for extended storage (*estore_seg_sz* and *num_estore_segs*) are no longer valid in Version 9.

The “ESTORE” column from the SYSCAT.BUFFERPOOLS catalog will also be removed in a future release or version.

Explanation:

Extended storage acted as an extended look-aside buffer for the main buffer pools. It allowed for memory performance improvements which took advantage of computers with large amounts of main memory. For computers with 64-bit environments, extended storage and other similar methods are no longer needed.

Resolution:

Extended storage should no longer be used. You should plan not to use the extended storage configuration parameters, nor the extended storage monitor elements.

Utilities and tools:

Desktop icon and folder making utility no longer supported (Linux):

Operating systems affected:

Only the Linux operating system is affected.

Change:

This release no longer includes a set of utilities for the creation of DB2 desktop folders and icons for launching commonly used product tools on the Gnome and KDE desktops for supported Intel[®]-based Linux distributions.

db2ilist command has deprecated options (Linux and UNIX):

Operating systems affected:

Only the Linux and supported UNIX operating systems are affected.

Change:

The **db2ilist** command has the following command options deprecated:

- *-w* (list the bitwidth for each instance)
- *-a* (list both regular and AFP[™] instances)
- *-p* (list the path for each instance)

Explanation:

In the past, the **db2ilist** command could be used to list all available instances on a system. Now, the **db2ilist** command only lists the instances related to the current installation path and only one type of instance on each UNIX or Linux platform.

Resolution:

The **db2ilist** command can still be used. The deprecated options on the command should not be used.

db2reg2large utility for converting DMS table space size is no longer available:

Operating systems affected:

All supported operating systems are affected.

Change:

The **db2reg2large** utility, which is used for converting Regular DMS table spaces to Large DMS table spaces has been discontinued in DB2 Version 9.

Resolution:

This utility has been replaced with a new CONVERT TO LARGE option on the ALTER TABLESPACE SQL statement.

db2profc and db2profp utilities are discontinued:

Operating systems affected:

All supported operating systems are affected.

Change:

In previous releases, db2profc was accepted as an alternative name for db2sqljcustomize, and db2profp was accepted as an alternative name for db2sqljprint. These alternative names are no longer accepted.

Explanation:

The DB2 JDBC Type 2 Driver originally used the name db2profc for the SQLJ profile customizer command, and the name db2profp for the SQLJ profile printer command.

Resolution:

For the IBM DB2 Driver for JDBC and SQKJ, the SQLJ profile customizer command is named db2sqljcustomize, and the SQLJ profile printer command is name db2sqljprint. Use these commands instead of db2profc and db2profp.

Set permissions for database objects (db2secv82) is deprecated:

Operating systems affected:

All supported operating systems are affected.

Change:

The set permissions for database objects (**db2secv82**) command is deprecated.

Explanation:

The name of the command suggested that it was only for use with Version 8.2 of the product (db2secv82). The new name will be for use in the current release and in future releases.

Resolution:

Use the set permissions for database objects (**db2extsec**) command in place of the set permissions for database objects (**db2secv82**) command. You should locate and change references to the **db2secv82** command in your applications and scripts and develop a plan to replace those references with ones to the **db2extsec** command.

db2look tool behavior changes:

Operating systems affected:

All supported operating systems are affected.

Change:

On systems using the Database Partitioning Feature (DPF), table space data definition language (DDL) may not be complete if some database partitions are not active. When requesting DDL on systems using DPF, a warning message is displayed in place of the DDL for table spaces that exist on inactive database partitions.

Explanation:

The use of automatic re-sizing and automatic storage across databases partitions, and the resulting need to gather data using a snapshot approach, requires that each database partition be active.

Resolution:

To ensure proper DDL is produced for all table spaces, all database partitions must be activated.

WordWidth parameter (-w option) of the db2icrt, db2ilist, and db2iupdt commands is ignored and deprecated:**Operating systems affected:**

All supported operating systems are affected.

Change:

The WordWidth (-w) option of the **db2icrt**, **db2ilist**, and **db2iupdt** commands is ignored when used and is being deprecated. This option provided the instance width in bits.

Resolution:

There is no effect if this option continues to be specified. The option is only valid on AIX 5L™, HP-UX, Linux, and the Solaris operating systems.

Manual installation:**Operating systems affected:**

Only Linux and UNIX operating systems are affected.

Change:

Manual installation, uninstallation, or querying of DB2 products using native Linux or UNIX operating system utilities such as pkgadd, rpm, SMIT, or swinstall is not supported.

Explanation:

To better manage and control the installation process, manual installation or uninstallation of DB2 products is no longer supported.

Resolution:

Use the **db2_install** command which has new parameters to support new function. The **db2_deinstall** command is part of the base installation image.

Support for Lock Object Name will be removed:

Operating systems affected:

All supported operating systems are affected.

Change:

The Lock Object Name that is part of the snapshot monitor sample output provides no value and contains redundant information from the Lock Name part of the output. The monitor element "lock_object_name" will be deprecated in a future release.

Explanation:

The output report from the snapshot monitor produces a list of locks. The Lock Name is the first item in the list. This information is taken from the monitor element "lockname". Later in the report, the Lock Object Name is shown. This information is taken from the monitor element "lock_object_name". The information presented as part of this item could also have been extracted from the value given for the monitor element "lockname".

Resolution:

The monitor element "lock_object_name" will be deprecated in a future release. The information it provides is also going to be removed from snapshot monitor output.

You should plan not to use the GET SNAPSHOT FOR LOCKS ON <dbname> command to return the "Lock Object Name" in any new or revised applications.

Remove raw log support:

Operating systems affected:

All supported operating systems are affected.

Change:

The raw device support for logging has been removed.

Explanation:

Increasing dedicated storage subsystems and full support of self-managing DMS storage is reducing the need for detailed storage management.

Resolution:

Do not use raw devices for logging. You may need to change the newlogpath database configuration parameter setting to a disk device instead of a raw device. Remember to stop and restart the database manager to make the new setting for the configuration parameter effective.

Changes to db2batch:

Operating systems affected:

All supported operating systems are affected.

Symptom:

The **db2batch** command now runs only in CLI mode. CLI mode used to be specified using the *-cli* option. Embedded dynamic SQL was the default mode, but this has been changed so that the command only runs in CLI mode. Also, scripts that perform REBIND or BIND on `db2batch.bnd` will fail because a “bnd” file is no longer shipped.

In addition, the *-p* option is not available.

Explanation:

The parallel option on the **db2batch** command is no longer supported.

Resolution:

You can continue to use the *-cli* option for backward compatibility only, but it has no effect. You can change the default isolation level by specifying the *TxnIsolation* configuration keyword in the `db2cli.ini` file. The new *-iso* option is used to specify the isolation level.

You can no longer use the *-p* option on the **db2batch** command.

Support for db2uiddl tool will be removed:

Operating systems affected:

All supported operating systems are affected.

Change:

Indexes which do not support deferred unique semantics will no longer be supported. The **db2uiddl - Prepare Unique Index Conversion to V5 Semantics** command will no longer be supported after DB2 Version 9.

Explanation:

In DB2 Universal Database (UDB) Version 5, the semantics of unique indexes were changed to deferred unique. To support this change, the **db2uiddl** tool was introduced to convert unique indexes to the new semantics. When databases using pre-Version 5 unique index semantics are migrated, all unique indexes are not automatically changed to Version 5 semantics because converting unique indexes is a very time-consuming operation, and you will want to manage the conversion based on your business needs.

Resolution:

Develop a plan to convert all unique indexes created prior to DB2 UDB Version 5 to the new deferred unique index semantics before support for **db2uiddl** is removed. The **db2uiddl** tool searches the system catalogs for indexes without

deferred unique semantics and writes CREATE UNIQUE INDEX statements for the indexes that require conversion. These statements are stored in a file which must be run after migration to Version 9 is successful. This will ensure that the indexes are converted prior to the deprecation of the **db2uiddl** tool.

Support for db2undgp tool will be removed:

Operating systems affected:

All supported operating systems are affected.

Change:

Execute privileges of routines (functions, procedures, and methods) are now controlled by the SYSCAT.ROUTINEAUTH system catalog view. The **db2undgp** command will no longer be available after DB2 Version 9.

Explanation:

In DB2 Universal Database (UDB) Version 8, a system catalog view, SYSCAT.ROUTINEAUTH, was added to control the EXECUTE privileges of routines (functions, procedures, and methods). During database migration to DB2 UDB Version 8, the EXECUTE privilege for all existing functions, methods, and external stored procedures are granted to all users (PUBLIC). This results in a security exposure for external stored procedures that access SQL data. The **db2undgp** command is used to prevent users from accessing SQL objects which they do not have privileges for.

Resolution:

Develop a plan to revoke the EXECUTE privilege from the PUBLIC group before the **db2undgp** tool is deprecated.

Support for re-creation of type-1 indexes will be removed:

Operating systems affected:

All supported operating systems are affected.

Change:

A new type of index was introduced in Version 8, called a type-2 index. With type-1 indexes, that is indexes created prior to Version 8, a key is physically removed from a leaf page as part of the deletion or update of a table row. With type-2 indexes, keys are marked as deleted when a row is deleted or updated, but they are not physically removed until after the deletion or update is committed. When support for re-creation of type-1 indexes is removed, you will not have to rebuild your indexes manually. Type-1 indexes will continue to function correctly. All actions that result in the re-creation of indexes will automatically convert type-1 indexes to type-2 indexes. In a future version, support for type-1 indexes will be removed.

Explanation:

Type-2 indexes have advantages over type-1 indexes:

- A type-2 index can be created on columns whose length is greater than 255 bytes

- The use of next-key locking is reduced to a minimum, which improves concurrency.

Resolution:

Develop a plan to convert your existing indexes to type-2 indexes over time. The Online Index Reorganization capability can help do this while minimizing availability outages. Increase index table space size if needed. Consider creating new indexes in large table spaces and moving existing indexes to large table spaces.

Note: If you convert pre-Version 5 indexes to type-2, you do not need to run the `db2uiddl` tool.

Connectivity and coexistence:

CLI keyword CLISHEMA no longer supported:

Operating systems affected:

All supported operating systems are affected.

Change:

For DB2 clients connecting to DB2 for Linux, UNIX, and Windows DB2 database servers, the CLISchema keyword is deprecated.

For DB2 clients connecting to DB2 for z/OS database servers, the CLISchema keyword is dropped.

Resolution:

DB2 clients should no longer use the CLISchema keyword. One keyword that is similar to CLISchema is SysSchema.

Data Warehouse Manager is no longer included:

Operating systems affected:

All supported operating systems are affected.

Change:

The DB2 Warehouse Manager Standard Edition is not available for this release. The Data Warehouse Center and the Information Catalog Center are not included in this release.

Resolution:

These products and centers are being developed and released separately from the base DB2 Version 9 product.

Text Extender is no longer supported:

Operating systems affected:

All supported operating systems are affected.

Change:

DB2 Text Extender is not supported in this release.

Resolution:

A direct replacement function is not available. However, there are other full-text search products capable of performing similar tasks. For example, there are DB2 Net Search Extender that is very like the Text Extender; and, WebSphere Information Integrator OmniFind™ Edition that provides an enterprise search solution for finding the most relevant corporate information. The information can be found not only in a relational database but searched across intranets, extranets, corporate public Web sites, and a wide range of content repositories.

Audio, Image, and Video (AIV) Extenders are no longer supported:

Operating systems affected:

All supported operating systems are affected.

Change:

Audio, Image, and Video (AIV) Extenders are no longer supported in this release.

Resolution:

You might consider implementing your own extensions similar to the AIV Extenders to enhance the DB2 functionality using DB2 user-defined functions and third party software.

Platform support changes for the DB2 Administration Tools:

Operating systems affected:

Only supported Windows and Linux operating systems are affected.

Change:

In previous releases, the DB2 Administration Tools, including the Control Center, were supported on all platforms. In Version 9, the DB2 Administration Tools, are supported only on Windows x86, Windows x64 (AMD64 or EM64T), Linux on x86, and Linux on AMD64 or EM64T.

32-bit instance support changes:

Operating systems affected:

All supported operating systems are affected.

Change:

In response to market demand, a priority is being placed on DB2 database server support for 64-bit hardware and operating systems. The number of supported

32-bit platforms is being reduced. Support for 32-bit Windows and Linux platforms will continue since those platforms are often preferred for building or running small and medium business applications.

Configuration parameters and registry variables:

Deprecated configuration parameters:

Operating systems affected:

All supported operating systems are affected.

Change:

The following configuration parameters are deprecated:

- *estore_seg_sz*
- *num_estore_segs*
- *min_priv_mem*
- *priv_mem_thresh*; Use the DB2MEMMAXFREE registry variable in its place.
- *fcnum_rqb*
- *fcnum_anchors*
- *fcnum_connect*

A value may be set for each of these configuration parameters but the value is ignored. (That is, the value will have no effect.)

Deprecated registry variables:

Operating systems affected:

All supported operating systems are affected.

Change:

The following registry variables are deprecated:

- DB2_FORCE_FCM_BP; the default value is changed from “No” to “Yes”.
- DB2_LGPAGE_BP; Use the DB2_LARGE_PAGE_MEM registry variable in its place.
- DB2LINUXAIO

The following registry variables are deprecated:

- DB2_SCATTERED_IO; Default is to always read from disk on Linux.

Related reference:

- “Deprecated SQL administrative routines and their replacement routines or views” in *Administrative SQL Routines and Views*

Version 9 incompatibilities with previous releases and changed behaviors

This section describes current incompatibilities that users of DB2 database systems should keep in mind when coding new applications, or when modifying existing applications. This will facilitate your current application development and future planning to move to newer versions of DB2. An incompatibility typically involves a change in defaults for product functions and feature; or, it involves a different requirement or outcome from what would have occurred in the previous version of the DB2 product. For example, if you used an SQL statement last version or release, you expect a certain behavior or result. If you used the same SQL statement this version and you receive a different behavior or result that is not expected, then there is an incompatibility between the last version (or release) and the current version. Those differences or “incompatibilities” are documented here.

Note: Although an attempt has been made to list all of the currently known product incompatibilities, there may be more recent incompatibilities documented in the product release notes.

System catalog information:

SETTING fields will be changing:

Windows	UNIX
---------	------

Change:

There is a change in type (and length) of the SETTING fields in specific catalogs.

Symptom:

Existing application programs referencing SETTING in the ORDER BY clause, or WHERE SETTING IN... clause, or WHERE SETTING= clause will fail.

Explanation:

The SETTING files in the following catalogs have changed from VARCHAR(255) to CLOB(32K):

- SYSCAT.TABOPTIONS
- SYSCAT.COLOPTIONS

Application programs that SELECT the SETTING fields from these catalogs will need to be rewritten because of the restrictions on large objects (LOBs) for SQL statements.

Resolution:

Rewrite your application programs that SELECT the SETTING fields from the catalogs specified.

Application Programming:

Application ID format changed:

Windows	UNIX
---------	------

Change:

The format of the application ID has changed.

Explanation:

The new format presents the port number and IP address in a readable form that also accommodates the longer IPv6 addresses.

Resolution:

If you have scripts that parse output that contains the application ID, you will need to modify the parsing conditions to account for the new format. For example, you may parse the output from the LIST APPLICATIONS command.

DB2 Embedded Application Server updated:

Windows	UNIX
---------	------

Change:

The DB2 Embedded Application Server enables you to run the Web applications supplied with DB2 without needing to purchase an application server. In Version 8, the DB2 Embedded Application Server was also referred to as the *application server for DB2 UDB*.

The XML Metadata Repository (XMR) application is not longer supported as one of the applications with DB2 Embedded Application Server.

Resolution:

For users of the XMR application in Version 8, it is necessary to uninstall XMR and find a replacement product. WebSphere offers several suitable replacement products.

IBM Software Development Kit (SDK) for Java 5.x is now supported:

Windows	UNIX
---------	------

Change:

The IBM Software Development Kit (SDK) for Java 5.x is now supported on the following operating system platforms: AIX 5, Linux on x86, Linux on AMD64/EM64T, Linux on zSeries[®], Linux on POWER[™], Windows x86 and Windows x64.

Resolution:

The IBM SDK is automatically installed on the server. If the client tools are installed, the IBM SDK is also installed on the client. If you are using the JDBC drivers with your own applications, you need to ensure the IBM SDK is installed.

Application and routine feature support changes:

Windows	UNIX
---------	------

Change:

The removal of support for most 32-bit database instances has resulted in changes in support for application and routines.

Symptom:

Client applications using DB2 Version 6 or Version 7 client instances cannot connect to DB2 Version 9 database servers.

There are new environment variable values within the client application environment.

32-bit unfenced routines (stored procedures and user-defined functions) created in DB2 Universal Database Version 8 will no longer work on 64-bit DB2 database servers in the AIX, HP, SUN, Linux on POWER, Linux for AMD64 and Intel EM64T, and Linux on zSeries environments. Migrating these routines to DB2 Version 9 requires that you rebuild them on the target 64-bit database server.

SQL procedures that you created for 32-bit instances of DB2 Universal Database Version 8 with any of FixPak 1 through FixPak 6 will not run on 64-bit instances of DB2 Version 9. To successfully migrate these SQL procedures to DB2 Version 9, you must drop and recreate the SQL procedures using the target 64-bit database server. SQL procedures created for 32-bit instances of DB2 Universal Database Version 7 or Version 8 with any FixPak will continue to work on the supported 32-bit instances of DB2 Version 9.

Explanation:

The removal of support for most 32-bit database instances has resulted in changes in support for application and routines.

Resolution:

You will need to consider whether you need to remain using a 32-bit database instance as a result of the product changes, or if you should move to a 64-bit database instance.

For example, only a 64-bit JVM is provided with 64-bit DB2 database servers. A 32-bit JVM is provided only for the Linux x86 and Windows on x86 operating systems. Finally, Java external routines require a 32-bit JVM for 32-bit DB2 database servers and a 64-bit JVM for 64-bit DB2 database servers.

New shipped functions and procedures:

Windows	UNIX
---------	------

Change:

New functions, function signatures for existing functions, or procedures have been added to the set of routines shipped with the product.

Symptom:

If a user-defined function or user-defined procedure has the same name and signature as a new shipped function or procedure, an unqualified reference to that function or procedure in a dynamic SQL statement will now execute the shipped function or procedure and not the user-defined one. Note that this does not affect static SQL in packages or SQL objects such as views, triggers, or SQL functions which will continue to execute the user-defined function or procedure until an explicit bind of the package or drop and create of the SQL object.

Explanation:

The default SQL path contains the schemas SYSIBM, SYSFUN, SYSPROC, and SYSIBMADM before the schema name which is the value of the USER special register. These system schemas are also usually included in the SQL path when it is explicitly set using the SET PATH statement or the FUNCPATH bind option. When function resolution and procedure resolution is performed, the shipped functions and procedures in these schemas will be considered before user-defined functions and user-defined procedures.

In Version 9, the following functions and procedures were added to the set of shipped functions and procedures:

- CHARACTER_LENGTH
- OCTET_LENGTH
- POSITION
- SECLABEL
- SECLABEL_BY_NAME
- SECLABEL_TO_CHAR
- STRIP
- SUBSTRING
- TRIM
- XMLCOMMENT
- XMLDOCUMENT
- XMLQUERY
- XMLTEXT
- XMLVALIDATE
- XMLXSROBJECTID

In Version 9, new administrative functions and procedures were added. Since the naming convention used for these functions and procedures make it more unlikely that a user-defined function or user-defined procedure would have the same name, they are not listed here. See the *Administrative SQL Routines and Views* for a list of these functions and procedures.

Resolution:

Rename the user-defined function or user-defined procedure; or, fully qualify the name to invoke it. Otherwise, you will be using the shipped function or procedure. Alternatively, the schema in which the user-defined function or user defined procedure exists can be placed in the SQL path before the schema in which the shipped function or procedure exists. However, doing this will increase the time it will take to resolve to all shipped functions and procedures since the schemas before the system schemas will be considered first.

Change to LIST APPLICATIONS output:

Windows	UNIX
---------	------

Change:

There are two new agents that will be included as part of the LIST APPLICATIONS command.

Explanation:

There are two new agents (db2stmm and db2taskd) requiring connection to the database at all times. As a result, there are two new agents that will be included as part of the LIST APPLICATIONS command. If you have any scripts designed to monitor the output from the LIST APPLICATIONS command, they will need to be modified based on these two new agents.

Resolution:

Modify any scripts designed to monitor the output from the LIST APPLICATIONS command to account for the presence of the two new agents.

Default size of DMS table spaces:

Windows	UNIX
---------	------

Change:

The new default size for DMS table spaces is "large".

Symptom:

There may be an increase in the amount of storage used if you have scripts that are used to create DMS table spaces and that do not explicitly specify the size (whether regular or large). The old default was "regular", the new default is "large".

Explanation:

The default size for DMS table spaces is "large" which takes up more space than a "regular" table space.

Resolution:

For those scripts which you use to create table spaces and have, in the past, simply accepted the default, you should consider modifying the scripts by adding explicit requests for a "regular" table space size if that is what you wish.

SQL:

SQL procedures can no longer use cursor blocking:

Windows	UNIX
---------	------

Change:

Cursor blocking can no longer be used for SQL procedures, regardless of the value that you specify for the BLOCKING bind option. The data is always received one row at a time.

Explanation:

There is a new limitation that applies to FETCH statements as well as FETCH statements that are implicitly contained in FOR loops.

Resolution:

Review those applications where you use cursor blocking. These applications might need to be modified based on this change in behavior.

Lock lists require additional space:

Windows	UNIX
---------	------

Change:

The space required by each lock in a lock list has changed such that a lock list of a given size can no longer represent as many locks as it once did.

Explanation:

Lock sizes have changed as follows:

- On 32-bit platforms, each lock requires 48 bytes of the lock list to record a lock on an object that has an existing lock on it. The lock requirement was 40 bytes.
- On 64-bit HP-UX/PA-RISC systems, each lock requires 80 bytes of the lock list to record a lock on an object that has an existing lock on it. The lock requirement was 64 bytes.

Resolution:

You may need to modify your lock list size.

New function SYSIBM.LOCATE replacing SYSFUN.LOCATE:

Windows	UNIX
---------	------

Change:

A new function SYSIBM.LOCATE is shipped in Version 9.1 that extends the functionality present in SYSFUN.LOCATE.

Symptom:

If the LOCATE function is used in an application that is compiled in Version 9.1 without the schema name qualification, the new SYSIBM.LOCATE will be invoked. SYSIBM.LOCATE may return results that are different from SYSFUN.LOCATE in some cases.

Explanation:

The SYSIBM.LOCATE extends the functionality of SYSFUN.LOCATE by adding character semantics to the LOCATE function and accepting graphic string arguments. Though the existing syntax supported by SYSFUN.LOCATE continues to work and have the same semantics using SYSIBM.LOCATE when used with OCTETS or without CODEUNITS specification, there are few cases where the results could be different.

One such case occurs when a search is performed on a graphic string data type. In releases before Version 9.1, in a Unicode database, graphic strings would be converted to character strings before the function was invoked. The position at which the search succeeded is counted in terms of bytes for SYSFUN.LOCATE, whereas it is counted in terms of units of TWO bytes for SYSIBM.LOCATE when no CODEUNITS specification is specified.

For example, the following returns the value 2:

```
VALUES SYSIBM.LOCATE(GX'0040', GX'D8000040')
```

whereas the following returns the value 4:

```
VALUES SYSFUN.LOCATE(GX'0040', GX'D8000040')
```

Another difference occurs when a 2-byte graphic character search string that may occur as a byte pattern that spans two “real characters” in the source string. For example searching for GX'2233 in a string GX'11223344 succeeds with SYSFUN.LOCATE, but returns 0 (NOT FOUND) with SYSIBM.LOCATE. This is because SYSFUN.LOCATE does a byte-based search and SYSIBM.LOCATE does a character-based search. The characters in the source string are “1122” and “3344”. There is no character “2233”, it is just a byte pattern that is present; but straddles two characters.

Another difference is that SYSIBM.LOCATE does some character validations that are not done by SYSFUN.LOCATE and may give different results when there are invalid characters in the search string.

Resolution:

If the difference in results caused by illegal characters is acceptable, users can use OCTETS as the codeunit specification to get the same behavior as SYSFUN.LOCATE. If the exact behavior of the SYSFUN.LOCATE function is required, applications can use the function name qualified with the schema name. Our recommendation is to adapt the application to use the new SYSIBM.LOCATE as it offers more functionality.

SQL functions processing based on specified units:

Windows	UNIX
---------	------

Change:

Not all SQL functions that operate on character strings are limited to processing “bytes”.

Explanation:

The CHARACTER_LENGTH, LENGTH, LOCATE, POSITION, and SUBSTRING functions include a parameter that allows you to specify a predefined set of string units. This means that the functions can process strings using the specified units instead of bytes or double bytes.

For example, the SYSIBM.LENGTH and SYSIBM.LOCATE functions process input as “characters” as compared to the SYSFUN version of LENGTH and LOCATE which use “bytes”. This may result in different behavior being exhibited when each handles illegal characters.

Resolution:

Take some care in the selection of the functions you use when you may encounter illegal characters in the data being processed. Different results could be expected based on the function used.

New scan default when creating an index:

Windows	UNIX
---------	------

Change:

When creating new primary keys, unique keys, or indexes (except extended index), ALLOW REVERSE SCANS is the default. Consequently, the access plan may change and query execution times may improve because the optimizer may be able to use the reverse index scan in some SQL statements. The exception is when working with extended index types. In the previous release, the default used to be DISALLOW REVERSE SCANS.

Explanation:

The new default allows the optimizer to consider both forward and reverse scans through the index.

Note: If you create two indexes on the same table, one specifying ascending order (ASC) and the other specifying descending order (DESC), and if you do not specify the DISALLOW REVERSE SCANS option in the CREATE INDEX statement, the two indexes will default to ALLOW REVERSE SCANS. As a result, the latter index will not be created and a duplicate index warning message is issued.

Resolution:

In prior versions, you may have created one forward scan index and one reverse scan index to speed up the application. Unfortunately, this requires the maintenance of two indexes. Now that reverse scans is enabled by default, the two indexes can be replaced with a single one that is enabled for reverse scans.

If you do not want to allow reverse scans on the indexes you are creating, then you must explicitly request that they be created with the DISALLOW REVERSE SCANS option.

New features enabled by default when a new database is created:

Windows	UNIX
---------	------

Change:

When a new database is created, self tuning memory, the Configuration Advisor, and automated RUNSTATS are enabled by default.

Explanation:

After creating a new database, you may see different query plans or workload behavior resulting from changes caused by the new defaults for these autonomic features. Existing applications or scripts that rely on previous DB2 default behavior and database configuration values, may see changes because some configuration values will have changed.

Resolution:

If you do not these features to be enabled by default, they can be disabled by actions involving the respective features:

- Configuration Advisor: Before creating the database, set the DB2_ENABLED_AUTOCONFIG_DEFAULT registry variable to "NO" using **db2set**.
- Self tuning memory: After the database is created, update the *self_tuning_mem* database configuration parameter by turning it "OFF".
- Automated RUNSTATS: After the database is created, update the *auto_runstats* database configuration parameter by turning it "OFF".

Disallowing multiple changes to the same buffer pool within one unit:

Windows	UNIX
---------	------

Change:

No longer allow multiple ALTER BUFFERPOOL statements within the same unit of work.

Explanation:

The addition of the self tuning memory manager in Version 9, increases the complexity in the actions that may be made on the characteristics of buffer pools. To limit this complexity, multiple alterations to the characteristics of the same buffer pool within a single unit of work is disallowed.

Resolution:

Attempting to make multiple ALTER BUFFERPOOL statements within the same unit of work will be disallowed.

Database security and tuning:

SET SESSION AUTHORIZATION requires SETSESSIONUSER privilege:

Windows	UNIX
---------	------

Change:

In DB2 Version 9, changing the session authorization ID to a new value using the SET SESSION AUTHORIZATION statement requires that the authorization ID of the SQL statement have the SETSESSIONUSER privilege. This privilege can be granted by a security administrator (SECADM) using the new GRANT SETSESSIONUSER statement.

Explanation:

In DB2 UDB Version 8, users with DBADM or SYSADM authority could assume different authorization IDs on the same connection using the SET SESSION AUTHORIZATION statement. In DB2 Version 9, the new SETSESSIONUSER privilege, which can only be granted by a security administrator (SECADM), is required to perform this task.

Resolution:

For backward compatibility, and to avoid loss of existing user privileges, any authorization ID that explicitly holds DBADM authority (as recorded in the SYSCAT.DBAUTH catalog view) is automatically granted the SETSESSIONUSER privilege upon migration to DB2 Version 9. A user who acquires DBADM authority after migration to DB2 Version 9 will not be able to change the session authorization ID unless they are explicitly granted the SETSESSIONUSER privilege.

TSM filtering changes involving management class:

Windows	UNIX
---------	------

Change:

Prior to DB2 Version 9, restore and log retrieval could search for objects based on a management class, if it was specified. Because the management class can change, filtering based on management class could produce incorrect results. Consequently, management class is no longer used as a basis for filtering.

Explanation:

Management class is a Tivoli Storage Manager (TSM) concept that helps with the management of objects according to defined storage policies. When a backup image, a load copy image, or a log file is written to TSM, a particular management class is associated with that object. After a log file is written to, or a backup image is stored, the management class may be changed through TSM.

Resolution:

Management class is no longer used as a basis for filtering.

Privileges and authorities changes to bring tables out of set integrity pending:

Windows	UNIX
---------	------

Change:

The SET INTEGRITY and REFRESH TABLE statements require specific authorities and privileges to work on the tables affected by these statements. The list of authorities and privileges that may be held by the authorization ID has changed from Version 8 to Version 9.

Resolution:

Bringing tables out of the set integrity pending state and performing the needed integrity processing requires specific authorities and privileges. The authorities and privileges held by the authorization ID of the statement must include at least one of the following:

- CONTROL privilege on:
 - The tables on which integrity processing is performed and, if exception tables are provided for one or more of those tables, INSERT privilege on the exception tables
 - All descendent foreign key tables, descendent immediate materialized query tables, and descendent immediate staging tables that will implicitly be placed in set integrity pending state by the statement
- LOAD authority (with conditions). The following conditions must all be met before LOAD authority can be considered as providing valid privileges:
 - The required integrity processing does not involve the following actions:
 - Refreshing a materialized query table
 - Propagation to a staging table
 - Updates of a generated or identity column
 - If exception tables are provided for one or more tables, the required access is granted for the duration of the integrity processing to the tables on which integrity processing is performed, and to the associated exception tables. That is:
 - SELECT and DELETE privilege on each table on which integrity processing is performed; and,
 - INSERT privilege on the exception tables
- SYSADM or DBADM authority

Index changes that could cause an error:

Windows	UNIX
---------	------

Change:

The maximum number of columns in an index has been increased from 16 to 64; and so has the maximum size of an index key which depends on the index page size. Sort heap overflows require a system temporary table space that has a large enough page size to be used by the sort.

Explanation:

In the previous release, the default temporary system table space using 4 KB pages may have been sufficient. However, in the current release, the key size plus record identifier plus the page header size may be greater than the 4 KB pages and result in error message SQL1584N.

Resolution:

Applications that may encounter this error message should be updated to detect it and to react in accordance with the message.

Extended storage is no longer supported:

Windows	UNIX
---------	------

Change:

Extended storage is no longer supported. With DB2 products moving to 64-bit environments, the need for extended storage is removed.

Explanation:

Extended storage acted as an extended look-aside buffer for the main buffer pools. It allowed for memory performance improvements which took advantage of computers with large amounts of main memory. For computers with 64-bit environments, extended storage and other similar methods are no longer needed.

Resolution:

Extended storage should no longer be used. If you are using Windows and you want to use more memory, you should consider moving to a 64-bit operating system. However, if you have to stay on a Windows 32-bit operating system, you can use Address Windowing Extensions (AWE) to overcome the 32-bit space limitation. AWE is controlled by the registry variable DB2_AWE.

Databases are created as automatic storage by default:

Windows	UNIX
---------	------

Change:

The CREATE DATABASE command and sqlcrea() API have been changed in Version 9. They will now create automatic storage-enabled databases by default. You will have to explicitly specify non-automatic storage to use the old behavior.

Symptom:

The SYSCATSPACE, TEMPSPACE1, and USERSPACE1 will all be created as automatic storage table spaces. This means that the database manager will manage the storage for these table spaces. Container operations are not valid against automatic storage table spaces. TableIDs (FIDs) will also change, although you may not be interested in these values. Redirected restore operations also act differently with automatic storage (that is, you redefine the storage paths instead of individual table space containers).

Explanation:

Automatic storage manages the containers for automatic storage table spaces, so there are certain operations which cannot be performed on those table spaces, such as container operations and redirected restore. Note that table spaces which are explicitly created as SMS or DMS will not be affected by this change. Databases migrated from previous releases are also unaffected. This may affect your scripts which rely on characteristics of the default table spaces.

Due to the change of table space type, disk requirements will increase. By default, non-temporary automatic storage table spaces increase by 32 MB at a time, so small databases may take more disk space. This space will be used as the database grows. Similarly, empty tables will consume more space. An empty table and index will consume 512 KB. This can be reduced by changing the extent size for the table space, either explicitly or by modifying the default extent size (DFT_EXTENT_SZ) in the database configuration. For small databases, an extent size of 4 is suggested. The extent size can only be chosen when the table space is created.

Resolution:

If necessary, you can create a non-automatic storage database by calling CREATE DATABASE with the AUTOMATIC STORAGE NO clause, or sqlecrea() with SQL_AUTOMATIC_STORAGE_NO. The application may also be updated to use the new table space type properly; for example, by redefining storage paths on restore instead of issuing SET CONTAINERS as a part of a redirected restore.

Utilities and tools:

Autoloader utility (db2atld) is no longer supported:

Windows	UNIX
---------	------

Change:

The Autoloader utility (db2atld) is no longer supported.

Explanation:

The load utility is now recommended for distributing and loading data within partitioned database environments.

Resolution:

Use the load utility for distributing and loading data within partitioned database environments.

Load from cursor:

Windows	UNIX
---------	------

Change:

You cannot use the distributed data files from a previous release when performing a load operation using the CURSOR file type and the PARTITION_ONLY partitioned database configuration load option in the current release.

Explanation:

The distributed data files are not compatible with the new DB2 server. The reverse is also true; that is, the distributed data files from the current release cannot be used when performing a load operation using the CURSOR file type and the PARTITION_ONLY partitioned database configuration load option.

Resolution:

When performing a load operation on a Version 9 DB2 server using the CURSOR file type and the PARTITION_ONLY partitioned database configuration load option, you must use the set of distributed data files created using DB2 Version 9.

Vendor load API (sqluvtd) is no longer valid:

Windows	UNIX
---------	------

Change:

The Vendor load API (sqluvtd) is no longer available for use.

Explanation:

The load utility is now recommended for distributing and loading data.

Resolution:

The load utility is the only supported bulk loader. The load utility can be run using the db2Load API.

Changes to db2batch:

Windows	UNIX
---------	------

Change:

Embedded dynamic SQL was the default mode, but this has been changed so that the command only runs in CLI mode. The output provided by the **db2batch** command is improved by including additional information such as time stamps and clearer messages. The output is also in a new format.

Explanation:

The parallel option (*-p*) on the **db2batch** command is no longer supported.

Resolution:

You can continue to use the *-cli* option for backward compatibility only, but it has no effect. You can change the default isolation level by specifying the *TxnIsolation* configuration keyword in the *db2cli.ini* file. The new *-iso* option is used to specify the isolation level.

You can no longer use the *-p* option on the **db2batch** command.

LOAD command set of distributed data files are changed:

Windows	UNIX
---------	------

Change:

The created set of distributed data files created using a LOAD command in a previous version using the CURSOR file type and the PARTITION_ONLY partitioned database configuration load option is specified cannot be used as input to the LOAD command in Version 9. That is, the distributed data files created previously are not compatible with a LOAD command using the CURSOR file type and the LOAD_ONLY partitioned database configuration load option.

Explanation:

The format of the distributed data files has changed in Version 9.

Resolution:

When creating a set of distributed data files, you should partition the data and load the data using the same version of the DB2 product.

Changes to the db2ckmig tool:

Windows	UNIX
---------	------

Change:

If an SQLCODE exists for messages returned by the **db2ckmig** tool, the **db2ckmig** log file now includes both the SQLCODE and the SQL message text.

Explanation:

In previous releases, the **db2ckmig** tool reported errors using message text from its own message file. However, in some cases, existing SQLCODEs also describe the errors. Having the SQLCODE in the **db2ckmig** log file, means that you can refer to the messages documentation for a more detailed explanation of the problem and possible user responses.

Resolution:

Any tools built on the exact message text of the **db2ckmig** log file might require changes to parse SQLCODEs.

REORGCHK command output changes:

Windows	UNIX
---------	------

Change:

The output generated as part of the REORGCHK command is changed for Version 9.

The SCHEMA and NAME columns are concatenated into one column (SCHEMA.NAME). In addition, the SCHEMA.NAME for each table and index is broken into two rows, one of the actual fully qualified name of the table, and one for the fully qualified name of each index on that table. The actual data for the remaining columns follows each index name.

Resolution:

You may have to take into account the changes made to the output from the REORGCHK command.

Changes to migration support tools and commands:

Windows	UNIX
---------	------

Change:

The database tools, utilities, and commands that are provided support migration from Version 8 but not from Version 7.

Explanation:

The breadth and complexity of the changes from Version 7 to Version 8; and from Version 8 to Version 9 make the migration path from Version 7 to Version 9 too difficult to be done with one set of migration tools and commands.

Resolution:

Use the migration information to plan for your migration to the current version and release. This may involve migrating to Version 8 before attempting to migrate to the current version and release.

New naming convention for backup images:

Windows	UNIX
---------	------

Change:

The naming convention for backup images stored on Windows operating systems has changed to match the naming convention used for all other operating systems.

Explanation:

File names for backup images created on disk will now consist of a concatenation of several elements, separated by periods:

`DB_alias.Type.Inst_name.NODEnnnn.CATNnnnn.timestamp.Seq_num`

Resolution:

Use this new naming convention for backup images.

Note: Backup images from earlier versions of the product that use the previous naming structure can still be restored on V9.1 DB2 database systems.

Change to db2look output:

Linux	UNIX
-------	------

Change:

In the output generated by the **db2look** command, the value displayed for the identity collating sequence is now IDENTITY.

Explanation:

In previous releases, the value BINARY was displayed for the identity collating sequence in the output generated by the **db2look** command and the GET DATABASE CONFIGURATION command. The collating sequence itself has not changed.

Changes to data movement utilities:

Linux	UNIX
-------	------

Change:

The following changes have been made to the load, import, and export utilities:

- When recreating tables using the IXF file format, if a feature cannot be recreated during the import process using the CREATE option, you will receive a warning during the export process and an error during the import process. In some cases, you can force the creation of tables from IXF files by specifying the file type modifier FORCECREATE. This new behavior only affects files exported using DB2 Version 9.1.
- The extension for an exported LOB file is now .lob. For example, filename.001.lob, filename.002.lob. The default name of the lob file is the input data file name. For example, <datafile>.001.lob, <datafile>.002.lob. If the input data file is generated in DB2 UDB Version 8, the DB2 Version 9.1 import utility can read it correctly.
- When moving LOB data, the default paths and the order in which the load, import, and export utilities search for these paths have changed.
- When exporting and importing LOB data, the LOBSINFILE keyword is specified automatically if you specify the LOBS TO or LOBFILE options in the EXPORT command, or the LOBS FROM option in the IMPORT command. In DB2 UDB Version 8, if the LOBSINFILE file type modifier was not specified, the LOBS TO, LOBS FROM, and LOBFILE options were ignored.

Changes to the db2mtrk command:

Linux	UNIX
-------	------

Change:

The -d option, which shows database level memory, is now supported on Windows platforms. The -i option, which shows instance level memory, no longer shows the database level memory.

Explanation:

Since the -d option is now available on Windows platforms, it should be used to display the database level memory. When the -i option is used, only the instance level memory is displayed.

Resolution:

On Windows platforms, use the -d option of the **db2mtrk** command to see the database level memory.

Changes to location of diagnostic messages for automatic maintenance:

Linux	UNIX
-------	------

Change:

The following changes have been made to the diagnostic level and location of messages related to automatic maintenance:

- A diagnostic record is written in the db2diag.log file whenever automatic maintenance health indicators are evaluated. If a maintenance operation occurs as a result of these evaluations, a diagnostic record is written in both the db2diag.log file and the notification log.
- The diagnostic records associated with automatic maintenance are classified as "info" records.
- These diagnostic records will only be written when the diagnostic level (diaglevel) or notification level (notifylevel) of the instance is set to a value of 4.

Explanation:

In DB2 Universal Database Version 8, whenever automatic maintenance health indicators were evaluated, a diagnostic record was written in db2diag.log file. Whenever a maintenance operation occurred as a result of these evaluations, another entry was written in the db2diag.log file. These diagnostic records were classified as "event" records and would appear when the diagnostic level of the instance (as specified in the diaglevel database manager configuration parameter) was set to values of 3 or 4.

Resolution:

To ensure that diagnostic records ("event" records) appear in the db2diag.log file and the notification log, set the diagnostic level (diaglevel) or notification level (notifylevel) of the instance to 4.

Restrictions for table space point in time rollforward operations:

Linux	UNIX
-------	------

Change:

For DB2 Version 9 clients, all table space rollforward recovery operation must be done to a point in time.

Resolution:

Ensure that all clients have been migrated to DB2 Version 9 and that you specify a point in time when initiating a rollforward operation.

Write-to-table event monitor changes:

Linux	UNIX
-------	------

Change:

In a partitioned database environment, a write-to-table event monitor will only be active on database partitions where the table space containing the event monitor table exists. When the target table space for an active event monitor does not exist on a particular database partition, the event monitor will be deactivated on that database partition, and an error is written to the db2diag.log file.

Explanation:

In earlier versions of DB2, the event monitor would be active and would appear as an active event monitor process on these database partitions but would not write any data.

Connectivity and coexistence:

Increased log, table space, and memory requirements:

Linux	UNIX
-------	------

Change:

Record identifier (RID) size was increased to support LARGE table spaces. The growth rate for log files and the size of log records also increases. Each RID now requires 8 bytes of memory in a single-partition environment and 16 bytes of memory in a partitioned database environment.

Explanation:

With the related change concerned with larger record identifiers (RID), there are increased requirements for logs, table spaces, and memory. Larger record identifiers (RID) allow more data pages per table object and more records per page. This increase in the number of pages and records also changes the required amount of memory and space used by log files and system temporary table spaces.

Resolution:

If the row size in your results sets is close to the maximum row length limit for your existing system temporary table space with the largest page size, you might need to create a system temporary table space with a larger page size. Another alternative is to reduce the length of the information retrieved by your query, or to split the query.

Databases require additional space:

Linux	UNIX
-------	------

Change:

To accommodate changes in the DB2 product requires that you allocate more space for database objects when compared to the same objects in a prior version.

Explanation:

Changes in this version of the DB2 product means that additional space is required for logs, table spaces, indexes, system catalog tables, and user table data.

Resolution:

Review the changes to the database objects so that you will understand the increased space requirements before creating those objects.

DB2 install images on Linux and UNIX have package format changes:

Linux	UNIX
-------	------

Change:

The DB2 install images on Linux and UNIX not longer use the operating system formats.

You can no longer use Linux and UNIX operating system utilities such as pkgadd, rpm, SMIT, or swinstall.

Explanation:

To enable you to install multiple DB2 copies on the same system, all DB2 install images for Linux and UNIX are compressed in a tar.gz format.

Resolution:

You should use the DB2 installation programs to ensure that your DB2 products are deployed and set up correctly. If you have scripts that you used in the past to install DB2 products using operating system commands, you must modify them to call DB2 installation programs (**db2setup** or **db2_install**) instead.

You can only use the **db2ls** command to query the installation of a DB2 product. If you used scripts containing operating system commands to query DB2 installation packages, you must modify them to use **db2ls**.

NetBIOS and SNA no longer supported:

Windows	UNIX
---------	------

Change:

NetBIOS and SNA are no longer supported as methods of communications between database systems.

Explanation:

NetBIOS and SNA are no longer supported.

Resolution:

Do not plan to use either NetBIOS or SNA as a future method of communication between and among database clients and servers. Remove the NetBIOS and SNA keywords from the DB2COMM registry variable to prevent the generation of an error when you start the instance. An error will also be returned when you use the CATALOG NETBIOS NODE, CATALOG APPC NODE, or CATALOG APPN NODE commands.

DB2 products no longer supported during installation:

Windows	UNIX
---------	------

Change:

The following products are no longer supported as installation options or as prerequisite components:

- DB2 Data Warehouse Center
- DB2 Data Warehouse Manager
- DB2 Information Catalog Center
- DB2 Data Links Manager
- DB2 Datajoiner

Explanation:**Resolution:**

If any of these products is installed on your system, they must be uninstalled before migrating your DB2 database system to Version 9. Instance migration will fail should any of these products be installed.

Also, any database objects created by these products (such as user-defined types, user-defined functions, and stored procedures) will remain in the database following the uninstall of the DB2 products. You should remove these objects from the databases before migration because they may cause the migration to fail.

Data Links Manager no longer supported:

Windows	UNIX
---------	------

Change:

DB2 Data Links Manager is no longer supported. This non-support includes several components of a Data Links server:

- Data Links File Manager (DLFM)
- Data Links Filesystem Filter (DLFF) controlling a Data Links File System (DLFS)
- DB2 Logging Manager

Explanation:

DB2 Data Links Manager is no longer supported.

Resolution:

Do not create any new database objects with the DATALINK data type or any new database objects that reference DATALINK columns.

VM/VSE objects no longer supported in the Control Center:

Windows	UNIX
---------	------

Change:

From the DB2 Control Center you can no longer connect or disconnect from VM/VSE databases. You can only display the cataloged VM and VSE databases. When adding an instance, the VM and VSE operating systems will no longer be available for selection.

Resolution:

Although you may display the cataloged VM and VSE databases, you will have to connect to them independent of the DB2 Control Center.

Maximum number of connections change in use:

Windows	UNIX
---------	------

Change:

There are two new agents requiring connection to the database at all times.

Explanation:

There are two new agents (db2stmm and db2taskd) requiring connection to the database at all times. These connection requirements will mean that in a tightly configured environment, two of the connections identified by the *max_connections* database manager configuration parameter will always be in use. As a result, you may run out of available connections.

Resolution:

If you are working in a tightly configured environment, you should consider increasing the *max_connections* database manager configuration parameter value by two.

Configuration parameters and registry variables:

Configuration parameters default value changes:

Windows	UNIX
---------	------

Change:

The default values for the following configuration parameters have changed between Version 8.2 and Version 9 of the DB2 database.

Table 41. Configuration parameters with changed default values

Parameter	V8.2 Default Value	V9.1 Default Value
app_ctl_heap_sz - Application control heap size configuration parameter	<p>Database server with local and remote clients: 128</p> <p>Database server with local clients:</p> <ul style="list-style-type: none"> • 64 (for non-UNIX platforms) • 128 (for UNIX-based platforms) <p>Partitioned database server with local and remote clients: 512</p>	<p>Database server with local and remote clients:</p> <ul style="list-style-type: none"> • 128 when INTRA_PARALLEL is not enabled • 512 when INTRA_PARALLEL is enabled <p>Database server with local clients:</p> <ul style="list-style-type: none"> • 64 (for non-UNIX platforms) when INTRA_PARALLEL is not enabled • 512 (for non-UNIX platforms) when INTRA_PARALLEL is enabled • 128 (for UNIX-based platforms) when INTRA_PARALLEL is not enabled • 512 (for UNIX-based platforms) when INTRA_PARALLEL is enabled <p>Partitioned database server with local and remote clients: 512</p>
auto_maint - Automatic maintenance configuration parameter	OFF	ON
auto_runstats (See auto_maint - Automatic maintenance configuration parameter for details.)	OFF	ON
auto_tbl_maint (See auto_maint - Automatic maintenance configuration parameter for details.)	OFF	ON
avg_appls - Average number of active applications configuration parameter	1	Automatic
database_memory - Database shared memory size configuration parameter	Automatic	<ul style="list-style-type: none"> • AIX and Windows: Automatic • Linux, HP-UX, Solaris Operating System: Computed
java_heap_sz - Maximum Java interpreter heap size configuration parameter	512	<ul style="list-style-type: none"> • 32-bit platforms: 512 • 64-bit platforms: 1024

Table 41. Configuration parameters with changed default values (continued)

Parameter	V8.2 Default Value	V9.1 Default Value
locklist - Maximum storage for lock list configuration parameter	<ul style="list-style-type: none"> • UNIX: 100 • Windows Database server with local and remote clients: 50 • Windows 64-bit Database server with local clients: 50 • Windows 32-bit Database server with local clients: 25 	Automatic
maxlocks - Maximum percent of lock list before escalation configuration parameter	<ul style="list-style-type: none"> • UNIX: 10 • Windows: 22 	Automatic
num_iocleaners - Number of asynchronous page cleaners configuration parameter	1	Automatic
num_ioservers - Number of I/O servers configuration parameter	3	Automatic
pckcachesz - Package cache size configuration parameter	-1	Automatic
sheapthres - Sort heap threshold configuration parameter	<ul style="list-style-type: none"> • UNIX 32-bit platforms: 20 000 • Windows 32-bit platforms: 10 000 • 64-bit platforms: 20 000 	0
sheapthres_shr - Sort heap threshold for shared sorts configuration parameter	<i>sheapthres</i>	Automatic
sortheap - Sort heap size configuration parameter	256	Automatic
userexit - User exit enable configuration parameter	No	Off

Configuration parameter changes:

Windows	UNIX
---------	------

Change:

The following configuration parameters are no longer valid in Version 9:

- *estore_seg_sz*
- *num_estore_segs*
- *min_priv_mem* (Windows only)
- *priv_mem_thresh* (Windows only)
- *fcm_num_rqb*
- *fcm_num_anchors*
- *fcm_num_connect*

The following changes have been made to configuration parameter content and meaning:

- *avg_appls*; the average number of active applications database configuration parameter has a new default. The default for the configuration parameter is set to one (the old default) unless an SAP environment is detected. If it is an SAP environment, the default is three active applications. In both cases, the defaults being set may not conflict with the setting for the *maxappls* database configuration parameter.

The new format presents the port number and IP address in a readable form that also accommodates the longer IPv6 addresses.

If you have scripts that parse output that contains the application ID, you will need to modify the parsing conditions to account for the new format. For example, you may parse the output from the LIST APPLICATIONS command.

- *database_memory*; the meaning of “AUTOMATIC” has changed. What was known as “AUTOMATIC” before Version 9 has been renamed to “COMPUTED”. To maintain the behavior for this configuration parameter as it was used before Version 9, set the *database_memory* configuration parameter to “COMPUTED”. By setting the *database_memory* configuration parameter to “AUTOMATIC” in Version 9, self tuning memory is enabled, and total database memory usage is automatically tuned.
- *sheapthres_shr*; the types of sorts using the shared memory has changed. Before Version 9, only sorts in a symmetric multi-processor (SMP) environment, or when using the concentrator, would use the shared memory. In Version 9, by setting the *sheapthres* instance configuration parameter to zero (0), and the *sheapthres_shr* database configuration parameter to a non-zero value, all sort memory consumers for the database use the database shared memory instead of private sort memory. Also, the default value for the *sheapthres_shr* database configuration parameter has changed from the value of *sheapthres* to 5000 4 KB pages.
- *dyn_query_mgmt*; during migration from Version 8 to Version 9, the default value for this configuration parameter is changed from “Enable” to “Disable”. Once migration is complete, and once Query Patroller is installed, then you must set this configuration parameter to “Enable” manually. You can set configuration parameter values using the UPDATE DATABASE CONFIGURATION command.
- *num_iocleaners* and *num_ioservers*; the default values set for these two configuration parameters is changed to “AUTOMATIC”. This means that the number of prefetchers and page cleaners started is based on environment characteristics such as the number of CPUs, the number of database partitions, and the parallelism settings of the table spaces in the database.

For existing databases, you can take advantage of this feature by setting the values of *num_iocleaners* and *num_ioservers* to “AUTOMATIC”.

The following changes have been made to registry variable content and meaning:

- *DB2_ALLOCATION_SIZE*; the default value is changed from 8388608 to 131072. This registry variable specifies the size of the memory allocation for buffer pools.
- *DB2_FORCE_FCM_BP*; the default value is changed from “No” to “Yes”. This registry variable specifies the memory allocation for FCM buffers.

Related concepts:

- “Migration overview for DB2 servers” in *Migration Guide*
- “Migration recommendations for DB2 servers” in *Migration Guide*
- “Migration essentials for DB2 clients” in *Migration Guide*
- “About the Release Notes” in *Release notes*
- “What’s new for V9.1: Administration changes summary” in *What’s New*

- “What's new for V9.1: Application development changes summary” in *What's New*
- “What's new for V9.1: Changes in existing functionality summary” in *What's New*
- “What's new for V9.1: Database setup changes summary” in *What's New*

Related reference:

- “Deprecated and discontinued features” on page 243

Version 8 incompatibilities with previous releases

System catalog information:

IMPLEMENTED column in catalog tables:

Windows	UNIX
---------	------

Change:

In previous versions, the column IMPLEMENTED in SYSIBM.SYSFUNCTIONS and SYSCAT.SYSFUNCTIONS had values of Y, M, H, and N. In Version 8, the values are Y and N.

Resolution:

Recode your applications to use only the values Y and N.

OBJCAT views renamed to SYSCAT views:

Windows	UNIX
---------	------

Change:

The following OBJCAT views have been renamed to SYSCAT views: TRANSFORMS, INDEXEXTENSIONS, INDEXEXTENSIONMETHODS, INDEXEXTENSIONDEP, INDEXEXTENSIONPARMS, PREDICATESPECS, INDEXEXPLOITRULES.

Resolution:

Recode your applications to use the SYSCAT views.

SYSCAT views are now read-only:

Windows	UNIX
---------	------

Change:

As of Version 8, the SYSCAT views are read-only.

Symptom:

An UPDATE or INSERT operation on a view in the SYSCAT schema now fails.

Explanation:

The SYSSTAT views are the recommended way to update the system catalog information. Some SYSCAT views were unintentionally updatable and this has now been fixed.

Resolution:

Change your applications to reference the updatable SYSSTAT views instead.

Application programming:**Audit context records statement size has grown:**

Windows	UNIX
---------	------

Change:

The statement limit has been raised to 2 MB.

Symptom:

The audit context record statement text is too large to fit into the table.

Explanation:

The existing tables used to record auditing context records only allow 32 KB for the statement text. The new statement limit is 2 MB. If you do not use long statement lengths, this will not affect you.

Resolution:

Create a new table to hold audit context records with a CLOB(2M) value for the statement text column. If desired, populate the new table with data from the old table, then drop the old table and use the new one. The new table may be renamed to the same name as the old table. Rebind any applications that use the new table.

Applications run multithreaded by default:

Windows	UNIX
---------	------

Change:

In Version 8, applications run in multithreaded mode by default. In previous versions, the default was to run applications in single-threaded mode. This change means that calls to the `sqlSetTypeCtx` API will have no effect.

The Version 8 multithreaded mode is equivalent to calling the `sqlSetTypeCtx` API with the `SQL_CTX_MULTI_MANUAL` option in a pre-Version 8 application. A Version 7 client can still run an application in single-threaded mode.

Explanation:

In Version 7, if you wanted to run an application in multithreaded mode, you had to call context APIs and manage the contexts. In Version 8, this is not necessary since DB2 Database for Linux, UNIX, and Windows will manage contexts

internally. However, in Version 8 you are still able to manage contexts for applications if you want to, through external context APIs.

SQL0818N error not returned when using VERSION option:

Windows	UNIX
---------	------

Change:

If you use the new VERSION option on the PRECOMPILE, BIND, REBIND, and DROP PACKAGE commands, requests to execute may now return an SQL0805N error instead of an SQL0818N error.

Symptom:

Applications coded to react to an SQL0818N error may not behave as before.

Resolution:

Recode your applications to react to both SQL0805N and SQL0818N errors.

SQL0306N error not returned to the precompiler when a host variable is not defined:

Windows	UNIX
---------	------

Change:

If a host variable is not declared in the BEGIN DECLARE section and is used in the EXEC SQL section, SQL0306N will not be returned by the precompiler. If the variable is declared elsewhere in the application, application runtime will return SQL0804N. If the variable is not declared anywhere in the application, the compiler will return an error at compilation time.

Symptom:

Applications coded to react to an SQL0306N error at precompilation time may not behave as before.

Resolution:

Host variables should be declared in the BEGIN DECLARE section. If host variables are declared in a section other than the BEGIN DECLARE section, you should recode your application to handle SQL0804 return codes.

Data types not supported for use with scrollable cursors:

Windows	UNIX
---------	------

Change:

Scrollable cursors using LONG VARCHAR, LONG VARGRAPHIC, DATALINK and LOB types, distinct types on any of these types, or structured types will not be supported in Version 8. Any of these data types supported for Version 7 scrollable cursors will no longer be supported.

Symptom:

If any columns with these data types are specified in the select list of a scrollable cursor, SQL0270N Reason Code 53 is returned.

Resolution:

Modify the select-list of the scrollable cursor so it does not include a column with any of these types.

Euro version of code page conversion tables:

Windows	UNIX
---------	------

Change:

The Version 8 code page conversion tables, which provide support for the euro symbol, are slightly different from the conversion tables supplied with previous versions of DB2.

Resolution:

If you want to use the pre-Version 8 code page conversion tables, they are provided in the directory `sql11ib/conv/v7`.

Switching between a LOB locator and a LOB value:

Windows	UNIX
---------	------

Change:

The ability to switch between a large object (LOB) locator and a LOB value has been changed during bindout on a cursor statement. When an application is bound with SQLRULES DB2 (the default behavior), the user will not be able to switch between LOB locators and LOB values.

Resolution:

If you want to switch between a LOB locator and a LOB value during bindout of a cursor statement, precompile your application with SQLRULES STD.

Uncommitted units of work on UNIX platforms:

	UNIX
--	------

Change:

In Version 8, all application terminations implicitly roll back outstanding units of work. Windows-based applications will not change as they already perform an

implicit ROLLBACK for normal or abnormal application termination. Prior to version 8, UNIX applications that did not use either explicit or implicit context support would commit an outstanding unit of work if the application terminated normally without directly invoking either a CONNECT RESET, COMMIT, or ROLLBACK statement. CLI, ODBC, and Java-based applications (implicit context support) and applications that would explicitly create application contexts would always roll back any outstanding unit of work if the application terminated. Abnormal application termination would also lead to an implicit ROLLBACK for the outstanding unit of work.

Resolution:

In order to ensure that transactions are committed, the application should perform either an explicit COMMIT or a CONNECT RESET before terminating.

Change to savepoint naming:

Windows	UNIX
---------	------

Change:

Savepoint names can no longer start with "SYS".

Symptom:

Creating a savepoint with a name that starts with "SYS" will fail with error SQL0707N.

Explanation:

Savepoint names that start with "SYS" are reserved for use by the system.

Resolution:

Rename any savepoints that start with "SYS" to another name that does not start with "SYS".

Code page conversion errors and byte substitution:

Windows	UNIX
---------	------

Change:

Character data in input host variables will be converted to the database code page, when necessary, before being used in the SQL statement where the host variable appears. During code page conversion, data expansion may occur. Previously, when code page conversion was detected for data in a host variable, the actual length assumed for the host variable was increased to handle the expansion. This assumed increase in length is no longer performed, to mitigate the impact of the change of the data type length on other SQL operations.

Note: None of this applies to host variables that are used in the context of FOR BIT DATA. The data in these host variables will not be converted before being used as for bit data.

Symptom:

If the host variable is not large enough to hold the expanded length after code page conversion, an error is returned (SQLSTATE 22001, SQLCODE -302).

Explanation:

Since expansion or contraction can occur during code page conversion, operations that depend on the length of the data in the host variable can produce different results or an error situation.

Resolution:

Alternatives that can be considered include:

- Coding the application to handle the possibility of code page conversion causing the length of the data to change by increasing the length of character host variables
- Changing the data to avoid characters that cause expansion
- Changing the application code page to match the database code page so that code page conversion does not occur.

Code page conversion for host variables:

Windows	UNIX
---------	------

Change:

Code page conversion, when necessary, will now be performed during the bind in phase.

Symptom:

Different results.

Explanation:

Now that code page conversion, when necessary, will always be done for host variables, predicate evaluation will always occur in the database code page and not the application code page. For example,

```
SELECT * FROM table WHERE :hv1 > :hv2
```

will be done using the database code page rather than the application code page. The collation used continues to be the database collation.

Resolution:

Verify that the results in previous versions were indeed the desired results. If they were, then change the predicate to produce the desired result given that the database collation and code page are used. Alternatively, change the application code page or the database code page.

Expansion and contraction of data in host variables:

Windows	UNIX
---------	------

Change:

Code page conversion, when necessary, will now be performed during a bind operation.

Symptom:

Data from host variables have a different length.

Explanation:

Since expansion or contraction can occur during code page conversion, operations that depend on the length of the data in the host variable can produce different results or an error situation.

Resolution:

Change the data, the application code page or the database code page so that code page conversion does not produce changes in length of the converted data, or code the application to handle the possibility of code page conversion causing the length of the data to change.

Length of host variables after code page conversion:

Windows	UNIX
---------	------

Change:

Code page conversion will no longer cause result length to increase for host variables or parameter markers due to expansion.

Symptom:

Data truncation errors.

Explanation:

The length of the character data type determined for the untyped parameter marker is no longer increased to account for potential expansion from code page conversion. The result length will be shorter for operations that determine result length using the length of the untyped parameter marker. For example, given that C1 is a CHAR(10) column:

```
VALUES CONCAT (?, C1)
```

no longer has a result data type and length of CHAR(40) for a database where 3 times expansion is possible when converting from the application code page to the database code page, but will have a result data type and length of CHAR(20).

Resolution:

Use a CAST to give the untyped parameter marker the type desired or change the operand that determines the type of the untyped parameter marker to a data type or length that would accommodate the expansion of the data due to code page conversion.

Change to output of DESCRIBE statement:

Windows	UNIX
---------	------

Change:

Code page conversion will no longer cause result length to increase for host variables or parameter markers due to expansion.

Symptom:

Output from DESCRIBE statement changes.

Explanation:

Since the result length is not increased due to potential expansion on code page conversion, the output of a DESCRIBE statement that describes such a result length will now be different.

Resolution:

If necessary, change the application to handle the new values returned from the DESCRIBE statement.

Error when using SUBSTR function with host variables:

Windows	UNIX
---------	------

Change:

Code page conversion will no longer cause result length to increase for host variables or parameter markers due to expansion.

Symptom:

Error SQL0138N from SUBSTR.

Explanation:

Potential expansion due to code page conversion was taken into account by increasing the length set aside for the host variable. This allowed, for example, SUBSTR (:hv,19,1) to work successfully for a host variable with a length of 10. This will no longer work.

Resolution:

Increase the length of the host variable to account for the length of the converted data or change the SUBSTR invocation to specify positions within the length of the host variable.

Non-thread safe libraries are no longer supported on Solaris:

	UNIX
--	------

Change:

The non-thread safe library `libdb2_noth.so` is no longer available.

Symptom:

Tools or applications that require `libdb2_noth.so` will not work.

Explanation:

Since support for the obsolete non-thread safe libraries is no longer required, the `libdb2_noth.so` library is not included with IBM DB2 Version 9.1 for Solaris.

Resolution:

Change the tool or application to use the thread-safe `libdb2.so` library instead. Re-link your applications with the `-mt` parameter.

Importing or exporting a DBCLOB when connected to a Unicode database:

Windows	UNIX
---------	------

Change:

Prior to Version 8, if you exported data that contained a DBCLOB from a Unicode database (UTF-8), and used the LOBSINFILE file type modifier, the DBCLOB would be exported in code page 1200 (the Unicode graphic code page). If you imported data that contained a DBCLOB, and used the LOBSINFILE file type modifier, the DBCLOB would be imported in code page 1200 (the Unicode graphic code page). This behavior is maintained in Version 8 if you set the `DB2GRAPHICUNICODESERVER` registry variable to ON.

In Version 8, the default setting of the `DB2GRAPHICUNICODESERVER` registry variable is OFF. If you export data containing a DBCLOB and using the LOBSINFILE file type modifier, the DBCLOB will be exported in the application's graphic code page. If you import data containing a DBCLOB and using the LOBSINFILE file type modifier, the DBCLOB will be imported in the application's graphic code page. If your application code page is IBM-eucJP (954) or IBM-eucTW (964), and you export data containing a DBCLOB and using the LOBSINFILE file type modifier, the DBCLOB will be exported in the application's character code page. If you import data containing a DBCLOB and using the LOBSINFILE file type modifier, the DBCLOB will be imported in the application's character code page.

Symptom:

When importing data with the LOBSINFILE file type modifier into a Unicode database, the character data will be converted correctly, but the DBCLOB data is corrupted.

Resolution:

If you are moving data between a Version 8 database and an earlier database, set the `DB2GRAPHICUNICODESERVER` registry variable to ON to retain the previous behavior.

SQL:

DROPPED TABLE RECOVERY default changed for the CREATE TABLESPACE statement:

Windows	UNIX
---------	------

Change:

The DROPPED TABLE RECOVERY default changed for the CREATE TABLESPACE statement from OFF to ON.

Symptom:

Forward recovery performance may be affected by this change in defaults. The performance impact may be noticeable when there are many drop table operations to recover; or, when the history table is large. In the later case, during forward recovery when redoing SQLP_PRT_DROP_TABLE_RECOVERY pending list action, the database manager needs to read and update the history file for the related dropped recovery entry. If the history file is large, this action could take time while searching the file for the correct entry and then updating the entry.

Resolution:

If you will be dropping many tables, and if you are either using circular logging or do not wish to recover those tables, then you should consider disabling this feature. To disable the feature for new table spaces, you can use the CREATE TABLESPACE statement and explicitly use the DROPPED TABLE RECOVERY clause with the value set to OFF. For existing table spaces, you can use the ALTER TABLESPACE statement and explicitly use the DROPPED TABLE RECOVERY clause to change the value from the default (ON) to OFF. Dropped tables will no longer be recovered as part of the forward recovery process.

Identical specific names not permitted for functions and procedures:

Windows	UNIX
---------	------

Change:

The name space for SPECIFICNAME has been unified. Previous versions of DB2 would allow a function and a procedure to have the same specific name, but Version 8 does not allow this.

Symptom:

If you are migrating a database to Version 8, the db2ckmig utility will check for functions and procedures with the same specific name. If duplicate names are encountered during migration, the migration will fail.

Resolution:

Drop the procedure and recreate it with a different specific name.

EXECUTE privilege on functions and procedures:

Windows	UNIX
---------	------

Change:

Previously, a user only had to create a routine for others to be able to use it. Now after creating a routine, a user has to GRANT EXECUTE on it first before others can use it.

In previous versions, there were no authorization checks on procedures, but the invoker had to have EXECUTE privilege on any package invoked from the procedure. For an embedded application precompiled with CALL_RESOLUTION IMMEDIATE in Version 8, and for a CLI cataloged procedure, the invoker has to have EXECUTE privilege on the procedure and only the definer of the procedure has to have EXECUTE privilege on any packages.

Symptom:

1. An application may not work correctly.
2. An existing procedure that is made up of multiple packages, and for which the definer of the procedure does not have access to all the packages, will not work correctly.

Resolution:

1. Issue the required GRANT EXECUTE statements. If all the routines are in a single schema, the privileges for each type of routine can be granted with a single statement, for example:

```
GRANT EXECUTE ON FUNCTION schema1.* TO PUBLIC
```

2. If one package is usable by everyone but another package is restricted to a few privileged users, a stored procedure that uses both packages will watch for an authority error when it tries to access the second package. If it sees the authority error, it knows that the user is not a privileged user and the procedure bypasses part of its logic.

You can resolve this in several ways:

- a. When precompiling a program, CALL_RESOLUTION DEFERRED should be set to indicate that the program will be executed as an invocation of the deprecated sqlproc() API when the precompiler fails to resolve a procedure on a CALL statement.
- b. The CLI keyword UseOldStpCall can be added to the db2cli.ini file to control the way in which procedures are invoked. It can have two values: A value of 0 means procedures will not be invoked using the old call method, while a value of 1 means procedures will be invoked using the old call method.
- c. Grant EXECUTE privilege to everyone who executes the package.

Adding a foreign key constraint to a table:

Windows	UNIX
---------	------

Change:

In previous versions, if you created a foreign key constraint that referenced a table in check pending state, the dependent table would also be put into check pending state. In Version 8, if you create a foreign key constraint that references a table in check pending state, there are two possible results:

1. If the foreign key constraint is added upon creation of the dependent table, the creation of the table and the addition of the constraint will be successful because the table will be created empty, and therefore no rows will violate the constraint.
2. If a foreign key is added to an existing table, you will receive error SQL0668N.

Resolution:

Use the SET INTEGRITY ... IMMEDIATE CHECKED statement to turn on integrity checking for the table that is in check pending state, before adding the foreign key that references the table.

Change to SET INTEGRITY ... IMMEDIATE CHECKED:

Windows	UNIX
---------	------

Change:

In previous releases, a table that had the SET INTEGRITY ... UNCHECKED statement issued on it (i.e. with some 'U' bytes in the const_checked column of SYSCAT.TABLES) would by default be fully processed upon the next SET INTEGRITY ... IMMEDIATE CHECKED statement, meaning all records would be checked for constraint violations. You had to explicitly specify INCREMENTAL to avoid full processing.

In Version 8, when the SET INTEGRITY ... IMMEDIATE CHECKED statement is issued, the default is to leave the unchecked data alone (i.e. keeping the 'U' bytes) by doing only incremental processing. (A warning will be returned that old data remains unverified.)

Explanation:

This change is made to avoid having the default behavior be a constraint check of all records, which usually consumes more resources.

Resolution:

You will have to explicitly specify NOT INCREMENTAL to force full processing.

Decimal separator for CHAR function:

Windows	UNIX
---------	------

Change:

Dynamic applications that run on servers with a locale that uses the comma as the decimal separator and include unqualified invocations of the CHAR function with an argument of type REAL or DOUBLE, will return a period as the separator

character in the result of the CHAR(double) function. This incompatibility will also be visible when objects like views and triggers are re-created in Version 8 or when static packages are explicitly rebound.

Explanation:

This is a result of resolving to the new SYSIBM.CHAR(double) function signature instead of the SYSFUN.CHAR(double) signature.

Resolution:

To maintain the behavior from earlier versions of DB2, the application will need to explicitly invoke the function with SYSFUN.CHAR instead of allowing function resolution to select the SYSIBM.CHAR signature.

Changes to CALL statement:

Windows	UNIX
---------	------

Change:

In Version 8, an application precompiled with CALL_RESOLUTION IMMEDIATE and a CLI cataloged procedure have several key differences compared to previous versions:

- Host variable support has been replaced by support for dynamic CALL.
- Support for compilation of applications that call uncataloged stored procedures has been removed. Uncataloged stored procedure support will be removed entirely in a future version of DB2.
- Variable argument list stored procedure support has been deprecated.
- There are different rules for loading the stored procedure library.

Resolution:

The CALL statement as supported prior to Version 8 will continue to be available and can be accessed using the CALL_RESOLUTION DEFERRED option on the PRECOMPILE PROGRAM command.

Existing applications (built prior to Version 8) will continue to work. If applications are re-precompiled without the CALL_RESOLUTION DEFERRED option, then source code changes may be necessary.

Support for the CALL_RESOLUTION DEFERRED statement will be removed in a future version.

Output from UDFs returning fixed-length strings:

Windows	UNIX
---------	------

Change:

A UDF (scalar or table function) can be defined to return a fixed-length string (CHAR(n) or GRAPHIC(n)). In previous versions, if the returned value contains an imbedded null character, the result would simply be n bytes (or 2n bytes for GRAPHIC data types) including the null character and any bytes to the right of

the null character. In Version 8, DB2 UDB looks for the null character and returns blanks from that point (the null character) to the end of the value.

Resolution:

If you want to continue the pre-Version 8 behavior, change the definition of the returned value from CHAR(n) to CHAR(n) FOR BIT DATA. There is no method to continue the pre-Version 8 behavior for GRAPHIC data.

Change in database connection behavior:

Windows	UNIX
---------	------

Change:

In Version 7, if you use embedded SQL to connect to a database, and then attempt a connection to a non-existent database, the attempt to connect to the non-existent database will fail with SQL1013N. The connection to the first database still exists. In Version 8, the attempt to connect to the non-existent database will result in a disconnection from the first database. This will result in the application being left with no connection.

Resolution:

Code your embedded SQL to reconnect to the initial database following an unsuccessful attempt to connect to another database.

Revoking CONTROL on packages:

Windows	UNIX
---------	------

Change:

A user can grant privileges on a package using the CONTROL privilege. In DB2 UDB Version 8, the WITH GRANT OPTION provides a mechanism to determine a user's authorization to grant privileges on packages to other users. This mechanism is used in place of CONTROL to determine whether a user may grant privileges to others. When CONTROL is revoked, users will continue to be able to grant privileges to others.

Symptom:

A user can still grant privileges on a package, following the revocation of CONTROL privilege.

Resolution:

If a user should no longer be authorized to grant privileges on packages to others, revoke all privileges on the package and grant only those required.

Error when casting a FOR BIT DATA character string to a CLOB:

Windows	UNIX
---------	------

Change:

Casting a character string defined as FOR BIT DATA to a CLOB (using the CAST specification or the CLOB function) now returns an error (SQLSTATE 42846).

Symptom:

Casting to a CLOB now returns an error where previously it did not.

Explanation:

FOR BIT DATA is not supported for the CLOB data type. The result of using the CAST specification or the CLOB function when a FOR BIT DATA string is given as an argument is not defined. This situation is now caught as an error.

Resolution:

Change the argument to the CAST specification or the CLOB function so that it is not a FOR BIT DATA string. This can be done by using the CAST specification to cast the FOR BIT DATA string to a FOR SBCS DATA string or a FOR MIXED DATA string. For example, if C1FBD is a VARCHAR(20) column declared as FOR BIT DATA, in a non-DBCS database, the following would be a valid argument to the CLOB function:

```
CAST (C1FBD AS VARCHAR(20) FOR SBCS DATA)
```

Output from CHR function:

Windows	UNIX
---------	------

Change:

CHR(0) returns a blank (X'20') instead of the character with code point X'00'.

Symptom:

Output from the CHR function with X'00' as the argument returns different results.

Explanation:

String handling when invoking and returning from user-defined functions interprets X'00' as end of string.

Resolution:

Change the application code to handle the new output value. Alternatively, define a user-defined function that returns CHAR(1) FOR BIT DATA which is sourced from the definition of the SYSFUN CHR function, and place this function before SYSFUN on the SQL path.

For example, to find the source definition for SYSFUN.CHRC located in column IMPLEMENTATION:

```
SELECT IMPLEMENTATION, ROUTINENAME FROM SYSIBM.SYSROUTINES
      WHERE ROUTINENAME LIKE '%CHR%';
```

```
IMPLEMENTATION      ROUTINENAME
-----
db2clifn!CLI_udfCHAR  CHR
```

Then, you could create a new user-defined function from the definition *db2clifn!CLI_udfCHAR* returned above.

```
CREATE FUNCTION DBS.CHR(INTEGER) RETURNS CHARACTER(1) FOR BIT DATA NOT FENCED
      LANGUAGE C PARAMETER STYLE DB2SQL NO DBINFO EXTERNAL NAME 'db2clifn!CLI_udfCHAR'
```

TABLE_NAME and TABLE_SCHEMA functions cannot be used in generated columns or check constraints:

Windows	UNIX
---------	------

Change:

The definitions for the TABLE_NAME and TABLE_SCHEMA functions have been corrected, and can now not be used in generated columns or check constraints.

Symptom:

The bind will fail with an SQLCODE -548/SQLSTATE 42621 stating that TABLE_NAME or TABLE_SCHEMA is invalid in the context of a check constraint.

Explanation:

The TABLE_NAME and TABLE_SCHEMA functions retrieve data from catalog views. They are of the class READS SQL DATA; functions of class READS SQL DATA are not permitted in GENERATED COLUMN expressions and check constraints, since DB2 cannot enforce the correctness of the constraint over time.

Resolution:

Update any columns that contain generated column expressions and check constraints to remove the use of TABLE_NAME and TABLE_SCHEMA. To alter a generated column, use the ALTER TABLE statement to SET a new expression. To remove a check constraint, use the ALTER TABLE statement with the DROP CONSTRAINT clause. This will allow you to BIND and continue accessing the tables that contain the affected columns.

Database security and tuning:

Authority for CREATE FUNCTION, CREATE METHOD and CREATE PROCEDURE statements:

Windows	UNIX
---------	------

Change:

The CREATE_EXTERNAL_ROUTINE authority is introduced in Version 8.

Symptom:

CREATE FUNCTION, CREATE METHOD and CREATE PROCEDURE statements with the EXTERNAL option may fail.

Resolution:

Grant CREATE_EXTERNAL_ROUTINE authority to users who issue CREATE FUNCTION, CREATE METHOD, and CREATE PROCEDURE statements with the EXTERNAL option.

Utilities and tools:

Changes with the DB2 Administration Server (DAS):

Windows	UNIX
---------	------

Symptom:

When migrating, errors occurred and were recorded in the db2diag.log at DB2 startup.

Explanation:

In Version 7, the DB2 Administration Server (DAS) was its own instance. In Version 8, the DAS is no longer an instance but is a control point used to assist with DB2 server tasks.

Resolution:

If you have attempted to migrate from Version 7 to Version 8 and encountered problems relating to the DAS, then use the “db2admin drop” command to stop and drop the DAS. Following the removal, then use the “db2admin create” command to create the DAS.

Changes when monitoring performance using the Control Center:

Windows	UNIX
---------	------

Symptom:

When looking within the Control Center, you do not find any references to the performance monitor.

Explanation:

The performance monitor capability of the Control Center has been removed.

Resolution:

When working with IBM DB2 Version 9.1 for Windows, there are tools that can be used to monitor performance:

- **DB2 Performance Expert**

The separately purchased DB2 Performance Expert for Multiplatforms, Version 1.1 consolidates, reports, analyzes and recommends self-managing and resource tuning changes based on DB2 performance-related information.

- **DB2 Health Center**

The functions of the Health Center provide you with different methods to work with performance-related information. These functions somewhat replace the performance monitor capability of the Control Center.

- **Windows Performance Monitor**

The Windows Performance Monitor enables you to monitor both database and system performance, retrieving information from any of the performance data providers registered with the system. Windows also provides performance information data on all aspects of machine operation including:

- CPU usage
- Memory utilization
- Disk activity
- Network activity

Running online utilities at the same time:

Windows	UNIX
---------	------

Symptom:

When online utilities are used at the same time, the utilities may take a long time to complete.

Explanation:

The locks required by one utility affect the progress of the other utilities running at the same time.

Resolution:

When there is a potential for conflict between the locking requirements of utilities that are being run at the same time, you should consider altering your scheduling for the utilities you wish to run. The utilities (like online backup table space, load table, or inplace reorganization of tables) use locking mechanisms to prevent conflicts between the utilities. The utilities use table locks, table space locks, and table space states at different times to control what needs to be done in the database. When locks are held by a utility, the other utilities requesting similar or related locks must wait until the locks are released.

For example, the last phase of an inplace table reorganization cannot start while an online backup is running that includes the table being reorganized. You can pause the reorganization request if you require the backup to complete.

In another example, the online load utility will not work with another online load request on the same table. If different tables are being loaded, then the load requests will not block each other.

Changes to db2move summary output:

Windows	UNIX
---------	------

Change:

In Version 8.2, the summary output generated by **db2move** is improved by being made more descriptive. However, the change in the summary output may cause a script written to analyze the old output to fail.

Symptom:

A script written to analyze the old output generated by **db2move** fails.

Explanation:

The summary output generated by **db2move** is improved.

When **db2move** is run with the "IMPORT" option, the old output appeared as:

```
IMPORT: -Rows read:      5; -Rows committed:      5; Table "DSCIARA2"."T20"
```

The new output appears as:

```
* IMPORT: table "DSCIARA2"."T20"  
-Rows read:      5  
-Inserted:      4  
-Rejected:      1  
-Committed:     5
```

When **db2move** is run with the "LOAD" option, the old output appeared as:

```
* LOAD: table "DSCIARA2"."T20"  
-Rows read:      5; -Loaded:      4; -Rejected  1 -Deleted  0 -Committed  5
```

The new output appears as:

```
* IMPORT: table "DSCIARA2"."T20"  
-Rows read:      5  
-Loaded:      4  
-Rejected:      1  
-Deleted:      0  
-Committed:     5
```

Resolution:

Your script used to analyze the **db2move** output will need to be modified to account for the changes in the layout and content.

Changes to the explain facility tables:

Windows	UNIX
---------	------

Change:

In Version 8, there are changes to the existing explain facility tables including two new tables: ADVISE_MQT and ADVISE_PARTITION.

Symptom:

The DB2 Design Advisor, when asked to make recommendations for materialized query tables (MQTs), or for database partitions, will return error messages if the explain tables have not been created.

Explanation:

The new tables ADVISE_MQT and ADVISE_PARTITION have not been created.

Resolution:

Use the **db2exmig** command to move the Version 7 and Version 8.1 explain tables to Version 8.2. This command has the necessary EXPLAIN DLL to create all of the needed explain facility tables.

Changes to the db2diag.log message format:

Windows	UNIX
---------	------

Change:

In Version 8, the db2diag.log message format is changed.

Symptom:

You will notice that the format has changed when reviewing the db2diag.log messages. The changes include the following examples: each message will have a diagnostic log record header, record fields will be preceded by the field name and column, and message and data portions of the logging record will be clearly marked. All of the changes to the format will make the logging record easier to use and to understand.

Explanation:

The DB2 diagnostic logs are being reworked. The db2diag.log file will be parsable.

Downlevel CREATE DATABASE and DROP DATABASE not supported:

Windows	UNIX
---------	------

Change:

In Version 8, the CREATE DATABASE and DROP DATABASE commands are not supported from downlevel clients or to downlevel servers.

Symptom:

You will receive error SQL0901N when you issue one of these commands.

Explanation:

The CREATE DATABASE and DROP DATABASE commands are both only supported from Version 8 clients to Version 8 servers. You cannot issue these commands from a Version 6 or Version 7 client to a Version 8 server. You cannot issue these commands from a Version 8 client to a Version 7 server.

Resolution:

Create or drop a Version 8 database from a Version 8 client. Create or drop a Version 7 database from a Version 6 or Version 7 client.

Mode change to tables after a load:

Windows	UNIX
---------	------

Change:

In previous versions, a table that has been loaded with the INSERT option and has immediate materialized query tables (also known as summary tables) would be in Normal (Full Access) state after a subsequent SET INTEGRITY IMMEDIATE CHECKED statement on it. In Version 8, the table will be in No Data Movement mode after the SET INTEGRITY IMMEDIATE CHECKED statement.

Explanation:

Access to a table in No Data Movement mode is very similar to a table in Normal (Full Access) mode, except for some statements and utilities that involve data movement within the table itself.

Resolution:

You can force the base table that has been loaded and has dependent immediate summary tables to bypass the No Data Movement mode and to go directly into Full Access mode by issuing a SET INTEGRITY ... IMMEDIATE CHECKED FULL ACCESS statement on the base table. However, use of this option is not recommended as it will force a full refresh of the dependent immediate materialized query tables (also known as summary tables).

Load utility in insert or replace mode:

Windows	UNIX
---------	------

Change:

In previous versions, when using the load utility in insert or replace mode, the default option was CASCADE IMMEDIATE when integrity checking was turned off; when the table was put into check pending state, all of its dependent foreign key tables and dependent materialized query tables (also known as summary tables) were also immediately put into check pending state.

For Version 8, when using the load utility in insert or replace mode, the default is CASCADE DEFERRED when integrity checking has been turned off.

Resolution:

You can put dependent foreign key tables and dependent materialized query tables into check pending state along with their parent tables by using the CHECK PENDING CASCADE IMMEDIATE option of the LOAD command.

DB2_LIKE_VARCHAR does not control collection of sub-element statistics:

Windows	UNIX
---------	------

Change:

In Version 7, the DB2_LIKE_VARCHAR registry variable controlled collection of sub-element statistics as well as the use of these statistics. In Version 8, DB2_LIKE_VARCHAR does not control collection of sub-element statistics; instead, collection of sub-element statistics is controlled by the LIKE STATISTICS option of the RUNSTATS command or the DB2RUNSTATS_COLUMN_LIKE_STATS value of the iColumnflags parameter of the db2Runstats API.

Symptom:

After invoking the RUNSTATS command or calling the db2Runstats API, sub-element statistics are set to -1 (the default) in the system catalog; this can be observed with a query like the following:

```
SELECT SUBSTR(TABSCHEMA,1,18), SUBSTR(TABNAME,1,18),
       SUBSTR(COLNAME,1,18), COLCARD, AVGCOLLEN, SUB_COUNT, SUB_DELIM_LENGTH
FROM SYSSTAT.COLUMNS
WHERE COLNAME IN ('P_TYPE', 'P_NAME')
ORDER BY 1,2,3
```

(Replace P_TYPE and P_NAME with the appropriate column names.)

If the result for a column has a non-negative value for COLCARD and AVGCOLLEN but a value of -1 for SUB_COUNT and SUB_DELIM_LENGTH, this indicates that basic statistics have been gathered for the column, but sub-element statistics have not been gathered.

Resolution:

If you specified DB2_LIKE_VARCHAR=?,Y (where ? is any value) in Version 7, then you should specify the LIKE STATISTICS option on the RUNSTATS command or DB2RUNSTATS_COLUMN_LIKE_STATS on the db2Runstats API to collect these statistics for appropriate columns.

Connectivity and coexistence:

Down level server support:

Windows	UNIX
---------	------

Change:

As you move your environment from Version 7 to Version 8, if you are in a situation where you migrate your client machines to Version 8 before you migrate all of your servers to Version 8, there are several restrictions and limitations. These restrictions and limitations are not associated with DB2 Connect; nor with zSeries, OS/390, or iSeries database servers.

Resolution:

For Version 8 clients to work with Version 7 servers, you need to configure/enable the use of DRDA Application Server capability on the Version 7 server. For information on how to do this, refer to the Version 7 *Installation and Configuration Supplement*.

To avoid the known restrictions and limitations, you should migrate all of your servers to Version 8 before you migrate any of your client machines to Version 8. If

this is not possible, then you should know that when accessing Version 7 servers from Version 8 clients, there is no support available for:

- Some data types:
 - Large object (LOB) data types.
 - User-defined distinct types (UDTs).
 - DATALINK data types.
- Some security capabilities:
 - Authentication type SERVER_ENCRYPT.
 - Changing passwords. You are not able to change passwords on the DB2 UDB Version 7 server from a DB2 UDB Version 8 client.
- Certain connections and communication protocols:
 - Instance requests that require an ATTACH instead of a connection.
 - The ATTACH statement is not supported from a DB2 UDB Version 8 client to a DB2 UDB Version 7 server.
 - The only supported network protocol is TCP/IP.
 - Other network protocols like SNA, NetBIOS, IPX/SPX, and others are not supported.
- Some application features and tasks:
 - The DESCRIBE INPUT statement is not supported with one exception for ODBC/JDBC applications. In order to support DB2 UDB Version 8 clients running ODBC/JDBC applications accessing DB2 UDB Version 7 servers, a fix for DESCRIBE INPUT support must be applied to all DB2 UDB Version 7 servers where this type of access is required. This fix is associated with APAR IY30655 and will be available before the DB2 UDB Version 8 General Availability date. Use the “Contacting IBM” information in any DB2 document to find out how to get the fix associated with APAR IY30655. The DESCRIBE INPUT statement is a performance and usability enhancement to allow an application requestor to obtain a description of input parameter markers in a prepared statement. For a CALL statement, this includes the parameter markers associated with the IN and INOUT parameters for the stored procedure.
 - Using `Result.getObject(1)` will return a `BigDecimal` instead of a `Java Long` datatype as required by the JDBC specification. The DB2 UDB Version 7 DRDA server maps `BIGINT` to `DEC(19,0)` when it responds to a `DESCRIBE INPUT` request and when it retrieves data. This behavior occurs because the DB2 UDB Version 7 server operates at a DRDA level where `BIGINT` is not defined.
 - Query interrupts are not supported. This affects the `CLI/ODBC SQL_QUERY_TIMEOUT` connection attribute as well as the interrupt APIs.
 - Two-phase commit. The DB2 UDB Version 7 server cannot be used as a transaction manager database when using coordinated transactions that involve DB2 UDB Version 8 clients. Nor can a DB2 UDB Version 7 server participate in a coordinated transaction where a DB2 UDB Version 8 server may be the transaction manager database.
 - XA-compliant transaction managers. An application using a DB2 UDB Version 8 client cannot use a DB2 UDB Version 7 server as an XA resource. This includes WebSphere, Microsoft COM+/MTS, BEA WebLogic, and others that are part of a transaction management arrangement.
 - Monitoring. Monitor functions are not supported from a DB2 UDB Version 8 client to a DB2 UDB Version 7 server.

- Utilities. Those utilities that can be initiated by a client to a server are not supported when:
 1. The client is at DB2 UDB Version 8 and the server is at DB2 UDB Version 7.
 2. SQL statements greater than 32 KB in size.
- Query interrupts are not supported. This affects the CLI/ODBC SQL_QUERY_TIMEOUT connection attribute as well as the interrupt APIs.

In addition to these limitations and restrictions for DB2 UDB Version 8 clients working with DB2 UDB Version 7 servers, there are also similar limitations and restrictions for DB2 UDB Version 8 tools working with DB2 UDB Version 7 servers. The following DB2 UDB Version 8 tools support only DB2 UDB Version 8 servers:

- Control Center
- Task Center
- Journal
- Satellite Administration Center
- Information Catalog Center (including the Web-version of this center)
- Health Center (including the Web-version of this center)
- License Center
- Spatial Extender
- Tools Settings
- Development Center. You should use Stored Procedure Builder to develop server objects on pre-Version 8 servers.

The following DB2 UDB Version 8 tools support DB2 UDB Version 7 servers (with some restrictions) and DB2 UDB Version 8 servers:

- Configuration Assistant

It is possible to discover a DB2 UDB Version 7 server and catalog it. However, even though cataloged, no function will work if attempting to access the DB2 UDB Version 7 server. Also, you are able to import a DB2 UDB Version 7 profile to a DB2 UDB Version 8 server, or import a DB2 UDB Version 8 profile to a DB2 UDB Version 7 server. However, all other Configuration Assistant functions will not work with DB2 UDB Version 7 servers.
- Data Warehouse Center
- Replication Center
- Command Editor (the replacement for the Command Center, including the Web-version of this center)

Importing and saving scripts to and from DB2 UDB Version 7 servers is not possible. Any utility requiring an ATTACH will not work.
- SQL Assist
- Visual Explain

In general, any DB2 UDB Version 8 tool that is only launched from within the navigation tree of the Control Center, or any details view based on these tools, will not be available or accessible to DB2 UDB Version 7 and earlier servers. You should consider using the DB2 UDB Version 7 tools when working with DB2 UDB Version 7 or earlier servers.

Scrollable cursor support:

Windows	UNIX
---------	------

Change:

In Version 8, scrollable cursor functionality will not be supported from a Version 8 DB2 UDB for Unix and Windows client to a Version 7 DB2 UDB for Unix and Windows server. Support for scrollable cursors will only be available from a Version 8 DB2 UDB for Unix and Windows client to a DB2 UDB for Unix and Windows Version 8 server or to a DB2 UDB for z/OS and OS/390 Version 7 server. DB2 UDB for Unix and Windows Version 7 clients will continue to support existing scrollable cursor functionality to Version 8 DB2 UDB for Unix and Windows servers.

Resolution:

Upgrade servers to Version 8.

Version 7 server access via a DB2 Connect Version 8 server:

Windows	UNIX
---------	------

Change:

In Version 8, access from a DB2 UDB for Unix and Windows client to a Version 7 DB2 UDB server will not be supported through a Version 8 server, where the functionality is provided either by DB2 Connect Enterprise Edition Version 8 or by DB2 UDB Enterprise Server Edition Version 8.

Resolution:

Upgrade servers to Version 8.

Type 1 connection with CLP and embedded SQL:

Windows	UNIX
---------	------

Change:

In previous versions of DB2, when using the Command Line Processor (CLP) or embedded SQL and connected to a database with a Type 1 connection, an attempt to connect to another database during a unit of work would fail with an SQL0752N error. In Version 8, the unit of work is committed, the connection is reset, and the connection to the second database is allowed. The unit of work will be committed and the connection will be reset even if AUTOCOMMIT is off.

Messages:

DB2 Connect messages returned instead of DB2 messages:

Windows	UNIX
---------	------

Change:

In Version 8, conditions that would have returned a DB2 message in previous releases may now return a DB2 Connect message.

The messages affected by this change are related to bind, connection, or security errors. SQL errors for queries and other SQL requests are not affected by this change.

Examples:

- SQLCODE -30081 will be returned instead of SQLCODE -1224
- SQLCODE -30082 will be returned instead of SQLCODE -1403
- SQLCODE -30104 will be returned instead of SQLCODE -4930

Symptom:

Applications coded to react to DB2 messages may not behave as before.

Configuration parameters:

Obsolete database manager configuration parameters:

Windows	UNIX
---------	------

Change:

The following database manager configuration parameters are obsolete:

- *backbufsz*: In previous versions you could perform a backup operation using a default buffer size, and the value of *backbufsz* would be taken as the default. In Version 8 you should explicitly specify the size of your backup buffers when you use the backup utility.
- *dft_client_adpt*: DCE directory services are no longer supported
- *dft_client_comm*: DCE directory services are no longer supported
- *dir_obj_name*: DCE directory services are no longer supported
- *dir_path_name*: DCE directory services are no longer supported
- *dir_type*: DCE directory services are no longer supported
- *dos_rqrioblk*
- *drda_heap_sz*
- *fcm_num_anchors*, *fcm_num_connect*, and *fcm_num_rqb*: DB2 will now adjust message anchors, connection entries, and request blocks dynamically and automatically, so you will not have to adjust these parameters
- *fileserv*: IPX/SPX is no longer supported
- *initdari_jvm*: Java stored procedures will now run multithreaded by default, and are run in separate processes from other language routines, so this parameter is no longer supported
- *ipx_socket*: IPX/SPX is no longer supported
- *jdk11_path*: replaced by *jdk_path* database manager configuration parameter
- *keepdari*: replaced by *keepfenced* database manager configuration parameter
- *max_logicagents*: replaced by *max_connections* database manager configuration parameter
- *maxdari*: replaced by *fenced_pool* database manager configuration parameter

- *num_initdaris*: replaced by *num_initfenced* database manager configuration parameter
- *objectname*: IPX/SPX is no longer supported
- *restbufsz*: In previous versions you could perform a restore operation using a default buffer size, and the value of *restbufsz* would be taken as the default. In Version 8 you should explicitly specify the size of your restore buffers when use restore utility.
- *route_obj_name*: DCE directory services are no longer supported
- *ss_logon*: this is an OS/2[®] parameter, and OS/2 is no longer supported
- *udf_mem_sz*: UDFs no longer pass data in shared memory, so this parameter is not supported

Resolution:

Remove all references to these parameters from your applications.

Obsolete database configuration parameters:

Windows	UNIX
---------	------

Change:

The following database configuration parameters are obsolete:

- *buffpage*: In previous versions, you could create or alter a buffer pool using a default size, and the value of *buffpage* would be taken as the default. In Version 8, you should explicitly specify the size of your buffer pools, using the SIZE keyword on the ALTER BUFFERPOOL or CREATE BUFFERPOOL statements.
- *copyprotect*
- *indexsort*

Resolution:

Remove all references to these parameters from your applications.

Related concepts:

- “Version 9 incompatibilities with previous releases and changed behaviors” on page 260

Related reference:

- “Deprecated and discontinued features” on page 243

Appendix B. National language support (NLS)

This section contains information about the national language support (NLS) provided by DB2 databases, including information about territories, languages, and code pages (code sets) supported, and how to configure and use DB2 NLS features in your databases and applications.

National language versions

DB2 Database for Linux, UNIX, and Windows Version 9.1 is available in Simplified Chinese, Traditional Chinese, Czech, Danish, English, Finnish, French, German, Italian, Japanese, Korean, Norwegian, Polish, Brazilian Portuguese, Russian, Spanish, and Swedish.

The DB2 Run-Time Client is available in these additional languages: Arabic, Bulgarian, Croatian, Dutch, Greek, Hebrew, Hungarian, Portuguese, Romanian, Slovak, Slovenian, and Turkish.

Related reference:

- “Supported territory codes and code pages” on page 313

Supported territory codes and code pages

The following tables show the languages and code sets supported by the database servers, and how these values are mapped to territory code and code page values that are used by the database manager.

Note: When creating a database, you can use any supported code page on any supported platform.

The following is an explanation of the columns in the tables:

- **Code page** shows the IBM-defined code page as mapped from the operating system code set.
- **Group** shows whether a code page is single-byte (“S”), double-byte (“D”), or neutral (“N”). The “-n” is a number used to create a letter-number combination. Matching combinations show where connection and conversion is allowed by DB2 database systems. For example, all “S-1” groups can work together. However, if the group is neutral, then connection and conversion with any other code page listed is allowed.
- **Code set** shows the code set associated with the supported language. The code set is mapped to the DB2 code page.
- **Territory code** shows the code that is used by the database manager internally to provide region-specific support.
- **Locale** shows the locale values supported by the database manager.
- **Operating system** shows the operating system that supports the languages and code sets. When used in this column, the word “host” refers to an operating system such as z/OS that supports the EBCDIC code pages natively. Note that Linux on z/OS is not a host platform. You cannot use DB2 database manager to create a database in a host code page, but you can use DB2 database manager to connect to a host database in a supported host code page.

Table 42. Unicode

Code page	Group	Code set	Territory code	Locale	Operating system
1200	N-1	16-bit Unicode	Any	Any	Any
1208	N-1	UTF-8 encoding of Unicode	Any	Any	Any

Table 43. Albania, territory identifier: AL

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	355	sq_AL	AIX
850	S-1	IBM-850	355	-	AIX
923	S-1	ISO8859-15	355	sq_AL.8859-15	AIX
1208	N-1	UTF-8	355	SQ_AL	AIX
37	S-1	IBM-37	355	-	Host
1140	S-1	IBM-1140	355	-	Host
819	S-1	iso88591	355	-	HP-UX
923	S-1	iso885915	355	-	HP-UX
1051	S-1	roman8	355	-	HP-UX
437	S-1	IBM-437	355	-	OS/2
850	S-1	IBM-850	355	-	OS/2
819	S-1	ISO8859-1	355	-	Solaris
923	S-1	ISO8859-15	355	-	Solaris
1252	S-1	1252	355	-	Windows

Table 44. Arabic countries/regions, territory identifier: AA

Code page	Group	Code set	Territory code	Locale	Operating system
1046	S-6	IBM-1046	785	Ar_AA	AIX
1089	S-6	ISO8859-6	785	ar_AA	AIX
1208	N-1	UTF-8	785	AR_AA	AIX
420	S-6	IBM-420	785	-	Host
425	S-6	IBM-425	785	-	Host
1089	S-6	iso88596	785	ar_SA.iso88596	HP-UX
864	S-6	IBM-864	785	-	OS/2
1256	S-6	1256	785	-	Windows

Table 45. Australia, territory identifier: AU

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	61	en_AU	AIX
850	S-1	IBM-850	61	-	AIX
923	S-1	ISO8859-15	61	en_AU.8859-15	AIX
1208	N-1	UTF-8	61	EN_AU	AIX
37	S-1	IBM-37	61	-	Host
1140	S-1	IBM-1140	61	-	Host
819	S-1	iso88591	61	-	HP-UX
923	S-1	iso885915	61	-	HP-UX
1051	S-1	roman8	61	-	HP-UX
437	S-1	IBM-437	61	-	OS/2

Table 45. Australia, territory identifier: AU (continued)

Code page	Group	Code set	Territory code	Locale	Operating system
850	S-1	IBM-850	61	-	OS/2
819	S-1	ISO8859-1	61	en_AU	SCO
819	S-1	ISO8859-1	61	en_AU	Solaris
923	S-1	ISO8859-15	61	-	Solaris
1252	S-1	1252	61	-	Windows

Table 46. Austria, territory identifier: AT

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	43	-	AIX
850	S-1	IBM-850	43	-	AIX
923	S-1	ISO8859-15	43	-	AIX
1208	N-1	UTF-8	43	-	AIX
37	S-1	IBM-37	43	-	Host
1140	S-1	IBM-1140	43	-	Host
819	S-1	iso88591	43	-	HP-UX
923	S-1	iso885915	43	-	HP-UX
1051	S-1	roman8	43	-	HP-UX
819	S-1	ISO-8859-1	43	de_AT	Linux
923	S-1	ISO-8859-15	43	de_AT@euro	Linux
437	S-1	IBM-437	43	-	OS/2
850	S-1	IBM-850	43	-	OS/2
819	S-1	ISO8859-1	43	de_AT	SCO
819	S-1	ISO8859-1	43	de_AT	Solaris
923	S-1	ISO8859-15	43	de_AT.ISO8859-15	Solaris
1252	S-1	1252	43	-	Windows

Table 47. Belarus, territory identifier: BY

Code page	Group	Code set	Territory code	Locale	Operating system
1167	S-5	KOI8-RU	375	-	-
915	S-5	ISO8859-5	375	be_BY	AIX
1208	N-1	UTF-8	375	BE_BY	AIX
1025	S-5	IBM-1025	375	-	Host
1154	S-5	IBM-1154	375	-	Host
915	S-5	ISO8859-5	375	-	OS/2
1131	S-5	IBM-1131	375	-	OS/2
1251	S-5	1251	375	-	Windows

Table 48. Belgium, territory identifier: BE

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	32	fr_BE	AIX
819	S-1	ISO8859-1	32	nl_BE	AIX
850	S-1	IBM-850	32	Fr_BE	AIX
850	S-1	IBM-850	32	Nl_BE	AIX
923	S-1	ISO8859-15	32	fr_BE.8859-15	AIX
923	S-1	ISO8859-15	32	nl_BE.8859-15	AIX
1208	N-1	UTF-8	32	FR_BE	AIX

Table 48. Belgium, territory identifier: BE (continued)

Code page	Group	Code set	Territory code	Locale	Operating system
1208	N-1	UTF-8	32	NL_BE	AIX
274	S-1	IBM-274	32	-	Host
500	S-1	IBM-500	32	-	Host
1148	S-1	IBM-1148	32	-	Host
819	S-1	iso88591	32	-	HP-UX
923	S-1	iso885915	32	-	HP-UX
819	S-1	ISO-8859-1	32	fr_BE	Linux
819	S-1	ISO-8859-1	32	nl_BE	Linux
923	S-1	ISO-8859-15	32	fr_BE@euro	Linux
923	S-1	ISO-8859-15	32	nl_BE@euro	Linux
437	S-1	IBM-437	32	-	OS/2
850	S-1	IBM-850	32	-	OS/2
819	S-1	ISO8859-1	32	fr_BE	SCO
819	S-1	ISO8859-1	32	nl_BE	SCO
819	S-1	ISO8859-1	32	fr_BE	Solaris
819	S-1	ISO8859-1	32	nl_BE	Solaris
923	S-1	ISO8859-15	32	fr_BE.ISO8859-15	Solaris
923	S-1	ISO8859-15	32	nl_BE.ISO8859-15	Solaris
1252	S-1	1252	32	-	Windows

Table 49. Bulgaria, territory identifier: BG

Code page	Group	Code set	Territory code	Locale	Operating system
915	S-5	ISO8859-5	359	bg_BG	AIX
1208	N-1	UTF-8	359	BG_BG	AIX
1025	S-5	IBM-1025	359	-	Host
1154	S-5	IBM-1154	359	-	Host
915	S-5	iso88595	359	bg_BG.iso88595	HP-UX
855	S-5	IBM-855	359	-	OS/2
915	S-5	ISO8859-5	359	-	OS/2
1251	S-5	1251	359	-	Windows

Table 50. Brazil, territory identifier: BR

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	55	pt_BR	AIX
850	S-1	IBM-850	55	-	AIX
923	S-1	ISO8859-15	55	pt_BR.8859-15	AIX
1208	N-1	UTF-8	55	PT_BR	AIX
37	S-1	IBM-37	55	-	Host
1140	S-1	IBM-1140	55	-	Host
819	S-1	ISO8859-1	55	-	HP-UX
923	S-1	ISO8859-15	55	-	HP-UX
819	S-1	ISO-8859-1	55	pt_BR	Linux
923	S-1	ISO-8859-15	55	-	Linux
850	S-1	IBM-850	55	-	OS/2
819	S-1	ISO8859-1	55	pt_BR	SCO
819	S-1	ISO8859-1	55	pt_BR	Solaris
923	S-1	ISO8859-15	55	-	Solaris
1252	S-1	1252	55	-	Windows

Table 51. Canada, territory identifier: CA

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	1	fr_CA	AIX
850	S-1	IBM-850	1	Fr_CA	AIX
923	S-1	ISO8859-15	1	fr_CA.8859-15	AIX
1208	N-1	UTF-8	1	FR_CA	AIX
37	S-1	IBM-37	1	-	Host
1140	S-1	IBM-1140	1	-	Host
819	S-1	iso88591	1	fr_CA.iso88591	HP-UX
923	S-1	iso885915	1	-	HP-UX
1051	S-1	roman8	1	fr_CA.roman8	HP-UX
819	S-1	ISO-8859-1	1	en_CA	Linux
923	S-1	ISO-8859-15	1	-	Linux
850	S-1	IBM-850	1	-	OS/2
819	S-1	ISO8859-1	1	en_CA	SCO
819	S-1	ISO8859-1	1	fr_CA	SCO
819	S-1	ISO8859-1	1	en_CA	Solaris
923	S-1	ISO8859-15	1	-	Solaris
1252	S-1	1252	1	-	Windows

Table 52. Canada (French), territory identifier: CA

Code page	Group	Code set	Territory code	Locale	Operating system
863	S-1	IBM-863	2	-	OS/2

Table 53. China (PRC), territory identifier: CN

Code page	Group	Code set	Territory code	Locale	Operating system
1383	D-4	IBM-eucCN	86	zh_CN	AIX
1386	D-4	GBK	86	Zh_CN.GBK	AIX
1208	N-1	UTF-8	86	ZH_CN	AIX
935	D-4	IBM-935	86	-	Host
1388	D-4	IBM-1388	86	-	Host
1383	D-4	hp15CN	86	zh_CN.hp15CN	HP-UX
1386	D-4	GBK	86	zh_CN.GBK	Linux
1381	D-4	IBM-1381	86	-	OS/2
1386	D-4	GBK	86	-	OS/2
1383	D-4	eucCN	86	zh_CN	SCO
1383	D-4	eucCN	86	zh_CN.eucCN	SCO
1383	D-4	gb2312	86	zh	Solaris
1208	N-1	UTF-8	86	zh.UTF-8	Solaris
1381	D-4	IBM-1381	86	-	Windows
1386	D-4	GBK	86	-	Windows
1392/5488	D-4		86	-	

See note 1 on page 332.

Table 54. Croatia, territory identifier: HR

Code page	Group	Code set	Territory code	Locale	Operating system
912	S-2	ISO8859-2	385	hr_HR	AIX
1208	N-1	UTF-8	385	HR_HR	AIX
870	S-2	IBM-870	385	-	Host
1153	S-2	IBM-1153	385	-	Host
912	S-2	iso88592	385	hr_HR.iso88592	HP-UX
912	S-2	ISO-8859-2	385	hr_HR	Linux
852	S-2	IBM-852	385	-	OS/2
912	S-2	ISO8859-2	385	hr_HR.ISO8859-2	SCO
1250	S-2	1250	385	-	Windows

Table 55. Czech Republic, territory identifier: CZ

Code page	Group	Code set	Territory code	Locale	Operating system
912	S-2	ISO8859-2	421	cs_CZ	AIX
1208	N-1	UTF-8	421	CS_CZ	AIX
870	S-2	IBM-870	421	-	Host
1153	S-2	IBM-1153	421	-	Host
912	S-2	iso88592	421	cs_CZ.iso88592	HP-UX
912	S-2	ISO-8859-2	421	cs_CZ	Linux
852	S-2	IBM-852	421	-	OS/2
912	S-2	ISO8859-2	421	cs_CZ.ISO8859-2	SCO
1250	S-2	1250	421	-	Windows

Table 56. Denmark, territory identifier: DK

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	45	da_DK	AIX
850	S-1	IBM-850	45	Da_DK	AIX
923	S-1	ISO8859-15	45	da_DK.8859-15	AIX
1208	N-1	UTF-8	45	DA_DK	AIX
277	S-1	IBM-277	45	-	Host
1142	S-1	IBM-1142	45	-	Host
819	S-1	iso88591	45	da_DK.iso88591	HP-UX
923	S-1	iso885915	45	-	HP-UX
1051	S-1	roman8	45	da_DK.roman8	HP-UX
819	S-1	ISO-8859-1	45	da_DK	Linux
923	S-1	ISO-8859-15	45	-	Linux
850	S-1	IBM-850	45	-	OS/2
819	S-1	ISO8859-1	45	da	SCO
819	S-1	ISO8859-1	45	da_DA	SCO
819	S-1	ISO8859-1	45	da_DK	SCO
819	S-1	ISO8859-1	45	da	Solaris
923	S-1	ISO8859-15	45	da.ISO8859-15	Solaris
1252	S-1	1252	45	-	Windows

Table 57. Estonia, territory identifier: EE

Code page	Group	Code set	Territory code	Locale	Operating system
922	S-10	IBM-922	372	Et_EE	AIX

Table 57. Estonia, territory identifier: EE (continued)

Code page	Group	Code set	Territory code	Locale	Operating system
1208	N-1	UTF-8	372	ET_EE	AIX
1122	S-10	IBM-1122	372	-	Host
1157	S-10	IBM-1157	372	-	Host
922	S-10	IBM-922	372	-	OS/2
1257	S-10	1257	372	-	Windows

Table 58. Finland, territory identifier: FI

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	358	fi_FI	AIX
850	S-1	IBM-850	358	Fi_FI	AIX
923	S-1	ISO8859-15	358	fi_FI.8859-15	AIX
1208	N-1	UTF-8	358	FI_FI	AIX
278	S-1	IBM-278	358	-	Host
1143	S-1	IBM-1143	358	-	Host
819	S-1	iso88591	358	fi_FI.iso88591	HP-UX
923	S-1	iso885915	358	-	HP-UX
1051	S-1	roman8	358	fi-FI.roman8	HP-UX
819	S-1	ISO-8859-1	358	fi_FI	Linux
923	S-1	ISO-8859-15	358	fi_FI@euro	Linux
437	S-1	IBM-437	358	-	OS/2
850	S-1	IBM-850	358	-	OS/2
819	S-1	ISO8859-1	358		SCO
819	S-1	ISO8859-1	358	fi_FI	SCO
819	S-1	ISO8859-1	358	sv_FI	SCO
819	S-1	ISO8859-1	358	-	Solaris
923	S-1	ISO8859-15	358	fi.ISO8859-15	Solaris
1252	S-1	1252	358	-	Windows

Table 59. FYR Macedonia, territory identifier: MK

Code page	Group	Code set	Territory code	Locale	Operating system
915	S-5	ISO8859-5	389	mk_MK	AIX
1208	N-1	UTF-8	389	MK_MK	AIX
1025	S-5	IBM-1025	389	-	Host
1154	S-5	IBM-1154	389	-	Host
915	S-5	iso88595	389	-	HP-UX
855	S-5	IBM-855	389	-	OS/2
915	S-5	ISO8859-5	389	-	OS/2
1251	S-5	1251	389	-	Windows

Table 60. France, territory identifier: FR

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	33	fr_FR	AIX
850	S-1	IBM-850	33	Fr_FR	AIX
923	S-1	ISO8859-15	33	fr_FR.8859-15	AIX
1208	N-1	UTF-8	33	FR_FR	AIX
297	S-1	IBM-297	33	-	Host

Table 60. France, territory identifier: FR (continued)

Code page	Group	Code set	Territory code	Locale	Operating system
1147	S-1	IBM-1147	33	-	Host
819	S-1	iso88591	33	fr_FR.iso88591	HP-UX
923	S-1	iso885915	33	-	HP-UX
1051	S-1	roman8	33	fr_FR.roman8	HP-UX
819	S-1	ISO-8859-1	33	fr_FR	Linux
923	S-1	ISO-8859-15	33	fr_FR@euro	Linux
437	S-1	IBM-437	33	-	OS/2
850	S-1	IBM-850	33	-	OS/2
819	S-1	ISO8859-1	33		SCO
819	S-1	ISO8859-1	33	fr_FR	SCO
819	S-1	ISO8859-1	33		Solaris
923	S-1	ISO8859-15	33	fr.ISO8859-15	Solaris
1208	N-1	UTF-8	33	fr.UTF-8	Solaris
1252	S-1	1252	33	-	Windows

Table 61. Germany, territory identifier: DE

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	49	de_DE	AIX
850	S-1	IBM-850	49	De_DE	AIX
923	S-1	ISO8859-15	49	de_DE.8859-15	AIX
1208	N-1	UTF-8	49	DE_DE	AIX
273	S-1	IBM-273	49	-	Host
1141	S-1	IBM-1141	49	-	Host
819	S-1	iso88591	49	de_DE.iso88591	HP-UX
923	S-1	iso885915	49	-	HP-UX
1051	S-1	roman8	49	de_DE.roman8	HP-UX
819	S-1	ISO-8859-1	49	de_DE	Linux
923	S-1	ISO-8859-15	49	de_DE@euro	Linux
437	S-1	IBM-437	49	-	OS/2
850	S-1	IBM-850	49	-	OS/2
819	S-1	ISO8859-1	49		SCO
819	S-1	ISO8859-1	49	de_DE	SCO
819	S-1	ISO8859-1	49		Solaris
923	S-1	ISO8859-15	49	de.ISO8859-15	Solaris
1208	N-1	UTF-8	49	de.UTF-8	Solaris
1252	S-1	1252	49	-	Windows

Table 62. Greece, territory identifier: GR

Code page	Group	Code set	Territory code	Locale	Operating system
813	S-7	ISO8859-7	30	el_GR	AIX
1208	N-1	UTF-8	30	EL_GR	AIX
423	S-7	IBM-423	30	-	Host
875	S-7	IBM-875	30	-	Host
813	S-7	iso88597	30	el_GR.iso88597	HP-UX
813	S-7	ISO-8859-7	30	el_GR	Linux
813	S-7	ISO8859-7	30	-	OS/2
869	S-7	IBM-869	30	-	OS/2
813	S-7	ISO8859-7	30	el_GR.ISO8859-7	SCO

Table 62. Greece, territory identifier: GR (continued)

Code page	Group	Code set	Territory code	Locale	Operating system
737	S-7	737	30	-	Windows
1253	S-7	1253	30	-	Windows

Table 63. Hungary, territory identifier: HU

Code page	Group	Code set	Territory code	Locale	Operating system
912	S-2	ISO8859-2	36	hu_HU	AIX
1208	N-1	UTF-8	36	HU_HU	AIX
870	S-2	IBM-870	36	-	Host
1153	S-2	IBM-1153	36	-	Host
912	S-2	iso88592	36	hu_HU.iso88592	HP-UX
912	S-2	ISO-8859-2	36	hu_HU	Linux
852	S-2	IBM-852	36	-	OS/2
912	S-2	ISO8859-2	36	hu_HU.ISO8859-2	SCO
1250	S-2	1250	36	-	Windows

Table 64. Iceland, territory identifier: IS

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	354	is_IS	AIX
850	S-1	IBM-850	354	Is_IS	AIX
923	S-1	ISO8859-15	354	is_IS.8859-15	AIX
1208	N-1	UTF-8	354	IS_IS	AIX
871	S-1	IBM-871	354	-	Host
1149	S-1	IBM-1149	354	-	Host
819	S-1	iso88591	354	is_IS.iso88591	HP-UX
923	S-1	iso885915	354	-	HP-UX
1051	S-1	roman8	354	is_IS.roman8	HP-UX
819	S-1	ISO-8859-1	354	is_IS	Linux
923	S-1	ISO-8859-15	354	-	Linux
850	S-1	IBM-850	354	-	OS/2
819	S-1	ISO8859-1	354		SCO
819	S-1	ISO8859-1	354	is_IS	SCO
819	S-1	ISO8859-1	354	-	Solaris
923	S-1	ISO8859-15	354	-	Solaris
1252	S-1	1252	354	-	Windows

Table 65. India, territory identifier: IN

Code page	Group	Code set	Territory code	Locale	Operating system
806	S-13	IBM-806	91	hi_IN	-
1137	S-13	IBM-1137	91	-	Host

Table 66. Indonesia, territory identifier: ID

Code page	Group	Code set	Territory code	Locale	Operating system
1252	S-1	1252	62	-	Windows

Table 67. Ireland, territory identifier: IE

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	353	-	AIX
850	S-1	IBM-850	353	-	AIX
923	S-1	ISO8859-15	353	-	AIX
1208	N-1	UTF-8	353	-	AIX
285	S-1	IBM-285	353	-	Host
1146	S-1	IBM-1146	353	-	Host
819	S-1	iso88591	353	-	HP-UX
923	S-1	iso885915	353	-	HP-UX
1051	S-1	roman8	353	-	HP-UX
819	S-1	ISO-8859-1	353	en_IE	Linux
923	S-1	ISO-8859-15	353	en_IE@euro	Linux
437	S-1	IBM-437	353	-	OS/2
850	S-1	IBM-850	353	-	OS/2
819	S-1	ISO8859-1	353	en_IE.ISO8859-1	SCO
819	S-1	ISO8859-1	353	en_IE	Solaris
923	S-1	ISO8859-15	353	en_IE.ISO8859-15	Solaris
1252	S-1	1252	353	-	Windows

Table 68. Israel, territory identifier: IL

Code page	Group	Code set	Territory code	Locale	Operating system
856	S-8	IBM-856	972	Iw_IL	AIX
916	S-8	ISO8859-8	972	iw_IL	AIX
1208	N-1	UTF-8	972	HE-IL	AIX
916	S-8	ISO-8859-8	972	iw_IL	Linux
424	S-8	IBM-424	972	-	Host
862	S-8	IBM-862	972	-	OS/2
1255	S-8	1255	972	-	Windows

Table 69. Italy, territory identifier: IT

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	39	it_IT	AIX
850	S-1	IBM-850	39	It_IT	AIX
923	S-1	ISO8859-15	39	it_IT.8859-15	AIX
1208	N-1	UTF-8	39	It_IT	AIX
280	S-1	IBM-280	39	-	Host
1144	S-1	IBM-1144	39	-	Host
819	S-1	iso88591	39	it_IT.iso88591	HP-UX
923	S-1	iso885915	39	-	HP-UX
1051	S-1	roman8	39	it_IT.roman8	HP-UX
819	S-1	ISO-8859-1	39	it_IT	Linux
923	S-1	ISO-8859-15	39	it_IT@euro	Linux
437	S-1	IBM-437	39	-	OS/2
850	S-1	IBM-850	39	-	OS/2
819	S-1	ISO8859-1	39	-	SCO
819	S-1	ISO8859-1	39	it_IT	SCO
819	S-1	ISO8859-1	39	-	Solaris
923	S-1	ISO8859-15	39	it.ISO8859-15	Solaris
1208	N-1	UTF-8	39	it.UTF-8	Solaris

Table 69. Italy, territory identifier: IT (continued)

Code page	Group	Code set	Territory code	Locale	Operating system
1252	S-1	1252	39	-	Windows

Table 70. Japan, territory identifier: JP

Code page	Group	Code set	Territory code	Locale	Operating system
932	D-1	IBM-932	81	Ja_JP	AIX
943	D-1	IBM-943	81	Ja_JP	AIX
See note 2 on page 332.					
954	D-1	IBM-eucJP	81	ja_JP	AIX
1208	N-1	UTF-8	81	JA_JP	AIX
930	D-1	IBM-930	81	-	Host
939	D-1	IBM-939	81	-	Host
5026	D-1	IBM-5026	81	-	Host
5035	D-1	IBM-5035	81	-	Host
1390	D-1		81	-	Host
1399	D-1		81	-	Host
954	D-1	eucJP	81	ja_JP.eucJP	HP-UX
5039	D-1	SJIS	81	ja_JP:SJIS	HP-UX
954	D-1	EUC-JP	81	ja_JP	Linux
932	D-1	IBM-932	81	-	OS/2
942	D-1	IBM-942	81	-	OS/2
943	D-1	IBM-943	81	-	OS/2
954	D-1	eucJP	81	ja	SCO
954	D-1	eucJP	81	ja_JP	SCO
954	D-1	eucJP	81	ja_JP.EUC	SCO
954	D-1	eucJP	81	ja_JP.eucJP	SCO
943	D-1	IBM-943	81	ja_JP.PCK	Solaris
954	D-1	eucJP	81	ja	Solaris
1208	N-1	UTF-8	81	ja_JP.UTF-8	Solaris
943	D-1	IBM-943	81	-	Windows
1394	D-1		81	-	
See note 3 on page 332.					

Table 71. Kazakhstan, territory identifier: KZ

Code page	Group	Code set	Territory code	Locale	Operating system
1251	S-5	1251	7	-	Windows

Table 72. Korea, South, territory identifier: KR

Code page	Group	Code set	Territory code	Locale	Operating system
970	D-3	IBM-eucKR	82	ko_KR	AIX
1208	N-1	UTF-8	82	KO_KR	AIX
933	D-3	IBM-933	82	-	Host
1364	D-3	IBM-1364	82	-	Host
970	D-3	eucKR	82	ko_KR.eucKR	HP-UX
970	D-3	EUC-KR	82	ko_KR	Linux
949	D-3	IBM-949	82	-	OS/2
970	D-3	eucKR	82	ko_KR.eucKR	SGI

Table 72. Korea, South, territory identifier: KR (continued)

Code page	Group	Code set	Territory code	Locale	Operating system
970	D-3	5601	82	ko	Solaris
1208	N-1	UTF-8	82	ko.UTF-8	Solaris
1363	D-3	1363	82	-	Windows

Table 73. Latin America, territory identifier: Lat

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	3	-	AIX
850	S-1	IBM-850	3	-	AIX
923	S-1	ISO8859-15	3	-	AIX
1208	N-1	UTF-8	3	-	AIX
284	S-1	IBM-284	3	-	Host
1145	S-1	IBM-1145	3	-	Host
819	S-1	iso88591	3	-	HP-UX
923	S-1	iso885915	3	-	HP-UX
1051	S-1	roman8	3	-	HP-UX
819	S-1	ISO-8859-1	3	-	Linux
923	S-1	ISO-8859-15	3	-	Linux
437	S-1	IBM-437	3	-	OS/2
850	S-1	IBM-850	3	-	OS/2
819	S-1	ISO8859-1	3	-	Solaris
923	S-1	ISO8859-15	3	-	Solaris
1252	S-1	1252	3	-	Windows

Table 74. Latvia, territory identifier: LV

Code page	Group	Code set	Territory code	Locale	Operating system
921	S-10	IBM-921	371	Lv_LV	AIX
1208	N-1	UTF-8	371	LV_LV	AIX
1112	S-10	IBM-1112	371	-	Host
1156	S-10	IBM-1156	371	-	Host
921	S-10	IBM-921	371	-	OS/2
1257	S-10	1257	371	-	Windows

Table 75. Lithuania, territory identifier: LT

Code page	Group	Code set	Territory code	Locale	Operating system
921	S-10	IBM-921	370	Lt_LT	AIX
1208	N-1	UTF-8	370	LT_LT	AIX
1112	S-10	IBM-1112	370	-	Host
1156	S-10	IBM-1156	370	-	Host
921	S-10	IBM-921	370	-	OS/2
1257	S-10	1257	370	-	Windows

Table 76. Malaysia, territory identifier: ID

Code page	Group	Code set	Territory code	Locale	Operating system
1252	S-1	1252	60	-	Windows

Table 77. Netherlands, territory identifier: NL

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	31	nl_NL	AIX
850	S-1	IBM-850	31	NL_NL	AIX
923	S-1	ISO8859-15	31	nl_NL.8859-15	AIX
1208	N-1	UTF-8	31	NL_NL	AIX
37	S-1	IBM-37	31	-	Host
1140	S-1	IBM-1140	31	-	Host
819	S-1	iso88591	31	nl_NL.iso88591	HP-UX
923	S-1	iso885915	31	-	HP-UX
1051	S-1	roman8	31	nl_NL.roman8	HP-UX
819	S-1	ISO-8859-1	31	nl_NL	Linux
923	S-1	ISO-8859-15	31	nl_NL@euro	Linux
437	S-1	IBM-437	31	-	OS/2
850	S-1	IBM-850	31	-	OS/2
819	S-1	ISO8859-1	31	nl	SCO
819	S-1	ISO8859-1	31	nl_NL	SCO
819	S-1	ISO8859-1	31	nl	Solaris
923	S-1	ISO8859-15	31	nl.ISO8859-15	Solaris
1252	S-1	1252	31	-	Windows

Table 78. New Zealand, territory identifier: NZ

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	64	-	AIX
850	S-1	IBM-850	64	-	AIX
923	S-1	ISO8859-15	64	-	AIX
1208	N-1	UTF-8	64	-	AIX
37	S-1	IBM-37	64	-	Host
1140	S-1	IBM-1140	64	-	Host
819	S-1	ISO8859-1	64	-	HP-UX
923	S-1	ISO8859-15	64	-	HP-UX
850	S-1	IBM-850	64	-	OS/2
819	S-1	ISO8859-1	64	en_NZ	SCO
819	S-1	ISO8859-1	64	en_NZ	Solaris
923	S-1	ISO8859-15	64	-	Solaris
1252	S-1	1252	64	-	Windows

Table 79. Norway, territory identifier: NO

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	47	no_NO	AIX
850	S-1	IBM-850	47	No_NO	AIX
923	S-1	ISO8859-15	47	no_NO.8859-15	AIX
1208	N-1	UTF-8	47	NO_NO	AIX
277	S-1	IBM-277	47	-	Host
1142	S-1	IBM-1142	47	-	Host
819	S-1	iso88591	47	no_NO.iso88591	HP-UX
923	S-1	iso885915	47	-	HP-UX
1051	S-1	roman8	47	no_NO.roman8	HP-UX

Table 79. Norway, territory identifier: NO (continued)

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO-8859-1	47	no_NO	Linux
923	S-1	ISO-8859-15	47	-	Linux
850	S-1	IBM-850	47	-	OS/2
819	S-1	ISO8859-1	47	no	SCO
819	S-1	ISO8859-1	47	no_NO	SCO
819	S-1	ISO8859-1	47	no	Solaris
923	S-1	ISO8859-15	47	-	Solaris
1252	S-1	1252	47	-	Windows

Table 80. Poland, territory identifier: PL

Code page	Group	Code set	Territory code	Locale	Operating system
912	S-2	ISO8859-2	48	pl_PL	AIX
1208	N-1	UTF-8	48	PL_PL	AIX
870	S-2	IBM-870	48	-	Host
1153	S-2	IBM-1153	48	-	Host
912	S-2	iso88592	48	pl_PL.iso88592	HP-UX
912	S-2	ISO-8859-2	48	pl_PL	Linux
852	S-2	IBM-852	48	-	OS/2
912	S-2	ISO8859-2	48	pl_PL.ISO8859-2	SCO
1250	S-2	1250	48	-	Windows

Table 81. Portugal, territory identifier: PT

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	351	pt_PT	AIX
850	S-1	IBM-850	351	Pt_PT	AIX
923	S-1	ISO8859-15	351	pt_PT.8859-15	AIX
1208	N-1	UTF-8	351	PT_PT	AIX
37	S-1	IBM-37	351	-	Host
1140	S-1	IBM-1140	351	-	Host
819	S-1	iso88591	351	pt_PT.iso88591	HP-UX
923	S-1	iso885915	351	-	HP-UX
1051	S-1	roman8	351	pt_PT.roman8	HP-UX
819	S-1	ISO-8859-1	351	pt_PT	Linux
923	S-1	ISO-8859-15	351	pt_PT@euro	Linux
850	S-1	IBM-850	351	-	OS/2
860	S-1	IBM-860	351	-	OS/2
819	S-1	ISO8859-1	351	pt	SCO
819	S-1	ISO8859-1	351	pt_PT	SCO
819	S-1	ISO8859-1	351	pt	Solaris
923	S-1	ISO8859-15	351	pt.ISO8859-15	Solaris
1252	S-1	1252	351	-	Windows

Table 82. Romania, territory identifier: RO

Code page	Group	Code set	Territory code	Locale	Operating system
912	S-2	ISO8859-2	40	ro_RO	AIX
1208	N-1	UTF-8	40	RO_RO	AIX

Table 82. Romania, territory identifier: RO (continued)

Code page	Group	Code set	Territory code	Locale	Operating system
870	S-2	IBM-870	40	-	Host
1153	S-2	IBM-1153	40	-	Host
912	S-2	iso88592	40	ro_RO.iso88592	HP-UX
912	S-2	ISO-8859-2	40	ro_RO	Linux
852	S-2	IBM-852	40	-	OS/2
912	S-2	ISO8859-2	40	ro_RO.ISO8859-2	SCO
1250	S-2	1250	40	-	Windows

Table 83. Russia, territory identifier: RU

Code page	Group	Code set	Territory code	Locale	Operating system
915	S-5	ISO8859-5	7	ru_RU	AIX
1208	N-1	UTF-8	7	RU_RU	AIX
1025	S-5	IBM-1025	7	-	Host
1154	S-5	IBM-1154	7	-	Host
915	S-5	iso88595	7	ru_RU.iso88595	HP-UX
878	S-5	KOI8-R	7	ru_RU.koi8-r	Linux, Solaris
915	S-5	ISO-8859-5	7	ru_RU	Linux
866	S-5	IBM-866	7	-	OS/2
915	S-5	ISO8859-5	7	-	OS/2
915	S-5	ISO8859-5	7	ru_RU.ISO8859-5	SCO
1251	S-5	1251	7	-	Windows

Table 84. Serbia/Montenegro, territory identifier: SP

Code page	Group	Code set	Territory code	Locale	Operating system
915	S-5	ISO8859-5	381	sr_SP	AIX
1208	N-1	UTF-8	381	SR_SP	AIX
1025	S-5	IBM-1025	381	-	Host
1154	S-5	IBM-1154	381	-	Host
915	S-5	iso88595	381	-	HP-UX
855	S-5	IBM-855	381	-	OS/2
915	S-5	ISO8859-5	381	-	OS/2
1251	S-5	1251	381	-	Windows

Table 85. Slovakia, territory identifier: SK

Code page	Group	Code set	Territory code	Locale	Operating system
912	S-2	ISO8859-2	422	sk_SK	AIX
1208	N-1	UTF-8	422	SK_SK	AIX
870	S-2	IBM-870	422	-	Host
1153	S-2	IBM-1153	422	-	Host
912	S-2	iso88592	422	sk_SK.iso88592	HP-UX
852	S-2	IBM-852	422	-	OS/2
912	S-2	ISO8859-2	422	sk_SK.ISO8859-2	SCO
1250	S-2	1250	422	-	Windows

Table 86. Slovenia, territory identifier: SI

Code page	Group	Code set	Territory code	Locale	Operating system
912	S-2	ISO8859-2	386	sl_SI	AIX
1208	N-1	UTF-8	386	SL_SI	AIX
870	S-2	IBM-870	386	-	Host
1153	S-2	IBM-1153	386	-	Host
912	S-2	iso88592	386	sl_SI.iso88592	HP-UX
912	S-2	ISO-8859-2	386	sl_SI	Linux
852	S-2	IBM-852	386	-	OS/2
912	S-2	ISO8859-2	386	sl_SI.ISO8859-2	SCO
1250	S-2	1250	386	-	Windows

Table 87. South Africa, territory identifier: ZA

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	27	en_ZA	AIX
850	S-1	IBM-850	27	En_ZA	AIX
923	S-1	ISO8859-15	27	en_ZA.8859-15	AIX
1208	N-1	UTF-8	27	EN_ZA	AIX
285	S-1	IBM-285	27	-	Host
1146	S-1	IBM-1146	27	-	Host
819	S-1	iso88591	27	-	HP-UX
923	S-1	iso885915	27	-	HP-UX
1051	S-1	roman8	27	-	HP-UX
437	S-1	IBM-437	27	-	OS/2
850	S-1	IBM-850	27	-	OS/2
819	S-1	ISO8859-1	27	en_ZA.ISO8859-1	SCO
819	S-1	ISO8859-1	27	-	Solaris
923	S-1	ISO8859-15	27	-	Solaris
1252	S-1	1252	27	-	Windows

Table 88. Spain, territory identifier: ES

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	34	es_ES	AIX
850	S-1	IBM-850	34	Es_ES	AIX
923	S-1	ISO8859-15	34	es_ES.8859-15	AIX
1208	N-1	UTF-8	34	ES_ES	AIX
284	S-1	IBM-284	34	-	Host
1145	S-1	IBM-1145	34	-	Host
819	S-1	iso88591	34	es_ES.iso88591	HP-UX
923	S-1	iso885915	34	-	HP-UX
1051	S-1	roman8	34	es_ES.roman8	HP-UX
819	S-1	ISO-8859-1	34	es_ES	Linux
923	S-1	ISO-8859-15	34	es_ES@euro	Linux
437	S-1	IBM-437	34	-	OS/2
850	S-1	IBM-850	34	-	OS/2
819	S-1	ISO8859-1	34	es	SCO
819	S-1	ISO8859-1	34	es_ES	SCO
819	S-1	ISO8859-1	34	es	Solaris
923	S-1	ISO8859-15	34	es.ISO8859-15	Solaris
1208	N-1	UTF-8	34	es.UTF-8	Solaris

Table 88. Spain, territory identifier: ES (continued)

Code page	Group	Code set	Territory code	Locale	Operating system
1252	S-1	1252	34	-	Windows

Table 89. Spain (Catalan), territory identifier: ES

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	34	ca_ES	AIX
850	S-1	IBM-850	34	Ca_ES	AIX
923	S-1	ISO8859-15	34	ca_ES.8859-15	AIX
1208	N-1	UTF-8	34	CA_ES	AIX

Table 90. Sweden, territory identifier: SE

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	46	sv_SE	AIX
850	S-1	IBM-850	46	Sv_SE	AIX
923	S-1	ISO8859-15	46	sv_SE.8859-15	AIX
1208	N-1	UTF-8	46	SV_SE	AIX
278	S-1	IBM-278	46	-	Host
1143	S-1	IBM-1143	46	-	Host
819	S-1	iso88591	46	sv_SE.iso88591	HP-UX
923	S-1	iso885915	46	-	HP-UX
1051	S-1	roman8	46	sv_SE.roman8	HP-UX
819	S-1	ISO-8859-1	46	sv_SE	Linux
923	S-1	ISO-8859-15	46	-	Linux
437	S-1	IBM-437	46	-	OS/2
850	S-1	IBM-850	46	-	OS/2
819	S-1	ISO8859-1	46	sv	SCO
819	S-1	ISO8859-1	46	sv_SE	SCO
819	S-1	ISO8859-1	46	sv	Solaris
923	S-1	ISO8859-15	46	sv.ISO8859-15	Solaris
1208	N-1	UTF-8	46	sv.UTF-8	Solaris
1252	S-1	1252	46	-	Windows

Table 91. Switzerland, territory identifier: CH

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	41	de_CH	AIX
850	S-1	IBM-850	41	De_CH	AIX
923	S-1	ISO8859-15	41	de_CH.8859-15	AIX
1208	N-1	UTF-8	41	DE_CH	AIX
500	S-1	IBM-500	41	-	Host
1148	S-1	IBM-1148	41	-	Host
819	S-1	iso88591	41	-	HP-UX
923	S-1	iso885915	41	-	HP-UX
1051	S-1	roman8	41	-	HP-UX
819	S-1	ISO-8859-1	41	de_CH	Linux
923	S-1	ISO-8859-15	41	-	Linux
437	S-1	IBM-437	41	-	OS/2
850	S-1	IBM-850	41	-	OS/2

Table 91. Switzerland, territory identifier: CH (continued)

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	41	de_CH	SCO
819	S-1	ISO8859-1	41	fr_CH	SCO
819	S-1	ISO8859-1	41	it_CH	SCO
819	S-1	ISO8859-1	41	de_CH	Solaris
923	S-1	ISO8859-15	41	-	Solaris
1252	S-1	1252	41	-	Windows

Table 92. Taiwan, territory identifier: TW

Code page	Group	Code set	Territory code	Locale	Operating system
950	D-2	big5	88	Zh_TW	AIX
See note 8 on page 333.					
964	D-2	IBM-eucTW	88	zh_TW	AIX
1208	N-1	UTF-8	88	ZH_TW	AIX
937	D-2	IBM-937	88	-	Host
1371	D-2	IBM-1371	88	-	Host
950	D-2	big5	88	zh_TW.big5	HP-UX
964	D-2	eucTW	88	zh_TW.eucTW	HP-UX
950	D-2	BIG5	88	zh_TW	Linux
938	D-2	IBM-938	88	-	OS/2
948	D-2	IBM-948	88	-	OS/2
950	D-2	big5	88	-	OS/2
950	D-2	big5	88	zh_TW.BIG5	Solaris
964	D-2	cns11643	88	zh_TW	Solaris
1208	N-1	UTF-8	88	zh_TW.UTF-8	Solaris
950	D-2	big5	88	-	Windows
See note 8 on page 333.					

Table 93. Thailand, territory identifier: TH

Code page	Group	Code set	Territory code	Locale	Operating system
874	S-20	TIS620-1	66	th_TH	AIX
1208	N-1	UTF-8	66	TH_TH	AIX
838	S-20	IBM-838	66	-	Host
1160	S-20	IBM-1160	66	-	Host
874	S-20	tis620	66	th_TH.tis620	HP-UX
874	S-20	TIS620-1	66	-	OS/2
874	S-20	TIS620-1	66	-	Windows

Table 94. Turkey, territory identifier: TR

Code page	Group	Code set	Territory code	Locale	Operating system
920	S-9	ISO8859-9	90	tr_TR	AIX
1208	N-1	UTF-8	90	TR_TR	AIX
1026	S-9	IBM-1026	90	-	Host
1155	S-9	IBM-1155	90	-	Host
920	S-9	iso88599	90	tr_TR.iso88599	HP-UX
920	S-9	ISO-8859-9	90	tr_TR	Linux
857	S-9	IBM-857	90	-	OS/2

Table 94. Turkey, territory identifier: TR (continued)

Code page	Group	Code set	Territory code	Locale	Operating system
920	S-9	ISO8859-9	90	tr_TR.ISO8859-9	SCO
1254	S-9	1254	90	-	Windows

Table 95. United Kingdom, territory identifier: GB

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	44	en_GB	AIX
850	S-1	IBM-850	44	En_GB	AIX
923	S-1	ISO8859-15	44	en_GB.8859-15	AIX
1208	N-1	UTF-8	44	EN_GB	AIX
285	S-1	IBM-285	44	-	Host
1146	S-1	IBM-1146	44	-	Host
819	S-1	iso88591	44	en_GB.iso88591	HP-UX
923	S-1	iso885915	44	-	HP-UX
1051	S-1	roman8	44	en_GB.roman8	HP-UX
819	S-1	ISO-8859-1	44	en_GB	Linux
923	S-1	ISO-8859-15	44	-	Linux
437	S-1	IBM-437	44	-	OS/2
850	S-1	IBM-850	44	-	OS/2
819	S-1	ISO8859-1	44	en_GB	SCO
819	S-1	ISO8859-1	44	en	SCO
819	S-1	ISO8859-1	44	en_GB	Solaris
923	S-1	ISO8859-15	44	en_GB.ISO8859-15	Solaris
1252	S-1	1252	44	-	Windows

Table 96. Ukraine, territory identifier: UA

Code page	Group	Code set	Territory code	Locale	Operating system
1124	S-12	IBM-1124	380	Uk_UA	AIX
1208	N-1	UTF-8	380	UK_UA	AIX
1123	S-12	IBM-1123	380	-	Host
1158	S-12	IBM-1158	380	-	Host
1168	S-12	KOI8-U	380	uk_UA.koi8u	Linux
1125	S-12	IBM-1125	380	-	OS/2
1251	S-12	1251	380	-	Windows

Table 97. United States of America, territory identifier: US

Code page	Group	Code set	Territory code	Locale	Operating system
819	S-1	ISO8859-1	1	en_US	AIX
850	S-1	IBM-850	1	En_US	AIX
923	S-1	ISO8859-15	1	en_US.8859-15	AIX
1208	N-1	UTF-8	1	EN_US	AIX
37	S-1	IBM-37	1	-	Host
1140	S-1	IBM-1140	1	-	Host
819	S-1	iso88591	1	en_US.iso88591	HP-UX
923	S-1	iso885915	1	-	HP-UX
1051	S-1	roman8	1	en_US.roman8	HP-UX
819	S-1	ISO-8859-1	1	en_US	Linux

Table 97. United States of America, territory identifier: US (continued)

Code page	Group	Code set	Territory code	Locale	Operating system
923	S-1	ISO-8859-15	1	-	Linux
437	S-1	IBM-437	1	-	OS/2
850	S-1	IBM-850	1	-	OS/2
819	S-1	ISO8859-1	1	en_US	SCO
819	S-1	ISO8859-1	1	en_US	SGI
819	S-1	ISO8859-1	1	en_US	Solaris
923	S-1	ISO8859-15	1	en_US.ISO8859-15	Solaris
1208	N-1	UTF-8	1	en_US.UTF-8	Solaris
1252	S-1	1252	1	-	Windows

Table 98. Vietnam, territory identifier: VN

Code page	Group	Code set	Territory code	Locale	Operating system
1129	S-11	IBM-1129	84	Vi_VN	AIX
1208	N-1	UTF-8	84	VI_VN	AIX
1130	S-11	IBM-1130	84	-	Host
1164	S-11	IBM-1164	84	-	Host
1129	S-11	IBM-1129	84	-	OS/2
1258	S-11	1258	84	-	Windows

Notes:

1. CCSIDs 1392 and 5488 (GB 18030) can only be used with the load or import utilities to move data from CCSIDs 1392 and 5488 to a DB2 Unicode database, or to export from a DB2 Unicode database to CCSIDs 1392 or 5488.
2. On AIX 4.3 or later the code page is 943. If you are using AIX 4.2 or earlier, the code page is 932.
3. Code page 1394 (Shift JIS X0213) can only be used with the load or import utilities to move data from code page 1394 to a DB2 Unicode database, or to export from a DB2 Unicode database to code page 1394.
4. The following map to Arabic Countries/Regions (AA):
 - Arabic (Saudi Arabia)
 - Arabic (Iraq)
 - Arabic (Egypt)
 - Arabic (Libya)
 - Arabic (Algeria)
 - Arabic (Morocco)
 - Arabic (Tunisia)
 - Arabic (Oman)
 - Arabic (Yemen)
 - Arabic (Syria)
 - Arabic (Jordan)
 - Arabic (Lebanon)
 - Arabic (Kuwait)
 - Arabic (United Arab Emirates)
 - Arabic (Bahrain)
 - Arabic (Qatar)

5. The following map to English (US):
 - English (Jamaica)
 - English (Caribbean)
6. The following map to Latin America (Lat):
 - Spanish (Mexican)
 - Spanish (Guatemala)
 - Spanish (Costa Rica)
 - Spanish (Panama)
 - Spanish (Dominican Republic)
 - Spanish (Venezuela)
 - Spanish (Colombia)
 - Spanish (Peru)
 - Spanish (Argentina)
 - Spanish (Ecuador)
 - Spanish (Chile)
 - Spanish (Uruguay)
 - Spanish (Paraguay)
 - Spanish (Bolivia)
7. The following Indic scripts are supported through Unicode: Hindi, Gujarati, Kannada, Konkani, Marathi, Punjabi, Sanskrit, Tamil and Telugu.
8. Code page 950 is also known as Big5. Microsoft code page 950 differs from IBM code page 950 in the following ways:

Range	Description	IBM	Microsoft	Difference
X'8140' to X'8DFE'	User defined characters	User defined area	User defined area	Same
X'8E40' to X'A0FE'	User defined characters	User defined area	User defined area	Same
X'A140' to X'A3BF'	Special symbols	System characters	System characters	Same
X'A3C0' to X'A3E0'	Control symbols	System characters	Empty	Different
X'A3E1' to X'A3FE'	Reserved	Empty	Empty	Same
X'A440' to X'C67E'	Primary use characters	System characters	System characters	Same
X'C6A1' to X'C878'	Eten added symbols	System characters	User defined area	Different
X'C879' to X'C8CC'	Eten added symbols	Empty	User defined area	Different
X'C8CD' to X'C8D3'	Eten added symbols	System characters	User defined area	Different
X'C8D4' to X'C8FD'	Reserved	System characters	User defined area	Different
X'C8FE'	Invalid/undefined character	System characters	User defined area	Different
X'C940' to X'F9D5'	Secondary use characters	System characters	System characters	Same
X'F9D6' to X'F9FE'	Eten extension for Big-5	User defined area	System characters	Different

Range	Description	IBM	Microsoft	Difference
X'FA40' to X'FEFE'	User defined characters	User defined area	User defined area	Same
X'8181' to X'8C82'	User defined characters	User defined area	Empty	Different
X'F286' to X'F9A0'	IBM select characters	System characters	Empty	Different
Total characters		14 060	13 502	
Total user defined characters		6 204	6 217	
Total defined code points		20 264	19 719	

Related tasks:

- “Installing the previous tables for converting between code page 1394 and Unicode” on page 366

Availability of Asian fonts (Linux)

IBM offers additional font packages for Linux that contain additional double-byte character set (DBCS) support for Asian characters. These font packages are necessary with some versions of Linux which install only the fonts required to display the country-specific or region-specific characters.

If you notice missing characters when you use the DB2 Setup wizard or the DB2 GUI tools (post-installation), install the necessary fonts provided with the DB2 product then re-run the **db2setup** command or restart the DB2 GUI tools you were using. The Asian fonts are found in the `java_fonts` directory on the National Language Pack CD-ROM (NLPACK CD) for your Linux operating system.

In this directory, there are two typefaces available: Times New Roman WorldType and Monotype Sans Duospace WorldType. For each typeface, there is a country- or region-specific font. The following table lists the eight fonts provided in compressed format in the `java_fonts` directory.

Font typeface	Font file name	Country/Region
Times New Roman WT J	<code>tnrwt_j.zip</code>	Japan and other countries/regions
Times New Roman WT K	<code>tnrwt_k.zip</code>	Korea
Times New Roman WT SC	<code>tnrwt_s.zip</code>	China (Simplified Chinese)
Times New Roman WT TC	<code>tnrwt_t.zip</code>	Taiwan (Traditional Chinese)
Monotype Sans Duospace WT J	<code>mtsansdj.zip</code>	Japan and other countries/regions
Monotype Sans Duospace WT K	<code>mtsansdk.zip</code>	Korea
Monotype Sans Duospace WT SC	<code>mtsansds.zip</code>	China (Simplified Chinese)
Monotype Sans Duospace WT TC	<code>mtsansdt.zip</code>	Taiwan (Traditional Chinese)

Note: These fonts do not replace the system fonts. These fonts are to be used in conjunction with or for use with DB2. You cannot engage in the general or unrestricted sale or distribution of these fonts.

To install a font:

1. Unzip the font package.
2. Copy the font package to the /opt/jre/lib/fonts directory. You need to create the directory if it does not already exist.
3. Enter the following command: **export JAVA_FONTS=/opt/jre/lib/fonts**

Note: Optionally, you can copy the Asian font package into the java directory in the DB2 installation path. For example, <DB2 installation path>/java/jdk32/jre/lib/fonts, or <DB2 installation path>/java/jdk64/jre/lib/fonts.

As a minimum you need to install one font of each typeface for your country or region. If you are in either the China, Korea, or Taiwan territory, use the territory-specific or region-specific versions; otherwise, use the Japanese version of the fonts. If you have space on your system, it is recommended that you install all eight fonts.

Simplified Chinese locale coding set

IBM AIX and some distributions of Linux have changed the code set bound to the simplified Chinese locale from GBK (code page 1386) to GB18030 (code page 5488 or 1392). For example, the Zh_CN locale on AIX is now bound to the GB18030 code set since:

- AIX Version 5.1.0000.0011
- AIX Version 5.1.0 with maintenance level 2

DB2 database manager supports the GBK code set natively and the GB18030 code set only through Unicode. DB2 database manager will default the locale's code set to ISO 8859-1 (code page 819), and in some operations will also default the locale's territory to the United States (US). To work around this limitation, you have two options:

1. You can override the locale's code set from GB18030 to GBK; and the territory from US to China (whose territory identifier is CN and territory code is 86).
2. You can use a different simplified Chinese locale.

If you choose to use the first option, issue the following commands:

```
db2set DB2CODEPAGE=1386
db2set DB2TERRITORY=86
db2 terminate
db2stop
db2start
```

If you choose to use the second option on AIX, issue either of the following commands:

```
export LANG=zh_CN
export LANG=ZH_CN
```

The code set associated with zh_CN is eucCN (code page 1383), and with ZH_CN is UTF-8 (code page 1208).

If you choose to use the second option on Linux, issue one of the following commands:

```
export LANG=zh_CN.gbk
export LANG=zh_CN
export LANG=zh_CN.utf8
```

The code set associated with zh_CN is eucCN (code page 1383), and with zh_CN.utf8 is UTF-8 (code page 1208).

Displaying Indic characters in the DB2 GUI tools

If you have problems displaying Indic characters when using the DB2 GUI tools on Linux or UNIX operating systems, you might not have the required fonts installed on your system.

DB2 has packaged the following IBM TrueType and OpenType proportional Indic language fonts for your use. You can find these fonts in the `java_fonts` directory on the National Language Pack CD-ROM (NLPACK CD) for the Linux and UNIX operating systems.

These fonts are to be used in conjunction with DB2. You cannot engage in the general or unrestricted sale or distribution of these fonts:

Table 99. Indic fonts packaged with DB2

Typeface	Weight	Font File Name
Devanagari MT for IBM	Medium	devamt.ttf
Devanagari MT for IBM	Bold	devamtb.ttf
Tamil	Medium	TamilMT.ttf
Tamil	Bold	TamilMTB.ttf
Telugu	Medium	TeluguMT.ttf
Telugu	Bold	TeleguMTB.ttf

Detailed instructions on how to install the fonts and modify the `font.properties` file can be found in the Internationalization section of the Java documentation.

In addition, some Microsoft products also come with Indic fonts that can be used with our GUI tools.

Enabling and disabling euro symbol support

DB2 Database for Linux, UNIX, and Windows provides support for the euro currency symbol. The euro symbol has been added to numerous code pages.

Microsoft ANSI code pages have been modified to include the euro currency symbol in position X'80'. Code page 850 has been modified to replace the character DOTLESS I (found at position X'D5') with the euro currency symbol. DB2 internal code page conversion routines use these revised code page definitions as the default to provide euro symbol support.

However, if you want to use the non-euro definitions of the code page conversion tables, follow the procedure below after installation is complete.

Prerequisites:

For replacing existing external code page conversion table files, you may want to back up the current files before copying the non-euro versions over them.

The files are located in the directory `sql1lib/conv/`. On UNIX, `sql1lib/conv/` is linked to the install path of the DB2 database system.

Procedure:

To disable euro-symbol support:

1. Stop the DB2 instance.
2. Download the appropriate conversion table files, in binary:
 - For big-endian platforms from `ftp://ftp.software.ibm.com/ps/products/db2/info/vr8/conv/BigEndian/`. This ftp server is anonymous, so if you are connecting via the command line, log in as user "anonymous" and use your e-mail address as your password. After logging in, change to the conversion tables directory: `cd ps/products/db2/info/vr8/conv/BigEndian/`
 - For little-endian platforms from `ftp://ftp.software.ibm.com/ps/products/db2/info/vr8/conv/LittleEndian/`. This ftp server is anonymous, so if you are connecting via the command line, log in as user "anonymous" and use your e-mail address as your password. After logging in, change to the conversion tables directory: `cd ps/products/db2/info/vr8/conv/LittleEndian`
3. Copy the files to your `sql1lib/conv/` directory.
4. Restart the DB2 instance.

Code pages 819 and 1047:

For code pages 819 (ISO 8859-1 Latin 1 ASCII) and 1047 (Latin 1 Open System EBCDIC), the euro replacement code pages, 923 (ISO 8859-15 Latin 9 ASCII) and 924 (Latin 9 Open System EBCDIC) respectively, contain not just the euro symbol but also several new characters. DB2 Database for Linux, UNIX, and Windows continues to use the old (non-euro) definitions of these two code pages and conversion tables, namely 819 and 1047, by default. There are two ways to activate the new 923/924 code page and the associated conversion tables:

- Create a new database that uses the new code page. For example,
`DB2 CREATE DATABASE dbname USING CODESET ISO8859-15 TERRITORY US`
- Copy the 923 or 924 conversion table files from the `sql1lib/conv/alt/` directory to the `sql1lib/conv/` directory and rename them to 819 or 1047, respectively.

Related concepts:

- "Character conversion" in *SQL Reference, Volume 1*

Related reference:

- "Conversion table files for euro-enabled code pages" on page 339
- "Conversion tables for code pages 923 and 924" on page 343

Character-conversion guidelines

Data conversion might be required to map data between application and database code pages when your application and database do not use the same code page. Because mapping and data conversion require additional overhead application performance improves if the application and database use the same code page or the identity collating sequence.

Character conversion occurs in the following circumstances:

- When a client or application runs in a code page that is different from the code page of the database that it accesses.

The conversion occurs on the database server machine that receives the data. If the database server receives the data, character conversion is from the application code page to the database code page. If the application machine receives the data, conversion is from the database code page to the application code page.

- When a client or application that imports or loads a file runs in a code page different from the file being imported or loaded.

Character conversion does not occur for the following objects:

- File names.
- Data targeted for or coming from a column for which the FOR BIT DATA attribute is assigned, or data that is used in an SQL operation whose result is FOR BIT or BLOB data.
- A DB2 product or platform for which no supported conversion function to or from EUC or UCS-2 is installed. Your application receives an SQLCODE -332 (SQLSTATE 57017) error in this case.

The conversion function and conversion tables or DBCS conversion APIs that the database manager uses when it converts multi-byte code pages depends on the operating system environment.

Note: Character string conversions between multi-byte code pages, such as DBCS with EUC, might increase or decrease length of a string. In addition, code points assigned to different characters in the PC DBCS, EUC, and UCS-2 code sets might produce different results when same characters are sorted.

Extended UNIX Code (EUC) Code Page Support

Host variables that use graphic data in C or C++ applications require special considerations that include special precompiler, application performance, and application design issues.

Many characters in both the Japanese and Traditional Chinese EUC code pages require special methods of managing database and client application support for graphic data, which require double byte characters. Graphic data from these EUC code pages is stored and manipulated using the UCS-2 code set.

Related concepts:

- “Guidelines for analyzing where a federated query is evaluated” in *Performance Guide*

Related reference:

- “Conversion table files for euro-enabled code pages” on page 339
- “Conversion tables for code pages 923 and 924” on page 343

Conversion table files for euro-enabled code pages

The following tables list the conversion tables that have been enhanced to support the euro currency symbol. If you want to disable euro symbol support, download the conversion table file indicated in the column titled “Conversion table file”.

Arabic:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
864, 17248	1046, 9238	08641046.cnv, 10460864.cnv, IBM00864.ucs
864, 17248	1256, 5352	08641256.cnv, 12560864.cnv, IBM00864.ucs
864, 17248	1200, 1208, 13488, 17584	IBM00864.ucs
1046, 9238	864, 17248	10460864.cnv, 08641046.cnv, IBM01046.ucs
1046, 9238	1089	10461089.cnv, 10891046.cnv, IBM01046.ucs
1046, 9238	1256, 5352	10461256.cnv, 12561046.cnv, IBM01046.ucs
1046, 9238	1200, 1208, 13488, 17584	IBM01046.ucs
1089	1046, 9238	10891046.cnv, 10461089.cnv
1256, 5352	864, 17248	12560864.cnv, 08641256.cnv, IBM01256.ucs
1256, 5352	1046, 9238	12561046.cnv, 10461256.cnv, IBM01256.ucs
1256, 5352	1200, 1208, 13488, 17584	IBM01256.ucs

Baltic:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
921, 901	1257	09211257.cnv, 12570921.cnv, IBM00921.ucs
921, 901	1200, 1208, 13488, 17584	IBM00921.ucs
1257, 5353	921, 901	12570921.cnv, 09211257.cnv, IBM01257.ucs
1257, 5353	922, 902	12570922.cnv, 09221257.cnv, IBM01257.ucs
1257, 5353	1200, 1208, 13488, 17584	IBM01257.ucs

Belarus:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
1131, 849	1251, 5347	11311251.cnv, 12511131.cnv
1131, 849	1283	11311283.cnv

Cyrillic:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
855, 872	866, 808	08550866.cnv, 08660855.cnv
855, 872	1251, 5347	08551251.cnv, 12510855.cnv
866, 808	855, 872	08660855.cnv, 08550866.cnv
866, 808	1251, 5347	08661251.cnv, 12510866.cnv
1251, 5347	855, 872	12510855.cnv, 08551251.cnv, IBM01251.ucs
1251, 5347	866, 808	12510866.cnv, 08661251.cnv, IBM01251.ucs
1251, 5347	1124	12511124.cnv, 11241251.cnv, IBM01251.ucs
1251, 5347	1125, 848	12511125.cnv, 11251251.cnv, IBM01251.ucs
1251, 5347	1131, 849	12511131.cnv, 11311251.cnv, IBM01251.ucs
1251, 5347	1200, 1208, 13488, 17584	IBM01251.ucs

Estonia:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
922, 902	1257	09221257.cnv, 12570922.cnv, IBM00922.ucs
922, 902	1200, 1208, 13488, 17584	IBM00922.ucs
1122, 1157	1257, 5353	11221257.cnv

Greek:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
813, 4909	869, 9061	08130869.cnv, 08690813.cnv, IBM00813.ucs
813, 4909	1253, 5349	08131253.cnv, 12530813.cnv, IBM00813.ucs
813, 4909	1200, 1208, 13488, 17584	IBM00813.ucs
869, 9061	813, 4909	08690813.cnv, 08130869.cnv
869, 9061	1253, 5349	08691253.cnv, 12530869.cnv

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
1253, 5349	813, 4909	12530813.cnv, 08131253.cnv, IBM01253.ucs
1253, 5349	869, 9061	12530869.cnv, 08691253.cnv, IBM01253.ucs
1253, 5349	1200, 1208, 13488, 17584	IBM01253.ucs

Hebrew:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
856, 9048	862, 867	08560862.cnv, 08620856.cnv, IBM0856.ucs
856, 9048	916	08560916.cnv, 09160856.cnv, IBM0856.ucs
856, 9048	1255, 5351	08561255.cnv, 12550856.cnv, IBM0856.ucs
856, 9048	1200, 1208, 13488, 17584	IBM0856.ucs
862, 867	856, 9048	08620856.cnv, 08560862.cnv, IBM00862.ucs
862, 867	916	08620916.cnv, 09160862.cnv, IBM00862.ucs
862, 867	1255, 5351	08621255.cnv, 12550862.cnv, IBM00862.ucs
862, 867	1200, 1208, 13488, 17584	IBM00862.ucs
916	856, 9048	09160856.cnv, 08560916.cnv
916	862, 867	09160862.cnv, 08620916.cnv
1255, 5351	856, 9048	12550856.cnv, 08561255.cnv, IBM01255.ucs
1255, 5351	862, 867	12550862.cnv, 08621255.cnv, IBM01255.ucs
1255, 5351	1200, 1208, 13488, 17584	IBM01255.ucs

Latin-1:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
437	850, 858	04370850.cnv, 08500437.cnv
500, 1148	437	05000437.cnv, IBM00500.ucs
850, 858	437	08500437.cnv, 04370850.cnv
850, 858	860	08500860.cnv, 08600850.cnv
850, 858	1114, 5210	08501114.cnv, 11140850.cnv
850, 858	1275	08501275.cnv, 12750850.cnv
860	850, 858	08600850.cnv, 08500860.cnv
1275	850, 858	12750850.cnv, 08501275.cnv

Latin-2:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
852, 9044	1250, 5346	08521250.cnv, 12500852.cnv
1250, 5346	852, 9044	12500852.cnv, 08521250.cnv, IBM01250.ucs
1250, 5346	1200, 1208, 13488, 17584	IBM01250.ucs

Simplified Chinese:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
837, 935, 1388	1200, 1208, 13488, 17584	1388ucs2.cnv
1386	1200, 1208, 13488, 17584	1386ucs2.cnv, ucs21386.cnv

Traditional Chinese:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
937, 835, 1371	950, 1370	09370950.cnv, 0937ucs2.cnv
937, 835, 1371	1200, 1208, 13488, 17584	0937ucs2.cnv
1114, 5210	850, 858	11140850.cnv, 08501114.cnv

Thailand:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
874, 1161	1200, 1208, 13488, 17584	IBM00874.ucs

Turkish:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
857, 9049	1254, 5350	08571254.cnv, 12540857.cnv
1254, 5350	857, 9049	12540857.cnv, 08571254.cnv, IBM01254.ucs
1254, 5350	1200, 1208, 13488, 17584	IBM01254.ucs

Ukraine:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
1124	1251, 5347	11241251.cnv, 12511124.cnv
1125, 848	1251, 5347	11251251.cnv, 12511125.cnv

Unicode:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
1200, 1208, 13488, 17584	813, 4909	IBM00813.ucs
1200, 1208, 13488, 17584	862, 867	IBM00862.ucs
1200, 1208, 13488, 17584	864, 17248	IBM00864.ucs
1200, 1208, 13488, 17584	874, 1161	IBM00874.ucs
1200, 1208, 13488, 17584	921, 901	IBM00921.ucs
1200, 1208, 13488, 17584	922, 902	IBM00922.ucs
1200, 1208, 13488, 17584	1046, 9238	IBM01046.ucs
1200, 1208, 13488, 17584	1250, 5346	IBM01250.ucs
1200, 1208, 13488, 17584	1251, 5347	IBM01251.ucs
1200, 1208, 13488, 17584	1253, 5349	IBM01253.ucs
1200, 1208, 13488, 17584	1254, 5350	IBM01254.ucs
1200, 1208, 13488, 17584	1255, 5351	IBM01255.ucs
1200, 1208, 13488, 17584	1256, 5352	IBM01256.ucs
1200, 1208, 13488, 17584	1386	ucs21386.cnv, 1386ucs2.cnv

Vietnamese:

Database server CCSIDs/CPGIDs	Database client CCSIDs/CPGIDs	Conversion table files
1258, 5354	1129, 1163	12581129.cnv

Related concepts:

- “Character conversion” in *SQL Reference, Volume 1*

Related tasks:

- “Enabling and disabling euro symbol support” on page 336

Conversion tables for code pages 923 and 924

The following is a list of all the code page conversion table files that are associated with code pages 923 and 924. Each file is of the form XXXXYYYY.cnv or ibmZZZZZ.ucs, where XXXXX is the source code page number and YYYYY is the target code page number. The file ibmZZZZZ.ucs supports conversion between code page ZZZZZ and Unicode.

To activate a particular code page conversion table, copy the conversion table file from the `sqllib/conv/alt/` directory to the `sqllib/conv/` directory and rename that conversion table file as shown in the second column.

For example, to support the euro symbol when connecting a 8859-1/15 (Latin 1/9) client to a Windows 1252 database, you need to copy and rename the following code page conversion table files:

- `sqllib/conv/alt/09231252.cnv` to `sqllib/conv/08191252.cnv`
- `sqllib/conv/alt/12520923.cnv` to `sqllib/conv/12520819.cnv`

- sqllib/conv/alt/ibm00923.ucs to sqllib/conv/ibm00819.ucs

923 and 924 conversion table files in the sqllib/conv/alt/ directory	New name in the sqllib/conv/ directory
04370923.cnv	04370819.cnv
08500923.cnv	08500819.cnv
08600923.cnv	08600819.cnv
08630923.cnv	08630819.cnv
09230437.cnv	08190437.cnv
09230850.cnv	08190850.cnv
09230860.cnv	08190860.cnv
09231043.cnv	08191043.cnv
09231051.cnv	08191051.cnv
09231114.cnv	08191114.cnv
09231252.cnv	08191252.cnv
09231275.cnv	08191275.cnv
09241252.cnv	10471252.cnv
10430923.cnv	10430819.cnv
10510923.cnv	10510819.cnv
11140923.cnv	11140819.cnv
12520923.cnv	12520819.cnv
12750923.cnv	12750819.cnv
ibm00923.ucs	ibm00819.ucs

Related concepts:

- “Character conversion” in *SQL Reference, Volume 1*

Related tasks:

- “Enabling and disabling euro symbol support” on page 336

Choosing a language for your database

When you create a database, you have to decide what language your data will be stored in. When you create a database, you can specify the territory and code set. The territory and code set may be different from the current operating system settings. If you do not explicitly choose a territory and code set at database creation time, the database will be created using the current locale. When you are choosing a code set, make sure it can encode all the characters in the language you will be using.

Another option is to store data in a Unicode database, which means that you do not have to choose a specific language; Unicode encoding includes characters from almost all of the living languages in the world.

Locale setting for the DB2 Administration Server

Ensure that the locale of the DB2 Administration Server instance is compatible with the locale of the DB2 instance. Otherwise, the DB2 instance cannot communicate with the DB2 Administration Server.

If the LANG environment variable is not set in the user profile of the DB2 Administration Server, the DB2 Administration Server will be started with the default system locale. If the default system locale is not defined, the DB2 Administration Server will be started with code page 819. If the DB2 instance uses one of the DBCS locales, and the DB2 Administration Server is started with code page 819, the instance will not be able to communicate with the DB2 Administration Server. The locale of the DB2 Administration Server and the locale of the DB2 instance must be compatible.

For example, on a Simplified Chinese Linux system, LANG=zh_CN should be set in the DB2 Administration Server's user profile.

Related tasks:

- “Changing the DB2 interface language (Linux and UNIX)” in *Quick Beginnings for DB2 Servers*
- “Changing the DB2 interface language (Windows)” in *Quick Beginnings for DB2 Servers*

Enabling bidirectional support

Bidirectional layout transformations are implemented in DB2 Database for Linux, UNIX, and Windows using the new Coded Character Set Identifier (CCSID) definitions. For the new bidirectional-specific CCSIDs, layout transformations are performed instead of, or in addition to, code page conversions. To use this support, the DB2BIDI registry variable must be set to YES. By default, this variable is not set. It is used by the server for all conversions, and can only be set when the server is started. Setting DB2BIDI to YES may have some performance impact because of additional checking and layout transformations.

Restrictions:

The following restrictions apply:

- If you select a CCSID that is not appropriate for the code page or string type of your client platform, you may get unexpected results. If you select an incompatible CCSID (for example, the Latin-1 CCSID for connection to an Arabic database), or if DB2BIDI has not been set for the server, you will receive an error message when you try to connect.
- The DB2 Command Line Processor on the Windows operating system does not have bidirectional support.
- CCSID override is not supported for cases where the HOST EBCDIC platform is the client, and DB2 Database is the server.

When converting from one Arabic CCSID to another Arabic CCSID, DB2 employs the following logic to deshape (or expand) the lam-alef ligature. Deshaping will occur when the Text Shaping attribute of the source Arabic CCSID is shaped but the Text Shaping attribute of the target Arabic CCSID is unshaped.

The logic to deshape the lam-alef ligature is:

1. If the *last* character of the data stream is a blank character, then every character after the lam-alef ligature will be shifted to the end of the data stream, therefore making available an empty position for the current lam-alef ligature to be deshaped (expanded) into its two constituent characters: lam and alef.
2. Otherwise, if the *first* character of the data stream is a blank character, then every character before the lam-alef ligature will be shifted to the beginning of the data stream, therefore making available an empty position for the current lam-alef ligature to be deshaped (expanded) into its two constituent characters: lam and alef.
3. Otherwise, there is no blank character at the beginning and end of the data stream, and the lam-alef ligature cannot be deshaped. If the target CCSID does have the lam-alef ligature, then the lam-alef ligature remains as is; otherwise, the lam-alef ligature is replaced by the target CCSID's SUBstitution character.

Conversely when converting from an Arabic CCSID whose Text Shaping attribute is unshaped to an Arabic CCSID whose Text Shaping attribute is shaped, the source lam and alef characters will be contracted to one ligature character, and a blank character is inserted at the end of the target area data stream.

Procedure:

To specify a particular bidirectional CCSID in a non-DRDA environment:

- Ensure the DB2BIDI registry variable is set to YES.
- Select the CCSID that matches the characteristics of your client, and set DB2CODEPAGE to that value.
- If you already have a connection to the database, you must issue a TERMINATE command, and then reconnect to allow the new setting for DB2CODEPAGE to take effect.

For DRDA environments, if the HOST EBCDIC platform also supports these bidirectional CCSIDs, you only need to set the DB2CODEPAGE value. Note that you must not further specify the same CCSID on the BIDI parameter in the PARMS field of the DCS database directory entry for the server database, otherwise an extra bidi layout conversion would occur, and render any Arabic data to be incorrectly reversed. However, if the HOST platform does not support these CCSIDs, you must also specify a CCSID override for the HOST database server to which you are connecting. This is accomplished through the use of the BIDI parameter in the PARMS field of the DCS database directory entry for the server database. The override is necessary because, in a DRDA environment, code page conversions and layout transformations are performed by the receiver of data. However, if the HOST server does not support these bidirectional CCSIDs, it does not perform layout transformation on the data that it receives from DB2. If you use a CCSID override, the DB2 client performs layout transformation on the outbound data as well.

Related concepts:

- "Bidirectional support with DB2 Connect" on page 349
- "Handling BiDi data" in *DB2 Connect User's Guide*

Related reference:

- "Bidirectional-specific CCSIDs" on page 347
- "General registry variables" in *Administration Guide: Implementation*

Bidirectional-specific CCSIDs

The following bidirectional attributes are required for correct handling of bidirectional data on different platforms:

- Text type
- Numeric shaping
- Orientation
- Text shaping
- Symmetric swapping

Because default values on different platforms are not the same, problems can occur when DB2 data is moved from one platform to another. For example, the Windows operating system uses LOGICAL UNSHAPED data, while z/OS and OS/390 usually use SHAPED VISUAL data. Therefore, without support for bidirectional attributes, data sent from DB2 Universal Database for z/OS and OS/390 to DB2 on Windows 32-bit operating systems may display incorrectly.

DB2 Database for Linux, UNIX, and Windows supports bidirectional data attributes through special bidirectional Coded Character Set Identifiers (CCSIDs). The following bidirectional CCSIDs have been defined and are implemented with DB2 as shown in Table 100. CDRA string types are defined as shown in Table 101 on page 349.

Table 100. Bidirectional CCSIDs

CCSID	Code Page	String Type
420	420	4
424	424	4
856	856	5
862	862	4
864	864	5
867	862	4
916	916	5
1046	1046	5
1089	1089	5
1200	1200	10
1208	1208	10
1255	1255	5
1256	1256	5
5351	1255	5
5352	1256	5
8612	420	5
8616	424	10
9048	856	5
9238	1046	5
12712	424	4
13488	13488	10

Table 100. Bidirectional CCSIDs (continued)

CCSID	Code Page	String Type
16804	420	4
17248	864	5
62208	856	4
62209	862	10
62210	916	4
62211	424	5
62213	862	5
62215	1255	4
62218	864	4
62220	856	6
62221	862	6
62222	916	6
62223	1255	6
62224	420	6
62225	864	6
62226	1046	6
62227	1089	6
62228	1256	6
62229	424	8
62230	856	8
62231	862	8
62232	916	8
62233	420	8
62234	420	9
62235	424	6
62236	856	10
62237	1255	8
62238	916	10
62239	1255	10
62240	424	11
62241	856	11
62242	862	11
62243	916	11
62244	1255	11
62245	424	10
62246	1046	8
62247	1046	9
62248	1046	4
62249	1046	12
62250	420	12

Table 101. CDRA string types

String type	Text type	Numeric shaping	Orientation	Text shaping	Symmetrical swapping
4	Visual	Passthrough	LTR	Shaped	Off
5	Implicit	Arabic	LTR	Unshaped	On
6	Implicit	Arabic	RTL	Unshaped	On
7*	Visual	Passthrough	Contextual*	Unshaped ligature	Off
8	Visual	Passthrough	RTL	Shaped	Off
9	Visual	Passthrough	RTL	Shaped	On
10	Implicit	Arabic	Contextual LTR	Unshaped	On
11	Implicit	Arabic	Contextual RTL	Unshaped	On
12	Implicit	Arabic	RTL	Shaped	Off

Note: * String orientation is left-to-right (LTR) when the first alphabetic character is a Latin character, and right-to-left (RTL) when it is an Arabic or Hebrew character. Characters are unshaped, but LamAlef ligatures are kept and are not broken into constituents.

Related concepts:

- “Bidirectional support with DB2 Connect” on page 349

Related tasks:

- “Enabling bidirectional support” on page 345

Bidirectional support with DB2 Connect

When data is exchanged between DB2 Connect and a database on the server, it is usually the receiver that performs conversion on the incoming data. The same convention would normally apply to bidirectional layout transformations, and is in addition to the usual code page conversion. DB2 Connect has the optional ability to perform bidirectional layout transformation on data it is about to send to the server database, in addition to data received from the server database.

In order for DB2 Connect to perform bidirectional layout transformation on outgoing data for a server database, the bidirectional CCSID of the server database must be overridden. This is accomplished through the use of the BIDI parameter in the PARMS field of the DCS database directory entry for the server database.

Note: If you want DB2 Connect to perform layout transformation on the data it is about to send to the DB2 host or iSeries database, even though you do not have to override its CCSID, you must still add the BIDI parameter to the PARMS field of the DCS database directory. In this case, the CCSID that you should provide is the default DB2 host or iSeries database CCSID.

The BIDI parameter is to be specified as the ninth parameter in the PARMS field, along with the bidirectional CCSID with which you want to override the default server database bidirectional CCSID:

```
",,,,,,,,,BIDI=xyz"
```

where *xyz* is the CCSID override.

Note: The registry variable DB2BIDI must be set to YES for the BIDI parameter to take effect.

The use of this feature is best described with an example.

Suppose you have a Hebrew DB2 client running CCSID 62213 (bidirectional string type 5), and you want to access a DB2 host or iSeries database running CCSID 00424 (bidirectional string type 4). However, you know that the data contained in the DB2 host or iSeries database is based on CCSID 08616 (bidirectional string type 6).

There are two problems here: The first is that the DB2 host or iSeries database does not know the difference in the bidirectional string types with CCSIDs 00424 and 08616. The second problem is that the DB2 host or iSeries database does not recognize the DB2 client CCSID (62213). It only supports CCSID 00862, which is based on the same code page as CCSID 62213.

You will need to ensure that data sent to the DB2 host or iSeries database is in bidirectional string type 6 format to begin with, and also let DB2 Connect know that it has to perform bidirectional transformation on data it receives from the DB2 host or iSeries database. You will need to use following catalog command for the DB2 host or iSeries database:

```
db2 catalog dcs database nydb1 as telaviv parms ",,,,,,,,,BIDI=08616"
```

This command tells DB2 Connect to override the DB2 host or iSeries database CCSID of 00424 with 08616. This override includes the following processing:

1. DB2 Connect connects to the DB2 host or iSeries database using CCSID 00862.
2. DB2 Connect performs bidirectional layout transformation on the data it is about to *send* to the DB2 host or iSeries database. The transformation is from CCSID 62213 (bidirectional string type 5) to CCSID 62221 (bidirectional string type 6).
3. DB2 Connect performs bidirectional layout transformation on data it *receives* from the DB2 host or iSeries database. This transformation is from CCSID 08616 (bidirectional string type 6) to CCSID 62213 (bidirectional string type 5).

Note: In some cases, use of a bidirectional CCSID may cause the SQL query itself to be modified in such a way that it is not recognized by the DB2 server. Specifically, you should avoid using IMPLICIT CONTEXTUAL and IMPLICIT RIGHT-TO-LEFT CCSIDs when a different string type can be used. CONTEXTUAL CCSIDs can produce unpredictable results if the SQL query contains quoted strings. Avoid using quoted strings in SQL statements; use host variables whenever possible.

If a specific bidirectional CCSID is causing problems that cannot be rectified by following these recommendations, set DB2BIDI to NO.

Related concepts:

- “Handling BiDi data” in *DB2 Connect User’s Guide*

Related reference:

- “Bidirectional-specific CCSIDs” on page 347

Collating sequences

The database manager compares character data using a *collating sequence*. This is an ordering for a set of characters that determines whether a particular character sorts higher, lower, or the same as another.

Note: Character string data defined with the FOR BIT DATA attribute, and BLOB data, is sorted using the binary sort sequence.

For example, a collating sequence can be used to indicate that lowercase and uppercase versions of a particular character are to be sorted equally.

The database manager allows databases to be created with custom collating sequences. The following sections help you determine and implement a particular collating sequence for a database.

The database manager allows databases to be created with custom collating sequences. For Unicode databases, the various collating sequences supported are described in the “Unicode implementation in the DB2 database” topic. The following sections help you determine and implement a particular collating sequence for a database.

Each single-byte character in a database is represented internally as a unique number between 0 and 255 (in hexadecimal notation, between X'00' and X'FF'). This number is referred to as the *code point* of the character; the assignment of numbers to characters in a set is collectively called a *code page*. A collating sequence is a mapping between the code point and the desired position of each character in a sorted sequence. The numeric value of the position is called the *weight* of the character in the collating sequence. In the simplest collating sequence, the weights are identical to the code points. This is called the *identity sequence*.

For example, suppose the characters B and b have the code points X'42' and X'62', respectively. If (according to the collating sequence table) they both have a sort weight of X'42' (B), they collate the same. If the sort weight for B is X'9E', and the sort weight for b is X'9D', b will be sorted before B. The collation sequence table specifies the weight of each character. The table is different from a code page, which specifies the code point of each character.

Consider the following example. The ASCII characters A through Z are represented by X'41' through X'5A'. To describe a collating sequence in which these characters are sorted consecutively (no intervening characters), you can write: X'41', X'42', ... X'59', X'5A'.

The hexadecimal value of a multi-byte character is also used as the weight. For example, suppose the code points for the double-byte characters A and B are X'8260' and X'8261' respectively, then the collation weights for X'82', X'60', and X'61' are used to sort these two characters according to their code points.

The weights in a collating sequence need not be unique. For example, you could give uppercase letters and their lowercase equivalents the same weight.

Specifying a collating sequence can be simplified if the collating sequence provides weights for all 256 code points. The weight of each character can be determined using the code point of the character.

In all cases, the DB2 database uses the collation table that was specified at database creation time. If you want the multi-byte characters to be sorted the way that they appear in their code point table, you must specify `IDENTITY` as the collation sequence when you create the database.

Note: For Unicode databases, the various collating sequences supported are described in the “Unicode implementation in the DB2 database” topic.

Once a collating sequence is defined, all future character comparisons for that database will be performed with that collating sequence. Except for character data defined as `FOR BIT DATA` or `BLOB` data, the collating sequence will be used for all SQL comparisons and `ORDER BY` clauses, and also in setting up indexes and statistics.

Potential problems can occur in the following cases:

- An application merges sorted data from a database with application data that was sorted using a different collating sequence.
- An application merges sorted data from one database with sorted data from another, but the databases have different collating sequences.
- An application makes assumptions about sorted data that are not true for the relevant collating sequence. For example, numbers collating lower than alphabets might or might not be true for a particular collating sequence.

A final point to remember is that the results of any sort based on a direct comparison of character code points will only match query results that are ordered using an identity collating sequence.

Related concepts:

- “Character comparisons based on collating sequences” in *Developing SQL and External Routines*
- “Character conversion” in *SQL Reference, Volume 1*
- “Unicode implementation in DB2 Database for Linux, UNIX, and Windows” on page 357

Collating Thai characters

Thai contains special vowels (“leading vowels”), tonal marks and other special characters that are not sorted sequentially.

Restrictions:

You must either create your database with a Thai locale and code set, or create a Unicode database.

Procedure:

When you create a database using Thai and corresponding code set, use the `COLLATE USING NLSCHAR` clause on the `CREATE DATABASE` command. When you create a Unicode database, use the `COLLATE USING UCA400_LTH` clause on the `CREATE DATABASE` command.

Related concepts:

- “Collating sequences” on page 351

Related reference:

- “CREATE DATABASE command” in *Command Reference*

Date and time formats by territory code

The character string representation of date and time formats is the default format of datetime values associated with the territory code of the application. This default format can be overridden by specifying the DATETIME format option when the program is precompiled or bound to the database.

Following is a description of the input and output formats for date and time:

- Input Time Format
 - There is no default input time format
 - All time formats are allowed as input for all territory codes.
- Output Time Format
 - The default output time format is equal to the local time format.
- Input Date Format
 - There is no default input date format
 - Where the local format for date conflicts with an ISO, JIS, EUR, or USA date format, the local format is recognized for date input. For example, see the UK entry in Table 102.
- Output Date Format
 - The default output date format is shown in Table 102.

Note: Table 102 also shows a listing of the string formats for the various territory codes.

Table 102. Date and Time Formats by Territory Code

Territory Code	Local Date Format	Local Time Format	Default Output Date Format	Input Date Formats
355 Albania	yyyy-mm-dd	JIS	LOC	LOC, USA, EUR, ISO
785 Arabic	dd/mm/yyyy	JIS	LOC	LOC, EUR, ISO
001 Australia (1)	mm-dd-yyyy	JIS	LOC	LOC, USA, EUR, ISO
061 Australia	dd-mm-yyyy	JIS	LOC	LOC, USA, EUR, ISO
032 Belgium	dd/mm/yyyy	JIS	LOC	LOC, EUR, ISO
055 Brazil	dd.mm.yyyy	JIS	LOC	LOC, EUR, ISO
359 Bulgaria	dd.mm.yyyy	JIS	EUR	LOC, USA, EUR, ISO
001 Canada	mm-dd-yyyy	JIS	USA	LOC, USA, EUR, ISO
002 Canada (French)	dd-mm-yyyy	ISO	ISO	LOC, USA, EUR, ISO

Table 102. Date and Time Formats by Territory Code (continued)

Territory Code	Local Date Format	Local Time Format	Default Output Date Format	Input Date Formats
385 Croatia	yyyy-mm-dd	JIS	ISO	LOC, USA, EUR, ISO
042 Czech Republic	yyyy-mm-dd	JIS	ISO	LOC, USA, EUR, ISO
045 Denmark	dd-mm-yyyy	ISO	ISO	LOC, USA, EUR, ISO
358 Finland	dd/mm/yyyy	ISO	EUR	LOC, EUR, ISO
389 FYR Macedonia	dd.mm.yyyy	JIS	EUR	LOC, USA, EUR, ISO
033 France	dd/mm/yyyy	JIS	EUR	LOC, EUR, ISO
049 Germany	dd/mm/yyyy	ISO	ISO	LOC, EUR, ISO
030 Greece	dd/mm/yyyy	JIS	LOC	LOC, EUR, ISO
036 Hungary	yyyy-mm-dd	JIS	ISO	LOC, USA, EUR, ISO
354 Iceland	dd-mm-yyyy	JIS	LOC	LOC, USA, EUR, ISO
091 India	dd/mm/yyyy	JIS	LOC	LOC, EUR, ISO
972 Israel	dd/mm/yyyy	JIS	LOC	LOC, EUR, ISO
039 Italy	dd/mm/yyyy	JIS	LOC	LOC, EUR, ISO
081 Japan	mm/dd/yyyy	JIS	ISO	LOC, USA, EUR, ISO
082 Korea	mm/dd/yyyy	JIS	ISO	LOC, USA, EUR, ISO
001 Latin America (1)	mm-dd-yyyy	JIS	LOC	LOC, USA, EUR, ISO
003 Latin America	dd-mm-yyyy	JIS	LOC	LOC, EUR, ISO
031 Netherlands	dd-mm-yyyy	JIS	LOC	LOC, USA, EUR, ISO
047 Norway	dd/mm/yyyy	ISO	EUR	LOC, EUR, ISO
048 Poland	yyyy-mm-dd	JIS	ISO	LOC, USA, EUR, ISO
351 Portugal	dd/mm/yyyy	JIS	LOC	LOC, EUR, ISO
086 China	mm/dd/yyyy	JIS	ISO	LOC, USA, EUR, ISO
040 Romania	yyyy-mm-dd	JIS	ISO	LOC, USA, EUR, ISO
007 Russia	dd/mm/yyyy	ISO	LOC	LOC, EUR, ISO
381 Serbia/Montenegro	yyyy-mm-dd	JIS	ISO	LOC, USA, EUR, ISO
042 Slovakia	yyyy-mm-dd	JIS	ISO	LOC, USA, EUR, ISO
386 Slovenia	yyyy-mm-dd	JIS	ISO	LOC, USA, EUR, ISO

Table 102. Date and Time Formats by Territory Code (continued)

Territory Code	Local Date Format	Local Time Format	Default Output Date Format	Input Date Formats
034 Spain	dd/mm/yyyy	JIS	LOC	LOC, EUR, ISO
046 Sweden	dd/mm/yyyy	ISO	ISO	LOC, EUR, ISO
041 Switzerland	dd/mm/yyyy	ISO	EUR	LOC, EUR, ISO
088 Taiwan	mm-dd-yyyy	JIS	ISO	LOC, USA, EUR, ISO
066 Thailand (2)	dd/mm/yyyy	JIS	LOC	LOC, EUR, ISO
090 Turkey	dd/mm/yyyy	JIS	LOC	LOC, EUR, ISO
044 UK	dd/mm/yyyy	JIS	LOC	LOC, EUR, ISO
001 USA	mm-dd-yyyy	JIS	USA	LOC, USA, EUR, ISO
084 Vietnam	dd/mm/yyyy	JIS	LOC	LOC, EUR, ISO
Notes:				
1. Countries/Regions using the default C locale are assigned territory code 001.				
2. yyyy in Buddhist era is equivalent to Gregorian + 543 years (Thailand only).				

Related reference:

- “BIND command” in *Command Reference*
- “PRECOMPILE command” in *Command Reference*

Unicode character encoding

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.

Information on Unicode can be found in the latest edition of The Unicode Standard book, and from The Unicode Consortium web site (www.unicode.org).

Unicode uses two encoding forms: 8-bit and 16-bit. The default encoding form is 16-bit, that is, each character is 16 bits (two bytes) wide, and is usually shown as U+hhhh, where hhhh is the hexadecimal code point of the character. While the resulting 65 000+ code elements are sufficient for encoding most of the characters of the major languages of the world, the Unicode standard also provides an extension mechanism that allows the encoding of as many as one million more characters. The extension mechanism uses a pair of high and low surrogate characters to encode one extended or supplementary character. The first (or high) surrogate character has a code value between U+D800 and U+DBFF, and the second (or low) surrogate character has a code value between U+DC00 and U+DFFF.

UCS-2

The International Standards Organization (ISO) and the International Electrotechnical Commission (IEC) standard 10646 (ISO/IEC 10646) specifies the Universal Multiple-Octet Coded Character Set (UCS) that has a 16-bit (two-byte) version (UCS-2) and a 32-bit (four-byte) version (UCS-4). UCS-2 is identical to the Unicode 16-bit form without surrogates. UCS-2 can encode all the (16-bit) characters defined in the Unicode version 3.0 repertoire. Two UCS-2 characters — a

high followed by a low surrogate — are required to encode each of the new supplementary characters introduced starting in Unicode version 3.1. These supplementary characters are defined outside the original 16-bit Basic Multilingual Plane (BMP or Plane 0).

UTF-8

Sixteen-bit Unicode characters pose a major problem for byte-oriented ASCII-based applications and file systems. For example, non-Unicode aware applications may misinterpret the leading 8 zero bits of the uppercase character 'A' (U+0041) as the single-byte ASCII NULL character.

UTF-8 (UCS Transformation Format 8) is an algorithmic transformation that transforms fixed-length Unicode characters into variable-length ASCII-safe byte strings. In UTF-8, ASCII and control characters are represented by their usual single-byte codes, and other characters become two or more bytes long. UTF-8 can encode both non-supplementary and supplementary characters.

UTF-16

ISO/IEC 10646 also defines an extension technique for encoding some UCS-4 characters using two UCS-2 characters. This extension, called UTF-16, is identical to the Unicode 16-bit encoding form with surrogates. In summary, the UTF-16 character repertoire consists of all the UCS-2 characters plus the additional one million characters accessible via the surrogate pairs.

When serializing 16-bit Unicode characters into bytes, some processors place the most significant byte in the initial position (known as big-endian order), while others place the least significant byte first (known as little-endian order). The default byte ordering for Unicode is big-endian.

The number of bytes for each UTF-16 character in UTF-8 format can be determined from Table 103.

Table 103. UTF-8 Bit Distribution

Code Value (binary)	UTF-16 (binary)	1st byte (binary)	2nd byte (binary)	3rd byte (binary)	4th byte (binary)
00000000 0xxxxxxx	00000000 0xxxxxxx	0xxxxxxx			
00000yyy yyxxxxxx	00000yyy yyxxxxxx	110yyyyy	10xxxxxx		
zzzzyyyy yyxxxxxx	zzzzyyyy yyxxxxxx	1110zzzz	10yyyyyy	10xxxxxx	
uuuuu zzzzyyyy yyxxxxxx	110110ww wwzzzzyy 11011yy yyxxxxxx	11110uuu (where uuuuu = wwww+1)	10uuzzzz	10yyyyyy	10xxxxxx

In each of the above, the series of u's, w's, x's, y's, and z's is the bit representation of the character. For example, U+0080 transforms into 11000010 10000000 in binary, and the surrogate character pair U+D800 U+DC00 becomes 11110000 10010000 10000000 10000000 in binary.

Related concepts:

- “Unicode handling of data types” on page 360
- “Unicode implementation in DB2 Database for Linux, UNIX, and Windows” on page 357
- “Unicode literals” on page 364

Related tasks:

- “Creating a Unicode database” on page 362

Unicode implementation in DB2 Database for Linux, UNIX, and Windows

DB2 Database for Linux, UNIX, and Windows supports UTF-8 and UCS-2.

When a Unicode database is created, CHAR, VARCHAR, LONG VARCHAR, and CLOB data are stored in UTF-8 form, and GRAPHIC, VARGRAPHIC, LONG VARGRAPHIC, and DBCLOB data are stored in UCS-2 big-endian form.

In versions of DB2 products prior to Version 7.2 FixPak 4, DB2 treats the two characters in a surrogate pair as two independent Unicode characters. Therefore transforming the pair from UTF-16/UCS-2 to UTF-8 results in two three-byte sequences. Starting in DB2 Universal Database Version 7.2 FixPak 4, DB2 recognizes surrogate pairs when transforming between UTF-16/UCS-2 and UTF-8, thus a pair of UTF-16 surrogates will become one UTF-8 four-byte sequence. In other usages, DB2 continues to treat a surrogate pair as two independent UCS-2 characters. You can safely store supplementary characters in DB2 Unicode databases, provided you know how to distinguish them from the non-supplementary characters.

DB2 treats each Unicode character, including those (non-spacing) characters such as the COMBINING ACUTE ACCENT character (U+0301), as an individual character. Therefore DB2 would not recognize that the character LATIN SMALL LETTER A WITH ACUTE (U+00E1) is canonically equivalent to the character LATIN SMALL LETTER A (U+0061) followed by the character COMBINING ACUTE ACCENT (U+0301).

The default collating sequence for a UCS-2 Unicode database is IDENTITY, which orders the characters by their code points. Therefore, by default, all Unicode characters are ordered and compared according to their code points. For non-supplementary Unicode characters, their binary collation orders when encoded in UTF-8 and UCS-2 are the same. But if you have any supplementary character that requires a pair of surrogate characters to encode, then in UTF-8 encoding the character will be collated towards the end, but in UCS-2 encoding the same character will be collated somewhere in the middle, and its two surrogate characters can be separated. The reason is the extended character, when encoded in UTF-8, has a four-byte binary code value of 11110xxx 10xxxxxx 10xxxxxx 10xxxxxx, which is greater than the UTF-8 encoding of U+FFFF, namely X'EFBFBF'. But in UCS-2, the same supplementary character is encoded as a pair of UCS-2 high and

low surrogate characters, and has the binary form of 1101 1000 xxxx xxxx 1101 1100 xxxx xxxx, which is less than the UCS-2 encoding of U+FFFF.

A Unicode database can also be created with the `IDENTITY_16BIT` collation option. The `IDENTITY_16BIT` collator 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 (www.unicode.org). CESU-8 is binary identical to UTF-8 except for the Unicode supplementary characters, that is, those characters that are defined outside the 16-bit Basic Multilingual Plane (BMP or Plane 0). In UTF-8 encoding, a supplementary character is represented by one four-byte sequence, but the same character in CESU-8 requires two three-byte sequences. Using the `IDENTITY_16BIT` collation option will yield the same collation order for both character and graphic data.

DB2 UDB Version 8.2 supports three new collation sequence keywords for Unicode databases: `UCA400_NO`, `UCA400_LSK`, and `UCA400_LTH`. The `UCA400_NO` collators implements the UCA (Unicode Collation Algorithm) based on the Unicode Standard version 4.00 with normalization implicitly set to on. The `UCA400_LSK` and `UCA400_LTH` collator also implement the UCA version 4.00. `UCA400_LSK` will sort all Slovakian characters in the appropriate order, and `UCA400_LTH` will sort all Thai characters as per the Royal Thai Dictionary order. Details of the UCA can be found in the Unicode Technical Standard #10 available at the Unicode Consortium web site (www.unicode.org).

All culturally sensitive parameters, such as date or time format, decimal separator, and others, are based on the current territory of the client.

A Unicode database allows connection from every code page supported by DB2. The database manager automatically performs code page conversion for character and graphic strings between the client's code page and Unicode.

Every client is limited by the character repertoire, the input method, and the fonts supported by its environment, but the UCS-2 database itself accepts and stores all UCS-2 characters. Therefore, every client usually works with a subset of UCS-2 characters, but the database manager allows the entire repertoire of UCS-2 characters.

When characters are converted from a local code page to Unicode, there may be expansion in the number of bytes. Prior to Version 8, based on the semantics of SQL statements, character data may have been marked as being encoded in the client's code page, and the database server would have manipulated the entire statement in the client's code page. This manipulation could have resulted in potential expansion of the data. Starting in Version 8, once an SQL statement enters the database server, it operates only on the database server's code page. In this case there is no size change. However, specifying string units for some string functions might result in internal codepage conversions. If this occurs, the size of the data string might change.

AIX, UNIX, and Linux distributions and code pages

Newer versions of AIX, some UNIX platforms, and many Linux distributions use Unicode (UTF-8) as the default code page instead of traditional non-Unicode code pages. If the operating system is upgraded on a system and the upgrade includes this change in the default code page, then:

- Applications that used to run may fail because the default active code page is modified.

- Any new database created after the operating system upgrade is created using the UTF-8 Unicode code page unless a code page is explicitly specified when creating a new database. All existing databases retain their original code page settings; that is, the setting established during database creation.

To determine the active code page the system is running on Linux, run:

```
locale
```

Not all of the information displayed from running this command is important or relevant, however the DB2 database manager uses the following items in the order presented to determine the active code page:

- LC_ALL
- LC_CTYPE
- LANG

To determine which code page a database is using, run:

```
db2 get db cfg for <database name>
```

and check the value for the "Database code page" parameter.

Code Page/CCSID Numbers

Within IBM, the UCS-2 code page has been registered as code page 1200, with a growing character set; that is, when new characters are added to a code page, the code page number does not change. Code page 1200 always refers to the current version of Unicode.

A specific version of the UCS standard, as defined by Unicode 2.0 and ISO/IEC 10646-1, has also been registered within IBM as CCSID 13488. This CCSID has been used internally by DB2 for storing graphic string data in IBM eucJP (Japan) and IBM eucTW (Taiwan) databases. CCSID 13488 and code page 1200 both refer to UCS-2, and are handled the same way, except for the value of their "double-byte" (DBCS) space:

CP/CCSID	Single-byte (SBCS) space	Double-byte (DBCS) space
1200	N/A	U+0020
13488	N/A	U+3000

Note: In a UCS-2 database, U+3000 has no special meaning.

Regarding the conversion tables, since code page 1200 is a superset of CCSID 13488, the same (superset) tables are used for both.

Within IBM, UTF-8 has been registered as CCSID 1208 with growing character set (sometimes also referred to as code page 1208). As new characters are added to the standard, this number (1208) will not change.

The MBCS code page number is 1208, which is the database code page number, and the code page of character string data within the database. The double-byte code page number for UCS-2 is 1200, which is the code page of graphic string data within the database.

Thai and Unicode collation algorithm differences

The collation algorithm used in a Thai Industrial Standard (TIS) TIS620-1 (code page 874) Thai database with the NLSCHAR collation option is similar, but not identical to, the collation algorithm used in a Unicode database with the UCA400_LTH collation option. The differences are as follows:

- When sorting TIS620-1 data, each character only has one weight, and that weight is used to compare with another character's weight during collation. When sorting Unicode data, each character has several weights, and all the weights of that character can be used during collation.
- When sorting TIS620-1 data, the space character X'20', hyphen character X'2D', and full stop character X'2E' all have smaller weights than all the Thai characters. When sorting Unicode data, however, those three characters are considered as punctuation marks; and are used for comparison only when all other characters in the two strings being compared are equal.
- The Paiyannoi character X'CF' and the Maiyamok character X'E6' in a TIS620-1 database are treated as punctuation marks when they follow other Thai characters, and as normal characters, with their own weights, when they appear at the beginning of a string. The same two characters in a Unicode database (U+0E2F and U+0E46 respectively) are always treated as punctuation marks, and will be used for comparison when all other characters in the two strings being compared are equal.

More information on Thai characters can be found in chapter 10.1 Thai of the Unicode Standard book, version 4.0, ISBN 0-321-18578-1.

Related concepts:

- "Unicode character encoding" on page 355
- "Unicode handling of data types" on page 360
- "Unicode literals" on page 364

Related tasks:

- "Creating a Unicode database" on page 362

Related reference:

- "Character strings" in *SQL Reference, Volume 1*

Unicode handling of data types

All data types supported by DB2 Database for Linux, UNIX, and Windows are also supported in a UCS-2 database. In particular, graphic string data is supported for a UCS-2 database, and is stored in UCS-2/Unicode. Every client, including SBCS clients, can work with graphic string data types in UCS-2/Unicode when connected to a UCS-2 database.

A UCS-2 database is like any MBCS database where character string data is measured in number of bytes. When working with character string data in UTF-8, one should not assume that each character is one byte. In multibyte UTF-8 encoding, each ASCII character is one byte, but non-ASCII characters take two to four bytes each. This should be taken into account when defining CHAR fields. Depending on the ratio of ASCII to non-ASCII characters, a CHAR field of size n bytes can contain anywhere from $n/4$ to n characters.

Using character string UTF-8 encoding versus the graphic string UCS-2 data type also has an impact on the total storage requirements. In a situation where the majority of characters are ASCII, with some non-ASCII characters in between, storing UTF-8 data may be a better alternative, because the storage requirements are closer to one byte per character. On the other hand, in situations where the majority of characters are non-ASCII characters that expand to three- or four-byte UTF-8 sequences (for example ideographic characters), the UCS-2 graphic-string format may be a better alternative, because every three-byte UTF-8 sequence becomes a 16-bit UCS-2 character, while each four-byte UTF-8 sequence becomes two 16-bit UCS-2 characters.

In MBCS environments, SQL functions that operate on character strings, such as SUBSTR, POSSTR, MAX, MIN, and the like, operate on the number of "bytes" rather than number of "characters". The behavior is the same in a UCS-2 database, but you should take extra care when specifying offsets and lengths for a UCS-2 database, because these values are always defined in the context of the database code page. That is, in the case of a UCS-2 database, these offsets should be defined in UTF-8. Since some single-byte characters require more than one byte in UTF-8, SUBSTR indexes that are valid for a single-byte database may not be valid for a UCS-2 database. If you specify incorrect indexes, SQLCODE -191 (SQLSTATE 22504) is returned.

Note: Not all SQL functions that operate on character strings are limited to processing "bytes". The CHARACTER_LENGTH, LENGTH, LOCATE, POSITION, and SUBSTRING functions include a parameter that allows you to specify a predefined set of string units. This means that the functions can process strings using the specified units instead of bytes or double bytes.

SQL CHAR data types are supported (in the C language) by the char data type in user programs. SQL GRAPHIC data types are supported by sqldbchar in user programs. Note that, for a UCS-2 database, sqldbchar data is always in big-endian (high byte first) format. When an application program is connected to a UCS-2 database, character string data is converted between the application code page and UTF-8, and graphic string data is converted between the application graphic code page and UCS-2 by DB2.

When retrieving data from a Unicode database to an application that does not use an SBCS, EUC, or Unicode code page, the defined substitution character is returned for each blank padded to a graphic column. DB2 pads fixed-length Unicode graphic columns with ASCII blanks (U+0200), a character that has no equivalent in pure DBCS code pages. As a result, each ASCII blank used in the padding of the graphic column is converted to the substitution character on retrieval. Similarly, in a DATE, TIME or TIMESTAMP string, any SBCS character that does not have a pure DBCS equivalent is also converted to the substitution character when retrieved from a Unicode database to an application that does not use an SBCS, EUC, or Unicode code page.

Note: Prior to Version 8, graphic string data was always assumed to be in UCS-2. To provide backward compatibility to applications that depend on the previous behavior of DB2, the registry variable DB2GRAPHICUNICODESERVER has been introduced. Its default value is OFF. Changing the value of this variable to ON will cause DB2 to use its earlier behavior and assume that graphic string data is always in UCS-2. Additionally, the DB2 server will check the version of DB2 running on the client, and will simulate DB2 Universal Database Version 7 behavior if the client is running DB2 UDB Version 7.

Related concepts:

- “Unicode character encoding” on page 355
- “Unicode implementation in DB2 Database for Linux, UNIX, and Windows” on page 357

Creating a Unicode database

By default, databases are created in the code page of the application creating them. Therefore, if you create your database from a Unicode (UTF-8) client, your database will be created as a Unicode database. Alternatively, you can explicitly specify “UTF-8” as the CODESET name, and use any valid TERRITORY code supported by DB2 Database for Linux, UNIX, and Windows.

In a future release of the DB2 database manager, the default code set will be changed to UTF-8 when creating a database, regardless of the application code page.

Procedure:

To create a Unicode database with the territory code for the United States of America:

```
DB2 CREATE DATABASE dbname USING CODESET UTF-8 TERRITORY US
```

To create a Unicode database using the `sqlcrea` API, you should set the values in `sqldbterritoryinfo` accordingly. For example, set `SQLDBCODESET` to UTF-8, and `SQLDBLOCALE` to any valid territory code (for example, US).

Related concepts:

- “Unicode implementation in DB2 Database for Linux, UNIX, and Windows” on page 357

Related tasks:

- “Converting non-Unicode databases to Unicode” on page 362

Related reference:

- “sqlcrea API - Create database” in *Administrative API Reference*
- “CREATE DATABASE command” in *Command Reference*
- “Supported territory codes and code pages” on page 313

Converting non-Unicode databases to Unicode

There are some cases where you might need to convert an existing non-Unicode database to a Unicode database. For example, because XML columns are only supported in Unicode databases, if you want to add an XML column to an existing non-Unicode database, you will need to convert the database to a Unicode database before you can add the XML column.

Prerequisites:

You must have enough free disk space to export the data from the non-Unicode database. Also, if you are not reusing the existing table spaces, you will need enough free disk space to create new table spaces for the data.

Restrictions:

XML data can only be stored in single-partition databases defined with the UTF-8 code set.

Procedure:

The following steps illustrate how to convert an existing non-Unicode database to a Unicode database:

1. Export your data using the **db2move** command:

```
cd <export-dir>
db2move sample export
```

where <export-dir> is the directory to which you want to export your data and SAMPLE is the existing database name.

2. Generate a DDL script for your existing database using the **db2look** command:

```
db2look -d sample -e -o unidb.ddl -l -x -f
```

where SAMPLE is the existing database name and unidb.ddl is the file name for the generated DDL script. The -l option generates DDL for user defined table spaces, database partition groups and buffer pools, the -x option generates authorization DDL, and the -f option generates an update command for database configuration parameters.

3. Create the Unicode database:

```
CREATE DATABASE UNIDB USING CODESET UTF-8 TERRITORY US
```

where UNIDB is the name of the Unicode database.

4. Edit the unidb.ddl script and change all occurrences of the database name to the new Unicode database name:

```
CONNECT TO UNIDB
```

To keep the existing database, you must also change the file name specification for table spaces in the unidb.ddl file. Otherwise, you can drop the existing database and use the same table space files:

```
DROP DATABASE SAMPLE
```

5. Recreate your database structure by running the DDL script that you edited:

```
db2 -tvf unidb.ddl
```

6. Import your data into the new Unicode database using the **db2move** command:

```
cd <export-dir>
db2move unidb import
```

where <export-dir> is the directory where you exported your data and UNIDB is the Unicode database name.

Related concepts:

- “Unicode implementation in DB2 Database for Linux, UNIX, and Windows” on page 357
- “Native XML data store overview” in *XML Guide*
- “XML data type” in *XML Guide*

Related tasks:

- “Creating a Unicode database” on page 362

Related reference:

- “db2look - DB2 statistics and DDL extraction tool command” in *Command Reference*
- “db2move - Database movement tool command” in *Command Reference*
- “DROP DATABASE command” in *Command Reference*
- “CONNECT (Type 1) statement” in *SQL Reference, Volume 2*
- “CONNECT (Type 2) statement” in *SQL Reference, Volume 2*

Unicode literals

Unicode literals can be specified in two ways:

- As a graphic string constant, using the G'...' or N'....' format. Any literal specified in this way will be converted by the database manager from the application code page to 16-bit Unicode.
- As a Unicode hexadecimal string, using the UX'...' or GX'....' format. The constant specified between the quotation marks after UX or GX must be a multiple of four hexadecimal digits in big-endian order. Each four-digit group represents one 16-bit Unicode code point. Note that surrogate characters always appear in pairs, therefore you need two four-digit groups to represent the high and low surrogate characters.

When using the command line processor (CLP), the first method is easier if the UCS-2 character exists in the local application code page (for example, when entering any code page 850 character from a terminal that is using code page 850). The second method should be used for characters that are outside of the application code page repertoire (for example, when specifying Japanese characters from a terminal that is using code page 850).

Related concepts:

- “Unicode character encoding” on page 355
- “Unicode implementation in DB2 Database for Linux, UNIX, and Windows” on page 357

Related reference:

- “Constants” in *SQL Reference, Volume 1*

String comparisons in a Unicode database

Pattern matching is one area where the behavior of existing MBCS databases is slightly different from the behavior of a UCS-2 database.

For MBCS databases in DB2 Database for Linux, UNIX, and Windows, the current behavior is as follows: If the match-expression contains MBCS data, the pattern can include both SBCS and non-SBCS characters. The special characters in the pattern are interpreted as follows:

- An SBCS halfwidth underscore refers to one SBCS character.
- A non-SBCS fullwidth underscore refers to one non-SBCS character.
- A percent (either SBCS halfwidth or non-SBCS fullwidth) refers to zero or more SBCS or non-SBCS characters.

In a Unicode database, there is really no distinction between "single-byte" and "non-single-byte" characters. Although the UTF-8 format is a "mixed-byte" encoding of Unicode characters, there is no real distinction between SBCS and non-SBCS characters in UTF-8. Every character is a Unicode character, regardless of the number of bytes in UTF-8 format. In a Unicode graphic column, every non-supplementary character, including the halfwidth underscore (U+005F) and halfwidth percent (U+0025), is two bytes in width. For Unicode databases, the special characters in the pattern are interpreted as follows:

- For character strings, a halfwidth underscore (X'5F') or a fullwidth underscore (X'EFBCBF') refers to one Unicode character. A halfwidth percent (X'25') or a fullwidth percent (X'EFBC85') refers to zero or more Unicode characters.
- For graphic strings, a halfwidth underscore (U+005F) or a fullwidth underscore (U+FF3F) refers to one Unicode character. A halfwidth percent (U+0025) or a fullwidth percent (U+FF05) refers to zero or more Unicode characters.

Note: You need two underscores to match a Unicode supplementary graphic character because such a character is represented by two UCS-2 characters in a GRAPHIC column. Only one underscore is needed to match a Unicode supplementary character in a CHAR column.

For the optional "escape expression", which specifies a character to be used to modify the special meaning of the underscore and percent sign characters, the expression can be specified by any one of:

- A constant
- A special register
- A host variable
- A scalar function whose operands are any of the above
- An expression concatenating any of the above

with the restrictions that:

- No element in the expression can be of type LONG VARCHAR, CLOB, LONG VARCHAR, or DBCLOB. In addition, it cannot be a BLOB file reference variable.
- For CHAR columns, the result of the expression must be one character or a binary string containing exactly one (1) byte (SQLSTATE 22019). For GRAPHIC columns, the result of the expression must be one character (SQLSTATE 22019).

Related concepts:

- "Unicode character encoding" on page 355
- "Unicode implementation in DB2 Database for Linux, UNIX, and Windows" on page 357

Related reference:

- "Character strings" in *SQL Reference, Volume 1*
- "Graphic strings" in *SQL Reference, Volume 1*

Installing the previous tables for converting between code page 1394 and Unicode

The conversion tables for code page 1394 (also known as Shift JIS X0213) and Unicode have been enhanced. The conversion between Japanese Shift JIS X0213 (1394) and Unicode now conforms to the final ISO/IEC 10646-1:2000 Amendment-1 for JIS X0213 characters. The previous version of the conversion tables is available via FTP from <ftp://ftp.software.ibm.com/ps/products/db2/info/vr8/conv/>.

Procedure:

To install the previous definitions for converting between Shift JIS X0213 and Unicode:

1. Stop the DB2 Database for Linux, UNIX, and Windows instance.
2. Point your Web browser to <ftp://ftp.software.ibm.com/ps/products/db2/info/vr8/conv/> or use FTP to connect to the <ftp.software.ibm.com> site. This FTP server is anonymous.
3. If you are connecting via the command line, log in by entering anonymous as your user ID and your e-mail address as your password.
4. After logging in, change to the conversion tables directory:

```
cd ps/products/db2/info/vr8/conv
```
5. Copy the two files, `1394ucs4.cnv` and `ucs41394.cnv`, in binary form to your `sql1lib/conv/` directory.
6. Restart the DB2 instance.

Related concepts:

- “Unicode implementation in DB2 Database for Linux, UNIX, and Windows” on page 357

Related reference:

- “Supported territory codes and code pages” on page 313

Alternative Unicode conversion table for the coded character set identifier (CCSID) 943

There are several IBM coded character set identifiers (CCSIDs) for Japanese code pages. CCSID 943 is registered as the Microsoft Japanese Windows Shift-JIS code page. You might encounter the following two problems when converting characters between CCSID 943 and Unicode. The problems are the result of differences between the IBM code page conversion tables and the Microsoft code page conversion tables.

Problem 1::

For historical reasons, over 300 characters in the CCSID 943 code page are represented by two or three code points each. The use of input method editors (IMEs) and code page conversion tables cause only one of these equivalent code points to be entered. For example, the lower case character for Roman numeral one (“i”) has two equivalent code points: `X'EEEE'` and `X'FA40'`. Microsoft Windows IMEs always generate `X'FA40'` when “i” is entered. In general, IBM and Microsoft use the same primary code point to represent the character, except for the following 13 characters:

Table 104. CCSID 943 Shift-JIS code point conversion

Character name (Unicode code point)	IBM primary Shift-JIS code point	Microsoft Shift-JIS primary code point
Roman numeral one (U+2160)	X'FA4A'	X'8754'
Roman numeral two (U+2161)	X'FA4B'	X'8755'
Roman numeral three (U+2162)	X'FA4C'	X'8756'
Roman numeral four (U+2163)	X'FA4D'	X'8757'
Roman numeral five (U+2164)	X'FA4E'	X'8758'
Roman numeral six (U+2165)	X'FA4F'	X'8759'
Roman numeral seven (U+2166)	X'FA50'	X'875A'
Roman numeral eight (U+2167)	X'FA51'	X'875B'
Roman numeral nine (U+2168)	X'FA52'	X'875C'
Roman numeral ten (U+2169)	X'FA53'	X'875D'
Parenthesized ideograph stock (U+3231)	X'FA58'	X'878A'
Numero sign (U+2116)	X'FA59'	X'8782'
Telephone sign (U+2121)	X'FA5A'	X'8784'

IBM products such as DB2 database manager primarily use IBM code points, for example X'FA4A', to present the upper case Roman numeral "I", but Microsoft products use X'8754' to represent the same character. A Microsoft ODBC application can insert the "I" character as X'8754' into a DB2 database of CCSID 943, and the DB2 Control Center can insert the same character as X'FA4A' into the same CCSID 943 database. However, Microsoft ODBC applications can find only those rows that have "I" encoded as X'8754', and the DB2 Control Center can locate only those rows that have encoded "I" as X'FA4A'. To enable the DB2 Control Center to select "I" as X'8754', you need to replace the default IBM conversion tables from Unicode to CCSID 943 with the alternate Microsoft conversion table provided by the DB2 database manager.

Problem 2::

The following list of characters, when converted from CCSID 943 to Unicode, will result in different code points depending on whether the IBM conversion table or the Microsoft conversion table is used. For these characters, the IBM conversion table conforms to the character names as specified in the Japanese Industry Standard JISX0208, JISX0212, and JISX0221.

Table 105. CCSID 943 to Unicode code point conversion

Shift-JIS code point (character name)	IBM primary code point (Unicode name)	Microsoft primary code point (Unicode name)
X'815C' (EM Dash)	U+2014 (EM Dash)	U+2015 (Horizontal Bar)

Table 105. CCSID 943 to Unicode code point conversion (continued)

Shift-JIS code point (character name)	IBM primary code point (Unicode name)	Microsoft primary code point (Unicode name)
X'8160' (Wave Dash)	U+301C (Wave Dash)	U+FF5E (Fullwidth Tilde)
X'8161' (Double vertical line)	U+2016 (Double vertical line)	U+2225 (Parallel To)
X'817C' (Minus sign)	U+2212 (Minus sign)	U+FF0D (Fullwidth hyphen-minus)
X'FA55' (Broken bar)	U+00A6 (Broken bar)	U+FFE4 (Fullwidth broken bar)

For example, the character EM dash with the CCSID 943 code point of X'815C' is converted to the Unicode code point U+2014 when using the IBM conversion table, but is converted to U+2015 when using the Microsoft conversion table. This can create potential problems for Microsoft ODBC applications because they would treat U+2014 as an invalid code point. To avoid these potential problems, you need to replace the default IBM conversion table from CCSID 943 to Unicode with the alternate Microsoft conversion table provided by the DB2 database manager.

The use of the alternate Microsoft conversion tables between CCSID 943 and Unicode should be restricted to closed environments, where the DB2 clients and the DB2 databases that are running CCSID 943 and are all using the same alternate Microsoft conversion tables. If you have a DB2 client using the default IBM conversion tables and another client using the alternate Microsoft conversion tables, and both clients are inserting data to the same DB2 database of CCSID 943, the same character may be stored as different code points in the database.

Related concepts:

- “Unicode character encoding” on page 355

Related tasks:

- “Replacing the Unicode conversion tables for coded character set identifier (CCSID) 943 with Microsoft conversion tables” on page 368

Replacing the Unicode conversion tables for coded character set identifier (CCSID) 943 with Microsoft conversion tables

When you convert between coded character set identifier (CCSID) 943 and Unicode, the DB2 Database for Linux, UNIX, and Windows database manager default code page conversion tables are used. If you want to use a different version of the conversion tables, such as the Microsoft version, you must manually override the default conversion tables.

Prerequisites:

If the code page conversion table file you want to override already exists in the conv subdirectory of the sql11ib directory, you should back up that file in case you want to revert to the default table.

Restrictions:

For conversion table replacement to be effective, the conversion table on the database server and all of its clients must be changed.

Procedure:

To replace the DB2 default conversion tables for converting between CCSID 943 and Unicode:

1. When replacing conversion tables on the client, stop all the applications that are using the database. If you have any CLP sessions running, issue the TERMINATE command for each session. When replacing conversion tables on the database server, stop all instances on all nodes by issuing the **db2stop** command.
2. Copy `sql1lib/conv/ms/0943ucs2.cnv` to `sql1lib/conv/0943ucs2.cnv`.
3. Copy `sql1lib/conv/ms/ucs20943.cnv` to `sql1lib/conv/ucs20943.cnv`.
4. Restart all the applications.

Related concepts:

- “Alternative Unicode conversion table for the coded character set identifier (CCSID) 943” on page 366

Alternative Unicode conversion table for the coded character set identifier (CCSID) 954

There are several IBM coded character set identifiers (CSIDs) for Japanese code pages. CCSID 954 is registered as the Japanese EUC code page. CCSID 954 is a common encoding for Japanese UNIX and Linux platforms. When using Microsoft ODBC applications to connect to a DB2 database using CCSID 954, you might encounter potential problems when converting data in CCSID 954 to Unicode. The problems are the result of differences between IBM’s code page conversion table and Microsoft’s code page conversion table.

The following list of characters, when converted from CCSID 954 to Unicode, will result in different code points depending on which conversion table (IBM or Microsoft) is used. For these characters, the IBM conversion table conforms to the character names as specified in the Japanese Industry Standard (JIS) JISX0208, JISX0212, and JISX0221.

Table 106. CCSID 954 to Unicode code point conversion

EUC-JP code point (character name)	IBM primary code point (Unicode name)	Microsoft primary code point (Unicode name)
X'A1BD' (EM Dash)	U+2014 (EM Dash)	U+2015 (Horizontal Bar)
X'A1C1' (Wave Dash)	U+301C (Wave Dash)	U+FF5E (Fullwidth Tilde)
X'A1C2' (Double vertical line)	U+2016 (Double vertical line)	U+2225 (Parallel To)
X'A1DD' (Minus sign)	U+2212 (Minus sign)	U+FF0D (Fullwidth hyphen-minus)
X'8FA2C3' (Broken bar)	U+00A6 (Broken bar)	U+FFE4 (Fullwidth broken bar)

For example, the character EM dash with the CCSID 954 code point of X'A1BD' is converted to the Unicode code point U+2014 when using the IBM conversion table, but is converted to U+2015 when using the Microsoft conversion table. This can create potential problems for Microsoft ODBC applications because they would treat U+2014 as an invalid code point. To avoid these potential problems, you need

to replace the default IBM conversion table from CCSID 954 to Unicode with the alternate Microsoft conversion table provided by the DB2 database manager.

Related concepts:

- “Replacing the Unicode conversion table for coded character set identifier (CCSID) 954 with the Microsoft conversion table” on page 370
- “Unicode character encoding” on page 355

Replacing the Unicode conversion table for coded character set identifier (CCSID) 954 with the Microsoft conversion table

When you convert from coded character set identifier (CCSID) 954 to Unicode, the DB2 database manager default code page conversion table is used. If you want to use a different version of the conversion table such as the Microsoft version, you must manually override the default conversion table.

Prerequisites::

If the code page conversion table file you want to override already exists in the *conv* subdirectory of the *sqllib* directory, you should back up that file in case you want to revert to the default table.

Restrictions::

For conversion table replacement to be effective, every DB2 client that connects to the same database must have its conversion table changed. If your client is Japanese Windows whose ANSI code page is Shift-JIS (CCSID 943), you will also need to change the default conversion tables between CCSID 943 and Unicode to the Microsoft version. Otherwise, the different clients might store the same character using different code points.

Procedure::

To replace the DB2 default conversion table for converting from CCSID 954 to Unicode, follow these steps:

1. When replacing conversion tables on the client, stop all the applications that are using the database. If you have any CLP sessions running, issue the **TERMINATE** command for each session. When replacing conversion tables on the database server, stop all instances on all nodes by issuing the **db2stop** command.
2. Copy *sqllib/conv/ms/0954ucs2.cnv* to *sqllib/conv/0954ucs2.cnv*.
3. Restart all the applications.

To replace the DB2 default conversion tables for converting between CCSID 943 and Unicode, follow these steps:

1. When replacing conversion tables on the client, stop all the applications that are using the database. If you have any CLP sessions running, issue the **TERMINATE** command for each session. When replacing conversion tables on the database server, stop all instances on all nodes by issuing the **db2stop** command.
2. Copy *sqllib/conv/ms/0943ucs2.cnv* to *sqllib/conv/0943ucs2.cnv*.
3. Copy *sqllib/conv/ms/ucs20943.cnv* to *sqllib/conv/ucs20943.cnv*.
4. Restart all the applications.

Related concepts:

- “Alternative Unicode conversion table for the coded character set identifier (CCSID) 954” on page 369
- “Unicode character encoding” on page 355

Alternative Unicode conversion table for the coded character set identifier (CCSID) 5026

There are several IBM coded character set identifiers (CSIDs) for Japanese code pages. CCSID 5026 is registered as a Japanese EBCDIC code page. When using Microsoft ODBC applications to connect to a DB2 host database of CCSID 5026, you might encounter potential problems when converting data in CCSID 5026 to Unicode. The problems are the result of differences between IBM’s code page conversion table and Microsoft’s code page conversion table. The following list of characters, when converted from CCSID 5026 to Unicode, will result in different code points depending on which conversion table (IBM or Microsoft) is used. For these characters, the IBM conversion table conforms to the character names as specified in the Japanese Industry Standard (JIS) JISX0208, JISX0212, and JISX0221.

Table 107. CCSID 5026 to Unicode code point conversion

EBCDIC code point (character name)	IBM primary code point (Unicode name)	Microsoft primary code point (Unicode name)
X'444A' (EM Dash)	U+2014 (EM Dash)	U+2015 (Horizontal Bar)
X'43A1' (Wave Dash)	U+301C (Wave Dash)	U+FF5E (Fullwidth Tilde)
X'447C' (Double vertical line)	U+2016 (Double vertical line)	U+2225 (Parallel To)
X'4260' (Minus sign)	U+2212 (Minus sign)	U+FF0D (Fullwidth hyphen-minus)
X'426A' (Broken bar)	U+00A6 (Broken bar)	U+FFE4 (Fullwidth broken bar)

For example, the character EM dash with the CCSID 5026 code point of X'444A' is converted to the Unicode code point U+2014 when using the IBM conversion table, but is converted to U+2015 when using the Microsoft conversion table. This can create potential problems for Microsoft ODBC applications because they would treat U+2014 as an invalid code point. To avoid these potential problems, you need to replace the default IBM conversion table from CCSID 5026 to Unicode with the alternate Microsoft conversion table provided by the DB2 database manager.

Related concepts:

- “Replacing the Unicode conversion table for coded character set identifier (CCSID) 5026 with the Microsoft conversion table” on page 371
- “Unicode character encoding” on page 355

Replacing the Unicode conversion table for coded character set identifier (CCSID) 5026 with the Microsoft conversion table

When you convert from coded character set identifier (CCSID) 5026 to Unicode, the DB2 database manager default code page conversion table is used. If you want to use a different version of the conversion table such as the Microsoft version, you must manually override the default conversion table.

Prerequisites::

If the code page conversion table file you want to override already exists in the *conv* subdirectory of the *sqllib* directory, you should back up that file in case you want to revert to the default table.

Restrictions::

For conversion table replacement to be effective, every DB2 client that connects to the same database must have its conversion table changed.

This Microsoft conversion table is only for data encoded in CCSID 5026 or 930, and cannot be used for data encoded in CCSID 1390. Since the DB2 database manager uses the same conversion table for data encoded in CCSIDs 5026, 930, and 1390, this means that once the default IBM conversion table has been replaced with the Microsoft conversion table, you should not select any data that is encoded in CCSID 1390.

Activating this alternate Microsoft conversion table does not change the code page conversion behavior of graphic data encoded in 5026 to Unicode. To enable graphic data encoded in 5026 conversion to Unicode using the alternate Microsoft conversion table, you must also copy the file *sqllib/conv/ms/0939ucs2.cnv* to *sqllib/conv/1399ucs2.cnv* in addition to the procedure outlined below. Once you complete these steps, the conversion of both character data and graphic data to Unicode from the following CCSIDs will also use the Microsoft conversion table: 5026, 930, 1390, 5035, 939, and 1399.

Procedure::

To replace the DB2 default conversion table for converting from CCSID 5026 to Unicode, follow these steps:

1. When replacing conversion tables on the client, stop all the applications that are using the database. If you have any CLP sessions running, issue the *db2 terminate* command for each session.
2. Copy *sqllib/conv/ms/0930ucs2.cnv* to *sqllib/conv/1390ucs2.cnv*.
3. Restart all the applications.

Related concepts:

- “Alternative Unicode conversion table for the coded character set identifier (CCSID) 5026” on page 371

Alternative Unicode conversion table for the coded character set identifier (CCSID) 5035

There are several IBM coded character set identifiers (CCSIDs) for Japanese code pages. CCSID 5035 is registered as a Japanese EBCDIC code page. When using Microsoft ODBC applications to connect to a DB2 host database of CCSID 5035, you might encounter potential problems when converting data in CCSID 5035 to Unicode. The problems are the result of differences between IBM’s code page conversion table and Microsoft’s code page conversion table. The following list of characters, when converted from CCSID 5035 to Unicode, will result in different code points depending on which conversion table (IBM or Microsoft) is used. For these characters, the IBM conversion table conforms to the character names as specified in the Japanese Industry Standard (JIS) JISX0208, JISX0212, and JISX0221.

Table 108. CCSID 5035 to Unicode code point conversion

EBCDIC code point (character name)	IBM primary code point (Unicode name)	Microsoft primary code point (Unicode name)
X'444A' (EM Dash)	U+2014 (EM Dash)	U+2015 (Horizontal Bar)
X'43A1' (Wave Dash)	U+301C (Wave Dash)	U+FF5E (Fullwidth Tilde)
X'447C' (Double vertical line)	U+2016 (Double vertical line)	U+2225 (Parallel To)
X'4260' (Minus sign)	U+2212 (Minus sign)	U+FF0D (Fullwidth hyphen-minus)
X'426A' (Broken bar)	U+00A6 (Broken bar)	U+FFE4 (Fullwidth broken bar)

For example, the character EM dash with the CCSID 5035 code point of X'444A' is converted to the Unicode code point U+2014 when using the IBM conversion table, but is converted to U+2015 when using the Microsoft conversion table. This can create potential problems for Microsoft ODBC applications because they would treat U+2014 as an invalid code point. To avoid these potential problems, you need to replace the default IBM conversion table from CCSID 5035 to Unicode with the alternate Microsoft conversion table provided by the DB2 database manager.

Related concepts:

- “Unicode character encoding” on page 355
- “Replacing the Unicode conversion table for coded character set identifier (CCSID) 5035 with the Microsoft conversion table” on page 373

Replacing the Unicode conversion table for coded character set identifier (CCSID) 5035 with the Microsoft conversion table

When you convert from coded character set identifier (CCSID) 5035 to Unicode, the DB2 database manager default code page conversion table is used. If you want to use a different version of the conversion table such as the Microsoft version, you must manually override the default conversion table.

Prerequisites::

If the code page conversion table file you want to override already exists in the *conv* subdirectory of the *sqliib* directory, you should back up that file in case you want to revert to the default table.

Restrictions::

For conversion table replacement to be effective, every DB2 client that connects to the same database must have its conversion table changed.

This Microsoft conversion table is only for data encoded in CCSID 5039 or 939, and cannot be used for data encoded in CCSID 1399. Since the DB2 database manager uses the same conversion table for data encoded in CCSIDs 5035, 939, and 1399, this means that once the default IBM conversion table has been replaced with the Microsoft conversion table, you should not select any data that is encoded in CCSID 1399.

Once you have replaced the default IBM conversion table with the Microsoft conversion table, the conversion of graphic data to Unicode from the following CCSIDs will also use this Microsoft conversion table: 930, 1390, 939, and 1399.

Procedure::

To replace the DB2 default conversion table for converting from CCSID 5035 to Unicode, follow these steps:

1. When replacing conversion tables on the client, stop all the applications that are using the database. If you have any CLP sessions running, issue the **TERMINATE** command for each session.
2. Copy *sqllib/conv/ms/0939ucs2.cnv* to *sqllib/conv/1399ucs2.cnv*.
3. Restart all the applications.

Related concepts:

- “Alternative Unicode conversion table for the coded character set identifier (CCSID) 5035” on page 372
- “Unicode character encoding” on page 355

Alternative Unicode conversion table for the coded character set identifier (CCSID) 5039

There are several IBM coded character set identifiers (CCSIDs) for Japanese code pages. CCSID 943 is registered as the the Microsoft Japanese Windows Shift-JIS code page. However, the Shift-JIS code page on the HP-UX platform is registered as CCSID 5039. CCSID 5039 contains only Japanese Industry Standard (JIS) characters, and does not have any vendor-defined characters. When using Microsoft ODBC applications, you might encounter potential problems when converting data in CCSID 5039 to Unicode. The problems are the result of differences between IBM’s code page conversion table and Microsoft’s code page conversion table.

The following list of characters, when converted from CCSID 5039 to Unicode, will result in different code points depending on which conversion table (IBM or Microsoft) is used. For these characters, the IBM conversion table conforms to the character names as specified in the Japanese Industry Standard (JIS) JISX0208, and JISX0221.

Table 109. CCSID 5039 to Unicode code point conversion

Shift-JIS code point (character name)	IBM primary code point (Unicode name)	Microsoft primary code point (Unicode name)
X'815C' (EM Dash)	U+2014 (EM Dash)	U+2015 (Horizontal Bar)
X'8160' (Wave Dash)	U+301C (Wave Dash)	U+FF5E (Fullwidth Tilde)
X'8161' (Double vertical line)	U+2016 (Double vertical line)	U+2225 (Parallel To)
X'817C' (Minus sign)	U+2212 (Minus sign)	U+FF0D (Fullwidth hyphen-minus)

For example, the character EM dash with the CCSID 5039 code point of X'815C' is converted to the Unicode code point U+2014 when using the IBM conversion table, but is converted to U+2015 when using the Microsoft conversion table. This can create potential problems for Microsoft ODBC applications because they would treat U+2014 as an invalid code point. To avoid these potential problems, you need

to replace the default IBM conversion table from CCSID 5039 to Unicode with the alternate Microsoft conversion table provided by the DB2 database manager.

Related concepts:

- “Replacing the Unicode conversion table for coded character set identifier (CCSID) 5039 with the Microsoft conversion table” on page 375
- “Unicode character encoding” on page 355

Replacing the Unicode conversion table for coded character set identifier (CCSID) 5039 with the Microsoft conversion table

When you convert from coded character set identifier (CCSID) 5039 to Unicode, the DB2 database manager default code page conversion table is used. If you want to use a different version of the conversion table such as the Microsoft version, you must manually override the conversion table.

Prerequisites::

If the code page conversion table file you want to override already exists in the *conv* subdirectory of the *sqllib* directory, you should back up that file in case you want to revert to the default table.

Restrictions::

For conversion table replacement to be effective, every DB2 client that connects to the same database must have its conversion table changed.

Procedure::

To replace the DB2 default conversion table for converting from CCSID 5039 to Unicode, follow these steps:

1. When replacing conversion tables on the client, stop all the applications that are using the database. If you have any CLP sessions running, issue the **TERMINATE** command for each session.
2. Copy *sqllib/conv/ms/5039ucs2.cnv* to *sqllib/conv/5039ucs2.cnv*.
3. Restart all the applications.

Related concepts:

- “Alternative Unicode conversion table for the coded character set identifier (CCSID) 5039” on page 374
- “Unicode character encoding” on page 355

Appendix C. DB2 Database technical information

Overview of the DB2 technical information

DB2 technical information is available through the following tools and methods:

- DB2 Information Center
 - Topics
 - Help for DB2 tools
 - Sample programs
 - Tutorials
- DB2 books
 - PDF files (downloadable)
 - PDF files (from the DB2 PDF CD)
 - printed books
- Command line help
 - Command help
 - Message help
- Sample programs

IBM periodically makes documentation updates available. If you access the online version on the DB2 Information Center at ibm.com[®], you do not need to install documentation updates because this version is kept up-to-date by IBM. If you have installed the DB2 Information Center, it is recommended that you install the documentation updates. Documentation updates allow you to update the information that you installed from the *DB2 Information Center CD* or downloaded from Passport Advantage as new information becomes available.

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 Redbooks™ online at 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 we can improve the DB2 documentation, send an e-mail 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 e-mail 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.

Related concepts:

- “Features of the DB2 Information Center” in *Online DB2 Information Center*
- “Sample files” in *Samples Topics*

Related tasks:

- “Invoking command help from the command line processor” in *Command Reference*
- “Invoking message help from the command line processor” in *Command Reference*
- “Updating the DB2 Information Center installed on your computer or intranet server” on page 383

Related reference:

- “DB2 technical library in hardcopy or PDF format” on page 378

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. DB2 Version 9 manuals in PDF format can be downloaded from www.ibm.com/software/data/db2/udb/support/manualsv9.html.

Although the tables identify books available in print, the books might not be available in your country or region.

The information in these books is fundamental to all DB2 users; you will find this information useful whether you are a programmer, a database administrator, or someone who works with DB2 Connect or other DB2 products.

Table 110. DB2 technical information

Name	Form Number	Available in print
<i>Administration Guide: Implementation</i>	SC10-4221	Yes
<i>Administration Guide: Planning</i>	SC10-4223	Yes
<i>Administrative API Reference</i>	SC10-4231	Yes
<i>Administrative SQL Routines and Views</i>	SC10-4293	No
<i>Call Level Interface Guide and Reference, Volume 1</i>	SC10-4224	Yes
<i>Call Level Interface Guide and Reference, Volume 2</i>	SC10-4225	Yes
<i>Command Reference</i>	SC10-4226	No
<i>Data Movement Utilities Guide and Reference</i>	SC10-4227	Yes
<i>Data Recovery and High Availability Guide and Reference</i>	SC10-4228	Yes
<i>Developing ADO.NET and OLE DB Applications</i>	SC10-4230	Yes
<i>Developing Embedded SQL Applications</i>	SC10-4232	Yes

Table 110. DB2 technical information (continued)

Name	Form Number	Available in print
<i>Developing SQL and External Routines</i>	SC10-4373	No
<i>Developing Java Applications</i>	SC10-4233	Yes
<i>Developing Perl and PHP Applications</i>	SC10-4234	No
<i>Getting Started with Database Application Development</i>	SC10-4252	Yes
<i>Getting started with DB2 installation and administration on Linux and Windows</i>	GC10-4247	Yes
<i>Message Reference Volume 1</i>	SC10-4238	No
<i>Message Reference Volume 2</i>	SC10-4239	No
<i>Migration Guide</i>	GC10-4237	Yes
<i>Net Search Extender Administration and User's Guide</i> Note: HTML for this document is not installed from the HTML documentation CD.	SH12-6842	Yes
<i>Performance Guide</i>	SC10-4222	Yes
<i>Query Patroller Administration and User's Guide</i>	GC10-4241	Yes
<i>Quick Beginnings for DB2 Clients</i>	GC10-4242	No
<i>Quick Beginnings for DB2 Servers</i>	GC10-4246	Yes
<i>Spatial Extender and Geodetic Data Management Feature User's Guide and Reference</i>	SC18-9749	Yes
<i>SQL Guide</i>	SC10-4248	Yes
<i>SQL Reference, Volume 1</i>	SC10-4249	Yes
<i>SQL Reference, Volume 2</i>	SC10-4250	Yes
<i>System Monitor Guide and Reference</i>	SC10-4251	Yes
<i>Troubleshooting Guide</i>	GC10-4240	No
<i>Visual Explain Tutorial</i>	SC10-4319	No
<i>What's New</i>	SC10-4253	Yes
<i>XML Extender Administration and Programming</i>	SC18-9750	Yes
<i>XML Guide</i>	SC10-4254	Yes
<i>XQuery Reference</i>	SC18-9796	Yes

Table 111. DB2 Connect-specific technical information

Name	Form Number	Available in print
<i>DB2 Connect User's Guide</i>	SC10-4229	Yes

Table 111. DB2 Connect-specific technical information (continued)

Name	Form Number	Available in print
Quick Beginnings for DB2 Connect Personal Edition	GC10-4244	Yes
Quick Beginnings for DB2 Connect Servers	GC10-4243	Yes

Table 112. WebSphere Information Integration technical information

Name	Form Number	Available in print
WebSphere Information Integration: Administration Guide for Federated Systems	SC19-1020	Yes
WebSphere Information Integration: ASNCLP Program Reference for Replication and Event Publishing	SC19-1018	Yes
WebSphere Information Integration: Configuration Guide for Federated Data Sources	SC19-1034	No
WebSphere Information Integration: SQL Replication Guide and Reference	SC19-1030	Yes

Note: The DB2 Release Notes provide additional information specific to your product's release and fix pack level. For more information, see the related links.

Related concepts:

- "Overview of the DB2 technical information" on page 377
- "About the Release Notes" in *Release notes*

Related tasks:

- "Ordering printed DB2 books" on page 380

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 CD* 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 CD 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 CD 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/db2help/>.

Procedure:

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:
 - 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.
 - When you call, specify that you want to order a DB2 publication.
 - Provide your representative with the titles and form numbers of the books that you want to order.

Related concepts:

- "Overview of the DB2 technical information" on page 377

Related reference:

- "DB2 technical library in hardcopy or PDF format" on page 378

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.

Procedure:

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.

Related tasks:

- "Invoking command help from the command line processor" in *Command Reference*
- "Invoking message help from the command line processor" in *Command Reference*

Accessing different versions of the DB2 Information Center

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/>.

Related tasks:

- “Updating the DB2 Information Center installed on your computer or intranet server” on page 383

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.

Procedure:

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 **Tools** → **Options** → **Languages** button. 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.

Related concepts:

- “Overview of the DB2 technical information” on page 377

Updating the DB2 Information Center installed on your computer or intranet server

If you have a locally-installed DB2 Information Center, updated topics can be available for download. The 'Last updated' value found at the bottom of most topics indicates the current level for that topic.

To determine if there is an update available for the entire DB2 Information Center, look for the 'Last updated' value on the Information Center home page. Compare the value in your locally installed home page to the date of the most recent downloadable update at <http://www.ibm.com/software/data/db2/udb/support/icupdate.html>. You can then update your locally-installed Information Center if a more recent downloadable update is available.

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 download and apply updates.
2. Use the Update feature to determine if update packages are available from IBM.

Note: Updates are also available on CD. For details on how to configure your Information Center to install updates from CD, see the related links. If update packages are available, use the Update feature to download the packages. (The Update feature is only available in stand-alone mode.)

3. Stop the stand-alone Information Center, and restart the DB2 Information Center service on your computer.

Procedure:

To update the DB2 Information Center installed on your computer or intranet server:

1. Stop the DB2 Information Center service.
 - 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/db2icdv9 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 `C:\Program Files\IBM\DB2 Information Center\Version 9` directory.
 - c. Run the `help_start.bat` file using the fully qualified path for the DB2 Information Center:
`<DB2 Information Center dir>\doc\bin\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 directory.
- b. Run the help_start script using the fully qualified path for the DB2 Information Center:

```
<DB2 Information Center dir>/doc/bin/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 download process, check the selections you want to download, then click **Install Updates**.
5. After the download and installation process has completed, click **Finish**.
6. Stop the stand-alone Information Center.
 - On Windows, run the help_end.bat file using the fully qualified path for the DB2 Information Center:

```
<DB2 Information Center dir>\doc\bin\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, run the help_end script using the fully qualified path for the DB2 Information Center:

```
<DB2 Information Center dir>/doc/bin/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 service.
 - 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/db2icdv9 start
```

The updated DB2 Information Center displays the new and updated topics.

Related concepts:

- “DB2 Information Center installation options” in *Quick Beginnings for DB2 Servers*

Related tasks:

- “Installing the DB2 Information Center using the DB2 Setup wizard (Linux)” in *Quick Beginnings for DB2 Servers*
- “Installing the DB2 Information Center using the DB2 Setup wizard (Windows)” in *Quick Beginnings for DB2 Servers*

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.

Native XML data store

Set up a DB2 database to store XML data and to perform basic operations with the native XML data store.

Visual Explain Tutorial

Analyze, optimize, and tune SQL statements for better performance using Visual Explain.

Related concepts:

- “Visual Explain overview” in *Administration Guide: Implementation*

DB2 troubleshooting information

A wide variety of troubleshooting and problem determination information is available to assist you in using DB2 products.

DB2 documentation

Troubleshooting information can be found in the DB2 Troubleshooting Guide or the Support and Troubleshooting 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 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/udb/support.html>

Related concepts:

- “Introduction to problem determination” in *Troubleshooting Guide*
- “Overview of the DB2 technical information” on page 377

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 D. Notices

IBM may not offer the products, services, or features discussed in this document in all 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.

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

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

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

Company, product, or service names identified in the documents of the DB2 Version 9 documentation library may be trademarks or service marks of International Business Machines Corporation or other companies. Information on the trademarks of IBM Corporation in the United States, other countries, or both is located at <http://www.ibm.com/legal/copytrade.shtml>.

The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the documents in the DB2 documentation library:

Microsoft, Windows, Windows NT[®], and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Itanium[®], Pentium[®], and Xeon[®] are trademarks of Intel Corporation in the United States, other countries, or both.

Java and all Java-based trademarks 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.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Index

A

- alternative unicode conversion tables
 - CCSID 5026 371
- APIs
 - heuristic 230
- append mode tables 167
- application design
 - collating sequences, guidelines 351
- applications
 - incompatibility 285
- Asian fonts
 - Linux 334
- audit activities 71
- audit context records
 - incompatibility 285
- authentication
 - about 20
 - description 20
- authority
 - incompatibility 285
- authorization
 - about 20
 - database design considerations 71
 - description 21
- automatic client reroute 219
- automatic features
 - automatic reorganization 32
 - enabled by default 30
 - statistics collection 33
- automatic maintenance 29, 31
 - about 29
 - backup 25
 - maintenance windows 35
 - offline 36
 - online 36
- automatic reorganization
 - description 32
 - enabling 32
- automatic statistics collection
 - description 30, 33
 - storage for 35
- automatic statistics profiling
 - description 34
 - enabling 34
 - storage for 35
- automatic storage
 - description 30

B

- backups
 - automated 25
 - automatic 31
- BEA Tuxedo, configuring 238
- bidirectional CCSID support
 - DB2 345
 - DB2 Connect 349
 - list of CCSIDs 347
- block identifier (BID) 177

- block indexes
 - benefits 175
 - composite 183
 - insert using 183
 - MDC table considerations 188
 - query performance 180
- block maps 185
- blocks
 - multidimensional clustering (MDC) 177
- buffer pools
 - description 3
 - IBMDEFAULTBP 145
- business rules
 - description 17
 - transitional 70

C

- CALL statement
 - incompatibility 285
- capacity
 - for each environment 42
- casting FOR BIT DATA
 - incompatibility 285
- catalog table spaces 112, 163
- CCSID
 - 5026 371
- CCSID (coded character set identifier) 366, 368
 - bidirectional support
 - DB2 345
 - DB2 Connect 349
 - types listed 347
- CCSID 5026
 - alternative unicode conversion tables 371
 - replacing unicode conversion tables 371
- CCSID 5035
 - Microsoft conversion table 373
 - unicode conversion table 372
- CCSID 5039
 - Microsoft conversion table 375
 - unicode conversion table 374
- CCSID 954
 - Microsoft conversion table 370
 - unicode conversion table 369
- changed behavior
 - from previous releases 260
- CHAR function
 - incompatibility 285
- character conversion
 - effect on application performance 338
- character strings
 - Unicode 360
- check constraints
 - as business rules 17
- choosing
 - extent size 144
- choosing (*continued*)
 - multidimensional table
 - dimensions 189
 - table spaces 112
- CHR function
 - incompatibility 285
- client reroute
 - automatic 219
- clustering
 - automatic 183
- clustering, data 172
- code page 950
 - IBM and Microsoft differences 313
- code page conversion
 - incompatibility 285
- code pages
 - 923 and 924 336, 343
 - converting 1394 to Unicode, previous conversion tables 366
 - converting Shift JIS X0213 to Unicode, previous conversion tables 366
 - DB2 supported 313
 - with euro symbol 336, 339
- code point 351, 366
- code sets
 - DB2 supported 313
- coded character set identifier
 - 5026 371
- coded character set identifier 5035
 - Microsoft conversion table 373
 - unicode conversion table 372
- coded character set identifier 5039
 - Microsoft conversion table 375
 - unicode conversion table 374
- coded character set identifier 943
 - considerations when using 366
- coded character set identifier 954
 - Microsoft conversion table 370
 - unicode conversion table 369
- Coded character set identifiers 5026 371
- collating algorithm differences
 - Thai and Unicode 357
- collating sequences
 - code point 351
 - concerns, general 351
 - identity sequence 351
 - multi-byte characters 351
 - overview 351
 - Thai characters 352
 - Unicode 357
- collocation, table 91
- column expressions, multidimensional tables 197
- columns
 - defining for a table 56
- commit
 - errors during two-phase 213
 - two-phase 210
- comparison of indexes
 - clustering and block-based 173

- comparison of tables
 - regular and multidimensional clustering 173
 - compatibility
 - partition 91
 - composite block index 188
 - composite keys
 - primary keys 58
 - Configuration Advisor
 - description 30
 - configuration files
 - description 12
 - location 12
 - configuration parameters
 - DB2 transaction manager considerations 207
 - description 12
 - incompatibility 285
 - configurations
 - multiple partition 42
 - configure automatic maintenance wizard 31
 - connection failure
 - automatic client reroute 219
 - constants
 - Unicode 364
 - constraints
 - check 17
 - foreign key 17
 - informational 17, 65
 - NOT NULL 17
 - primary key 17
 - referential 65
 - table check 65
 - unique 17, 65
 - contacting IBM 393
 - containers
 - description 3
 - DMS table spaces
 - addition of containers to 129
 - dropping containers from 137
 - extension of containers in 129
 - reduction of containers in 137
 - CONTROL privilege on packages
 - incompatibility 285
 - conversions
 - Unicode to CCSID 943 366, 368
 - coordinator partition 41
 - CREATE TABLE
 - OVERFLOW clause 171
 - creating
 - multidimensional tables 197
 - Unicode databases 362
- D**
- data
 - distribution 41
 - large object (LOB) 79
 - long field 78
 - security 20
 - data distribution
 - table 99
 - data organization
 - table 99
 - data organization schemes
 - combining 99
 - data organization schemes (*continued*)
 - comparison 105
 - description 105
 - data partitioning
 - see table partitioning 93
 - data partitions
 - description 92
 - data types
 - database design considerations 71
 - Unicode handling 360
 - data types and scrollable cursors
 - incompatibility 285
 - database connection
 - incompatibility 285
 - database design
 - additional considerations 71
 - logical 53
 - physical 73
 - database directories
 - structure described 73
 - database objects
 - database partition groups 3
 - databases 3
 - indexes 3
 - instances 3
 - recovery history file 25
 - recovery log file 25
 - schemas 3
 - system catalog tables 3
 - table space change history file 25
 - table spaces 3
 - tables 3
 - views 3
 - database partition groups
 - collocation 87
 - description 3, 85
 - designing 87
 - determining data location 88
 - IBM CATGROUP 112
 - IBM DEFAULTGROUP 112
 - IBM TEMPGROUP 112
 - database partitioning
 - database 99
 - database partitions
 - database 41
 - description 41
 - database-managed space (DMS)
 - containers 129
 - description 120, 123
 - overview 3
 - reducing containers 137
 - databases
 - about 3
 - accessing in a single transaction 204
 - description 3
 - distributed 22
 - estimating size requirements 75
 - host system 204
 - language, selecting 344
 - nonrecoverable 25
 - recoverable 25
 - dates
 - formats 353
 - DB2 Connect
 - for multisite updates 204
 - incompatibility 285
 - DB2 Information Center
 - updating 383
 - versions 382
 - viewing in different languages 382
 - DB2 sync point manager (SPM) 210
 - DB2 transaction manager 206
 - DB2_LIKE_VARCHAR
 - incompatibility 285
 - DB2_NO_MPFA_FOR_NEW_DB 117, 144, 197
 - DB2_OPT_MAX_TEMP_SIZE 83
 - DB2_PARALLEL_IO registry
 - variable 164
 - DB2_SMS_TRUNC
 - _TMPTABLE_THRESH 83
 - DB2_SMS_TRUNC_TMPTABLE_THRESH 162
 - DB2_USE_PAGE_CONTAINER_TAG
 - environment variable 164
 - db2empfa command 119
 - db2empfa utility 117, 144, 197
 - db2set command 14
 - declustering
 - partial 41
 - defining
 - columns 56
 - delete rule
 - with referential constraint 65
 - dependent row 65
 - dependent table 65
 - deprecated features 243
 - descendent row 65
 - descendent table 65
 - DESCRIBE statement output
 - incompatibility 285
 - designing
 - database partition groups 87
 - tables spaces 112
 - dimension block index 177
 - dimension block indexes 180
 - dimension values
 - updating 187
 - dimensions
 - multidimensional tables 189
 - disabling
 - euro symbol support 336, 339
 - disaster recovery
 - high availability feature 23
 - discontinued features 243
 - displaying
 - Indic characters 336
 - distributed relational databases
 - units of work 22
 - distributed transaction processing
 - application program 215
 - configuration considerations 232
 - database connection considerations 219
 - error handling 227
 - resource manager 215
 - security considerations 231
 - transaction manager 215
 - updating host and iSeries databases 227
 - distributing data
 - description 41

- distribution keys
 - description 89
- distribution maps
 - description 88
- DMS (database managed space) 3, 120
- DMS device
 - buffering behavior 124
 - caching behavior 124
- DMS table spaces
 - adding containers 129
 - compared to SMS table spaces 140
 - dropping containers 137
 - extending containers 129
 - reducing containers 137
- documentation 377, 378
 - terms and conditions of use 386
- downlevel servers, tools, and clients
 - incompatibility 285
- DTP (distributed transaction processing) 215

E

- entities, database 53
- environment variables
 - profile registry 14
- estimating size requirements
 - index space 80
 - large object (LOB) data 79
 - log file space 82
 - long field data 78
- euro code page conversion tables
 - incompatibility 285
- euro symbol
 - conversion table files 339
 - enabling and disabling 336
- EXECUTE privilege
 - incompatibility 285
- extent size
 - choosing 144
 - database objects 3
 - description 112
- extents
 - extent map pages (EMP) for DMS table spaces 123
 - for SMS table spaces 119

F

- first normal form 61
- first-fit order 77
- foreign key constraint
 - incompatibility 285
- foreign key constraints
 - enforcing business rules 17
- foreign keys
 - constraints 65
- fourth normal form 61
- fragment by expression
 - comparison with table partitioning 105
- functions and procedures
 - incompatibility 285

G

- global level profile registry 14
- graphic strings
 - Unicode 360

H

- hardware environments 42
 - logical database partitions 42
 - partitions with multiple processors 42
 - partitions with one processor 42
 - single partition, multiple processors 42
 - single partition, single processor 42
 - types of parallelism 42
- health monitor
 - description 30
- help
 - displaying 382
 - for SQL statements 381
- heuristic decisions 227
- heuristic operations
 - resolving indoubt transactions 227
- high availability disaster recovery (HADR)
 - database design considerations 71
 - overview 23
- historical data, design considerations 71
- host databases
 - updating with XA transaction managers 227
- host variables
 - incompatibility 285

I

- I/O considerations
 - table space 141
- I/O parallelism 37
 - using RAID devices 164
- IBM TXSeries CICS
 - configuring 236
- IBM TXSeries Encina
 - configuring 236
- IBMCAITGROUP 112
- IBMDEFAULTGROUP 112
- IBMTMPGROUP 112
- identifying candidate key columns 58
- identity columns
 - overview 60
- identity sequence 351
- IMPLEMENTED column
 - incompatibility 285
- incompatibilities
 - COLNAMES (planned) 243
 - description 243
 - FK_COLNAMES (planned) 243
 - PK_COLNAMES (planned) 243
 - planned 243
 - Version 8 285
 - with previous releases 260
- index keys 3
- index space
 - estimating size requirements for 80

- indexes
 - block-based 175
 - description 3
 - dimension block 177
 - unique 3
- Indic characters
 - displaying 336
- indoubt transactions
 - recovering 213, 215
 - resolving 227
 - resynchronizing 213
- Information Center
 - updating 383
 - versions 382
 - viewing in different languages 382
- informational constraints
 - description 65
- insert rule with referential constraint 65
- instance level profile registry 14
- instance profile registry 14
- instances
 - description 3
- inter-partition parallelism
 - used with intra-partition parallelism 37
- inter-query parallelism 37
- intra-partition parallelism
 - used with inter-partition parallelism 37
- intra-query parallelism 37
- iSeries databases
 - updating with XA transaction managers 227

J

- joins
 - paths 54

K

- key columns
 - identifying 58
- keys
 - description 58
 - distribution 89
 - foreign 65
 - parent 65
 - partitioning, table 96
 - table partitioning 96
 - unique 65

L

- languages
 - available 313
 - compatibility between DAS and instance 344
 - DB2 supported 313
- large object (LOB) data types
 - caching behavior 124
 - column definition 56
 - estimating data size requirements 79
- Linux
 - Asian fonts 334

- LIST INDOUBT TRANSACTIONS
 - command 227
- literals
 - Unicode 364
- load utility
 - incompatibility 285
- loading data
 - multidimensional clustering tables 188
- LOB (large object) data types
 - caching behavior 124
 - column definition 56
 - estimating size requirements 79
- LOB locator switching
 - incompatibility 285
- locale coding set
 - simplified chinese 335
- locales
 - compatibility between DAS and instance 344
- locking
 - discrete 172
- log file space
 - estimating size requirements 82
- logging
 - MDC table updates 188
- logical database design
 - deciding what data to record 53
 - defining tables 54
 - relationships 54
- logical database partitions 42
- long fields
 - caching behavior 124
 - estimating data size requirements for 78

M

- maintenance
 - automatic 29
- maintenance windows
 - about 35
- map pages
 - extent 123
 - space 123
- mapping
 - table spaces to buffer pools 145
 - table spaces to database partition groups 146
 - tables to table spaces 166
- maps
 - table space 125
- materialized query tables (MQTs)
 - database design considerations 71
 - replicated 111
- MDC (multidimensional clustering) 172
- MDC (multidimensional clustering) tables 197
 - choosing dimensions 189
- messages
 - incompatibility 285
- Microsoft conversion table
 - CCSID 5035 373
 - CCSID 5039 375
 - CCSID 954 370
- mode change to tables
 - incompatibility 285

- monotonicity 197
- moving a DBCLOB
 - incompatibility 285
- moving data
 - to multidimensional tables 197
- MPP environment 42
- MQTs (materialized query tables)
 - database design considerations 71
 - replicated 111
- multi-partition database partition group 85
- multidimensional clustering (MDC) tables
 - block index considerations 188
 - block maps 185
 - choosing dimensions 189
 - creating 197
 - deletion of records 187
 - density of values 189
 - in SMS table spaces 197
 - load considerations 188
 - logging considerations 188
 - moving data to 197
 - table types 167
 - updating 187
 - using column expressions as dimensions 197
 - working with 177
- multipage_alloc configuration parameter
 - effect on memory 119
 - setting for SMS table spaces 119
- multiple partition configurations 42
- multisite updates
 - host or iSeries applications accessing a DB2 server 210
 - multiple databases 205
 - single database 204

N

- national language support (NLS)
 - bidirectional CCSIDs 347
- national languages
 - available 313
- NLS (national language support)
 - bidirectional CCSIDs 347
- node level profile registry 14
- non-Unicode databases
 - converting to Unicode 362
- nonrecoverable databases
 - backup and recovery 25
- nonthread safe library support, incompatibilities 285
- normalizing tables 61
- NOT NULL constraints 17
- notices 387
- NULL value
 - in column definitions 56

O

- OBJCAT views
 - incompatibility 285
- offline maintenance
 - about 36
- online maintenance
 - about 36

- ordering DB2 books 380
- organizing data
 - approaches 99

P

- parallelism
 - and different hardware environments 42
 - and index creation 37
 - database backup and restore utilities 37
 - I/O 37, 164
 - inter-partition 37
 - intra-partition
 - description 37
 - load utility 37
 - overview 41
 - query 37
 - utility 37
- parent key 65
- parent row 65
- parent table 65
- partial declustering 41
- partitioned databases
 - description 41
 - transaction access 219
- partitioned tables
 - description 104
 - restrictions 104
 - table partitioning 93
- partitioning data
 - table partitioning 93
- partitioning keys, table
 - description 96
 - guidelines 96
- partitions
 - compatibility 91
 - with multiple processors 42
 - with one processor 42
- partitions, data
 - description 92
- pattern matching
 - Unicode databases 364
- performance
 - table space 164
- permissions 21
- physical database design 73
- precompiler and host variable
 - incompatibility 285
- primary indexes 58
- primary keys
 - constraints 17
 - description 58
 - generating unique values 60
- printed books
 - ordering 380
- privileges
 - planning 21
- problem determination
 - online information 385
 - tutorials 385
- profile registry 14

Q

- queries
 - multidimensional clustering 189
 - parallelism 37
- query performance
 - block indexes 180

R

- RAID (Redundant Array of Independent Disks) devices
 - optimizing table space
 - performance 164
- range partition
 - see data partition 92
- range partitioning
 - see table partitioning 93
- range-clustered tables
 - advantages 168
 - comparison with other table types 167
 - description 168
 - out-of-range record keys 171
 - table locks 172
- range-partitioned tables
 - partitioned tables 104
- record deletion
 - from an MDC table 187
- recoverable databases
 - description 25
- recovery
 - strategy overview 25
- Redundant Array of Independent Disks (RAID)
 - optimizing performance 164
- reference types
 - description 56
- referential constraints
 - description 65
- referential integrity
 - constraints 65
- registry variables
 - DB2_NO_MPFA_FOR
 - _NEW_DB 117, 144, 197
 - DB2_OPT_MAX_TEMP_SIZE 83
 - DB2_SMS_TMPTABLE
 - _THRESH 161
 - DB2_SMS_TRUNC
 - _TMPTABLE_THRESH 83
 - DB2_SMS_TRUNC_TMPTABLE
 - _THRESH 162
 - environment variables 14
- regular tables 167
- relationships
 - many-to-many 54
 - many-to-one 54
 - one-to-many 54
 - one-to-one 54
- release to release incompatibilities
 - description 243
- remote unit of work
 - updating a single database 203
- reorganization
 - automatic 31
- replacing unicode conversion tables
 - CCSID 5026 371

- replicated materialized query tables 111
- resolving indoubt transactions 227
- resource managers (RM)
 - described 215
 - setting up a database as 219
- root types 56
- rows
 - dependent 65
 - descendent 65
 - parent 65
 - self-referencing 65

S

- savepoint naming
 - incompatibility 285
- scalability 42
- schemas
 - description 3
- scope
 - reference type 56
- scrollable cursors
 - incompatibility 285
- second normal form 61
- security
 - authentication 20
 - data 20
 - database design considerations 71
- self tuning memory
 - description 30
- self-referencing row 65
- self-referencing table 65
- SET INTEGRITY
 - incompatibility 285
- set integrity pending state 65
- Shift JIS X0213 code page
 - previous conversion tables 366
- simplified Chinese
 - locale coding set 335
- single partition
 - multiple processor environment 42
 - single processor environment 42
- size requirements
 - estimating 75
 - temporary tables
 - estimating 83
- SMP cluster environment 42
- SMS (system managed space) 3
 - table spaces
 - compared to DMS table spaces 140
 - descriptions 117
- SNA (Systems Network Architecture)
 - updating databases 210
- snapshots
 - storage 146
- space map pages (SMP), DMS table spaces 123
- SPM (sync point manager) 207
- SQL optimizer 3
- SQL statements
 - displaying help 381
- SQLDBCON configuration file 12
- statistics collection
 - automatic 31, 33
- statistics profiling
 - automatic 31
- STMG_CONTAINER table 148
- STMG_CURR_THRESHOLD table 148
- STMG_DATABASE table 148
- STMG_DBPARTITION table 148
- STMG_DBPGROUP table 148
- STMG_HIST_THRESHOLD table 148
- STMG_INDEX table 148
- STMG_OBJECT table 148
- STMG_OBJECT_TYPE table 148
- STMG_ROOT_OBJECT table 148
- STMG_TABLE table 148
- STMG_TABLESPACE table 148
- STMG_TBPARTITION table 148
- STMG_THRESHOLD_REGISTRY table 148
- storage management snapshots 146
- storage management tool
 - storage management view 146
 - stored procedures 147
- storage management view
 - tables in 148
- storage management thresholds 160
- storage objects
 - buffer pools 3
 - container 3
 - table spaces 3
- storage requirements
 - XML documents 84
- stored procedures
 - for storage management tool 147
- strings
 - Unicode comparisons 364
- stripe sets 125
- structured types
 - database design considerations 71
 - in column definitions 56
- SUBSTR function
 - incompatibility 285
- subtypes
 - inheritance 56
- supertypes
 - in structured type hierarchies 56
- surrogate characters
 - Unicode 355, 357
- sync point manager (SPM)
 - description 207
- SYSCAT views
 - incompatibility 285
- SYSCATSPACE table spaces 112
- SYSPROC.CAPTURE_STORAGEEMGMT_INFO stored procedure 147
- SYSPROC.CREATE_STORAGEEMGMT_TABLES stored procedure 147
- SYSPROC.DROP_STORAGEEMGMT_TABLES stored procedure 147
- system catalog tables
 - description 3
 - estimating initial size 76
- system managed space (SMS) 3, 117
 - described 119
- system temporary table spaces 112
- Systems Network Architecture (SNA) 210
- SYSTOOLSPACE table spaces
 - uses 115

SYSTOOLSTMPSPACE table spaces
uses 115

T

table partitioning
benefits 93
description 93

table partitioning keys
description 96
guidelines 96

table spaces
catalogs 112, 163
choice by optimizer 112
database managed space (DMS) 120
description 3
design
description 112
OLTP workload 143
query workload 143
workload considerations 143

disk I/O considerations 141

DMS 123

mapping to buffer pools 145

mapping to database partition groups 146

maps 125

OLTP workload 143

performance 164

query workload 143

SYSCATSPACE 112

system managed space (SMS) 117

temporary 112, 161

TEMPSPACE1 112

types
SMS or DMS 140
user 112

USERSPACE1 112

workload considerations 143

tables
append mode 167
check constraints
types 65
collocation 91
dependent 65
descendent 65
description 3
estimating size requirements 75
introduction 167
mapping to table spaces 166
multidimensional clustering 167
multidimensional clustering (MDC) 173
normalization 61
parent 65
partitioned 93, 167
partitioned tables 104
range-clustered 167, 168
regular 167
self-referencing 65
system catalog 76
temporary 162
transition 70
user 77

targets
rows 56
tables 56

targets (*continued*)
types 56
views 56

temporary table spaces
design 112
recommendations 161

temporary tables
size requirements 83
SMS table spaces 162

temporary work spaces
size requirements 83

TEMPSPACE1 table space 112

terms and conditions
use of publications 386

territory codes
DB2 supported 313

Thai characters
sorting 352

third normal form 61

thresholds
about 160

time
formats
description 353

TPM values 221

TPMONNAME values 221

transaction managers
BEA Tuxedo 238
DB2 transaction manager 206
distributed transaction
processing 215
IBM TXSeries CICS 236
IBM TXSeries Encina 236
IBM WebSphere Application Server 236
multiple database updates 205
problem determination 235
XA architecture 233

transaction processing monitors
BEA Tuxedo 238
configuration considerations 232
IBM TXSeries CICS 236
IBM TXSeries Encina 236
security considerations 231

transactions
accessing partitioned databases 219
description 22
global 215
loosely coupled 215
non-XA 215
tightly coupled 215
two-phase commit 215

triggers
business rules for data 17
cascading 70
description 70

troubleshooting
online information 385
tutorials 385

tutorials
troubleshooting and problem determination 385
Visual Explain 385

Tuxedo
configuring 238

two-phase commit
error handling 213

two-phase commit (*continued*)
process 210
updating
a single database in a
multi-database transaction 204
multiple databases 205

TXSeries CICS 236

TXSeries Encina 236

type 1 connection
incompatibility 285

type hierarchy 56

typed tables
database design considerations 71
description 56

typed views
description 56

U

UCS-2
see Unicode (UCS-2) 355

UDFs (user-defined functions)
description 56

uncommitted units of work on UNIX
incompatibility 285

Unicode (UCS-2) 355
CCSID 357
character strings 360
code page 357
constants 364
conversion tables 368
converting code page 1394 to
previous conversion tables 366
converting Shift JIS X0213 to
previous conversion tables 366
database 362
DB2 supported 357
graphic strings 360
literals 364
pattern matching 364
string comparisons 364
surrogate characters 355

unicode conversion table
CCSID 5035 372
CCSID 5039 374
CCSID 954 369

uniprocessor environment 42

unique constraints
about 17
definition 65

unique keys
description 58, 65

units of work (UOW) 22
remote 203

update rule, with referential constraints 65

updates
DB2 Information Center 383
Information Center 383
user table page limits 77
user table spaces 112
user temporary table spaces
designing 112

user-defined functions (UDFs)
description 56
incompatibility 285

- user-defined types (UDTs)
 - column definition 56
- USERSPACE1 table space 112
- UTF-16 355
- UTF-8 355, 357
- utility parallelism 37
- utility throttling
 - description 30

V

- variables
 - transition 70
- VERSION option
 - incompatibility 285
- views
 - description 3
- Visual Explain
 - tutorial 385

W

- WebSphere Application Server
 - configuring 236
- weight, definition 351

X

- X/Open distributed transaction processing (DTP) model 215
- XA interface
 - distributed transaction processing model 215
- XA specification 233
- XA switch 233
- XA transaction managers
 - configuration considerations 232
 - security considerations 231
 - troubleshooting 235
 - updating host and iSeries databases 227
- XML documents
 - storage 84
 - storage requirements 84
- XML storage object
 - overview 84

Contacting IBM

To contact IBM in your country or region, check the IBM Directory of Worldwide Contacts at <http://www.ibm.com/planetwide>

To learn more about DB2 products, go to <http://www.ibm.com/software/data/db2/>.



Printed in USA

SC10-4223-00



Spine information:

IBM DB2 DB2 Version 9

Administration Guide: Planning

