DB2 Universal Database Server for OS/390

What's New?

Version 6

GC26-9017-03

Note!

Before using this information and the product it supports, be sure to read the general information under “Chapter 7. Notices” on page 67.

Fourth Edition (Softcopy Only May 2000)

This edition applies to Version 6 of DB2 Universal Database Server for OS/390, 5645-DB2, and to any subsequent releases until otherwise indicated in new editions. Make sure you are using the correct edition for the level of the product.

This softcopy version is based on the printed edition of the book. This softcopy version of the manual contains information that was not available when the book was published.

© Copyright International Business Machines Corporation 1998, 1999. All rights reserved.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Chapter 1. About this book 1

Chapter 2. What's new in Version 6? 3
From OLTP to business intelligence and beyond 4
Web-ready and universally accessible 4
DB2 for OS/390 and business intelligence 5
More capacity for business intelligence and OLTP 6
16-terabyte tables 7
Virtual storage constraint relief 7
Better performance and availability 8
Data partitions rebalanced more easily 8
Checkpoint frequency changed dynamically 9
Maximum number of open data sets increased 9
Variable-length column lengths altered more easily 9
Better performance for SQL that uses temporary tables 10
Faster, more parallel, and more usable utilities 10
Enhanced query performance 14
Performance and availability improved even more 15
Data sharing enhancements 17
Continuous availability with group buffer pool duplexing 17
More caching options 18
Turning off space map copy maintenance on selected table spaces 19
Easier to manage 19
Predictive governing 19
Statement cost estimation 19
More information available for monitoring DB2 19
Better integration of DB2 and Workload Manager delay reporting 20
New special register for decimal precision 20
Ability to specify a default buffer pool for user data and indexes 20
Easier recovery management 20
Use of pattern-matching characters in DB2 commands 20
Batch processing of dynamic SQL 21
DB2 family compatibility and compliance with standards 21
Information to help you use DB2 22
Network computing and e-business enhancements 23
Better performance for Java applications with SQLJ 23
DRDA support for three-part names 24
Support for encrypted passwords for DRDA clients 24
Improved management of stored procedures 24
More portable applications with extensions to DB2 ODBC support 25
ODBC access to DB2 for OS/390 catalog data 26
Improved support for ODBC applications 26
Multithreading and multiple-context support 26
Improvements for dynamically prepared SQL statements 26
DDF connection pooling 27
Improved monitoring and accounting of DDF client/server systems 27
Object-relational extensions and active data 27
Object-relational extensions for new data types 28
User-defined data types for better business modeling 29
Support for user-defined functions 30
More built-in functions 31
More automation and control of business processes with triggers 32
More flexibility with SQL 33

© Copyright IBM Corp. 1998, 1999

Powerful business applications 33
Multimedia object support with DB2 Extenders 34
Enhanced business partnerships with independent software developers 34

Chapter 3. Great tools for managing your enterprise 37
DB2 Management Tools Package 37
Managing DB2 for OS/390 from your workstation with Control Center 37
Building DB2 stored procedures from your workstation 38
Installing DB2 from a workstation 39
Using workstation views of DB2 Explain output 39
Estimating DB2 performance with DB2 Estimator 40
DB2 Web applications and Net.Data for OS/390. 41
Enterprise data that is enabled for the World Wide Web 41
Enhancements to Net.Data for OS/390 41
Serving business knowledge to your enterprise with QMF 42
QMF Version 6 42
QMF High Performance Option 43
QMF for Windows 43
Data replication with DB2 DataPropagator 44
What you can do with DB2 DataPropagator 44
DB2 DataPropagator and IBM's data replication solution 46
DB2 DataPropagator Version 6 46
Tuning for optimum performance with DB2 Performance Monitor 47
What you can do with DB2 PM 47
DB2 PM enhancements for Version 6 47
Tuning DB2 with the DB2 Buffer Pool Tool 48
Managing your DB2 subsystems with the DB2 Administration Tool 49
What you can do with DB2 Admin 49
Version 6 enhancements to DB2 Admin 50

Chapter 4. What's coming next? 51
More ways to extend the power of your database 51
Improvements to performance, scalability, and availability 53

Chapter 5. Migration considerations 55

Chapter 6. Prerequisites of Version 6 of DB2 for OS/390 57
DB2 for OS/390 Version 6 prerequisites. 57
Hardware requirements 57
Program requirements and optional programs 58
Virtual storage requirements 63
Prerequisites of features of DB2 for OS/390 Version 6 63
DB2 Installer requirements 63
Visual Explain requirements 64
DB2 Estimator requirements 64
Net.Data requirements 65
QMF requirements 65
DB2 Performance Monitor requirements. 66
Workstation-Based Analysis and Tuning. 66

Chapter 7. Notices 67
Trademarks 68

Glossary 71

Chapter 1. About this book

This book provides an executive overview of Version 6 of IBM® DATABASE 2™ Universal Database Server for OS/390®.

This book refers to IBM DATABASE 2 Universal Database Server for OS/390 as "DB2® for OS/390." In cases where the context makes the meaning clear, this book uses the term "DB2" in place of the longer reference, DB2 for OS/390.

Although the focus of this book is on Version 6 of DB2 for OS/390, IBM has also made many of the functions and capabilities available to existing Version 5 users. This book identifies these functions and capabilities with the following pair of graphic symbols.

V5

Information about some function or capability that is available not only in Version 6, but also in Version 5.

Users of DB2 for OS/390 Version 5 or Version 6 can subscribe to iSource, an IBM service that sends notification of enhancements to the DB2 licensed program. For more information about how to subscribe to this service, see "Chapter 4. What's coming next?" on page 51.

Version 5 users can learn how to obtain these capabilities now, prior to migrating to Version 6, by visiting the following Web site:

http://www.ibm.com/software/db2os390/v5apar.html

The term "MVS™" represents the MVS/Enterprise Systems Architecture (MVS/ESA™) element of OS/390. The term "CICS®" represents CICS/ESA® and CICS Transaction Server for OS/390. "IMS™" represents IMS/ESA®. "C" and "C language" are terms that represent the C programming language.

Read "Chapter 5. Migration considerations" on page 55 to plan for your migration to Version 6. "Chapter 6. Prerequisites of Version 6 of DB2 for OS/390" on page 57 outlines the hardware and software requirements for installing and operating DB2 for OS/390.

© Copyright IBM Corp. 1998, 1999

Chapter 2. What's new in Version 6?

IBM is extending universal database to the System/390® platform. Now IBM combines the power of PC and UNIX application development and query tools with the industrial-strength, reliability, and performance of DB2 Universal Database™ Server for OS/390. (This book usually refers to DB2 Universal Database Server for OS/390 as "DB2 for OS/390" or "DB2".)

Recently you probably heard of a new class of databases—universal databases. With universal databases, you can store and query not only alphanumeric data, but also text documents, images, audio, video, and other complex objects. Universal databases are important if you want your databases to be:

- **Universally applicable**—to data warehousing, decision support, business intelligence, online transaction processing, and enterprise resource planning
- **Universally scalable**—from laptops to desktops, to workgroups, and to the largest enterprise systems
- **Universally available**—across all leading platforms, such as System/390, AS/400®, Windows™, UNIX®, and AIX®
- **Universally manageable**—through great new graphical tools for controlling database functions
- **Web-ready and universally accessible**—throughout the enterprise, over the Internet, and across computer platforms
- **Universally extensible**—with DB2 Extenders™ for text, image, audio, video, and more

Perhaps your company, like most companies, uses several different databases to manage your business. Maybe you've had difficulty finding one database management system that handles all types of applications well. Now, the DB2 family of database products makes it simple to consolidate your server strategy. With enhancements available in Version 6 of DB2 for OS/390, IBM provides compatible DB2 solutions for all major IBM and non-IBM operating systems.

You want the security of knowing that the database you're investing in for tomorrow's applications will continue to serve today's business needs. In this book, read about major new enhancements in DB2 for OS/390 Version 6, and see how your current database environment can benefit.

Version 6 gives you new data types to deploy multimedia at your own pace. However, your existing applications will benefit from new functions from the moment you turn the switch for this release. You can take advantage of performance improvements for utilities, faster restart and recovery, better query performance, greater capacity, and more built-in functions. You will enjoy a more powerful network-computing environment with better-performing Java™ applications, which can now use static SQL, with improvements to stored procedures, and with enhanced Open Database Connectivity (ODBC) support.

With Version 6 of DB2 Universal Database Server for OS/390, your DB2 database environment is easier to manage than ever before. IBM delivers many optional features for implementing Internet data connectivity, querying data, programming applications, managing and tuning databases, doing capacity planning, and replicating data. Workstation components with easy-to-use graphical interfaces simplify system management. DB2 for OS/390 Version 6 includes some no-charge features and other features that you can use as part of a "Try and Buy" program.

© Copyright IBM Corp. 1998, 1999

IBM also offers you enterprise resource planning solutions. Deploying applications on the System/390 platform is easier than ever with solutions from independent software development companies that take advantage of the power and capabilities of DB2. Use DB2 for OS/390 for your transaction processing, data warehousing, business intelligence, and Web applications.

From OLTP to business intelligence and beyond

Power your business with DB2, a universal database for today that will take your business into the next century and millenium.

Since DB2 entered the world in 1983 on the MVS operating system, it has evolved to meet the needs of businesses worldwide. At that time, DB2's mission—and strength—was to deliver the power of the SQL language for decision support applications. Shortly thereafter, however, DB2 began to prove itself capable of delivering the performance that was required for online transaction processing. DB2 manages the ever-increasing workloads of your core business, for which subsecond response time is always the demand.At the same time, DB2 handles strategic decision support and data-mining queries that take full advantage of the powerful constructs of the Structured Query Language, or SQL. DB2 has the capacity to deliver information wherever you want it in a timely manner.

Over the years, IBM delivered DB2 on more hardware platforms in more operating environments. In 1997, IBM extended the power of DB2 when it delivered DB2 Universal Database for Windows, UNIX, and OS/2®. In 1999, the DB2 Universal Database Server family expanded again by adding the DB2 UDB name and capabilities to DB2 UDB for AS/400 Version 4 Release 4 and to DB2 UDB for OS/390 Version 6.

Web-ready and universally accessible

DB2 for OS/390 continues to enhance its role as a true enterprise database server. Version 6 makes it even easier to deliver mission-critical information to applications that run on PCs, UNIX, and Apple Macintosh workstations. Store large objects in DB2 for OS/390 and manipulate them on your workstations, as shown in Figure 1 on page 5.

Combine the power of your PC and UNIX application development and query tools with the industrial strength, reliability, and performance of DB2 for OS/390. You do not need to sacrifice security or manageability to achieve this powerful combination. Improved management of stored procedures and enhancements to DB2 Open Database Connectivity (ODBC) support combine with high-performance features. These changes let you prioritize work and efficiently allocate resources to support thousands or even hundreds of thousands of workstation users. Better connection pooling improves throughput across the network.

You have probably heard the term e-business, which refers to doing business by using the Internet. Ideally you combine the broad reach of the Internet with the vast resources of traditional information technology systems. E-business is dynamic and interactive, and its scope is broad—spanning the spectrum from private intranets, through shared extranets, to the public Internet. E-business uses the Web to bring together customers, vendors, suppliers, and employees in a way that was never before possible. In short, e-business efficiently connects valuable information to the people who need it. Real challenges exist, too—like security, scalability, and reliability—but they are surmountable. And you don't need to start from scratch. IBM offers e-business solutions that use DB2. See "Network computing and e-business

enhancements” on page 23 for a discussion of new enhancements in DB2 for OS/390 Version 6 that make doing e-business easier than ever.

Organizations throughout the world are making Java the foundation for their e-business strategies. DB2 for OS/390 provides Java Database Connectivity (JDBC), which delivers the promise of Java as the means of developing platform-independent e-business solutions. DB2 Connect™ delivers access for your remote Java applications. You can now develop full-function Java applications, which use either dynamic or static SQL, for fast access to your DB2 data. For more information, see “Better performance for Java applications with SQLJ” on page 23.

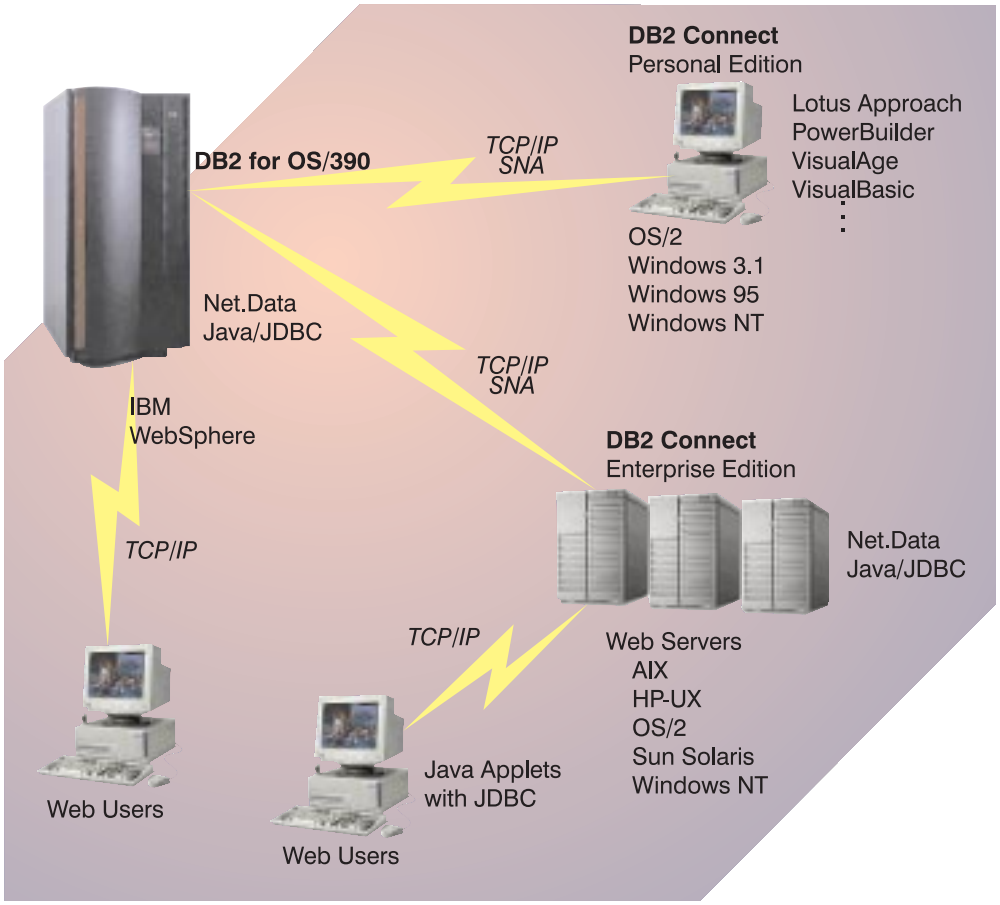

Figure 1. The power of DB2. Data where you want it.

Doing business on the Web also means being able to run Java applets from your customers’ browsers. DB2 for OS/390, with the Java applet support of DB2 Connect, delivers business data and processes transactions from any Java-enabled Web browser.Your customers can access data and process transactions on the Web without installing any special software on their machines.

DB2 for OS/390 and business intelligence

DB2 for OS/390 gives you a competitive edge for your business intelligence systems. DB2 is strongly integrated with the System/390 platform, which provides it advantages such as hardware compression and the powerful workload balancing capability of Workload Manager (WLM). DB2 has demonstrated that it is highly scalable in its efficient use of clustering CMOS technology in a Parallel Sysplex® environment.

In Version 6, the following enhancements demonstrate DB2's continued commitment to System/390 as a premiere platform for your business intelligence applications:

- The ability to have a single DB2 table that is up to 16 terabytes in size, as described in "16-terabyte tables" on page 7. This enhancement also increases the size for large objects *(LOBs)* to 4000 terabytes for each LOB column, as described in "Object-relational extensions for new data types" on page 28.
- Additional built-in functions that make it easier for DB2 to handle queries, as described in "More built-in functions" on page 31.
- Query enhancements such as:
  - Faster outer join
  - Enhanced Cartesian join
  - Ability to use indexes in more cases
  - Greater opportunities for parallelism for:
    - Outer joins
    - Queries that use nonpartitioned tables

  DB2 provides leading-edge query technology with its robust cost-based optimizer, which handles automatic query rewriting and advanced techniques for selecting indexes. In Version 6, DB2 continues to respond to the need for efficient processing for complex queries. See "Enhanced query performance" on page 14.
- Performance improvements and parallelism enhancements for the LOAD, REORG, RUNSTATS, and COPY utilities, as described in "Faster, more parallel, and more usable utilities" on page 10.
- The ability to prevent a long-running query from using too much system resource, as described in "Predictive governing" on page 19.
- The ability to put buffer pools in data spaces. Because of this enhancement, you can move the virtual buffer pools out of the DB2 address space. By moving buffers into data spaces, more virtual storage is available for processing your large, complex applications, as described in "Virtual storage constraint relief" on page 7.
- Removal of the restriction on updating values that are part of a partitioning key. Generally, performance is better when you choose partitioning key values that are not, or that are rarely, updated. However, some software products that work with DB2 operate more effectively if the partitioning key values can be updated. Removal of this restriction lets you benefit from the advantages of partitioned table spaces while you can effectively use software programs that support your business needs.

When it was introduced in 1983, DB2 was positioned primarily for queries. Over the years and in each of its ten releases, DB2 has expanded its ability to scale with larger numbers of queries, more complexity in the queries, and larger amounts of data.

More capacity for business intelligence and OLTP

Many businesses are discovering the benefits of analyzing very large amounts of data, such as historical data or point-of-sale data from a huge customer population. With DB2 for OS/390, you can grow the amount of data you can collect and query. This growth is possible because of an increase in capacity of large tables and very large memory support.

16-terabyte tables

Version 6 of DB2 for OS/390 greatly expands the capacity to store data in a single table space. As shown in Figure 2, DB2 increases the limit for storing data in a single table space to 16 terabytes. This limit is up from 1 terabyte in Version 5 and 64 gigabytes in prior releases. If your data includes large objects (LOBs), which are described in “Object-relational extensions and active data” on page 27, you can store up to 4000 terabytes in a LOB column.

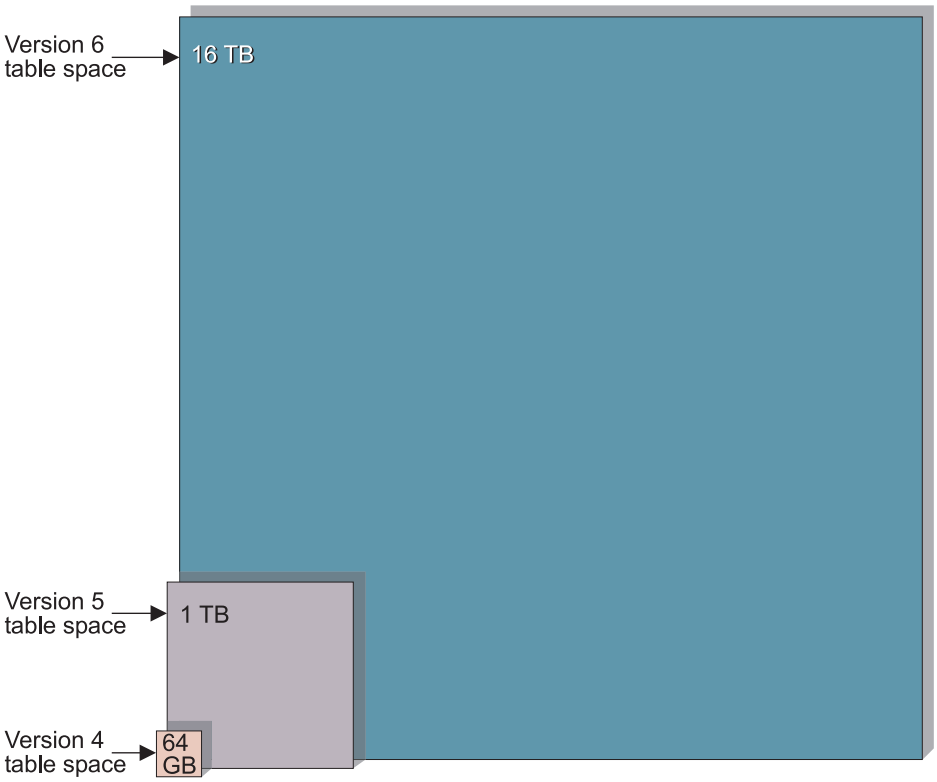

Figure 2. New limits for large table spaces

You can create tables that can be up to 16 terabytes in size, either in compressed or uncompressed format, assuming that sufficient disk space is available.

Virtual storage constraint relief

Companies today want to run larger and larger application workloads for processing online transactions and business intelligence queries. DB2 is highly scalable, making DB2 the database of choice for many of the world’s largest applications.

IBM maximizes the capacity of Version 6 of DB2 for OS/390 by providing the ability to put buffer pools in data spaces. You can make valuable DB2 virtual storage available to help alleviate problems that are caused by limited storage. DB2 can support up to 256 gigabytes of virtual buffer storage in data spaces. These larger buffer pool sizes will facilitate effective use of very large real memories in future System/390 processors.

Better performance and availability

Every release of DB2 provides more function to help you deliver mission-critical applications. Version 6 offers the following performance and availability enhancements:

- You can rebalance data partitions more easily.
- You can dynamically change checkpoint frequency.
- The maximum number of open data sets is increased.
- You can alter the length of variable-length columns more easily.
- SQL statements that use temporary tables perform better.
- Utilities can run faster, with increased parallelism, and they are easier to use.
- Queries can run faster, with increased parallelism.
- Performance for both queries and transactions continues to improve.

Data partitions rebalanced more easily

When data in a partitioned table space becomes heavily skewed, performance can be negatively affected because of contention for I/O and other resources. Using prior versions of DB2 to shift data among partitions, you needed to drop and re-create the table space. In Version 6, you can alter partitioning key values while letting applications continue to access data that is not affected by the change.

See Figure 3, which shows a table space that is partitioned by year.In the partitions at the top of the figure, the data is severely unbalanced. The result is I/O contention on the data in the partitions, which causes application response-time problems. Issue the ALTER INDEX statement as shown in the example to rebalance the partitions (shown at the bottom of the figure).

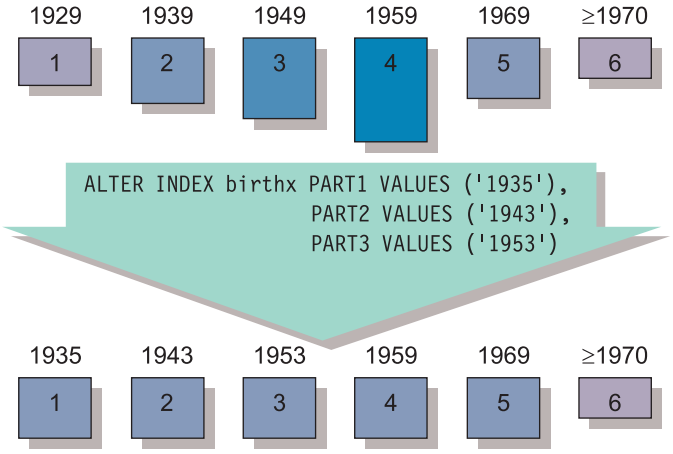

Figure 3. Rebalancing data partitions using ALTER INDEX

This enhancement can be of strategic use for table spaces that you want to grow. Define dummy partitions that have key values that will not initially be populated. Later, alter the partitioning key values for those dummy partitions to values that cause them to be populated as your data grows.

Checkpoint frequency changed dynamically

To change the checkpoint frequency in prior releases of DB2, you needed to:

1. Change the LOGLOAD subsystem parameter.
2. Stop DB2, making the DB2 subsystem unavailable for some period of time.
3. Start DB2.

As a result, specifying different checkpoint frequencies for different times of the day wasn't feasible, nor could you easily force a checkpoint when you wanted one.

In Version 6, you can issue a DB2 DISPLAY LOG command to determine the current LOGLOAD setting. Then you can issue the new SET LOG command to change the LOGLOAD subsystem parameter dynamically. You can initiate a checkpoint at any time by specifying 0 for the LOGLOAD value. This enhancement improves the availability of your DB2 subsystem.

DB2 bases its system checkpoint frequency on how many log records are written. At installation time, you tell DB2 that you want it to take a checkpoint every *'n'* number of log records that are written. Restart time is directly affected by the number of log records that are written after the latest system checkpoint. The greater the number of log records, the longer the restart time will be. Generally, a LOGLOAD value in the range of 50 000 to 800 000 provides both good performance and fast DB2 restarts with DB2 Version 6.

During prime shift, you might have a lower logging rate, but fast restart may be critical. Therefore, you might choose a lower LOGLOAD value (which means more frequent checkpoints) during prime shift. If you run batch jobs during off-shift hours, the logging rate might increase significantly. If you keep the same LOGLOAD value, your system takes checkpoints more frequently. Frequent checkpoints might not be necessary for these jobs, because restart time is not as critical during off-shift periods. When you alter the LOGLOAD value to a higher value, checkpoints happen less frequently.

V5

Maximum number of open data sets increased

DB2 installations that run in the OS/390 Version 2 Release 6 environment can now have as many as (approximately) 25 000 open DB2 data sets at one time. The maximum number of open data sets in earlier releases of OS/390 is 10 000.

Variable-length column lengths altered more easily

Now you can easily alter the length of variable-length character columns. This enhancement is particularly useful when you migrate to a new release. During migration, you sometimes need to store more data, requiring you to define larger variable-length character columns on existing tables.

Previously, altering the length of variable-length columns required several steps:

1. Creating a new table with the larger columns
2. Moving data from the existing table to the new table
3. Dropping the existing table
4. Renaming the new table
5. Re-creating authorization and dependent objects

This enhancement eliminates these steps. Now you can change the definition of an existing VARCHAR column by using the new ALTER COLUMN clause of the ALTER TABLE statement.

Better performance for SQL that uses temporary tables

In Version 6, DB2 now maintains statistics about global temporary tables, which were introduced in earlier releases. For dynamically prepared SQL statements that refer to temporary tables, the optimizer can use these statistics to choose the optimal access path. The statistics that DB2 maintains are the cardinality (number of rows) and the number of pages in the temporary tables.

Faster, more parallel, and more usable utilities

With Version 5, utility enhancements to DB2 for OS/390 demonstrated IBM's commitment to providing faster and more parallel utilities that are easier to use. Feedback from DB2 Version 5 users is very positive. The excitement continues with Version 6 as changes to several utilities improve availability of your data.

Important

All performance results are based on tests that were done in a controlled environment with an early level of code. Results you might see will vary based on your actual system configuration and other factors.

Faster backup and recovery

Improvements to several utilities reduce the amount of time that is required to back up and recover data; the result is improved data availability and faster recoveries.

Parallelism for COPY and RECOVER utilities: With Version 6, COPY and RECOVER can perform tasks on a list of image copies that are stored on DASD. The utilities can process the list in parallel, reducing the elapsed time of these jobs.

Reduced time to recover indexes: RECOVER INDEX is renamed to REBUILD INDEX. In Versions 4 and 5 of DB2, an alias REBUILD INDEX helps you prepare for the change. With Version 6, you can use the RECOVER utility to rebuild both indexes and table spaces at the same time by using image copies.You can copy and perform log-apply operations for indexes now, and you can specify a list of objects that are to be recovered. The RECOVER utility recovers indexes by restoring a copy and then applying the logs. With these enhancements, you can:

- Take a full image copy or a concurrent copy of an index
- Specify a list of indexes and table spaces to copy
- Generate a report of recovery information for indexes
- Remove the recovery information for indexes from the DB2 catalog and directory

DB2 processes the logs for all the indexes and table spaces in one pass to make your data available sooner.

Reduced elapsed time for LOAD and REORG of table spaces and REBUILD INDEX with parallel index build

DB2 Version 6 sorts, builds, and rebuilds indexes and index keys in parallel when you have defined more than one index. The parallel activity can reduce the elapsed time for LOAD or REORG of table spaces or partitions of table spaces.

Figure 4 shows how DB2 sorts and builds indexes and reorganizes a table space without SORTKEYS. In this example, DB2 uses three unload subtasks and three subtasks in which companion indexes are built. The sorting and building are done sequentially after DB2 unloads and reloads all rows and extracts all keys.

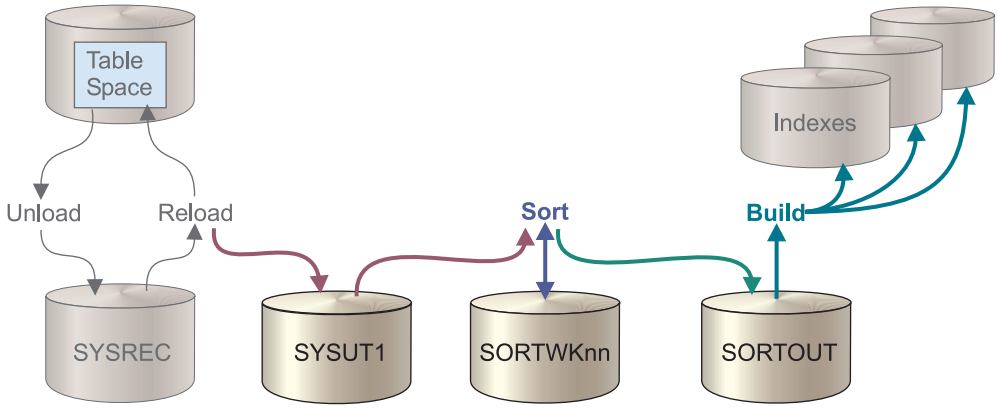

Figure 4. Rebuilding indexes and index keys without SORTKEYS

Figure 5 shows a REORG TABLESPACE utility flow with SORTKEYS. This example shows how DB2 Version 5 sorts the keys sequentially and, in parallel, reloads the rows into the table space.

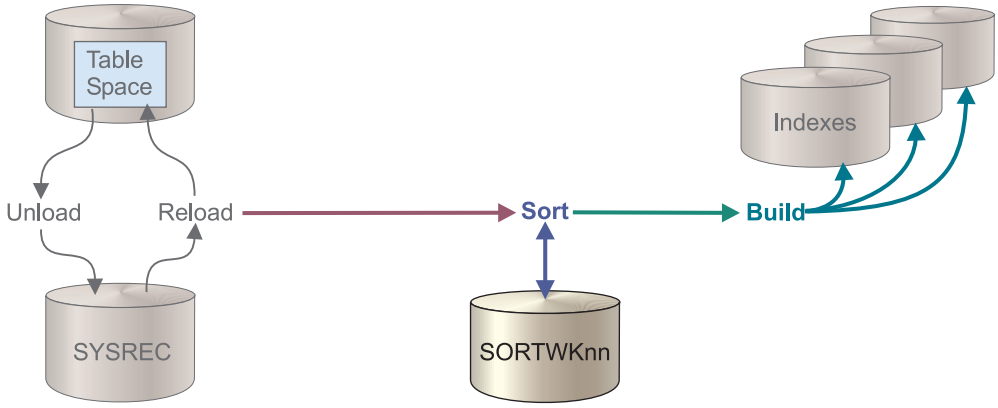

Figure 5. Rebuilding indexes and index keys with SORTKEYS

Figure 6 on page 12 shows a REORG TABLESPACE flow with parallel index build, which requires SORTKEYS. In this example, DB2 Version 6 starts multiple subtasks to sort and build the indexes in parallel.

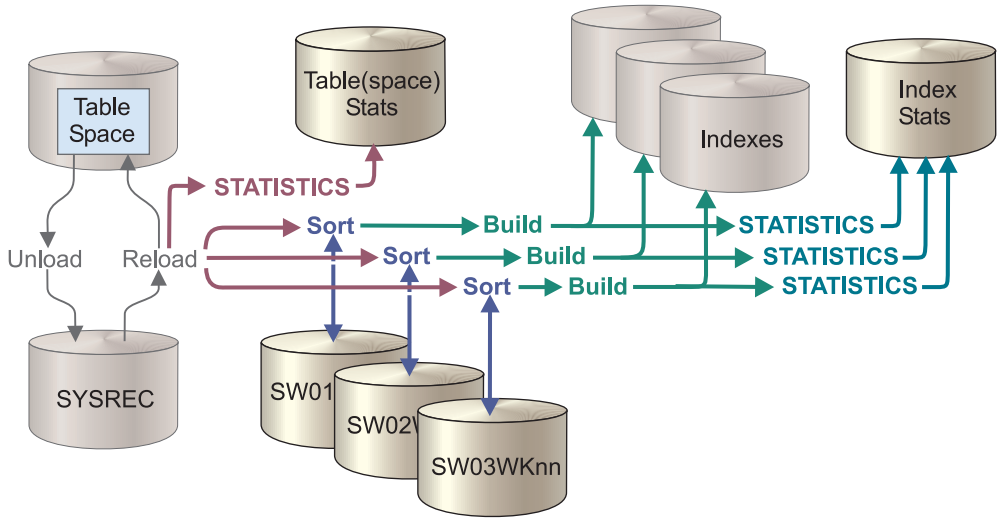

Figure 6. Rebuilding indexes and index keys using Version 6 parallel index build

Inline statistics

Prior releases of DB2 require the DBA or system administrator to update statistics by executing RUNSTATS after common utility operations, such as LOAD and REORG on table spaces. Version 6 lets you include RUNSTATS as part of the following utility jobs:

- LOAD REPLACE
- REORG TABLESPACE
- REORG INDEX
- REBUILD INDEX

Collecting statistics within utility operations avoids the need for separate RUNSTATS jobs and uses less processing resource by making fewer passes over the data.

Ability to determine if you should run REORG

With Version 6, you can avoid the cost of running REORG. You can specify threshold limits for relevant statistics from the DB2 catalog to indicate when the REORG TABLESPACE utility or the REORG INDEX utility is to run.

V5

Discard during REORG, plus faster UNLOAD

With the new options of the REORG utility, you can:

- Select rows to discard during a REORG and, optionally, write the discarded records to a file.
- Perform a faster external UNLOAD than the DSNTIAUL sample program provides.

- Specify the rows that you want to unload or discard with more flexibility than in previous releases of DB2. You can use more predicates, such as IN, NOT IN, BETWEEN, NOT BETWEEN, LIKE, NOT LIKE, IS NULL, and IS NOT NULL. You can also use more comparison operators, such as <, ≤, >, ≥, ¬=, <>, ¬<, and ¬>. Comparisons with DATE or TIMESTAMP columns can also involve simple expressions with an offset from CURRENT DATE or CURRENT TIMESTAMP. For example:
  – CURRENT DATE − 30 DAYS
  – CURRENT TIMESTAMP − 2 DAYS + 3 HOURS

Improved performance for DB2-managed data sets

Previously, to run the REBUILD, RECOVER, REORG, or LOAD utility on DB2-managed indexes or table spaces, you needed to reset the objects by deleting and redefining data sets. Now, when you run these utilities, DB2 provides an option to logically reset and reuse DB2-managed objects. This option lets you bypass the delete and redefine steps. In addition to improving performance, this option facilitates use of the storage management subsystem (SMS) to manage your data sets.

Performance test results for LOAD, REORG, and REBUILD INDEX utilities

Preliminary tests of several utilities demonstrate substantial reductions in elapsed times.

Tests of LOAD and REORG utilities: Two sets of tests measured the Version 6 LOAD and REORG utilities. The first set of tests included the following new capabilities:
- Building indexes in parallel
- Collecting statistics inline

The second set of tests excluded these new capabilities. All tests involved a table space with 10 partitions and 6 indexes. When the LOAD and REORG utilities ran while building indexes in parallel and while collecting statistics inline, the results were impressive:
- Up to a 6 times improvement in elapsed time for LOAD
- Up to a 2.5 times improvement in elapsed time for REORG

Tests of REBUILD utility: A similar comparison of the Version 6 REBUILD utility involved a variety of database configurations that DB2 installations commonly use. Tests included the rebuilding of:
- The partitioning index only
- A single nonpartitioning index only
- The partitioning index and five nonpartitioning indexes

Using the capability to build the indexes in parallel while collecting statistics inline, tests of the REBUILD utility demonstrated excellent reductions in elapsed time:
- Up to 5 times faster for a partitioning index
- Over 3 times faster for a nonpartitioning index
- Almost 8 times faster for six indexes

Decreased elapsed and processor time for online REORG

Processor time and elapsed time for online reorganizations (REORG with SHRLEVEL CHANGE) may be greatly reduced. The more disorganized the data is,

the greater the benefits, especially in the RELOAD phase. Tests in a controlled environment show that some online reorganizations take half the elapsed time as they did in Version 5.

Enhanced query performance

Version 6 of DB2 for OS/390 enhances query performance with more parallel processing.

> **Important**
> All performance results are based on tests that were done in a controlled environment with an early level of code. Results you might see will vary based on your actual system configuration and other factors.

Outer join performance improvements

The combination of join operator transformations, removal of work file steps, more aggressive view and table expression merging, and more efficient predicate processing can significantly reduce elapsed and CPU times. Early tests with customer queries and data resulted in the overall reduction of CPU and elapsed times by a factor of 30.

V5

Faster joins for string columns of different lengths

DB2 for OS/390 can now use indexes for joins on string columns that have different lengths, which is especially important for joins of tables with a million rows or more. In laboratory measurements, this enhancement improved the processing time by a factor of up to 100. This enhancement is crucial for customers who use PeopleSoft application programs.

Index support for noncorrelated IN subqueries

Version 6 of DB2 for OS/390 can use an index to access predicates with noncorrelated IN subqueries. For the following example, previous versions of DB2 were not able to use TAB1's index on PROG to access TAB1:

```
UPDATE TAB1
SET SDATE = ?, STIME = ?
 WHERE PROG IN
          (SELECT MASTER
           FROM TAB2
           WHERE INCLUDE = ?)
```

In Version 6, DB2 can use the matching index.

More query parallelism

Query parallelism is available for more queries, such as outer joins, queries that use nonpartitioned tables, and queries that use IN-list index access for the inner table of join processing.

Better workload balancing in a Parallel Sysplex environment

With Version 6, Sysplex query parallelism is much smarter at handling Parallel Sysplex systems that are composed of unlike processor models. It also performs additional analysis at run time to redistribute work when one of the following situations occurs:

- A different number of members in the group exists.
- Shortages of buffer pool resource occur.

More powerful optimizing when fetching large amounts of data

You can improve performance when fetching a large number of data rows by specifying OPTIMIZE FOR ″n″ ROWS. Enhanced support lets you request multiple DRDA® (Distributed Relational Database Architecture™) query blocks when performing high-volume operations. This enhancement is particularly beneficial for DB2 DataPropagator™, a feature of DB2 for OS/390 that replicates data.

You can optimize query answer sets that are returned in a stored procedure that a DB2 server issues. You can also optimize SELECT statements that are issued by a DRDA client.

Other query performance enhancements

Changes to the way DB2 optimizes queries provide much better performance for sophisticated applications for data warehousing and business intelligence. Among other improvements, Version 6 delivers:

- Enhanced Cartesian join processing
- Easier statistical modeling of run-time data

V5

- Index screening for list prefetch access path, which:
  - Uses less RID pool storage
  - Reduces CPU cost for sorting the RID list
  - Reduces I/O cost for accessing data

Performance and availability improved even more

Additional features in Version 6 improve performance and keep your data available.

Faster restart and recovery

A new installation option lets you postpone some of the back-out work that is traditionally performed during system restart. By postponing back-out work, you can bring up DB2 more quickly than with prior releases.

A faster log-apply process improves restart and recovery times.The new process requires fewer I/O operations and can reduce CPU time, improving the performance of:

- START DATABASE commands
- RECOVER utility
- Automatic recovery from group buffer pool outages in a data sharing environment

Important

All performance results are based on tests that were done in a controlled environment with an early level of code. Results you might see will vary based on your actual system configuration and other factors.

Tests measured the elapsed time of the log-apply processes of the Version 5 and Version 6 RECOVER utilities. The elapsed time for the Version 6 test was almost 10 time faster than for the Version 5 test.

More flexibility with 8-KB and 16-KB page sizes

In prior releases of DB2, data rows that are too large for a single 4-KB page must use a 32-KB page. Limited choices of page sizes can result in higher I/O costs, especially for random processing. For example, you could have a table with a 5-KB row that is usually randomly accessed. However, a row size of 32 KB would be too large, because the larger row size increases I/O activity and coupling facility overhead. New 8-KB and 16-KB pages give you choices that are just right to improve the balance for different processing requirements. Now you can store larger rows more efficiently and improve performance. This enhancement is beneficial for both non-data sharing and data sharing environments.

When sharing data, control the amount of traffic to the coupling facility for specific types of workloads. Controlling traffic is especially important when an application uses both random online and sequential batch processing. Reduce trips to the coupling facility for lock and group buffer pool requests by setting a large page size for sequential processing. Improve read and I/O performance by setting a small page size for random processing.

Row IDs for direct access to data rows

A new data type, ROWID, can give you the ability to re-access a row directly, without using the index or scanning the table. When an application selects a ROWID column from a table, the returned value can be saved in a host variable. Then you can use that value as the search condition of subsequent selects, deletes, or updates, which results in direct access to the row.

You can partition a table space using the ROWID column as a partitioning key, enabling you to distribute data randomly among partitions of your table space.

Ability to retain prior access path

When you rebind a statement, you almost always get the same or a better access path. When you do not, you may want to continue using an access path from a previous bind rather than the new access path. Version 6 of DB2 for OS/390 lets you retain the access path from a prior BIND by using rows in an Explain table as input to optimization.

Increased log output buffer size

DB2 increases the log output buffer size from the previous threshold of 1000 4-KB buffers to 100 000 4-KB buffers and raises the default sizes.Using larger buffer sizes can improve performance, because the log is read from the buffer rather than from log data sets. This enhancement is especially beneficial when replicating data by using DB2 DataPropagator, when performing recovery operations, during rollback, and during online REORG processing.

V5

Tracker site for disaster recovery

You can set up a tracker site that shadows the activity of a primary site. You can keep shadow data current by running periodic LOGONLY recoveries. The tracker site can be on a separate DB2 subsystem or data sharing group. Because the tracker site is shadowing the activity on the primary site, you do not need to constantly ship image copies. If a disaster occurs, the tracker site can take over more quickly than the traditional remote site, because the data at the tracker site is more current.

Option for index-only retrieval of varying-length data

Previously, DB2 could not retrieve varying-length data without going to the data page. Now, you can specify YES on a new subsystem parameter called RETVLCFK, and DB2 can return data to your application by using only the index. When RETVLCFK is YES, data is returned from the index and is padded with blanks to the maximum length of the column. If your application is sensitive to these blanks, keep the default value, NO.

Data sharing enhancements

With the release of DB2 Version 4, IBM broke new ground by providing data sharing in a Parallel Sysplex environment, improving scalability and availability. As your data grows, you can add another DB2 subsystem onto another central processor complex and access the same data through the new DB2 subsystem. Data sharing and Parallel Sysplex technology also eliminate many of the single points of failure that exist in a single-system environment. The initial delivery of the data sharing function was only the first step in realizing the full potential of high-availability characteristics. With each successive release, IBM improves on the high availability and performance characteristics of DB2 data sharing.

Version 6 is no exception. DB2 for OS/390 now delivers group buffer pool duplexing, which supports continuous availability to DB2 data in the event of a coupling facility failure. New caching options provide more ways to improve performance.

V5

Continuous availability with group buffer pool duplexing

Prior releases of DB2 rely on DASD and the merged recovery logs to recover group buffer pool (GBP) data that is lost if a coupling facility fails. Now you can choose to duplex group buffer pools. With group buffer pool duplexing, DB2 writes changed pages to both a *primary GBP* and a *secondary GBP.* Overlapped writes to the GBPs provide good performance and eliminate the writes to DASD. In the rare event of a GBP failure, DB2 can quickly recover by switching over to the secondary GBP. With this enhancement, DB2 provides continuous availability to DB2 applications that are involved in a GBP failure.

Figure 7 shows you how group buffer pool duplexing works. DB2B is a DB2 subsystem within a data sharing group. DB2 writes the changed page in buffer pool BP4 to the primary group buffer pool, GBP4, in the coupling facility. DB2 also writes the changed page to the secondary group buffer pool, GBP4-SEC. Overlapped writes to the group buffer pools (GBPs) provide good performance and eliminate the need to recover the data in a duplexed GBP if one coupling facility fails.

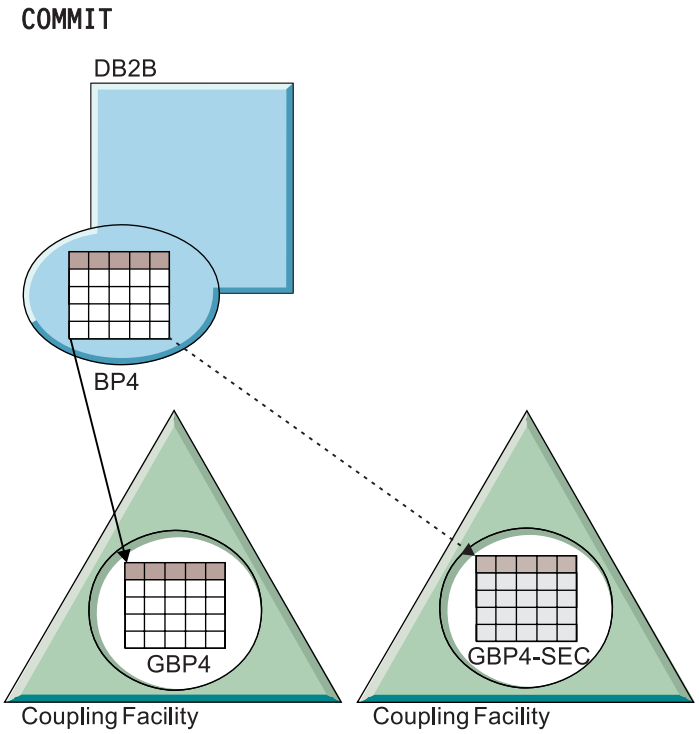

Figure 7. Duplexing group buffer pools. DB2 writes the changed data to both the primary and secondary group buffer pools.

If you have configured your coupling facilities for high availability, you don't need to obtain additional coupling facility storage. You can begin using group buffer pool duplexing right away. (To configure your coupling facilities for high availability, ensure that they have enough reserved capacity to absorb the work of any coupling facility that might fail.)

DB2 now uses advanced OS/390 and coupling facility features, which enable DB2 members to shut down faster. Use of these features also reduces occurrences of problems (coupling facility usage spikes and transient response-time delays). These problems sometimes occurred when certain DB2 system activities were triggered.

More caching options

DB2 gives you more options for using the coupling facility. One new option lets you use the coupling facility strictly for cross-invalidation purposes, not for caching data. Data is written synchronously to DASD at commit time, so using the new option results in some performance cost. However, for applications that do sequential

batch updates to a large table space, use of this option can reduce overhead costs for the coupling facility. Another option enables you to cache specific information from LOB table spaces.

Turning off space map copy maintenance on selected table spaces

You can process applications faster by choosing not to track changed pages in the space map page of the table. DB2 tracks changed pages to optimize the performance of incremental image copies. Maintaining this information in the space map may cause locking and contention. You can benefit from this option on selected table spaces. However, if your site regularly uses incremental image copies, the elapsed time for the copy will increase.

Easier to manage

IBM delivers new function to save you time when using DB2. Exercise more control over system resources with predictive governing, statement cost estimation, and improvements in monitoring. Manage your environment more easily with a new special register for decimal precision, the ability to specify a default buffer pool, and enhancements that help you manage the recovery environment. Use pattern-matching characters in some DB2 commands. Do batch-mode SQL processing with the popular DSNTEP2 sample application, without needing a PL/I license. Benefit from greater DB2 family compatibility and compliance with standards.

Predictive governing

Previously, you used DB2's resource limit facility to specify processing limits for dynamic statements. However, the facility did not stop the statement until after it had reached the limit. Therefore, the statement had already used processing resources. Now, you can use DB2's resource limit facility to stop a long-running query before it starts. Predictive governing can work in conjunction with DB2's traditional governing to help you minimize run-away queries.

Statement cost estimation

Now you can use the EXPLAIN statement to populate a new table called DSN_STATEMNT_TABLE. From this table, you can get DB2's estimate of how much processing resource is needed for any SQL statement, dynamic or static. Use this information to help you decide which statements need tuning and to determine error and warning thresholds for governing. The accuracy of the estimates depends greatly on the reliability of your catalog statistics and on other application considerations. DB2 provides help with information about the accuracy of your statistics.

More information available for monitoring DB2

Data set I/O activity that was previously available only in output from a DISPLAY BUFFERPOOL command is now available in traces, both for batch reporting and online monitors.You use a new parameter (DSSTIME) to set the interval for collecting data set statistics. As a result of this enhancement, problem analysis and resolution are more efficient, because you can now collect and review data set statistics on shorter intervals than in previous releases.

Better integration of DB2 and Workload Manager delay reporting

SMF record 72 contains information about delays in DB2, and RMF™ includes that record in reports for a particular service class or reporting period. Record 72 now contains information about:

- Locks
- Queuing of user-defined functions and stored procedures
- I/O waits
- Network delays

V5

New special register for decimal precision

The new special register provides better granularity, so that applications that need different values for the decimal precision can run in the same DB2 subsystem. Users of many PeopleSoft applications can benefit from this enhancement.

Ability to specify a default buffer pool for user data and indexes

Tune the system for application needs without impact to the DB2 catalog and directory. At installation time, define a default buffer pool for user data and indexes. This enhancements helps you ensure that the user data is stored in a buffer pool other than the one that is used for the catalog and directory. You can make better tuning decisions when user data is separate from system data.

Easier recovery management

Using several enhancements that make managing the recovery environment easier, you can:

- Recover a list of objects, as described in “Faster backup and recovery” on page 10.
- Achieve a single point of recovery more easily by specifying a set of table spaces to be quiesced.
- Specify in an installation option that COPY2 logs are to be read first for easier recovery at a remote site.
- Request that REPORT is to show data sets for one or both archive log copies, so you can determine what data sets are available at your recovery site.

Use of pattern-matching characters in DB2 commands

Use Version 6 of DB2 for OS/390 to more easily obtain valuable information about your databases. The DISPLAY, START, and STOP DATABASE commands provide extended support for pattern-matching characters (*). Using DISPLAY DATABASE, you can also filter the output by the state of the spaces. For example, you can display only those spaces that have locks, or you can show only those spaces that have claimers.

Batch processing of dynamic SQL

The DSNTEP2 sample application is a popular tool for users who want to process dynamic SQL in batch mode. However, before availability of DB2 Version 6, use of DSNTEP2 required a PL/I license, so some installations couldn't benefit from DSNTEP2.

Now, an object form of DSNTEP2 is shipped with DB2 for OS/390, so a PL/I license is no longer required. This means that more installations can benefit from this productivity aid.

DB2 family compatibility and compliance with standards

DB2 for OS/390 Version 6 makes a substantial improvement in DB2 family compatibility. Many of the capabilities that this book describes are available in other members of the DB2 family. For example, DB2 for OS/390 now joins other DB2 family members in enabling you to do an INSERT with a subselect from the same table. This lets you quickly replicate the rows in the same table or replicate some columns of data during an INSERT.

Other functions that improve compatibility, some of which are described later in this chapter, include:

- Large objects
- User-defined functions
- User-defined data types (called distinct types)
- Triggers
- Additional and enhanced built-in functions
- A VALUES clause of INSERT that can contain any expression
- A VALUES INTO statement
- An IN predicate that supports any expression
- The INCLUDE SQLDA, which is now supported for COBOL applications
- Extensions to outer join support to allow local predicates in the ON clause
- Other SQL changes, such as:
  - Optimization improvements
  - ORDER BY for a column that is not in the SELECT list
- JDBC enhancements
- SQLJ support
- Extensions to DB2 ODBC support

The functions are compliant with international and national standards, where applicable.

V5

In addition, to simplify migration of applications from other operating systems to OS/390, DB2 for OS/390 now treats data that is char(255) or varchar(255) as a short string. That is, when a character column is defined with a length of 255, the column can now participate in an index, or it can participate in a predicate evaluation.

Information to help you use DB2

The DB2 for OS/390 Version 6 library is designed to help you perform the various tasks associated with a database management system. The books in the Version 6 library are listed below.

SC26-9003	*DB2 UDB for OS/390 Administration Guide Volume 1: Design and Implementation* and *DB2 UDB for OS/390 Administration Guide Volume 2: Performance*
SC26-9004	*DB2 UDB for OS/390 Application Programming and SQL Guide*
SC26-9018	*DB2 UDB for OS/390 Application Programming Guide and Reference for Java*
SC26-9006	*DB2 UDB for OS/390 Command Reference*
SC26-9007	*DB2 UDB for OS/390 Data Sharing: Planning and Administration*
SX26-3943	*DB2 UDB for OS/390 Data Sharing Quick Reference*
LY36-3736	*DB2 UDB for OS/390 Diagnosis Guide and Reference*
LY36-3737	*DB2 UDB for OS/390 Diagnostic Quick Reference*
SC26-9650	*DB2 UDB for OS/390 Image, Audio, and Visual Extenders Administration and Programming*
GC26-9008	*DB2 UDB for OS/390 Installation Guide*
SC26-9010	*DB2 UDB for OS/390 Master Index*
GC26-9011	*DB2 UDB for OS/390 Messages and Codes*
SC26-9005	*DB2 UDB for OS/390 ODBC Guide and Reference*
SC26-9012	*DB2 UDB for OS/390 Reference for Remote DRDA Requesters and Servers*
SX26-3944	*DB2 UDB for OS/390 Reference Summary*
SC26-9013	*DB2 UDB for OS/390 Release Planning Guide*
SC26-9014	*DB2 UDB for OS/390 SQL Reference*
SC26-9651	*DB2 UDB for OS/390 Text Extender Administration and Programming*
SC26-9015	*DB2 UDB for OS/390 Utility Guide and Reference*
GC26-9017	*DB2 UDB for OS/390 What's New?*

Search DB2 books for information about new function, and stay current with the latest information about DB2 on the World Wide Web. The DB2 books are available on the Web in both PDF and HTML formats. News items keep you informed about the latest enhancements to the product. Technical articles, fact sheets, service information, product announcements, and press releases help you plan your database management strategy. Discover how DB2 customers are maximizing their own business opportunities by using DB2. View the DB2 site on the World Wide Web:

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

Network computing and e-business enhancements

With Version 5, IBM enhanced DB2 for OS/390 to support online transaction processing and mixed workload applications in client/server and network computing environments. To facilitate network computing, IBM delivered:

- Native TCP/IP (Transmission Control Protocol/Internet Protocol)
- Functions for prioritizing and scheduling
- Asynchronous transmission of query result sets for better performance
- Workload balancing
- Integration with desktop tools
- Improvements to connectivity facilities and the underlying operating system to reduce network traffic

Developing and maintaining client/server applications and doing calls from popular desktop tools are easier than ever with Version 6 of DB2 for OS/390. Use Version 6 to bring the benefits of System/390 to your workstation.

Better performance for Java applications with SQLJ

Organizations throughout the world are making Java the foundation of their e-business strategies. Take advantage of popular Java applications that are available for running on any platform, or write your own Java applications. Combine the power of Java with the latest universal database features that are available in the OS/390 environment. Java applications are not limited to Web usage. Any application that requires platform portability can capitalize on the benefits of Java.

V5

In Version 5 of DB2 for OS/390, IBM delivered Java Database Connectivity (JDBC) support to the System/390 platform. The IBM JDBC support included all the standard JDBC interfaces and enabled applications to use dynamic SQL. The dynamic SQL can access local DB2 data or remote relational data that is stored on a DRDA application server.

With Version 6, IBM 's newest Java implementation for the OS/390 environment is *SQLJ.* SQLJ is Structured Query Language (SQL) that is embedded in the Java programming language. Enjoy the benefits of Java for applications with static SQL that were previously available only for processing dynamic SQL statements. Your Java programs can take advantage of the superior performance, manageability, and authorization that is associated with static SQL. SQLJ includes support for:

- INSERT, UPDATE, and DELETE
- A SELECT that returns a single row and cursor-based SELECT
- Calls to stored procedures, including result sets

- COMMIT and ROLLBACK
- Methods for CONNECT and DISCONNECT

Advantages of using SQLJ include:

- Portable applications across platforms and database management systems
- Strong typing, with compile and bind-time schema checking to ensure that applications are well designed for the database
- The good performance and authorization checking of static SQL

  Static SQL provides control over authorization checking. You cannot manage table privileges for applications that use dynamic SQL, because every end user must have table privileges. You can avoid the problem completely by using static SQL. The application uses the table privileges of the package owner.

DRDA support for three-part names

Version 6 of DB2 for OS/390 improves the capability and performance of client/server applications with enhancements to DRDA support. Application programs using three-part names for remote access can now use the DRDA protocol. In prior releases of DB2, DRDA access was limited to using the CONNECT statement. Applications that used access by three-part names had to use DB2 private protocol.

With Version 6, IBM continues to support the DB2 private protocol, but is not enhancing the protocol with new distributed processing features. Now you can migrate your applications from the DB2 private protocol to DRDA and enjoy these DRDA benefits:

- Better performance at run time because binding occurs when the package is bound, not at run time
- Less preparing and rebinding of statements after a commit
- Access to any DRDA site through three-part names or aliases
- Access to stored procedures through three-part names
- Simplified coding for programs
- An increase in the number of remote connections to 150 000

Support for encrypted passwords for DRDA clients

DRDA level 4 provides a security mechanism that encrypts passwords by using an industry-standard encryption algorithm. Previously, if a client connected by using DRDA, the password would flow in clear-text form. Now DB2 for OS/390 supports this new security mechanism, which means that any DRDA client that chooses to do so can encrypt passwords to avoid exposing OS/390 passwords to the network. In conjunction with this support, DB2 Connect Version 5.2 (and subsequent releases) now provides support for this new security mechanism. As a result of this enhancement, you can now more easily deploy distributed applications in an untrusted network with your DB2 for OS/390 system.

Improved management of stored procedures

Use new stored procedures functions in much the same way as you use user-defined functions. Capabilities include:

- New SQL support for stored procedures:

  - New CREATE and ALTER SQL statements and enhancements to the DROP statement for creating, changing, and deleting stored procedure definitions
  - New SQL GRANT and REVOKE statements to manage execution privileges for stored procedures
  - New CURRENT PATH special register and PATH bind option to implicitly qualify stored procedure names in CALL statements
- Improved data transfer, which allows the DRDA server to send multiple DRDA query blocks. This enhancement is especially helpful with products, such as DB2 DataPropagator, that perform high-volume download operations. See “More powerful optimizing when fetching large amounts of data” on page 15 for more information.
- Support for nested calls for stored procedures and user-defined functions. For example, a stored procedure can now call another stored procedure.
- SQL CALL statements that are embedded in application programs or dynamically invoked from the DB2 Open Database Connectivity (ODBC) driver.
- New data types on the SQL CALL statement.

V5

- DB2 stored procedures now support IMS Open Database Access (ODBA). Now a DB2 stored procedure can directly connect to IMS DBCTL and access IMS data. The stored procedure does this by issuing DL/I calls via a new callable interface.
- Ability to invoke utilities from a stored procedure, which means you can now invoke utilities from an application that uses the SQL CALL statement.

You can now use the Stored Procedure Builder to develop, install, and test stored procedures by using a graphical user interface. The Stored Procedure Builder is an element of the DB2 Management Tools Package, which is a feature of DB2 for OS/390 Version 6. For more information, see “Building DB2 stored procedures from your workstation” on page 38.

More portable applications with extensions to DB2 ODBC support

Version 5 of DB2 for OS/390 introduced support for Open Database Connectivity (ODBC) and X/Open Call Level Interface. Many popular tools use this standard interface, which requires only one version of an application to access many data sources. You could issue ODBC calls or CLI calls from within a stored procedure.

With Version 6, IBM extends DB2 ODBC to support the new universal database functions (which are explained later in this chapter), including:
- Large objects (LOBs)
- User-defined functions
- Newly added built-in functions
- User-defined data types, which are called distinct types in DB2

Gain all the benefits of stored procedures by issuing ODBC calls from within your stored procedures that access data in DB2 for OS/390 and throughout your data management system.

V5

ODBC access to DB2 for OS/390 catalog data

Many of today's popular desktop tools use ODBC to access and collect DB2 catalog data. The DB2 catalog can contain information for hundreds or even thousands of DB2 tables. Most ODBC end users require access to a very small subset of the tables that are defined in the DB2 catalog. To improve performance of ODBC catalog queries, redirect the queries to shadow copies of your DB2 catalog tables.

Learn how you can use DB2 DataPropagator to capture the data and synchronize updates to the shadow copies of the DB2 catalog with the actual DB2 catalog. For a description of the DB2 DataPropagator feature, see "Data replication with DB2 DataPropagator" on page 44.

Improved support for ODBC applications

DB2 adds DRDA support for the DESCRIBE INPUT statement. This enhancement improves performance for many ODBC applications (such as PowerBuilder applications) and enables other ODBC applications.Many network messages are exchanged when applications execute dynamic SQL with input host variables without knowing in advance the correct data type of the input host variables. This enhancement reduces the number of network messages that need to be exchanged in this circumstance.

Multithreading and multiple-context support

Now, you can write multithreaded DB2 ODBC applications, and restrictions on connection switching no longer exist.

The multithreading capability makes DB2 ODBC applications threadsafe by providing private storage areas for each thread, and by properly serializing shared (global) storage areas. With this support, you can write multithreaded DB2 ODBC applications in a POSIX run-time environment.

With multiple-context support, DB2 ODBC follows the ODBC Version 2 connection model. Previously, DB2 ODBC did not support connection switching among multiple data sources unless the transaction on the connection was committed or rolled back. Now, DB2 ODBC supports context switching among multiple data sources, regardless of the commit status of the transaction.

Improvements for dynamically prepared SQL statements

A new CURRENT PATH special register and PATH bind option implicitly qualify names of user-defined data types (called distinct types), user-defined functions, and stored procedures in dynamically prepared SQL statements. The new register and bind option can be very useful for test purposes. Use them to distinguish between test and production versions of a single stored procedure within one DB2 subsystem.

DDF connection pooling

A new type of thread, a *type 2 inactive thread*, uses less storage for each thread, improving throughput in a distributed environment. DDF connection pooling provides the following benefits:

- DB2 can perform connection pooling between thousands of TCP/IP or SNA connections and a small number of database access threads.
- The CPU cost of connecting a new DDF user to DB2 is greatly reduced. This is particularly important for workloads, such as Web Common Gateway Interface (CGI) applications, that connect and disconnect to DB2 very frequently.
- DB2 lets new DRDA transactions reuse database threads.
- You can change the authorization identifier of the thread at transaction boundaries.
- DDF can now support a very large number of network connections with much less storage than is required by prior releases of DB2.

V5

Improved monitoring and accounting of DDF client/server systems

DB2 now identifies end users and end user applications, displays currently running SQL statements in DISPLAY THREAD output, and records the elapsed time that a DB2 server thread is suspended. This enhancement is particularly useful for DB2 Connect client applications.

Object-relational extensions and active data

The explosive growth of the World Wide Web demands effective use of multimedia. The Web requires that network-savvy applications, as well as image, text, voice, and video data, become a major part of your information base. To gain business advantage, you need to process your new information assets as effectively as you process your traditional data.

Now you can store large multimedia objects in DB2 tables, defining them as large objects. Use new and enhanced SQL statements and other enhancements to define and manipulate these large string values. DB2 object extensions include support for:

- Large objects (LOBs)
- User-defined functions
- New and extended built-in functions
- User-defined data types, called distinct types
- Triggers

New universal database capabilities support more than multimedia applications. These mechanisms that DB2 provides are also very useful to support all kinds of traditional applications. Use user-defined functions, user-defined data types (distinct types), built-in functions, and triggers to model any database application.

Object-relational extensions for new data types

With Version 6 of DB2 for OS/390, business professionals can use and manipulate large objects on their workstations and store the large objects in DB2 tables. DRDA protocols support the transfer of LOBs across the network. After the client application manipulates the object, the application returns the object to the server for storage.

DB2 capacity for storing LOBs

DB2 Version 6 provides capacity characteristics that exceed what most companies consider a large amount of data today. A single LOB column can be 2 gigabytes (GB), and the collection of all LOB values for a given LOB column can be up to 4000 terabytes (TB). Figure 8, which contains a list of some representative object sizes for comparison purposes, shows the importance of storage capacity.

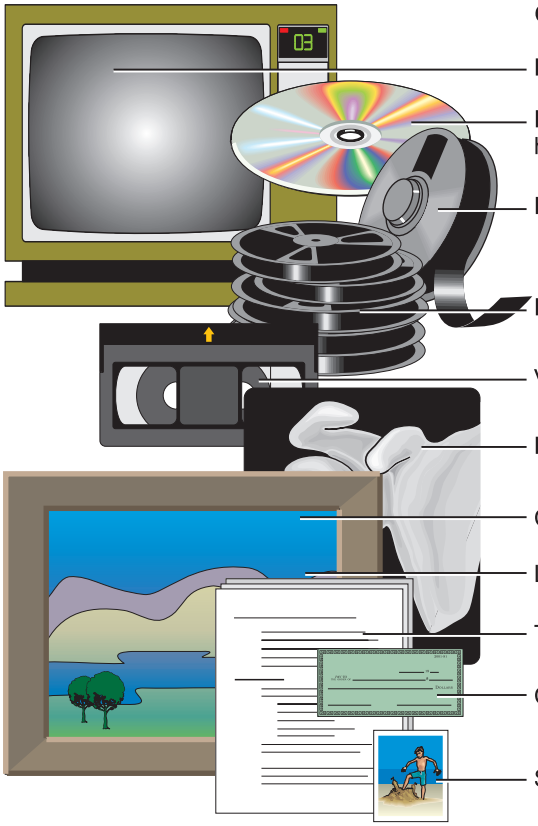

Object	Typical Size
High-definition television	200 MB/second
Feature-length, high-resolution movie	5-6 GB
High-resolution video	3 GB/hour
Feature-length movie	2 GB
Video	1 GB/hour
Radiologic image	40-60 MB
Color image	20-40 MB
Large image	200 KB-3 MB
Text	30-40 KB/page
Check image	45 KB
Small image	30-40 KB

Figure 8. Typical object sizes

Data types of LOBs

DB2 provides three new data types for these data objects:

- Binary large objects (BLOBs), for storing pictures, spatial objects, and mixed media
- Character large objects (CLOBs), for storing text-oriented information that is written in single-byte character strings or mixed character strings
- Double-byte character large objects (DBCLOBs), for storing text-oriented information that is written in double-byte character sets

LOB manipulation

Refer to and manipulate LOBs using host variables, just as you would any other data type. *LOB locators* are mechanisms that allow an application program to manipulate a large object value in the database system. When you use LOB locators, the value that is associated with a given LOB locator is not materialized until absolutely necessary. When the value that is associated with the LOB locator is materialized, it is materialized only at the server. Therefore the resulting LOB value may never need to be transmitted between the server and the client.

Your application can use LOB values or LOB locators in any of the string built-in functions that DB2 for OS/390 supports. An application that cannot acquire large enough buffers to contain the entire LOB value can use the SUBSTR function to read pieces of the LOB.

User-defined data types for better business modeling

Gain the advantages of encapsulating business semantics into the database so that the database can manage performance with user-defined data types, which are called *distinct types* in DB2 for OS/390. With distinct types, you are no longer limited to modeling business data with built-in data types. DB2 provides the capability to define a unique data type, but one that shares its internal representation with an existing type, called the *source* type. You can enforce business rules on the distinct type. You can issue GRANT and REVOKE statements to manage privileges for distinct types. Distinct types are useful for modeling both object-oriented and traditional applications.

Distinct types in applications

Image, video, and audio objects are stored in the database as BLOBs. A distinct type provides a way to differentiate one BLOB from another. You create distinct types with the CREATE statement. For example, suppose that you are developing an application that processes geographic features on maps. You can create a distinct type that is named *map* for map objects as follows:

```
CREATE DISTINCT TYPE map AS BLOB (1M)
```

The map data type is represented initially as a BLOB of 1 MB in length. The map data type is a distinct type. DB2 enforces *strong typing* against distinct types. Strong typing guarantees that only user-defined functions and operations that are defined on a particular distinct type can be applied to that type.

Distinct types in traditional applications

Distinct types are also useful in more traditional business modeling to ensure consistent behavior. For example, the euro currency in the European Common Market is one of the most important challenges that face commercial and government organizations of all sizes in coming years. You are probably still considering how you will handle currency exchange in relation to your business. However, you will likely have applications that contain both the euro and local currencies. Distinct typing can prevent costly mistakes. For example, a traditional application might erroneously allow addition of the euro and the German mark. Figure 9 on page 30 shows how to avoid the error by creating distinct types for both euro and marks. With distinct types, you cannot erroneously add euro and marks.

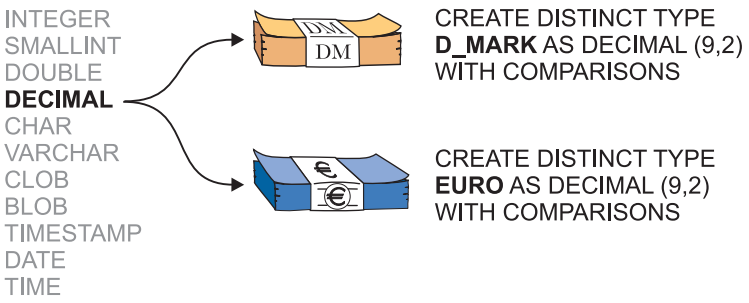

Figure 9. Distinct types in traditional applications. Distinct types, in combination with user-defined functions, ensure consistent behavior in traditional applications.

Support for user-defined functions

To define your own functions, you can use DB2 Extenders, add functions from software vendors, or write your own extensions to the capability that SQL provides. (For more information on DB2 Extenders, see “Multimedia object support with DB2 Extenders” on page 34.)

A user-defined function is very similar to a stored procedure. To define your own functions, use an SQL CREATE FUNCTION statement. The statement specifies, among other things, the input data types. For example, the following statement creates a user-defined function that is named *map_scale*, which calculates the scale of a map:

```
CREATE FUNCTION map_scale (map)
  RETURNS SMALLINT
  PARAMETER STYLE DB2SQL
  LANGUAGE C
  EXTERNAL NAME 'MAPSCALE'
  CONTAINS SQL;
```

In the example, the user-defined function:
- Identifies map as the input data type
- Is written in the C language
- Has a load module name of MAPSCALE
- Contains SQL statements that do not query or change tables

You can use a user-defined function in an SQL statement in the same way as you use built-in functions. In the following example, an SQL statement uses the user-defined function, *map_scale*. This user-defined function returns the scale of a map that is stored in the table column, *grid*:

```
SELECT map_scale (grid)
  FROM places
  WHERE location='SAN JOSE, CALIFORNIA'
```

DB2 provides a dozen sample user-defined functions, including sample user-defined functions for date and time formats. You can use these samples as starting points for writing your own functions. You can write an external user-defined function in C, C++, COBOL, OO COBOL, PL/I, or Assembler. You can also write a user-defined function with any tool that generates one of these languages, such as IBM’s VisualGen® or QMF™ High Performance Option (HPO). All programs must be compiled to run under Language Environment®.

The user-defined function runs in the stored procedures address space that is established in the Workload Manager (WLM) environment.

Advantages of user-defined functions

User-defined functions offer many advantages. You can:

- Invoke user-defined functions from dynamic SQL
- Customize the function specifically for your applications
- Encourage reuse of user-defined functions among programmers and authorized users
- Issue SQL GRANT and REVOKE statements to manage execution privileges for user-defined functions
- Improve performance by executing functions on the server rather than on the client
- Use user-defined functions in combination with LOBs and distinct types for processing objects that are tailored to your business needs
- Allow clients to benefit from the powerful Workload Manager environment

User-defined functions on System/390

With Version 6 of DB2 for OS/390, IBM makes the powerful capabilities of the System/390 platform available to your clients. You can:

- Issue SQL from within the user-defined function that is run in the MVS address spaces
- Have Workload Manager (WLM) manage the priority of user-defined functions
- Access non-SQL resources with the two-phase commit support of the OS/390 resource recovery services (RRS)
- Execute the user-defined function under the caller's RACF® (Resource Access Control Facility) profile

Now your clients can run more concurrent transactions while you let WLM optimize performance. You gain the advantage of having WLM assign individual priorities, and, at the same time, you improve security for your system.

Capturing external data for processing by DB2

Process data that is not stored in the relational format, and make full use of the query capabilities of the DB2 relational database. A user-defined function can construct and return a table to DB2. You can then use the returned table in the FROM clause of an SQL statement. Table functions make it easier to access external data in SQL queries by eliminating the need to INSERT the external data into a DB2 table prior to issuing the query.

More built-in functions

Perhaps you need some additional SQL functions to be able to ask your question. Like other DB2 users, you want functions to obtain the absolute value, modulus, square root, calendar, and standard deviation, or you want them to manipulate text.DB2 provides over 60 new built-in functions, which makes your queries more flexible and powerful. For example, the new square root function returns the actual value, so that you don't need to get the data and then perform the function on it.

Among the new built-in functions are some that handle advanced trigonometry functions, such as sine and cosine, and other advanced math functions, such as degree, exponential, and power functions. IBM has improved the power of the SQL language so you can answer questions easily. The new and extended SQL functions enhance DB2 Universal Database family compatibility across operating systems. Table 1 on page 32 shows a partial list of built-in functions.

Table 1. More built-in functions. These are just some of the popular functions, now in Version 6, that enhance compatibility across the DB2 Universal Database family.

ABS	MOD
CEILING	POSSTR
CONCAT	QUARTER
DAYOFMONTH	REPEAT
DAYOFWEEK	REPLACE
DAYOFYEAR	ROUND
FLOOR	SIGN
IFNULL	SQRT
INSERT	STDDEV
JULIAN_DAY	TRANSLATE
LOCATE	TRUNCATE
LOWER	UPPER
LTRIM	VARIANCE
MIDNIGHT_SECONDS	

More automation and control of business processes with triggers

With the addition of triggers, DB2 helps you bring application logic into the database. A trigger defines a set of actions that are to execute when a specific SQL data change operation occurs on a specified base table. The SQL data change operation includes actions that are initiated by SQL INSERT, UPDATE, and DELETE statements, as well as by actions of referential integrity constraints.

Advantages of using triggers

Using triggers, you can:

- Enforce transitional business rules that involve different states of the data. For example, use a trigger to enforce a rule that prevents a salary column from increasing by more than ten percent. The trigger invokes a comparison of the salary values before and after the increase.
- Automatically update summary data in one table when changes occur in another table.
- Initiate “alerts” outside the database, such as sending an e-mail notification or scheduling a job on another system.
- Automatically generate values for newly inserted rows.

Processing non-SQL resources with triggers

A trigger can invoke a stored procedure or a user-defined function that takes advantage of the OS/390 resource recovery services (RRS). The stored procedure or user-defined function can then initiate actions that affect non-SQL resources. If RRS controls the external resource, RRS also manages the changes as part of the same unit of work as the triggering SQL operation. For example, a trigger can invoke a stored procedure that accesses DL/I data, as shown in Figure 10 on page 33. RRS coordinates the commit between DB2 and IMS.

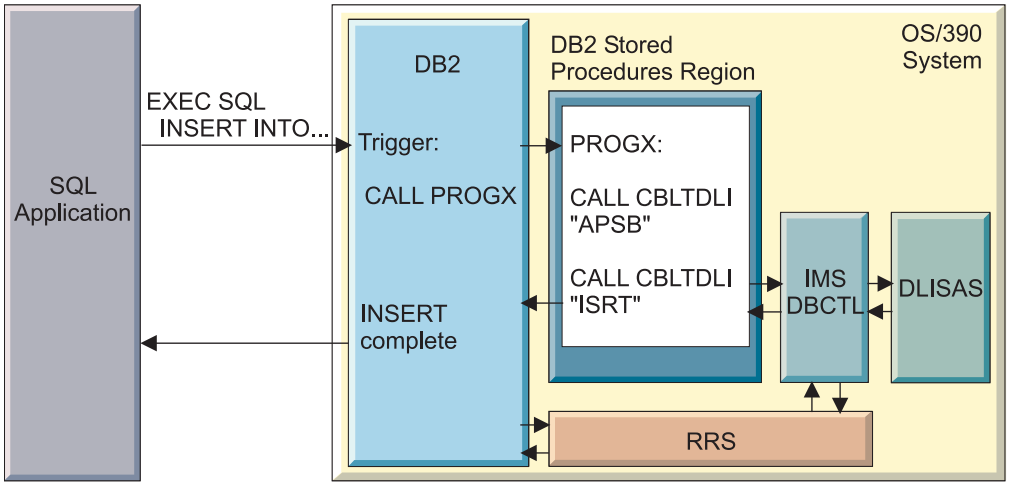

Figure 10. Process non-SQL resources with triggers

More flexibility with SQL

You have new choices when using SQL.

More tables allowed in SQL statements

In prior releases of DB2, the maximum number of base tables in a view was 15.In Version 6, the number of tables that a view can support is 225. You can also specify 225 tables in SELECT, UPDATE, INSERT, and DELETE statements. (The number of tables and views in a subselect is unchanged.)

SET assignment statement for expression values

IBM provides a new SET assignment statement, which you can use to assign a value of an expression (or the value of null) to a host variable. For example:

```
SET:lastmon = CURRENT DATE - 30 DAYS
```

Powerful business applications

Use object-oriented extensions and active data to power your business. Active data is a mechanism whereby an SQL statement can invoke an action that is not explicitly specified by the SQL statement.

For example, keep your company documents (such as images, video, and audio) in DB2 tables. Write queries to find their titles, tables of contents, and indexes. Because the documents are very large and have a complex internal structure, represent the documents by means of a distinct type, whose source type is a LOB. Define a set of user-defined functions on your document type, such as functions to extract the title of the documents, tables of contents, dates, and so on. Define functions you can use to search the content of documents.

To populate your DB2 tables, use new extensions to built-in functions. Define triggers to update summary reports when the data changes in your tables.

To perform these tasks, you can use DB2 functions, write your own functions, or use a new feature of DB2 for OS/390, DB2 Extenders.

Multimedia object support with DB2 Extenders

You can easily leverage DB2's support for large multimedia objects by using DB2 Extenders. Use DB2 Extenders to store and manipulate image, audio, video, and text objects.

The DB2 Extenders comprise a separate Image Extender, Audio Extender, Video Extender, and Text Extender. Each extender defines:

- A user-defined data type (distinct type). For example, the Image Extender defines a distinct type that is named DB2IMAGE for image objects.
- A set of user-defined functions for use with objects of its distinct type. For example, the Audio Extender defines user-defined functions to store, retrieve, and replace audio objects of a distinct type named DB2AUDIO.

The distinct types, user-defined functions, and triggers are defined when a DB2 subsystem, a DB2 table, and a column within the table are enabled for an extender. The extender does the work, eliminating the need to define distinct types, user-defined functions, and triggers for objects that you store in the column.

The extenders also automatically capture and maintain a variety of attribute information about each object that you store. The information ranges from the number of colors in an image object, to the number of recorded channels in an audio object, to the compression format of a video object, to the format of a text document. Use the user-defined functions that the extender defines to get this information.

The extenders also provide a rich body of application programming interfaces (APIs) that can take your applications to higher levels of sophistication. You can use these APIs to perform operations such as:

- Searching for documents that contain a specific word or phrase, and highlighting the occurrences, or "hits".
- Displaying images or playing audio and video clips.
- Searching for images by their visual characteristics, such as by their distribution of colors or by patterns within the image. To perform the search, the Image Extender uses IBM's leading-edge Query-by-Image-Content technology.
- Creating video storyboards that summarize the content of a video clip. Here, the Video Extender detects scene changes in a video object and identifies video frames that represent each distinct shot in the video clip.

Enhanced business partnerships with independent software developers

In today's competitive business environment, deploying enterprise applications requires that organizations build on a solid database foundation. You need applications that are capable of evolving to meet the changing needs of your business. You can develop these applications yourself, or you can purchase and implement applications from a variety of software companies. If your enterprise data is in DB2, however, you want these applications to take advantage of the power, reliability, and industrial-strength performance of DB2.

Working closely with IBM, leading software developers create enterprise solutions that are designed to leverage the reliability and industrial-strength performance of DB2. Several software developers (such as Baan, Dassault, SAP, Siebel Systems, and PeopleSoft) work closely with IBM's DB2 development team to identify new

functions and capabilities for DB2 users. Many of the new Version 6 enhancements (listed below), which you've read about in this chapter are the result of understanding the technical requirements of the applications of these independent software development companies.

- **Scalability**
  - Faster, more parallel utilities and SQL
  - Data spaces for buffers, more page sizes
  - 16-terabyte tables, partitions up to 64 GB
  - Optimization improvements
- **Availability**
  - ALTER partition range and VARCHAR length
  - Group buffer pool duplexing
  - Faster, more consistent restart and recovery
- **Management**
  - Enhanced visual tools, more features
  - Improved monitoring of data
- **Access**
  - Java static SQL (SQLJ)
- **Extensibility**
  - Triggers
  - Outer join extension (local predicates in ON clause)
  - Insert with self-referencing subselect

For more information about the IBM relationships with independent software developers, see this Web site:

http://www.ibm.com/software/data/partners/ae1partners

Chapter 3. Great tools for managing your enterprise

DB2 for OS/390 Version 6 includes more tools than prior releases so your administration and application programming tasks are easier. Make DB2 data accessible on the Web. Query and propagate your data, and manage and tune your system with optional features of the server. The tools are optional features of the server that are shipped to you automatically when you order DB2 Universal Database Server for OS/390.

The following no-charge features are available:

- DB2 Management Tools Package, which includes several workstation-based tools to help you manage your DB2 for OS/390 environment
- Net.Data® for OS/390

Priced features are available in the “Try and Buy” program. This program lets you experience the power and utility of these additional features of DB2 for OS/390. Install and use them for up to 90 days without paying license charges. To obtain unrestricted access for a priced feature, simply order the “Buy” component, and install it onto your DB2 server. The following features are available in the “Try and Buy” program:

- Query Management Facility
- DB2 DataPropagator
- DB2 Performance Monitor
- DB2 Buffer Pool Tool
- DB2 Administration Tool

You can find more information about all of the features of DB2 for OS/390 Version 6 by visiting the following Web site:

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

DB2 Management Tools Package

The DB2 Management Tools Package is a new, no-charge feature of DB2 for OS/390 Version 6. The Management Tools Package is a collection of workstation-based tools you can use to work with and manage your DB2 for OS/390 environment. The elements of the DB2 Management Tools Package are:

- DB2 UDB Control Center
- DB2 Stored Procedure Builder
- DB2 Installer
- DB2 Visual Explain
- DB2 Estimator

Managing DB2 for OS/390 from your workstation with Control Center

Users of DB2 for OS/390 can now manage data in a new way.

The Control Center capability of IBM’s DB2 Universal Database Version 6 for Windows, UNIX, and OS/2 now extends support to DB2 for OS/390 Version 6. As a Java-based tool, the Control Center can run on many types of workstations and on many different operating systems. Users can now use the same tool, with its easy-to-use graphical user interface (GUI), to manage DB2 databases on OS/390, as well as on workstation servers. The GUI supports DB2 for OS/390 SQL statements (such as CREATE and ALTER), DB2 commands (such as DISPLAY, START, and STOP), and utilities (such as REORG and LOAD).

© Copyright IBM Corp. 1998, 1999

With the Control Center, users can manage the family of DB2 databases on many different operating systems. DB2 for OS/390 objects are displayed on the Control Center main window along with objects of other members of the DB2 UDB family. To initiate an action or utility to manage these objects, users select the object. For example, a user can list the table spaces of a particular database and perform the following actions, and others, on one of the selected table spaces:

- Alter (ALTER TABLESPACE statement)
- Drop (DROP statement)
- Copy (COPY utility)
- Run statistics (RUNSTATS utility)
- Check data (CHECK DATA utility)
- Get a report (REPORT utility)
- Modify (MODIFY utility)
- Load data (LOAD utility)
- Reorganize (REORG utility)
- Recover (RECOVER utility)
- Display (DISPLAY command)

The Control Center can run either as a Java application or as an application on your Web server that your Web browser can access. Because the Control Center requires DB2 Connect, the DB2 Management Tools Package provides a restricted-use copy of DB2 Connect Version 6 to satisfy this functional dependency. Control Center is packaged along with DB2 Connect on the same CD.

The Control Center approach to managing DB2 is now extended to the System/390 platform. Sites that have multiple DB2 subsystems, on the same or different operating systems, can use the Control Center as a central point of administration. Users who are more experienced in the workstation environment can manage DB2 for OS/390 more easily as a result of the graphical user interface.

Building DB2 stored procedures from your workstation

The IBM DB2 Stored Procedure Builder (SPB), an element of the DB2 Management Tools Package, provides an easy-to-use development environment for creating, installing, and testing stored procedures. With the DB2 Stored Procedure Builder, you can focus on creating your stored procedure logic rather than on the details of registering, building, and installing stored procedures on a DB2 server. The Stored Procedure Builder provides a single development environment that supports the entire DB2 family ranging from the workstation to System/390.

With DB2 Stored Procedure Builder, you can develop stored procedures on one operating system and deploy them on other server operating systems. The Stored Procedure Builder supports two commonly used languages for stored procedures—Java and SQL procedure language. For more information about DB2 for OS/390 support for these stored procedure languages, see “Chapter 4. What’s coming next?” on page 51.

The DB2 Stored Procedure Builder has a graphical user interface that guides you through the tasks with the help of basic design patterns, SQL assistants, and costing information. Use the Stored Procedure Builder to perform a variety of tasks that are associated with stored procedures, such as:

- Viewing existing stored procedures
- Modifying existing stored procedures
- Creating new stored procedures
- Running existing stored procedures
- Copying and pasting stored procedures across connections

- One-step building of stored procedures on target databases
- Customizing the settings to enable remote debugging of installed stored procedures

Because the Stored Procedure Builder requires DB2 Connect, the DB2 Management Tools Package provides a restricted-use copy of DB2 Connect Version 6 to satisfy this functional dependency. The DB2 Stored Procedure Builder is packaged along with DB2 Connect on the same CD.

Installing DB2 from a workstation

DB2 Installer is a workstation client that is delivered as an element of the DB2 Management Tools Package. DB2 Installer enhances your productivity whether you are installing DB2 for OS/390 for the first time or you are an experienced installer. From your workstation, you can:

- Install, migrate, or update DB2 for OS/390 from a graphical interface. The graphical interface illustrates the overall installation process and keeps a graphical record of how each subsystem is defined.
- Customize your DB2 subsystem as much or as little as you need to. You can define a basic subsystem quickly, or you can customize every installation option. The main windows display those parameters that you must specify, and secondary windows display the advanced options.
- Easily control DB2 parameters and run SMP/E, installation, migration, update, fallback, and sample jobs if you have a TCP/IP connection to the DB2 Universal Database Server for OS/390. You receive job status dynamically, and you can edit JCL and examine job output from the workstation.
- Install optional features of DB2 for OS/390 Version 6, including DB2 Performance Monitor and DB2 DataPropagator.

Using workstation views of DB2 Explain output

DB2 Visual Explain, a workstation client that is an element of the DB2 Management Tools Package, is an easy-to-use workstation tool that presents the output from DB2 EXPLAIN and dynamic EXPLAIN in a graphical format. Relationships between database objects, such as tables and indexes, are instantly clear as are various operations, such as table space scans and sorts. DB2 Visual Explain also includes a browser for viewing DB2 subsystem parameters.

As a result of enhancements to Visual Explain Version 6, you can:

- View statement costs in milliseconds and service units. The graph of the access page includes the DB2-estimated cost. You can include the cost when viewing explainable statements, and then you can either sort your statements by cost or filter out statements based on their costs.
- Filter explainable SQL statements from multiple plans and packages. You can list explainable statements from many plans and packages and filter them by criteria that you specify. The criteria include statement cost, access path steps (such as table space scans and sorts), and tables or indexes that your SQL statements reference. You can save filter criteria specifications for later retrieval or modification.
- Quickly generate customized reports. The report wizard lets you create customized reports on one or more explainable statements. You choose how detailed a report you want—including the statement cost, a description of the access path, and catalog statistics on the indexes and tables. You can print the report or save it as a text file.

- Specify your own qualifier for catalog tables. You can avoid using your DB2 catalog for Visual Explain queries by copying the catalog tables. Then you can specify your own qualifier to access the copied tables when Visual Explain retrieves DB2 catalog information.

If you are using Control Center, you can launch Visual Explain directly from the Control Center.

Because DB2 Visual Explain requires DB2 Connect, the DB2 Management Tools Package provides a restricted-use copy of DB2 Connect Version 6 to satisfy this functional dependency.

Estimating DB2 performance with DB2 Estimator

DB2 Estimator is an easy-to-use, stand-alone tool for estimating the performance of applications for DB2 for OS/390. DB2 Estimator is one of the elements of the DB2 Management Tools Package, which is provided with DB2 for OS/390 Version 6. Run it on your desktop personal computer, or take it with you on your portable notebook computer.

With DB2 Estimator, you can model a partial DB2 application or a complete real or planned DB2 application without requiring an actual DB2 system. From simple table sizings to a detailed performance analysis of an entire DB2 application, DB2 Estimator saves time, lowers costs, and reduces risk. You can use DB2 Estimator to investigate the impact to your production system of new or modified applications before you implement them. You can do what-if analysis to assess the impact of changes you're considering. Answer many questions, such as:

- What is the impact on my system if the transaction volume doubles?
- What is the impact if my databases increase in size?
- What is the effect on response time if I use a faster processor?
- Is my batch window large enough for my utilities?
- How much storage do I need for the new table and its indexes?

Use DB2 Estimator during all life-cycle phases of a DB2 application. During the initial design phase, you can easily:

- Determine whether your design is optimal and feasible
- Investigate alterative database designs
- Assess the impact of using triggers and different ways of structuring queries and transactions

When creating database objects, use the models you specified in DB2 Estimator as a guide for naming columns and for specifying attribute values. You can model your system using actual DB2 information by importing information from the DB2 catalog and, if available, from EXPLAIN or DB2 Performance Monitor. Eliminate problems early in the design phase.

When your application is in production, use DB2 Estimator in conjunction with tools such as DB2 PM to solve application performance problems. You can evaluate alternative SQL designs, without any risk to your production environment, before changing any production database objects. DB2 Estimator also helps you determine the impact of hardware or workload changes.

You can use DB2 Estimator for both Version 5 and Version 6 of DB2 for OS/390. It runs in any environment that supports Windows (Windows 3.1, Windows 95, Windows 98, and Windows NT). Use DB2 Estimator on any data that has been

imported from DB2 for OS/390, or you can model an application for which none of the tables, SQL, transactions, or configuration exist.

In Version 6, DB2 Estimator expands support for utilities and SQL statements to help you approximate your working environment more closely.

DB2 Web applications and Net.Data for OS/390

Net.Data for OS/390 can access your business data for the World Wide Web. With Net.Data for OS/390, you can build high-performance, business-critical Web applications with your DB2 data, as well as other data, throughout your enterprise.

Enterprise data that is enabled for the World Wide Web

Net.Data takes advantage of the System/390 capabilities as a premier platform for electronic commerce and Internet technology. Net.Data Web applications provide continuous application availability, scalability, security, and high performance. Net.Data for OS/390 has four key advantages for Web-based applications in the enterprise:

- **Publishing interactive data and providing universal access to dynamic data**

  Web applications need access to the key data in the enterprise. Net.Data provides dynamic access to DB2 family data, DRDA-enabled data sources, and HFS (hierarchical file system) flat-file data.
- **High-performance Web-enabled data applications**

  Net.Data on OS/390 benefits from the proximity of Web applications to the DB2 data source, eliminating the need to transfer large volumes of data over a network to a workstation server. Net.Data applications can call DB2 stored procedures to take advantage of DB2 performance and to reduce network traffic.
- **Availability, scalability, and security**

  Net.Data relies on the strength of the System/390 platform as a highly available, scalable network server. Net.Data works with DB2 and the Web server (IBM Internet Connection Secure Server and Domino™ Go Webserver™) for secure Web applications.
- **Robust application development function**

  Net.Data provides a powerful macro capability for robust Web application development. The Net.Data application can call DB2 stored procedures, as well as scripts and programs that are written in Perl, REXX, and C/C++. The application can also call Java applets and JavaScripts™. Net.Data provides an Open Database Connectivity (ODBC) interface for DB2.

Enhancements to Net.Data for OS/390

Net.Data Version 2.2, which is a no-charge feature of DB2 for OS/390 Version 6, now uses the following Web server application programming interfaces (APIs):

- From IBM, Internet Connection API
- From Lotus™, Domino Go Webserver API

Applications that use the Web server APIs perform more efficiently than Common Gateway Interface (CGI) applications.

With Net.Data for OS/390, your high-performance, business-critical Web applications can more efficiently use the data in your enterprise. Enhancements improve Net.Data's performance, scalability, and usability.

Better performance and scalability

Net.Data can now cache its macros, which can improve application performance, especially for large macro files. The performance of the Net.Data built-in functions has also improved. Scalability enhancements involve the processing of large result sets. The capability in Net.Data for establishing and reusing persistent connections to the local DB2 subsystem now works for applications that use the ODBC interface.

Improved usability

With the Net.Data macro language, you can now receive multiple result sets from a stored procedure. Additional built-in functions let Net.Data macros manipulate the contents of table variables, get and set Netscape cookies, and send e-mail. Net.Data can now execute as a servlet, which enables the execution of existing Net.Data macros in the IBM WebSphere™ environment. As a result, your Net.Data applications can co-exist in the WebSphere Application Server environment. Net.Data also supports the euro currency symbol.

Serving business knowledge to your enterprise with QMF

Query Management Facility (QMF), an optional priced feature of DB2 for OS/390, is the tightly integrated, powerful, and reliable tool for query and reporting within IBM's DB2 family. QMF provides new capabilities that help you access and present mission-critical data better than ever before. With QMF, you can work with data all over the enterprise. Data can be anywhere from DB2 for OS/390 and DB2 for VSE and VM, to workstation servers that run OS/2, Windows NT™, AIX and other UNIX operating systems, and to large parallel processors. When coupled with DB2 DataJoiner®, QMF allows access to non-relational and other vendor data sources.

The QMF family of integrated tools—QMF, QMF High Performance Option (HPO), and QMF for Windows—offers a total solution for you. With these tools, you can access large amounts of data, share central repositories of queries and enterprise reports, and implement tightly controlled distributed or client/server solutions. You can publish reports to the Web and view them with your favorite Web browser.

QMF Version 6

QMF Version 6 enhancements include support for:

For users

- The euro currency symbol.
- The user's current authorization ID, displayed prominently on the new QMF home panel.
- Enhanced SAVE command syntax. QMF now recognizes object name and type in context.
- A new global variable to set the default for the QMF scroll amount.
- A new edit code to replace zeros with blanks (zero suppression) in a QMF report.
- An additional edit code and associated global variable for currency formatting, enabling two different currency symbols in the same report.
- Table Editor support for column defaults in CHANGE mode.
- The ERASE TABLE command, which you can now use to erase a database alias.

For administrators and application developers

- Formal QMF administrative authority, allowing administrators unrestricted ability to display, save, erase, import, and export QMF objects that are owned by any user.
- A new confirmation panel that works with the DB2 for OS/390 Version 6 predictive governor. When a long-running query is detected, the user can choose in advance whether to run or cancel the query.
- The ISPF LIBDEF service that you can use to allocate QMF program libraries.
- User-written edit codes using Language Environment.
- Access to data on DB2 Universal Database for UNIX, Windows, and OS/2 Version 5.
- DB2 special registers.
- Consolidated installation and management information in a single book for each operating environment.

QMF High Performance Option

QMF HPO is an optional priced feature of DB2 for OS/390 Version 6. QMF HPO includes QMF for Windows (described in the next section), the QMF HPO/Manager, and the QMF HPO/Compiler.

QMF HPO/Manager consists of a group of functions that improve governing and object management capabilities, including a pre-emptive governor to analyze QMF queries. The governing capabilities enable you to establish controls that protect production applications while delivering on-demand information. QMF HPO/Manager includes many governing parameters, such as:

- Time of day
- Day of week
- Maximum number of rows to fetch
- Allowed and disallowed SQL verbs and QMF commands
- Different settings for batch and online
- Control over resource consumption based on use of QMF commands and SQL statements

QMF HPO/Compiler lets you automatically convert queries and reports into efficient programs in COBOL. The converted programs use static SQL in place of dynamic SQL, which reduces CPU consumption, DB2 catalog contention, DB2 optimizer overhead, and security concerns.

QMF for Windows

For customers with DB2 databases of many sizes, QMF for Windows provides a Windows-based, point-and-click query tool. QMF for Windows delivers several benefits in one package, which includes an intuitive quick-start user interface. Using QMF for Windows, you can automate tasks and develop powerful native Windows applications. QMF for Windows:

- Includes a full-screen table editor you can use to update enterprise data
- Offers centralized control over resource consumption
- Provides Web publishing capabilities
- Includes a robust Windows-based API to automate database query, update, and report distribution tasks

Based on DRDA, QMF for Windows provides access to DB2 database servers on all platforms via TCP/IP or SNA. To provide this access, QMF uses direct DRDA connectivity without database gateways, middleware, or ODBC drivers.

Enhancements to QMF for Windows Version 6 include support for the following functions:

- DB2 stored procedures
- DB2 predictive governor
- QMF linear procedures, which you can create and run on the QMF host and on Windows
- QMF forms, which you can now publish on the Web in HTML format
- Advanced Web features, which you can build into your QMF report with specialized variables
- DB2 UDB for AS/400 data access support
- Server load balancing for DB2 for OS/390
- Global variables
- QMF forms calculations
- IXF format, which you can now use to import and export data
- Database logon passwords, which you can now change from within QMF for Windows
- Command line parameters to automate start-up procedures and integrate with other applications, such as Visual Warehouse™
- Enhanced object tracking to locate unused objects, determine frequently accessed data sources, and identify potential problem areas
- DB2 UDB for UNIX, Windows, and OS/2 Version 5.2 (or later) Call Level Interface (CLI) support
- Support for English Wizard (from Linguistic Technologies, an IBM Business Partner), which adds a natural language query interface if English Wizard is installed
- Support for the Windows “Send To” function, which lets you use a simple menu command to place reports directly into Windows applications (such as spreadsheets and word processing documents) and on Web servers

QMF for Windows is included with the QMF HPO feature. You can also order QMF for Windows as a feature of QMF or as a feature of DB2 for OS/390 Version 6.

Data replication with DB2 DataPropagator

DB2 DataPropagator, an optional priced feature of DB2 for OS/390, provides a powerful replication capability for the DB2 family of databases. The need to keep multiple copies of the same data in separate physical databases grows as corporations build data warehouses, adopt client/server applications, or implement mobile computing. Data replication has evolved as the key technology for harnessing the full power and potential of such distributed database environments.

What you can do with DB2 DataPropagator

DB2 DataPropagator, the core component of IBM’s replication solution, lets you unite your distributed relational databases into a cohesive and integrated database solution. It automatically captures your data changes against a source database and propagates those changes to any specified target database, keeping the two consistent. With DB2 DataPropagator, you can:

- **Reengineer business processes**

  With DB2 DataPropagator, you can replicate transactional data to servers across your enterprise. You can also improve availability and responsiveness by moving data and applications to the point of each business transaction. Many replication products support data subsetting only according to information contained in the replicated data. Unlike these products, DB2 DataPropagator supports data subsetting based on join predicates or subselects, allowing you to distribute data efficiently from normalized databases.
- **Go mobile**

  DB2 DataPropagator supports the unique needs of mobile users and occasionally connected systems by accommodating the infrequent, unpredictable, and expensive connections from these systems. Specifically, DB2 DataPropagator enables on-demand replication, automates connection and disconnection, and minimizes connection time. DB2 DataPropagator also lets you initiate all data transfers from the mobile units and infrequently connected users, whether they are downloading from a central server or uploading for consolidated processing.
- **Build powerful distributed applications**

  DB2 DataPropagator supports mobile computing, which is one component of a new breed of distributed applications. The update-anywhere replication capability of DB2 DataPropagator provides rigorous conflict detection and automatic compensation for offending transactions. This capability helps you maintain integrity of your primary database and its many distributed replicas.

Improving your decision-making effectiveness

DB2 DataPropagator enables you to tailor data for maximum usability, letting you automate data enhancement as the data is copied to its target table. For example, you can:

- Derive data by using arithmetic, Boolean, or any valid SQL expression
- Aggregate data to produce sums or averages by using SQL column functions
- Convert data by translating encoded fields to descriptive fields
- Consolidate data through joins or unions
- Generate histories to support trend analysis

Manage replication with a graphical interface: An intuitive graphical user interface simplifies definition of replication scenarios including sources, targets, frequency and timing, replication-events, and pre- and post-change processing. DB2 DataPropagator automatically creates and loads target tables. You can add, delete, or change replication requests while the rest of the system continues to run.

Minimize impact on production systems and networks: DB2 DataPropagator uses a log-based change-capture technique that minimizes impact on transaction performance, thus avoiding contention with source tables and inline transaction processing. DB2 DataPropagator features optimizations to support various networked environments. You can specify distribution timing on a copy-by-copy basis, to minimize use of peak network periods or to take advantage of economy network prices.

Integrate mixed database environments: DB2 DataPropagator expands the range of solutions available to you by supporting an open architecture. In particular:

- DB2 DataPropagator is built on standard SQL to leverage the database engine functions for data enhancement, network connectivity, and data security.
- The data staging is designed to support interoperability among multi-vendor sources and targets, between relational and non-relational formats, and among products from independent software vendors.

- DB2 DataPropagator replication directly supports multi-vendor sources and targets through DataJoiner, IBM's multi-database server product.

DB2 DataPropagator and IBM's data replication solution

DB2 DataPropagator establishes the base architecture for IBM's data replication solution, which is based on individual components that work together. As changes occur in the source, the Capture component stores them in the staging table. The Apply component reads the staging area and applies those changes to targets or copies data directly from the source in full-refresh mode. The Administration component provides a user interface for defining replication requests.

DataRefresher™ and DataPropagator NonRelational (IMS) facilitate replication of non-relational data from enterprise servers. They can populate the DB2 DataPropagator's data staging area with IMS or VSAM data. DB2 DataPropagator then enhances, distributes, and applies data to the target tables to provide an end-to-end replication solution from legacy sources.

DB2 DataPropagator Version 6

In Version 6, DB2 DataPropagator delivers powerful functions:

- Highly efficient architecture for automatic capture and asynchronous propagation of data changes to the DB2 family of databases.
- High-performance log-based change-capture component which captures all DB2 updates without response-time impact or application changes.
- Powerful data enhancements that uses standard SQL, including multitable joins and stored procedures.
- Support for full-refresh, update-change replication, which includes update of denormalized copies.
- Seamless interoperation with DataPropagator NonRelational, DataRefresher, DataJoiner, and Lotus® NotesPump™. Because of this interoperation, DB2 DataPropagator helps you deliver robust, versatile replication among IMS, DB2, Oracle, Sybase, Informix, Microsoft™ SQL Server, and Lotus Notes™.
- Easier, intuitive administration from DB2 Universal Database Control Center, which runs on OS/2, Windows NT, and Windows 95.
- Subscription sets for transaction consistency.
- Replication that is based on DB2 views, including join views.
- Event-driven and continuous replication.
- Update-anywhere replication with robust conflict detection and compensation.
- On-demand replication for occasionally connected and mobile systems.

The following functions and capabilities are new with Version 6:

- DB2 Catalog replication for speeding up of ODBC-based queries
- Copying of large object data types (BLOB, CLOB, and DBCLOB)
- Replication for ROWID
- Easier installation with DB2 Installer
- Support for a mixed data sharing environment with different versions of DB2
- Better resource utilization for continuous replication scenarios
- Improved scalability, performance, and concurrency
- Support for many versions of DB2 including:
  - DB2 for OS/390 Version 6
  - DB2 for OS/390 Version 5
  - DB2 for MVS/ESA Version 4

Tuning for optimum performance with DB2 Performance Monitor

DB2 Performance Monitor for OS/390 (DB2 PM) is IBM's strategic tool for analyzing, controlling, and tuning the performance of DB2 and DB2 applications. DB2 PM is an optional, priced feature of DB2 for OS/390. It includes a real-time online monitor, a wide variety of reports for in-depth analysis, and an explain feature to analyze and optimize SQL statements. DB2 developers use it, too, for testing the performance of new DB2 versions before releasing them. So, whenever you move to a new version of DB2, you can be sure that DB2 PM supports the latest DB2 enhancements.

What you can do with DB2 PM

Use DB2 PM to:

- Monitor DB2 and DB2 applications online.
- Be alerted immediately when problems occur.
- View and examine the status of a DB2 subsystem and its applications while they are currently active.
- Investigate events and performance problems that happened in the past.
- Monitor individual data sharing members or entire data sharing groups.
- Monitor applications that are running in a Parallel Sysplex environment, even if the parallel tasks are running on different processors.
- Analyze performance problems online.
- Analyze performance problems by examining a vast set of detailed reports.
- Obtain tuning recommendations.
- Obtain information about the access paths that were used for your SQL statements. You can use this information to change your SQL for better performance.
- Recognize periodic exceptions and event exceptions and take appropriate actions by means of a user exit routine.
- Restrict user authorization to specific areas or specific threads.
- Recognize trends and anticipate potential bottlenecks.

Using these capabilities of DB2 PM in your daily business helps increase the productivity of DB2 system and application programmers. Most importantly, using DB2 PM can considerably improve the capacity, reliability, and availability of your DB2 system environment.

DB2 PM enhancements for Version 6

DB2 PM Version 6 supports full performance monitoring and problem analysis for all functions of DB2. DB2 PM even supports the new functions in Version 6 of DB2 for OS/390. For example, DB2 PM shows times and events for the newly introduced triggers and user-defined functions, and it monitors the behavior of large objects and the new ROWID data type.

If you are using Control Center, you can launch DB2 PM directly from the Control Center.

The following components and functions are new or enhanced for DB2 PM Version 6.

Workstation Online Monitor for Windows NT and OS/2

DB2 PM Version 5 introduced the Workstation Online Monitor for both the OS/2 and Windows NT environments. The Workstation Online Monitor:

- Eliminates the need to monitor through TSO
- Improves your efficiency with an easy-to-use graphical interface
- Lets you monitor multiple DB2 subsystems concurrently
- Offers improved tuning recommendations on a per-field basis
- Lets you submit DB2 commands from DB2 PM workstation windows
- Displays historical data in enhanced System Health graphics

For Version 6, the Workstation Online Monitor adds several new functions. Now you can:

- Qualify threads
- Sort threads
- Monitor the content of the cache for dynamic statements
- Report on statistics data that is calculated by delta and interval processing

Online monitoring of data sharing groups in a Parallel Sysplex environment

Now it is easier than ever to monitor your data sharing groups:

- Connect to a whole data sharing group.
- Display threads that are running in a data sharing group to obtain an overview of all threads, even if they run on several members.
- Display threads that are running on a single member.
- Display threads and statistics information for any member of a data sharing group.

This capability is planned to be delivered via service update to DB2 PM after the general availability of DB2 PM Version 6.

Additional improvements for monitoring DB2

The following enhancements make performance monitoring even more efficient:

- DDF threads are shown in more detail, such as information on the end user's user ID, the workstation name, and the transaction name. You use this information as ordering and filtering criteria to customize your report sets.
- You can monitor information on the DB2 DISPLAY BPOOL LSTATS command online or request it in statistics reports that include data set statistics for buffer pools.
- Accounting data for class 3 was grouped together under service task switch time in earlier releases. Now, the times are separated into more granular data, making analysis of the data much faster and easier.
- Statistics data is provided for cached SQL statements so you can analyze and evaluate their efficiency.
- You can install and customize DB2 PM from your workstation by using DB2 Installer.

Tuning DB2 with the DB2 Buffer Pool Tool

DB2 Buffer Pool Tool, an optional priced feature of DB2 for OS/390, helps performance analysts evaluate tuning alternatives and achieve performance gains by tuning DB2 buffer pools.

DB2 Buffer Pool Tool uses the standard DB2 IFI interface to collect specific internal performance data. The tool processes the collected data and provides information at the buffer pool and individual object levels. Obtain information such as:

- System and application hit ratios
- Getpage activity and I/O rates
- Random access, dynamic prefetch access, and list prefetch access activity
- Read and write I/O activity
- Average and maximum I/O elapsed times
- Random page hiperpool retrieval ratio

You can predict performance at buffer pool and object levels and at varying pool sizes. The tool also lets you remove objects from a pool or place them into a different pool so you can evaluate performance.

DB2 Buffer Pool Tool includes two optional workstation components. From your workstation, quickly display graphic illustrations of pool and object usage, and general pool-tuning recommendations. Generate member include and exclude lists for simulation runs, and use cluster analysis techniques to help you determine which objects to group together into buffer pools. For your data sharing environment, obtain initial sizing estimates for the coupling facility structures, using actual statistical data from your operational DB2 subsystems.

Managing your DB2 subsystems with the DB2 Administration Tool

DB2 Administration Tool (DB2 Admin), an optional priced feature of DB2 for OS/390, provides an easy way to perform queries of the DB2 catalog and to manage DB2 objects. The tool runs under Interactive System Productivity Facility (ISPF) and uses SQL to access DB2 catalog tables. The tool provides comprehensive information for system administrators, database administrators, and application developers. Use DB2 Admin to obtain a quick overview of a database, discover problems, copy tables from one DB2 to another, and perform many administration tasks with ease.

What you can do with DB2 Admin

Use DB2 Admin to display the following catalog information:

- Any object in the catalog
- Related DB2 objects using line commands
- Static SQL from application plans and packages
- Definitions of existing views

Perform many tasks by using DB2 Admin-provided ISPF panels:

- Execute dynamic SQL statements
- Issue DB2 commands for DB2 databases and table spaces
- Run most DB2 utilities

DB2 Admin helps you manage your system by letting you issue complex performance and space queries, use EXPLAIN functions, and manage SQL IDs. You can easily display threads and manage DB2's Distributed Data Facility.

With DB2 Admin, you can also reverse engineer the definitions of DB2 objects in the DB2 catalog. The starting point can be databases, table spaces, or tables. DB2 Admin generates all the dependent objects, such as:

- Table spaces
- Tables
- Indexes

- Views
- Synonyms
- Aliases
- Referential constraints
- Table checks
- Authorizations to all of the preceding objects

The generation of the SQL statements can run as a batch job or online in a TSO/ISPF session. The output from the reverse-engineering function is a ready-to-run SPUFI file. For each object that DB2 Admin reverse engineers (by extracting definitions from the DB2 catalog), the file contains the necessary SQL statements to re-create the object.

The reverse-engineering function helps you perform tasks such as copying application data from a production system to a test system (or vice versa), or capturing the current definition of the DB2 objects in your DB2 catalog.

Version 6 enhancements to DB2 Admin

With DB2 Admin Version 6, you can now also:

- Manage the new and changed DB2 for OS/390 Version 6 objects. For example, you can create, display, alter, drop, and reverse engineer schemas, user-defined functions, distinct types, stored procedures, and triggers.
- Work with tables that have LOB columns and work with the new BLOB, CLOB, DBCLOB, and ROWID data types.
- Manage predictive governor limits.
- Decide whether to continue processing of a dynamic SQL statement if it receives a predictive governor warning.

Chapter 4. What's coming next?

You've read about what is available in DB2 for OS/390 Version 6. You may also be interested to know that more DB2 capabilities are intended and that you won't need to wait until the next formal version or release of DB2 to use them. Some new capabilities have already been delivered.

Some information in this section describes IBM's preliminary plans and is subject to change at IBM's business discretion. You should not rely on such information to make business plans.

In 1998, IBM began refreshing the DB2 for OS/390 Version 5 product for its customers by delivering not only maintenance, but also significant new DB2 capabilities. Many of these capabilities are mentioned in this book (and marked with "V5" flags). You can read more about Version 5 in *DB2 for OS/390 Version 5 What's New?* An updated edition of this book is available at the following Web site:

http:///www.ibm.com.software/data/db2/os390/library.html

With DB2 for OS/390 Version 6, IBM continues to meet customer needs in a rapidly changing business environment. IBM plans to periodically refresh the DB2 for OS/390 Version 6 product with maintenance changes and some new enhancements that deliver even more power, performance, and flexibility. This approach will increase the value of your DB2 Version 6 license. Here are some of the enhancements that you can watch for.

More ways to extend the power of your database

Programmers have more choices when writing applications.

Available

SQL support for REXX

With this enhancement, application programmers can issue SQL statements from REXX programs. The SQL statements can be anywhere a REXX command can be. Programmers can also write DB2 stored procedures in REXX. The SQL interface to REXX supports almost all SQL statements that DB2 for OS/390 supports.

Available

Stored procedures that are written entirely in SQL

With SQL procedures, you can now write stored procedures consisting entirely of SQL statements. An SQL procedure can include declarations (of variables, conditions, cursors, and handlers), flow control, assignment statements, and traditional SQL for defining and manipulating relational data. These extensions provide a procedural language for writing stored procedures, and they are consistent with the PSM (Persistent Stored Modules) portion of the SQL standard.You can use the Stored Procedure Builder to build SQL stored procedures. The Stored Procedure Builder is an element of the DB2 Management Tools Package, which is a no-charge feature of DB2 for OS/390 Version 6. For more

© Copyright IBM Corp. 1998, 1999

information about the Stored Procedure Builder, see “Building DB2 stored procedures from your workstation” on page 38.

Support for Java stored procedures

With this enhancement, DB2 for OS/390 handles stored procedures that are written in the Java programming language. You can write Java stored procedures that contain either static SQL (by using SQLJ) or dynamic SQL (by using JDBC). Alternatively, you can use the Stored Procedure Builder (an element of the DB2 Management Tools Package) to build Java stored procedures. Java stored procedures that run on DB2 for OS/390 can also run on other operating systems and platforms, including those of other database vendors, without being recompiled or modified.

Available

More tables in a join

This enhancement enables you to perform joins with more than 15 tables.

Available

Savepoints

Your application can set a savepoint within a transaction. Application logic can undo the data and schema changes made since the application set the savepoint without affecting the overall outcome of the transaction. Using savepoints makes coding applications more efficient. You no longer need to include contingency and what-if logic in your applications.

Available

Support for defining identity columns

When a column of a table is an *identity column,* the column has an attribute that enables DB2 to automatically generate a unique numeric value for each row that is inserted into the table. Identity columns are ideally suited to the task of generating unique primary key values. Applications that use identity columns may be able to avoid concurrency and performance problems that sometimes occur when applications implement their own unique counters.

Available

Ability to define declared temporary tables

Declared temporary tables complement the existing created temporary tables available in Version 5 of DB2 for OS/390. Declared temporary tables do not have descriptions in the catalog tables. The tables support indexes, UPDATE statements, and DELETE statements. You can implicitly define the columns and use the result table from a SELECT.

Available

Update with subselect

You can use a subselect to determine the value that is to be used in the SET clause of an UPDATE statement. This enhancement improves DB2 family compatibility.

Improvements to performance, scalability, and availability

Several DB2 capabilities improve performance, scalability, and availability of your database.

Available

Better optimization for complex queries

Using a new star join method, DB2 for OS/390 can provide better optimization and execution performance for queries that join tables together in a star schema. A *star schema* consists of a fact table and a number of dimensions, each of which consists of one or more tables. In addition to improving execution performance, this enhancement enables DB2 to handle more complex star schemas and more tables in a join.

Available

Deferred definitions of data sets

You can defer defining data sets when you have many tables that do not have data in them. This enhancement will provide faster installation of applications and make data set management easier for SAP R/3 and PeopleSoft products.

Available

DDF suspend and resume

You can use DDF suspend and resume commands at a server to temporarily suspend activity from requesters without terminating connections. Suspending requester activity enables data definition operations at the server to complete.

Available

Log suspend and resume for external copy

The LOG SUSPEND command suspends update activity and logging while you take an external copy of your production system. The LOG RESUME command causes update activity and logging to resume. During the brief suspension, you can take a copy using a fast-disk copy facility, such as Enterprise Storage Server FlashCopy or RAMAC Virtual Array SnapShot.

You can be among the first to learn about the latest information about these and other enhancements as they are integrated into the DB2 for OS/390 Version 6

licensed program. Take advantage of iSource, an IBM service that lets you create a personalized subscription to IBM offering announcements. For information about this subscription service, visit the following Web site:

http://www.ibm.com/isource

Chapter 5. Migration considerations

Prepare for migration to DB2 for OS/390 Version 6. As stated by IBM in prior releases of the product, some capabilities from earlier versions can receive less support from the product in the future. You can migrate to Version 6 only from a Version 5 subsystem.

Important

Information is provided for planning purposes only. When migrating your subsystem, see the specific migration considerations listed in *DB2 UDB for OS/390 Version 6 Program Directory*, which is available to you as part of DB2 for OS/390 Version 6.

IBM adds many universal database capabilities to DB2 Version 6 and removes support for some capabilities. As you prepare to migrate your subsystems to Version 6, you should be aware of the following changes:

- IBM removes support for type 1 indexes. DB2 requires type 2 indexes for every index. Convert all indexes to type 2 before migrating to Version 6.
- IBM removes support for password protection in Version 6. DB2 subsystems should protect data sets by using a security subsystem, such as RACF, which is an element of the OS/390 Security Server, rather than by passwords. Remove all passwords from all indexes and table spaces before migrating to Version 6.
- Shared read-only data access is replaced by more substantial, more usable data sharing.Another alternative is to use distributed data. IBM eliminates shared read-only data in Version 6. Convert or drop all shared read-only databases before migrating to Version 6.
- IBM will probably remove support for the private-protocol access of the Distributed Data Facility in a future release of DB2; IBM has no plans for additional enhancements. DRDA replaces the private-protocol access. Private protocol does not support stored procedures and TCP/IP, both of which are available with Version 5 of DB2 for OS/390. Private protocol also does not support the network computing enhancements in Version 6.

  Recommendation: Convert your private-protocol applications to DRDA.
- Host variables in SQL statements now require a preceding colon. In previous releases of DB2, the colons are optional.
- RECOVER INDEX is renamed to REBUILD INDEX. Versions 4 and 5 of DB2 provide an alias REBUILD INDEX so you can prepare for the change. Convert utility jobs to use the REBUILD INDEX syntax.
- Using prior releases of DB2, you register stored procedures in the SYSIBM.SYSPROCEDURES catalog table. You control access by using the AUTHID and LUNAME of the caller. Using Version 6, you register stored procedures in a new catalog table by using the CREATE PROCEDURE statement. You can map your stored procedure definitions that use AUTHID to the schema and CURRENT PATH support. You must eliminate rows that control access with LUNAME prior to migrating to Version 6.

For more detailed information on migrating to Version 6, see *DB2 UDB for OS/390 Version 6 Release Planning Guide.*

© Copyright IBM Corp. 1998, 1999

Chapter 6. Prerequisites of Version 6 of DB2 for OS/390

This chapter identifies prerequisites and optional programs of DB2 for OS/390 Version 6 and for the features that are delivered with it. See “Prerequisites of features of DB2 for OS/390 Version 6” on page 63 for information about the requirements for the features of DB2 for OS/390 Version 6. This chapter identifies the minimum level of hardware or software that is required; unless otherwise noted, subsequent versions and releases of these products are acceptable.

DB2 for OS/390 Version 6 prerequisites

DB2 for OS/390 Version 6 has requirements for processors, other programs, and virtual storage.

Hardware requirements

DB2 for OS/390 Version 6 operates on any processor that is supported by OS/390 Version 1 Release 3 or subsequent releases. The processor must have enough real storage to satisfy the combined requirements of DB2, OS/390, the appropriate Data Facility Product, the appropriate access methods, telecommunications, batch, and other customer-required applications.

OS/390 runs on the hardware listed below:

- All models of the S/390® Parallel Enterprise Servers or S/390 Parallel Transaction Servers (IBM 9672)
- All models of the IBM ES/9000® Processor Unit 9021, the 9121, or the 9221
- PC Server S/390 or RS/6000® with S/390 Server-on-Board
- S/390 Multiprise® 2000

The configuration must include sufficient I/O devices to support the requirements for system output, system residence, and system data sets. Sufficient direct access storage (DASD) must be available to satisfy the user’s information storage requirements and can consist of any direct-access facility that is supported by the system configuration and the programming system.

In addition to listing auxiliary storage and data communications devices, this section identifies function-dependent hardware requirements and virtual storage requirements.

Auxiliary storage

DB2 is independent of both DASD and tape device type. You can use any magnetic, optical, or tape device that is supported by the data facilities component of DFSMS/MVS® for the DB2 data sets. See Table 2 for a list of device types that are supported for DB2 data sets.

Table 2. Auxiliary storage

Data set type	Device type
Active recovery log data sets	DASD
Archive recovery log data sets	DASD, tape
Image copy data sets	DASD, tape
Bootstrap data set	DASD
User data sets	DASD

© Copyright IBM Corp. 1998, 1999

Table 2. Auxiliary storage (continued)

Data set type	Device type
DB2 catalog data sets	DASD
Work data sets (for utilities)	DASD, tape

If these data sets are on DASD that is shared with other OS/390 systems, you should use global resource serialization to prevent concurrent access by more than one OS/390 system.

The minimum DASD space requirement, based on installing DB2 using the panel default values, is approximately 600 MB. Users need additional DASD space for their data.

If you use dual logging and tape for the log archiving device, you need at least two tape drives.

Data communication devices

Control DB2 operations from:

- The system console
- Authorized IMS/ESA Transaction Manager terminals
- Authorized CICS terminals
- TSO terminals (by authorized users)

For information about the data communication devices that are supported by IMS/ESA Transaction Manager, CICS, and TSO, see the documentation for these products.

Function-dependent hardware requirements

DB2 has the following function-dependent requirements.

COPY and RECOVER INDEX: Use of the COPY utility to take DFSMS™ concurrent copies of indexes requires a 9390, a 3990 Model 6, or a 3990 Model 3 with the Extended Platform installed or equivalent function.

Use of the RECOVER utility to restore DFSMS concurrent copies of indexes also requires a 9390, a 3990 Model 6, or a 3990 Model 3 3390 Model 3 or 3990 Model 6 controller at the with the Extended Platform installed or equivalent function.

DB2 table spaces that are larger than 1 terabyte: DB2 table spaces that are larger than 1 terabyte require one of the following products:

- 3990 Model 3 or Model 6 controllers with extended support
- 9340 DASD array

Program requirements and optional programs

This section lists licensed programs that are required in the DB2 for OS/390 environment. You can use subsequent versions or releases of these programs, unless the description for a given program states otherwise. Check the RETAIN® Preventative Service Planning (PSP) Facility for the most current information about APARs you must install to run DB2 for OS/390 and its optional features.

This section also identifies the requirements associated with specific DB2 capabilities, as well as optional programs that you can use with DB2 for OS/390 Version 6.

Operating system and support programs

For an OS/390 environment, DB2 requires the function that is provided by the following licensed programs or their equivalents; subsequent versions or releases of any product are acceptable.

- OS/390 Version 1 Release 3 Base Services (5645-001)
- OS/390 Version 1 Release 3 Application Enablement optional feature for DFSORT™

Function-dependent program requirements

DB2 for OS/390 has the following function-specific program requirements. For specific software requirements for the required products, refer to the respective product announcements.

User-defined functions: User-defined functions require that you include Language Environment mapping macros and macros that generate a Language Environment-conforming prolog and epilog. Language Environment is part of the OS/390 Version 1 Release 3 Application Enablement base feature.

See *OS/390 Language Environment for OS/390 & VM Programming Guide* for details. See "Programming languages" on page 62 for specific levels of Assembler, COBOL, C, C++, or PL/I.

SQL procedures: SQL procedures have the following requirements:

- Execution of SQL procedures requires Language Environment.

  Language Environment is part of the OS/390 Version 1 Release 3 Application Enablement base feature.
- Preparation of SQL procedures for execution using JCL requires the following compiler language levels, or higher:
  - IBM C/C++ for MVS/ESA Version 3 Release 2 (5655-121)
  - IBM SAA AD/Cycle C/370™ Version 1 Release 2 (5688-194)
- Preparation of SQL procedures for execution using the SQL procedure processor requires the following products:
  - The following compiler language levels, or higher:
    - IBM C/C++ for MVS/ESA Version 3 Release 2 (5655-121)
    - IBM SAA AD/Cycle C/370™ Version 1 Release 2 (5688-194)
  - IBM DATABASE 2 Universal Database Server for OS/390 REXX Language Support (5645-DB2)

Predictive governing: To take full advantage of DB2 for OS/390 Version 6 predictive governing capabilities, DB2 Connect workstation clients require DB2 Connect Version 5.2 (with appropriate service applied).

Language Environment built-in functions: Use of any of the following built-in functions requires OS/390 Version 2 Release 4 Application Enablement base element with APARs.

- Language Environment built-in functions, including the following advanced math functions:
  - DEGREES and RADIANS
  - RAND
  - EXP and POWER
  - LOG10, LOG, and LN
  - Trigonometry functions (ACOS, ASIN, ATAN, ATANH, ATAN2, COS, COSH, SIN, SINH, TAN, TANH)
- DB2 built-in functions UPPER, LOWER, and TRANSLATE with locale

Language Environment character conversion: Use of Language Environment character conversions requires OS/390 Version 2 Release 9.

More than 10 000 open data sets: Use of more than 10 000 open data sets requires OS/390 Version 2 Release 6.

DB2 table spaces that are larger than 1 terabyte: DB2 table spaces that are larger than 1 terabyte require one of the following products:

- VSAM Extended Addressability Linear Data Sets in OS/390 Version 2 Release 7 (5647-A01)
- DFSMS Version 1 Release 5 (5695-DF1)

Group buffer pool duplexing: Group buffer pool duplexing has the following requirements:

- A minimum coupling facility architectural level of CFLEVEL=5.
- OS/390 Version 2 Release 6, or OS/390 Release 3, 4, or 5 with APAR; OS/390 Version 2 Release 6 and above do not require the APAR.

Open Database Connectivity: Open Database Connectivity (ODBC) functions have the following requirements:

- Execution of ODBC component in the application address space requires OS/390 Version 1 Release 3 Application Enablement optional feature for C/C++.
- Customer applications are supported in the following compiler language levels, or higher:
  - IBM C/C++ for MVS/ESA Version 3 Release 2 (5655-121)
  - IBM SAA AD/Cycle C/370™ Version 1 Release 2 (5688-194)

Java Database Connectivity: Java Database Connectivity (JDBC) requires Java for OS/390 (5655-A46).

Java stored procedures: Java stored procedures require Enterprise ToolKit for OS/390, which is part of VisualAge for Java, Enterprise Edition for OS/390 (5655-JAV).

DB2 Extenders: Use of the DB2 Extenders requires OS/390 Version 2 Release 4 (5647-801).

To write extender applications, C or C++ for OS/390 is required. Use of DB2 Extenders does not require C or C++.

In addition, use of the DB2 Text Extender requires the IBM Text Search Engine Version 2 Release 1 (FMID HIMN210) or higher. If you are running OS/390 Version 2 (Releases 4 through 7), download and install the IBM Text Search Engine in SMP/E format from the following Web site:

http://www.ibm.com/software/iminer/fortext/

OS/390 Release 8 includes the IBM Text Search Engine as a base element of that operating system.

Sysplex workload balancing: For Sysplex workload balancing, workstation clients require one of the following products:

- DB2 Connect Version 6 (strongly recommended)
- DB2 Connect Version 5 Release 2

IEEE floating point: For DB2 for OS/390 to accept IEEE floating point (also called binary floating point) values in SQL statements or LOAD utility jobs, OS/390 Version 2 Release 6 is required.

31-digit decimal support in COBOL: DB2 COBOL application program use of 31-digit decimal variables requires IBM COBOL for OS/390 and VM Version 2 (5648-A25).

Accessing IMS databases through DB2 stored procedures: The use of DB2 stored procedures to access IMS databases by using the ODBA application programming interface requires IMS/ESA Version 6 (5655-158).

Optional programs

You can use the following optional licensed programs with DB2 for OS/390 Version 6. Unless otherwise specified, the release shown for a product and any subsequent release are acceptable. In some cases, earlier versions or releases of IBM licensed programs may also work with DB2 for OS/390, but IBM may not have tested them at the time this document was published. If you have questions, please check with your IBM representative.

Connectivity: In addition to any DRDA-compliant database management systems, DB2 for OS/390 Version 6 supports the following IBM relational database products:

- IBM DB2 Universal Database Extended Enterprise Edition Version 5 with the DB2 Connect component installed
- IBM DB2 Universal Database Enterprise Edition Version 5 with the DB2 Connect component installed
- IBM DB2 Universal Database for AS/400 Version 4 Release 2
- IBM Operating System/400® (OS/400®) Version 4 Release 1 with DB2 for AS/400 (5769-SS1)
- IBM DB2 Server for VM and VSE Version 5 (5648-158)
- DB2 DataJoiner Version 2 Release 1.1 (5231-200)

Net.Data for OS/390, a feature of DB2 for OS/390 Version 6, provides connectivity to DB2 from the Web.

Capacity planning: DB2 Estimator, a feature of DB2 for OS/390 Version 6, works with DB2 data to estimate application feasibility, to model application cost and performance, and to estimate required CPU and I/O capacity.

Transaction management: The following transaction management products work with DB2:

- Information Management System (IMS)
  - Information Management System/ESA (IMS/ESA) Version 6 (5655-158)
  - Information Management System/ESA (IMS/ESA) Version 5 (5695-176)
- Customer Information Control System (CICS)
  - CICS Transaction Server for OS/390 Release 1 (5655-147)
  - CICS/ESA Version 4 (5655-018)

Query support: The following query products work with DB2:

- Query Management Facility (QMF), a feature of DB2 for OS/390 Version 6
- QMF for Windows, a feature of DB2 for OS/390 Version 6
- QMF High Performance Option, a feature of DB2 for OS/390 Version 6
- The DB2 Extenders capability of DB2 for OS/390 Version 6

Application development: The following application development programs work with DB2:

- C/C++ Productivity Tools for OS/390 (5655-B85)
- IBM VisualAge® Generator Version 3 (see software announcement 297-395)
- IBM VisualAge for Java, Enterprise Edition for OS/390, Version 2 (5655-JAV)
- IBM VisualAge COBOL Version 2 (see software announcement 297-371)
- IBM VisualAge PL/I Version 2 (see software announcement 297-372)
- Application System (AS) Version 4 Release 2 (5648-092)

Operational support: The following programs provide operational support for DB2:

- IBM DATABASE 2 Performance Monitor (DB2 PM), an optional feature of DB2 for OS/390 Version 6
- OS/390 Version 1 Release 3 System Services optional feature for DFSMS features
- OS/390 Version 1 Release 3 Security Server optional feature for RACF
- NetView® Version 3 for MVS/ESA™ (5655-007), or NetView Version 2 Release 4 (5685-111)
- Tivoli™ Performance Reporter for OS/390 Version 1 Release 4 (569510100) or Performance Reporter for MVS Version 1 Release 2 (5695-101)
- Library Readers included on the CD-ROMs for BookManager® books

Replication support: The following programs provide replication support for DB2:

- DB2 DataPropagator, an optional feature of DB2 for OS/390 Version 6; or one or both of the following products:
  - DataPropagator Relational Apply for MVS Version 5 Release 1 (5655-A22)
  - DataPropagator Relational Capture for MVS Version 5 Release 1 (5655-A23)
- DataPropagator NonRelational MVS/ESA Version 2 (5696-705)
- IBM DataRefresher Version 1 (5696-703)

Database administration and systems management support: The following programs support database systems management for DB2:

- DB2 Automated Utility Generator (DB2AUG) Version 1 Release 2 (5695-077)
- DB2 Administration Tool, a feature of DB2 for OS/390 Version 6
- DB2 Buffer Pool Tool, a feature of DB2 for OS/390 Version 6
- DB2 Control Center, an element of the DB2 Management Tools Package
- DB2 Installer, an element of the DB2 Management Tools Package
- Visual Explain, an element of the DB2 Management Tools Package
- DB2 Estimator, an element of the DB2 Management Tools Package

The DB2 Management Tools Package is a feature of DB2 for OS/390 Version 6.

Programming languages: You can use the following programming languages (in addition to High-Level Assembler, which is part of OS/390) to develop application programs for DB2 for OS/390 Version 6:

C	IBM AD/Cycle® C/370 Compiler Version 1 Release 2 (5688-216); IBM C/370 Library Version 2 Release 2 (5688-188)
C++	OS/390 Version 1 Release 3 Application Enablement optional feature for C/C++
COBOL	IBM COBOL for MVS and VM Version 1 Release 2 (5688-197), or

	IBM COBOL for OS/390 and VM Version 2 Release 1 (5648-A25), or VS COBOL II Compiler and Library Version 1 Release 4 (5668-958)
Fortran	VS Fortran Compiler, Library, and Interactive Debugger Version 2 Release 6 (5668-806)
Java	Java for OS/390 (5655-A46)
PL/I	IBM PL/I for MVS and VM Version 1 Release 1 (5688-235), or OS PL/I Compiler, Library, and Interactive Test Facility Version 2 Release 3 (5668-909)
REXX	IBM TSO Extensions for MVS REXX, which is part of OS/390
Smalltalk	IBM VisualAge Smalltalk Version 4 Release 5 (5655-B14 or 5802-AAR)

Virtual storage requirements

The amount of space needed for the common service area (CSA) below the 16-MB line is less than 40 KB for each DB2 subsystem and 24 KB for each IRLM when APAR PQ12390 and prior service is applied. High concurrent activity, parallelism, or high contention can require more CSA.

Most of the DB2 common data resides in the extended common service area (ECSA). Most modules, control blocks, and buffers reside in the extended private area. A DB2 subsystem with 200 concurrent users and 2000 open data sets should need less than 2 MB of virtual storage below the 16-MB line.

Prerequisites of features of DB2 for OS/390 Version 6

DB2 for OS/390 Version 6 includes many features, some of which have requirements of their own, above and beyond what DB2 for OS/390 Version 6 requires. This section identifies the requirements for using these features with DB2 for OS/390 Version 6, but it does not repeat those DB2 for OS/390 requirements that apply to the features. Also, some of these features can be used with prior releases of DB2; these requirements are not included in this section, but rather in the detailed installation information for the specific feature.

You can use subsequent versions or releases of the products mentioned in this section, unless otherwise noted.

Recommendation: Before using these features, refer to the installation information for these features to ensure that you have all required and recommended products.

DB2 Installer requirements

DB2 Installer is an element of the DB2 Management Tools Package, which is a feature of DB2 for OS/390 Version 6. DB2 Installer has hardware and program requirements.

Hardware requirements

DB2 Installer requires:

- 20 MB disk memory on the target drive and 2 MB of disk space for each subsystem that is defined
- A monitor that is capable of displaying 1024-by-768 resolution

Program requirements

DB2 Installer can run in either of the following environments, each of which has its own requirements:

- Windows NT, which requires Microsoft™ Windows NT Version 4.0
- OS/2, which requires:
  - OS/2 Version 4
  - TCP/IP on OS/2, Version 3.0

Both environments require TCP/IP in any of the following circumstances:

- To run jobs from the workstation
- To use the copy-jobs-to-host function from the workstation
- To use all the functions of DB2 Installer

If you don't have TCP/IP, you can use DB2 Installer to customize your installation jobs, but you need to use a method outside of DB2 Installer to move jobs from the workstation to OS/390 for execution.

Visual Explain requirements

The Visual Explain element of the DB2 Management Tools Package has hardware and program requirements.

Hardware requirements

The Visual Explain element of the DB2 Management Tools Package requires:

- A workstation with either OS/2 Version 4 or Windows NT 4.0
- Approximately 12 MB of hard disk space
- A high-resolution monitor

Program requirements

One of the following products must be installed on the DB2 Visual Explain workstation:

- DB2 Connect Personal Edition Version 5.2
- DB2 Client Application Enablers™ (CAE) Version 5.2, connected to a server that is running DB2 Connect Enterprise Edition Version 5.2

In addition, DB2 Visual Explain requires one of the following communication protocols:

- TCP/IP
- SNA communications using a product such as Communication Server 5.0, SNA Server Version 4.0, or the integrated SNA support in DB2 Universal Database Personal Edition

Visual Explain includes a browser that lets users view current values of subsystem parameters. To use this browser, your DB2 subsystem must have:

- Stored procedures capability
- The DSNWZP stored procedure enabled

DB2 Estimator requirements

DB2 Estimator has hardware and program requirements.

Hardware requirements

DB2 Estimator, which is an element of the DB2 Management Tools Package, requires:

- A workstation with one of the operating systems listed under "Program requirements" on page 65.

- Approximately 12 MB of hard disk space

Program requirements

The DB2 Estimator operates in the following environments:

- Windows 3.1 and subsequent releases
- Windows 98
- Windows 95
- Windows NT 4

Net.Data requirements

Net.Data for OS/390 Version 2 Release 2 requires an HTTP Web server that is installed on the same server as Net.Data and DB2 for OS/390. The Web server can be another HTTP-compliant Web server or one of the following servers:

- IBM Internet Connection Secure Server for OS/390 Version 2 Release 2 (5697-B14)
- Domino Go Webserver Version 4 Release 6 Modification 1 for OS/390 (5697-C58)

In addition, Net.Data servlets have the following requirements:

- OS/390 Version 2 Release 5 (5647-A01)
- Domino Go Webserver Version 5 (5697-D43) with Servlet Express

Net.Data for OS/390 has no hardware requirements.

QMF requirements

QMF, which is a feature of DB2 for OS/390 Version 6, has hardware and program requirements.

Hardware requirements

The following QMF features have hardware dependencies:

- QMF for OS/390 requires a display station that is supported by GDDM®.
- QMF High Performance Option (HPO) requires a display station that is supported by ISPF.
- QMF for Windows requires a workstation that supports:
  - A Windows (16-bit or 32-bit) operating system
  - Network connectivity

Program requirements

The following QMF features have program dependencies.

- QMF for OS/390 requires Graphical Data Display Manager (GDDM) Version 2 Release 3 (5665-356). (This requirement is satisfied by OS/390 Version 2.)

  Use of QMF forms calculations requires a Windows 32-bit operating system and IBM Object REXX Interpreter Edition Version 1.0 (5639-B73).
- QMF for Windows requires:
  - A Windows (16-bit or 32-bit) operating system
  - Network communication software on each user machine, plus one or both of the following programs:
    - An SNA product that provides a CPI-C interface
    - A TCP/IP product that provides a WinSock Version 1.1 interface

  Use of QMF for Windows with English Wizard natural language query requires English Wizard Release 3.1, from Linguistic Technologies.

The QMF for Windows Administrator module requires a Windows 32-bit operating system.

DB2 Performance Monitor requirements

DB2 Performance Monitor (DB2 PM), which is a feature of DB2 for OS/390 Version 6, has hardware and program requirements.

Hardware requirements

DB2 Performance Monitor has the following dependencies.

- For the host-based Online Monitor, a display station that is supported by Interactive System Productivity Facility (ISPF)
- For the host-based graphics facility, an IBM color graphics display station, or equivalent, that is supported by Graphical Data Display Manager (GDDM)

The DB2 PM Workstation-Based Online Monitor has the following dependencies:

- A high-resolution monitor
- A workstation that supports OS/2 Version 3 or Windows NT Version 4.0
- Approximately 20 MB of hard disk space

Program requirements

For the host-based Online Monitor and host-based graphics facility, DB2 Performance Monitor has no functional dependencies if you are monitoring DB2 for OS/390 Version 6. The DB2 Performance Monitor feature supports an environment of multiple DB2 releases, namely instrumentation, catalog, and PLAN_TABLE data from:

- DB2 UDB for OS/390 Version 6 (5645-DB2)
- DB2 for OS/390 Version 5 (5655-DB2)
- DB2 for MVS/ESA Version 4 (5695-DB2)

Refer to the software requirements for the specific version of DB2.

Workstation-Based Analysis and Tuning

The Workstation-Based Analysis and Tuning feature of DB2 for OS/390 Version 6 has hardware and program requirements.

Hardware requirements

Workstation-Based Analysis and Tuning has the following hardware requirements:

- A high-resolution monitor
- A workstation that supports OS/2 Version 3 or Windows NT Version 4.0

Program requirements

Workstation-Based Analysis and Tuning requires one of the following operating systems:

- OS/2, which requires both:
  - OS/2 Version 3
  - Personal Communications AS/400 and 3270 V4.1 (5622-972)
- Windows NT Version 4.0. Windows NT uses TCP/IP to communicate with the programs that run on OS/390, which requires TCP/IP for MVS Version 3 Release 2 (5655-HAL).

Chapter 7. Notices

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

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

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

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

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

IBM Corporation
J74/G4
555 Bailey Avenue
P.O. Box 49023
San Jose, CA 95161-9023
U.S.A.

© Copyright IBM Corp. 1998, 1999

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

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

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

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 is for planning purposes only. The information herein is subject to change before the products described become available.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, or other countries, or both:

AD/Cycle
AIX
AS/400
BookManager
CICS
CICS/ESA
C/370
DATABASE 2
DataJoiner
DataPropagator
DB2
DB2 Client Application Enablers
DB2 Connect
DB2 Extenders
DB2 Universal Database
DFSMS
DFSMS/MVS
DFSORT
Distributed Relational Database Architecture
DRDA
ES/9000
GDDM
IBM
IMS
IMS/ESA
Language Environment
Multiprise
MVS
MVS/ESA
Net.Data
Operating System/400
OS/2
OS/390
OS/400
Parallel Sysplex
PR/SM
QMF
RACF
RETAIN
RMF
RS/6000
S/390
System/390
VisualAge
VisualGen
Visual Warehouse
WebSphere

Lotus, Lotus Go Webserver, Lotus Notes, NotesPump, and Domino are trademarks of Lotus Development Corporation in the United States, or other countries, or both.

Tivoli and NetView are trademarks of Tivoli Systems Inc. in the United States, or other countries, or both.

The following terms are trademarks of other companies as follows:

- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.
- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.
- UNIX is a registered trademark in the United States and/or other countries licensed exclusively through X/Open Company Limited.

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

Glossary

The following terms and abbreviations are defined as they are used in the DB2 library. If you do not find the term you are looking for, refer to the index or to *IBM Dictionary of Computing*.

A

API. Application programming interface.

auxiliary table. A table that stores columns outside the table in which they are defined. Contrast with *base table*.

B

base table. (1) A table that is created by the SQL CREATE TABLE statement and that holds persistent data. Contrast with *result table* and *temporary table*.

(2) A table containing a LOB column definition. The actual LOB column data is not stored with the base table. The base table contains a row identifier for each row and an indicator column for each of its LOB columns. Contrast with *auxiliary table*.

binary large object (BLOB). A sequence of bytes, where the size of the value ranges from 0 bytes to 2 GB - 1. Such a string does not have an associated CCSID.

BLOB. Binary large object.

buffer pool. Main storage that is reserved to satisfy the buffering requirements for one or more table spaces or indexes.

built-in function. A function that DB2 supplies. Contrast with *user-defined function*.

C

call level interface (CLI). A callable application programming interface (API) for database access, which is an alternative to using embedded SQL. In contrast to embedded SQL, DB2 ODBC (which is based on the CLI architecture) does not require the user to precompile or bind applications, but instead provides a standard set of functions to process SQL statements and related services at run time.

character large object (CLOB). A sequence of bytes representing single-byte characters or a mixture of single- and double-byte characters where the size of the value can be up to 2 GB - 1. In general, character large object values are used whenever a character string might exceed the limits of the VARCHAR type.

CLOB. Character large object.

coupling facility. A special PR/SM™ LPAR logical partition that runs the coupling facility control program and provides high-speed caching, list processing, and locking functions in a Sysplex.

D

data sharing. The ability of two or more DB2 subsystems to directly access and change a single set of data.

data type. An attribute of columns, literals, host variables, special registers, and the results of functions and expressions.

DBCLOB. Double-byte character large object.

DBCS. Double-byte character set.

DB2 catalog. Tables that are maintained by DB2 and that contain descriptions of DB2 objects, such as tables, views, and indexes.

distinct type. A user-defined data type that is internally represented as an existing type (its source type), but is considered to be a separate and incompatible type for semantic purposes.

Distributed Relational Database Architecture (DRDA). A connection protocol for distributed relational database processing that is used by IBM's relational database products. DRDA includes protocols for communication between an application and a remote relational database management system, and for communication between relational database management systems.

double-byte character large object (DBCLOB). A sequence of bytes representing double-byte characters where the size of the values can be up to 2 GB. In general, double-byte character large object values are used whenever a double-byte character string might exceed the limits of the VARGRAPHIC type.

G

GB. Gigabyte (1 073 741 824 bytes).

GBP. Group buffer pool.

group buffer pool (GBP). A coupling facility cache structure that is used by a data sharing group to cache data and to ensure that the data is consistent for all members.

K

KB. Kilobyte (1024 bytes).

© Copyright IBM Corp. 1998, 1999

L

large object (LOB). A sequence of bytes representing bit data, single-byte characters, double-byte characters, or a mixture of single- and double-byte characters. A LOB can be up to 2 GB - 1 byte in length. See also *BLOB*, *CLOB*, and *DBCLOB*.

LOB. Large object.

LOB locator. A mechanism that allows an application program to manipulate a large object value in the database system. A LOB locator is a fullword integer value that represents a single LOB value. An application program retrieves a LOB locator into a host variable and can then apply SQL operations to the associated LOB value using the locator.

M

MB. Megabyte (1 048 576 bytes).

O

ODBC. Open Database Connectivity.

Open Database Connectivity (ODBC). A Microsoft™ database application programming interface (API) for C that allows access to database management systems by using callable SQL. ODBC does not require the use of an SQL preprocessor. In addition, ODBC provides an architecture that lets users add modules called *database drivers*, which link the application to their choice of database management systems at run time. This means that applications no longer need to be directly linked to the modules of all the database management systems that are supported.

S

SQLJ. Structured Query Language (SQL) that is embedded in the Java programming language.

stored procedure. A user-written application program, that can be invoked through the use of the SQL CALL statement.

strong typing. A process that guarantees that only user-defined functions and operations that are defined on a distinct type can be applied to that type. For example, you cannot directly compare two currency types, such as Canadian dollars and US dollars. But you can provide a user-defined function to convert one currency to the other and then do the comparison.

T

temporary table. A table that holds temporary data; for example, temporary tables are useful for holding or sorting intermediate results from queries that contain a large number of rows. The two kinds of temporary table, which are created by different SQL statements, are the created temporary table and the declared temporary table. Contrast with *result table*. See also *created temporary table* and *declared temporary table*.

three-part name. The full name of a table, view, or alias. It consists of a location name, authorization ID, and an object name, separated by a period.

trigger. A set of SQL statements that are stored in a DB2 database and executed when a certain event occurs in a DB2 table.

type 2 indexes. Indexes that are created on a release of DB2 after Version 5 or that are specified as type 2 indexes in Version 4 or Version 5.

U

user-defined data type (UDT). See *distinct type*.

user-defined function (UDF). A function that is defined to DB2 using the CREATE FUNCTION statement and that can be referenced thereafter in SQL statements. A user-defined function can be either an *external function* or a *sourced function*. Contrast with *built-in function*.

Program Number: 5645-DB2

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

GC26-9017-03