

DB2 Universal Database for OS/390



# ODBC Guide and Reference

*Version 6*

**Note!**

Before using this information and the product it supports, be sure to read the general information under Appendix F, "Notices" on page 489.

**Second Edition, Softcopy Only (April 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 and includes the changes indicated in the printed version by vertical bars. Additional changes made to this softcopy version of the manual since the hardcopy manual was published are indicated by the hash (#) symbol in the left-hand margin. Editorial changes that have no technical significance are not noted.

© **Copyright International Business Machines Corporation 1997, 1999. All rights reserved.**

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

---

# Contents

	<b>Chapter 1. Introduction to this book and the DB2 library</b>	1
	Who should use this book	1
	How this book is organized	1
#	Product terminology and citations	2
	How to use the DB2 library	2
	How to obtain DB2 information	4
	Summary of changes to DB2 UDB for OS/390 Version 6	6
	Summary of changes to this book	10
	<b>Chapter 2. Introduction to DB2 ODBC</b>	13
	DB2 ODBC background information	13
	Differences between DB2 ODBC and embedded SQL	15
	Advantages of using DB2 ODBC	18
	Deciding which interface to use	18
	Other information sources	19
	<b>Chapter 3. Writing a DB2 ODBC application</b>	21
	Initialization and termination	22
	Transaction processing	27
	Diagnostics	34
	Data types and data conversion	36
	Working with string arguments	43
	Querying environment and data source information	45
	<b>Chapter 4. Configuring DB2 ODBC and running sample applications</b>	49
	Installing DB2 ODBC	49
	DB2 ODBC runtime environment	49
	Setting up DB2 ODBC runtime environment	52
	Preparing and executing a DB2 ODBC application	56
	DB2 ODBC initialization file	60
#	DB2 ODBC migration considerations	72
	<b>Chapter 5. Functions</b>	73
	Function summary	74
	SQLAllocConnect - Allocate connection handle	79
	SQLAllocEnv - Allocate environment handle	83
	SQLAllocStmt - Allocate a statement handle	85
	SQLBindCol - Bind a column to an application variable	87
	SQLBindParameter - Binds a parameter marker to a buffer or LOB locator	93
	SQLCancel - Cancel statement	104
	SQLColAttributes - Get column attributes	106
	SQLColumnPrivileges - Get privileges associated with the columns of a table	112
	SQLColumns - Get column information for a table	117
	SQLConnect - Connect to a data source	123
	SQLDataSources - Get list of data sources	129
	SQLDescribeCol - Describe column attributes	132
	SQLDescribeParam - Describe parameter marker	138
	SQLDisconnect - Disconnect from a data source	140
	SQLDriverConnect - (Expanded) connect to a data source	142
	SQLError - Retrieve error information	148

SQLExecDirect - Execute a statement directly . . . . .	154
SQLExecute - Execute a statement . . . . .	159
SQLExtendedFetch - Extended fetch (fetch array of rows) . . . . .	162
SQLFetch - Fetch next row . . . . .	169
SQLForeignKeys - Get the list of foreign key columns . . . . .	175
SQLFreeConnect - Free connection handle . . . . .	184
SQLFreeEnv - Free environment handle . . . . .	186
SQLFreeStmt - Free (or reset) a statement handle . . . . .	188
SQLGetConnectOption - Returns current setting of a connect option . . . . .	191
SQLGetCursorName - Get cursor name . . . . .	193
SQLGetData - Get data from a column . . . . .	199
SQLGetEnvAttr - Returns current setting of an environment attribute . . . . .	212
SQLGetFunctions - Get functions . . . . .	214
SQLGetInfo - Get general information . . . . .	219
SQLGetLength - Retrieve length of a string value . . . . .	235
SQLGetPosition - Return starting position of string . . . . .	238
SQLGetSQLCA - Get SQLCA data structure . . . . .	242
SQLGetStmtOption - Returns current setting of a statement option . . . . .	249
SQLGetSubString - Retrieve portion of a string value . . . . .	251
SQLGetTypeInfo - Get data type information . . . . .	254
SQLMoreResults - Determine if there are more result sets . . . . .	262
SQLNativeSql - Get native SQL text . . . . .	266
SQLNumParams - Get number of parameters in a SQL statement . . . . .	269
SQLNumResultCols - Get number of result columns . . . . .	271
SQLParamData - Get next parameter for which a data value is needed . . . . .	273
SQLParamOptions - Specify an input array for a parameter . . . . .	275
SQLPrepare - Prepare a statement . . . . .	277
SQLPrimaryKeys - Get primary key columns of a table . . . . .	285
SQLProcedureColumns - Get input/output parameter information for a procedure . . . . .	290
SQLProcedures - Get list of procedure names . . . . .	300
SQLPutData - Passing data value for a parameter . . . . .	304
SQLRowCount - Get row count . . . . .	307
SQLSetColAttributes - Set column attributes . . . . .	309
SQLSetConnection - Set connection handle . . . . .	314
SQLSetConnectOption - Set connection option . . . . .	316
SQLSetCursorName - Set cursor name . . . . .	322
SQLSetEnvAttr - Set environment attribute . . . . .	325
SQLSetParam - Binds a parameter marker to a buffer . . . . .	328
SQLSetStmtOption - Set statement option . . . . .	334
SQLSpecialColumns - Get special (row identifier) columns . . . . .	339
SQLStatistics - Get index and statistics information for a base table . . . . .	344
SQLTablePrivileges - Get privileges associated with a table . . . . .	349
SQLTables - Get table information . . . . .	353
SQLTransact - Transaction management . . . . .	357
<b>Chapter 6. Using advanced features . . . . .</b>	<b>361</b>
Environment, connection, and statement options . . . . .	361
Distributed unit of work (coordinated distributed transactions) . . . . .	363
Querying system catalog information . . . . .	368
Sending/retrieving long data in pieces . . . . .	373
Using arrays to input parameter values . . . . .	375
Retrieving a result set into an array . . . . .	379
Using large objects . . . . .	383

Using distinct types . . . . .	386
Using stored procedures . . . . .	389
Writing multithreaded applications . . . . .	394
Mixing embedded SQL and DB2 ODBC . . . . .	401
Using vendor escape clauses . . . . .	404
Programming hints and tips . . . . .	407
<b>Chapter 7. Problem diagnosis . . . . .</b>	<b>413</b>
Tracing . . . . .	413
Debugging . . . . .	423
<b>Appendix A. DB2 ODBC and ODBC . . . . .</b>	<b>425</b>
DB2 ODBC and ODBC drivers . . . . .	425
ODBC APIs and data types . . . . .	426
Isolation levels . . . . .	428
<b>Appendix B. Extended scalar functions . . . . .</b>	<b>429</b>
String functions . . . . .	430
Date and time functions . . . . .	431
System functions . . . . .	431
<b>Appendix C. SQLSTATE cross reference . . . . .</b>	<b>433</b>
<b>Appendix D. Data conversion . . . . .</b>	<b>443</b>
Data type attributes . . . . .	444
Converting data from SQL to C data types . . . . .	449
Converting data from C to SQL data types . . . . .	459
<b>Appendix E. Example code . . . . .</b>	<b>469</b>
<b>Appendix F. Notices . . . . .</b>	<b>489</b>
Programming interface information . . . . .	490
Trademarks . . . . .	491
<b>Glossary . . . . .</b>	<b>493</b>
<b>Bibliography . . . . .</b>	<b>499</b>
<b>Index . . . . .</b>	<b>505</b>



---

# Chapter 1. Introduction to this book and the DB2 library

This book provides the information necessary to write applications using DB2 ODBC to access IBM® DATABASE 2™ servers as well as any database that supports DRDA® level 1 or DRDA level 2 protocols. This book should also be used as a supplement when writing portable ODBC applications that can be executed in a native DB2 for OS/390® environment using DB2 ODBC.

---

## Who should use this book

This book is intended for:

- DB2 application programmers with a knowledge of SQL and the C programming language.
- ODBC application programmers with a knowledge of SQL and the C programming language.

---

## How this book is organized

This book is divided into the following chapters:

- “Chapter 1. Introduction to this book and the DB2 library,” identifies the purpose, the audience, and the use of this book.
- “Chapter 2. Introduction to DB2 ODBC” on page 13, introduces DB2 ODBC and discusses the background of the interface and its relation to embedded SQL and Microsoft ODBC.
- “Chapter 3. Writing a DB2 ODBC application” on page 21, provides an overview of a typical DB2 ODBC application. This chapter discusses the basic tasks or steps within a simple DB2 ODBC application. General concepts are introduced as well as the basic functions and the interaction between them.
- “Chapter 4. Configuring DB2 ODBC and running sample applications” on page 49, contains information for setting up the necessary environment to compile and run DB2 ODBC applications. Sample applications are provided in order to verify your environment.
- “Chapter 5. Functions” on page 73, is a reference for the functions that make up DB2 ODBC.
- “Chapter 6. Using advanced features” on page 361, provides an overview of more advanced tasks and the functions used to perform them.
- “Chapter 7. Problem diagnosis” on page 413, explains how to work with traces and debug applications.

The appendixes contain the following information:

- Appendix A, “DB2 ODBC and ODBC” on page 425, discusses the differences between ODBC and DB2 ODBC.
- Appendix B, “Extended scalar functions” on page 429, describes the scalar functions that can be accessed as DB2 functions, or using ODBC vendor escape clauses.

- Appendix C, “SQLSTATE cross reference” on page 433, contains an SQLSTATE table that lists the functions that may generate each SQLSTATE. (Each function description in “Chapter 5. Functions” on page 73 lists the possible SQLSTATEs for each function.)
- Appendix D, “Data conversion” on page 443, contains information about SQL and C data types, and conversion between them.
- Appendix E, “Example code” on page 469, provides an extensive stored procedure example.

---

## # Product terminology and citations

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

References in this book to "DB2 UDB" relate to the DB2 Universal Database™ product that is available on the AIX®, OS/2®, and Windows NT™ operating systems. When this book refers to books about the DB2 UDB product, the citation includes the complete title and order number.

The following terms are used as indicated:

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

**C and C language**

Represent the C programming language.

**CICS®**

Represents CICS/ESA® and CICS Transaction Server for OS/390 Release 1.

**IMS™** Represents IMS/ESA®.

**MVS** Represents the MVS element of OS/390.

---

## How to use the DB2 library

Titles of books in the library begin with DB2 Universal Database for OS/390 Version 6. However, references from one book in the library to another are shortened and do not include the product name, version, and release. Instead, they point directly to the section that holds the information. For a complete list of books in the library, and the sections in each book, see the bibliography at the back of this book.

Throughout the library, the DB2 for OS/390 licensed program and a particular DB2 for MVS/ESA subsystem are each referred to as "DB2." In each case, the context makes the meaning clear.

The most rewarding task associated with a database management system is asking questions of it and getting answers, the task called *end use*. Other tasks are also necessary—defining the parameters of the system, putting the data in place, and so on. The tasks associated with DB2 are grouped into the following major categories



(but supplemental information relating to all of the below tasks for new releases of DB2 can be found in *DB2 Release Guide*):

**Installation:** If you are involved with DB2 only to install the system, *DB2 Installation Guide* might be all you need.

If you will be using data sharing then you also need *DB2 Data Sharing: Planning and Administration*, which describes installation considerations for data sharing.

**End use:** End users issue SQL statements to retrieve data. They can also insert, update, or delete data, with SQL statements. They might need an introduction to SQL, detailed instructions for using SPUFI, and an alphabetized reference to the types of SQL statements. This information is found in *DB2 Application Programming and SQL Guide* and *DB2 SQL Reference*.

End users can also issue SQL statements through the Query Management Facility (QMF™) or some other program, and the library for that program might provide all the instruction or reference material they need. For a list of the titles in the QMF library, see the bibliography at the end of this book.

**Application Programming:** Some users access DB2 without knowing it, using programs that contain SQL statements. DB2 application programmers write those programs. Because they write SQL statements, they need *DB2 Application Programming and SQL Guide*, *DB2 SQL Reference*, and *DB2 ODBC Guide and Reference* just as end users do.

Application programmers also need instructions on many other topics:

- How to transfer data between DB2 and a host program—written in COBOL, C, or FORTRAN, for example
- How to prepare to compile a program that embeds SQL statements
- How to process data from two systems simultaneously, say DB2 and IMS or DB2 and CICS®
- How to write distributed applications across platforms
- How to write applications that use DB2 ODBC to access DB2 servers
- How to write applications that use Open Database Connectivity (ODBC) to access DB2 servers
- How to write applications in the Java™ programming language to access DB2 servers

The material needed for writing a host program containing SQL is in *DB2 Application Programming and SQL Guide* and in *DB2 Application Programming Guide and Reference for Java™*. The material needed for writing applications that use DB2 ODBC or ODBC to access DB2 servers is in *DB2 ODBC Guide and Reference*. For handling errors, see *DB2 Messages and Codes*.

Information about writing applications across platforms can be found in *Distributed Relational Database Architecture™: Application Programming Guide*.

**System and Database Administration:** *Administration* covers almost everything else. *DB2 Administration Guide* divides those tasks among the following sections:

- Section 2 (Volume 1) of *DB2 Administration Guide* discusses the decisions that must be made when designing a database and tells how to bring the design into being by creating DB2 objects, loading data, and adjusting to changes.
- Section 3 (Volume 1) of *DB2 Administration Guide* describes ways of controlling access to the DB2 system and to data within DB2, to audit aspects of DB2 usage, and to answer other security and auditing concerns.
- Section 4 (Volume 1) of *DB2 Administration Guide* describes the steps in normal day-to-day operation and discusses the steps one should take to prepare for recovery in the event of some failure.
- Section 5 (Volume 2) of *DB2 Administration Guide* explains how to monitor the performance of the DB2 system and its parts. It also lists things that can be done to make some parts run faster.

In addition, the appendixes in *DB2 Administration Guide* contain valuable information on DB2 sample tables, National Language Support (NLS), writing exit routines, interpreting DB2 trace output, and character conversion for distributed data.

If you are involved with DB2 only to design the database, or plan operational procedures, you need *DB2 Administration Guide*. If you also want to carry out your own plans by creating DB2 objects, granting privileges, running utility jobs, and so on, then you also need:

- *DB2 SQL Reference*, which describes the SQL statements you use to create, alter, and drop objects and grant and revoke privileges
- *DB2 Utility Guide and Reference*, which explains how to run utilities
- *DB2 Command Reference*, which explains how to run commands

If you will be using data sharing, then you need *DB2 Data Sharing: Planning and Administration*, which describes how to plan for and implement data sharing.

Additional information about system and database administration can be found in *DB2 Messages and Codes*, which lists messages and codes issued by DB2, with explanations and suggested responses.

**Diagnosis:** Diagnosticians detect and describe errors in the DB2 program. They might also recommend or apply a remedy. The documentation for this task is in *DB2 Diagnosis Guide and Reference* and *DB2 Messages and Codes*.

---

## How to obtain DB2 information

### DB2 on the Web

Stay current with the latest information about DB2. View the DB2 home page on the World Wide Web. News items keep you informed about the latest enhancements to the product. Product announcements, press releases, fact sheets, and technical articles help you plan your database management strategy.

You can view and search DB2 publications on the Web, or you can download and print many of the most current DB2 books. Follow links to other Web sites with more information about DB2 family and OS/390 solutions. Access DB2 on the Web at the following address:

## DB2 publications

The DB2 publications for DB2 Universal Database Server for OS/390 are available in both hardcopy and softcopy format.

### BookManager® format

Using online books on CD-ROM, you can read, search across books, print portions of the text, and make notes in these BookManager books. With the appropriate BookManager READ product or IBM Library Readers, you can view these books in the OS/390, VM, OS/2, DOS, AIX, and Windows™ environments. You can also view many of the DB2 BookManager books on the Web.

### PDF format

Many of the DB2 books are available in Portable Document Format (PDF) for viewing or printing from CD-ROM or the Web. Download the PDF books to your intranet for distribution throughout your enterprise.

### CD-ROMs

Books for Version 6 of DB2 Universal Database Server for OS/390 are available on CD-ROMs:

- *DB2 UDB for OS/390 Version 6 Licensed Online Book*, LK3T-3519, containing *DB2 UDB for OS/390 Version 6 Diagnosis Guide and Reference* in BookManager format, for ordering with the product.
- *DB2 UDB Server for OS/390 Version 6 Online and PDF Library*, SK3T-3518, a collection of books for the DB2 server in BookManager and PDF formats.

Periodically, the books will be refreshed on subsequent editions of these CD-ROMs.

The books for Version 6 of DB2 UDB Server for OS/390 are also available on the following collection kits that contain online books for many IBM products:

- *Online Library Omnibus Edition OS/390 Collection*, SK2T-6700, in English
- *IBM Online Library MVS Collection Kit*, SK88-8002, in Japanese, for viewing on DOS and Windows operating systems.

## DB2 education

IBM Education and Training offers a wide variety of classroom courses to help you quickly and efficiently gain DB2 expertise. Classes are scheduled in cities all over the world. You can find class information, by country, at the IBM Learning Services Web site:

<http://www.ibm.com/services/learning/>

For more information, including the current local schedule, please contact your IBM representative.

Classes can also be taught at your location, at a time that suits your needs. Courses can even be customized to meet your exact requirements. The *All-in-One Education and Training Catalog* describes the DB2 curriculum in the United States.

You can inquire about or enroll in these courses by calling 1-800-IBM-TEACH (1-800-426-8322).

## How to order the DB2 library

You can order DB2 publications and CD-ROMs through your IBM representative or the IBM branch office serving your locality. If you are located within the United States or Canada, you can place your order by calling one of the toll-free numbers :

- In the U.S., call 1-800-879-2755.
- In Canada, call 1-800-565-1234.

To order additional copies of licensed publications, specify the SOFTWARE option. To order additional publications or CD-ROMs, specify the PUBLICATIONS and SLSS option. Be prepared to give your customer number, the product number, and the feature code(s) or order numbers you want.

---

## Summary of changes to DB2 UDB for OS/390 Version 6

DB2 UDB for OS/390 Version 6 delivers an enhanced relational database server solution for OS/390. This release focuses on greater capacity, performance improvements for utilities and queries, easier database management, more powerful network computing, and DB2 family compatibility with rich new object-oriented capability, triggers, and more built-in functions.

### Capacity improvements

**16-terabyte tables** provide a significant increase to table capacity for partitioned and LOB table spaces and indexes, and for nonpartitioning indexes.

**Buffer pools in data spaces** provide virtual storage constraint relief for the ssnmDBM1 address space, and data spaces increase the maximum amount of virtual buffer pool space allowed.

### Performance and availability

**Improved partition rebalancing** lets you redistribute partitioned data with minimal impact to data availability. One REORG of a range of partitions both reorganizes and rebalances the partitions.

You can **change checkpoint frequency dynamically** using the new SET LOG command and initiate checkpoints any time while your subsystem remains available.

**Utilities that are faster, more parallel, easier to use:**

- **Faster backup and recovery** enables COPY and RECOVER to process a list of objects in parallel, and recover indexes and table spaces at the same time from image copies and the log.
- **Parallel index build** reduces the elapsed time of LOAD and REORG jobs of table spaces, or partitions of table spaces, that have more than one index; the elapsed time of REBUILD INDEX jobs is also reduced.
- Tests show **decreased elapsed and processor time for online REORG.**
- **Inline statistics** embeds statistics collection into utility jobs, making table spaces available sooner.

- You can **determine when to run REORG** by specifying threshold limits for relevant statistics from the DB2 catalog.

**Query performance enhancements** include:

- **Query parallelism extensions** for complex queries, such as outer joins and queries that use nonpartitioned tables
- **Improved workload balancing in a Parallel Sysplex®** that reduces elapsed time for a single query that is split across active DB2 members
- **Improved data transfer** that lets you request multiple DRDA query blocks when performing high-volume operations
- The ability to use an **index to access predicates with noncorrelated IN subqueries**
- **Faster query processing** of queries that include join operations

**More performance and availability enhancements** include:

- **Faster restart and recovery** with the ability to postpone backout work during restart, and a faster log apply process
- **Increased flexibility with 8-KB and 16-KB page sizes** for balancing different workload requirements more efficiently, and for controlling traffic to the coupling facility for some workloads
- **Direct-row access** using the new ROWID data type to re-access a row directly without using the index or scanning the table
- **Ability to retain prior access path** when you rebind a statement. You almost always get the same or a better access path. For the exceptional cases, 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.
- An **increased log output buffer size** (from 1000 4-KB to 100000 4-KB buffers) that improves log read and write performance

## Data sharing enhancements

**More caching options** use the coupling facility to improve performance in a data sharing environment for some applications by writing changed pages directly to DASD.

**Control of space map copy maintenance** with a new option avoids tracking of page changes, thereby optimizing performance of data sharing applications.

## User productivity

**Predictive governing** capabilities enhance the resource limit facility to help evaluate resource consumption for queries that run against large volumes of data.

**Statement cost estimation** of processing resource that is needed for an SQL statement helps you to determine error and warning thresholds for governing, and to decide which statements need tuning.

A **default buffer pool** for user data and indexes isolates user data from the DB2 catalog and directory, and separating user data from system data helps you make better tuning decisions.

**More information available for monitoring DB2** includes data set I/O activity in traces, both for batch reporting and online monitors.

**Better integration of DB2 and Workload Manager delay reporting** enables DB2 to notify Workload Manager about the current state of a work request.

**More tables are allowed in SQL statements** SELECT, UPDATE, INSERT, and DELETE, and in views. DB2 increases the limit from 15 to 225 tables. The number of tables and views in a subselect is not changed.

**Improved DB2 UDB family compatibility** includes SQL extensions, such as:

- A VALUES clause of INSERT that supports any expression
- A new VALUES INTO statement

**Easier recovery management** lets you achieve a single point of recovery and recover data at a remote site more easily.

**Enhanced database commands** extend support for pattern-matching characters (\*) and let you filter display output.

You can easily **process dynamic SQL in batch mode** with the new object form of DSNTEP2 shipped with DB2 for OS/390.

## Network computing

**SQLJ**, the newest Java implementation for the OS/390 environment, supports SQL embedded in the Java programming language. With SQLJ, your Java programs benefit from the superior performance, manageability, and authorization available to static SQL, and they are easy to write.

**DRDA® support for three-part names** offers more functionality to applications using three-part names for remote access and improves the performance of client/server applications.

**Stored procedure enhancements** include the ability to create and modify stored procedure definitions, make nested calls for stored procedures and user-defined functions, and imbed CALL statements in application programs or dynamically invoke CALL statements from IBM's ODBC and CLI drivers.

**DB2 ODBC extensions** include new and modified APIs and new data types to support the object-relational extensions.

**ODBC access to DB2 for OS/390 catalog data** improves the performance of your ODBC catalog queries by redirecting them to shadow copies of DB2 catalog tables.

**Better performance for ODBC applications** reduces the number of network messages that are exchanged when an application executes dynamic SQL.

**Improvements for dynamically prepared SQL statements** include a new special register that you use to implicitly qualify names of distinct types, user-defined functions, and stored procedures.

**DDF connection pooling** uses a new type of inactive thread that improves performance for large volumes of inbound DDF connections.

## Object-relational extensions and active data

The object extensions of DB2 offer the benefits of object-oriented technology while increasing the strength of your relational database with an enriched set of data types and functions. Complementing these extensions is a powerful mechanism, triggers, that brings application logic into the database that governs the following new structures:

- **Large objects (LOBs)** are well suited to represent large, complex structures in DB2 tables. Now you can make effective use of multimedia by storing objects such as complex documents, videos, images, and voice. Some key elements of LOB support include:
  - LOB data types for storing byte strings up to 2 GB in size
  - LOB locators for easily manipulating LOB values in manageable pieces
  - Auxiliary tables (that reside in LOB table spaces) for storing LOB values
- **Distinct types** (which are sometimes called user-defined data types), like built-in data types, describe the data that is stored in columns of tables where the instances (or objects) of these data types are stored. They ensure that only those functions and operators that are explicitly defined on a distinct type can be applied to its instances.
- **User-defined functions**, like built-in functions or operators, support manipulation of distinct type instances (and built-in data types) in SQL queries.
- **New and extended built-in functions** improve the power of the SQL language with about 100 new built-in functions, extensions to existing functions, and sample user-defined functions.

**Triggers** automatically execute a set of SQL statements whenever a specified event occurs. These statements validate and edit database changes, read and modify the database, and invoke functions that perform operations inside and outside the database.

You can use the **DB2 Extenders** feature of DB2 for OS/390 to store and manipulate image, audio, video, and text objects. The extenders automatically capture and maintain object information and provide a rich body of APIs.

## More function

Some function and capability is available to both Version 6 and Version 5 users. Learn how to obtain these functions now, prior to migrating to Version 6, by visiting the following Web site:

<http://www.software.ibm.com/data/db2/os390/v5apar.html>

## Features of DB2 for OS/390

DB2 for OS/390 Version 6 offers a number of tools, which are optional features of the server, that are shipped to you automatically when you order DB2 Universal Database for OS/390:

- DB2 Management Tools Package, which includes the following elements:
  - DB2 UDB Control Center
  - DB2 Stored Procedures Builder
  - DB2 Installer

- DB2 Visual Explain
- DB2 Estimator
- Net.Data® for OS/390

You can install and use these features in a “Try and Buy” program for up to 90 days without paying license charges:

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

## Migration considerations

Migration to Version 6 eliminates all type 1 indexes, shared read-only data, data set passwords, use of host variables without the colon, and RECOVER INDEX usage. You can migrate to Version 6 only from a Version 5 subsystem.

---

## Summary of changes to this book

Listed below are the major changes to this edition of the book.

### Name changes

- In Version 5, DB2 for OS/390 introduced DB2 Call Level Interface (CLI). In Version 6, CLI is renamed ODBC (Open Database Connectivity), the name commonly known in the industry. This is a change in name only; the support and function remains the same.
- In OS/390 Release 6, the name OS/390 OpenEdition® is replaced with OS/390 UNIX System Services or OS/390 UNIX. Throughout this book, we use the name OS/390 UNIX.

### Improved performance

- The new DB2 ODBC shadow catalog, described in “Directing catalog queries to the DB2 ODBC shadow catalog” on page 371, provides fast catalog access for your ODBC catalog queries. Support for this enhancement includes:
  - The new DB2 ODBC initialization file keyword, CLISHEMA, described in “Initialization keywords” on page 63, lets you direct queries to the shadow catalog.
  - Changes to result set columns returned by `SQLColumns()`, `SQLForeignKeys()`, and `SQLProcedureColumns()`, described in “Chapter 5. Functions” on page 73.
- “SQLDescribeParam - Describe parameter marker” on page 138 describes a new API that obtains a description of input parameter markers.

### Support for object-relational extensions

Three new APIs are added for LOB support:

- “SQLGetLength - Retrieve length of a string value” on page 235.
- “SQLGetPosition - Return starting position of string” on page 238.



- “SQLGetSubString - Retrieve portion of a string value” on page 251.

Support for LOBs is also added to several APIs in “Chapter 5. Functions” on page 73.

Additional support for object-relational extensions includes:

- LOB data types: BLOB, CLOB, and DBCLOB, and the ROWID data type, described in “Data types and data conversion” on page 36 and “ODBC APIs and data types” on page 426.
- “Using large objects” on page 383 and “Optimizing for large object data” on page 411 provide guidance for using LOBs and LOB locators.
- The new FREE LOCATOR and HOLD LOCATOR SQL statements, described in “Using LOB locators” on page 384.
- Support for user-defined functions. The CURRENT PATH special register uses a new initialization file keyword, CURRENTFUNCTIONPATH, described in “DB2 ODBC initialization file” on page 60.
- Support for distinct types. “Using distinct types” on page 386 and “Casting distinct types” on page 411 provide guidance information.

### **Usability improvements**

- Changes to trace keywords, described in “DB2 ODBC initialization file” on page 60, make application and diagnostic traces easier to use.
- “Chapter 7. Problem diagnosis” on page 413 provides more detailed ODBC trace information

### **Organizational changes**

- All application and service diagnostic and debugging information is in a new chapter, “Chapter 7. Problem diagnosis” on page 413.
- Programming hints and tips are located in “Chapter 6. Using advanced features” on page 361. See “Programming hints and tips” on page 407.
- The “JDBC application support” chapter is removed. See *DB2 Application Programming Guide and Reference for Java™*.



---

## Chapter 2. Introduction to DB2 ODBC

DB2 Open Database Connectivity (ODBC) is IBM's callable SQL interface by the DB2 family of products. It is a 'C' and 'C++' application programming interface for relational database access, and it uses function calls to pass dynamic SQL statements as function arguments. It is an alternative to embedded dynamic SQL, but unlike embedded SQL, it does not require a precompiler.

DB2 ODBC is based on the Microsoft™ Open Database Connectivity (ODBC) specification, and the X/Open Call Level Interface specification. These specifications were chosen as the basis for the DB2 ODBC in an effort to follow industry standards and to provide a shorter learning curve for those application programmers already familiar with either of these *data source* interfaces. In addition, some DB2 specific extensions were added to help the DB2 application programmer specifically exploit DB2 features.

---

### DB2 ODBC background information

To understand DB2 ODBC or any callable SQL interface, it is helpful to understand what it is based on, and to compare it with existing interfaces.

The X/Open Company and the SQL Access Group jointly developed a specification for a callable SQL interface referred to as the X/Open Call Level Interface. The goal of this interface is to increase the portability of applications by enabling them to become independent of any one database product vendor's programming interface. Most of the X/Open Call Level Interface specification was accepted as part of the ISO Call Level Interface Draft International Standard (ISO CLI DIS).

Microsoft developed a callable SQL interface called Open Database Connectivity (ODBC) for Microsoft operating systems based on a preliminary draft of X/Open CLI. The Call Level Interface specifications in ISO, X/Open, ODBC, and DB2 ODBC continue to evolve in a cooperative manner to provide functions with additional capabilities.

The ODBC specification also includes an operating environment where data source specific ODBC drivers are dynamically loaded at run time by a driver manager based on the data source name provided on the connect request. The application is linked directly to a single driver manager library rather than to each DBMS's library. The driver manager mediates the application's function calls at run time and ensures they are directed to the appropriate DBMS specific ODBC driver.

The ODBC driver manager only knows about the ODBC-specific functions, that is, those functions supported by the DBMS for which no API is specified. Therefore, DBMS-specific functions cannot be directly accessed in an ODBC environment. However, DBMS-specific dynamic SQL statements are indirectly supported using a mechanism called the vendor escape clause. See "Using vendor escape clauses" on page 404 for detailed information.

ODBC is not limited to Microsoft operating systems. Other implementations are available, such as DB2 ODBC, or are emerging on various platforms.

## Differences between DB2 ODBC and ODBC version 2.0

DB2 ODBC is derived from the ISO Call Level Interface Draft International Standard (ISO CLI DIS) and ODBC Version 2.0.

If you port existing ODBC applications to DB2 for OS/390 or write a new application according to the ODBC specifications, you must comply with the specifications defined in this publication. However, before you write to any API, validate that the API is supported by DB2 for OS/390 and that the syntax and semantics are identical. If there are any differences, you must code to the APIs documented in this publication.

On the DB2 for OS/390 platform, no ODBC driver manager exists. Consequently, DB2 ODBC support is implemented as a CLI/ODBC driver/driver manager that is loaded at run time into the application address space. See “DB2 ODBC runtime environment” on page 49 for details about the DB2 ODBC runtime environment.

The DB2 UDB CLI executes on Windows and AIX as an ODBC driver, loaded by the Windows driver manager (Windows environment) or the Visi genic driver manager (UNIX platforms). In this context, DB2 ODBC support is limited to the ODBC specifications. Alternatively, an application can directly invoke the CLI application programming interfaces (APIs) including those not supported by ODBC. In this context, the set of APIs supported by DB2 UDB is referred to as the “Call Level Interface”. See *DB2 UDB Call Level Interface Guide and Reference*.

The use of DB2 ODBC in this publication refers to DB2 for OS/390 support of DB2 ODBC unless otherwise noted.

### ODBC features supported

DB2 ODBC support should be viewed as consisting of most of ODBC Version 2.0 as well as IBM extensions. Where differences exist, applications should be written to the specifications defined in this publication.

DB2 ODBC includes support of the following ODBC functions:

- All ODBC level 1 functions
- All ODBC level 2 functions with the following four exceptions:
  - SQLBrowseConnect()
  - SQLDrivers()
  - SQLSetPos()
  - SQLSetScrollOptions()
- Some X/Open CLI functions
- Some DB2 specific functions

For a complete list of supported functions, see “Function summary” on page 74.

The following DB2 features are available to both ODBC and DB2 ODBC applications:

- The double byte (graphic) data types
- Stored procedures
- Distributed unit of work (DUW) as defined by DRDA, two-phase commit
- Distinct types
- User-defined functions

DB2 ODBC contains extensions to access DB2 features that can not be accessed by ODBC applications:

- SQLCA access for detailed DB2 specific diagnostic information
- Control over null termination of output strings.
- Support of large objects (LOBs) and LOB locators

DB2 ODBC does not support the following functions (a deviation from the Microsoft ODBC Version 2.0 Specification):

- Asynchronous SQL
- Scrollable cursor support
- Interactive data source connection specified using `SQLBrowseConnect()` and `SQLDriverConnect()`.

For more information on the relationship between DB2 ODBC and ODBC, see Appendix A, “DB2 ODBC and ODBC” on page 425.

---

## Differences between DB2 ODBC and embedded SQL

An application that uses an embedded SQL interface requires a precompiler to convert the SQL statements into code, which is then compiled, bound to the data source, and executed. In contrast, a DB2 ODBC application does not have to be precompiled or bound, but instead uses a standard set of functions to execute SQL statements and related services at run time.

This difference is important because, traditionally, precompilers have been specific to each database product, which effectively ties your applications to that product. DB2 ODBC enables you to write portable applications that are independent of any particular database product. This independence means DB2 ODBC applications do not have to be recompiled or rebound to access different DB2 or DRDA data sources, but rather just connect to the appropriate data source at run time.

DB2 ODBC and embedded SQL also differ in the following ways:

- DB2 ODBC does not require the explicit declaration of cursors. They are generated by DB2 ODBC as needed. The application can then use the generated cursor in the normal cursor fetch model for multiple row `SELECT` statements and positioned `UPDATE` and `DELETE` statements.
- The `OPEN` statement is not used in DB2 ODBC. Instead, the execution of a `SELECT` automatically causes a cursor to be opened.
- Unlike embedded SQL, DB2 ODBC allows the use of parameter markers on the equivalent of the `EXECUTE IMMEDIATE` statement (the `SQLExecDirect()` function).
- A `COMMIT` or `ROLLBACK` in DB2 ODBC is issued using the `SQLTransact()` function call rather than by passing it as an SQL statement.
- DB2 ODBC manages statement related information on behalf of the application, and provides a *statement handle* to refer to it as an abstract object. This handle eliminates the need for the application to use product specific data structures.
- Similar to the statement handle, the *environment handle* and *connection handle* provide a means to refer to all global variables and connection specific information.

- DB2 ODBC uses the SQLSTATE values defined by the X/Open SQL CAE specification. Although the format and most of the values are consistent with values used by the IBM relational database products, there are differences. (There are also differences between ODBC SQLSTATES and the X/Open defined SQLSTATES). Refer to Table 150 on page 433 for a cross reference of all DB2 ODBC SQLSTATES.

Despite these differences, there is an important common concept between embedded SQL and DB2 ODBC:

DB2 ODBC can execute any SQL statement that can be prepared dynamically in embedded SQL.

Table 1 lists each DB2 for OS/390 SQL statement, and indicates whether or not it can be executed using DB2 ODBC.

Each DBMS might have additional statements that can be dynamically prepared, in which case DB2 ODBC passes them to the DBMS. There is one exception: COMMIT and ROLLBACK can be dynamically prepared by some DBMSs but are not passed. The SQLTransact() function should be used instead to specify either COMMIT or ROLLBACK.

Table 1 (Page 1 of 2). SQL statements

SQL statement	Dynamic <sup>a</sup>	DB2 ODBC <sup>c</sup>
ALTER TABLE	X	X
ALTER DATABASE	X	X
ALTER INDEX	X	X
ALTER STOGROUP	X	X
ALTER TABLESPACE	X	X
BEGIN DECLARE SECTION <sup>b</sup>		
CALL		X <sup>d</sup>
CLOSE		SQLFreeStmt()
COMMENT ON	X	X
COMMIT	X	SQLTransact()
CONNECT (type 1)		SQLConnect(), SQLDriverConnect()
CONNECT (type 2)		SQLConnect(), SQLDriverConnect()
CREATE { ALIAS, DATABASE, INDEX, STOGROUP, SYNONYM, TABLE, TABLESPACE, VIEW, DISTINCT TYPE }	X	X
DECLARE CURSOR <sup>b</sup>		SQLAllocStmt()
DECLARE STATEMENT		
DECLARE TABLE		
DELETE	X	X
DESCRIBE		SQLDescribeCol(), SQLColAttributes()
DROP	X	X
END DECLARE SECTION <sup>b</sup>		

Table 1 (Page 2 of 2). SQL statements

SQL statement	Dynamic <sup>a</sup>	DB2 ODBC <sup>c</sup>
EXECUTE		SQLExecute()
EXECUTE IMMEDIATE		SQLExecDirect()
EXPLAIN	X	X
FETCH		SQLFetch()
FREE LOCATOR <sup>d</sup>		X
GRANT	X	X
HOLD LOCATOR <sup>d</sup>		X
INCLUDE <sup>b</sup>		
INSERT	X	X
LABEL ON	X	X
LOCK TABLE	X	X
OPEN		SQLExecute(), SQLExecDirect()
PREPARE		SQLPrepare()
RELEASE		
REVOKE	X	X
ROLLBACK	X	SQLTransact()
select-statement	X	X
SELECT INTO		
SET CONNECTION		SQLSetConnection()
SET host_variable		
SET CURRENT DEGREE	X	X
SET CURRENT PACKAGESET		
SET CURRENT PATH	X	X
SET CURRENT SQLID	X	X
UPDATE	X	X
WHENEVER <sup>b</sup>		

**Note:**

- a All statements in this list can be coded as static SQL, but only those marked with X can be coded as dynamic SQL.
- b This statement is not executable.
- c An X indicates that this statement can be executed using either SQLExecDirect() or SQLPrepare() and SQLExecute(). If there is an equivalent DB2 ODBC function, the function name is listed.
- d Although this statement is not dynamic, DB2 ODBC allows the statement to be specified when calling either SQLExecDirect() or SQLPrepare() and SQLExecute().

---

## Advantages of using DB2 ODBC

DB2 ODBC provides a number of key features that offer advantages in contrast to embedded SQL. DB2 ODBC:

- Ideally suits the client-server environment in which the target data source is unknown when the application is built. It provides a consistent interface for executing SQL statements, regardless of which database server the application connects to.
- Lets you write portable applications that are independent of any particular database product. DB2 ODBC applications do not have to be recompiled or rebound to access different DB2 or DRDA data sources. Instead they connect to the appropriate data source at run time.
- Reduces the amount of management required for an application while in general use. Individual DB2 ODBC applications do not need to be bound to each data source. Bind files provided with DB2 ODBC need to be bound only once for all DB2 ODBC applications.
- Lets applications connect to multiple data sources from the same application.
- Allocates and controls data structures, and provides a handle for the application to refer to them. Applications do not have to control complex global data areas such as the SQLDA and SQLCA.
- Provides enhanced parameter input and fetching capability. You can specify arrays of data on input to retrieve multiple rows of a result set directly into an array. You can execute statements that generate multiple result sets.
- Lets you retrieve multiple rows and result sets generated from a call to a stored procedure.
- Provides a consistent interface to query catalog information that is contained in various DBMS catalog tables. The result sets that are returned are consistent across DBMSs. Application programmers can avoid writing version-specific and server-specific catalog queries.
- Provides extended data conversion which requires less application code when converting information between various SQL and C data types.
- Aligns with the emerging ISO CLI standard in addition to using the accepted industry specifications of ODBC and X/Open CLI.
- Allows application developers to apply their knowledge of industry standards directly to DB2 ODBC. The interface is intuitive for programmers who are familiar with function libraries but know little about product specific methods of embedding SQL statements into a host language.

---

## Deciding which interface to use

DB2 ODBC is ideally suited for query-based applications that require portability. Use the following guidelines to help you decide which interface meets your needs.



## Static and dynamic SQL

Only embedded SQL applications can use static SQL. Both static and dynamic SQL have advantages. Consider these factors:

- Performance

Dynamic SQL is prepared at run time. Static SQL is prepared at bind time. The preparation step for dynamic SQL requires more processing and might incur additional network traffic.

However, static SQL does not always perform better than dynamic SQL. Dynamic SQL can make use of changes to the data source, such as new indexes, and can use current catalog statistics to choose the optimal access plan.

- Encapsulation and security

In static SQL, authorization to objects is associated with a package and validated at package bind time. Database administrators can grant execute authority on a particular package to a set of users rather than grant explicit access to each database object.

In dynamic SQL, authorization is validated at run time on a per statement basis; therefore, users must be granted explicit access to each database object.

## Use both interfaces

An application can take advantage of both static and dynamic interfaces. An application programmer can create a stored procedure that contains static SQL. The stored procedure is called from within a DB2 ODBC application and executed on the server. After the stored procedure is created, any DB2 ODBC or ODBC application can call it.

## Write a mixed application

You can write a mixed application that uses both DB2 ODBC and embedded SQL. In this scenario, DB2 ODBC provides the base application, and you write key modules using static SQL for performance or security. Choose this option only if stored procedures do not meet your applications requirements.

---

## Other information sources

Application developers should refer to *Microsoft ODBC 3.0 Programmer's Reference and SDK Guide* as a supplement to this publication.

When writing DB2 ODBC applications, you also might need to reference information for the database servers that are being accessed, in order to understand any connectivity issues, environment issues, SQL language support issues, and other server-specific information. For DB2 for OS/390 versions, see *DB2 SQL Reference* and *DB2 Application Programming and SQL Guide*. If you are writing applications that access other DB2 server products, see *IBM SQL Reference* for information that is common to all products, including any differences.



---

## Chapter 3. Writing a DB2 ODBC application

This section introduces a conceptual view of a typical DB2 ODBC application.

A DB2 ODBC application can be broken down into a set of tasks. Some of these tasks are organized into discrete steps, while others might apply throughout the application. Each task is carried out by calling one or more DB2 ODBC functions.

Tasks described in this section are basic tasks that apply to all applications. More advanced tasks, such as using array insert, are described in “Chapter 6. Using advanced features” on page 361.

The functions are used in examples to illustrate their use in DB2 ODBC applications. See “Chapter 5. Functions” on page 73 for complete descriptions and usage information for each of the functions.




Figure 1. Conceptual view of a DB2 ODBC application

Every DB2 ODBC application contains the three main tasks shown in Figure 1.

### Initialization

This task allocates and initializes some resources in preparation for the main *transaction processing* task. See “Initialization and termination” on page 22 for details.

### Transaction processing

This is the main task of the application. SQL statements are passed to DB2 ODBC to query and modify the data. See “Transaction processing” on page 27 for details.

### Termination

This task frees allocated resources. The resources generally consist of data areas identified by unique handles. See “Initialization and termination” on page 22 for details.

In addition to the three tasks listed above, there are general tasks, such as handling diagnostic messages, which occur throughout an application.

## Initialization and termination

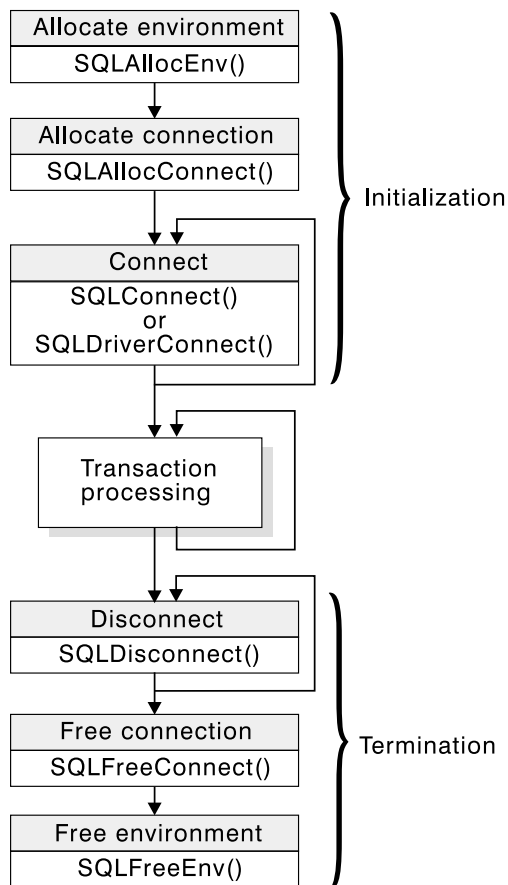


Figure 2. Conceptual view of initialization and termination tasks

Figure 2 shows the function call sequences for both the initialization and termination tasks. The transaction processing task in the middle of the diagram is shown in Figure 3 on page 27.

## Handles

The initialization task consists of the allocation and initialization of environment and connection handles (which are later freed in the termination task). An application then passes the appropriate handle when it calls other DB2 ODBC functions. A handle is a variable that refers to a data object controlled by DB2 ODBC. Using handles relieves the application from having to allocate and manage global variables or data structures, such as the SQLDA or SQLCA, used in IBM's embedded SQL interfaces.

There are three types of handles:

### Environment handle

The environment handle refers to the data object that contains information regarding the global state of the application, such as attributes and connections. This handle is allocated by calling `SQLAllocEnv()`, and freed by calling `SQLFreeEnv()`. An environment handle must be allocated before a connection handle can be allocated.

### **Connection handle**

A connection handle refers to a data object that contains information associated with a connection to a particular data source. This includes connection options, general status information, transaction status, and diagnostic information. Each connection handle is allocated by calling `SQLAllocConnect()` and freed by calling `SQLFreeConnect()`.

An application can be connected to several database servers at the same time. An application requires a connection handle for each concurrent connection to a database server. For information on multiple connections, see “Connecting to one or more data sources” on page 24.

Call `SQLGetInfo()` to determine if a user-imposed limit on the number of connection handles has been set.

### **Statement handles**

Statement handles are discussed in the next section, “Transaction processing” on page 27.

## **ODBC connection model**

The ODBC specifications support any number of concurrent connections, each of which is an independent transaction. That is, the application can issue `SQLConnect()` to X, perform some work, issue `SQLConnect()` to Y, perform some work, and then commit the work at X. ODBC supports multiple concurrent and independent transactions, one per connection.

### **DB2 ODBC restrictions on the ODBC connection model**

If the application is not using the `MULTICONTEXT=1` initialization file setting, there are restrictions on DB2 ODBC's support of the ODBC connection model. To obtain simulated support of the ODBC connection model, the application must specify a `CONNECT` type value of 1 (either by using the initialization file or the `SQLSetConnectOption()` API. See “Initialization keywords” on page 63 and “Specifying the connect type” on page 24.)

The application can then logically connect to any number of data sources. However, the DB2 ODBC driver maintains only a single physical connection, that of the last data source to which the application successfully connected or at which the last SQL statement was executed.

As a result, the application is affected as follows:

- When connected to one or more data sources so that the application has allocated some number of connect handles, any attempt to connect to a new data source commits the work on the current data source and terminates that connection. Therefore, the application cannot have cursors concurrently open at two data sources (including cursors `WITH HOLD`).
- If the application is currently connected to X and has performed work at X that has not yet been committed or rolled back, then any execution of an API to perform work on a valid statement handle Y results in committing the transaction at X and reestablishing the connection to Y.

With multiple context support, DB2 ODBC can fully support the ODBC connection model. See “DB2 ODBC support of multiple contexts” on page 396.

## Connect type 1 and type 2

Every IBM RDBMS supports both type 1 and type 2 CONNECT semantics. In both cases, there is only one transaction active at any time.

CONNECT (type 1) lets the application connect to only a single database at any time so that the single transaction is active on the current connection. This models the DRDA remote unit of work (RUW) processing.

Conversely, CONNECT (type 2) connect lets the application connect concurrently to any number of database servers, all of which participate in the single transaction. This models the DRDA distributed unit of work (DUW) processing.

ODBC does not support multiple connections participating in a distributed transaction.

### Specifying the connect type

**Important:** The connect type must be established prior to issuing `SQLConnect()`.

You can establish the connect type using either of the following methods:

- Specify the `CONNECTTYPE` keyword in the common section of the initialization file with a value of 1 (CONNECT (type 1)) or 2 (CONNECT (type 2)). The initialization file is described in “DB2 ODBC initialization file” on page 60.
- Invoke `SQLSetConnectOption()`. Specify `fOption = SQL_CONNECTTYPE` with a value of `SQL_CONCURRENT_TRANS` (CONNECT (type 1)) or a value of `SQL_COORDINATED_TRANS` (CONNECT (type 2)).

## Connecting to one or more data sources

DB2 ODBC supports connections to remote data sources through DRDA.

If the application is executing with CONNECT (type 1) and `MULTICONTEXT=0`, then DB2 ODBC allows an application to *logically* connect to multiple data sources; however, all transactions other than the transaction associated with the current connection, must be complete (committed or rolled back). If the application is executing with CONNECT (type 2), then the transaction is a distributed unit of work and all data sources participate in the disposition of the transaction (commit or rollback).

To connect concurrently to one or more data sources, an application calls `SQLAllocConnect()` once for each connection. The subsequent connection handle is used with `SQLConnect()` to request a data source connection and with `SQLAllocStmt()` to allocate statement handles for use within that connection. There is also an extended connect function, `SQLDriverConnect()`, which allows for additional connect options.

Unlike the distributed unit of work connections described in “Distributed unit of work (coordinated distributed transactions)” on page 363, there is no coordination between the statements that are executed on different connections.

## Initialization and connection example

```
/* ... */
/*****
**   - demonstrate basic connection to two datasources.
**   - error handling mostly ignored for simplicity
**
**   Functions used:
**
**   SQLAllocConnect  SQLDisconnect
**   SQLAllocEnv      SQLFreeConnect
**   SQLConnect        SQLFreeEnv
**   Local Functions:
**   DBconnect
**
*****/

#include <stdio.h>
#include <stdlib.h>
#include "sqlcli1.h"

int
DBconnect(SQLHENV henv,
          SQLHDBC * hdbc,
          char * server);

#define MAX_UID_LENGTH 18
#define MAX_PWD_LENGTH 30
#define MAX_CONNECTIONS 2

int
main( )
{
    SQLHENV      henv;
    SQLHDBC      hdbc[MAX_CONNECTIONS];
    char *      svr[MAX_CONNECTIONS] =
        {
            "KARACHI" ,
            "DAMASCUS"
        }

    /* allocate an environment handle */
    SQLAllocEnv(&henv);

    /* Connect to first data source */
    DBconnect(henv, &hdbc[0],
              svr[0]);

    /* Connect to second data source */
    DBconnect(henv, &hdbc[1],
              svr[1]);
}
```

```

/***** Start Processing Step *****/
/* allocate statement handle, execute statement, etc. */
/***** End Processing Step *****/

/*****
/* Commit work on connection 1. */
*****/

SQLTransact (henv,
             hdbc[0],
             SQL_COMMIT);

/*****
/* Commit work on connection 2. This has NO effect on the */
/* transaction active on connection 1. */
*****/

SQLTransact (henv,
             hdbc[1],
             SQL_COMMIT);

printf("\nDisconnecting ..... \n");

SQLDisconnect(hdbc[0]); /* disconnect first connection */

SQLDisconnect(hdbc[1]); /* disconnect second connection */
SQLFreeConnect(hdbc[0]); /* free first connection handle */
SQLFreeConnect(hdbc[1]); /* free second connection handle */
SQLFreeEnv(henv); /* free environment handle */

return (SQL_SUCCESS);
}

/*****
** Server is passed as a parameter. Note that USERID and PASSWORD**
** are always NULL. **
*****/

int
DBconnect(SQLHENV henv,
          SQLHDBC * hdbc,
          char * server)
{
    SQLRETURN rc;
    SQLCHAR buffer[255];
    SQLSMALLINT outlen;

    SQLAllocConnect(henv, hdbc); /* allocate a connection handle */

    rc = SQLConnect(*hdbc, server, SQL_NTS, NULL, SQL_NTS, NULL, SQL_NTS);
    if (rc != SQL_SUCCESS) {
        printf(">--- Error while connecting to database: %s ----- \n", server);
        return (SQL_ERROR);
    } else {
        printf(">Connected to %s \n", server);
        return (SQL_SUCCESS);
    }
}
/* ... */

```



## Transaction processing

The following figure shows the typical order of function calls in a DB2 ODBC application. Not all functions or possible paths are shown.

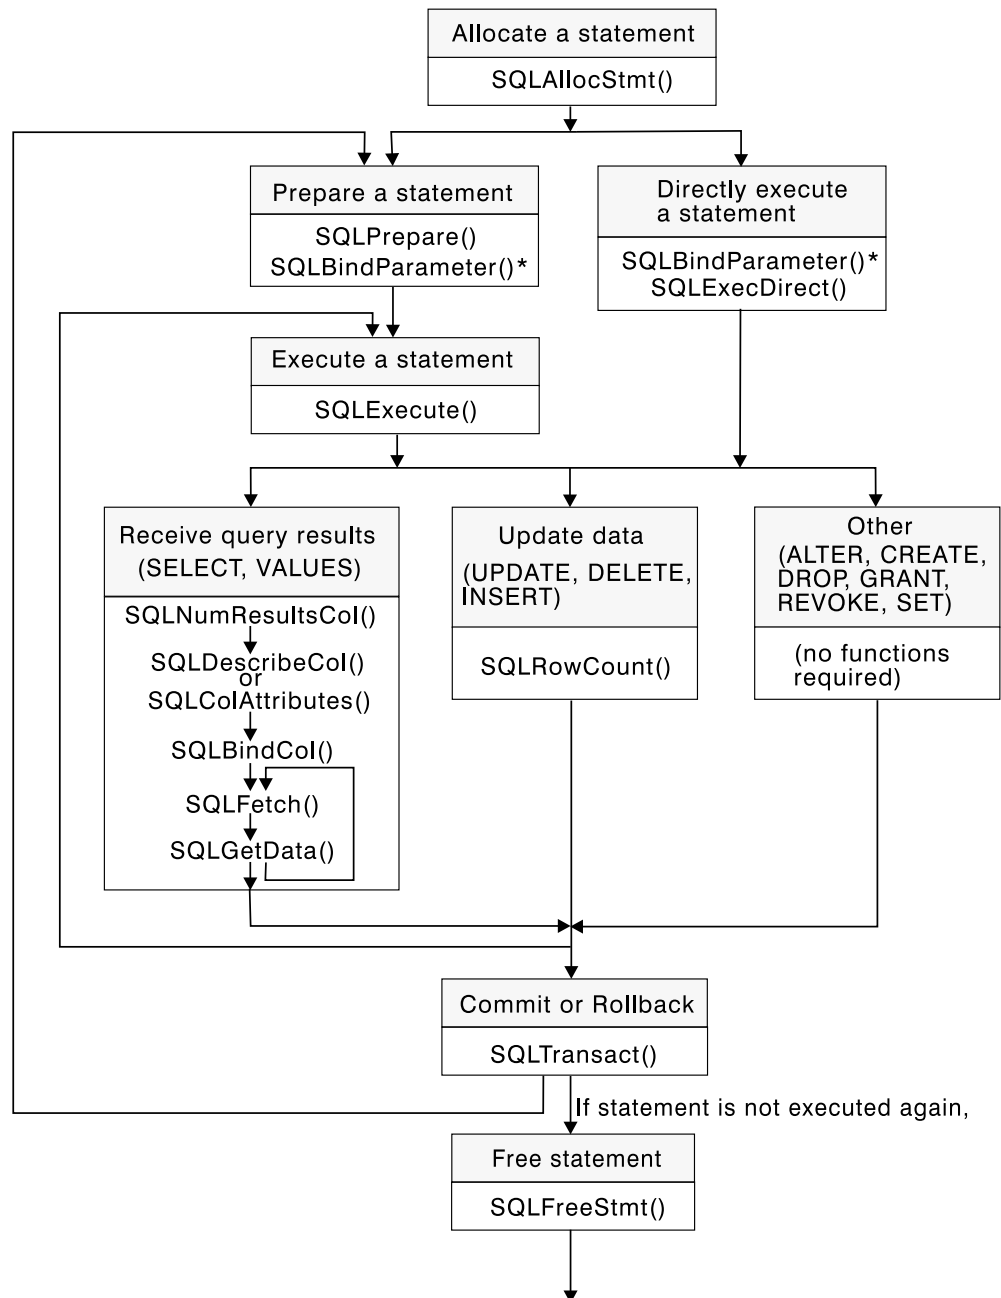


Figure 3. Transaction processing

Figure 3 shows the steps and the DB2 ODBC functions in the transaction processing task. This task contains five steps:

- Allocating statement handles
- Preparation and execution of SQL statements
- Processing result
- Commit or rollback

- Optionally, freeing statement handles if the statement is unlikely to be executed again.

Generally, you can use the `SQLSetParam()` function in place of the `SQLBindParameter()` function. `SQLSetParam()` is slightly simpler.

**Recommendation:** Use the `SQLBindParameter()` function which is more current.

## Allocating statement handles

`SQLAllocStmt()` allocates a statement handle. A statement handle refers to the data object that is used to track the execution of a single SQL statement. This includes information such as statement options, SQL statement text, dynamic parameters, cursor information, bindings for dynamic arguments and columns, result values and status information (these are discussed later). Each statement handle is associated with a unique connection handle.

A statement handle must be allocated before a statement can be executed. By default, the maximum number of statement handles that can be allocated at any one time is limited by the application heap size. The maximum number of statement handles that can actually be used, however, is defined by DB2 ODBC. Table 2 lists the number of statement handles allowed for each isolation level. If an application exceeds these limits, `SQLSTATE S1014` is returned on the call to `SQLPrepare()` or `SQLExecDirect()`.

*Table 2. Maximum number of statement handles allocated at one time*

Isolation level	Without hold	With hold	Total
Cursor stability	296	254	550
No commit	296	254	550
Repeatable read	296	254	550
Read stability	296	254	550
Uncommitted read	296	254	550

## Preparation and execution

After a statement handle is allocated, there are two methods of specifying and executing SQL statements:

1. Prepare then execute
  - a. Call `SQLPrepare()` with an SQL statement as an argument.
  - b. Call `SQLBindParameter()`, or `SQLSetParam()` if the SQL statement contains *parameter markers*.
  - c. Call `SQLExecute()`.
2. Execute direct
  - a. Call `SQLBindParameter()` or `SQLSetParam()` if the SQL statement contains *parameter markers*.
  - b. Call `SQLExecDirect()` with an SQL statement as an argument.

The first method splits the preparation of the statement from the execution. This method is used when:

- The statement is executed repeatedly (usually with different parameter values). This avoids having to prepare the same statement more than once. The

subsequent executions make use of the access plans already generated by the `prepare`.

- The application requires information about the columns in the result set, prior to statement execution.

The second method combines the `prepare` step and the `execute` step into one. This method is used when:

- The statement is executed only once. This avoids having to call two functions to execute the statement.
- The application does not require information about the columns in the result set, before the statement is executed.

DB2 for OS/390 and DB2 UDB provide *dynamic statement caching* at the database server. In DB2 ODBC terms this means that for a given statement handle, once a statement is prepared, it does not need to be prepared again (even after commits or rollbacks), as long as the statement handle is not freed. Applications that repeatedly execute the same SQL statement across multiple transactions, can save a significant amount of processing time and network traffic by:

1. Associating each such statement with its own statement handle, and
2. Preparing these statements once at the beginning of the application, then
3. Executing the statements as many times as is needed throughout the application.

## Binding parameters in SQL statements

Both of the execution methods described above, allow the use of parameter markers in place of an *expression* (or host variable in embedded SQL) in an SQL statement.

Parameter markers are represented by the '?' character and indicate the position in the SQL statement where the contents of application variables are to be substituted when the statement is executed. The parameter markers are referenced sequentially, from left to right, starting at 1. `SQLNumParams()` can be used to determine the number of parameters in a statement.

When an application variable is associated with a parameter marker it is *bound* to the parameter marker. The application must bind an application variable to each parameter marker in the SQL statement before it executes that statement. Binding is carried out by calling the `SQLBindParameter()` function with a number of arguments to indicate, the numerical position of the parameter, the SQL type of the parameter, the data type of the variable, a pointer to the application variable, and length of the variable.

The bound application variable and its associated length are called *deferred* input arguments since only the pointers are passed when the parameter is bound; no data is read from the variable until the statement is executed. Deferred arguments allow the application to modify the contents of the bound parameter variables, and repeat the execution of the statement with the new values.

Information for each parameter remains in effect until overridden, or until the application unbinds the parameter or drops the statement handle. If the application executes the SQL statement repeatedly without changing the parameter binding, then DB2 ODBC uses the same pointers to locate the data on each execution. The

application can also change the parameter binding to a different set of deferred variables. The application must not deallocate or discard variables used for deferred input fields between the time it binds the fields to parameter markers and the time DB2 ODBC accesses them at execution time.

It is possible to bind the parameters to a variable of a different type from that required by the SQL statement. The application must indicate the C data type of the source, and the SQL type of the parameter marker, and DB2 ODBC converts the contents of the variable to match the SQL data type specified. For example, the SQL statement might require an integer value, but your application has a string representation of an integer. The string can be bound to the parameter, and DB2 ODBC converts the string to the corresponding integer value when you execute the statement.

`SQLDescribeParam()` can be used to determine the data type of a parameter marker. If the application indicates an incorrect type for the parameter marker, either an extra conversion by the DBMS, or an error can occur. See “Data types and data conversion” on page 36 for more information about data conversion.

A parameter marker that is part of a predicate on a query and is associated with a distinct type must be cast to the built-in type in the predicate portion of the statement. Otherwise, an error occurs. For an example of casting distinct types, see “Casting distinct types” on page 411.

For information on more advanced methods of binding application storage to parameter markers, see “Using arrays to input parameter values” on page 375 and “Sending/retrieving long data in pieces” on page 373.

## Processing results

The next step after the statement has been executed depends on the type of SQL statement.

### Processing query (SELECT, VALUES) statements

If the statement is a query statement, the following steps are generally needed in order to retrieve each row of the result set:

1. Establish (describe) the structure of the result set, number of columns, column types and lengths
2. (Optionally) bind application variables to columns in order to receive the data
3. Repeatedly fetch the next row of data, and receive it into the bound application variables
4. (Optionally) retrieve columns that were not previously bound, by calling `SQLGetData()` after each successful fetch

Each of the above steps requires some diagnostic checks. “Chapter 6. Using advanced features” on page 361 discusses advanced techniques of using `SQLExtendedFetch()` to fetch multiple rows at a time.

#### Step 1

Analyze the executed or prepared statement. If the SQL statement was generated by the application, then this step might not be necessary since the application might know the structure of the result set and the data types of each column. If the application does know the structure of the entire result set, and if there are a very large number of columns to retrieve, then the application might wish to supply DB2 ODBC with the descriptor

information. This can reduce network traffic since DB2 ODBC does not have to retrieve the information from the server.

On the other hand, if the SQL statement was generated at runtime (for example, entered by a user), then the application has to query the number of columns, the type of each column, and perhaps the names of each column in the result set. This information can be obtained by calling `SQLNumResultCols()` and `SQLDescribeCol()` (or `SQLColAttributes()`) after preparing or after executing the statement.

## Step 2

The application retrieves column data directly into an application variable on the next call to `SQLFetch()`. For each column to be retrieved, the application calls `SQLBindCol()` to bind an application variable to a column in the result set. The application can use the information obtained from Step 1 to determine the C data type of the application variable and to allocate the maximum storage the column value could occupy. Similar to variables bound to parameter markers using `SQLBindParameter()` and `SQLSetParam()`, columns are bound to deferred arguments. This time the variables are deferred output arguments, as data is written to these storage locations when `SQLFetch()` is called.

If the application does not bind any columns, as in the case when it needs to retrieve columns of long data in pieces, it can use `SQLGetData()`. Both the `SQLBindCol()` and `SQLGetData()` techniques can be combined if some columns are bound and some are unbound. The application must not deallocate or discard variables used for deferred output fields between the time it binds them to columns of the result set and the time DB2 ODBC writes the data to these fields.

## Step 3

Call `SQLFetch()` to fetch the first or next row of the result set. If any columns are bound, the application variable is updated. There is also a method that allows the application to fetch multiple rows of the result set into an array, see “Retrieving a result set into an array” on page 379 for more information.

If data conversion was indicated by the data types specified on the call to `SQLBindCol()`, the conversion occurs when `SQLFetch()` is called. See “Data types and data conversion” on page 36 for an explanation.

## Step 4 (optional)

Call `SQLGetData()` to retrieve any unbound columns. All columns can be retrieved this way, provided they were not bound. `SQLGetData()` can also be called repeatedly to retrieve large columns in smaller pieces, which cannot be done with bound columns.

Data conversion can also be indicated here, as in `SQLBindCol()`, by specifying the desired target C data type of the application variable. See “Data types and data conversion” on page 36 for more information.

To unbind a particular column of the result set, use `SQLBindCol()` with a null pointer for the application variable argument (*rgbValue*). To unbind all of the columns with one function call, use `SQLFreeStmt()`.

Applications generally perform better if columns are bound rather than retrieved using `SQLGetData()`. However, an application can be constrained in the amount of

long data that it can retrieve and handle at one time. If this is a concern, then `SQLGetData()` might be the better choice.

For information on more advanced methods for binding application storage to result set columns, see “Retrieving a result set into an array” on page 379 and “Sending/retrieving long data in pieces” on page 373.

### **Processing UPDATE, DELETE and INSERT statements**

If the statement is modifying data (UPDATE, DELETE or INSERT), no action is required, other than the normal check for diagnostic messages. In this case, `SQLRowCount()` can be used to obtain the number of rows affected by the SQL statement.

If the SQL statement is a positioned UPDATE or DELETE, it is necessary to use a *cursor*. A cursor is a moveable pointer to a row in the result table of an active query statement. (This query statement must contain the FOR UPDATE OF clause to ensure that the query is not opened as read-only.) In embedded SQL, cursor names are used to retrieve, update or delete rows. In DB2 ODBC, a cursor name is needed only for positioned UPDATE or DELETE SQL statements as they reference the cursor by name.

To update a row that was fetched, the application uses two statement handles, one for the fetch and one for the update. The application calls `SQLGetCursorName()` to obtain the cursor name. The application generates the text of a positioned UPDATE or DELETE, including this cursor name, and executes that SQL statement using a second statement handle. The application cannot reuse the fetch statement handle to execute a positioned UPDATE or DELETE as it is still in use. You can also define your own cursor name using `SQLSetCursorName()`, but it is best to use the generated name, since all error messages reference the generated name, rather than the name defined by `SQLSetCursorName()`.

### **Processing other statements**

If the statement neither queries nor modifies the data, then there is no further action other than the normal check for diagnostic messages.

## **Commit or rollback**

A *transaction* is a recoverable unit of work, or a group of SQL statements that can be treated as one atomic operation. This means that all the operations within the group are guaranteed to be completed (committed) or undone (rolled back), as if they were a single operation. A transaction can also be referred to as a unit of work or a logical unit of work. When the transaction spans multiple connections, it is referred to as a distributed unit of work.

DB2 ODBC supports two commit modes: *auto-commit* and *manual-commit*.

In auto-commit mode, every SQL statement is a complete transaction, which is automatically committed. For a non-query statement, the commit is issued at the end of statement execution. For a query statement, the commit is issued after the cursor is closed. Given a single statement handle, the application must not start a second query before the cursor of the first query is closed.

In manual-commit mode, transactions are started implicitly with the first access to the data source using `SQLPrepare()`, `SQLExecDirect()`, `SQLGetTypeInfo()`, or any function that returns a result set, such as those described in “Querying system

catalog information” on page 368. At this point a transaction begins, even if the call failed. The transaction ends when you use `SQLTransact()` to either rollback or commit the transaction. This means that any statements executed (on the same connection) between these are treated as one transaction.

The default commit mode is auto-commit (except when participating in a coordinated transaction, see “Distributed unit of work (coordinated distributed transactions)” on page 363). An application can switch between manual-commit and auto-commit modes by calling `SQLSetConnectOption()`. Typically, a query-only application might wish to stay in auto-commit mode. Applications that need to perform updates to the data source should turn off auto-commit as soon as the data source connection is established.

When multiple connections exist, each connection has its own transaction (unless `CONNECT` (type 2) is specified). Special care must be taken to call `SQLTransact()` with the correct connection handle to ensure that only the intended connection and related transaction is affected. Unlike distributed unit of work connections (described in “Distributed unit of work (coordinated distributed transactions)” on page 363), there is no coordination between the transactions on each connection.

### **When to call `SQLTransact()`**

If the application is in auto-commit mode, it never needs to call `SQLTransact()`, a commit is issued implicitly at the end of each statement execution.

In manual-commit mode, `SQLTransact()` must be called before calling `SQLDisconnect()`. If distributed unit of work is involved, additional rules can apply. See “Distributed unit of work (coordinated distributed transactions)” on page 363 for details.

**Recommendation:** An application that performs updates should not wait until the disconnect before committing or rolling back a transaction.

The other extreme is to operate in auto-commit mode, which is also not recommended as this adds extra processing. The application can modify the auto-commit mode by invoking the `SQLSetConnectOption()` function. See “Environment, connection, and statement options” on page 361 and the `SQLSetConnectOption()` function for information about switching between auto-commit and manual-commit.

Consider the following when deciding where in the application to end a transaction:

- If using `CONNECT` (type 1) with `MULTICONTEXT=0`, only the current connection can have an outstanding transaction. If using `CONNECT` (type 2), all connections participate in a single transaction.
- If using `MULTICONTEXT=1`, each connection can have an outstanding transaction.
- Various resources can be held while you have an outstanding transaction. Ending the transaction releases the resources for use by other users.
- When a transaction is successfully committed or rolled back, it is fully recoverable from the system logs. Open transactions are not recoverable.

## Effects of calling SQLTransact()

When a transaction ends:

- All locks on DBMS objects are released, except those that are associated with a held cursor.
- Prepared statements are preserved from one transaction to the next if the data source supports statement caching (DB2 for OS/390 Version 5 does). After a statement is prepared on a specific statement handle, it does not need to be prepared again even after a commit or rollback, provided the statement continues to be associated with the same statement handle.
- Cursor names, bound parameters, and column bindings are maintained from one transaction to the next.
- By default, cursors are preserved after a commit (but not a rollback). All cursors are defined using the WITH HOLD clause (except when connected to DB2 Server for VSE & VM, which does not support the WITH HOLD clause). For information about changing the default behavior, see “SQLSetStmtOption - Set statement option” on page 334.

For more information and an example see “SQLTransact - Transaction management” on page 357.

## Freeing statement handles

Call `SQLFreeStmt()` to end processing for a particular statement handle. This function can be used to do one or more of the following:

- Unbind all columns of the result set
- Unbind all parameter markers
- Close any cursors and discard any pending results
- Drop the statement handle, and release all associated resources

The statement handle can be reused for other statements provided it is not dropped. If a statement handle is reused for another SQL statement string, any cached access plan for the original statement is discarded.

The columns and parameters should always be unbound before using the handle to process a statement with a different number or type of parameters or a different result set; otherwise application programming errors might occur.

---

## Diagnostics

Diagnostics refers to dealing with warning or error conditions generated within an application. There are two levels of diagnostics when calling DB2 ODBC functions:

- Return codes
- Detailed diagnostics (SQLSTATEs, messages, SQLCA)

Each DB2 ODBC function returns the function return code as a basic diagnostic. The `SQLError()` function provides more detailed diagnostic information. The `SQLGetSQLCA()` function provides access to the SQLCA, if the diagnostic is reported by the data source. This arrangement lets applications handle the basic flow control, and the SQLSTATEs allow determination of the specific causes of failure.

The `SQLError()` function returns three pieces of information:

- SQLSTATE



- Native error: if the diagnostic is detected by the data source, this is the SQLCODE; otherwise, this is set to -99999.
- Message text: this is the message text associated with the SQLSTATE.

For the detailed function information and example usage, see “SQLException - Retrieve error information” on page 148.

For diagnostic information about DB2 ODBC traces and debugging, see “Chapter 7. Problem diagnosis” on page 413.

## Function return codes

The following table lists all possible return codes for DB2 ODBC functions.

*Table 3. DB2 ODBC function return codes*

Return code	Explanation
SQL_SUCCESS	The function completed successfully, no additional SQLSTATE information is available.
SQL_SUCCESS_WITH_INFO	The function completed successfully, with a warning or other information. Call <code>SQLException()</code> to receive the SQLSTATE and any other informational messages or warnings. The SQLSTATE class is '01'. See Table 150 on page 433.
SQL_NO_DATA_FOUND	The function returned successfully, but no relevant data was found. When this is returned after the execution of an SQL statement, additional information might be available which can be obtained by calling <code>SQLException()</code> .
SQL_NEED_DATA	The application tried to execute an SQL statement but DB2 ODBC lacks parameter data that the application had indicated would be passed at execute time. For more information, see “Sending/retrieving long data in pieces” on page 373.
SQL_ERROR	The function failed. Call <code>SQLException()</code> to receive the SQLSTATE and any other error information.
SQL_INVALID_HANDLE	The function failed due to an invalid input handle (environment, connection or statement handle). This is a programming error. No further information is available.

## SQLSTATEs

Since different database servers often have different diagnostic message codes, DB2 ODBC provides a standard set of *SQLSTATEs* that are defined by the X/Open SQL CAE specification. This allows consistent message handling across different database servers.

SQLSTATEs are alphanumeric strings of 5 characters (bytes) with a format of `ccsss`, where `cc` indicates class and `sss` indicates subclass. Any SQLSTATE that has a class of:

- '01', is a warning.
- 'S1', is generated by the DB2 ODBC driver.

**Note:** X/Open has reserved class 'HY' for ODBC/CLI implementations, which is currently equivalent to the 'S1' class. This might be a consideration if you intend to follow the X/Open and/or ISO CLI standard in the future.

For some error conditions, DB2 ODBC returns SQLSTATEs that differ from those states listed in the *Microsoft ODBC 3.0 Programmer's Reference and SDK Guide*. This is a result of DB2 ODBC following the X/Open SQL CAE and SQL92 specifications.

DB2 ODBC SQLSTATEs include both additional IBM-defined SQLSTATEs that are returned by the database server, and DB2 ODBC defined SQLSTATEs for conditions that are not defined in the X/Open specification. This allows for the maximum amount of diagnostic information to be returned.

Follow these guidelines for using SQLSTATEs within your application:

- Always check the function return code before calling `SQLError()` to determine if diagnostic information is available.
- Use the SQLSTATEs rather than the native error code.
- To increase your application's portability, only build dependencies on the subset of DB2 ODBC SQLSTATEs that are defined by the X/Open specification, and return the additional ones as information only. (Dependencies refers to the application making logic flow decisions based on specific SQLSTATEs.)

**Note:** It might be useful to build dependencies on the class (the first 2 characters) of the SQLSTATEs.

- For maximum diagnostic information, return the text message along with the SQLSTATE (if applicable, the text message also includes the IBM-defined SQLSTATE). It is also useful for the application to print out the name of the function that returned the error.

See Table 150 on page 433 for a listing and description of the SQLSTATEs explicitly returned by DB2 ODBC.

## SQLCA

Embedded applications rely on the SQLCA for all diagnostic information. Although DB2 ODBC applications can retrieve much of the same information by using `SQLError()`, there still might be a need for the application to access the SQLCA related to the processing of a statement. (For example, after preparing a statement, the SQLCA contains the relative cost of executing the statement.) The SQLCA only contains meaningful information if there was an interaction with the data source on the previous request (for example: connect, prepare, execute, fetch, disconnect).

The `SQLGetSQLCA()` function is used to retrieve this structure. See “SQLGetSQLCA - Get SQLCA data structure” on page 242 for more information.

---

## Data types and data conversion

When writing a DB2 ODBC application it is necessary to work with both SQL data types and C data types. This is unavoidable since the DBMS uses SQL data types, and the application must use C data types. This means the application must match C data types to SQL data types when transferring data between the DBMS and the application (when calling DB2 ODBC functions).

To help address this, DB2 ODBC provides symbolic names for the various data types, and manages the transfer of data between the DBMS and the application. It also performs data conversion (from a C character string to an SQL INTEGER type, for example) if required. To accomplish this, DB2 ODBC needs to know both the source and target data type. This requires the application to identify both data types using symbolic names.

## C and SQL data types

Table 4 on page 39 lists each of the SQL data types, with its corresponding symbolic name, and the default C symbolic name. These data types represent the combination of the ODBC V2.0 minimum, core, and extended data types. The ODBC extended data type SQL\_BIGINT is not supported. In addition, DB2 ODBC supports SQL\_GRAPHIC, SQL\_VARGRAPHIC and SQL\_LONGVARGRAPHIC.

### SQL data type

This column contains the SQL data types as they would appear in an SQL CREATE DDL statement. The SQL data types are dependent on the DBMS.

### Symbolic SQL data type

This column contains SQL symbolic names that are defined (in `sqlcli1.h`) as an integer value. These values are used by various functions to identify the SQL data types listed in the first column. See “Example” on page 135 for an example using these values.

### Default C symbolic data type

This column contains C symbolic names, also defined as integer values. These values are used in various function arguments to identify the C data type as shown in Table 5 on page 40. The symbolic names are used by various functions, (such as `SQLBindParameter()`, `SQLGetData()`, `SQLBindCol()`) to indicate the C data types of the application variables. Instead of explicitly identifying the C data type when calling these functions, `SQL_C_DEFAULT` can be specified instead, and DB2 ODBC assumes a default C data type based on the SQL data type of the parameter or column as shown by this table. For example, the default C data type of `SQL_DECIMAL` is `SQL_C_CHAR`.

Table 4. SQL symbolic and default data types

SQL data type	Symbolic SQL data type	Default symbolic C data type
BLOB	SQL_BLOB	SQL_C_BINARY
BLOB LOCATOR <sup>a</sup>	SQL_BLOB_LOCATOR	SQL_C_BLOB_LOCATOR
CHAR	SQL_CHAR	SQL_C_CHAR
CHAR FOR BIT DATA	SQL_BINARY	SQL_C_BINARY
CLOB	SQL_CLOB	SQL_C_CHAR
CLOB LOCATOR <sup>a</sup>	SQL_CLOB_LOCATOR	SQL_C_CLOB_LOCATOR
DATE	SQL_DATE	SQL_C_DATE
DBCLOB	SQL_DBCLOB	SQL_C_DBCHAR
DBCLOB LOCATOR <sup>a</sup>	SQL_DBCLOB_LOCATOR	SQL_C_DBCLOB_LOCATOR
DECIMAL	SQL_DECIMAL	SQL_C_CHAR
DOUBLE	SQL_DOUBLE	SQL_C_DOUBLE
FLOAT	SQL_FLOAT	SQL_C_DOUBLE
GRAPHIC	SQL_GRAPHIC	SQL_C_DBCHAR
INTEGER	SQL_INTEGER	SQL_C_LONG
LONG VARCHAR <sup>b</sup>	SQL_LONGVARCHAR	SQL_C_CHAR
LONG VARCHAR FOR BIT DATA <sup>b</sup>	SQL_LONGVARBINARY	SQL_C_BINARY
LONG VARGRAPHIC <sup>b</sup>	SQL_LONGVARGRAPHIC	SQL_C_DBCHAR
NUMERIC <sup>c</sup>	SQL_NUMERIC <sup>c</sup>	SQL_C_CHAR
REAL <sup>d</sup>	SQL_REAL	SQL_C_FLOAT
ROWID	SQL_ROWID	SQL_C_CHAR
SMALLINT	SQL_SMALLINT	SQL_C_SHORT
TIME	SQL_TIME	SQL_C_TIME
TIMESTAMP	SQL_TIMESTAMP	SQL_C_TIMESTAMP
VARCHAR	SQL_VARCHAR	SQL_C_CHAR
VARCHAR FOR BIT DATA <sup>b</sup>	SQL_VARBINARY	SQL_C_BINARY
VARGRAPHIC	SQL_VARGRAPHIC	SQL_C_DBCHAR

**Note:**

**a** LOB locator types are not persistent SQL data types (columns cannot be defined by a locator type; instead, it describes parameter markers, or represents a LOB value). See “Using large objects” on page 383 for more information.

**b** Whenever possible, replace long data types and FOR BIT DATA data types with appropriate LOB types.

**c** NUMERIC is a synonym for DECIMAL on DB2 for OS/390, DB2 for VSE & VM and DB2 UDB.

**d** REAL is not valid for DB2 UDB or DB2 for OS/390.

The data types, DATE, DECIMAL, NUMERIC, TIME, and TIMESTAMP cannot be transferred to their default C buffer types without a conversion.

Table 5 on page 40 shows the generic C type definitions for each symbolic C type.

**C symbolic data type**

This column contains C symbolic names, defined as integer values. These values are used in various function arguments to identify the C data type shown in the last column. See “Example” on page 92 for an example using these values.

## C type

This column contains C defined types, defined in `sqlcli1.h` using a C typedef statement. The values in this column should be used to declare all DB2 ODBC related variables and arguments, in order to make the application more portable. See Table 7 on page 41 for a list of additional symbolic data types used for function arguments.

## Base C type

This column is shown for reference only. All variables and arguments should be defined using the symbolic types in the previous column. Some of the values are C structures that are described in Table 6.

Table 5. C data types

C symbolic data type	C type	Base C type
SQL_C_CHAR	SQLCHAR	unsigned char
SQL_C_BIT	SQLCHAR	unsigned char or char (Value 1 or 0)
SQL_C_TINYINT	SQLSCHAR	signed char (Range -128 to 127)
SQL_C_SHORT	SQLSMALLINT	short int
SQL_C_LONG	SQLINTEGER	long int
SQL_C_DOUBLE	SQLDOUBLE	double
SQL_C_FLOAT	SQLREAL	float
SQL_C_DATE	DATE_STRUCT	see Table 6
SQL_C_TIME	TIME_STRUCT	see Table 6
SQL_C_TIMESTAMP	TIMESTAMP_STRUCT	see Table 6
SQL_C_CLOB_LOCATOR	SQLINTEGER	long int
SQL_C_BINARY	SQLCHAR	unsigned char
SQL_C_BLOB_LOCATOR	SQLINTEGER	long int
SQL_C_DBCHAR	SQLDBCHAR	wchar_t
SQL_C_DBCLOB_LOCATOR	SQLINTEGER	long int

Table 6 (Page 1 of 2). C date, time, and timestamp structures

C type	Generic structure
DATE_STRUCT	<pre>typedef struct DATE_STRUCT {     SQLSMALLINT  year;     SQLUSMALLINT month;     SQLUSMALLINT day; } DATE_STRUCT;</pre>
TIME_STRUCT	<pre>typedef struct TIME_STRUCT {     SQLUSMALLINT  hour;     SQLUSMALLINT  minute;     SQLUSMALLINT  second; } TIME_STRUCT;</pre>

Table 6 (Page 2 of 2). C date, time, and timestamp structures

C type	Generic structure
TIMESTAMP_STRUCT	<pre>typedef struct TIMESTAMP_STRUCT {     SQLUSMALLINT  year;     SQLUSMALLINT  month;     SQLUSMALLINT  day;     SQLUSMALLINT  hour;     SQLUSMALLINT  minute;     SQLUSMALLINT  second;     SQLINTEGER    fraction; } TIMESTAMP_STRUCT;</pre>

See Table 7 for more information on the SQLUSMALLINT C data type.

## Other C data types

In addition to the data types that map to SQL data types, there are also C symbolic types used for other function arguments, such as pointers and handles. Both the generic and ODBC data types are shown below.

Table 7. C data types and base C data types

Defined C type	Base C type	Typical usage
SQLPOINTER	void *	Pointers to storage for data and parameters.
SQLHENV	long int	Handle referencing environment information.
SQLHDBC	long int	Handle referencing data source connection information.
SQLHSTMT	long int	Handle referencing statement information.
SQLUSMALLINT	unsigned short int	Function input argument for unsigned short integer values.
SQLUIINTEGER	unsigned long int	Function input argument for unsigned long integer values.
SQLRETURN	short int	Return code from DB2 ODBC functions.

## Data conversion

As mentioned previously, DB2 ODBC manages the transfer and any required conversion of data between the application and the DBMS. Before the data transfer actually takes place, the source, target or both data types are indicated when calling `SQLBindParameter()`, `SQLBindCol()` or `SQLGetData()`. These functions use the symbolic type names shown in Table 4 on page 39, to identify the data types involved.

For example, to bind a parameter marker that corresponds to an SQL data type of `DECIMAL(5,3)`, to an application's C buffer type of `double`, the appropriate `SQLBindParameter()` call would look like:

```
SQLBindParameter (hstmt, 1, SQL_PARAM_INPUT, SQL_C_DOUBLE,
                 SQL_DECIMAL, 5, 3, double_ptr, NULL);
```

Table 4 shows only the default data conversions. The functions mentioned in the previous paragraph can be used to convert data to other types, but not all data conversions are supported or make sense. Table 8 on page 42 shows all the conversions supported by DB2 ODBC.

The first column in Table 8 contains the data type of the SQL data type, the remaining columns represent the C data types. If the C data type columns contains:

- D** The conversion is supported and is the default conversion for the SQL data type.
- X** All IBM DBMSs support the conversion,
- Blank** No IBM DBMS supports the conversion.

For example, the table indicates that a char (or a C character string as indicated in Table 8) can be converted into a SQL\_C\_LONG (a signed long). In contrast, a LONGVARCHAR cannot be converted to a SQL\_C\_LONG.

See Appendix D, “Data conversion” on page 443 for information about required formats and the results of converting between data types.

Limits on precision, and scale, as well as truncation and rounding rules for type conversions follow rules specified in the *IBM SQL Reference* with the following exception; truncation of values to the right of the decimal point for numeric values returns a truncation warning, whereas truncation to the left of the decimal point returns an error. In cases of error, the application should call `SQLERROR()` to obtain the `SQLSTATE` and additional information on the failure. When moving and converting floating point data values between the application and DB2 ODBC, no correspondence is guaranteed to be exact as the values can change in precision and scale.

Table 8 (Page 1 of 2). Supported data conversions

SQL data type	SQL_C_CHAR	SQL_C_LONG	SQL_C_SHORT	SQL_C_TINYINT	SQL_C_FLOAT	SQL_C_DOUBLE	SQL_C_DATE	SQL_C_TIME	SQL_C_TIMESTAMP	SQL_C_BINARY	SQL_C_BIT	SQL_C_DBCHAR	SQL_C_CLOB_LOCATOR	SQL_C_BLOB_LOCATOR	SQL_C_DBCLOB_LOCATOR
SQL_BLOB	X									D				X	
SQL_CHAR	X	X	X	X	X	X	X	X	X	X	X				
SQL_CLOB	D									X			X		
SQL_DATE	X						D		X						
SQL_DBCLOB										X		D			X
SQL_DECIMAL	D	X	X	X	X	X					X				
SQL_DOUBLE	X	X	X	X	X	D					X				
SQL_FLOAT	X	X	X	X	X	D					X				
SQL_GRAPHIC	X											D			
SQL_INTEGER	X	D	X	X	X	X					X				
SQL_LONGVARCHAR	D						X		X	X					
SQL_LONGVARGRAPHIC	X									X		D			
SQL_NUMERIC	D	X	X	X	X	X					X				
SQL_REAL	X	X	X	X	D	X					X				
SQL_ROWID	X														



Table 8 (Page 2 of 2). Supported data conversions

SQL data type	SQL_C_CHAR	SQL_C_LONG	SQL_C_SHORT	SQL_C_TINYINT	SQL_C_FLOAT	SQL_C_DOUBLE	SQL_C_DATE	SQL_C_TIME	SQL_C_TIMESTAMP	SQL_C_BINARY	SQL_C_BIT	SQL_C_DBCHAR	SQL_C_CLOB_LOCATOR	SQL_C_BLOB_LOCATOR	SQL_C_DBCLOB_LOCATOR
SQL_SMALLINT	X	X	D	X	X	X					X				
SQL_TIME	X							D	X						
SQL_TIMESTAMP	X						X	X	D						
SQL_VARCHAR	D	X	X	X	X	X	X	X	X	X	X				
SQL_VARGRAPHIC	X											D			

**Note:**

Data is not converted to LOB locator types; locators represent a data value.

REAL is not supported by DB2 UDB.

NUMERIC is a synonym for DECIMAL on DB2 for OS/390, DB2 for VSE & VM, and DB2 UDB.

## Working with string arguments

The following conventions deal with the various aspects of working with string arguments in DB2 ODBC functions.

### Length of string arguments

Input string arguments have an associated length argument. This argument indicates to DB2 ODBC, either the exact length of the argument (not including the null terminator), the special value SQL\_NTS to indicate a null-terminated string, or SQL\_NULL\_DATA to pass a NULL value. If the length is set to SQL\_NTS, DB2 ODBC determines the length of the string by locating the null terminator.

Output string arguments have two associated length arguments, an input length argument to specify the length of the allocated output buffer, and an output length argument to return the actual length of the string returned by DB2 ODBC. The returned length value is the total length of the string available for return, regardless of whether it fits in the buffer or not.

For SQL column data, if the output is a null value, SQL\_NULL\_DATA is returned in the length argument and the output buffer is untouched.

If a function is called with a null pointer for an output length argument, DB2 ODBC does not return a length, and assumes that the data buffer is large enough to hold the data. When the output data is a NULL value, DB2 ODBC can not indicate that the value is NULL. If it is possible that a column in a result set can contain a NULL value, a valid pointer to the output length argument must always be provided. It is highly recommended that a valid output length argument always be used.

If the length argument (*pcbValue*) and the output buffer (*rgbValue*) are contiguous in memory, DB2 ODBC can return both values more efficiently, improving

application performance. For example, if the following structure is defined and `&buffer.pcbValue` and `buffer.rgbValue` are passed to `SQLBindCol()`, DB2 ODBC updates both values in one operation.

```
struct
{
    SQLINTEGER pcbValue;
    SQLCHAR    rgbValue [BUFFER_SIZE];
} buffer;
```

## Null-termination of strings

By default, every character string that DB2 ODBC returns is terminated with a null terminator (hex 00), except for strings returned from the graphic and DBCLOB data types into `SQL_C_CHAR` application variables. Graphic and DBCLOB data types that are retrieved into `SQL_C_DBCHAR` application variables are null terminated with a double byte null terminator. This requires that all buffers allocate enough space for the maximum number of bytes expected, plus the null-terminator.

It is also possible to use `SQLSetEnvAttr()` and set an environment attribute to disable null termination of variable length output (character string) data. In this case, the application allocates a buffer exactly as long as the longest string it expects. The application must provide a valid pointer to storage for the output length argument so that DB2 ODBC can indicate the actual length of data returned; otherwise, the application has no means to determine this. The DB2 ODBC default is to always write the null terminator.

## String truncation

If an output string does not fit into a buffer, DB2 ODBC truncates the string to the size of the buffer, and writes the null terminator. If truncation occurs, the function returns `SQL_SUCCESS_WITH_INFO` and an `SQLSTATE` of **01004** indicating truncation. The application can then compare the buffer length to the output length to determine which string was truncated.

For example, if `SQLFetch()` returns `SQL_SUCCESS_WITH_INFO`, and an `SQLSTATE` of **01004**, at least one of the buffers bound to a column is too small to hold the data. For each buffer that is bound to a column, the application can compare the buffer length with the output length and determine which column was truncated.

ODBC specifies that string data can be truncated on input or output with the appropriate `SQLSTATE`. As the data source, an IBM relational database (DB2) does not truncate data on input, but might truncate data on output to maintain data integrity. On input, DB2 rejects string truncation with a negative `SQLCODE` (-302) and an `SQLSTATE` of **22001**. On output, DB2 truncates the data and issues `SQL_SUCCESS_WITH_INFO` and an `SQLSTATE` of **01004**.

## Interpretation of strings

Normally, DB2 ODBC interprets string arguments in a case-sensitive manner and does not trim any spaces from the values. The one exception is the cursor name input argument on the `SQLSetCursorName()` function. In this case, if the cursor name is not delimited (enclosed by double quotes) the leading and trailing blanks are removed and case is preserved.

---

## Querying environment and data source information

There are many situations when an application requires information about the characteristics and capabilities of the current DB2 ODBC driver or the data source that it is connected to.

One of the most common situations involves displaying information for the user. Information such as the data source name and version, or the version of the DB2 ODBC driver might be displayed at connect time, or as part of the error reporting process.

These functions are also useful to generic applications that are written to adapt and take advantage of facilities that might be available from some, but not all database servers. The following DB2 ODBC functions provide data source specific information:

- “SQLDataSources - Get list of data sources” on page 129
- “SQLGetFunctions - Get functions” on page 214
- “SQLGetInfo - Get general information” on page 219
- “SQLGetTypeInfo - Get data type information” on page 254

## Querying environment information example

```
/*
*****
/* Querying environment and data source information */
*****
#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <sqlcli1.h>

void main()
{
    SQLHENV      hEnv;          /* Environment handle          */
    SQLHDBC      hDbc;          /* Connection handle          */
    SQLRETURN    rc;           /* Return code for API calls  */
    SQLSTMT      hStmt;        /* Statement handle           */
    SQLCHAR      dsname[30];    /* Data source name           */
    SQLCHAR      dsdescr[255]; /* Data source description    */
    SQLSMALLINT  dslen;        /* Length of data source      */
    SQLSMALLINT  descLen;     /* Length of dsdescr         */
    BOOL         found = FALSE;
    SQLSMALLINT  funcs[100];
    SQLINTEGER   rgbValue;

    /*
     * Initialize environment - allocate environment handle.
     */
    rc = SQLAllocEnv( &hEnv );
    rc = SQLAllocConnect( hEnv, &hDbc );

    /*
     * Use SQLDataSources to verify MVSDB2 does exist.
     */
    while( ( rc = SQLDataSources( hEnv,
                                SQL_FETCH_NEXT,
                                dsname,
                                SQL_MAX_DSN_LENGTH+1,
                                &dslen,
                                dsdescr,
                                &descLen ) ) != SQL_NO_DATA_FOUND )
    {
        if( !strcmp( dsname, "MVSDB2" ) ) /* data source exist */
        {
            found = TRUE;
            break;
        }
    }

    if( !found )
    {
        fprintf(stdout, "Data source %s does not exist...\n", dsname );
        fprintf(stdout, "program aborted.\n");
        exit(1);
    }
}
```

```

if( ( rc = SQLConnect( hDbc, dsname, SQL_NTS, "myid", SQL_NTS, "mypd", SQL_NTS ) )
    == SQL_SUCCESS )
{
    fprintf( stdout, "Connect to %s\n", dsname );
}

SQLAllocStmt( hDbc, &hStmt );

/*
 * Use SQLGetFunctions to store all APIs status.
 */
rc = SQLGetFunctions( hDbc, SQL_API_ALL_FUNCTIONS, funcs );

/*
 * Check whether SQLGetInfo is supported in this driver. If so,
 * verify whether DATE is supported for this data source.
 */
if( funcs[SQL_API_SQLGETINFO] == 1 )
{
    SQLGetInfo( hDbc, SQL_CONVERT_FUNCTIONS, (SQLPOINTER)&rgbValue, 255, &descLen );
    if( rgbValue & SQL_CVT_DATE )
    {
        SQLGetTypeInfo( hStmt, SQL_DATE );

        /* use SQLBindCol and SQLFetch to retrieve data ....*/
    }
}
}
}

```



---

## Chapter 4. Configuring DB2 ODBC and running sample applications

This section provides information about installing DB2 ODBC, the DB2 ODBC runtime environment, and the preparation steps needed to run a DB2 ODBC application.

- “Installing DB2 ODBC”
- “DB2 ODBC runtime environment”
- “Setting up DB2 ODBC runtime environment” on page 52
- “Preparing and executing a DB2 ODBC application” on page 56
- “DB2 ODBC initialization file” on page 60
- “DB2 ODBC migration considerations” on page 72

#

---

### Installing DB2 ODBC

The steps in this section describe the SMP/E jobs you must edit and run to install DB2 ODBC. Customize these jobs to specify data set names for your DB2 installation and SMP/E data sets. See the header notes in each job and to Section 2 of *DB2 Installation Guide* for details.

#### Step 1: Copy and edit the SMP/E jobs

For sample JCL to invoke the MVS utility IEBCOPY to copy the SMP/E jobs to DASD, see the *DB2 Program Directory*.

#### Step 2: Run the receive job: DSNTCJRC

DSNTCJRC invokes SMP/E to receive the FMIDs for DB2 ODBC into the SMP/E control data sets.

#### Step 3: Run the apply job: DSNTCJAP

DSNTCJAP invokes SMP/E to apply the FMIDs for DB2 ODBC to the DB2 target libraries.

The DSNTCJAP job can include objects for Language Environment® (LE) or Callable Services (CSS) libraries. If maintenance is later applied to LE or CSS libraries, you might need to run the SMP REPORT CALLLIBS command to generate a summary report. See *OS/390 SMP/E Reference* for detailed information.

#### Step 4: Run the accept job: DSNTCJAC

DSNTCJAC invokes SMP/E to accept the FMIDs for DB2 ODBC into the DB2 distribution libraries.

---

### DB2 ODBC runtime environment

DB2 ODBC does not support an ODBC driver manager. All API calls are routed through the single ODBC driver that is loaded at run time into the application address space. DB2 ODBC support is implemented as an IBM C/C++ Dynamic Load Library (DLL). By providing DB2 ODBC support using a DLL, DB2 ODBC applications do not need to link-edit any DB2 ODBC driver code with the application

load module. Instead, the linkage to the DB2 ODBC APIs is resolved dynamically at runtime by the IBM Language Environment (LE) runtime support.

The DB2 ODBC driver can use either the call attachment facility (CAF) or the Recoverable Resource Manager Services attachment facility (RRSAF) to connect to the DB2 for OS/390 address space.

- If the DB2 ODBC application is not running as a DB2 for OS/390 stored procedure, the `MVSATTACHTYPE` keyword in the DB2 ODBC initialization file determines the attachment facility that DB2 ODBC uses.
- If the DB2 ODBC application is running as a DB2 for OS/390 stored procedure, then DB2 ODBC uses the attachment facility that was specified for stored procedures.

When the DB2 ODBC application invokes the first ODBC function, `SQLAllocEnv()`, the DB2 ODBC driver DLL is loaded.

DB2 ODBC supports access to the local DB2 for OS/390 subsystems and any remote data source that is accessible using DB2 for OS/390 Version 6. This includes:

- Remote DB2 subsystems using specification of an alias or three-part name
- Remote DRDA-1 and DRDA-2 servers using LU 6.2 or TCP/IP.

The relationship between the application, the DB2 for OS/390 V6 ODBC driver and the DB2 for OS/390 subsystem are illustrated in Figure 4 on page 51.



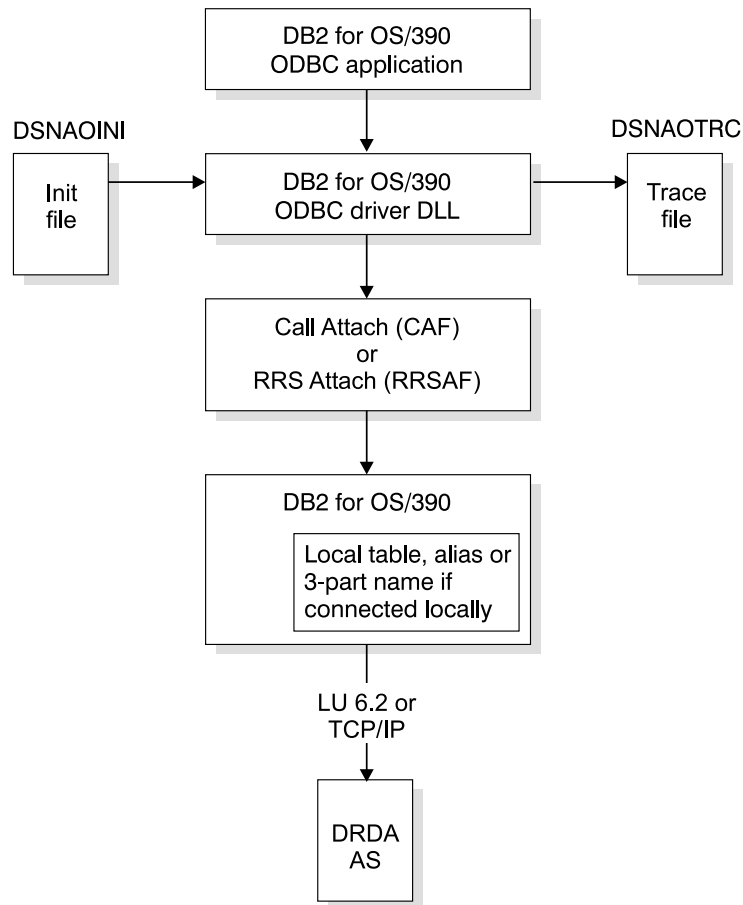


Figure 4. Relationship between DB2 for OS/390 V6 ODBC components

## Connectivity requirements

DB2 for OS/390 V6 ODBC has the following connectivity requirements:

- DB2 ODBC applications must execute on a machine on which Version 6 of DB2 for OS/390 is installed.
- If the application is executing with MULTICONTEXT=1, then there are multiple physical connections. Each connection corresponds to an independent transaction and DB2 thread.
- If the application is executing CONNECT (type 1) (described in “Connect type 1 and type 2” on page 24) and MULTICONTEXT=0, then there is only one current physical connection and one transaction on that connection. All transactions on logical connections (that is, with a valid connection handle) are rolled back by the application or committed by DB2 ODBC. This is a deviation from the ODBC connection model.

---

## Setting up DB2 ODBC runtime environment

This section describes the general setup required to enable DB2 ODBC applications. The steps in this section only need to be performed once, and are usually performed as part of the installation process for DB2 for OS/390.

The DB2 ODBC bind files must be bound to the data source. The following two bind steps are required:

- Create packages at every data source
- Create at least one plan to name those packages.

These bind steps are described in the following sections:

- “Bind DBRMs to packages”
- “Bind an application plan” on page 54

The online bind sample, DSNTIJCL, is available in DSN610.SDSNSAMP. We strongly recommend that you use this bind sample as a guide for binding DBRMS to packages and binding an application plan.

Special considerations for the OS/390 UNIX environment are described in the following section:

- “Setting up OS/390 UNIX environment” on page 55

In OS/390 Release 6, the name OS/390 OpenEdition is replaced with OS/390 UNIX System Services or OS/390 UNIX. Throughout this book, we use the name OS/390 UNIX.

## Bind DBRMs to packages

This section explains how to bind DBRMs to packages. Use the online bind sample, DSN610.SDSNSAMP(DSNTIJCL), for guidance.

For an application to access a data source using DB2 ODBC, the following IBM supplied DBRMs (shipped in DSN610.SDSNDBRM) must be bound to all data sources, including the local DB2 for OS/390 subsystem and all remote (DRDA) data sources.

- **DSNCLICS** bound with ISOLATION(CS)
- **DSNCLIRR** bound with ISOLATION(RR)
- **DSNCLIRS** bound with ISOLATION(RS)
- **DSNCLIUR** bound with ISOLATION(UR)
- **DSNCLINC** bound with ISOLATION(NC)
- **DSNCLIC1** bound with default options
- **DSNCLIC2** bound with default options
- **DSNCLIMS** bound with default options
- **DSNCLIVM** (for SQL/DS™ access only) bound with default options
- **DSNCLIAS** (for OS/400® access only) bound with default options
- **DSNCLIV1** (for DB2 UDB Version 1 access only) bound with default options
- **DSNCLIV2** (for DB2 UDB Version 2 access only) bound with default options

- **DSNCLIQR** bound to any site that supports DRDA query result sets.
- **DSNCLIF4** bound with default options

### Package bind options

For packages listed above that use the ISOLATION keyword, the impact of package bind options in conjunction with the DB2 ODBC initialization file keywords is as follows:

- ISOLATION

Packages must be bound with the isolation specified.

- DYNAMICRULES(BIND)

Binding the packages with this option offers encapsulation and security similar to that of static SQL. The recommendations and consequences for using this option are as follows:

1. Bind DB2 ODBC packages or plan with DYNAMICRULES(BIND) from a 'driver' authorization ID with table privileges.
2. Issue GRANT EXECUTE on each collection.package or plan name to individual users. Plans are differentiated by plan name; packages are differentiated by collection.
3. Select a plan or package by using the PLANNAME or COLLECTIONID keywords in the DB2 ODBC initialization file.
4. When dynamic SQL is issued, the statement is processed with the 'driver' authorization ID. Users need execute privileges; table privileges are not required.
5. The CURRENTSQLID keyword cannot be used in the DB2 ODBC initialization file. Use of this keyword results in an error at SQLConnect().

- SQLERROR(CONTINUE)

```
# Use this keyword to bind DB2 ODBC to Version 5 of DB2 for OS/390. The
# symptoms of binding to a downlevel server are:
#
#   - Binding DSNCLIMS results in SQLCODE -199 on the VALUES INTO
#     statement. Bind with the SQLERROR(CONTINUE) keyword to avoid this
#     error.
#
#   - Binding DSNCLIMS results in SQLCODE -199 on the DESCRIBE INPUT
#     statement. Apply APAR PQ24584 and retry the bind to avoid this error.
#     Alternatively, you can bind with the SQLERROR(CONTINUE) keyword,
#     however, the SQLDescribeParam() API will be unavailable to you at that
#     DB2 for OS/390 Version 5 server.
```

**Bind return codes:** A bind to DB2 for OS/390 receives several expected warnings:

- For all packages:  
WARNING, ONLY IBM-SUPPLIED COLLECTION-IDS SHOULD BEGIN WITH "DSN"
- For bind of DSNCLINC package to DB2 for OS/390:  
BIND WARNING - ISOLATION NC NOT SUPPORTED CHANGED TO ISOLATION UR

- For bind of DSNCLIF4 package to DB2 for OS/390 for SYSIBM.LOCATIONS naming differences between releases. For example when bound to a Version 5 system:

```
SYSIBM.SYSLOCATIONS IS NOT DEFINED
```

## # Bind packages at remote sites

```
# For an application to access a data source using DB2 ODBC, bind the DBRMs
# listed above to all data sources, including the local DB2 for OS/390 subsystem and
# all remote (DRDA) data sources. The SQLConnect() argument szDSN identifies the
# data source. The data source is the location in the DB2 SYSIBM.LOCATION
# catalog table. An application running under DB2 ODBC to a remote DB2 for
# OS/390, or another DBMS, does not need to be bound into the DB2 ODBC plan;
# rather it can be bound as a package at the remote site. Failure to bind the package
# at the remote site results in SQLCODE -805.
```

## Bind stored procedures

A stored procedure running under DB2 ODBC to a remote DB2 for OS/390, or another DBMS, does not need to be bound into the DB2 ODBC plan; rather it can be bound as a package at the remote site.

For a stored procedure that resides on the local DB2 for OS/390, the stored procedure package must be bound in the DB2 ODBC plan, using PKLIST. Stored procedures on remote servers only need to bind to that remote server as a package.

For example, DB2 ODBC must always be bound in a plan to a DB2 for OS/390 subsystem to which DB2 ODBC first establishes an affinity on the `SQLAllocEnv()` call. This is the local DB2. The scenario in this example is equivalent to specifying the `MVSDEFAULTSSID` keyword in the initialization file. If DB2 ODBC calls a stored procedure that resides at this local DB2 for OS/390, that stored procedure package must be in the DB2 ODBC plan, using PKLIST.

```
# This process is unique to DB2 for OS/390 stored procedure support. For more
# information about using stored procedures, see "Using stored procedures" on
# page 389.
```

```
# Include local, remote, and stored procedure packages in the PKLIST of the plan at
# the site where the client will execute.
```

## Bind an application plan

This section explains how to bind an application plan. Use the online bind sample, `DSN610.SDSNSAMP(DSNTIJCL)`, for guidance.

A DB2 plan must be created using the PKLIST keyword to name all packages listed in "Bind DBRMs to packages" on page 52. Any name can be selected for the plan; the default name is **DSNACLI**. If a name other than the default is selected, that name must be specified within the initialization file by using the `PLANNAME` keyword.

## Plan bind options

Use PLAN bind options as follows:

- DISCONNECT(EXPLICIT)

All DB2 ODBC plans are created using this option. DISCONNECT(EXPLICIT) is the default value; do not change it.

- CURRENTSERVER

Do not specify this keyword when binding plans.

## Setting up OS/390 UNIX environment

To use DB2 ODBC in the OS/390 UNIX environment, the DB2 ODBC definition side-deck must be available to OS/390 UNIX users.

The OS/390 UNIX compiler determines the contents of an input file based on the file extension. In the case of a file residing in an MVS partitioned data set (PDS), the last qualifier in the PDS name is treated as the file extension.

The OS/390 UNIX compiler recognizes the DB2 ODBC definition side-deck by these criteria:

- It must reside in an MVS PDS
- The last qualifier in the PDS name must be .EXP

Therefore, to make the DB2 ODBC definition side-deck available to OS/390 UNIX users, you should define an MVS data set alias that uses .EXP as the last qualifier in the name. This alias should relate to the SDSNMACS data set which is where the DB2 ODBC definition side-deck is installed.

For example, assume that DB2 is installed using DSN610 as the high level data set qualifier. You can define the alias using the following command:

```
DEFINE ALIAS(NAME('DSN610.SDSNC.EXP') RELATE('DSN610.SDSNMACS'))
```

This alias allows OS/390 UNIX users to directly reference the DB2 ODBC definition side-deck by specifying:

```
"//'DSN610.SDSNC.EXP(DSNAOCLI)'"
```

as one of the input files to the OS/390 UNIX c89 command.

As an alternative to defining a system alias for the ODBC side-deck, use the `_XSUFFIX_HOST` environmental variable that specifies the MVS data set suffix. The default value is EXP. For example, changing the default from EXP to SDSNMACS allows the link to work without a Define Alias.

For the c89 compiler issue:

```
export _C89_XSUFFIX_HOST="SDSNMACS"
```

For thecxx compiler issue:

```
export _CXX_XSUFFIX_HOST="SDSNMACS"
```

## Preparing and executing a DB2 ODBC application

This section provides an overview of the DB2 ODBC components and explains the steps you follow to prepare and execute a DB2 ODBC application.

Figure 5 shows the DB2 ODBC components used to build the DB2 ODBC DLL, and the process you follow to install and prepare a DB2 ODBC application. The shaded areas identify the components that are shipped.

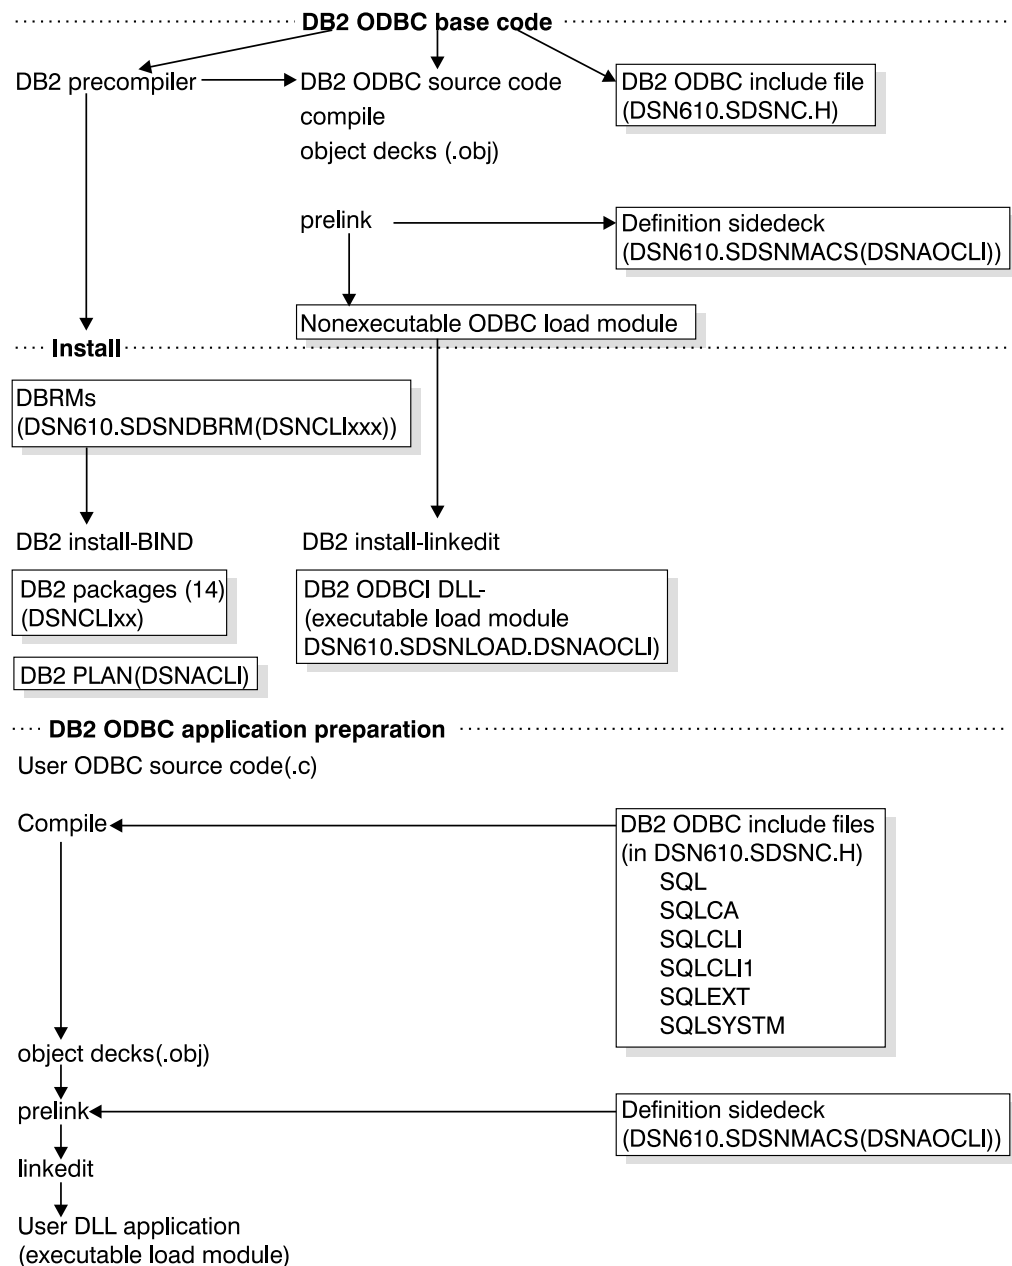


Figure 5. DB2 ODBC application development and execution

The following sections describe the requirements and steps that are necessary to run a DB2 ODBC application.

- “DB2 ODBC application requirements”
- “Application preparation and execution steps”

## DB2 ODBC application requirements

To successfully build a DLL application, you must ensure that the correct compile, pre-link, and link-edit options are used. In particular, your application must generate the appropriate DLL linkage for the exported DB2 ODBC DLL functions.

The C++ compiler always generates DLL linkage. However, the C compiler only generates DLL linkage if the DLL compile option is used. Failure to generate the necessary DLL linkage can cause the prelinker and linkage editor to issue warning messages for unresolved references to DB2 ODBC functions.

The minimum requirements for a DB2 ODBC application are as follows:

- OS/390 Version 1 Release 3 Application Enablement optional feature for C/C++.  
If the C compiler is used, then the DLL compiler option must be specified.
- OS/390 Release 3 Language Environment Application Enablement base feature.
- The DB2 ODBC application must be written and link-edited to execute with a 31-bit addressing mode, AMODE(31).

### Special considerations for OS/390 UNIX

A special consideration applies to DB2 ODBC product data set access. If you build a DB2 ODBC application in OS/390 UNIX, you can use the c89 compile command to compile your application. Although you compile your application under OS/390 UNIX, you can directly reference the non-HFS DB2 ODBC data sets in the c89 command. There is no need to copy the DB2 ODBC product files to HFS.

## Application preparation and execution steps

The following steps describe application preparation and execution:

- “Step 1. Compile the application” on page 58
- “Step 2. Pre-link and link-edit the application” on page 58
- “Step 3. Execute the application” on page 59

DB2 ODBC provides two online samples for installation verification:

### DSN8OIVP

A sample C application. You can use this sample to verify that your DB2 ODBC installation is correct.

### DSNTEJ8

Sample JCL. You can use this sample to compile, pre-link, link-edit, and execute the sample application DSN8OIVP.

The DSNTEJ8 and DSN8OIVP online samples are available in DSN610.SDSNSAMP. We strongly recommend that you use these samples for guidance when running an application.

**Using the ODBC sample in OS/390 UNIX:** To use the ODBC sample DSN8OIVP in OS/390 UNIX, copy DSN8OIVP from the sample data set to HFS. user/db2 is considered the user's directory. For example:

```
oput 'dsn610.sdsnsamp(dsn8oivp)' '/usr/db2/dsn8oivp.c' TEXT
```

## Step 1. Compile the application

Include the following statement in your DB2 ODBC application:

```
#include <sqlcli1.h>
```

The sqlcli1.h file includes all information that is required for compiling your DB2 ODBC application. All DB2 ODBC header files, including sqlcli1.h, that define the function prototypes, constants, and data structures that are needed for a DB2 ODBC application are shipped in the DSN610.SDSNC.H data set. Therefore, you must add this dataset to your SYSPATH concatenation when you compile your DB2 ODBC application.

For an example of a compile job, use the DSNTTEJ8 online sample in DSN610.SDSNSAMP.

**Compiling in OS/390 UNIX:** If you build a DB2 ODBC application in OS/390 UNIX, you can use the c89 command to compile your application. For example, to compile a C application named 'dsn8oivp.c' that resides in the current working directory, the c89 compile command might look like:

```
c89 -c -W 'c,dll,long,source,list' -  
      -I"//DSN610.SDSNC.H" \  
      dsn8oivp.c
```

Alternatively, if you write an application in C++, the cxx command might look like:

```
cxx -c -W 'c,long,source,list' -  
      -I"//DSN610.SDSNC.H" \  
      dsn8oivp.C
```

If your source code is in C, rather than C++, you must compile using the 'dll' option to enable use of the DB2 ODBC driver. This is a requirement even when using the cxx compile command to compile C parts.

## Step 2. Pre-link and link-edit the application

Before you can link-edit your DB2 ODBC application, you must pre-link your application with the DB2 ODBC definition side-deck provided with Version 6 of DB2 for OS/390.

The definition side-deck defines all of the exported functions in the DB2 ODBC dynamic load library, DSNAOCLI. It resides in the DSN610.SDSNMACS data set, as member DSNAOCLI. The definition side-deck should also be available under the alias data set name of DSN610.SDSNC.EXP as member DSNAOCLI (see "Setting up OS/390 UNIX environment" on page 55 for details). You must include the DSNAOCLI member as input to the Prelinker by specifying it in the pre-link SYSIN DD card concatenation.

For an example of pre-link and link-edit jobs, use the DSNTTEJ8 online sample in DSN610.SDSNSAMP.

For more information about DLL, see *OS/390 C/C++ Programming Guide*



**Pre-linking and link-editing in OS/390 UNIX:** If you build a DB2 ODBC application in OS/390 UNIX, you can use the `c89` command to pre-link and link-edit your application. You need to include the DB2 ODBC definition side-deck as one of the input data sets to the `c89` command and specify `'d11'` as one of the link-edit options.

For example, assume that you have already compiled a C application named `'myapp.c'` to create a `'myapp.o'` file in the current working directory. The `c89` command to pre-link and link-edit your application might look like:

```
c89 -W l,p,map,noer -W l,d11,AMODE=31,map \  
    -o dsn8oivp dsn8oivp.o "'DSN610.SDSNC.EXP(DSNAOCLI)'"
```

The following command references the ODBC side-deck as a C89 `_XSUFFIX_HOST` environmental variable (as described in “Setting up OS/390 UNIX environment” on page 55):

```
c89 -W l,p,map,noer -W l,d11,AMODE=31,map -o dsn8oivp dsn8oivp.o  
"'DSN610.SDSNMACS(dsnaocli)'"
```

### Step 3. Execute the application

DB2 ODBC applications must access the `DSN610.SDSNLOAD` data set at execution time. The `SDSNLOAD` data set contains both the DB2 ODBC dynamic load library and the attachment facility used to communicate with DB2.

In addition, the DB2 ODBC driver accesses the DB2 for OS/390 load module `DSNHDECP`. `DSNHDECP` contains, among other things, the coded character set ID (CCSID) information that DB2 for OS/390 uses.

A default `DSNHDECP` is shipped with DB2 for OS/390 in the `DSN610.SDSNLOAD` data set. However, if the values provided in the default `DSNHDECP` are not appropriate for your site, a new `DSNHDECP` can be created during the installation of DB2 for OS/390. If a site specific `DSNHDECP` is created during installation, you should concatenate the data set containing the new `DSNHDECP` before the `DSN610.SDSNLOAD` data set in your `STEPLIB` or `JOBLIB` DD card.

For an example of an execute job, use the `DSNTEJ8` online sample in `DSN610.SDSNSAMP`.

**Executing in OS/390 UNIX:** To execute a DB2 ODBC application in OS/390 UNIX, you need to include the `DSN610.SDSNEXIT` and `DSN610.SDSNLOAD` data sets in the data set concatenation of your `STEPLIB` environmental variable. The `STEPLIB` environmental variable can be set in your `.profile` with the statement:

```
export STEPLIB=DSN610.SDSNEXIT:DSN610.SDSNLOAD
```

### Defining a subsystem

There are two ways to define a DB2 subsystem to DB2 ODBC. You can identify the DB2 subsystem by specifying the `MVSDEFAULTSSID` keyword in the common section of initialization file. If the `MVSDEFAULTSSID` keyword does not exist in the initialization file, DB2 ODBC uses the default subsystem name specified in the `DSNHDECP` load module that was created when DB2 was installed. Therefore, you should ensure that DB2 ODBC can find the intended `DSNHDECP` when your application issues the `SQLAL1locEnv` call.

The DSNHDECP load module is usually link-edited into the DSN610.SDSNEXIT data set. In this case, your STEPLIB DD card includes:

```
//STEPLIB DD DSN=DSN610.SDSNEXIT,DISP=SHR
// DD DSN=DSN610.SDSNLOAD,DISP=SHR
...
```

---

## DB2 ODBC initialization file

A set of optional keywords can be specified in a DB2 ODBC *initialization file*, an EBCDIC file that stores default values for various DB2 ODBC configuration options. Because the initialization file has EBCDIC text, it can be updated using a file editor, such as the TSO editor.

For most applications, use of the DB2 ODBC initialization file is not necessary. However, to make better use of IBM RDBMS features, the keywords can be specified to:

- Help improve the performance or usability of an application.
- Provide support for applications written for a previous version of DB2 ODBC.
- Provide specific work-arounds for existing ODBC applications.

The following sections describe how to create the initialization file and define the keywords:

- “Using the initialization file”
- “Initialization keywords” on page 63

## Using the initialization file

The DB2 ODBC initialization file is read at application run time. The file can be specified by either a DSNA0INI DD card or by defining a DSNA0INI OS/390 UNIX environmental variable. DB2 ODBC opens the DSNAOINI data set allocated in your JCL first. If a DSNAOINI data set is not allocated, then DB2 ODBC opens the environmental variable data set.

The initialization file specified can be either a tradition MVS data set or an OS/390 UNIX HFS file. For MVS data sets, the record format of the initialization file can be either fixed or variable length.

The following JCL examples use a DSNA0INI JCL DD card to specify the DB2 ODBC initialization file types supported:

MVS sequential data set USER1.DB2ODBC.ODBCINI:

```
//DSNA0INI DD DSN=USER1.DB2ODBC.ODBCINI,DISP=SHR
```

MVS partitioned data set USER1.DB2ODBC.DATA, member ODBCINI:

```
//DSNA0INI DD DSN=USER1.DB2ODBC.DATA(ODBCINI),DISP=SHR
```

Inline JCL DSNA0INI DD specification:

```
//DSNA0INI DD *
  [COMMON]
  MVSDEFAULTSSID=V61A
/*
```

HFS file /u/user1/db2odbc/odbcini:

```
//DSNA0INI DD PATH='/u/user1/db2odbc/odbcini'
```

The following examples of OS/390 UNIX export statements define the DB2 ODBC DSNAOINI OS/390 UNIX environmental variable for the DB2 ODBC initialization file types supported:

HFS fully qualified file /u/user1/db2odbc/odbcini:

```
export DSNAOINI="/u/user1/db2odbc/odbcini"
```

HFS file ./db2odbc/odbcini, relative to the present working directory of the application:

```
export DSNAOINI="./db2odbc/odbcini"
```

MVS sequential data set USER1.ODBCINI:

```
export DSNAOINI="USER1.ODBCINI"
```

Redirecting to use a file specified by another DD card, MYDD, that is already allocated:

```
export DSNAOINI="//DD:MYDD"
```

MVS partitioned data set USER1.DB2ODBC.DATA, member ODBCINI:

```
export DSNAOINI="USER1.DB2ODBC.DATA(ODBCINI)"
```

When specifying an HFS file, the value of the DSNAOINI environmental variable must begin with either a single forward slash (/), or a period followed by a single forward slash (.). If a setting starts with any other characters, DB2 ODBC assumes that an MVS data set name is specified.

Allocation precedence: DB2 ODBC opens the DSNAOINI data set allocated in your JCL first. If a DSNAOINI data set is not allocated, then DB2 ODBC opens the environmental variable data set.

## Initialization file structure

The initialization file consists of the following three sections, or stanzas:

**Common section** Contains parameters that are global to all applications using this initialization file.

**Subsystem section** Contains parameter values unique to that subsystem.

### Data source sections

Contain parameter values to be used only when connected to that data source. You can specify zero or more data source sections.

Each section is identified by a syntactic identifier enclosed in square brackets. Specific guidelines for coding square brackets are described in the list item below marked 'Attention'.

The syntactic identifier is either the literal 'common', the subsystem ID or the data source (location name). For example:

```
[data-source-name]
```

This is the *section header*.

The parameters are set by specifying a keyword with its associated keyword value in the form:

**KeywordName** =*keywordValue*

- All the keywords and their associated values for each data source must be located below the data source section header.
- The keyword settings in each section apply only to the data source name in that section header.
- The keywords are **not** case sensitive; however, their values can be if the values are character based.
- For the syntax associated with each keyword, see “Initialization keywords” on page 63.
- If a data source name is not found in the DB2 ODBC initialization file, the default values for these keywords are in effect.
- Comment lines are introduced by having a semi-colon in the first position of a new line.
- Blank lines are also permitted. If duplicate entries for a keyword exist, the first entry is used (and no warning is given).
- **Attention:** You can avoid common errors by ensuring that the following contents of the initialization file are accurate:
  - Square brackets: The square brackets in the initialization file must consist of the correct EBCDIC characters. The open square bracket must use the hexadecimal characters X'AD'. The close square bracket must use the hexadecimal characters X'BD'. DB2 ODBC does not recognize brackets if coded differently.
  - Sequence numbers: The initialization file cannot accept sequence numbers. All sequence numbers must be removed.

The following is a sample DB2 ODBC initialization file with a common stanza, a subsystem stanza, and two data source stanzas.

```
; This is a comment line...
; Example COMMON stanza
[COMMON]
MVSDEFAULTSSID=V61A

; Example SUBSYSTEM stanza for V61A subsystem
[V61A]
MVSATTACHTYPE=CAF
PLANNNAME=DSNACLI

; Example DATA SOURCE stanza for STLEC1 data source
[STLEC1]
AUTOCOMMIT=0
CONNECTTYPE=2

; Example DATA SOURCE stanza for STLEC1B data source
[STLEC1B]
CONNECTTYPE=2
CURSORHOLD=0
```

## Initialization keywords

The initialization keywords are described in this section. The section (common, subsystem, or data source) in which each keyword must be defined is identified.

### **APPLTRACE = 0 | 1**

This keyword is placed in the common section.

The APPLTRACE keyword controls whether the DB2 ODBC application trace is enabled. The application trace is designed for diagnosis of application errors. If enabled, every call to any DB2 ODBC API from the application is traced, including input parameters. The trace is written to the file specified on the APPLTRACEFILENAME keyword.

- 0 = Disabled (default)
- 1 = Enabled

For more information about using the APPLTRACE keyword, see “Application trace” on page 413.

**Important:** This keyword is renamed. DB2 ignores the Version 5 keyword name CLITRACE.

### **APPLTRACEFILENAME = *dataset name***

This keyword is placed in the common section.

APPLTRACEFILENAME is only used if a trace is started by the APPLTRACE keyword. When APPLTRACE is set to 1, use the APPLTRACEFILENAME keyword to identify an MVS data set name or OS/390 UNIX HFS file name that records the DB2 ODBC application trace. “Diagnostic trace” on page 415 provides detailed information about specifying file name formats.

**Important:** This keyword is renamed. DB2 ignores the Version 5 keyword name TRACEFILENAME.

### **AUTOCOMMIT = 1 | 0**

This keyword is placed in the data source section.

To be consistent with ODBC, DB2 ODBC defaults with AUTOCOMMIT on, which means each statement is treated as a single, complete transaction. This keyword can provide an alternative default, but is only used if the application does not specify a value for AUTOCOMMIT as part of the program.

- 1 = on (default)
- 0 = off

Most ODBC applications assume the default of AUTOCOMMIT is on. Extreme care must be used when overriding this default during runtime as the application might depend on this default to operate properly.

This keyword also allows you to specify whether autocommit should be enabled in a distributed unit of work (DUW) environment. If a connection is part of a coordinated DUW, and AUTOCOMMIT is not set, the default does not apply; implicit commits arising from autocommit processing are suppressed. If AUTOCOMMIT is set to 1, and the connection is part of a coordinated DUW, the implicit commits are processed. This can result in severe performance degradations, and possibly other unexpected results elsewhere in the DUW system. However, some applications might not work at all unless this is enabled.

A thorough understanding of the transaction processing of an application is necessary, especially applications written by a third party, before applying it to a DUW environment.

**BITDATA = 1 | 0**

This keyword is placed in the data source section.

The BITDATA keyword allows you to specify whether ODBC binary data types, SQL\_BINARY, SQL\_VARBINARY, and SQL\_LONGVARBINARY, and SQL\_BLOB are reported as binary type data. IBM DBMSs support columns with binary data types by defining CHAR, VARCHAR and LONG VARCHAR columns with the FOR BIT DATA attribute.

Only set BITDATA = 0 if you are sure that all columns defined as FOR BIT DATA or BLOB contain only character data, and the application is incapable of displaying binary data columns.

1 = Report FOR BIT DATA and BLOB data types as binary data types. This is the default.

0 = Disabled.

**CLISHEMA = *schema\_name***

This keyword is placed in the data source section.

The CLISHEMA keyword lets you indicate the schema of the DB2 ODBC shadow catalog tables or views to search when you issue an ODBC catalog function call. For example, if you specify CLISHEMA=PAYROLL, the ODBC catalog functions that normally reference the DB2 system catalog tables (SYSIBM schema), will reference the following views of the DB2 ODBC shadow catalog tables:

- PAYROLL.COLUMNS
- PAYROLL.TABLES
- PAYROLL.COLUMNPRIVILIGES
- PAYROLL.TABLEPRIVILIGES
- PAYROLL.SPECIALCOLUMNS
- PAYROLL.PRIMARYKEYS
- PAYROLL.FOREIGNKEYS
- PAYROLL.TSTATISTICS
- PAYROLL.PROCEDURES

You must build the DB2 ODBC shadow catalog tables and optional views before using the CLISHEMA keyword. If this keyword is not specified, the ODBC catalog query APIs reference the DB2 (SYSIBM) system tables by default.

**COLLECTIONID = *collection\_id***

This keyword is placed in the data source section.

The COLLECTIONID keyword allows you to specify the collection identifier that is used to resolve the name of the package allocated at the server. This package supports the execution of subsequent SQL statements.

The value is a character string and must not exceed 18 characters. It can be overridden by executing the SET CURRENT PACKAGESET statement.

**CONNECTTYPE = 1 | 2**

This keyword is placed in the common section.

The CONNECTTYPE keyword allows you to specify the default connect type for all connections to data sources.

- 1 = Multiple concurrent connections, each with its own commit scope. If MULTICONTEXT=0 is specified, a new connection might not be added unless the current transaction on the current connection is on a transaction boundary (either committed or rolled back). This is the default.
- 2 = Coordinated connections where multiple data sources participate under the same distributed unit of work. CONNECTTYPE=2 is ignored if MULTICONTEXT=1 is specified.

**CURRENTFUNCTIONPATH = ""schema1', 'schema2' ,..."**

This keyword is placed in the data source section.

The CURRENTFUNCTIONPATH keyword defines the path used to resolve unqualified user-defined functions, distinct types, and stored procedure references that are used in dynamic SQL statements. It contains a list of one or more schema names, which are used to set the CURRENT PATH special register using the SET CURRENT PATH SQL statement upon connection to the data source. Each schema name in the keyword string must be delimited with single quotes and separated by commas. The entire keyword string must be enclosed in double quotes and must not exceed 254 characters.

The default value of the CURRENT PATH special register is:

"SYSIBM", "SYSFUN", "SYSPROC", X

where X is the value of the USER special register as a delimited identifier. The schemas SYSIBM, SYSFUN, and SYSPROC do not need to be specified. If any of these schemas is not included in the current path, then it is implicitly assumed at the beginning of the path in the order shown above. The order of the schema names in the path determines the order in which the names are resolved. For more detailed information on schema name resolution, see *DB2 SQL Reference*.

Unqualified user-defined functions, distinct types, and stored procedures are searched from the list of schemas specified in the CURRENTFUNCTIONPATH setting in the order specified. If the user-defined function, distinct type, or stored procedure is not found in a specified schema, the search continues in the schema specified next in the list. For example:

```
CURRENTFUNCTIONPATH=" 'USER01', 'PAYROLL', 'SYSIBM', 'SYSFUN', 'SYSPROC' "
```

results in searching schema "USER01", followed by schema "PAYROLL", followed by schema "SYSIBM", etc..

Although the SQL statement CALL is a static statement, the CURRENTFUNCTIONPATH setting affects a CALL statement if the stored procedure name is specified with a host variable (making the CALL statement a pseudo-dynamic SQL statement). This is always the case for a CALL statement processed by DB2 ODBC.

**CURRENTSQLID = current\_sqlid**

This keyword is placed in the data source section.

The CURRENTSQLID keyword is valid only for those DB2 DBMSs that support SET CURRENT SQLID (such as DB2 for OS/390). If this keyword is present, then a SET CURRENT SQLID statement is sent to the DBMS after a successful connect. This allows the end user and the application to name SQL objects without having to qualify by schema name.

Do not specify this keyword if you are binding the DB2 ODBC packages with DYNAMICRULES(BIND).

**CURSORHOLD = 1 | 0**

This keyword is placed in the data source section.

The CURSORHOLD keyword controls the effect of a transaction completion on open cursors.

1 = Cursor hold. The cursors are not destroyed when the transaction is committed. This is the default.

0 = Cursor no hold. The cursors are destroyed when the transaction is committed.

Cursors are always destroyed when transactions are rolled back.

This keyword can be used by an end user to improve performance. If the user is sure that the application:

1. Does not have behavior that is dependent on the SQL\_CURSOR\_COMMIT\_BEHAVIOR or the SQL\_CURSOR\_ROLLBACK\_BEHAVIOR information returned using SQLGetInfo(), and
2. Does not require cursors to be preserved from one transaction to the next,

then the value of this keyword can be set to **0**. The DBMS operates more efficiently as resources no longer need to be maintained after the end of a transaction.

**DBNAME = dbname**

This keyword is placed in the data source section.

The DBNAME keyword is only used when connecting to DB2 for OS/390, and only if (*base*) table catalog information is requested by the application.

If a large number of tables exist in the DB2 for OS/390 subsystem, a *dbname* can be specified to reduce the time it takes for the database to process the catalog query for table information, and reduce the number of tables returned to the application.

The value of the *dbname* keyword maps to the DBNAME column in the DB2 for OS/390 system catalog tables. If no value is specified, or if views, synonyms, system tables, or aliases are also specified using TABLETYPE, only table information is restricted; views, aliases, and synonyms are not restricted with DBNAME. This keyword can be used in conjunction with SCHEMALIST and TABLETYPE to further limit the number of tables for which information is returned.

**DIAGTRACE = 0 | 1**

This keyword is placed in the common section.

The DIAGTRACE keyword lets you enable the DB2 ODBC diagnostic trace.

0 = The DB2 ODBC diagnostic trace is not enabled. No diagnostic data is captured. This is the default.

You can enable the diagnostic trace using the DSNAOTRC command when the DIAGTRACE keyword is set to 0.

1 = The DB2 ODBC diagnostic trace is enabled. Diagnostic data is recorded in the application address space. If you include a DSNAOTRC DD statement in your job or TSO logon procedure that identifies an MVS data set or an OS/390 UNIX HFS file name, the trace is externalized at normal



program termination. You can format the trace using the DSNAOTRC trace formatting program.

For more information about using the DIAGTRACE keyword and the DSNAOTRC command, see “Diagnostic trace” on page 415.

**Important:** This keyword is renamed. DB2 ignores the Version 5 keyword name TRACE.

#### **DIAGTRACE\_BUFFER\_SIZE = *buffer size***

This keyword is placed in the common section.

The DIAGTRACE\_BUFFER\_SIZE keyword controls the size of the DB2 ODBC diagnostic trace buffer. This keyword is only used if a trace is started by using the DIAGTRACE keyword.

*buffer size* is an integer value that represents the number of bytes to allocate for the trace buffer. The buffer size is rounded down to a multiple of 65536 (64K). If the value specified is less than 65536, then 65536 is used. The default value for the trace buffer size is 65536.

If a trace is already active, this keyword is ignored.

**Important:** DB2 ignores the Version 5 keyword name TRACE\_BUFFER\_SIZE.

#### **DIAGTRACE\_NO\_WRAP = 0 | 1**

This keyword is placed in the common section.

The DIAGTRACE\_NO\_WRAP keyword controls the behavior of the DB2 ODBC diagnostic trace when the DB2 ODBC diagnostic trace buffer fills up. This keyword is only used if a trace is started by the DIAGTRACE keyword.

0 = The trace table is a wrap-around trace. In this case, the trace remains active to capture the most current trace records. This is the default.

1 = The trace stops capturing records when the trace buffer fills. The trace captures the initial trace records that were written.

If a trace is already active, this keyword is ignored.

**Important:** This keyword is renamed. DB2 ignores the Version 5 keyword name TRACE\_NO\_WRAP.

#### **GRAPHIC = 0 | 1 | 2 | 3**

This keyword is placed in the data source section.

The GRAPHIC keyword controls whether DB2 ODBC reports IBM GRAPHIC (double byte character support) as one of the supported data types when SQLGetTypeInfo() is called. SQLGetTypeInfo() lists the data types supported by the data source for the current connection. These are not native ODBC types but have been added to expose these types to an application connected to a DB2 family product.

0 = disabled (default)

1 = enabled

2 = report the length of graphic columns returned by DESCRIBE in number of bytes rather than DBCS characters. This applies to all DB2 ODBC and ODBC functions that return length or precision either on the output argument or as part of the result set.

3 = settings 1 and 2 combined; that is, GRAPHIC=3 achieves the combined effect of 1 and 2.

The default is that GRAPHIC is not returned since many applications do not recognize this data type and cannot provide proper handling.

#### **MAXCONN = 0 | positive number**

This keyword is placed in the common section.

The MAXCONN keyword is used to specify the maximum number of connections allowed for each DB2 ODBC application program. This can be used by an administrator as a governor for the maximum number of connections established by each application.

A value of 0 can be used to represent *no limit*; that is, an application is allowed to open up as many connections as permitted by the system resources. This is the default.

This parameter limits the number of SQLConnect() statements that the application can successfully issue. In addition, if the application is executing with CONNECT (type 1) semantics, then this value specifies the number of logical connections. There is only one physical connection to either the local DB2/MVS subsystem or a remote DB2 subsystem or remote DRDA-1 or DRDA-2 server.

#### **MULTICONTEXT = 0 | 1**

This keyword is placed in the common section.

The MULTICONTEXT keyword controls whether each connection in an application can be treated as a separate unit of work with its own commit scope that is independent of other connections.

0 = The DB2 ODBC code does not create an independent *context* for a data source connection. Connection switching among multiple data sources governed by the CONNECTTYPE=1 rules is not allowed unless the current transaction on the current connection is on a transaction boundary (either committed or rolled back). This is the default.

1 = The DB2 ODBC code creates an independent context for a data source connection at the connection handle level when SQLAllocConnect() is issued. Each connection to multiple data sources is governed by CONNECTTYPE=1 rules and is associated with an independent DB2 thread. Connection switching among multiple data sources is not prevented due to the commit status of the transaction; an application can use multiple connection handles without having to perform a commit or rollback on a connection before switching to another connection handle. The use of MULTICONTEXT=1 requires MVSATTACHTYPE=RRSAF and OS/390 Version 2 Release 5 or higher.

The application can use SQLGetInfo() with *infoType*=SQL\_MULTIPLE\_ACTIVE\_TXN to determine whether MULTICONTEXT=1 is supported.

MULTICONTEXT=1 is ignored if any of these conditions are true:

- The application created a DB2 thread before invoking DB2 ODBC. This is always the case for a stored procedure using DB2 ODBC.
- The application created and switched to a private context using OS/390 Context Services before invoking DB2 ODBC.

- The application started a unit of recovery with any RRS resource manager (for example, IMS) before invoking DB2 ODBC.
- MVSATTACHTYPE=CAF is specified in the initialization file.
- The OS/390 operating system level does not support Unauthorized Context Services.

**MVSATTACHTYPE = CAF | RRSAF**

This keyword is placed in the subsystem section.

The MVSATTACHTYPE keyword is used to specify the DB2 for OS/390 attachment type that DB2 ODBC uses to connect to the DB2 for OS/390 address space. This parameter is ignored if the DB2 ODBC application is running as a DB2 for OS/390 stored procedure. In that case, DB2 ODBC uses the attachment type that was defined for the stored procedure.

CAF: DB2 ODBC uses the DB2 for OS/390 call attachment facility (CAF). This is the default.

RRSAF: DB2 ODBC uses the DB2 for OS/390 Recoverable Resource Manager Services attachment facility (RRSAF).

**MVSDEFAULTSSID = ssid**

This keyword is placed in the common section.

The MVSDEFAULTSSID keyword specifies the default DB2 subsystem to which the application is connected when invoking the `SQLAllocConnect` function. Specify the DB2 subsystem name or group attachment name (if used in a data sharing group) to which connections will be made. The default subsystem is 'DSN'.

**OPTIMIZEFORNROWS = integer**

This keyword is placed in the data source section.

The OPTIMIZEFORNROWS keyword appends the "OPTIMIZE FOR n ROWS" clause to every select statement, where n is an integer larger than 0. The default action is not to append this clause.

For more information on the effect of the OPTIMIZE FOR n ROWS clause, see *DB2 SQL Reference*.

**PLANNAME = planname**

This keyword is placed in the subsystem section.

The PLANNAME keyword specifies the name of the DB2 for OS/390 PLAN that was created during installation. A PLAN name is required when initializing the application connection to the DB2 for OS/390 subsystem which occurs during the processing of the `SQLAllocEnv()` call.

If no PLANNAME is specified, the default value DSNACLI is used.

**SCHEMALIST = "'schema1', 'schema2' ,..."**

This keyword is placed in the data source section.

The SCHEMALIST keyword specifies a list of schemas in the data source.

If there are a large number of tables defined in the database, a schema list can be specified to reduce the time it takes for the application to query table information, and reduce the number of tables listed by the application. Each schema name is case sensitive, must be delimited with single quotes and separated by commas. The entire string must also be enclosed in double quotes, for example:

```
SCHEMALIST="'USER1', 'USER2', USER3'"
```

For DB2 for OS/390, CURRENT SQLID can also be included in this list, but without the single quotes, for example:

```
SCHEMALIST="'USER1', CURRENT SQLID, 'USER3'"
```

The maximum length of the keyword string is 256 characters.

This keyword can be used in conjunction with DBNAME and TABLETYPE to further limit the number of tables for which information is returned.

SCHEMALIST is used to provide a more restrictive default in the case of those applications that always give a list of every table in the DBMS. This improves performance of the table list retrieval in cases where the user is only interested in seeing the tables in a few schemas.

### **SYSSCHEMA = *sys schema***

This keyword is placed in the data source section. This keyword is placed in the data source section.

The SYSSCHEMA keyword indicates an alternative schema to be searched in place of the SYSIBM (or SYSTEM, QSYS2) schemas when the DB2 ODBC and ODBC catalog function calls are issued to obtain system catalog information.

Using this schema name, the system administrator can define a set of views consisting of a subset of the rows for each of the following system catalog tables:

- SYSCOLAUTH
- SYSCOLUMNS
- SYSDATABASE
- SYSFOREIGNKEYS
- SYSINDEXES
- SYSKEYS
- SYSPARMS
- SYSRELS
- SYSROUTINES
- SYSSYNONYMS
- SYSTABAUTH
- SYSTABLES

For example, if the set of views for the system catalog tables are in the ACME schema, then the view for SYSIBM.SYSTABLES is ACME.SYSTABLES; and SYSSCHEMA should then be set to ACME.

Defining and using limited views of the system catalog tables reduces the number of tables listed by the application, which reduces the time it takes for the application to query table information.

If no value is specified, the default is:

- SYSIBM on DB2 for OS/390 and OS/400
- SYSTEM on DB2 for VSE & VM
- QSYS2 on DB2 for AS/400®

This keyword can be used in conjunction with SCHEMALIST, TABLETYPE (and DBNAME on DB2 for OS/390) to further limit the number of tables for which information is returned.

**TABLETYPE=""TABLE' | , 'ALIAS' | , 'VIEW' | , ' | ,  
'SYSTEM TABLE' | , 'SYNONYM'"**

This keyword is placed in the data source section.

The TABLETYPE keyword specifies a list of one or more table types. If there are a large number of tables defined in the data source, a table type string can be specified to reduce the time it takes for the application to query table information, and reduce the number of tables listed by the application.

Any number of the values can be specified, but each type must be delimited with single quotes, separated by commas, and in upper case. The entire string must also be enclosed in double quotes, for example:

```
TABLETYPE=""TABLE' , 'VIEW'"
```

This keyword can be used in conjunction with DBNAME and SCHEMALIST to further limit the number of tables for which information is returned.

TABLETYPE is used to provide a default for the DB2 ODBC function that retrieves the list of tables, views, aliases, and synonyms in the data source. If the application does not specify a table type on the function call, and this keyword is not used, information about all table types is returned. If the application does supply a value for the *tabletype* on the function call, then that argument value overrides this keyword value.

If TABLETYPE includes any value other than TABLE, then the DBNAME keyword setting cannot be used to restrict information to a particular DB2 for OS/390 subsystem.

**THREADSAFE= 1 | 0**

This keyword is placed in the common section.

The THREADSAFE keyword controls whether DB2 ODBC uses *POSIX mutexes* to make the DB2 ODBC code *threadsafe* for multiple concurrent or parallel LE threads.

- 1 = The DB2 ODBC code is threadsafe if the application is executing in a POSIX(ON) environment. Multiple LE threads in the process can use DB2 ODBC. The threadsafe capability cannot be provided in a POSIX(OFF) environment. This is the default.
- 0 = The DB2 ODBC code is not threadsafe. This reduces the overhead of serialization code in DB2 ODBC for applications that are not *multithreaded*, but provides no protection for concurrent LE threads in applications that are multithreaded.

**TXNISOLATION = 1 | 2 | 4 | 8 | 32**

This keyword is placed in the data source section.

The TXNISOLATION keyword sets the isolation level to:

- 1 = Read uncommitted (uncommitted read)
- 2 = Read committed (cursor stability) (default)
- 4 = Repeatable read (read stability)
- 8 = Serializable (repeatable read)
- 32 = (No commit, DB2 for OS/400 only)

The words in round brackets are the DB2 equivalents for SQL92 isolation levels. Note that *no commit* is not an SQL92 isolation level and is supported only on DATABASE 2 for OS/400. See *DB2 Application Programming and SQL Guide* for more information on isolation levels.

## UNDERSCORE = 1 | 0

This keyword is placed in the data source section.

The UNDERSCORE keyword specifies whether the underscore character "\_" is to be used as a wildcard character (matching any one character, including no character), or to be used as itself. This parameter only affects catalog function calls that accept search pattern strings.

1 = "\_" acts as a wildcard (default)

The underscore is treated as a wildcard matching any one character or none. For example, if two tables are defined as follows:

```
CREATE TABLE "OWNER"."KEY_WORDS" (COL1 INT)
CREATE TABLE "OWNER"."KEYWORDS" (COL1 INT)
```

The DB2 ODBC catalog function call that returns table information (SQLTables()) returns both of these entries if "KEY\_WORDS" is specified in the table name search pattern argument.

0 = "\_" acts as itself

The underscore is treated as itself. If two tables are defined as shown in the example above, SQLTables() returns only the "KEY\_WORDS" entry if "KEY\_WORDS" is specified in the table name search pattern argument.

Setting this keyword to 0 can result in performance improvement in those cases where object names (owner, table, column) in the data source contain underscores.

---

## # DB2 ODBC migration considerations

# Changes to the columns of the result set for several DB2 ODBC catalog functions  
# are introduced with the DB2 for OS/390 Version 6 ODBC driver. These changes,  
# which include new and reordered columns, will affect existing applications that are  
# migrating from a DB2 for OS/390 Version 5 ODBC driver to a Version 6 ODBC  
# driver *and* are issuing the SQLColumns(), SQLForeignKeys(), or  
# SQLProcedureColumns() APIs. For a list of the columns returned for each of these  
# APIs, see:

- # • Table 30 on page 118
- # • Table 54 on page 176
- # • Table 106 on page 291

---

## Chapter 5. Functions

This section provides a description of each function. Each description has the following sections.

- Purpose
- Syntax
- Function arguments
- Usage
- Return codes
- Diagnostics
- Restrictions
- Example
- References

Each section is described below.

### Purpose

This section gives a brief overview of what the function does. It also indicates if any functions should be called before and after calling the function being described.

Each function also has a table, such as the one below that indicates which specification or standard the function conforms to. The first column indicates which version (1.0 or 2.0) of the ODBC specification the function was first provided. The second and third columns indicate if the function is included in the X/Open CLI CAE specification and the ISO CLI standard.

Table 9. Sample function specification table

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

This table indicates support of the function. Some functions use a set of options that do not apply to all specifications or standards. The restrictions section identifies any significant differences.

### Syntax

This section contains the generic 'C' prototype. If the function is defined by ODBC V2.0, then the prototype should be identical to that specified in *Microsoft ODBC 3.0 Programmer's Reference and SDK Guide*.

All function arguments that are pointers are defined using the FAR macro. This macro is defined out (set to a blank). This is consistent with the ODBC specification.

### Function arguments

This section lists each function argument, along with its data type, a description and whether it is an input or output argument.

Only SQLGetInfo() and SQLBindParameter() have parameters that are both input and output.

Some functions contain input or output arguments which are known as *deferred* or *bound* arguments. These arguments are pointers to buffers allocated by the application, and are associated with (or bound to) either a parameter in an SQL statement, or a column in a result set. The data areas specified by the function

are accessed by DB2 ODBC at a later time. It is important that these deferred data areas are still valid at the time DB2 ODBC accesses them.

### Usage

This section provides information about how to use the function, and any special considerations. Possible error conditions are not discussed here, but are listed in the diagnostics section instead.

### Return codes

This section lists all the possible function return codes. When `SQL_ERROR` or `SQL_SUCCESS_WITH_INFO` is returned, error information can be obtained by calling `SQLError()`.

See “Diagnostics” on page 34 for more information about return codes.

### Diagnostics

This section contains a table that lists the `SQLSTATE`s explicitly returned by DB2 ODBC (`SQLSTATE`s generated by the DBMS can also be returned) and indicates the cause of the error. These values are obtained by calling `SQLError()` after the function returns an `SQL_ERROR` or `SQL_SUCCESS_WITH_INFO`.

See “Diagnostics” on page 34 for more information about diagnostics.

### Restrictions

This section indicates any differences or limitations between DB2 ODBC and ODBC that can affect an application.

### Example

This section contains a code fragment that demonstrates the use of the function, using the generic data type definitions.

See “Chapter 4. Configuring DB2 ODBC and running sample applications” on page 49 for more information on setting up the DB2 ODBC environment and accessing the sample applications.

### References

This section lists related DB2 ODBC functions.

---

## Function summary

Table 10 provides a complete list of functions that DB2 ODBC and Microsoft ODBC 2.0 support. For each function, the table indicates whether ODBC or DB2 ODBC supports it, the ODBC conformance level, and a brief description of the function.

Table 10 (Page 1 of 5). Function list by category

Task function name	ODBC 2.0	DB2 ODBC support	Purpose
<b>Note:</b> Both Version 5 and Version 6 of DB2 ODBC support the functions indicated in the DB2 ODBC column, unless otherwise noted.			
Connecting to a data source			
<code>SQLAllocEnv</code>	Core	Yes	Obtains an environment handle. One environment handle is used for one or more connections.
<code>SQLAllocConnect</code>	Core	Yes	Obtains a connection handle.



Table 10 (Page 2 of 5). Function list by category

Task function name	ODBC 2.0	DB2 ODBC support	Purpose
SQLConnect	Core	Yes	Connects to specific driver by data source name, user ID, and password.
SQLDriverConnect	Lvl 1	Yes	Connects to a specific driver by connection string or requests that the driver manager and driver display connection dialogs for the user.  <b>Note:</b> This function is also extended by the additional IBM keywords supported in the ODBC.INI file in the DB2 UDB CLI environment. Within the DB2 for OS/390 CLI environment, there is no equivalent of the ODBC.INI file.
SQLSetConnection	No	Yes	Connects to a specific data source by connection string.
SQLBrowseConnect	Lvl 2	No	Returns successive levels of connection attributes and valid attribute values. When a value is specified for each connection attribute, connects to the data source.
Obtaining information about a driver and data source			
SQLDataSources	Lvl 2	Yes	Returns the list of available data sources.
SQLDrivers	Lvl 2	No	Returns the list of installed drivers and their attributes (ODBC 2.0).  <b>Note:</b> This function is implemented within the ODBC driver manager and is therefore not applicable within the DB2 for OS/390 ODBC environment.
SQLGetInfo	Lvl 1	Yes	Returns information about a specific driver and data source.
SQLGetFunctions	Lvl 1	Yes	Returns supported driver functions.
SQLGetTypeInfo	Lvl 1	Yes	Returns information about supported data types.
Setting and retrieving driver options			
SQLSetEnvAttr	No	Yes	Sets an environment option.
SQLGetEnvAttr	No	Yes	Returns the value of an environment option.
SQLSetConnectOption	Lvl 1	Yes	Sets a connection option.
SQLGetConnectOption	Lvl 1	Yes	Returns the value of a connection option.
SQLSetStmtOption	Lvl 1	Yes	Sets a statement option.
SQLGetStmtOption	Lvl 1	Yes	Returns the value of a statement option.
Preparing SQL Requests			
SQLAllocStmt	Core	Yes	Allocates a statement handle.
SQLPrepare	Core	Yes	Prepares an SQL statement for later execution.
SQLBindParameter	Lvl 1	Yes	Assigns storage for a parameter in an SQL statement (ODBC 2.0)
SQLSetParam	Core	Yes	Assigns storage for a parameter in an SQL statement (ODBC 2.0). In ODBC, SQLBindParameter() replaces this function.

Table 10 (Page 3 of 5). Function list by category

<b>Task function name</b>	<b>ODBC 2.0</b>	<b>DB2 ODBC support</b>	<b>Purpose</b>
SQLParamOptions	Lvl 2	Yes	Specifies the use of multiple values for parameters.
SQLGetCursorName	Core	Yes	Returns the cursor name associated with a statement handle.
SQLSetCursorName	Core	Yes	Specifies a cursor name.
SQLSetScrollOptions	Lvl 2	No	Sets options that control cursor behavior.
<b>Submitting Requests</b>			
SQLExecute	Core	Yes	Executes a prepared statement.
SQLExecDirect	Core	Yes	Executes a statement.
SQLNativeSql	Lvl 2	Yes	Returns the text of an SQL statement as translated by the driver.
SQLDescribeParam <sup>a</sup>	Lvl 2	Yes	Returns the description for a specific input parameter in a statement.
SQLNumParams	Lvl 2	Yes	Returns the number of parameters in a statement.
SQLParamData	Lvl 1	Yes	Used in conjunction with SQLPutData() to supply parameter data at execution time. (Useful for long data values.)
SQLPutData	Lvl 1	Yes	Send part or all of a data value for a parameter. (Useful for long data values.)
<b>Retrieving results and information about Results</b>			
SQLRowCount	Core	Yes	Returns the number of rows affected by an insert, update, or delete request.
SQLNumResultCols	Core	Yes	Returns the number of columns in the result set.
SQLDescribeCol	Core	Yes	Describes a column in the result set.
SQLColAttributes	Core	Yes	Describes attributes of a column in the result set.
SQLSetColAttributes	No	Yes	Sets attributes of a column in the result set.
SQLBindCol	Core	Yes	Assigns storage for a result column and specifies the data type.
SQLFetch	Core	Yes	Returns a result row.
SQLExtendedFetch	Lvl 2	Yes	Returns multiple result rows.
SQLGetData	Lvl 1	Yes	Returns part or all of one column of one row of a result set. (Useful for long data values.)
SQLSetPos	Lvl 2	No	Positions a cursor within a fetched block of data.
SQLMoreResults	Lvl 2	Yes	Determines whether there are more result sets available and, if so, initializes processing for the next result set.
SQLError	Core	Yes	Returns additional error or status information.
SQLGetSQLCA	No	Yes	Returns the SQLCA associated with a statement handle.
<b>Large object support</b>			
SQLGetLength <sup>a</sup>	No	Yes	Gets length of a string referenced by a LOB locator.

Table 10 (Page 4 of 5). Function list by category

<b>Task function name</b>	<b>ODBC 2.0</b>	<b>DB2 ODBC support</b>	<b>Purpose</b>
SQLGetPosition <sup>a</sup>	No	Yes	Gets the position of a string within a source string referenced by a LOB locator.
SQLGetSubString <sup>a</sup>	No	Yes	Creates a new LOB locator that references a substring within a source string (the source string is also represented by a LOB locator).
<b>Obtaining information about the data source's system tables (catalog functions)</b>			
SQLColumnPrivileges	Lvl 2	Yes	Returns a list of columns and associated privileges for a table.
SQLColumns	Lvl 1	Yes	Returns the list of column names in specified tables.
SQLForeignKeys	Lvl 2	Yes	Returns a list of column names that comprise foreign keys, if they exist for a specified table.
SQLPrimaryKeys	Lvl 2	Yes	Returns the list of column names that comprise the primary key for a table.
SQLProcedureColumns	Lvl 2	Yes	Returns the list of input and output parameters, as well as the columns that make up the result set for the specified procedures.
SQLProcedures	Lvl 2	Yes	Returns the list of procedure names stored in a specific data source.
SQLSpecialColumns	Lvl 1	Yes	Returns information about the optimal set of columns that uniquely identifies a row in a specified table, or the columns that are automatically updated when any value in the row is updated by a transaction.
SQLStatistics	Lvl 1	Yes	Returns statistics about a single table and the list of indexes associated with the table.
SQLTablePrivileges	Lvl 2	Yes	Returns a list of tables and the privileges associated with each table.
SQLTables	Lvl 1	Yes	Returns the list of table names stored in a specific data source.
<b>Terminating a Statement</b>			
SQLFreeStmt	Core	Yes	End statement processing and closes the associated cursor, discards pending results, and, optionally, frees all resources associated with the statement handle.
SQLCancel	Core	Yes	Cancels an SQL statement.
SQLTransact	Core	Yes	Commits or rolls back a transaction.
<b>Terminating a Connection</b>			
SQLDisconnect	Core	Yes	Closes the connection.
SQLFreeConnect	Core	Yes	Releases the connection handle.
SQLFreeEnv	Core	Yes	Releases the environment handle.

Table 10 (Page 5 of 5). Function list by category

---

<b>Task function name</b>	<b>ODBC 2.0</b>	<b>DB2 ODBC support</b>	<b>Purpose</b>
-------------------------------	---------------------	---------------------------------	----------------

---

**Note:**

a This function is supported by Version 6 of DB2 ODBC only.

ODBC functions not supported by DB2 ODBC:

- `SQLSetPos()` and `SQLBrowseConnect()` are not supported by Call Level Interface or DB2 ODBC.
- `SQLSetScrollOptions()` is not supported. It is superseded by the `SQL_CURSOR_TYPE`, `SQL_CONCURRENCY`, `SQL_KEYSET_SIZE`, and `SQL_ROWSET_SIZE` statement options.
- `SQLDrivers()` is implemented by the ODBC driver manager and is not supported by DB2 ODBC.

## SQLAllocConnect - Allocate connection handle

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLAllocConnect() allocates a connection handle and associated resources within the environment identified by the input environment handle. Call SQLGetInfo() with fInfoType set to SQL\_ACTIVE\_CONNECTIONS, to query the number of connections that can be allocated at any one time.

While this API is active, the DB2 ODBC driver establishes an affinity with the DB2 subsystem. Processing includes allocating a DB2 for OS/390 plan as a resource.

SQLAllocEnv() must be called before calling this function.

This function must be called before calling SQLConnect() or SQLDriverConnect().

### Syntax

```
SQLRETURN SQLAllocConnect (SQLHENV      henv,
                           SQLHDBC FAR *phdbc);
```

### Function arguments

Table 11. SQLAllocConnect arguments

Data type	Argument	Use	Description
SQLHENV	<i>henv</i>	input	Environment handle
SQLHDBC *	<i>phdbc</i>	output	Pointer to connection handle

### Usage

The output connection handle is used by DB2 ODBC to reference all information related to the connection, including general status information, transaction state, and error information.

If the pointer to the connection handle (*phdbc*) already points to a valid connection handle previously allocated by SQLAllocConnect(), then the original value is overwritten as a result of this call. This is an application programming error which is not detected by DB2 ODBC.

### Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

If SQL\_ERROR is returned, the *phdbc* argument is set to SQL\_NULL\_HDBC. The application should call SQLError() with the environment handle (*henv*) and with *hdbc* and *hstmt* arguments set to SQL\_NULL\_HDBC and SQL\_NULL\_HSTMT respectively.

## Diagnostics

Table 12. SQLAllocConnect SQLSTATES

SQLSTATE	Description	Explanation
58004	Unexpected system failure.	This could be a failure to establish the association with the DB2 for OS/390 subsystem or any other system related error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	<i>phdbc</i> was a null pointer.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1014	No more handles.	Returned if the MAXCONN keyword or SQL_MAXCONN attribute is set to a positive integer and the number of connections has reached that value. If MAXCONN is set to zero, there is no limit.  DB2 ODBC is not able to allocate a handle due to internal resources.

## Restrictions

None.

## Example

The following example shows a basic connect, with minimal error handling.

```

/* ... */
/*****
** - demonstrate basic connection to two data sources.
** - error handling mostly ignored for simplicity
**
** Functions used:
**
**   SQLAllocConnect  SQLDisconnect
**   SQLAllocEnv     SQLFreeConnect
**   SQLConnect      SQLFreeEnv
** Local Functions:
**   DBconnect
**
*****/

#include <stdio.h>
#include <stdlib.h>
#include "sqlcli1.h"
int
DBconnect(SQLHENV henv,
          SQLHDBC * hdbc,
          char * server);

#define MAX_UID_LENGTH 18
#define MAX_PWD_LENGTH 30
#define MAX_CONNECTIONS 2

```

```

int
main( )
{
    SQLHENV          henv;
    SQLHDBC          hdbc[MAX_CONNECTIONS];
    SERVER           svr[MAX_CONNECTIONS] =
        {
            "KARACHI" ,
            "DAMASSUSS"
        }

    /* allocate an environment handle */
    SQLAllocEnv(&henv);

    /* Connect to first data source */
    DBconnect(henv, &hdbc[0],
              svr[0]);

    /* Connect to second data source */
    DBconnect(henv, &hdbc[1],
              svr[1]);

    /****** Start Processing Step *****/
    /* allocate statement handle, execute statement, etc. */
    /****** End Processing Step *****/

    /******
    /* Commit work on connection 1. This has NO effect on the */
    /* transaction active on connection 2. */
    /******

    SQLTransact (henv,
                 hdbc[0],
                 SQL_COMMIT);

    /******
    /* Commit work on connection 2. This has NO effect on the */
    /* transaction active on connection 1. */
    /******

    SQLTransact (henv,
                 hdbc[1],
                 SQL_COMMIT);

    printf("\nDisconnecting ....\n");

    SQLDisconnect(hdbc[0]); /* disconnect first connection */

    SQLDisconnect(hdbc[1]); /* disconnect second connection */
    SQLFreeConnect(hdbc[0]); /* free first connection handle */
    SQLFreeConnect(hdbc[1]); /* free second connection handle */
    SQLFreeEnv(henv); /* free environment handle */

    return (SQL_SUCCESS);
}

```

## SQLAllocConnect

```

/*****
**  Server is passed as a parameter. Note that USERID and PASSWORD**
**  are always NULL.                                          **
*****/

int
DBconnect(SQLHENV henv,
          SQLHDBC * hdbc,
          char * server)
{
    SQLRETURN    rc;
    SQLCHAR      buffer[255];
    SQLSMALLINT  outlen;

    SQLAllocConnect(henv, hdbc); /* allocate a connection handle */

    rc = SQLConnect(*hdbc, server, SQL_NTS, NULL, SQL_NTS, NULL, SQL_NTS);
    if (rc != SQL_SUCCESS) {
        printf(">--- Error while connecting to database: %s -----\n", server);
        return (SQL_ERROR);
    } else {
        printf(">Connected to %s\n", server);
        return (SQL_SUCCESS);
    }
}
/* ... */

```

## References

- “SQLAllocEnv - Allocate environment handle” on page 83
- “SQLConnect - Connect to a data source” on page 123
- “SQLDriverConnect - (Expanded) connect to a data source” on page 142
- “SQLDisconnect - Disconnect from a data source” on page 140
- “SQLFreeConnect - Free connection handle” on page 184
- “SQLGetConnectOption - Returns current setting of a connect option” on page 191
- “SQLSetConnectOption - Set connection option” on page 316



## SQLAllocEnv - Allocate environment handle

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>	<b>X/OPEN CLI</b>	<b>ISO CLI</b>
-----------------------	-----------------	-------------------	----------------

SQLAllocEnv() allocates an environment handle and associated resources. There can be only one environment active at any one time per application.

An application must call this function prior to SQLAllocConnect() or any other DB2 ODBC functions. The *henv* value is passed in all subsequent function calls that require an environment handle as input.

### Syntax

```
SQLRETURN SQLAllocEnv (SQLHENV FAR *phenv);
```

### Function arguments

Table 13. SQLAllocEnv arguments

Data type	Argument	Use	Description
SQLHENV *	<i>phenv</i>	output	Pointer to environment handle

### Usage

There can be only one active environment at a time per application. Any subsequent calls to SQLAllocEnv() return the same handle as the first SQLAllocEnv() call.

SQLFreeEnv() must be called for each successful SQLAllocEnv() call before the resources associated with the handle are released. SQLFreeEnv() must also be called to free a restricted environment handle as described under 'Return Codes' below.

### Return codes

- SQL\_SUCCESS
- SQL\_ERROR

If SQL\_ERROR is returned and *phenv* is equal to SQL\_NULL\_HENV, then SQLERROR() cannot be called because there is no handle with which to associate additional diagnostic information.

If the return code is SQL\_ERROR and the pointer to the environment handle is not equal to SQL\_NULL\_HENV, then the handle is a *restricted handle*. This means the handle can only be used in a call to SQLERROR() to obtain more error information, or to SQLFreeEnv().

### Diagnostics

Table 14. SQLAllocEnv SQLSTATES

SQLSTATE	Description	Explanation
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.

### Restrictions

None.

### Example

See “Example” on page 80.

### References

- “SQLAllocConnect - Allocate connection handle” on page 79
- “SQLFreeEnv - Free environment handle” on page 186

## SQLAllocStmt - Allocate a statement handle

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLAllocStmt() allocates a new statement handle and associates it with the connection specified by the connection handle. There is no defined limit on the number of statement handles that can be allocated at any one time.

SQLConnect() or SQLDriverConnect() must be called before calling this function.

This function must be called before SQLBindParameter(), SQLPrepare(), SQLExecute(), SQLExecDirect(), or any other function that has a statement handle as one of its input arguments.

### Syntax

```
SQLRETURN SQLAllocStmt (SQLHDBC hdbc,
                        SQLHSTMT FAR *phstmt);
```

### Function arguments

Table 15. SQLAllocStmt arguments

Data type	Argument	Use	Description
SQLHDBC	<i>hdbc</i>	input	Connection handle
SQLHSTMT *	<i>phstmt</i>	output	Pointer to statement handle

### Usage

DB2 ODBC uses each statement handle to relate all the descriptors, attribute values, result values, cursor information, and status information to the SQL statement processed. Although each SQL statement must have a statement handle, you can reuse the handles for different statements.

A call to this function requires that *hdbc* references an active database connection.

To execute a positioned UPDATE or DELETE, the application must use different statement handles for the SELECT statement and the UPDATE or DELETE statement.

If the input pointer to the statement handle (*phstmt*) already points to a valid statement handle allocated by a previous call to SQLAllocStmt(), then the original value is overwritten as a result of this call. This is an application programming error that is not detected by DB2 ODBC.

## Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

If SQL\_ERROR is returned, the *phstmt* argument is set to SQL\_NULL\_HSTMT. The application should call SQLError() with the same *hdbc* and with the *hstmt* argument set to SQL\_NULL\_HSTMT.

## Diagnostics

Table 16. SQLAllocStmt SQLSTATEs

SQLSTATE	Description	Explanation
08003	Connection is closed.	The connection specified by the <i>hdbc</i> argument is not open. The connection must be established successfully (and the connection must be open) for the application to call SQLAllocStmt().
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	<i>phstmt</i> was a null pointer.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1014	No more handles.	DB2 ODBC is not able to allocate a handle due to internal resources.

## Restrictions

None.

## Example

See “Example” on page 173.

## References

- “SQLConnect - Connect to a data source” on page 123
- “SQLDriverConnect - (Expanded) connect to a data source” on page 142
- “SQLFreeStmt - Free (or reset) a statement handle” on page 188
- “SQLGetStmtOption - Returns current setting of a statement option” on page 249
- “SQLSetStmtOption - Set statement option” on page 334

## SQLBindCol - Bind a column to an application variable

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>	<b>X/OPEN CLI</b>	<b>ISO CLI</b>
-----------------------	-----------------	-------------------	----------------

SQLBindCol () is used to associate (bind) columns in a result set to:

- Application variables or arrays of application variables (storage buffers), for all C data types. In this case, data is transferred from the DBMS to the application when SQLFetch() or SQLExtendedFetch() is called. Data conversion can occur as the data is transferred.
- A LOB locator, for LOB columns. In this case a LOB locator, not the data itself, is transferred from the DBMS to the application when SQLFetch() is called. A LOB locator can represent the entire data or a portion of the data.

SQLBindCol () is called once for each column in the result set that the application needs to retrieve.

In general, SQLPrepare(), SQLExecDirect() or one of the schema functions is called before this function, and SQLFetch() or SQLExtendedFetch() is called after. Column attributes might also be needed before calling SQLBindCol(), and can be obtained using SQLDescribeCol() or SQLColAttributes().

### Syntax

```
SQLRETURN SQLBindCol (SQLHSTMT hstmt,
                      SQLUSMALLINT icol,
                      SQLSMALLINT fctype,
                      SQLPOINTER rgbValue,
                      SQLINTEGER cbValueMax,
                      SQLINTEGER FAR *pcbValue);
```

### Function arguments

Table 17 (Page 1 of 2). SQLBindCol arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle
SQLUSMALLINT	<i>icol</i>	input	Number identifying the column. Columns are numbered sequentially, from left to right, starting at 1.

Table 17 (Page 2 of 2). SQLBindCol arguments

Data type	Argument	Use	Description
SQLSMALLINT	<i>fCType</i>	input	<p>The C data type for column number <i>icol</i> in the result set. The following types are supported:</p> <ul style="list-style-type: none"> <li>• SQL_C_BINARY</li> <li>• SQL_C_BIT</li> <li>• SQL_C_BLOB_LOCATOR</li> <li>• SQL_C_CHAR</li> <li>• SQL_C_CLOB_LOCATOR</li> <li>• SQL_C_DATE</li> <li>• SQL_C_DBCHAR</li> <li>• SQL_C_DBCLOB_LOCATOR</li> <li>• SQL_C_DOUBLE</li> <li>• SQL_C_FLOAT</li> <li>• SQL_C_LONG</li> <li>• SQL_C_SHORT</li> <li>• SQL_C_TIME</li> <li>• SQL_C_TIMESTAMP</li> <li>• SQL_C_TINYINT</li> </ul> <p>The supported data types are based on the data source to which you are connected. Specifying SQL_C_DEFAULT causes data to be transferred to its default C data type. See Table 4 on page 39 for more information.</p>
SQLPOINTER	<i>rgbValue</i>	output (deferred)	<p>Pointer to buffer (or an array of buffers if using SQLExtendedFetch()) where DB2 ODBC is to store the column data or the LOB locator when the fetch occurs. If <i>rgbValue</i> is null, the column is unbound.</p>
SQLINTEGER	<i>cbValueMax</i>	input	<p>Size of <i>rgbValue</i> buffer in bytes available to store the column data or the LOB locator.</p> <p>If <i>fCType</i> denotes a binary or character string (either single or double byte) or is SQL_C_DEFAULT, then <i>cbValueMax</i> must be &gt; 0, or an error is returned. Otherwise, this argument is ignored.</p>
SQLINTEGER *	<i>pcbValue</i>	output (deferred)	<p>Pointer to value (or array of values) which indicates the number of bytes DB2 ODBC has available to return in the <i>rgbValue</i> buffer. If <i>fCType</i> is a LOB locator, the size of the locator is returned, not the size of the LOB data.</p> <p>SQLFetch() returns SQL_NULL_DATA in this argument if the data value of the column is null.</p> <p>This pointer value must be unique for each bound column, or NULL.</p> <p>SQL_NO_LENGTH can also be returned. See the 'Usage' section below for more information.</p>

**Note:**

- For this function, pointers *rgbValue* and *pcbValue* are deferred outputs, meaning that the storage locations they point to do not get updated until a result set row is fetched. As a result, the locations referenced by these pointers must remain valid until SQLFetch() or SQLExtendedFetch() is called. For example, if SQLBindCol() is called within a local function, SQLFetch() must be

called from within the same scope of the function or the *rgbValue* buffer must be allocated as static or global.

- DB2 ODBC performs better for all variable length data types if *rgbValue* is placed consecutively in memory after *pcbValue*. See the 'Usage' section for more details.

## Usage

The application calls `SQLBindCol()` once for each column in the result set for which it wishes to retrieve either the data, or optionally in the case of LOB columns, a LOB locator. Result sets are generated either by calling `SQLPrepare()`, `SQLExecDirect()`, `SQLGetTypeInfo()`, or one of the catalog functions. When `SQLFetch()` is called, the data in each of these *bound* columns is placed into the assigned location (given by the pointers *rgbValue* and *cbValue*). If *fCType* is a LOB locator, a locator value is returned (not the LOB data). The LOB locator references the entire data value in the LOB column.

`SQLExtendedFetch()` can be used in place of `SQLFetch()` to retrieve multiple rows from the result set into an array. In this case, *rgbValue* references an array. For more information, see “Retrieving a result set into an array” on page 379 and “`SQLExtendedFetch` - Extended fetch (fetch array of rows)” on page 162. Use of `SQLExtendedFetch()` and `SQLFetch()` cannot be mixed for the same result set.

Columns are identified by a number, assigned sequentially from left to right, starting at 1. The number of columns in the result set can be determined by calling `SQLNumResultCols()` or by calling `SQLColAttributes()` with the *fdescType* argument set to `SQL_COLUMN_COUNT`.

The application can query the attributes (such as data type and length) of the column by first calling `SQLDescribeCol()` or `SQLColAttributes()`. (As an alternative, see “Programming hints and tips” on page 407 for information about using `SQLSetColAttributes()` when the application has prior knowledge of the format of the result set.) This information can then be used to allocate a storage location of the correct data type and length to indicate data conversion to another data type, or in the case of LOB data types, optionally return a locator. See “Data types and data conversion” on page 36 for more information on default types and supported conversions.

An application can choose not to bind every column, or even not to bind any columns. Data in any of the columns can also be retrieved using `SQLGetData()` after the bound columns have been fetched for the current row. Generally, `SQLBindCol()` is more efficient than `SQLGetData()`. For a discussion of when to use one function over the other, refer to “Programming hints and tips” on page 407.

In subsequent fetches, the application can change the binding of these columns or bind previously unbound columns by calling `SQLBindCol()`. The new binding does not apply to data already fetched, it is used on the next fetch. To unbind a single column, call `SQLBindCol()` with the *rgbValue* pointer set to `NULL`. To unbind all the columns, the application should call `SQLFreeStmt()` with the *fOption* input set to `SQL_UNBIND`.

The application must ensure enough storage is allocated for the data to be retrieved. If the buffer is to contain variable length data, the application must allocate as much storage as the maximum length of the bound column requires; otherwise, the data might be truncated. If the buffer is to contain fixed length data,

DB2 ODBC assumes the size of the buffer is the length of the C data type. If data conversion is specified, the required size might be affected, see “Data types and data conversion” on page 36 for more information.

If string truncation does occur, `SQL_SUCCESS_WITH_INFO` is returned and `pcbValue` is set to the actual size of `rgbValue` available for return to the application.

Truncation is also affected by the `SQL_MAX_LENGTH` statement option (used to limit the amount of data returned to the application). The application can specify not to report truncation by calling `SQLSetStmtOption()` with `SQL_MAX_LENGTH` and a value for the maximum length to return for all variable length columns, and by allocating an `rgbValue` buffer of the same size (plus the null-terminator). If the column data is larger than the set maximum length, `SQL_SUCCESS` is returned when the value is fetched and the maximum length, not the actual length, is returned in `pcbValue`.

If the column to be bound is an `SQL_GRAPHIC`, `SQL_VARGRAPHIC` or `SQL_LONGVARGRAPHIC` type, then `fCType` can be set to `SQL_C_DBCHAR` or `SQL_C_CHAR`. If `fCType` is `SQL_C_DBCHAR`, the data fetched into the `rgbValue` buffer is null-terminated by a double byte null-terminator. If `fCType` is `SQL_C_CHAR`, then the data is not null-terminated. In both cases, the length of the `rgbValue` buffer (`cbValueMax`) is in units of bytes and should therefore be a multiple of 2.

When binding any variable length column, DB2 ODBC can write `pcbValue` and `rgbValue` in one operation if they are allocated contiguously. For example:

```
struct {
    SQLINTEGER  pcbValue;
    SQLCHAR     rgbValue[MAX_BUFFER];
} column;
```

**Note:** `SQL_NO_TOTAL` is returned in `pcbValue` if:

- The SQL type is a variable length type, and
- `pcbValue` and `rgbValue` are contiguous, and
- The column type is NOT NULLABLE, and
- String truncation occurred.

If the column to be bound is an `SQL_ROWID` type, then `fCType` can be set to `SQL_C_CHAR` or `SQL_C_DEFAULT`. The data fetched into the `rgbValue` buffer is null-terminated. The maximum length of a ROWID column in the DBMS is 40 bytes, therefore, the application must allocate an `rgbValue` buffer of at least 40 bytes for the data to be retrieved without truncation.

LOB locators are generally treated like any other data type, but there are some important differences:

- Locators are generated at the server when a row is fetched and a LOB locator C data type is specified on `SQLBindCol()`, or when `SQLGetSubString()` is called to define a locator on a portion of another LOB. Only the locator is transferred to the application.
- The value of the locator is only valid within the current transaction. You cannot store a locator value and use it beyond the current transaction, even if the cursor used to fetch the LOB locator has the WITH HOLD attribute.
- A locator can be freed before the end of the transaction using the `FREE LOCATOR` statement.



- When a locator is received, the application can use `SQLGetSubString()` to either receive a portion of the LOB value, or to generate another locator representing the substring. The locator value can also be used as input for a parameter marker (using `SQLBindParameter()`).

A LOB locator is not a pointer to a database position; rather it is a reference to a LOB value, a snapshot of that LOB value. There is no association between the current position of the cursor and the row from which the LOB value is extracted. Therefore, even after the cursor moves to a different row, the LOB locator {and thus the value that it represents) can still be referenced.

- `SQLGetPosition()` and `SQLGetLength()` can be used with `SQLGetSubString()` to define the substring.

A given LOB column in the result set can be bound to one of the following:

- Storage buffer, for holding the entire LOB data value
- LOB locator

The most recent bind column function call determines the type of binding that is in effect.

## Return codes

- `SQL_SUCCESS`
- `SQL_ERROR`
- `SQL_INVALID_HANDLE`

## Diagnostics

Table 18 (Page 1 of 2). *SQLBindCol* SQLSTATEs

SQLSTATE	Description	Explanation
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1002	Invalid column number.	The value specified for the argument <i>icol</i> was less than 1. The value specified for the argument <i>icol</i> exceeded the maximum number of columns supported by the data source.
S1003	Program type out of range.	<i>fCType</i> is not a valid data type or <code>SQL_C_DEFAULT</code> .
S1010	Function sequence error.	The function is called while in a data-at-execute ( <code>SQLParamData()</code> , <code>SQLPutData()</code> ) operation.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1090	Invalid string or buffer length.	The value specified for the argument <i>cbValueMax</i> is less than 1 and the argument <i>fCType</i> is either <code>SQL_C_CHAR</code> , <code>SQL_C_BINARY</code> or <code>SQL_C_DEFAULT</code> .
S1C00	Driver not capable.	DB2 ODBC recognizes, but does not support the data type specified in the argument <i>fCType</i> A LOB locator C data type is specified, but the connected server does not support LOB data types.

## SQLBindCol

Table 18 (Page 2 of 2). SQLBindCol SQLSTATES

SQLSTATE	Description	Explanation
<b>Note:</b> Additional diagnostic messages relating to the bound columns might be reported at fetch time.		

### Restrictions

None.

### Example

See “Example” on page 173.

### References

- “SQLFetch - Fetch next row” on page 169
- “SQLExtendedFetch - Extended fetch (fetch array of rows)” on page 162

## SQLBindParameter - Binds a parameter marker to a buffer or LOB locator

### Purpose

Specification:	ODBC 2.0		
----------------	----------	--	--

SQLBindParameter() associates (bind) parameter markers in an SQL statement to either:

- Application variables or arrays of application variables (storage buffers), for all C data types. In this case, data is transferred from the application to the DBMS when SQLExecute() or SQLExecDirect() is called. Data conversion can occur as the data is transferred.
- A LOB locator, for SQL LOB data types. In this case, the application transfers a LOB locator value (not the LOB data itself) to the server when the SQL statement is executed.

This function must also be used to bind application storage to a parameter of a stored procedure CALL statement where the parameter can be input, output or both. This function is essentially an extension of SQLSetParam().

### Syntax

```
SQLRETURN SQL_API SQLBindParameter(
    SQLHSTMT          hstmt,
    SQLUSMALLINT      ipar,
    SQLSMALLINT       fParamType,
    SQLSMALLINT       fCType,
    SQLSMALLINT       fSqlType,
    SQLINTEGER         cbColDef,
    SQLSMALLINT       ibScale,
    SQLPOINTER        rgbValue,
    SQLINTEGER         cbValueMax,
    SQLINTEGER        FAR *pcbValue);
```

### Function arguments

Table 19 (Page 1 of 5). SQLBindParameter arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle
SQLUSMALLINT	ipar	input	Parameter marker number, ordered sequentially left to right, starting at 1.

Table 19 (Page 2 of 5). SQLBindParameter arguments

Data type	Argument	Use	Description
SQLSMALLINT	fParamType	input	<p>The type of parameter. The supported types are:</p> <ul style="list-style-type: none"> <li>SQL_PARAM_INPUT: The parameter marker is associated with an SQL statement that is not a stored procedure CALL; or, it marks an input parameter of the CALLED stored procedure.</li> </ul> <p>When the statement is executed, actual data value for the parameter is sent to the server: the <i>rgbValue</i> buffer must contain valid input data values; the <i>pcbValue</i> buffer must contain the corresponding length value or SQL_NTS, SQL_NULL_DATA, or (if the value should be sent using SQLParamData() and SQLPutData()) SQL_DATA_AT_EXEC.</p> <ul style="list-style-type: none"> <li>SQL_PARAM_INPUT_OUTPUT: The parameter marker is associated with an input/output parameter of the CALLED stored procedure.</li> </ul> <p>When the statement is executed, actual data value for the parameter is sent to the server: the <i>rgbValue</i> buffer must contain valid input data values; the <i>pcbValue</i> buffer must contain the corresponding length value or SQL_NTS, SQL_NULL_DATA, or (if the value should be sent using SQLParamData() and SQLPutData()) SQL_DATA_AT_EXEC.</p> <ul style="list-style-type: none"> <li>SQL_PARAM_OUTPUT: The parameter marker is associated with an output parameter of the CALLED stored procedure or the return value of the stored procedure.</li> </ul> <p>After the statement is executed, data for the output parameter is returned to the application buffer specified by <i>rgbValue</i> and <i>pcbValue</i>, unless both are NULL pointers, in which case the output data is discarded.</p>
SQLSMALLINT	fCType	input	<p>C data type of the parameter. The following types are supported:</p> <ul style="list-style-type: none"> <li>SQL_C_BINARY</li> <li>SQL_C_BIT</li> <li>SQL_C_BLOB_LOCATOR</li> <li>SQL_C_CHAR</li> <li>SQL_C_CLOB_LOCATOR</li> <li>SQL_C_DATE</li> <li>SQL_C_DBCHAR</li> <li>SQL_C_DBCLOB_LOCATOR</li> <li>SQL_C_DOUBLE</li> <li>SQL_C_FLOAT</li> <li>SQL_C_LONG</li> <li>SQL_C_SHORT</li> <li>SQL_C_TIME</li> <li>SQL_C_TIMESTAMP</li> <li>SQL_C_TINYINT</li> </ul> <p>Specifying SQL_C_DEFAULT causes data to be transferred from its default C data type to the type indicated in <i>fSqlType</i>.</p>

Table 19 (Page 3 of 5). SQLBindParameter arguments

Data type	Argument	Use	Description
SQLSMALLINT	fSqlType	input	<p>SQL data type of the parameter. The supported types are:</p> <ul style="list-style-type: none"> <li>• SQL_BINARY</li> <li>• SQL_BLOB</li> <li>• SQL_BLOB_LOCATOR</li> <li>• SQL_CHAR</li> <li>• SQL_CLOB</li> <li>• SQL_CLOB_LOCATOR</li> <li>• SQL_DATE</li> <li>• SQL_DBCLOB</li> <li>• SQL_DBCLOB_LOCATOR</li> <li>• SQL_DECIMAL</li> <li>• SQL_DOUBLE</li> <li>• SQL_FLOAT</li> <li>• SQL_GRAPHIC</li> <li>• SQL_INTEGER</li> <li>• SQL_LONGVARIABLE</li> <li>• SQL_LONGVARCHAR</li> <li>• SQL_LONGVARGRAPHIC</li> <li>• SQL_NUMERIC</li> <li>• SQL_REAL</li> <li>• SQL_ROWID</li> <li>• SQL_SMALLINT</li> <li>• SQL_TIME</li> <li>• SQL_TIMESTAMP</li> <li>• SQL_VARBINARY</li> <li>• SQL_VARCHAR</li> <li>• SQL_VARGRAPHIC</li> </ul> <p><b>Note:</b> SQL_BLOB_LOCATOR, SQL_CLOB_LOCATOR, and SQL_DBCLOB_LOCATOR are application related concepts and do not map to a data type for column definition during a CREATE TABLE.</p>
SQLINTEGER	cbColDef	input	<p>Precision of the corresponding parameter marker. If <i>fSqlType</i> denotes:</p> <ul style="list-style-type: none"> <li>• A binary or single byte character string (for example, SQL_CHAR, SQL_BINARY), this is the maximum length in bytes for this parameter marker.</li> <li>• A double byte character string (for example, SQL_GRAPHIC), this is the maximum length in double-byte characters for this parameter.</li> <li>• SQL_DECIMAL, SQL_NUMERIC, this is the maximum decimal precision.</li> <li>• SQL_ROWID, this must be set to 40, the maximum length in bytes for this data type. Otherwise, an error is returned.</li> <li>• Otherwise, this argument is ignored.</li> </ul>
SQLSMALLINT	ibScale	input	<p>Scale of the corresponding parameter if <i>fSqlType</i> is SQL_DECIMAL or SQL_NUMERIC. If <i>fSqlType</i> is SQL_TIMESTAMP, this is the number of digits to the right of the decimal point in the character representation of a timestamp (for example, the scale of yyyy-mm-dd hh:mm:ss.fff is 3).</p> <p>Other than for the <i>fSqlType</i> values mentioned here, <i>ibScale</i> is ignored.</p>

Table 19 (Page 4 of 5). SQLBindParameter arguments

Data type	Argument	Use	Description
SQLPOINTER	rgbValue	input (deferred) and/or output (deferred)	<ul style="list-style-type: none"> <li>On input (<i>fParamType</i> set to SQL_PARAM_INPUT, or SQL_PARAM_INPUT_OUTPUT):           <p>At execution time, if <i>pcbValue</i> does not contain SQL_NULL_DATA or SQL_DATA_AT_EXEC, then <i>rgbValue</i> points to a buffer that contains the actual data for the parameter.</p> <p>If <i>pcbValue</i> contains SQL_DATA_AT_EXEC, then <i>rgbValue</i> is an application-defined 32-bit value that is associated with this parameter. This 32-bit value is returned to the application using a subsequent SQLParamData() call.</p> <p>If SQLParamOptions() is called to specify multiple values for the parameter, then <i>rgbValue</i> is a pointer to an input buffer array of <i>cbValueMax</i> bytes.</p> </li> <li>On output (<i>fParamType</i>) set to SQL_PARAM_OUTPUT, or SQL_PARAM_INPUT_OUTPUT):           <p><i>rgbValue</i> points to the buffer where the output parameter value of the stored procedure is stored.</p> <p>If <i>fParamType</i> is set to SQL_PARAM_OUTPUT, and both <i>rgbValue</i> and <i>pcbValue</i> are NULL pointers, then the output parameter value or the return value from the stored procedure call is discarded.</p> </li> </ul>
SQLINTEGER	cbValueMax	input	<p>For character and binary data, <i>cbValueMax</i> specifies the length of the <i>rgbValue</i> buffer (if treated as a single element) or the length of each element in the <i>rgbValue</i> array (if the application calls SQLParamOptions() to specify multiple values for each parameter). For non-character and non-binary data, this argument is ignored -- the length of the <i>rgbValue</i> buffer (if it is a single element) or the length of each element in the <i>rgbValue</i> array (if SQLParamOptions() is used to specify an array of values for each parameter) is assumed to be the length associated with the C data type.</p> <p>For output parameters, <i>cbValueMax</i> is used to determine whether to truncate character or binary output data in the following manner:</p> <ul style="list-style-type: none"> <li>For character data, if the number of bytes available to return is greater than or equal to <i>cbValueMax</i>, the data in <i>rgbValue</i> is truncated to <i>cbValueMax-1</i> bytes and is null-terminated (unless null-termination has been turned off).</li> <li>For binary data, if the number of bytes available to return is greater than <i>cbValueMax</i>, the data in <i>rgbValue</i> is truncated to <i>cbValueMax</i> bytes.</li> </ul>

Table 19 (Page 5 of 5). SQLBindParameter arguments

Data type	Argument	Use	Description
SQLINTEGER *	pcbValue	input (deferred) and/or output (deferred)	<p>- If this is an input or input/output parameter:</p> <p>This is the pointer to the location which contains (when the statement is executed) the length of the parameter marker value stored at <i>rgbValue</i>.</p> <p>To specify a null value for a parameter marker, this storage location must contain SQL_NULL_DATA.</p> <p>If <i>fCType</i> is SQL_C_CHAR, this storage location must contain either the exact length of the data stored at <i>rgbValue</i>, or SQL_NTS if the contents at <i>rgbValue</i> are null-terminated.</p> <p>If <i>fCType</i> indicates character data (explicitly, or implicitly using SQL_C_DEFAULT), and this pointer is set to NULL, it is assumed that the application always provides a null-terminated string in <i>rgbValue</i>. This also implies that this parameter marker never has a null value.</p> <p>If <i>fSqlType</i> denotes a graphic data type and the <i>fCType</i> is SQL_C_CHAR, the pointer to <i>pcbValue</i> can never be NULL and the contents of <i>pcbValue</i> can never hold SQL_NTS. In general for graphic data types, this length should be the number of octets that the double byte data occupies; therefore, the length should always be a multiple of 2. In fact, if the length is odd, then an error occurs when the statement is executed.</p> <p>When SQLExecute() or SQLExecDirect() is called, and <i>pcbValue</i> points to a value of SQL_DATA_AT_EXEC, the data for the parameter is sent with SQLPutData(). This parameter is referred to as a <i>data-at-execution</i> parameter.</p> <p>If SQLParamOptions() is used to specify multiple values for each parameter, <i>pcbValue</i> points to an array of SQLINTEGER values where each of the elements can be the number of bytes in the corresponding <i>rgbValue</i> element (excluding the null-terminator), or SQL_NULL_DATA.</p> <p>- If this is an output parameter (<i>fParamType</i> is set to SQL_PARAM_OUTPUT):</p> <p>This must be an output parameter or return value of a stored procedure CALL and points to one of the following, after the execution of the stored procedure:</p> <ul style="list-style-type: none"> <li>• number of bytes available to return in <i>rgbValue</i>, excluding the null-termination character.</li> <li>• SQL_NULL_DATA</li> <li>• SQL_NO_TOTAL if the number of bytes available to return cannot be determined.</li> </ul>

### Usage

A parameter marker is represented by a "?" character in an SQL statement and is used to indicate a position in the statement where an application supplied value is to be substituted when the statement is executed. This value can be obtained from:

- An application variable. `SQLBindParameter()` (or `SQLSetParam()`) is used to bind the application storage area to the parameter marker.
- A LOB value from the database server (by specifying a LOB locator).

`SQLBindParameter()` (or `SQLSetParam()`) is used to bind a LOB locator to the parameter marker. The LOB value itself is supplied by the database server, so only the LOB locator is transferred between the database server and the application.

An application can use a locator with `SQLGetSubstring()`, `SQLGetPosition()`, or `SQLGetLength()`. `SQLGetSubstring()` can either return another locator or the data itself. All locators remain valid until the end of the transaction in which they are created (even when the cursor moves to another row), or until it is freed using the `FREE LOCATOR` statement.

The application must bind a variable to each parameter marker in the SQL statement before executing the SQL statement. For this function, *rgbValue* and *pcbValue* are deferred arguments, the storage locations must be valid and contain input data values when the statement is executed. This means either keeping the `SQLExecDirect()` or `SQLExecute()` call in the same procedure scope as the `SQLBindParameter()` calls, or, these storage locations must be dynamically allocated or declared statically or globally.

`SQLSetParam()` can be called before `SQLPrepare()` if the columns in the result set are known; otherwise, the attributes of the result set can be obtained after the statement is prepared.

Parameter markers are referenced by number (*icol*) and are numbered sequentially from left to right, starting at 1.

All parameters bound by this function remain in effect until `SQLFreeStmt()` is called with either the `SQL_DROP` or `SQL_RESET_PARAMS` option, or until `SQLBindParameter()` is called again for the same parameter *ipar* number.

After the SQL statement is executed, and the results processed, the application might wish to reuse the statement handle to execute a different SQL statement. If the parameter marker specifications are different (number of parameters, length or type), `SQLFreeStmt()` should be called with `SQL_RESET_PARAMS` to reset or clear the parameter bindings.

The C buffer data type given by *fCType* must be compatible with the SQL data type indicated by *fSqlType*, or an error occurs.

An application can pass the value for a parameter either in the *rgbValue* buffer or with one or more calls to `SQLPutData()`. In the latter case, these parameters are data-at-execution parameters. The application informs DB2 ODBC of a data-at-execution parameter by placing the `SQL_DATA_AT_EXEC` value in the *pcbValue* buffer. It sets the *rgbValue* input argument to a 32-bit value which is returned on a subsequent `SQLParamData()` call and can be used to identify the parameter position.



Since the data in the variables referenced by *rgbValue* and *pcbValue* is not verified until the statement is executed, data content or format errors are not detected or reported until `SQLExecute()` or `SQLExecDirect()` is called.

`SQLBindParameter()` essentially extends the capability of the `SQLSetParam()` function by providing a method of:

- Specifying whether a parameter is input, input / output, or output, necessary for proper handling of parameters for stored procedures.
- Specifying an array of input parameter values when `SQLParamOptions()` is used in conjunction with `SQLBindParameter()`. `SQLSetParam()` can still be used to bind single element application variables to parameter markers that are not part of a stored procedure CALL statement.

The *fParamType* argument specifies the type of the parameter. All parameters in the SQL statements that do not call procedures are input parameters. Parameters in stored procedure calls can be input, input/output, or output parameters. Even though the DB2 stored procedure argument convention typically implies that all procedure arguments are input/output, the application programmer can still choose to specify the nature of input or output more exactly on the `SQLBindParameter()` to follow a more rigorous coding style.

**Note:**

- If an application cannot determine the type of a parameter in a procedure call, set *fParamType* to `SQL_PARAM_INPUT`; if the data source returns a value for the parameter, DB2 ODBC discards it.
- If an application has marked a parameter as `SQL_PARAM_INPUT_OUTPUT` or `SQL_PARAM_OUTPUT` and the data source does not return a value, DB2 ODBC sets the *pcbValue* buffer to `SQL_NULL_DATA`.
- If an application marks a parameter as `SQL_PARAM_OUTPUT`, data for the parameter is returned to the application after the CALL statement is processed. If the *rgbValue* and *pcbValue* arguments are both null pointers, DB2 ODBC discards the output value. If the data source does not return a value for an output parameter, DB2 ODBC sets the *pcbValue* buffer to `SQL_NULL_DATA`.
- For this function, *rgbValue* and *pcbValue* are deferred arguments. In the case where *fParamType* is set to `SQL_PARAM_INPUT` or `SQL_PARAM_INPUT_OUTPUT`, the storage locations must be valid and contain input data values when the statement is executed. This means either keeping the `SQLExecDirect()` or `SQLExecute()` call in the same procedure scope as the `SQLBindParameter()` calls, or, these storage locations must be dynamically allocated or statically or globally declared.  
Similarly, if *fParamType* is set to `SQL_PARAM_OUTPUT` or `SQL_PARAM_INPUT_OUTPUT`, the *rgbValue* and *pcbValue* buffer locations must remain valid until the CALL statement is executed.
- If *fSqlType* is `SQL_ROWID`, the *cbColDef* value must be set to 40, the maximum length in bytes for a ROWID data type. If *cbColDef* is not set to 40, the application receives `SQLSTATE=22001` when *cbColDef* is less than 40 and `SQLSTATE=S1104` when *cbColDef* is greater than 40.

For character and binary C data, the *cbValueMax* argument specifies the length of the *rgbValue* buffer if it is a single element; or, if the application calls `SQLParamOptions()` to specify multiple values for each parameter, *cbValueMax* is

## SQLBindParameter

the length of *each* element in the *rgbValue* array, INCLUDING the null-terminator. If the application specifies multiple values, *cbValueMax* is used to determine the location of values in the *rgbValue* array. For all other types of C data, the *cbValueMax* argument is ignored.

An application can pass the value for a parameter either in the *rgbValue* buffer or with one or more calls to `SQLPutData()`. In the latter case, these parameters are data-at-execution parameters. The application informs DB2 ODBC of a data-at-execution parameter by placing the `SQL_DATA_AT_EXEC` value in the *pcbValue* buffer. It sets the *rgbValue* input argument to a 32-bit value which is returned on a subsequent `SQLParamData()` call and can be used to identify the parameter position.

When `SQLBindParameter()` is used to bind an application variable to an output parameter for a stored procedure, DB2 ODBC can provide some performance enhancement if the *rgbValue* buffer is placed consecutively in memory after the *pcbValue* buffer. For example:

```
struct {  SQLINTEGER  pcbValue;
         SQLCHAR    rgbValue[MAX_BUFFER];
         } column;
*
```

## Return codes

- `SQL_SUCCESS`
- `SQL_SUCCESS_WITH_INFO`
- `SQL_ERROR`
- `SQL_INVALID_HANDLE`

## Diagnostics

Table 20 (Page 1 of 2). *SQLBindParameter* SQLSTATEs

SQLSTATE	Description	Explanation
07006	Invalid conversion.	The conversion from the data value identified by the <i>fCType</i> argument to the data type identified by the <i>fSqlType</i> argument is not a meaningful conversion. (For example, conversion from <code>SQL_C_DATE</code> to <code>SQL_DOUBLE</code> .)
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1003	Program type out of range.	The value specified by the argument <i>fCType</i> not a valid data type or <code>SQL_C_DEFAULT</code> .
S1004	SQL data type out of range.	The value specified for the argument <i>fSqlType</i> is not a valid SQL data type.
S1009	Invalid argument value.	The argument <i>rgbValue</i> is a null pointer; the argument <i>pcbValue</i> is a null pointer; and <i>fParamType</i> is not <code>SQL_PARAM_OUTPUT</code> .

Table 20 (Page 2 of 2). SQLBindParameter SQLSTATEs

SQLSTATE	Description	Explanation
S1010	Function sequence error.	Function is called after SQLExecute() or SQLExecDirect() returned SQL_NEED_DATA, but data have not been sent for all <i>data-at-execution</i> parameters.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1090	Invalid string or buffer length.	The value specified for the argument <i>cbValueMax</i> is less than 0.
S1093	Invalid parameter number.	The value specified for the argument <i>ipar</i> is less than 1 or greater than the maximum number of parameters supported by the server.
S1094	Invalid scale value.	The value specified for <i>fSqlType</i> is either SQL_DECIMAL or SQL_NUMERIC and the value specified for <i>ibScale</i> is less than 0 or greater than the value for the argument <i>cbParamDef</i> (precision).  The value specified for <i>fCType</i> is SQL_C_TIMESTAMP; the value for <i>fSqlType</i> is either SQL_CHAR or SQL_VARCHAR; and the value for <i>ibScale</i> is less than 0 or greater than 6.
S1104	Invalid precision value.	The value specified for <i>fSqlType</i> was either SQL_DECIMAL or SQL_NUMERIC and the value specified for <i>cbParamDef</i> was less than 1.
S1105	Invalid parameter type.	<i>fParamType</i> is not one of SQL_PARAM_INPUT, SQL_PARAM_OUTPUT, or SQL_PARAM_INPUT_OUTPUT.
S1C00	Driver not capable.	DB2 ODBC or data source does not support the conversion specified by the combination of the value specified for the argument <i>fCType</i> and the value specified for the argument <i>fSqlType</i> .  The value specified for the argument <i>fSqlType</i> is not supported by either DB2 ODBC or the data source.

## Restrictions

In ODBC 2.0, this function has replaced SQLSetParam().

A new value for *pcbValue*, SQL\_DEFAULT\_PARAM, was introduced in ODBC 2.0, to indicate that the procedure should use the default value of a parameter, rather than a value sent from the application. Since DB2 stored procedure arguments do not have the concept of default values, specification of this value for *pcbValue* argument results in an error when the CALL statement is executed since the SQL\_DEFAULT\_PARAM value is considered an invalid length.

ODBC 2.0 also introduced the SQL\_LEN\_DATA\_AT\_EXEC(*length*) macro to be used with the *pcbValue* argument. The macro is used to specify the sum total length of the entire data that is sent for character or binary C data using the subsequent SQLPutData() calls. Since the DB2 ODBC driver does not need this information, the macro is not needed. An ODBC application calls SQLGetInfo() with the SQL\_NEED\_LONG\_DATA\_LEN option to check if the driver needs this information. The DATABASE 2 ODBC driver returns 'N' to indicate that this information is not needed by SQLPutData().

## Example

The example shown below binds a variety of data types to a set of parameters. For an additional example see Appendix E, "Example code" on page 469.

```

/* ... */
SQLCHAR          stmt[] =
"INSERT INTO PRODUCT VALUES (?, ?, ?, ?, ?)";

SQLINTEGER       Prod_Num[NUM_PRODS] = {
    100110, 100120, 100210, 100220, 100510, 100520, 200110,
    200120, 200210, 200220, 200510, 200610, 990110, 990120,
    500110, 500210, 300100
};

SQLCHAR          Description[NUM_PRODS][257] = {
    "Aquarium-Glass-25 litres", "Aquarium-Glass-50 litres",
    "Aquarium-Acrylic-25 litres", "Aquarium-Acrylic-50 litres",
    "Aquarium-Stand-Small", "Aquarium-Stand-Large",
    "Pump-Basic-25 litre", "Pump-Basic-50 litre",
    "Pump-Deluxe-25 litre", "Pump-Deluxe-50 litre",
    "Pump-Filter-(for Basic Pump)",
    "Pump-Filter-(for Deluxe Pump)",
    "Aquarium-Kit-Small", "Aquarium-Kit-Large",
    "Gravel-Colored", "Fish-Food-Deluxe-Bulk",
    "Plastic-Tubing"
};

SQLDOUBLE        UPrice[NUM_PRODS] = {
    110.00, 190.00, 100.00, 150.00, 60.00, 90.00, 30.00,
    45.00, 55.00, 75.00, 4.75, 5.25, 160.00, 240.00,
    2.50, 35.00, 5.50
};

SQLCHAR          Units[NUM_PRODS][3] = {
    " ", " ", " ", " ", " ", " ", " ", " ", " ", " ", " ", " ",
    " ", " ", " ", " ", " ", " ", "kg", "kg", "m"
};

SQLCHAR          Combo[NUM_PRODS][2] = {
    "N", "N", "N", "N", "N", "N", "N", "N", "N",
    "N", "N", "N", "Y", "Y", "N", "N", "N"
};

SQLINTEGER       pirow = 0;
/* ... */

```

```

/* Prepare the statement */
rc = SQLPrepare(hstmt, stmt, SQL_NTS);

rc = SQLBindParameter(hstmt, 1, SQL_PARAM_INPUT, SQL_C_SLONG, SQL_INTEGER,
                      0, 0, Prod_Num, 0, NULL);

rc = SQLBindParameter(hstmt, 2, SQL_PARAM_INPUT, SQL_C_CHAR, SQL_VARCHAR,
                      257, 0, Description, 257, NULL);

rc = SQLBindParameter(hstmt, 3, SQL_PARAM_INPUT, SQL_C_DOUBLE, SQL_DECIMAL,
                      10, 2, UPrice, 0, NULL);

rc = SQLBindParameter(hstmt, 4, SQL_PARAM_INPUT, SQL_C_CHAR, SQL_CHAR,
                      3, 0, Units, 3, NULL);

rc = SQLBindParameter(hstmt, 5, SQL_PARAM_INPUT, SQL_C_CHAR, SQL_CHAR,
                      2, 0, Combo, 2, NULL);

rc = SQLParamOptions(hstmt, NUM_PRODS, &pirow);

rc = SQLExecute(hstmt);
printf("Inserted %ld Rows\n", pirow);
/* ... */

```

## References

- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLExecute - Execute a statement” on page 159
- “SQLParamData - Get next parameter for which a data value is needed” on page 273
- “SQLParamOptions - Specify an input array for a parameter” on page 275
- “SQLPutData - Passing data value for a parameter” on page 304

## SQLCancel - Cancel statement

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLCancel () can be used to prematurely terminate the *data-at-execution* sequence described in “Sending/retrieving long data in pieces” on page 373.

### Syntax

```
SQLRETURN SQLCancel (SQLHSTMT hstmt);
```

### Function arguments

Table 21. SQLCancel arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle

### Usage

After SQLExecDirect () or SQLExecute () returns SQL\_NEED\_DATA to solicit for values for data-at-execution parameters, SQLCancel () can be used to cancel the data-at-execution sequence described in “Sending/retrieving long data in pieces” on page 373. SQLCancel () can be called any time before the final SQLParamData () in the sequence. After the cancellation of this sequence, the application can call SQLExecute () or SQLExecDirect () to re-initiate the data-at-execution sequence.

If an application calls SQLCancel () on an hstmt not associated with a data-at-execution sequence, SQLCancel () has the same effect as SQLFreeStmt () with the SQL\_CLOSE option. Applications should not call SQLCancel () to close a cursor; but rather SQLFreeStmt () should be used.

### Return codes

- SQL\_SUCCESS
- SQL\_INVALID\_HANDLE
- SQL\_ERROR

### Diagnostics

Table 22. SQLCancel SQLSTATEs

SQLSTATE	Description	Explanation
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.

## Restrictions

DB2 ODBC does not support asynchronous statement execution.

## Example

See “Example” on page 306.

## References

- “SQLPutData - Passing data value for a parameter” on page 304
- “SQLParamData - Get next parameter for which a data value is needed” on page 273

## SQLColAttributes - Get column attributes

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI <sup>1</sup>	ISO CLI <sup>1</sup>
----------------	----------	-------------------------	----------------------

SQLColAttributes() is used to get an attribute for a column of the result set, and can also be used to determine the number of columns. SQLColAttributes() is a more extensible alternative to the SQLDescribeCol() function, but can only return one attribute per call.

Either SQLPrepare() or SQLExecDirect() must be called before calling this function.

If the application does not know the various attributes (such as, data type and length) of the column, this function (or SQLDescribeCol()) must be called before binding, using SQLBindCol(), to any columns.

**Note: 1** - X/Open and ISO define this function with a singular name, SQLColAttribute().

### Syntax

```
SQLRETURN SQLColAttributes (SQLHSTMT      hstmt,
                            SQLUSMALLINT   icol,
                            SQLUSMALLINT   fDescType,
                            SQLPOINTER     rgbDesc,
                            SQLSMALLINT    cbDescMax,
                            SQLSMALLINT FAR pcbDesc,
                            SQLINTEGER FAR  pfDesc);
```

### Function arguments

Table 23. SQLColAttributes arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle
SQLUSMALLINT	<i>icol</i>	input	Column number in result set (must be between 1 and the number of columns in the results set inclusive). This argument is ignored when SQL_COLUMN_COUNT is specified.
SQLUSMALLINT	<i>fDescType</i>	input	The supported values are described in Table 24 on page 107.
SQLCHAR *	<i>rgbDesc</i>	output	Pointer to buffer for string column attributes
SQLSMALLINT	<i>cbDescMax</i>	input	Length of <i>rgbDesc</i> descriptor buffer.
SQLSMALLINT *	<i>pcbDesc</i>	output	Actual number of bytes returned in <i>rgbDesc</i> buffer. If this argument contains a value equal to or greater than the length specified in <i>cbDescMax</i> , truncation occurred. The column attribute value is then truncated to <i>cbDescMax</i> bytes minus the size of the null-terminator (or to <i>cbDescMax</i> bytes if null termination is off).
SQLINTEGER *	<i>pfDesc</i>	output	Pointer to integer which holds the value of numeric column attributes.



The following values can be specified for the *fDescType* argument:

Table 24 (Page 1 of 3). *fdesc*type descriptor types

Descriptor	Description
SQL_COLUMN_AUTO_INCREMENT	<p>Indicates if the column data type is an auto increment data type.</p> <p>FALSE is returned in <i>pfDesc</i> for all DB2 SQL data types.</p>
SQL_COLUMN_CASE_SENSITIVE	<p>Indicates if the column data type is a case sensitive data type.</p> <p>Either TRUE or FALSE is returned in <i>pfDesc</i> depending on the data type.</p> <p>Case sensitivity does not apply to graphic data types, FALSE is returned.</p> <p>FALSE is returned for non-character data types.</p>
SQL_COLUMN_CATALOG_NAME (SQL_COLUMN_QUALIFIER_NAME)	<p>The catalog of the table that contains the column is returned in <i>rgbDesc</i>. An empty string is returned since DB2 ODBC only supports two-part naming for a table.</p> <p>SQL_COLUMN_QUALIFIER_NAME is defined for compatibility with ODBC. DB2 ODBC applications should use SQL_COLUMN_CATALOG_NAME.</p>
SQL_COLUMN_COUNT	<p>The number of columns in the result set is returned in <i>pfDesc</i>.</p>
SQL_COLUMN_DISPLAY_SIZE	<p>The maximum number of bytes needed to display the data in character form is returned in <i>pfDesc</i>.</p> <p>See Table 154 on page 448 for the display size of each of the column types.</p>
SQL_COLUMN_DISTINCT_TYPE	<p>The distinct type data type name of the column is returned in <i>rgbDesc</i>. If the column is a built-in SQL type and not a distinct type, an empty string is returned.</p> <p><b>Note:</b> This is an IBM-defined extension to the list of descriptor attributes defined by ODBC.</p>
SQL_COLUMN_LABEL	<p>The column label is returned in <i>rgbDesc</i>. If the column does not have a label, the column name or the column expression is returned. If the column is unlabeled and unnamed, an empty string is returned.</p>

Table 24 (Page 2 of 3). *fdesc* type descriptor types

Descriptor	Description
SQL_COLUMN_LENGTH	<p>The number of <i>bytes</i> of data associated with the column is returned in <i>pfDesc</i>. This is the length in bytes of data transferred on the fetch or <code>SQLGetData()</code> for this column if <code>SQL_C_DEFAULT</code> is specified as the C data type. See Table 153 on page 447 for the length of each of the SQL data types.</p> <p>If the column identified in <i>icol</i> is a fixed length character or binary string, (for example, <code>SQL_CHAR</code> or <code>SQL_BINARY</code>) the actual length is returned.</p> <p>If the column identified in <i>icol</i> is a variable length character or binary string, (for example, <code>SQL_VARCHAR</code>) the maximum length is returned.</p>
SQL_COLUMN_MONEY	<p>Indicates if the column data type is a money data type.</p> <p>FALSE is returned in <i>pfDesc</i> for all DB2 SQL data types.</p>
SQL_COLUMN_NAME	<p>The name of the column <i>icol</i> is returned in <i>rgbDesc</i>. If the column is an expression, then the result returned is product specific.</p>
SQL_COLUMN_NULLABLE	<p>If the column identified by <i>icol</i> can contain nulls, then <code>SQL_NULLABLE</code> is returned in <i>pfDesc</i>.</p> <p>If the column is constrained not to accept nulls, then <code>SQL_NO_NULLS</code> is returned in <i>pfDesc</i>.</p>
SQL_COLUMN_PRECISION	<p>The precision in units of digits is returned in <i>pfDesc</i> if the column is <code>SQL_DECIMAL</code>, <code>SQL_NUMERIC</code>, <code>SQL_DOUBLE</code>, <code>SQL_FLOAT</code>, <code>SQL_INTEGER</code>, <code>SQL_REAL</code> or <code>SQL_SMALLINT</code>.</p> <p>If the column is a character SQL data type, then the precision returned in <i>pfDesc</i>, indicates the maximum number of <i>characters</i> the column can hold.</p> <p>If the column is a graphic SQL data type, then the precision returned in <i>pfDesc</i>, indicates the maximum number of double-byte <i>characters</i> the column can hold.</p> <p>See Table 151 on page 445 for the precision of each of the SQL data types.</p>
SQL_COLUMN_SCALE	<p>The scale attribute of the column is returned. See Table 152 on page 446 for the scale of each of the SQL data types.</p>
SQL_COLUMN_SCHEMA_NAME (SQL_COLUMN_OWNER_NAME)	<p>The schema of the table that contains the column is returned in <i>rgbDesc</i>. An empty string is returned as DB2 ODBC is unable to determine this attribute.</p> <p><code>SQL_COLUMN_OWNER_NAME</code> is defined for compatibility with ODBC. DB2 ODBC applications should use <code>SQL_COLUMN_SCHEMA_NAME</code>.</p>

Table 24 (Page 3 of 3). *fdesc*type descriptor types

Descriptor	Description
SQL_COLUMN_SEARCHABLE	<p>Indicates if the column data type is searchable:</p> <ul style="list-style-type: none"> <li>• SQL_UNSEARCHABLE if the column cannot be used in a WHERE clause.</li> <li>• SQL_LIKE_ONLY if the column can be used in a WHERE clause only with the LIKE predicate.</li> <li>• SQL_ALL_EXCEPT_LIKE if the column can be used in a WHERE clause with all comparison operators except LIKE.</li> <li>• SQL_SEARCHABLE if the column can be used in a WHERE clause with any comparison operator.</li> </ul>
SQL_COLUMN_TABLE_NAME	The name of the table that contains the column is returned in <i>rgbDesc</i> . An empty string is returned as DB2 ODBC cannot determine this attribute.
SQL_COLUMN_TYPE	The SQL data type of the column identified in <i>icol</i> is returned in <i>pfDesc</i> . The possible values for <i>pfSqlType</i> are listed in Table 4 on page 39.
SQL_COLUMN_TYPE_NAME	<p>The type of the column (as entered in an SQL statement) is returned in <i>rgbDesc</i>.</p> <p>For information on each data type see the TYPE_NAME attribute found in “Data types and data conversion” on page 36.</p>
SQL_COLUMN_UNSIGNED	<p>Indicates if the column data type is an unsigned type or not.</p> <p>TRUE is returned in <i>pfDesc</i> for all non-numeric data types, FALSE is returned for all numeric data types.</p>
SQL_COLUMN_UPDATABLE	<p>Indicates if the column data type is an updateable data type.</p> <p>SQL_ATTR_READWRITE_UNKNOWN is returned in <i>pfDesc</i> for all DB2 SQL data types.</p> <p>SQL_ATTR_READONLY is returned if the column is obtained from a catalog function call.</p>

## Usage

Instead of returning a specific set of attributes like `SQLDescribeCol()`, `SQLColAttributes()` allows you to specify which attribute you wish to receive for a specific column. If the desired information is a string, it is returned in *rgbDesc*. If the desired information is a number, it is returned in *pfDesc*.

`SQLColAttributes()` is an extensible alternative to `SQLDescribeCol()`, which is used to return a fixed set of commonly used column attribute information.

If an *fDescType* descriptor type does not apply to the database server, an empty string is returned in *rgbDesc* or zero is returned in *pfDesc*, depending on the expected result of the descriptor.

Columns are identified by a number (numbered sequentially from left to right starting with 1) and can be described in any order.

## SQLColAttributes

Calling `SQLColAttributes()` with *fDescType* set to `SQL_COLUMN_COUNT` is an alternative to calling `SQLNumResultCols()` to determine whether any columns can be returned.

## Return codes

- `SQL_SUCCESS`
- `SQL_SUCCESS_WITH_INFO`
- `SQL_ERROR`
- `SQL_INVALID_HANDLE`

## Diagnostics

Table 25. *SQLColAttributes* SQLSTATEs

SQLSTATE	Description	Explanation
01004	Data truncated.	The character string returned in the argument <i>rgbDesc</i> is longer than the value specified in the argument <i>cbDescMax</i> . The argument <i>pcbDesc</i> contains the actual length of the string to be returned. (Function returns <code>SQL_SUCCESS_WITH_INFO</code> .)
07005	The statement did not return a result set.	The statement associated with the <i>hstmt</i> did not return a result set. There are no columns to describe.  To prevent encountering this error, call <code>SQLNumResultCols()</code> before calling <code>SQLColAttributes()</code> .
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1002	Invalid column number.	The value specified for the argument <i>icol</i> is less than 1.  The value specified for the argument <i>icol</i> is greater than the number of columns in the result set. Not returned if <code>SQL_COLUMN_COUNT</code> is specified.
S1010	Function sequence error.	The function is called prior to calling <code>SQLPrepare()</code> or <code>SQLExecDirect()</code> for the <i>hstmt</i> .  The function is called while in a data-at-execute ( <code>SQLParamData()</code> , <code>SQLPutData()</code> ) operation.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1090	Invalid string or buffer length.	The length specified in the argument <i>cbDescMax</i> is less than 0 and <i>fDescType</i> requires a character string be returned in <i>rgbDesc</i> .
S1091	Descriptor type out of range.	The value specified for the argument <i>fDescType</i> was not equal to a value specified in Table 24 on page 107.
S1C00	Driver not capable.	The SQL data type returned by the database server for column <i>icol</i> is not recognized by DB2 ODBC.

## Restrictions

None.

## Example

See “Example” on page 135

## References

- “SQLDescribeCol - Describe column attributes” on page 132
- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLPrepare - Prepare a statement” on page 277
- “SQLSetColAttributes - Set column attributes” on page 309

## SQLColumnPrivileges - Get privileges associated with the columns of a table

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>		
-----------------------	-----------------	--	--

SQLColumnPrivileges() returns a list of columns and associated privileges for the specified table. The information is returned in an SQL result set, which can be retrieved using the same functions that are used to process a result set generated from a query.

### Syntax

```
SQLRETURN SQLColumnPrivileges (SQLHSTMT      hstmt,
                                SQLCHAR      FAR *szCatalogName,
                                SQLSMALLINT   cbCatalogName,
                                SQLCHAR      FAR *szSchemaName,
                                SQLSMALLINT   cbSchemaName,
                                SQLCHAR      FAR *szTableName,
                                SQLSMALLINT   cbTableName,
                                SQLCHAR      FAR *szColumnName,
                                SQLSMALLINT   cbColumnName);
```

### Function arguments

Table 26. SQLColumnPrivileges arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle.
SQLCHAR *	szCatalogName	input	Catalog qualifier of a 3-part table name. This must be a NULL pointer or a zero length string.
SQLSMALLINT	cbCatalogName	input	Length of <i>szCatalogName</i> . This must be set to 0.
SQLCHAR *	szSchemaName	input	Schema qualifier of table name.
SQLSMALLINT	cbSchemaName	input	Length of <i>szSchemaName</i> .
SQLCHAR *	szTableName	input	Table name.
SQLSMALLINT	cbTableName	input	Length of <i>szTableName</i> .
SQLCHAR *	szColumnName	input	Buffer that can contain a <i>pattern-value</i> to qualify the result set by column name.
SQLSMALLINT	cbColumnName	input	Length of <i>szColumnName</i> .

### Usage

The results are returned as a standard result set containing the columns listed in Table 27 on page 113. The result set is ordered by TABLE\_CAT, TABLE\_SCHEM, TABLE\_NAME, COLUMN\_NAME, and PRIVILEGE. If multiple privileges are associated with any given column, each privilege is returned as a separate row. A typical application might wish to call this function after a call to SQLColumns() to determine column privilege information. The application should use the character

strings returned in the TABLE\_SCHEM, TABLE\_NAME, COLUMN\_NAME columns of the SQLColumns() result set as input arguments to this function.

Since calls to SQLColumnPrivileges() in many cases map to a complex and thus expensive query against the system catalog, they should be used sparingly, and the results saved rather than repeating the calls.

The VARCHAR columns of the catalog functions result set are declared with a maximum length attribute of 128 to be consistent with SQL92 limits. Since DB2 names are less than 128, the application can choose to always set aside 128 characters (plus the null-terminator) for the output buffer, or alternatively, call SQLGetInfo() with the SQL\_MAX\_CATALOG\_NAME\_LEN, SQL\_MAX\_SCHEMA\_NAME\_LEN, SQL\_MAX\_TABLE\_NAME\_LEN, and SQL\_MAX\_COLUMN\_NAME\_LEN to determine respectively the actual lengths of the TABLE\_CAT, TABLE\_SCHEM, TABLE\_NAME, and COLUMN\_NAME columns supported by the connected DBMS.

Note that the *szColumnName* argument accepts a search pattern. For more information about valid search patterns, see “Input arguments on catalog functions” on page 369.

Although new columns might be added and the names of the existing columns changed in future releases, the position of the current columns does not change.

Table 27 (Page 1 of 2). Columns returned by SQLColumnPrivileges

Column number/name	Data type	Description
1 TABLE_CAT	VARCHAR(128)	This is always NULL.
2 TABLE_SCHEM	VARCHAR(128)	The name of the schema containing TABLE_NAME.
3 TABLE_NAME	VARCHAR(128) not NULL	Name of the table or view.
4 COLUMN_NAME	VARCHAR(128) not NULL	Name of the column of the specified table or view.
5 GRANTOR	VARCHAR(128)	Authorization ID of the user who granted the privilege.
6 GRANTEE	VARCHAR(128)	Authorization ID of the user to whom the privilege is granted.

## SQLColumnPrivileges

Table 27 (Page 2 of 2). Columns returned by SQLColumnPrivileges

Column number/name	Data type	Description
7 PRIVILEGE	VARCHAR(128)	<p>The column privilege. This can be:</p> <ul style="list-style-type: none"><li>• ALTER</li><li>• CONTROL</li><li>• DELETE</li><li>• INDEX</li><li>• INSERT</li><li>• REFERENCES</li><li>• SELECT</li><li>• UPDATE</li></ul> <p>Supported privileges are based on the data source to which you are connected.</p> <p><b>Note:</b> Most IBM RDBMSs do not offer column level privileges at the column level. DB2 for OS/390 and DB2 for VSE &amp; VM support the UPDATE column privilege; there is one row in this result set for each updateable column. For all other privileges for DB2 for OS/390 and DB2 for VSE &amp; VM, and for all privileges for other IBM RDBMSs, if a privilege has been granted at the table level, a row is present in this result set.</p>
8 IS_GRANTABLE	VARCHAR(3)	<p>Indicates whether the grantee is permitted to grant the privilege to other users.</p> <p>Either "YES", "NO".</p>

**Note:** The column names used by DB2 ODBC follow the X/Open CLI CAE specification style. The column types, contents and order are identical to those defined for the SQLColumnPrivileges() result set in ODBC.

If there is more than one privilege associated with a column, then each privilege is returned as a separate row in the result set.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 28 (Page 1 of 2). SQLColumnPrivileges SQLSTATES

SQLSTATE	Description	Explanation
24000	Invalid cursor state.	A cursor is already opened on the statement handle.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	<i>szTableName</i> is NULL.
S1014	No more handles.	DB2 ODBC is not able to allocate a handle due to internal resources.



Table 28 (Page 2 of 2). SQLColumnPrivileges SQLSTATES

SQLSTATE	Description	Explanation
S1090	Invalid string or buffer length.	The value of one of the name length arguments is less than 0, but not equal to SQL_NTS.
S1C00	Driver not capable.	DB2 ODBC does not support <i>catalog</i> as a qualifier for table name.

## Restrictions

None.

## Example

```

/* ... */
SQLRETURN
list_column_privileges(SQLHDBC hdbc, SQLCHAR *schema, SQLCHAR *tablename )
{
/* ... */

    rc = SQLColumnPrivileges(hstmt, NULL, 0, schema, SQL_NTS,
                             tablename, SQL_NTS, columnname.s, SQL_NTS);

    rc = SQLBindCol(hstmt, 4, SQL_C_CHAR, (SQLPOINTER) columnname.s, 129,
                   &columnname.ind);

    rc = SQLBindCol(hstmt, 5, SQL_C_CHAR, (SQLPOINTER) grantor.s, 129,
                   &grantor.ind);

    rc = SQLBindCol(hstmt, 6, SQL_C_CHAR, (SQLPOINTER) grantee.s, 129,
                   &grantee.ind);

    rc = SQLBindCol(hstmt, 7, SQL_C_CHAR, (SQLPOINTER) privilege.s, 129,
                   &privilege.ind);

    rc = SQLBindCol(hstmt, 8, SQL_C_CHAR, (SQLPOINTER) is_grantable.s, 4,
                   &is_grantable.ind);

    printf("Column Privileges for %s.%s\n", schema, tablename);
    /* Fetch each row, and display */
    while ((rc = SQLFetch(hstmt)) == SQL_SUCCESS) {
        sprintf(cur_name, " Column: %s\n", columnname.s);
        if (strcmp(cur_name, pre_name) != 0) {
            printf("\n%s\n", cur_name);
            printf("  Grantor      Grantee      Privilege Grantable\n");
            printf("  -----      -----      ----- ---\n");
        }
        strcpy(pre_name, cur_name);
        printf("    %-15s", grantor.s);
        printf("    %-15s", grantee.s);
        printf("    %-10s", privilege.s);
        printf("    %-3s\n", is_grantable.s);
    }
    /* endwhile */
}
/* ... */

```

### References

- “SQLColumns - Get column information for a table” on page 117
- “SQLTables - Get table information” on page 353

## SQLColumns - Get column information for a table

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	
----------------	----------	------------	--

SQLColumns() returns a list of columns in the specified tables. The information is returned in an SQL result set, which can be retrieved using the same functions that are used to fetch a result set generated by a query.

### Syntax

```
SQLRETURN SQLColumns(
    (SQLHSTMT hstmt,
    SQLCHAR FAR *szCatalogName,
    SQLSMALLINT cbCatalogName,
    SQLCHAR FAR *szSchemaName,
    SQLSMALLINT cbSchemaName,
    SQLCHAR FAR *szTableName,
    SQLSMALLINT cbTableName,
    SQLCHAR FAR *szColumnName,
    SQLSMALLINT cbColumnName);
```

### Function arguments

Table 29. SQLColumns arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle.
SQLCHAR *	szCatalogName	input	Buffer that can contain a <i>pattern-value</i> to qualify the result set. <i>Catalog</i> is the first part of a 3-part table name. This must be a NULL pointer or a zero length string.
SQLSMALLINT	cbCatalogName	input	Length of <i>szCatalogName</i> . This must be set to 0.
SQLCHAR *	szSchemaName	input	Buffer that can contain a <i>pattern-value</i> to qualify the result set by schema name.
SQLSMALLINT	cbSchemaName	input	Length of <i>szSchemaName</i> .
SQLCHAR *	szTableName	input	Buffer that can contain a <i>pattern-value</i> to qualify the result set by table name.
SQLSMALLINT	cbTableName	input	Length of <i>szTableName</i> .
SQLCHAR *	szColumnName	input	Buffer that can contain a <i>pattern-value</i> to qualify the result set by column name.
SQLSMALLINT	cbColumnName	input	Length of <i>szColumnName</i> .

### Usage

This function is called to retrieve information about the columns of either a table or a set of tables. A typical application might wish to call this function after a call to SQLTables() to determine the columns of a table. The application should use the character strings returned in the TABLE\_SCHEMA and TABLE\_NAME columns of the SQLTables() result set as input to this function.

## SQLColumns

SQLColumns() returns a standard result set, ordered by TABLE\_CAT, TABLE\_SCHEM, TABLE\_NAME, and ORDINAL\_POSITION. Table 30 on page 118 lists the columns in the result set.

The *szSchemaName*, *szTableName*, and *szColumnName* arguments accept search patterns. For more information about valid search patterns, see “Input arguments on catalog functions” on page 369.

Since calls to SQLColumns() in many cases map to a complex and thus expensive query against the system catalog, they should be used sparingly, and the results saved rather than repeating calls.

The VARCHAR columns of the catalog functions result set are declared with a maximum length attribute of 128 to be consistent with SQL92 limits. Since DB2 names are less than 128, the application can choose to always set aside 128 characters (plus the null-terminator) for the output buffer, or alternatively, call SQLGetInfo() with the SQL\_MAX\_CATALOG\_NAME\_LEN, SQL\_MAX\_OWNER\_SCHEMA\_LEN, SQL\_MAX\_TABLE\_NAME\_LEN, and SQL\_MAX\_COLUMN\_NAME\_LEN to determine respectively the actual lengths of the TABLE\_CAT, TABLE\_SCHEM, TABLE\_NAME, and COLUMN\_NAME columns supported by the connected DBMS.

Although new columns might be added and the names of the existing columns changed in future releases, the position of the current columns does not change.

Table 30 (Page 1 of 3). Columns returned by SQLColumns

Column number/name	Data type	Description
1 TABLE_CAT	VARCHAR(128)	This is always NULL.
2 TABLE_SCHEM	VARCHAR(128)	The name of the schema containing TABLE_NAME.
3 TABLE_NAME	VARCHAR(128) NOT NULL	Name of the table, view, alias, or synonym.
4 COLUMN_NAME	VARCHAR(128) NOT NULL	Column identifier. Name of the column of the specified table, view, alias, or synonym.
5 DATA_TYPE	SMALLINT NOT NULL	SQL data type of column identified by COLUMN_NAME. This is one of the values in the Symbolic SQL Data Type column in Table 4 on page 39.
6 TYPE_NAME	VARCHAR(128) NOT NULL	Character string representing the name of the data type corresponding to DATA_TYPE.
7 COLUMN_SIZE	INTEGER	If the DATA_TYPE column value denotes a character or binary string, then this column contains the maximum length in characters for the column.  For date, time, timestamp data types, this is the total number of characters required to display the value when converted to character.  For numeric data types, this is either the total number of digits, or the total number of bits allowed in the column, depending on the value in the NUM_PREC_RADIX column in the result set.  See also, Table 151 on page 445.

Table 30 (Page 2 of 3). Columns returned by SQLColumns

Column number/name	Data type	Description
8 BUFFER_LENGTH	INTEGER	<p>The maximum number of bytes for the associated C buffer to store data from this column if SQL_C_DEFAULT is specified on the SQLBindCol(), SQLGetData() and SQLBindParameter() calls. This length does not include any null-terminator. For exact numeric data types, the length accounts for the decimal and the sign.</p> <p>See also, Table 153 on page 447.</p>
9 DECIMAL_DIGITS	SMALLINT	<p>The scale of the column. NULL is returned for data types where scale is not applicable.</p> <p>See also, Table 152 on page 446.</p>
10 NUM_PREC_RADIX	SMALLINT	<p>Either 10 or 2 or NULL. If DATA_TYPE is an approximate numeric data type, this column contains the value 2, then the COLUMN_SIZE column contains the number of bits allowed in the column.</p> <p>If DATA_TYPE is an exact numeric data type, this column contains the value 10 and the COLUMN_SIZE contains the number of decimal digits allowed for the column.</p> <p>For numeric data types, the DBMS can return a NUM_PREC_RADIX of either 10 or 2.</p> <p>NULL is returned for data types where radix is not applicable.</p>
11 NULLABLE	SMALLINT NOT NULL	<p>SQL_NO_NULLS if the column does not accept NULL values.</p> <p>SQL_NULLABLE if the column accepts NULL values.</p>
12 REMARKS	VARCHAR(254)	Might contain descriptive information about the column.
13 COLUMN_DEF	VARCHAR(254)	<p>The column's default value. If the default value is a numeric literal, then this column contains the character representation of the numeric literal with no enclosing single quotes. If the default value is a character string, then this column is that string enclosed in single quotes. If the default value is a <i>pseudo-literal</i>, such as for DATE, TIME, and TIMESTAMP columns, then this column contains the keyword of the pseudo-literal (for example, CURRENT DATE) with no enclosing quotes.</p> <p>If NULL was specified as the default value, then this column returns the word NULL, not enclosed in quotes. If the default value cannot be represented without truncation, then this column contains TRUNCATED with no enclosing single quotes. If no default value was specified, then this column is NULL.</p>
14 SQL_DATA_TYPE	SMALLINT NOT NULL	<p>The SQL data type. This column is the same as the DATA_TYPE column. For datetime data types, the SQL_DATA_TYPE field in the result set is SQL_DATETIME, and the SQL_DATETIME_SUB field returns the subcode for the specific datetime data type (SQL_CODE_DATE, SQL_CODE_TIME or SQL_CODE_TIMESTAMP).</p>

## SQLColumns

Table 30 (Page 3 of 3). Columns returned by SQLColumns

Column number/name	Data type	Description
15 SQL_DATETIME_SUB	SMALLINT	The subtype code for datetime data types: <ul style="list-style-type: none"><li>• SQL_CODE_DATE</li><li>• SQL_CODE_TIME</li><li>• SQL_CODE_TIMESTAMP</li></ul> For all other data types, this column returns a NULL.
16 CHAR_OCTET_LENGTH	INTEGER	Contains the maximum length in octets for a character data type column. For Single Byte character sets, this is the same as COLUMN_SIZE. For all other data types it is NULL.
17 ORDINAL_POSITION	INTEGER NOT NULL	The ordinal position of the column in the table. The first column in the table is number 1.
18 IS_NULLABLE	VARCHAR(254)	Contains the string 'NO' if the column is known to be not nullable; and 'YES' otherwise.

**Note:** This result set is identical to the X/Open CLI Columns() result set specification, which is an extended version of the SQLColumns() result set specified in ODBC V2. The ODBC SQLColumns() result set includes every column in the same position up to the REMARKS column.

# DB2 ODBC applications that issue SQLColumns() against a DB2 for OS/390 server, Version 5 or later, should  
# expect the result set columns listed in the table above. Revision bars identify the new and changed columns.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 31. SQLColumns SQLSTATES

SQLSTATE	Description	Explanation
24000	Invalid cursor state.	A cursor is already opened on the statement handle.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1014	No more handles.	DB2 ODBC is not able to allocate a handle due to internal resources.
S1090	Invalid string or buffer length.	The value of one of the name length arguments is less than 0, but not equal SQL_NTS.
S1C00	Driver not capable.	DB2 ODBC does not support <i>catalog</i> as a qualifier for table name.

## Restrictions

None.

## Example

```

/* ... */
SQLRETURN
list_columns(SQLHDBC hdbc, SQLCHAR *schema, SQLCHAR *tablename )
{
/* ... */

    rc = SQLColumns(hstmt, NULL, 0, schema, SQL_NTS,
                   tablename, SQL_NTS, "%", SQL_NTS);

    rc = SQLBindCol(hstmt, 4, SQL_C_CHAR, (SQLPOINTER) column_name.s, 129,
                   &column_name.ind);

    rc = SQLBindCol(hstmt, 6, SQL_C_CHAR, (SQLPOINTER) type_name.s, 129,
                   &type_name.ind);

    rc = SQLBindCol(hstmt, 7, SQL_C_LONG, (SQLPOINTER) &length,
                   sizeof(length), &length_ind);

    rc = SQLBindCol(hstmt, 9, SQL_C_SHORT, (SQLPOINTER) &scale,
                   sizeof(scale), &scale_ind);

    rc = SQLBindCol(hstmt, 12, SQL_C_CHAR, (SQLPOINTER) remarks.s, 129,
                   &remarks.ind);

    rc = SQLBindCol(hstmt, 11, SQL_C_SHORT, (SQLPOINTER) &nullable,
                   sizeof(nullable), &nullable_ind);

    printf("Schema: %s Table Name: %s\n", schema, tablename);
    /* Fetch each row, and display */
    while ((rc = SQLFetch(hstmt)) == SQL_SUCCESS) {
        printf("  %s", column_name.s);
        if (nullable == SQL_NULLABLE) {
            printf(", NULLABLE");
        } else {
            printf(", NOT NULLABLE");
        }
        printf(", %s", type_name.s);
        if (length_ind != SQL_NULL_DATA) {
            printf(" (%ld", length);
        } else {
            printf("\n");
        }
        if (scale_ind != SQL_NULL_DATA) {
            printf(", %d)\n", scale);
        } else {
            printf("\n");
        }
    }
    /* endwhile */
}
/* ... */

```

## SQLColumns

## References

- “SQLTables - Get table information” on page 353
- “SQLColumnPrivileges - Get privileges associated with the columns of a table” on page 112
- “SQLSpecialColumns - Get special (row identifier) columns” on page 339



## SQLConnect - Connect to a data source

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLConnect() establishes a connection to the target database. The application must supply a target SQL database.

SQLAllocConnect() must be called before calling this function.

This function must be called before calling SQLAllocStmt().

### Syntax

```
SQLRETURN SQLConnect(
    (SQLHDBC          hdbc,
     SQLCHAR FAR     *szDSN,
     SQLSMALLINT     cbDSN,
     SQLCHAR FAR     *szUID,
     SQLSMALLINT     cbUID,
     SQLCHAR FAR     *szAuthStr,
     SQLSMALLINT     cbAuthStr);
```

### Function arguments

Table 32. SQLConnect arguments

Data type	Argument	Use	Description
SQLHDBC	<i>hdbc</i>	input	Connection handle
SQLCHAR *	<i>szDSN</i>	input	Data Source: The name or alias-name of the database.
SQLSMALLINT	<i>cbDSN</i>	input	Length of contents of <i>szDSN</i> argument
SQLCHAR *	<i>szUID</i>	input	Authorization-name (user identifier). This parameter is validated but not used.
SQLSMALLINT	<i>cbUID</i>	input	Length of contents of <i>szUID</i> argument. This parameter is validated but not used.
SQLCHAR *	<i>szAuthStr</i>	input	Authentication-string (password). This parameter is validated but not used.
SQLSMALLINT	<i>cbAuthStr</i>	input	Length of contents of <i>szAuthStr</i> argument. This parameter is validated but not used.

### Usage

#  
#  
#  
#  
#  
#  
#  
#  
#

The target database (also known as *data source*) for IBM RDBMSs is the location name as defined in SYSIBM.LOCATIONS when DDF has been configured in the DB2 subsystem. The application can obtain a list of databases available to connect to by calling SQLDataSources(). For many applications, a local database is being accessed (DDF is not being used). The local database name is the name that was set during DB2 installation as 'DB2 LOCATION NAME' on the DSNTIPR installation panel for the DB2 subsystem. Your local DB2 administration staff can provide you with this name, or the application can use a 'null connect', as described below, to connect to the default local database without supplying a database name.

The input length arguments to `SQLConnect()` (*cbDSN*, *cbUID*, *cbAuthStr*) can be set to the actual length of their associated data (not including any null-terminating character) or to `SQL_NTS` to indicate that the associated data is null-terminated.

The *szuid* and *szauthstr* argument values must not contain any blanks. If these values are specified, they are ignored. The semantics of *szDSN* are as follows:

- If the *szDSN* argument pointer is NULL or the *cbDSN* argument value is 0, this is a null `SQLConnect()`. (The function performed by a null `SQLConnect()` is referred to as a 'null connect'.)

A null `SQLConnect` still requires that `SQLAllocEnv()` and `SQLAllocConnect()` be called first. Reasons for coding a null `SQLConnect` include:

- The DB2 ODBC application needs to connect to the default data source. (The default data source is the DB2 subsystem specified by the `MVSDEFAULTSSID` initialization file setting.)
  - The DB2 ODBC application is mixing embedded SQL and DB2 ODBC calls, and the application already connected to a data source before invoking DB2 ODBC. In this case, the application must issue a null `SQLConnect()`.
  - The DB2 ODBC application is running as a stored procedure. DB2 ODBC applications running as stored procedures must issue a null `SQLConnect()`.
- If the *szDSN* argument pointer is not NULL and the *cbDSN* argument value is not 0, DB2 ODBC issues a `CONNECT` to the data source.

Use the more extensible `SQLDriverConnect()` function to connect when the application needs to override any or all of the keyword values specified for this data source in the initialization file.

Various connection characteristics (options) can be specified by the end user in the section of the initialization file associated with the *szDSN* data source argument or set by the application using `SQLSetConnectOption()`. The extended connect function, `SQLDriverConnect()`, can be called with additional connect options and can also perform a null connect.

For a null `SQLConnect()`, the `CONNECT` type defaults to the value of the `CONNECTTYPE` keyword specified in the common section of the initialization file. The DB2 ODBC application can override the `CONNECT` type by specifying the parameter of the `SQL_CONNECTTYPE` option using:

- `SQLSetConnectOption()` before the null `SQLConnect()` is issued, or
- `SQLSetEnvAttr` before the null `SQLConnect()` is issued.

## Return codes

- `SQL_SUCCESS`
- `SQL_SUCCESS_WITH_INFO`
- `SQL_ERROR`
- `SQL_INVALID_HANDLE`

## Diagnostics

Table 33. SQLConnect SQLSTATEs

SQLSTATE	Description	Explanation
08001	Unable to connect to data source.	DB2 ODBC is not able to establish a connection with the data source (server).  The connection request is rejected because an existing connection established using embedded SQL already exists.
08002	Connection in use.	The specified <i>hdbc</i> is already being used to establish a connection with a data source and the connection is still open.
08004	The application server rejected establishment of the connection.	The data source (server) rejected the establishment of the connection.  The number of connections specified by the MAXCONN keyword has been reached.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	A non-matching double quote (") is found in either the <i>szDSN</i> , <i>szUID</i> , or <i>szAuthStr</i> argument.
S1090	Invalid string or buffer length.	The value specified for argument <i>cbDSN</i> is less than 0, but not equal to SQL_NTS and the argument <i>szDSN</i> is not a null pointer.  The value specified for argument <i>cbUID</i> is less than 0, but not equal to SQL_NTS and the argument <i>szUID</i> is not a null pointer.  The value specified for argument <i>cbAuthStr</i> is less than 0, but not equal to SQL_NTS and the argument <i>szAuthStr</i> is not a null pointer.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1501	Invalid data source name.	An invalid data source name is specified in argument <i>szDSN</i> .

## Restrictions

The implicit connection (or default database) option for IBM RDBMSs is not supported. SQLConnect() must be called before any SQL statements can be executed.

## Example

## SQLConnect

```
/* ... */
/* Global Variables for user id and password, defined in main module.
   To keep samples simple, not a recommended practice.
   The INIT_UID_PWD macro is used to initialize these variables.
*/
extern  SQLCHAR  server[SQL_MAX_DSN_LENGTH + 1];
/*****/
SQLRETURN
DBconnect(SQLHENV henv,
          SQLHDBC * hdbc)
{
    SQLRETURN      rc;
    SQLSMALLINT    outlen;

    /* allocate a connection handle */
    if (SQLAllocConnect(henv, hdbc) != SQL_SUCCESS) {
        printf(">---ERROR while allocating a connection handle-----\n");
        return (SQL_ERROR);
    }
}
```

```

/* Set AUTOCOMMIT OFF */
rc = SQLSetConnectOption(*hdbc, SQL_AUTOCOMMIT, SQL_AUTOCOMMIT_OFF);
if (rc != SQL_SUCCESS) {
    printf(">---ERROR while setting AUTOCOMMIT OFF -----\n");
    return (SQL_ERROR);
}
rc = SQLConnect(*hdbc, server, SQL_NTS, NULL, SQL_NTS, NULL, SQL_NTS);
if (rc != SQL_SUCCESS) {
    printf(">--- Error while connecting to database: %s -----\n", server);
    SQLDisconnect(*hdbc);
    SQLFreeConnect(*hdbc);
    return (SQL_ERROR);
} else {
    /* Print Connection Information */
    printf(">Connected to %s\n", server);
}
return (SQL_SUCCESS);
}

/*****
/* DBconnect2 - Connect with connect type */
/* Valid connect types SQL_CONCURRENT_TRANS, SQL_COORDINATED_TRANS */
*****/
SQLRETURN DBconnect2(SQLHENV henv,
    SQLHDBC * hdbc, SQLINTEGER contype)
    SQLHDBC * hdbc, SQLINTEGER contype, SQLINTEGER conphase)
{
    SQLRETURN rc;
    SQLSMALLINT outlen;

    /* allocate a connection handle */
    if (SQLAllocConnect(henv, hdbc) != SQL_SUCCESS) {
        printf(">---ERROR while allocating a connection handle-----\n");
        return (SQL_ERROR);
    }
    /* Set AUTOCOMMIT OFF */
    rc = SQLSetConnectOption(*hdbc, SQL_AUTOCOMMIT, SQL_AUTOCOMMIT_OFF);
    if (rc != SQL_SUCCESS) {
        printf(">---ERROR while setting AUTOCOMMIT OFF -----\n");
        return (SQL_ERROR);
    }
    rc = SQLSetConnectOption(hdbc[0], SQL_CONNECTTYPE, contype);
    if (rc != SQL_SUCCESS) {
        printf(">---ERROR while setting Connect Type -----\n");
        return (SQL_ERROR);
    }
    if (contype == SQL_COORDINATED_TRANS ) {
        rc = SQLSetConnectOption(hdbc[0], SQL_SYNC_POINT, conphase);
        if (rc != SQL_SUCCESS) {
            printf(">---ERROR while setting Syncpoint Phase -----\n");
            return (SQL_ERROR);
        }
    }
}

```

## SQLConnect

```
rc = SQLConnect(*hdbc, server, SQL_NTS, NULL, SQL_NTS, NULL, SQL_NTS);
if (rc != SQL_SUCCESS) {
    printf(">--- Error while connecting to database: %s -----\n", server);
    SQLDisconnect(*hdbc);
    SQLFreeConnect(*hdbc);
    return (SQL_ERROR);
} else {
    /* Print Connection Information */
    printf(">Connected to %s\n", server);
}
return (SQL_SUCCESS);
}
/* ... */
```

## References

- “SQLAllocConnect - Allocate connection handle” on page 79
- “SQLDriverConnect - (Expanded) connect to a data source” on page 142
- “SQLSetConnectOption - Set connection option” on page 316
- “SQLGetConnectOption - Returns current setting of a connect option” on page 191
- “SQLAllocStmt - Allocate a statement handle” on page 85
- “SQLDataSources - Get list of data sources” on page 129
- “SQLDisconnect - Disconnect from a data source” on page 140

## SQLDataSources - Get list of data sources

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLDataSources() returns a list of target databases available, one at a time.

SQLDataSources() is usually called before a connection is made, to determine the databases that are available to connect to.

### Syntax

```
SQLRETURN SQLDataSources (SQLHENV          henv,
                          SQLUSMALLINT     fDirection,
                          SQLCHAR FAR      *szDSN,
                          SQLSMALLINT      cbDSNMax,
                          SQLSMALLINT FAR  *pcbDSN,
                          SQLCHAR FAR      *szDescription,
                          SQLSMALLINT      cbDescriptionMax,
                          SQLSMALLINT FAR  *pcbDescription);
```

### Function arguments

Table 34. SQLDataSources arguments

Data Type	Argument	Use	Description
SQLHENV	<i>henv</i>	input	Environment handle.
SQLUSMALLINT	<i>fDirection</i>	input	Used by application to request the first data source name in the list or the next one in the list. <i>fDirection</i> can take on only the following values: <ul style="list-style-type: none"> <li>SQL_FETCH_FIRST</li> <li>SQL_FETCH_NEXT</li> </ul>
SQLCHAR *	<i>szDSN</i>	output	Pointer to buffer to hold the data source name retrieved.
SQLSMALLINT	<i>cbDSNMax</i>	input	Maximum length of the buffer pointed to by <i>szDSN</i> . This should be less than or equal to SQL_MAX_DSN_LENGTH + 1.
SQLSMALLINT *	<i>pcbDSN</i>	output	Pointer to location where the maximum number of bytes available to return in the <i>szDSN</i> are stored.
SQLCHAR *	<i>szDescription</i>	output	Pointer to buffer where the description of the data source is returned. DB2 ODBC returns the <b>Comment</b> field associated with the database cataloged to the DBMS.

**Note:** IBM RDBMSs always return blank padded to 30 bytes.

SQLSMALLINT	<i>cbDescriptionMax</i>	input	Maximum length of the <i>szDescription</i> buffer. DB2 for OS/390 always returns NULL.
SQLSMALLINT *	<i>pcbDescription</i>	output	Pointer to location where this function returns the actual number of bytes available to return for the description of the data source. DB2 for OS/390 always returns zero.

## Usage

The application can call this function any time with *fDirection* set to either SQL\_FETCH\_FIRST or SQL\_FETCH\_NEXT.

If SQL\_FETCH\_FIRST is specified, the first database in the list is always returned.

If SQL\_FETCH\_NEXT is specified:

- Directly following a SQL\_FETCH\_FIRST call, the second database in the list is returned
- Before any other SQLDataSources() call, the first database in the list is returned
- When there are no more databases in the list, SQL\_NO\_DATA\_FOUND is returned. If the function is called again, the first database is returned.
- Any other time, the next database in the list is returned.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE
- SQL\_NO\_DATA\_FOUND

## Diagnostics

Table 35. SQLDataSources SQLSTATEs

SQLSTATE	Description	Explanation
01004	Data truncated.	The data source name returned in the argument <i>szDSN</i> is longer than the value specified in the argument <i>cbDSNMax</i> . The argument <i>pcbDSN</i> contains the length of the full data source name. (Function returns SQL_SUCCESS_WITH_INFO.)  The data source name returned in the argument <i>szDescription</i> is longer than the value specified in the argument <i>cbDescriptionMax</i> . The argument <i>pcbDescription</i> contains the length of the full data source description. (Function returns SQL_SUCCESS_WITH_INFO.)
58004	Unexpected system failure.	Unrecoverable system error.
S1000	General error.	An error occurred for which there is no specific SQLSTATE and for which no specific SQLSTATE is defined. The error message returned by <code>SQLError()</code> in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1090	Invalid string or buffer length.	The value specified for argument <i>cbDSNMax</i> is less than 0.  The value specified for argument <i>cbDescriptionMax</i> is less than 0.
S1103	Direction option out of range.	The value specified for the argument <i>fDirection</i> is not equal to SQL_FETCH_FIRST or SQL_FETCH_NEXT.



## Restrictions

None.

## Example

```

/* ... */
/*****
**   - demonstrate SQLDataSource function
**   - list available servers
**     (error checking has been ignored for simplicity)
**
** Functions used:
**
**   SQLAllocEnv      SQLFreeEnv
**   SQLDataSources
*****/

#include <stdio.h>
#include <stdlib.h>
#include "sqlcli1.h"

int
main()
{
    SQLRETURN      rc;
    SQLHENV        henv;
    SQLCHAR        source[SQL_MAX_DSN_LENGTH + 1], description[255];
    SQLSMALLINT    buff1, des1;

    SQLAllocEnv(&henv);          /* allocate an environment handle */

    /* list the available data sources (servers) */
    printf("The following data sources are available:\n");
    printf("ALIAS NAME          Comment(Description)\n");
    printf("-----\n");

    while ((rc = SQLDataSources(henv, SQL_FETCH_NEXT, source,
                                SQL_MAX_DSN_LENGTH + 1, &buff1, description, 255, &des1))
           != SQL_NO_DATA_FOUND) {
        printf("%-30s  %s\n", source, description);
    }

    SQLFreeEnv(henv);

    return (SQL_SUCCESS);
}
/* ... */

```

## References

None.

## SQLDescribeCol - Describe column attributes

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLDescribeCol() returns a set of commonly used descriptor information (column name, type, precision, scale, nullability) for the indicated column in the result set generated by a query.

If the application needs only one attribute of the descriptor information, or needs an attribute not returned by SQLDescribeCol(), the SQLColAttributes() function can be used in place of SQLDescribeCol(). See "SQLColAttributes - Get column attributes" on page 106 for more information.

Either SQLPrepare() or SQLExecDirect() must be called before calling this function.

This function (or SQLColAttributes()) is usually called before a bind column function SQLBindCol() to determine the attributes of a column before binding it to an application variable.

### Syntax

```
SQLRETURN SQLDescribeCol (SQLHSTMT hstmt,
                          SQLUSMALLINT icol,
                          SQLCHAR FAR *szColName,
                          SQLSMALLINT cbColNameMax,
                          SQLSMALLINT FAR *pcbColName,
                          SQLSMALLINT FAR *pfSqlType,
                          SQLINTEGER FAR *pcbColDef,
                          SQLSMALLINT FAR *piScale,
                          SQLSMALLINT FAR *pfNullable);
```

### Function arguments

Table 36 (Page 1 of 2). SQLDescribeCol arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle.
SQLUSMALLINT	<i>icol</i>	input	Column number to be described. Columns are numbered sequentially from left to right, starting at 1.
SQLCHAR *	<i>szColName</i>	output	Pointer to column name buffer. Set this to NULL if column name is not needed.
SQLSMALLINT	<i>cbColNameMax</i>	input	Size of <i>szColName</i> buffer.
SQLSMALLINT *	<i>pcbColName</i>	output	Bytes available to return for <i>szColName</i> argument. Truncation of column name ( <i>szColName</i> ) to <i>cbColNameMax</i> - 1 bytes occurs if <i>pcbColName</i> is greater than or equal to <i>cbColNameMax</i> .

Table 36 (Page 2 of 2). SQLDescribeCol arguments

Data type	Argument	Use	Description
SQLSMALLINT *	<i>pfSqlType</i>	output	Base SQL data type of column. To determine if there is a distinct type associated with the column, call <code>SQLColAttributes()</code> with <i>fDescType</i> set to <code>SQL_COLUMN_DISTINCT_TYPE</code> . See the symbolic SQL data type column of Table 4 on page 39 for the data types that are supported.
SQLINTEGER *	<i>pcbColDef</i>	output	Precision of column as defined in the database.  If <i>fSqlType</i> denotes a graphic or DBCLOB SQL data type, then this variable indicates the maximum number of double-byte characters the column can hold.
SQLSMALLINT *	<i>pibScale</i>	output	Scale of column as defined in the database (only applies to <code>SQL_DECIMAL</code> , <code>SQL_NUMERIC</code> , <code>SQL_TIMESTAMP</code> ). See Table 152 on page 446 for the scale of each of the SQL data types.
SQLSMALLINT *	<i>pfNullable</i>	output	Indicates whether NULLS are allowed for this column <ul style="list-style-type: none"> <li>• <code>SQL_NO_NULLS</code></li> <li>• <code>SQL_NULLABLE</code></li> </ul>

## Usage

Columns are identified by a number, are numbered sequentially from left to right starting with 1, and can be described in any order.

If a null pointer is specified for any of the pointer arguments, DB2 ODBC assumes that the information is not needed by the application and nothing is returned.

If the column is a distinct type, `SQLDescribeCol()` only returns the built-in type in *pfSqlType*. Call `SQLColAttributes()` with *fDescType* set to `SQL_COLUMN_DISTINCT_TYPE` to obtain the distinct type.

## Return codes

- `SQL_SUCCESS`
- `SQL_SUCCESS_WITH_INFO`
- `SQL_ERROR`
- `SQL_INVALID_HANDLE`

## Diagnostics

If `SQLDescribeCol()` returns either `SQL_ERROR`, or `SQL_SUCCESS_WITH_INFO`, one of the following `SQLSTATES` can be obtained by calling the `SQLError()` function.

Table 37 (Page 1 of 2). SQLDescribeCol SQLSTATES

SQLSTATE	Description	Explanation
01004	Data truncated.	The column name returned in the argument <i>szColName</i> is longer than the value specified in the argument <i>cbColNameMax</i> . The argument <i>pcbColName</i> contains the length of the full column name. (Function returns <code>SQL_SUCCESS_WITH_INFO</code> .)

## SQLDescribeCol

Table 37 (Page 2 of 2). SQLDescribeCol SQLSTATES

SQLSTATE	Description	Explanation
07005	The statement did not return a result set.	The statement associated with the <i>hstmt</i> did not return a result set. There are no columns to describe. (Call <code>SQLNumResultCols()</code> first to determine if there are any rows in the result set.)
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1002	Invalid column number.	The value specified for the argument <i>icol</i> is less than 1.  The value specified for the argument <i>icol</i> is greater than the number of columns in the result set.
S1090	Invalid string or buffer length.	The length specified in argument <i>cbColNameMax</i> is less than 1.
S1010	Function sequence error.	The function is called prior to calling <code>SQLPrepare()</code> or <code>SQLExecDirect()</code> for the <i>hstmt</i> .  The function is called while in a data-at-execute ( <code>SQLParamData()</code> , <code>SQLPutData()</code> ) operation.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1C00	Driver not capable.	The SQL data type of column <i>icol</i> is not recognized by DB2 ODBC.

## Restrictions

The ODBC defined data type `SQL_BIGINT` is not supported.

## Example

```

/* ... */
/*****
** process_stmt
** - allocates a statement handle
** - executes the statement
** - determines the type of statement
** - if there are no result columns, therefore non-select statement
**   - if rowcount > 0, assume statement was UPDATE, INSERT, DELETE
**   else
**     - assume a DDL, or Grant/Revoke statement
**   else
**     - must be a select statement.
**     - display results
** - frees the statement handle
*****/

int
process_stmt(SQLHENV henv,
            SQLHDBC hdbc,
            SQLCHAR * sqlstr)
{
    SQLHSTMT      hstmt;
    SQLSMALLINT   nresultcols;
    SQLINTEGER    rowcount;
    SQLRETURN     rc;

    SQLAllocStmt(hdbc, &hstmt); /* allocate a statement handle */

    /* execute the SQL statement in "sqlstr" */

    rc = SQLExecDirect(hstmt, sqlstr, SQL_NTS);
    if (rc != SQL_SUCCESS)
        if (rc == SQL_NO_DATA_FOUND) {
            printf("\nStatement executed without error, however,\n");
            printf("no data was found or modified\n");
            return (SQL_SUCCESS);
        } else
            check_error(henv, hdbc, hstmt, rc, __LINE__, __FILE__);
}

```

```

rc = SQLNumResultCols(hstmt, &nresultcols);

/* determine statement type */
if (nresultcols == 0) { /* statement is not a select statement */
    rc = SQLRowCount(hstmt, &rowcount);
    if (rowcount > 0) { /* assume statement is UPDATE, INSERT, DELETE */
        printf("Statement executed, %ld rows affected\n", rowcount);
    } else { /* assume statement is GRANT, REVOKE or a DLL
        * statement */
        printf("Statement completed successful\n");
    }
} else { /* display the result set */
    display_results(hstmt, nresultcols);
} /* end determine statement type */

rc = SQLFreeStmt(hstmt, SQL_DROP); /* free statement handle */

return (0);
} /* end process_stmt */

/*****
** display_results
**
** - for each column
**   - get column name
**   - bind column
** - display column headings
** - fetch each row
**   - if value truncated, build error message
**   - if column null, set value to "NULL"
**   - display row
**   - print truncation message
** - free local storage
*****/
display_results(SQLHSTMT hstmt,
                SQLSMALLINT nresultcols)
{
    SQLCHAR        colname[32];
    SQLSMALLINT    coltype;
    SQLSMALLINT    colnamelen;
    SQLSMALLINT    nullable;
    SQLINTEGER     collen[MAXCOLS];
    SQLINTEGER     precision;
    SQLSMALLINT    scale;
    SQLINTEGER     outlen[MAXCOLS];
    SQLCHAR        *data[MAXCOLS];
    SQLCHAR        errmsg[256];
    SQLRETURN      rc;
    SQLINTEGER     i;
    SQLINTEGER     x;
    SQLINTEGER     displaysize;

    for (i = 0; i < nresultcols; i++) {
        SQLDescribeCol(hstmt, i + 1, colname, sizeof(colname),
                      &colnamelen, &coltype, &precision, &scale, NULL);
        collen[i] = precision; /* Note, assignment of unsigned int to signed */

        /* get display length for column */
        SQLColAttributes(hstmt, i + 1, SQL_COLUMN_DISPLAY_SIZE, NULL, 0,
                        NULL, &displaysize);
    }
}

```

```

/*
 * set column length to max of display length, and column name
 * length. Plus one byte for null terminator
 */
collen[i] = max(displaysize, strlen((char *) colname)) + 1;

printf("%-*.*s", collen[i], collen[i], colname);

/* allocate memory to bind column */
data[i] = (SQLCHAR *) malloc(collen[i]);

/* bind columns to program vars, converting all types to CHAR */
rc = SQLBindCol(hstmt, i + 1, SQL_C_CHAR, data[i], collen[i], &outlen[i]);
}
printf("\n");

/* display result rows */
while ((rc = SQLFetch(hstmt)) != SQL_NO_DATA_FOUND) {
    errmsg[0] = '\0';
    for (i = 0; i < nresultcols; i++) {
        /* Build a truncation message for any columns truncated */
        if (outlen[i] >= collen[i]) {
            sprintf((char *) errmsg + strlen((char *) errmsg),
                "%ld chars truncated, col %d\n",
                outlen[i] - collen[i] + 1, i + 1);
            sprintf((char *) errmsg + strlen((char *) errmsg),
                "Bytes to return = %ld size of buffer\n",
                outlen[i], collen[i]);
        }
        if (outlen[i] == SQL_NULL_DATA)
            printf("%-*.*s", collen[i], collen[i], "NULL");
        else
            printf("%-*.*s", collen[i], collen[i], data[i]);
    }
    /* for all columns in this row */

    printf("\n%s", errmsg); /* print any truncation messages */
} /* while rows to fetch */

/* free data buffers */
for (i = 0; i < nresultcols; i++) {
    free(data[i]);
}

} /* end display_results */
/* ... */

```

## References

- “SQLColAttributes - Get column attributes” on page 106
- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLNumResultCols - Get number of result columns” on page 271
- “SQLPrepare - Prepare a statement” on page 277

## SQLDescribeParam - Describe parameter marker

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLDescribeParam() retrieves the description of a parameter marker associated with a prepared statement.

Either SQLPrepare() or SQLExecDirect() must be called before calling this function.

### Syntax

```
SQLRETURN SQLDescribeParam (SQLHSTMT          hstmt,
                             SQLUSMALLINT      ipar,
                             SQLSMALLINT FAR   *pfSqlType,
                             SQLINTEGER FAR    *pcbColDef,
                             SQLSMALLINT FAR   *pibScale,
                             SQLSMALLINT FAR   *pfNullable);
```

### Function arguments

Table 38. SQLDescribeParam arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle.
SQLUSMALLINT	ipar	input	Parameter marker number ordered sequentially left to right in prepared SQL statement, starting from 1.
SQLSMALLINT *	pfSqlType	output	Base SQL data type.
SQLINTEGER *	pcbColDef	output	Precision of the parameter marker. See Appendix D, "Data conversion" on page 443 for more details on precision, scale, length, and display size.
SQLSMALLINT *	pibScale	output	Scale of the parameter marker. See Appendix D, "Data conversion" on page 443 for more details on precision, scale, length, and display size.
SQLSMALLINT *	pfNullable	output	Indicates whether the parameter allows NULL values. Returns one of the following values: <ul style="list-style-type: none"> <li>SQL_NO_NULLS: The parameter does not allow NULL values (this is the default).</li> <li>SQL_NULLABLE: The parameter allows NULL values.</li> <li>SQL_NULLABLE_UNKNOWN: The driver cannot determine if the parameter allows NULL values.</li> </ul>

### Usage

For distinct types, SQLDescribeParam() returns both base data types for the input parameter.

For information about a parameter marker associated with the SQL CALL statement, use the SQLProcedureColumns() function.



## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 39. SQLDescribeParam SQLSTATEs

SQLSTATE	Description	Explanation
01000	Warning.	Informational message indicating an internal commit was issued on behalf of the application as part of the processing to set the specified connection option.
S1000	General error.	An error occurred for which there is no specific SQLSTATE and for which no specific SQLSTATE is defined. The error message returned by <code>SQLError()</code> in the argument <code>szErrorMsg</code> describes the error and its cause.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute ( <code>SQLParamData()</code> , <code>SQLPutData()</code> ) operation.
S1093	Invalid parameter number.	The value specified for the argument <code>ipar</code> is less than 1 or greater than the maximum number of parameters supported by the server.
S1C00	Driver not capable.	The data source does not support the description of input parameters.

## Restrictions

None.

## References

- “SQLBindParameter - Binds a parameter marker to a buffer or LOB locator” on page 93
- “SQLCancel - Cancel statement” on page 104
- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLExecute - Execute a statement” on page 159
- “SQLPrepare - Prepare a statement” on page 277

---

## SQLDisconnect - Disconnect from a data source

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>	<b>X/OPEN CLI</b>	<b>ISO CLI</b>
-----------------------	-----------------	-------------------	----------------

SQLDisconnect() closes the connection associated with the database connection handle.

SQLTransact() must be called before calling SQLDisconnect() if an outstanding transaction exists on this connection.

After calling this function, either call SQLConnect() to connect to another database, or call SQLFreeConnect().

### Syntax

```
SQLRETURN SQLDisconnect (SQLHDBC          hdbc);
```

### Function arguments

Table 40. SQLDisconnect arguments

Data type	Argument	Use	Description
SQLHDBC	<i>hdbc</i>	input	Connection handle

### Usage

If an application calls SQLDisconnect() before it has freed all the statement handles associated with the connection, DB2 ODBC frees them after it successfully disconnects from the database.

If SQL\_SUCCESS\_WITH\_INFO is returned, it implies that even though the disconnect from the database is successful, additional error or implementation specific information is available. For example, a problem was encountered on the clean up subsequent to the disconnect, or if there is no current connection because of an event that occurred independently of the application (such as communication failure).

After a successful SQLDisconnect() call, the application can re-use *hdbc* to make another SQLConnect() or SQLDriverConnect() request.

### Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 41. SQLDisconnect SQLSTATEs

SQLSTATE	Description	Explanation
01002	Disconnect error.	An error occurred during the disconnect. However, the disconnect succeeded. (Function returns SQL_SUCCESS_WITH_INFO.)
08003	Connection is closed.	The connection specified in the argument <i>hdbc</i> is not open.
25000 25501	Invalid transaction state.	A transaction is in process on the connection specified by the argument <i>hdbc</i> . The transaction remains active, and the connection cannot be disconnected.  <b>Note:</b> This error does not apply to stored procedures written in DB2 ODBC.
25501	Invalid transaction state.	A transaction is in process on the connection specified by the argument <i>hdbc</i> . The transaction remains active, and the connection cannot be disconnected.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.

## Restrictions

None.

## Example

See “Example” on page 80

## References

- “SQLAllocConnect - Allocate connection handle” on page 79
- “SQLConnect - Connect to a data source” on page 123
- “SQLDriverConnect - (Expanded) connect to a data source” on page 142
- “SQLTransact - Transaction management” on page 357

## SQLDriverConnect - (Expanded) connect to a data source

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>		
-----------------------	-----------------	--	--

SQLDriverConnect() is an alternative to SQLConnect(). Both functions establish a connection to the target database, but SQLDriverConnect() supports additional connection parameters.

Use SQLDriverConnect() when you want to pass any or all keyword values defined in the DB2 ODBC initialization file.

When a connection is established, the completed connection string is returned. Applications can store this string for future connection requests. This allows you to override any or all keyword values in the DB2 ODBC initialization file.

### Syntax

#### Generic

```
SQLRETURN SQLDriverConnect (SQLHDBC          hdbc,
                             SQLHWND         hwnd,
                             SQLCHAR          FAR *szConnStrIn,
                             SQLSMALLINT     cbConnStrIn,
                             SQLCHAR          FAR *szConnStrOut,
                             SQLSMALLINT     cbConnStrOutMax,
                             SQLSMALLINT     FAR *pcbConnStrOut,
                             SQLSMALLINT     fDriverCompletion);
```

### Function arguments

Table 42 (Page 1 of 2). SQLDriverConnect arguments


Data type	Argument	Use	Description
SQLHDBC	<i>hdbc</i>	input	Connection handle.
SQLHWND	<i>hwindow</i>	input	NULL value. Not used.
SQLCHAR *	<i>szConnStrIn</i>	input	A full, partial or empty (null pointer) connection string (see syntax and description below).
SQLSMALLINT	<i>cbConnStrIn</i>	input	Length of <i>szConnStrIn</i> .
SQLCHAR *	<i>szConnStrOut</i>	output	Pointer to buffer for the completed connection string.  If the connection is established successfully, this buffer contains the completed connection string. Applications should allocate at least SQL_MAX_OPTION_STRING_LENGTH bytes for this buffer.
SQLSMALLINT	<i>cbConnStrOutMax</i>	input	Maximum size of the buffer pointed to by <i>szConnStrOut</i> .
SQLCHAR *	<i>pcbConnStrOut</i>	output	Pointer to the number of bytes available to return in the <i>szConnStrOut</i> buffer.  If the value of <i>pcbConnStrOut</i> is greater than or equal to <i>cbConnStrOutMax</i> , the completed connection string in <i>szConnStrOut</i> is truncated to <i>cbConnStrOutMax</i> - 1 bytes.

Table 42 (Page 2 of 2). SQLDriverConnect arguments

Data type	Argument	Use	Description
SQLUSMALLINT	<i>fDriverCompletion</i>	input	Indicates when DB2 ODBC should prompt the user for more information.  Possible values: <ul style="list-style-type: none"> <li>• SQL_DRIVER_PROMPT</li> <li>• SQL_DRIVER_COMPLETE</li> <li>• SQL_DRIVER_COMPLETE_REQUIRED</li> <li>• SQL_DRIVER_NOPROMPT</li> </ul> However, DB2 for OS/390 supports SQL_DRIVER_NOPROMPT only.

## Usage

The connection string is used to pass one or more values needed to complete a connection.



Each keyword above has an attribute that is equal to the following:

**DSN** Data source name. The name or alias-name of the database. Required if *fDriverCompletion* is equal to SQL\_DRIVER\_NOPROMPT.

**UID** Authorization-name (user identifier). This value is ignored.

**PWD** The password corresponding to the authorization name. If there is no password for the user ID, an empty string is specified (PWD=;). This value is ignored.

The list of DB2 ODBC defined keywords and their associated attribute values are discussed in “Initialization keywords” on page 63. Any one of the keywords in that section can be specified on the connection string. If any keywords are repeated in the connection string, the value associated with the first occurrence of the keyword is used.

If any keywords exist in the DB2 ODBC initialization file, the keywords and their respective values are used to augment the information passed to DB2 ODBC in the connection string. If the information in the DB2 ODBC initialization file contradicts information in the connection string, the values in connection string take precedence.

The application receives an error on any value of *fDriverCompletion* as follows:

**SQL\_DRIVER\_PROMPT:**

DB2 ODBC returns SQL\_ERROR.

**SQL\_DRIVER\_COMPLETE:**

DB2 ODBC returns SQL\_ERROR.

### SQL\_DRIVER\_COMPLETE\_REQUIRED:

DB2 ODBC returns SQL\_ERROR.

### SQL\_DRIVER\_NOPROMPT:

The user is not prompted for any information. A connection is attempted with the information contained in the connection string. If there is not enough information, SQL\_ERROR is returned.

When a connection is established, the complete connection string is returned.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_NO\_DATA\_FOUND
- SQL\_INVALID\_HANDLE
- SQL\_ERROR

## Diagnostics

All of the diagnostics generated by “SQLConnect - Connect to a data source” on page 123 can be returned here as well. The following table shows the additional diagnostics that can be returned.

Table 43. SQLDriverConnect SQLSTATEs

SQLSTATE	Description	Explanation
01004	Data truncated.	The buffer <i>szConnstrOut</i> is not large enough to hold the entire connection string. The argument <i>pcbConnStrOut</i> contains the actual length of the connection string available for return. (Function returns SQL_SUCCESS_WITH_INFO)
01S00	Invalid connection string attribute.	An invalid keyword or attribute value is specified in the input connection string, but the connection to the data source is successful because one of the following occurred: <ul style="list-style-type: none"> <li>• The unrecognized keyword is ignored.</li> <li>• The invalid attribute value is ignored, the default value is used instead.</li> </ul> (Function returns SQL_SUCCESS_WITH_INFO)
01S02	Option value changed.	SQL_CONNECTTYPE changed to SQL_CONCURRENT_TRANS when MULTICONTEXT=1 in use.
S1090	Invalid string or buffer length.	The value specified for <i>cbConnStrIn</i> is less than 0, but not equal to SQL_NTS.  The value specified for <i>cbConnStrOutMax</i> is less than 0.
S1110	Invalid driver completion.	The value specified for the argument <i>fCompletion</i> is not equal to one of the valid values.

## Restrictions

See restrictions described above for *fDriverCompletion* and *SQLHWND* parameters.

## Example

```

/*****
/* DB2 for OS/390 example:
/* Issues SQLDriverConnect to pass a string of initialization
/* parameters to compliment the connection to the data source.
/*****

#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include "sqlcli1.h"

/*****
/* SQLDriverConnect -----
/*****

int main( )
{
    SQLHENV          hEnv    = SQL_NULL_HENV;
    SQLHDBC          hDbc    = SQL_NULL_HDBC;
    SQLRETURN        rc      = SQL_SUCCESS;
    SQLINTEGER       RETCODE = 0;

    char             *ConnStrIn =
        "dsn=STLEC1;connecttype=2;bitdata=2;optimizeforrows=30";

    char             ConnStrOut [200];
    SQLSMALLINT      cbConnStrOut;
    int              i;
    char             *token;

    (void) printf ("*** Entering CLIP10.\n\n");

    /*****
    /* CONNECT to DB2 for OS/390
    /*****

    rc = SQLAllocEnv(&hEnv);

    if( rc != SQL_SUCCESS )
        goto dberror;

    /*****
    /* Allocate Connection Handle to DSN
    /*****

    RETCODE = SQLAllocConnect(hEnv,
                              &hDbc);

    if( RETCODE != SQL_SUCCESS ) // Could not get a Connect Handle
        goto dberror;

    /*****
    /* Invoke SQLDriverConnect -----
    /*****

    RETCODE = SQLDriverConnect (hDbc
                               ,
                               NULL
                               ,
                               (SQLCHAR *)ConnStrIn
                               ,
                               strlen(ConnStrIn)
                               ,
                               (SQLCHAR *)ConnStrOut,

```

```

        sizeof(ConnStrOut)    ,
        &cbConnStrOut        ,
        SQL_DRIVER_NOPROMPT);
if( RETCODE != SQL_SUCCESS ) // Could not get a Connect Handle
{
    (void) printf ("**** Driver Connect Failed. rc = %d.\n", RETCODE);
    goto dberror;
}

/*****
/* Enumerate keywords and values returned from SQLDriverConnect */
*****/

(void) printf ("**** ConnStrOut = %s.\n", ConnStrOut);

for (i = 1, token = strtok (ConnStrOut, ";");
     (token != NULL);
     token = strtok (NULL, ";"), i++)
    (void) printf ("**** Keyword # %d is: %s.\n", i, token);

/*****
/* DISCONNECT from data source */
*****/

RETCODE = SQLDisconnect(hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Deallocate Connection Handle */
*****/

RETCODE = SQLFreeConnect (hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Disconnect from data sources in Connection Table */
*****/

SQLFreeEnv(hEnv);          /* free the environment handle */

goto exit;

dberror:
RETCODE=12;

exit:

(void) printf ("**** Exiting CLIP10.\n\n");

return(RETCODE);
}

```



## References

- “SQLAllocConnect - Allocate connection handle” on page 79
- “SQLConnect - Connect to a data source” on page 123

## SQLError - Retrieve error information

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLError() returns the diagnostic information (both errors and warnings) associated with the most recently invoked DB2 ODBC function for a particular statement, connection or environment handle.

The information consists of a standardized SQLSTATE and native error code. See “Diagnostics” on page 34 for more information.

Call SQLError() after receiving a return code of SQL\_ERROR or SQL\_SUCCESS\_WITH\_INFO from another function call.

**Note:** Some database servers provide product-specific diagnostic information after returning SQL\_NO\_DATA\_FOUND from the execution of a statement.

### Syntax

```
SQLRETURN SQLError
            (SQLHENV      henv,
             SQLHDBC      hdbc,
             SQLHSTMT     hstmt,
             SQLCHAR FAR  *szSqlState,
             SQLINTEGER FAR *pfNativeError,
             SQLCHAR FAR  *szErrorMsg,
             SQLSMALLINT  cbErrorMsgMax,
             SQLSMALLINT FAR *pcbErrorMsg);
```

### Function arguments

Table 44 (Page 1 of 2). SQLError arguments

Data type	Argument	Use	Description
SQLHENV	<i>henv</i>	input	Environment handle. To obtain diagnostic information associated with an environment, pass a valid environment handle. Set <i>hdbc</i> and <i>hstmt</i> to SQL_NULL_HDBC and SQL_NULL_HSTMT respectively.
SQLHDBC	<i>hdbc</i>	input	Database connection handle. To obtain diagnostic information associated with a connection, pass a valid database connection handle, and set <i>hstmt</i> to SQL_NULL_HSTMT. The <i>henv</i> argument is ignored.
SQLHSTMT	<i>hstmt</i>	input	Statement handle. To obtain diagnostic information associated with a statement, pass a valid statement handle. The <i>henv</i> and <i>hdbc</i> arguments are ignored.
SQLCHAR *	<i>szSqlState</i>	output	SQLSTATE as a string of 5 characters terminated by a null character. The first 2 characters indicate error class; the next 3 indicate subclass. The values correspond directly to SQLSTATE values defined in the X/Open SQL CAE specification and the ODBC specification, augmented with IBM specific and product specific SQLSTATE values.

Table 44 (Page 2 of 2). SQLError arguments

Data type	Argument	Use	Description
SQLINTEGER *	<i>pfNativeError</i>	output	Native error code. In DB2 ODBC, the <i>pfNativeError</i> argument contains the SQLCODE value returned by the DBMS. If the error is generated by DB2 ODBC and not the DBMS, then this field is set to -99999.
SQLCHAR *	<i>szErrorMsg</i>	output	<p>Pointer to buffer to contain the implementation defined message text. If the error is detected by DB2 ODBC, then the error message is prefaced by:</p> <p>[DB2 for OS/390][CLI Driver]</p> <p>to indicate that it is DB2 ODBC that detected the error and there is no database connection yet.</p> <p>The error location, ERRLOC x:y:z, keyword value is embedded in the buffer also. This is an internal error code for diagnostics.</p> <p>If the error is detected while there is a database connection, then the error message returned from the DBMS is prefaced by:</p> <p>[DB2 for OS/390][CLI Driver][DBMS-name]</p> <p>where DBMS-name is the name returned by SQLGetInfo() with SQL_DBMS_NAME information type.</p> <p>For example,</p> <pre>DB2 DB2/6000 Vendor</pre> <p>Vendor indicates a non-IBM DRDA DBMS.</p> <p>If the error is generated by the DBMS, the IBM-defined SQLSTATE is appended to the text string.</p>
SQLSMALLINT	<i>cbErrorMsgMax</i>	input	The maximum (that is, the allocated) length of the buffer <i>szErrorMsg</i> . The recommended length to allocate is SQL_MAX_MESSAGE_LENGTH + 1.
SQLSMALLINT *	<i>pcbErrorMsg</i>	output	Pointer to total number of bytes available to return to the <i>szErrorMsg</i> buffer. This does not include the null termination character.

## Usage

The SQLSTATEs are those defined by the X/OPEN SQL CAE and the X/Open SQL CLI CAE, augmented with IBM specific and product specific SQLSTATE values.

To obtain diagnostic information associated with:

- An environment, pass a valid environment handle. Set *hdbc* and *hstmt* to SQL\_NULL\_HDBC and SQL\_NULL\_HSTMT respectively.
- A connection, pass a valid database connection handle, and set *hstmt* to SQL\_NULL\_HSTMT. The *henv* argument is ignored.
- A statement, pass a valid statement handle. The *henv* and *hdbc* arguments are ignored.

## SQL\_Error

If diagnostic information generated by one DB2 ODBC function is not retrieved before a function other than `SQL_Error()` is called with the same handle, the information for the previous function call is lost. This is true whether or not diagnostic information is generated for the second DB2 ODBC function call.

Multiple diagnostic messages might be available after a given DB2 ODBC function call. These messages can be retrieved one at a time by repeatedly calling `SQL_Error()`. For each message retrieved, `SQL_Error()` returns `SQL_SUCCESS` and removes it from the list of messages available. When there are no more messages to retrieve, `SQL_NO_DATA_FOUND` is returned, the `SQLSTATE` is set to "00000", `pfNativeError` is set to 0, and `pcbErrorMsg` and `szErrorMsg` are undefined.

Diagnostic information stored under a given handle is cleared when a call is made to `SQL_Error()` with that handle, or when another DB2 ODBC function call is made with that handle. However, information associated with a given handle type is not cleared by a call to `SQL_Error()` with an associated but different handle type: for example, a call to `SQL_Error()` with a connection handle input does not clear errors associated with any statement handles under that connection.

`SQL_SUCCESS` is returned even if the buffer for the error message (`szErrorMsg`) is too short since the application is not able to retrieve the same error message by calling `SQL_Error()` again. The actual length of the message text is returned in the `pcbErrorMsg`.

To avoid truncation of the error message, declare a buffer length of `SQL_MAX_MESSAGE_LENGTH + 1`. The message text is never longer than this.

## Return codes

- `SQL_SUCCESS`
- `SQL_ERROR`
- `SQL_INVALID_HANDLE`
- `SQL_NO_DATA_FOUND`

`SQL_NO_DATA_FOUND` is returned if no diagnostic information is available for the input handle, or if all of the messages are retrieved by calls to `SQL_Error()`.

## Diagnostics

`SQLSTATE`s are not defined, since `SQL_Error()` does not generate diagnostic information for itself.

## Restrictions

Although ODBC also returns X/Open SQL CAE `SQLSTATE`s, only DB2 ODBC (and the DB2 ODBC driver) returns the additional IBM-defined `SQLSTATE`s. For more information on ODBC specific `SQLSTATE`s see *Microsoft ODBC 3.0 Programmer's Reference and SDK Guide*.

Because of this, you should only build dependencies on the standard `SQLSTATE`s. This means any branching logic in the application should only rely on the standard `SQLSTATE`s. The augmented `SQLSTATE`s are most useful for debugging purposes.

**Note:** It might be useful to build dependencies on the class (the first 2 characters) of the `SQLSTATE`s.

## Example

This example shows several utility functions used by most of the other DB2 ODBC examples.

```

/* ... */
/*****
** - print_error - call SQLError(), display SQLSTATE and message
**               - called by check_error, see below
*****/

SQLRETURN
print_error(SQLHENV henv,
            SQLHDBC hdbc,
            SQLHSTMT hstmt,
            SQLRETURN frc, /* Return code to be included with error msg */
            SQLINTEGER line, /* Used for output message, indicate where */
            SQLCHAR * file) /* the error was reported from */
{
    SQLCHAR      buffer[SQL_MAX_MESSAGE_LENGTH + 1];
    SQLCHAR      sqlstate[SQL_SQLSTATE_SIZE + 1];
    SQLINTEGER   sqlcode;
    SQLSMALLINT  length;

    printf(">--- ERROR -- RC= %ld Reported from %s, line %ld -----\n",
           frc, file, line);
    while (SQLError(henv, hdbc, hstmt, sqlstate, &sqlcode, buffer,
                   SQL_MAX_MESSAGE_LENGTH + 1, &length) == SQL_SUCCESS) {
        printf("      SQLSTATE: %s\n", sqlstate);
        printf("Native Error Code: %ld\n", sqlcode);
        printf("%s \n", buffer);
    };
    printf(">-----\n");
    return (SQL_ERROR);
}

/*****
** - check_error - call print_error(), checks severity of return code
*****/
SQLRETURN
check_error(SQLHENV henv,
            SQLHDBC hdbc,
            SQLHSTMT hstmt,
            SQLRETURN frc,
            SQLINTEGER line,
            SQLCHAR * file)
{
    SQLRETURN    rc;

    print_error(henv, hdbc, hstmt, frc, line, file);

    switch (frc) {
    case SQL_SUCCESS:
        break;
    case SQL_INVALID_HANDLE:
        printf("\n>----- ERROR Invalid Handle -----\n");
    }
}

```

## SQL\_Error

```
case SQL_ERROR:
    printf("\n>--- FATAL ERROR, Attempting to rollback transaction --\n");
    rc = SQLTransact(henv, hdbc, SQL_ROLLBACK);
    if (rc != SQL_SUCCESS)
        printf(">Rollback Failed, Exiting application\n");
    else
        printf(">Rollback Successful, Exiting application\n");
    exit(0);
    break;
case SQL_SUCCESS_WITH_INFO:
    printf("\n> ----- Warning Message, application continuing ----- \n");
    break;
case SQL_NO_DATA_FOUND:
    printf("\n> ----- No Data Found, application continuing ----- \n");
    break;
default:
    printf("\n> ----- Invalid Return Code ----- \n");
    printf("> ----- Attempting to rollback transaction ----- \n");
    SQLTransact(henv, hdbc, SQL_ROLLBACK);
    exit(0 );
    break;
}
return (SQL_SUCCESS);
}
/* ... */

}
/*****
* The following macros use check_error
*
* {check_error(henv, SQL_NULL_HDBC, SQL_NULL_HSTMT, RC, __LINE__, __FILE__);}
*
* {check_error(SQL_NULL_HENV, hdbc, SQL_NULL_HSTMT, RC, __LINE__, __FILE__);}
*
* {check_error(SQL_NULL_HENV, SQL_NULL_HDBC, hstmt, RC, __LINE__, __FILE__);}
*
*****/

/*****
** - check_error - call print_error(), checks severity of return code
*****/
SQLRETURN
check_error(SQLHENV henv,
            SQLHDBC hdbc,
            SQLHSTMT hstmt,
            SQLRETURN frc,
            SQLINTEGER line,
            SQLCHAR * file)
{
    SQLRETURN rc;

    print_error(henv, hdbc, hstmt, frc, line, file);
}
```

```

switch (frc) {
case SQL_SUCCESS:
    break;
case SQL_INVALID_HANDLE:
    printf("\n>----- ERROR Invalid Handle ----- \n");
case SQL_ERROR:
    printf("\n>--- FATAL ERROR, Attempting to rollback transaction -- \n");
    rc = SQLTransact(henv, hdbc, SQL_ROLLBACK);
    if (rc != SQL_SUCCESS)
        printf(">Rollback Failed, Exiting application \n");
    else
        printf(">Rollback Successful, Exiting application \n");
    exit(terminate(henv, frc));
    break;
case SQL_SUCCESS_WITH_INFO:
    printf("\n> ----- Warning Message, application continuing ----- \n");
    break;
case SQL_NO_DATA_FOUND:
    printf("\n> ----- No Data Found, application continuing ----- \n");
    break;
default:
    printf("\n> ----- Invalid Return Code ----- \n");
    printf("> ----- Attempting to rollback transaction ----- \n");
    SQLTransact(henv, hdbc, SQL_ROLLBACK);
    exit(terminate(henv, frc));
    break;
}
return (SQL_SUCCESS);
}
/* ... */
/* end check_error */

```

## References

- “SQLGetSQLCA - Get SQLCA data structure” on page 242

## SQLExecDirect - Execute a statement directly

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLExecDirect() directly executes the specified SQL statement. The statement can only be executed once. Also, the connected database server must be able to dynamically prepare statement. (For more information about supported SQL statements see Table 1 on page 16.)

### Syntax

```
SQLRETURN SQLExecDirect (SQLHSTMT hstmt,
                          SQLCHAR FAR *szSqlStr,
                          SQLINTEGER cbSqlStr);
```

### Function arguments

Table 45. SQLExecDirect arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle. There must not be an open cursor associated with <i>hstmt</i> , see “SQLFreeStmt - Free (or reset) a statement handle” on page 188 for more information.
SQLCHAR *	<i>szSqlStr</i>	input	SQL statement string. The connected database server must be able to prepare the statement, see Table 1 on page 16 for more information.
SQLINTEGER	<i>cbSqlStr</i>	input	Length of contents of <i>szSqlStr</i> argument. The length must be set to either the exact length of the statement, or if the statement is null-terminated, set to SQL_NTS.

### Usage

If the SQL statement text contains vendor escape clause sequences, DB2 ODBC first modifies the SQL statement text to the appropriate DB2-specific format before submitting it for preparation and execution. If the application does not generate SQL statements that contain vendor escape clause sequences ( “Using vendor escape clauses” on page 404), then it should set the SQL\_NO\_SCAN statement option to SQL\_NOSCAN\_ON at the connection level so that each statement passed to DB2 ODBC does not incur the performance impact of scanning for vendor escape clauses.

The SQL statement cannot be a COMMIT or ROLLBACK. Instead, SQLTransact() must be called to issue COMMIT or ROLLBACK. For more information about supported SQL statements see Table 1 on page 16.

The SQL statement string can contain parameter markers. A parameter marker is represented by a “?” character, and is used to indicate a position in the statement where an application supplied value is to be substituted when SQLExecDirect() is called. This value can be obtained from:

- An application variable.



SQLSetParam() or SQLBindParameter() is used to bind the application storage area to the parameter marker.

- A LOB value residing at the server referenced by a LOB locator.

SQLBindParameter() or SQLSetParam() is used to bind a LOB locator to a parameter marker. The actual value of the LOB is kept at the server and does not need to be transferred to the application before being used as the input parameter value for another SQL statement.

All parameters must be bound before calling SQLExecDirect().

If the SQL statement is a query, SQLExecDirect() generates a cursor name, and open the cursor. If the application has used SQLSetCursorName() to associate a cursor name with the statement handle, DB2 ODBC associates the application generated cursor name with the internally generated one.

If a result set is generated, SQLFetch() or SQLExtendedFetch() retrieves the next row (or rows) of data into bound variables. Data can also be retrieved by calling SQLGetData() for any column that was not bound.

If the SQL statement is a positioned DELETE or a positioned UPDATE, the cursor referenced by the statement must be positioned on a row and must be defined on a **separate** statement handle under the same connection handle.

There must not already be an open cursor on the statement handle.

If SQLParamOptions() is called to specify that an array of input parameter values is bound to each parameter marker, then the application needs to call SQLExecDirect() only once to process the entire array of input parameter values.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE
- SQL\_NEED\_DATA
- SQL\_NO\_DATA\_FOUND

SQL\_NEED\_DATA is returned when the application requests to input data-at-execution parameter values by calling SQLParamData() and SQLPutData().

SQL\_SUCCESS is returned if the SQL statement is a searched UPDATE or searched DELETE and no rows satisfy the search condition. Use SQLRowCount() to determine the number of rows in a table that were affected by an UPDATE, INSERT, or DELETE statement executed against the table, or a view of the table.

## Diagnostics

Table 46 (Page 1 of 4). SQLExecDirect SQLSTATEs

SQLSTATE	Description	Explanation
01504	The UPDATE or DELETE statement does not include a WHERE clause.	szSqlStr contains an UPDATE or DELETE statement but no WHERE clause. (Function returns SQL_SUCCESS_WITH_INFO or SQL_NO_DATA_FOUND if there are no rows in the table).

Table 46 (Page 2 of 4). SQLExecDirect SQLSTATES

SQLSTATE	Description	Explanation
07001	Wrong number of parameters.	The number of parameters bound to application variables using SQLBindParameter() is less than the number of parameter markers in the SQL statement contained in the argument <i>szSqlStr</i> .
07006	Invalid conversion.	Transfer of data between DB2 ODBC and the application variables would result in incompatible data conversion.
21S01	Insert value list does not match column list.	<i>szSqlStr</i> contains an INSERT statement and the number of values to be inserted did not match the degree of the derived table.
21S02	Degrees of derived table does not match column list.	<i>szSqlStr</i> contains a CREATE VIEW statement and the number of names specified is not the same degree as the derived table defined by the query specification.
22001	String data right truncation.	A character string assigned to a character type column exceeded the maximum length of the column.
22003	Numeric value out of range.	A numeric value assigned to a numeric type column caused truncation of the whole part of the number, either at the time of assignment or in computing an intermediate result.  <i>szSqlStr</i> contains an SQL statement with an arithmetic expression which caused division by zero.
22005	Error in assignment.	<i>szSqlStr</i> contains an SQL statement with a parameter or literal and the value or LOB locator was incompatible with the data type of the associated table column.  The length associated with a parameter value (the contents of the <i>pcbValue</i> buffer specified on SQLBindParameter()) is not valid.  The argument <i>fSQLType</i> used in SQLBindParameter() or SQLSetParam(), denoted an SQL graphic data type, but the deferred length argument ( <i>pcbValue</i> ) contains an odd length value. The length value must be even for graphic data types.
22007	Invalid datetime format.	<i>szSqlStr</i> contains an SQL statement with an invalid datetime format; that is, an invalid string representation or value was specified, or the value was an invalid date.
22008	Datetime field overflow.	Datetime field overflow occurred; for example, an arithmetic operation on a date or timestamp has a result that is not within the valid range of dates, or a datetime value cannot be assigned to a bound variable because it is too small.
22012	Division by zero is invalid.	<i>szSqlStr</i> contains an SQL statement with an arithmetic expression that caused division by zero.
23000	Integrity constraint violation.	The execution of the SQL statement is not permitted because the execution would cause integrity constraint violation in the DBMS.
24000	Invalid cursor state.	A cursor was already opened on the statement handle.
24504	The cursor identified in the UPDATE, DELETE, SET, or GET statement is not positioned on a row.	Results were pending on the <i>hstmt</i> from a previous query or a cursor associated with the <i>hstmt</i> had not been closed.
34000	Invalid cursor name.	<i>szSqlStr</i> contains a positioned DELETE or a positioned UPDATE and the cursor referenced by the statement being executed was not open.

Table 46 (Page 3 of 4). SQLExecDirect SQLSTATES

SQLSTATE	Description	Explanation
37xxx <sup>a</sup>	Invalid SQL syntax.	<i>szSqlStr</i> contains one or more of the following: <ul style="list-style-type: none"> <li>• a COMMIT</li> <li>• a ROLLBACK</li> <li>• an SQL statement that the connected database server could not prepare</li> <li>• a statement containing a syntax error</li> </ul>
40001	Transaction rollback.	The transaction to which this SQL statement belongs is rolled back due to a deadlock or timeout.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
42xxx	Syntax error or access rule violation	<b>425xx</b> indicates the authorization ID does not have permission to execute the SQL statement contained in <i>szSqlStr</i> .  Other <b>42xxx</b> SQLSTATES indicate a variety of syntax or access problems with the statement.
42895	The value of a host variable in the EXECUTE or OPEN statement cannot be used because of its data type	The LOB locator type specified on the bind parameter function call does not match the LOB data type of the parameter marker.  The argument <i>fSQLType</i> used on the bind parameter function specified a LOB locator type but the corresponding parameter marker is not a LOB.
44000	Integrity constraint violation.	<i>szSqlStr</i> contains an SQL statement with a parameter or literal. This parameter value is NULL for a column defined as NOT NULL in the associated table column, or a duplicate value is supplied for a column constrained to contain only unique values, or some other integrity constraint is violated.
58004	Unexpected system failure.	Unrecoverable system error.
S0001	Database object already exists.	<i>szSqlStr</i> contains a CREATE TABLE or CREATE VIEW statement and the table name or view name specified already exists.
S0002	Database object does not exist.	<i>szSqlStr</i> contains an SQL statement that references a table name or view name which does not exist.
S0011	Index already exists.	<i>szSqlStr</i> contains a CREATE INDEX statement and the specified index name already exists.
S0012	Index not found.	<i>szSqlStr</i> contains a DROP INDEX statement and the specified index name does not exist.
S0021	Column already exists.	<i>szSqlStr</i> contains an ALTER TABLE statement and the column specified in the ADD clause is not unique or identifies an existing column in the base table.
S0022	Column not found.	<i>szSqlStr</i> contains an SQL statement that references a column name that does not exist.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	<i>szSqlStr</i> is a null pointer.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1014	No more handles.	DB2 ODBC is not able to allocate a handle due to internal resources.
S1090	Invalid string or buffer length.	The argument <i>cbSqlStr</i> is less than 1 but not equal to SQL_NTS.

## SQLExecDirect

Table 46 (Page 4 of 4). SQLExecDirect SQLSTATES

SQLSTATE	Description	Explanation
<b>Note:</b>		
a	xxx refers to any SQLSTATE with that class code. Example, <b>37xxx</b> refers to any SQLSTATE in the <b>37</b> class.	

## Restrictions

None.

## Example

See “Example” on page 173.

## References

- “SQLBindParameter - Binds a parameter marker to a buffer or LOB locator” on page 93
- “SQLExecute - Execute a statement” on page 159
- “SQLExtendedFetch - Extended fetch (fetch array of rows)” on page 162
- “SQLFetch - Fetch next row” on page 169
- “SQLExtendedFetch - Extended fetch (fetch array of rows)” on page 162
- “SQLParamData - Get next parameter for which a data value is needed” on page 273
- “SQLPutData - Passing data value for a parameter” on page 304
- “SQLSetParam - Binds a parameter marker to a buffer” on page 328

## SQLExecute - Execute a statement

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLExecute() executes a statement, that is successfully prepared using SQLPrepare(), once or multiple times. The statement is executed using the current value of any application variables that are bound to parameter markers by SQLBindParameter() or SQLSetParam() .

### Syntax

```
SQLRETURN SQLExecute (SQLHSTMT hstmt);
```

### Function arguments

Table 47. SQLExecute arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle. There must not be an open cursor associated with hstmt, see "SQLFreeStmt - Free (or reset) a statement handle" on page 188 for more information.

### Usage

The SQL statement string can contain parameter markers. A parameter marker is represented by a "?" character, and is used to indicate a position in the statement where an application supplied value is to be substituted when SQLExecute() is called. This value can be obtained from:

- An application variable.  
SQLSetParam() or SQLBindParameter() is used to bind the application storage area to the parameter marker.
- A LOB value residing at the server referenced by a LOB locator.  
SQLBindParameter() or SQLSetParam() is used to bind a LOB locator to a parameter marker. The actual value of the LOB is kept at the server and does not need to be transferred to the application before being used as the input parameter value for another SQL statement.

You must bind all parameters before calling SQLExecute().

After the application processes the results from the SQLExecute() call, it can execute the statement again with new (or the same) parameter values.

A statement executed by SQLExecDirect() cannot be re-executed by calling SQLExecute(); SQLPrepare() must be called first.

If the prepared SQL statement is a query, SQLExecute() generates a cursor name, and open the cursor. If the application uses SQLSetCursorName() to associate a cursor name with the statement handle, DB2 ODBC associates the application generated cursor name with the internally generated one.

To execute a query more than once, the application must close the cursor by calling `SQLFreeStmt()` with the `SQL_CLOSE` option. There must not be an open cursor on the statement handle when calling `SQLExecute()`.

If a result set is generated, `SQLFetch()` or `SQLExtendedFetch()` retrieves the next row (or rows) of data into bound variables or LOB locators. Data can also be retrieved by calling `SQLGetData()` for any column that was not bound.

If the SQL statement is a positioned DELETE or a positioned UPDATE, the cursor referenced by the statement must be positioned on a row at the time `SQLExecute()` is called, and must be defined on a separate statement handle under the same connection handle.

If `SQLParamOptions()` is called to specify that an array of input parameter values is bound to each parameter marker, then the application needs to call `SQLExecDirect()` only once to process the entire array of input parameter values. If the executed statement returns multiple result sets (one for each set of input parameters), then `SQLMoreResults()` should be used to advance to the next result set when processing on the current result set is complete. See “SQLMoreResults - Determine if there are more result sets” on page 262 for more information.

## Return codes

- `SQL_SUCCESS`
- `SQL_SUCCESS_WITH_INFO`
- `SQL_ERROR`
- `SQL_INVALID_HANDLE`
- `SQL_NEED_DATA`
- `SQL_NO_DATA_FOUND`

`SQL_NEED_DATA` is returned when the application requests to input data-at-execution parameter values by calling `SQLParamData()` and `SQLPutData()`.

`SQL_SUCCESS` is returned if the SQL statement is a searched UPDATE or searched DELETE and no rows satisfy the search condition. Use `SQLRowCount()` to determine the number of rows in a table that were affected by an UPDATE, INSERT, or DELETE statement executed against the table, or a view of the table.

## Diagnostics

The `SQLSTATE`s for `SQLExecute()` include all those for `SQLExecDirect()` (see Table 46 on page 155) except for **S1009**, **S1090** and with the addition of the `SQLSTATE` in the table below.

Table 48. `SQLExecute SQLSTATE`s

SQLSTATE	Description	Explanation
40001	Transaction rollback.	The transaction to which this SQL statement belongs is rolled back due to a deadlock or timeout.
S1010	Function sequence error.	The specified <i>hstmt</i> is not in prepared state. <code>SQLExecute()</code> is called without first calling <code>SQLPrepare()</code> .

## Restrictions

None.

## Example

See “Example” on page 280.

## References

- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLExecute - Execute a statement” on page 159
- “SQLExtendedFetch - Extended fetch (fetch array of rows)” on page 162
- “SQLPrepare - Prepare a statement” on page 277
- “SQLFetch - Fetch next row” on page 169
- “SQLSetParam - Binds a parameter marker to a buffer” on page 328
- “SQLParamOptions - Specify an input array for a parameter” on page 275
- “SQLBindParameter - Binds a parameter marker to a buffer or LOB locator” on page 93

## SQLExtendedFetch - Extended fetch (fetch array of rows)

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>		
-----------------------	-----------------	--	--

SQLExtendedFetch() extends the function of SQLFetch() by returning a block of data containing multiple rows (called a *rowset*), in the form of an array, for each bound column. The size of the rowset is determined by the SQL\_ROWSET\_SIZE option on an SQLSetStmtOption() call.

To fetch one row of data at a time, an application should call SQLFetch().

For more description on block or array retrieval, see “Retrieving a result set into an array” on page 379.

### Syntax

```
SQLRETURN SQLExtendedFetch (SQLHSTMT          hstmt,
                             SQLUSMALLINT      fFetchType,
                             SQLINTEGER        irow,
                             SQLUINTEGER FAR   *pcrow,
                             SQLUSMALLINT FAR  *rgfRowStatus);
```

### Function arguments

Table 49 (Page 1 of 2). SQLExtendedFetch arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	Input	Statement handle.
SQLUSMALLINT	fFetchType	Input	Direction and type of fetch. DB2 ODBC only supports the fetch direction SQL_FETCH_NEXT; that is, forward only cursor direction. The next array (rowset) of data is retrieved.
SQLINTEGER	irow	Input	Reserved for future use.
SQLUINTEGER *	pcrow	Output	Number of the rows actually fetched. If an error occurs during processing, <i>pcrow</i> points to the ordinal position of the row (in the rowset) that precedes the row where the error occurred. If an error occurs retrieving the first row <i>pcrow</i> points to the value 0.



Table 49 (Page 2 of 2). SQLExtendedFetch arguments

Data type	Argument	Use	Description
SQLUSMALLINT *	rgfRowStatus	Output	<p>An array of status values. The number of elements must equal the number of rows in the rowset (as defined by the SQL_ROWSET_SIZE option). A status value for each row fetched is returned:</p> <ul style="list-style-type: none"> <li>• SQL_ROW_SUCCESS</li> </ul> <p>If the number of rows fetched is less than the number of elements in the status array (i.e. less than the rowset size), the remaining status elements are set to SQL_ROW_NOROW.</p> <p>DB2 ODBC cannot detect whether a row has been updated or deleted since the start of the fetch. Therefore, the following ODBC-defined status values are not reported:</p> <ul style="list-style-type: none"> <li>• SQL_ROW_DELETED</li> <li>• SQL_ROW_UPDATED</li> </ul>

## Usage

SQLExtendedFetch() performs an array fetch of a set of rows. An application specifies the size of the array by calling SQLSetStmtOption() with the SQL\_ROWSET\_SIZE option.

Before SQLExtendedFetch() is called the first time, the cursor is positioned before the first row. After SQLExtendedFetch() is called, the cursor is positioned on the row in the result set corresponding to the last row element in the rowset just retrieved.

For any columns in the result set that are bound using the SQLBindCol() function, DB2 ODBC converts the data for the bound columns as necessary and stores it in the locations bound to these columns. As mentioned in section “Retrieving a result set into an array” on page 379, the result set can be bound in a column-wise or row-wise fashion.

- For column-wise binding of application variables:

To bind a result set in column-wise fashion, an application specifies SQL\_BIND\_BY\_COLUMN for the SQL\_BIND\_TYPE statement option. (This is the default value.) Then the application calls the SQLBindCol() function.

When the application calls SQLExtendedFetch(), data for the first row is stored at the start of the buffer. Each subsequent row of data is stored at an offset of *cbValueMax* bytes (argument on SQLBindCol() call) or, if the associated C buffer type is fixed width (such as SQL\_C\_LONG), at an offset corresponding to that fixed length from the data for the previous row.

For each bound column, the number of bytes available to return for each element is stored in the *pcbValue* array buffer (deferred output argument on SQLBindCol()) buffer bound to the column. The number of bytes available to return for the first row of that column is stored at the start of the buffer, and the number of bytes available to return for each subsequent row is stored at an offset of *sizeof(SQLINTEGER)* bytes from the value for the previous row. If the

data in the column is NULL for a particular row, the associated element in the *pcbValue* array is set to `SQL_NULL_DATA`.

- For row-wise binding of application variables:

The application needs to first call `SQLSetStmtOption()` with the `SQL_BIND_TYPE` option, with the *vParam* argument set to the size of the structure capable of holding a single row of retrieved data and the associated data lengths for each column data value.

For each bound column, the first row of data is stored at the address given by the *rgbValue* supplied on the `SQLBindCol()` call for the column and each subsequent row of data at an offset of *vParam* bytes (used on the `SQLSetStmtOption()` call) from the data for the previous row.

For each bound column, the number of bytes available to return for the first row is stored at the address given by the *pcbValue* argument supplied on the `SQLBindCol()` call, and the number of bytes available to return for each subsequent row at an offset of *vParam* bytes from address containing the value for the previous row.

If `SQLExtendedFetch()` returns an error that applies to the entire rowset, the `SQL_ERROR` function return code is reported with the appropriate `SQLSTATE`. The contents of the rowset buffer are undefined and the cursor position is unchanged.

If an error occurs that applies to a single row:

- The corresponding element in the *rgbRowStatus* array for the row is set to `SQL_ROW_ERROR`
- An `SQLSTATE` of `01S01` is added to the list of errors that can be obtained using `SQLError()`
- Zero or more additional `SQLSTATE`s, describing the error for the current row, are added to the list of errors that can be obtained using `SQLError()`

An `SQL_ROW_ERROR` in the *rgbRowStatus* array only indicates that there was an error with the corresponding element; it does not indicate how many `SQLSTATE`s were generated. Therefore, `SQLSTATE 01S01` is used as a separator between the resulting `SQLSTATE`s for each row. DB2 ODBC continues to fetch the remaining rows in the rowset and returns `SQL_SUCCESS_WITH_INFO` as the function return code. After `SQLExtendedFetch()` returns, for each row encountering an error there is an `SQLSTATE` of `01S01` and zero or more additional `SQLSTATE`s indicating the errors for the current row, retrievable using `SQLError()`. Individual errors that apply to specific rows do not affect the cursor which continues to advance.

The number of elements in the *rgbRowStatus* array output buffer must equal the number of rows in the rowset (as defined by the `SQL_ROWSET_SIZE` statement option). If the number of rows fetched is less than the number of elements in the status array, the remaining status elements are set to `SQL_ROW_NOROW`.

An application cannot mix `SQLExtendedFetch()` with `SQLFetch()` calls.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE
- SQL\_NO\_DATA\_FOUND

## Diagnostics

Table 50 (Page 1 of 2). SQLExtendedFetch SQLSTATES

SQLSTATE	Description	Explanation
01004	Data truncated.	The data returned for one <i>or more</i> columns is truncated. (Function returns SQL_SUCCESS_WITH_INFO.)
01S01	Error in row.	An error occurred while fetching one or more rows. (Function returns SQL_SUCCESS_WITH_INFO.)
07002	Too many columns.	A column number specified in the binding for one or more columns is greater than the number of columns in the result set.  The application has used SQLSetColAttributes() to inform DB2 ODBC of the descriptor information of the result set, but it did not provide this for every column in the result set.
07006	Invalid conversion.	The data value could not be converted in a meaningful manner to the data type specified by <i>fCType</i> in SQLBindCol().
22002	Invalid output or indicator buffer specified.	The pointer value specified for the argument <i>pcbValue</i> in SQLBindCol() is a null pointer and the value of the corresponding column is null. There is no means to report SQL_NULL_DATA.
22003	Numeric value out of range.	Returning the numeric value (as numeric or string) for one or more columns causes the whole part of the number to be truncated either at the time of assignment or in computing an intermediate result.  A value from an arithmetic expression is returned which results in division by zero.
22005	Error in assignment.	A returned value is incompatible with the data type of the bound column.  A returned LOB locator was incompatible with the data type of the bound column.
22007	Invalid datetime format.	Conversion from character a string to a datetime format is indicated, but an invalid string representation or value is specified, or the value is an invalid date.  The value of a date, time, or timestamp does not conform to the syntax for the specified data type.
22008	Datetime field overflow.	Datetime field overflow occurred; for example, an arithmetic operation on a date or timestamp has a result that is not within the valid range of dates, or a datetime value cannot be assigned to a bound variable because it is too small.
22012	Division by zero is invalid.	A value from an arithmetic expression is returned which results in division by zero.
24000	Invalid cursor state.	The previous SQL statement executed on the statement handle is not a query.

## SQLExtendedFetch

Table 50 (Page 2 of 2). SQLExtendedFetch SQLSTATEs

SQLSTATE	Description	Explanation
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	SQLExtendedFetch() is called for an hstmt after SQLFetch() is called and before SQLFreeStmt() is called with the SQL_CLOSE option.  The function is called prior to calling SQLPrepare() or SQLExecDirect() for the hstmt.  The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1106	Fetch type out of range.	The value specified for the argument <i>fFetchType</i> is not recognized.
S1C00	Driver not capable.	DB2 ODBC or the data source does not support the conversion specified by the combination of the <i>fCType</i> in SQLBindCol() and the SQL data type of the corresponding column.  A call to SQLBindCol() is made for a column data type which is not supported by DB2 ODBC.  The specified fetch type is recognized, but not supported.

## Restrictions

None.

## Example

```
/* ... */
"SELECT deptnumb, deptname, id, name FROM staff, org \
    WHERE dept=deptnumb AND job = 'Mgr'";

/* Column-Wise */
SQLINTEGER    deptnumb[ROWSET_SIZE];

SQLCHAR       deptname[ROWSET_SIZE][15];
SQLINTEGER    deptname_1[ROWSET_SIZE];

SQLSMALLINT   id[ROWSET_SIZE];

SQLCHAR       name[ROWSET_SIZE][10];
SQLINTEGER    name_1[ROWSET_SIZE];
```

```

/* Row-Wise (Includes buffer for both column data and length) */
struct {
    SQLINTEGER    deptnumb_1; /* length */
    SQLINTEGER    deptnumb; /* value */
    SQLINTEGER    deptname_1;
    SQLCHAR       deptname[15];
    SQLINTEGER    id_1;
    SQLSMALLINT   id;
    SQLINTEGER    name_1;
    SQLCHAR       name[10];
}
R[ROWSET_SIZE];

SQLUSMALLINT    Row_Stat[ROWSET_SIZE];
SQLUIINTEGER    pcrow;
int             i;
/* ... */
/*****
/* Column-Wise Binding */
*****/
rc = SQLAllocStmt(hdbc, &hstmt);

rc = SQLSetStmtOption(hstmt, SQL_ROWSET_SIZE, ROWSET_SIZE);

rc = SQLExecDirect(hstmt, stmt, SQL_NTS);

rc = SQLBindCol(hstmt, 1, SQL_C_LONG, (SQLPOINTER) deptnumb, 0, NULL);

rc = SQLBindCol(hstmt, 2, SQL_C_CHAR, (SQLPOINTER) deptname, 15, deptname_1);

rc = SQLBindCol(hstmt, 3, SQL_C_SSHORT, (SQLPOINTER) id, 0, NULL);

rc = SQLBindCol(hstmt, 4, SQL_C_CHAR, (SQLPOINTER) name, 10, name_1);

/* Fetch ROWSET_SIZE rows at a time, and display */
printf("\nDEPTNUMB DEPTNAME          ID          NAME\n");
printf("-----\n");
while ((rc = SQLExtendedFetch(hstmt, SQL_FETCH_NEXT, 0, &pcrow, Row_Stat))
    == SQL_SUCCESS) {
    for (i = 0; i < pcrow; i++) {
        printf("%8ld %-14s %8ld %-9s\n", deptnumb[i], deptname[i], id[i], name[i]);
    }
    if (pcrow < ROWSET_SIZE)
        break;
}
/* endwhile */

if (rc != SQL_NO_DATA_FOUND && rc != SQL_SUCCESS)
    check_error(henv, hdbc, hstmt, rc, __LINE__, __FILE__);

rc = SQLFreeStmt(hstmt, SQL_DROP);

```

## SQLExtendedFetch

```
/*
*****
*/
/* Row-Wise Binding */
*****
rc = SQLAllocStmt(hdbc, &hstmt);
if (rc != SQL_SUCCESS)
    check_error(henv, hdbc, SQL_NULL_HSTMT, rc, __LINE__, __FILE__);

/* Set maximum number of rows to receive with each extended fetch */
rc = SQLSetStmtOption(hstmt, SQL_ROWSET_SIZE, ROWSET_SIZE);
if (rc != SQL_SUCCESS)
    check_error(henv, hdbc, hstmt, rc, __LINE__, __FILE__);

/*
* Set vparam to size of one row, used as offset for each bindcol
* rgbValue
*/
/* ie. &(R[0].deptnum) + vparam = &(R[1].deptnum) */
rc = SQLSetStmtOption(hstmt, SQL_BIND_TYPE, sizeof(R) / ROWSET_SIZE);

rc = SQLExecDirect(hstmt, stmt, SQL_NTS);

rc = SQLBindCol(hstmt, 1, SQL_C_LONG, (SQLPOINTER) &R[0].deptnum, 0,
                &R[0].deptnum_1);

rc = SQLBindCol(hstmt, 2, SQL_C_CHAR, (SQLPOINTER) R[0].deptname, 15,
                &R[0].deptname_1);

rc = SQLBindCol(hstmt, 3, SQL_C_SSHORT, (SQLPOINTER) &R[0].id, 0,
                &R[0].id_1);

rc = SQLBindCol(hstmt, 4, SQL_C_CHAR, (SQLPOINTER) R[0].name, 10, &R[0].name_1);

/* Fetch ROWSET_SIZE rows at a time, and display */
printf("\nDEPTNUMB DEPTNAME ID NAME\n");
printf("-----\n");
while ((rc = SQLExtendedFetch(hstmt, SQL_FETCH_NEXT, 0, &pcrow, Row_Stat))
      == SQL_SUCCESS) {
    for (i = 0; i < pcrow; i++) {
        printf("%8ld %-14s %8ld %-9s\n", R[i].deptnum, R[i].deptname,
            R[i].id, R[i].name);
    }
    if (pcrow < ROWSET_SIZE)
        break;
} /* endwhile */

if (rc != SQL_NO_DATA_FOUND && rc != SQL_SUCCESS)
    check_error(henv, hdbc, hstmt, rc, __LINE__, __FILE__);
/* Free handles, commit, exit */
/* ... */
```

## References

- “SQLExecute - Execute a statement” on page 159
- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLFetch - Fetch next row” on page 169

## SQLFetch - Fetch next row

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>	<b>X/OPEN CLI</b>	<b>ISO CLI</b>
-----------------------	-----------------	-------------------	----------------

SQLFetch() advances the cursor to the next row of the result set, and retrieves any bound columns.

Columns can be bound to:

- Application storage
- LOB locators

When SQLFetch() is called, the appropriate data transfer is performed, along with any data conversion if conversion was indicated when the column was bound. The columns can also be received individually after the fetch, by calling SQLGetData().

SQLFetch() can only be called after a result set is generated (using the same statement handle) by either executing a query, calling SQLGetTypeInfo() or calling a catalog function.

To retrieve multiple rows at a time, use SQLExtendedFetch().

### Syntax

```
SQLRETURN SQLFetch (SQLHSTMT hstmt);
```

### Function arguments

Table 51. SQLFetch arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle

### Usage

SQLFetch() can only be called after a result set is generated on the same statement handle. Before SQLFetch() is called the first time, the cursor is positioned before the start of the result set.

The number of application variables bound with SQLBindCol() must not exceed the number of columns in the result set or SQLFetch() fails.

If SQLBindCol() has not been called to bind any columns, then SQLFetch() does not return data to the application, but just advances the cursor. In this case, SQLGetData() can be called to obtain all of the columns individually. Data in unbound columns is discarded when SQLFetch() advances the cursor to the next row. For fixed length data types, or small variable length data types, binding columns provides better performance than using SQLGetData().

Columns can be bound to:

- Application storage

`SQLBindCol()` is used to bind application storage to the column. Data is transferred from the server to the application at fetch time. Length of the available data to return is also set.

- LOB locators

`SQLBindCol()` is used to bind LOB locators to the column. Only the LOB locator (4 bytes) is transferred from the server to the application at fetch time.

When an application receives a locator, it can use the locator in `SQLGetSubString()`, `SQLGetPosition()`, `SQLGetLength()`, or as the value of a parameter marker in another SQL statement. `SQLGetSubString()` can either return another locator, or the data itself. All locators remain valid until the end of the transaction in which they are created (even when the cursor moves to another row), or until they are freed using the `FREE LOCATOR` statement.

If LOB values are too large to retrieve in one fetch, they can be retrieved in pieces by either using `SQLGetData()` (which can be used for any column type), or by binding a LOB locator, and using `SQLGetSubString()`.

If any bound storage buffers are not large enough to hold the data returned by `SQLFetch()`, the data is truncated. If character data is truncated, `SQL_SUCCESS_WITH_INFO` is returned, and an `SQLSTATE` is generated indicating truncation. The `SQLBindCol()` deferred output argument *pcbValue* contains the actual length of the column data retrieved from the server. The application should compare the actual output length to the input buffer length (*pcbValue* and *cbValueMax* arguments from `SQLBindCol()`) to determine which character columns are truncated.

Truncation of numeric data types is reported as a warning if the truncation involves digits to the right of the decimal point. If truncation occurs to the left of the decimal point, an error is returned (see the diagnostics section).

Truncation of graphic data types is treated the same as character data types, except that the *rgbValue* buffer is filled to the nearest multiple of two bytes that is still less than or equal to the *cbValueMax* specified in `SQLBindCol()`. Graphic (DBCS) data transferred between DB2 ODBC and the application is not null-terminated if the C buffer type is `SQL_C_CHAR`. If the buffer type is `SQL_C_DBCHAR`, then null-termination of graphic data does occur.

Truncation is also affected by the `SQL_MAX_LENGTH` statement option. The application can specify that DB2 ODBC should not report truncation by calling `SQLSetStmtOption()` with `SQL_MAX_LENGTH` and a value for the maximum length to return for any one column, and by allocating an *rgbValue* buffer of the same size (plus the null-terminator). If the column data is larger than the set maximum length, `SQL_SUCCESS` is returned and the maximum length, not the actual length is returned in *pcbValue*.

When all the rows are retrieved from the result set, or the remaining rows are not needed, `SQLFreeStmt()` should be called to close the cursor and discard the remaining data and associated resources.

To retrieve multiple rows at a time, use `SQLExtendedFetch()`. An application cannot mix `SQLFetch()` with `SQLExtendedFetch()` calls on the same statement handle.



## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE
- SQL\_NO\_DATA\_FOUND

SQL\_NO\_DATA\_FOUND is returned if there are no rows in the result set, or previous SQLFetch() calls have fetched all the rows from the result set.

If all the rows were fetched, the cursor is positioned after the end of the result set.

## Diagnostics

Table 52 (Page 1 of 2). SQLFetch SQLSTATES

SQLSTATE	Description	Explanation
01004	Data truncated.	The data returned for one or more columns is truncated. String values or numeric values are right truncated. (SQL_SUCCESS_WITH_INFO is returned if no error occurred.)
07002	Too many columns.	A column number specified in the binding for one or more columns is greater than the number of columns in the result set.  The application used SQLSetColAttributes() to inform DB2 ODBC of the descriptor information of the result set, but it did not provide this for every column in the result set.
07006	Invalid conversion.	The data value cannot be converted in a meaningful manner to the data type specified by <i>fCType</i> in SQLBindCol()
22002	Invalid output or indicator buffer specified.	The pointer value specified for the argument <i>pcbValue</i> in SQLBindCol() is a null pointer and the value of the corresponding column is null. There is no means to report SQL_NULL_DATA.
22003	Numeric value out of range.	Returning the numeric value (as numeric or string) for one or more columns causes the whole part of the number to be truncated either at the time of assignment or in computing an intermediate result.  A value from an arithmetic expression is returned which results in division by zero.  <b>Note:</b> The associated cursor is undefined if this error is detected by DB2 for OS/390. If the error is detected by DB2 UDB or by other IBM RDBMSs, the cursor remains open and continues to advance on subsequent fetch calls.
22005	Error in assignment.	A returned value is incompatible with the data type of binding.  A returned LOB locator is incompatible with the data type of the bound column.
22007	Invalid datetime format.	Conversion from character a string to a datetime format is indicated, but an invalid string representation or value is specified, or the value is an invalid date.  The value of a date, time, or timestamp does not conform to the syntax for the specified data type.

## SQLFetch

Table 52 (Page 2 of 2). SQLFetch SQLSTATES

SQLSTATE	Description	Explanation
22008	Datetime field overflow.	Datetime field overflow occurred; for example, an arithmetic operation on a date or timestamp has a result that is not within the valid range of dates, or a datetime value cannot be assigned to a bound variable because it is too small.
22012	Division by zero is invalid.	A value from an arithmetic expression is returned which results in division by zero.
24000	Invalid cursor state.	The previous SQL statement executed on the statement handle is not a query.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
54028	The maximum number of concurrent LOB handles has been reached.	Maximum LOB locator assigned. The maximum number of concurrent LOB locators has been reached. A new locator can not be assigned.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1002	Invalid column number.	The specified column is less than 0 or greater than the number of result columns. The specified column is 0, but DB2 ODBC does not support ODBC bookmarks ( <i>icol</i> = 0). SQLExtendedFetch() is called for this result set.
S1010	Function sequence error.	SQLFetch() is called for an <i>hstmt</i> after SQLExtendedFetch() is called and before SQLFreeStmt() had been called with the SQL_CLOSE option. The function is called prior to calling SQLPrepare() or SQLExecDirect() for the <i>hstmt</i> . The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1C00	Driver not capable.	DB2 ODBC or the data source does not support the conversion specified by the combination of the <i>fCType</i> in SQLBindCol() and the SQL data type of the corresponding column. A call to SQLBindCol() was made for a column data type which is not supported by DB2 ODBC.

## Restrictions

None.

## Example

```

/* ... */
/*****
** main
*****/
int
main( int argc, char * argv[] )
{
    SQLHENV      henv;
    SQLHDBC      hdbc;
    SQLHSTMT     hstmt;
    SQLRETURN    rc;
    SQLCHAR      sqlstmt[] = "SELECT deptname, location from org where
                                division = 'Eastern'";

    struct { SQLINTEGER ind;
             SQLCHAR  s[15];
             } deptname, location;

/* macro to initialize server, uid and pwd */
    INIT_UID_PWD;

    rc = SQLAllocEnv(&henv); /* allocate an environment handle */
    if (rc != SQL_SUCCESS)
        return (terminate(henv, rc));
    rc = DBconnect(henv, &hdbc);/* allocate a connect handle, and connect */
    if (rc != SQL_SUCCESS)
        return (terminate(henv, rc));
    rc = SQLAllocStmt(hdbc, &hstmt);
    rc = SQLExecDirect(hstmt, sqlstmt, SQL_NTS);
    rc = SQLBindCol(hstmt, 1, SQL_C_CHAR, (SQLPOINTER) deptname.s, 15,
                    &deptname.ind);
    rc = SQLBindCol(hstmt, 2, SQL_C_CHAR, (SQLPOINTER) location.s, 15,
                    &location.ind);

    printf("Departments in Eastern division:\n");
    printf("DEPTNAME      Location\n");
    printf("-----\n");

    while ((rc = SQLFetch(hstmt)) == SQL_SUCCESS) {
        printf("%-14.14s %-14.14s \n", deptname.s, location.s);
    }
    if (rc != SQL_NO_DATA_FOUND)
        check_error(henv, hdbc, hstmt, rc, __LINE__, __FILE__);

    rc = SQLFreeStmt(hstmt, SQL_DROP);
    rc = SQLTransact(henv, hdbc, SQL_COMMIT);

    printf("Disconnecting ....\n");

    rc = SQLDisconnect(hdbc);
    rc = SQLFreeConnect(hdbc);
    rc = SQLFreeEnv(henv);
    if (rc != SQL_SUCCESS)
        return (terminate(henv, rc));
}
/* ... */

```

### References

- “SQLExtendedFetch - Extended fetch (fetch array of rows)” on page 162
- “SQLExecute - Execute a statement” on page 159
- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLGetData - Get data from a column” on page 199

## SQLForeignKeys - Get the list of foreign key columns

### Purpose

Specification:	ODBC 1.0		
----------------	----------	--	--

SQLForeignKeys() returns information about foreign keys for the specified table. The information is returned in an SQL result set which can be processed using the same functions that are used to retrieve a result generated by a query.

### Syntax

```
SQLRETURN SQLForeignKeys (SQLHSTMT          hstmt,
                          SQLCHAR           FAR *szPkCatalogName,
                          SQLSMALLINT      cbPkCatalogName,
                          SQLCHAR           FAR *szPkSchemaName,
                          SQLSMALLINT      cbPkSchemaName,
                          SQLCHAR           FAR *szPkTableName,
                          SQLSMALLINT      cbPkTableName,
                          SQLCHAR           FAR *szFkCatalogName,
                          SQLSMALLINT      cbFkCatalogName,
                          SQLCHAR           FAR *szFkSchemaName,
                          SQLSMALLINT      cbFkSchemaName,
                          SQLCHAR           FAR *szFkTableName,
                          SQLSMALLINT      cbFkTableName);
```

### Function arguments

Table 53. SQLForeignKeys arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle.
SQLCHAR *	szPkCatalogName	input	Catalog qualifier of the primary key table. This must be a NULL pointer or a zero length string.
SQLSMALLINT	cbPkCatalogName	input	Length of <i>szPkCatalogName</i> . This must be set to 0.
SQLCHAR *	szPkSchemaName	input	Schema qualifier of the primary key table.
SQLSMALLINT	cbPkSchemaName	input	Length of <i>szPkSchemaName</i> .
SQLCHAR *	szPkTableName	input	Name of the table name containing the primary key.
SQLSMALLINT	cbPkTableName	input	Length of <i>szPkTableName</i> .
SQLCHAR *	szFkCatalogName	input	Catalog qualifier of the table containing the foreign key. This must be a NULL pointer or a zero length string.
SQLSMALLINT	cbFkCatalogName	input	Length of <i>szFkCatalogName</i> . This must be set to 0.
SQLCHAR *	szFkSchemaName	input	Schema qualifier of the table containing the foreign key.
SQLSMALLINT	cbFkSchemaName	input	Length of <i>szFkSchemaName</i> .
SQLCHAR *	szFkTableName	input	Name of the table containing the foreign key.
SQLSMALLINT	cbFkTableName	input	Length of <i>szFkTableName</i> .

## Usage

If *szPkTableName* contains a table name, and *szFkTableName* is an empty string, `SQLForeignKeys()` returns a result set containing the primary key of the specified table and all of the foreign keys (in other tables) that refer to it.

If *szFkTableName* contains a table name, and *szPkTableName* is an empty string, `SQLForeignKeys()` returns a result set containing all of the foreign keys in the specified table and the primary keys (in other tables) to which they refer.

If both *szPkTableName* and *szFkTableName* contain table names, `SQLForeignKeys()` returns the foreign keys in the table specified in *szFkTableName* that refer to the primary key of the table specified in *szPkTableName*. This should be one key at the most.

If the schema qualifier argument associated with a table name is not specified, then the schema name defaults to the one currently in effect for the current connection.

Table 54 lists the columns of the result set generated by the `SQLForeignKeys()` call. If the foreign keys associated with a primary key are requested, the result set is ordered by `FKTABLE_CAT`, `FKTABLE_SCHEM`, `FKTABLE_NAME`, and `ORDINAL_POSITION`. If the primary keys associated with a foreign key are requested, the result set is ordered by `PKTABLE_CAT`, `PKTABLE_SCHEM`, `PKTABLE_NAME`, and `ORDINAL_POSITION`.

The `VARCHAR` columns of the catalog functions result set are declared with a maximum length attribute of 128 to be consistent with SQL92 limits. Since DB2 names are less than 128, the application can choose to always set aside 128 characters (plus the null-terminator) for the output buffer, or alternatively, call `SQLGetInfo()` with the `SQL_MAX_CATALOG_NAME_LEN`, `SQL_MAX_SCHEMA_NAME_LEN`, `SQL_MAX_TABLE_NAME_LEN`, and `SQL_MAX_COLUMN_NAME_LEN` to determine respectively the actual lengths of the associated `TABLE_CAT`, `TABLE_SCHEM`, `TABLE_NAME`, and `COLUMN_NAME` columns supported by the connected DBMS.

Although new columns might be added and the names of the existing columns changed in future releases, the position of the current columns does not change.

Table 54 (Page 1 of 2). Columns returned by `SQLForeignKeys`

Column number/name	Data type	Description
1 PKTABLE_CAT	VARCHAR(128)	This is always NULL.
2 PKTABLE_SCHEM	VARCHAR(128)	The name of the schema containing PKTABLE_NAME.
3 PKTABLE_NAME	VARCHAR(128) NOT NULL	Name of the table containing the primary key.
4 PKCOLUMN_NAME	VARCHAR(128) NOT NULL	Primary key column name.
5 FKTABLE_CAT	VARCHAR(128)	This is always NULL.
6 FKTABLE_SCHEM	VARCHAR(128)	The name of the schema containing FKTABLE_NAME.
7 FKTABLE_NAME	VARCHAR(128) NOT NULL	The name of the table containing the foreign key.
8 FKCOLUMN_NAME	VARCHAR(128) NOT NULL	Foreign key column name.

Table 54 (Page 2 of 2). Columns returned by SQLForeignKeys

Column number/name	Data type	Description
9 KEY_SEQ	SMALLINT NOT NULL	The ordinal position of the column in the key, starting at 1.
10 UPDATE_RULE	SMALLINT	Action to be applied to the foreign key when the SQL operation is UPDATE: <ul style="list-style-type: none"> <li>• SQL_RESTRICT</li> <li>• SQL_NO_ACTION</li> </ul> <p>The update rule for IBM DB2 DBMSs is always either RESTRICT or SQL_NO_ACTION. However, ODBC applications might encounter the following UPDATE_RULE values when connected to non-IBM RDBMSs:</p> <ul style="list-style-type: none"> <li>• SQL_CASCADE</li> <li>• SQL_SET_NULL</li> </ul>
11 DELETE_RULE	SMALLINT	Action to be applied to the foreign key when the SQL operation is DELETE: <ul style="list-style-type: none"> <li>• SQL_CASCADE</li> <li>• SQL_NO_ACTION</li> <li>• SQL_RESTRICT</li> <li>• SQL_SET_DEFAULT</li> <li>• SQL_SET_NULL</li> </ul>
12 FK_NAME	VARCHAR(128)	Foreign key identifier. NULL if not applicable to the data source.
13 PK_NAME	VARCHAR(128)	Primary key identifier. NULL if not applicable to the data source.
14 DEFERRABILITY	SMALLINT	One of the following: <ul style="list-style-type: none"> <li>• SQL_INITIALLY_DEFERRED</li> <li>• SQL_INITIALLY_IMMEDIATE</li> <li>• SQL_NOT_DEFERRABLE</li> </ul> <p>DB2 ODBC returns NULL.</p>

**Note:** The column names used by DB2 ODBC follow the X/Open CLI CAE specification style. The column types, contents and order are identical to those defined for the SQLForeignKeys() result set in ODBC.

# DB2 ODBC applications that issue SQLForeignKeys() against a DB2 for OS/390 server, Version 5 or later,  
# should expect the result set columns listed in the table above. Revision bars identify the new and changed  
# columns.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 55 (Page 1 of 2). SQLForeignKeys SQLSTATEs

SQLSTATE	Description	Explanation
24000	Invalid cursor state.	A cursor is already opened on the statement handle.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.

## SQLForeignKeys

Table 55 (Page 2 of 2). SQLForeignKeys SQLSTATEs

SQLSTATE	Description	Explanation
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	The arguments <i>szPkTableName</i> and <i>szFkTableName</i> are both NULL pointers.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1014	No more handles.	DB2 ODBC is not able to allocate a handle due to internal resources.
S1090	Invalid string or buffer length.	The value of one of the name length arguments was less than 0, but not equal SQL_NTS.  The length of the table or owner name is greater than the maximum length supported by the server. See "SQLGetInfo - Get general information" on page 219.
S1C00	Driver not capable.	DB2 ODBC does not support <i>catalog</i> as a qualifier for table name.

## Restrictions

None.

## Example

```
/******  
/* DB2 for OS/390 example:                               */  
/*     Invokes SQLForeignKeys against PARENT Table. Find all */  
/*     tables that contain foreign keys on PARENT.         */  
/******  
#include <stdio.h>  
#include <string.h>  
#include <stdlib.h>  
#include <sqlca.h>  
#include "cli.h"  
#include "sqlcli1.h"  
#include "sqlcli1.h"  
  
int main( )  
{  
    SQLHENV      hEnv   = SQL_NULL_HENV;  
    SQLHDBC      hDbc   = SQL_NULL_HDBC;  
    SQLHSTMT     hStmt  = SQL_NULL_HSTMT;  
    SQLRETURN    rc     = SQL_SUCCESS;  
    SQLINTEGER   RETCODE = 0;  
    char         pTable [200];  
    char         *pDSN = "STLEC1";  
    SQLSMALLINT  update_rule;  
    SQLSMALLINT  delete_rule;  
    SQLINTEGER   update_rule_ind;  
    SQLINTEGER   delete_rule_ind;  
    char         update [25];  
    char         delet [25];
```



```

typedef struct varchar // define VARCHAR type
{
    SQLSMALLINT length;
    SQLCHAR name [128];
    SQLINTEGER ind;

} VARCHAR;
VARCHAR pktable_schem;
VARCHAR pktable_name;
VARCHAR pkcolumn_name;
VARCHAR fhtable_schem;
VARCHAR fhtable_name;
VARCHAR fkcolumn_name;

(void) printf ("**** Entering CLIP02.\n\n");

/*****
/* Allocate Environment Handle */
*****/

RETCODE = SQLAllocEnv(&hEnv);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Allocate Connection Handle to DSN */
*****/

RETCODE = SQLAllocConnect(hEnv,
                          &hDbc);

if( RETCODE != SQL_SUCCESS ) // Could not get a connect handle
    goto dberror;

/*****
/* CONNECT TO data source (STLEC1) */
*****/

RETCODE = SQLConnect(hDbc, // Connect handle
                    (SQLCHAR *) pDSN, // DSN
                    SQL_NTS, // DSN is nul-terminated
                    NULL, // Null UID
                    0,
                    NULL, // Null Auth string
                    0);

if( RETCODE != SQL_SUCCESS ) // Connect failed
    goto dberror;

/*****
/* Allocate Statement Handle */
*****/

rc = SQLAllocStmt (hDbc,
                  &hStmt);

if (rc != SQL_SUCCESS)
    goto exit;

```

## SQLForeignKeys

```
/* *****  
/* Invoke SQLForeignKeys against PARENT Table, specifying NULL */  
/* for table with foreign key. */  
/* *****  
  
rc = SQLForeignKeys (hStmt,  
                    NULL,  
                    0,  
                    (SQLCHAR *) "ADMFO01",  
                    SQL_NTS,  
                    (SQLCHAR *) "PARENT",  
                    SQL_NTS,  
                    NULL,  
                    0,  
                    NULL,  
                    SQL_NTS,  
                    NULL,  
                    SQL_NTS);  
  
if (rc != SQL_SUCCESS)  
{  
    (void) printf ("**** SQLForeignKeys Failed.\n");  
    goto dberror;  
}  
  
/* *****  
/* Bind following columns of answer set: */  
/* */  
/* 2) phtable_schem */  
/* 3) phtable_name */  
/* 4) pkcolumn_name */  
/* 6) fhtable_schem */  
/* 7) fhtable_name */  
/* 8) fkcolumn_name */  
/* 10) update_rule */  
/* 11) delete_rule */  
/* */  
/* *****  
  
rc = SQLBindCol (hStmt,          // bind phtable_schem  
                2,  
                SQL_C_CHAR,  
                (SQLPOINTER) phtable_schem.name,  
                128,  
                &phtable_schem.ind);  
  
rc = SQLBindCol (hStmt,          // bind phtable_name  
                3,  
                SQL_C_CHAR,  
                (SQLPOINTER) phtable_name.name,  
                128,  
                &phtable_name.ind);  
  
rc = SQLBindCol (hStmt,          // bind pkcolumn_name  
                4,  
                SQL_C_CHAR,  
                (SQLPOINTER) pkcolumn_name.name,  
                128,  
                &pkcolumn_name.ind);
```

```

rc = SQLBindCol (hStmt,          // bind fktable_schem
                6,
                SQL_C_CHAR,
                (SQLPOINTER) fktable_schem.name,
                128,
                &fktable_schem.ind);

rc = SQLBindCol (hStmt,          // bind fktable_name
                7,
                SQL_C_CHAR,
                (SQLPOINTER) fktable_name.name,
                128,
                &fktable_name.ind);

rc = SQLBindCol (hStmt,          // bind fkcolumn_name
                8,
                SQL_C_CHAR,
                (SQLPOINTER) fkcolumn_name.name,
                128,
                &fkcolumn_name.ind);

rc = SQLBindCol (hStmt,          // bind update_rule
                10,
                SQL_C_SHORT,
                (SQLPOINTER) &update_rule;
                0,
                &update_rule_ind);

rc = SQLBindCol (hStmt,          // bind delete_rule
                11,
                SQL_C_SHORT,
                (SQLPOINTER) &delete_rule,
                0,
                &delete_rule_ind);

/*****
/* Retrieve all tables with foreign keys defined on PARENT */
*****/

while ((rc = SQLFetch (hStmt)) == SQL_SUCCESS)
{
    (void) printf ("**** Primary Table Schema is %s. Primary Table Name is %s.\n",
                  pktable_schem.name, pktable_name.name);
    (void) printf ("**** Primary Table Key Column is %s.\n",
                  pkcolumn_name.name);
    (void) printf ("**** Foreign Table Schema is %s. Foreign Table Name is %s.\n",
                  fktable_schem.name, fktable_name.name);
    (void) printf ("**** Foreign Table Key Column is %s.\n",
                  fkcolumn_name.name);

    if (update_rule == SQL_RESTRIC) // isolate update rule
        strcpy (update, "RESTRICT");
    else
        if (update_rule == SQL_CASCADE)
            strcpy (update, "CASCADE");
        else
            strcpy (update, "SET NULL");
}

```

## SQLForeignKeys

```
    if (delete_rule == SQL_RESTRIC)        // isolate delete rule
        strcpy (delet, "RESTRICT");
    else
    if (delete_rule == SQL_CASCADE)
        strcpy (delet, "CASCADE");
    else
    if (delete_rule == SQL_NO_ACTION)
        strcpy (delet, "NO ACTION");
    else
        strcpy (delet, "SET NULL");

    (void) printf ("**** Update Rule is %s. Delete Rule is %s.\n",
                  update, delet);
}

/*****
/* Deallocate Statement Handle
*****/

rc = SQLFreeStmt (hStmt,
                  SQL_DROP);

/*****
/* DISCONNECT from data source
*****/

RETCODE = SQLDisconnect(hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Deallocate Connection Handle
*****/

RETCODE = SQLFreeConnect (hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Free Environment Handle
*****/

RETCODE = SQLFreeEnv (hEnv);

if (RETCODE == SQL_SUCCESS)
    goto exit;

dberror:
RETCODE=12;

exit:

(void) printf ("**** Exiting CLIP02.\n\n");

return RETCODE;
}
```

## References

- “SQLPrimaryKeys - Get primary key columns of a table” on page 285
- “SQLStatistics - Get index and statistics information for a base table” on page 344

## SQLFreeConnect - Free connection handle

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>	<b>X/OPEN CLI</b>	<b>ISO CLI</b>
-----------------------	-----------------	-------------------	----------------

SQLFreeConnect() invalidates and frees the connection handle. All DB2 ODBC resources associated with the connection handle are freed.

SQLDisconnect() must be called before calling this function.

### Syntax

```
SQLRETURN SQLFreeConnect (SQLHDBC          hdbc);
```

### Function arguments

Table 56. SQLFreeConnect arguments

Data type	Argument	Use	Description
SQLHDBC	<i>hdbc</i>	input	Connection handle

### Usage

If this function is called when a connection still exists, SQL\_ERROR is returned, and the connection handle remains valid.

To continue termination, call SQLFreeEnv(), or, if a new connection handle is required, call SQLAllocConnect().

### Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE
- SQL\_SUCCESS\_WITH\_INFO

### Diagnostics

Table 57. SQLFreeConnect SQLSTATEs

SQLSTATE	Description	Explanation
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called prior to SQLDisconnect() for the <i>hdbc</i> .
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.

## Restrictions

None.

## Example

See “Example” on page 80.

## References

- “SQLDisconnect - Disconnect from a data source” on page 140
- “SQLFreeEnv - Free environment handle” on page 186

## SQLFreeEnv - Free environment handle

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLFreeEnv() invalidates and frees the environment handle. All DB2 ODBC resources associated with the environment handle are freed.

SQLFreeConnect() must be called before calling this function.

This function is the last DB2 ODBC step an application needs to do before terminating.

### Syntax

```
SQLRETURN SQLFreeEnv (SQLHENV henv);
```

### Function arguments

Table 58. SQLFreeEnv arguments

Data type	Argument	Use	Description
SQLHENV	<i>henv</i>	input	Environment handle

### Usage

If this function is called when there is still a valid connection handle, SQL\_ERROR is returned, and the environment handle remains valid.

The number of SQLFreeEnv() calls must equal the number of SQLAllocEnv() calls before the environment information is reset.

### Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

### Diagnostics

Table 59. SQLFreeEnv SQLSTATEs

SQLSTATE	Description	Explanation
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	There is an <i>hdbc</i> which is in allocated or connected state. Call SQLDisconnect() and SQLFreeConnect() for the <i>hdbc</i> before calling SQLFreeEnv().
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.



## Restrictions

None.

## Example

See “Example” on page 80.

## References

- “SQLFreeConnect - Free connection handle” on page 184

---

## SQLFreeStmt - Free (or reset) a statement handle

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLFreeStmt() ends processing on the statement referenced by the statement handle. Use this function to:

- Close a cursor
- Drop the statement handle and free the DB2 ODBC resources associated with the statement handle.

SQLFreeStmt() is called after executing an SQL statement and processing the results.

### Syntax

```
SQLRETURN SQLFreeStmt (SQLHSTMT hstmt,
                        SQLUSMALLINT fOption);
```

### Function arguments

Table 60. SQLFreeStmt arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle
SQLUSMALLINT	<i>fOption</i>	input	Option which specified the manner of freeing the statement handle. The option must have one of the following values: <ul style="list-style-type: none"> <li>• SQL_CLOSE</li> <li>• SQL_DROP</li> <li>• SQL_UNBIND</li> <li>• SQL_RESET_PARAMS</li> </ul>

### Usage

SQLFreeStmt() can be called with the following options:

**SQL\_CLOSE** The cursor (if any) associated with the statement handle (*hstmt*) is closed and all pending results are discarded. The application can reopen the cursor by calling SQLExecute() or SQLExecDirect() with the same or different values in the application variables (if any) that are bound to *hstmt*. The cursor name is retained until the statement handle is dropped or the next successful SQLSetCursorName() call. If a cursor is not associated with the statement handle, this option has no effect (no warning or error is generated).

**SQL\_DROP** DB2 ODBC resources associated with the input statement handle are freed, and the handle is invalidated. The open cursor, if any, is closed and all pending results are discarded.

**SQL\_UNBIND** All the columns bound by previous `SQLBindCol()` calls on this statement handle are released (the association between application variables or file references and result set columns is broken).

**SQL\_RESET\_PARAMS**

All the parameters set by previous `SQLBindParameter()` calls on this statement handle are released (the association between application variables or file references and parameter markers in the SQL statement for the statement handle is broken).

`SQLFreeStmt()` has no effect on LOB locators. Call `SQLExecDirect()` with the `FREE LOCATOR` statement to free a locator. See “Using large objects” on page 383 for more information on using LOBs.

You can reuse a statement handle to execute a different statement. If the handle is:

- Associated with a query, catalog function, or `SQLGetTypeInfo()`, you must close the cursor.
- Bound with a different number or type of parameters, the parameters must be reset.
- Bound with a different number or type of column bindings, the columns must be unbound.

Alternatively, you can drop the statement handle and allocate a new one.

## Return codes

- `SQL_SUCCESS`
- `SQL_SUCCESS_WITH_INFO`
- `SQL_ERROR`
- `SQL_INVALID_HANDLE`

`SQL_SUCCESS_WITH_INFO` is not returned if *fOption* is set to `SQL_DROP`, since there would be no statement handle to use when `SQLError()` is called.

## Diagnostics

Table 61. `SQLFreeStmt` `SQLSTATEs`

SQLSTATE	Description	Explanation
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute ( <code>SQLParamData()</code> , <code>SQLPutData()</code> ) operation.
S1092	Option type out of range.	The value specified for the argument <i>fOption</i> is not <code>SQL_CLOSE</code> , <code>SQL_DROP</code> , <code>SQL_UNBIND</code> , or <code>SQL_RESET_PARAMS</code> .
S1506	Error closing a file.	Error encountered while trying to close a temporary file.

## SQLFreeStmt

### Restrictions

None.

### Example

See “Example” on page 173.

### References

- “SQLAllocStmt - Allocate a statement handle” on page 85
- “SQLBindParameter - Binds a parameter marker to a buffer or LOB locator” on page 93
- “SQLExtendedFetch - Extended fetch (fetch array of rows)” on page 162
- “SQLFetch - Fetch next row” on page 169
- “SQLSetParam - Binds a parameter marker to a buffer” on page 328

## SQLGetConnectOption - Returns current setting of a connect option

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>	<b>X/OPEN CLI</b>	
-----------------------	-----------------	-------------------	--

SQLGetConnectOption() returns the current settings for the specified connection option.

These options are set using the SQLSetConnectOption() function.

### Syntax

```
SQLRETURN SQLGetConnectOption (
                                SQLHDBC          hdbc,
                                SQLUSMALLINT     fOption,
                                SQLPOINTER       pvParam);
```

### Function arguments

Table 62. SQLGetConnectOption arguments

Data type	Argument	Use	Description
HDBC	hdbc	input	Connection handle.
SQLUSMALLINT	fOption	input	Option to set. See Table 120 on page 317 for the complete list of connection options and their descriptions.
SQLPOINTER	pvParam	input/output	Value associated with <i>fOption</i> . Depending on the value of <i>fOption</i> , this can be a 32-bit integer value, or a pointer to a null terminated character string. The maximum length of any character string returned is SQL_MAX_OPTION_STRING_LENGTH bytes (excluding the null-terminator).

### Usage

If SQLGetConnectOption() is called, and the specified *fOption* has not been set using SQLSetConnectOption and does not have a default, then SQLGetConnectOption() returns SQL\_NO\_DATA\_FOUND.

Although SQLSetConnectOption() can be used to set statement options, SQLGetConnectOption() cannot be used to retrieve statement options, use SQLGetStmtOption() instead.

For a list of valid connect options, see Table 120 on page 317, in the function description for SQLSetConnectOption().

### Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

### Diagnostics

Table 63. SQLGetConnectOption SQLSTATEs

SQLSTATE	Description	Explanation
40003 08S01	Communication link failure.	The function is called after the communication link source to which DB2 ODBC is connected, failed during the processing of a previous request.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	The <i>pvParam</i> argument is NULL.
S1092	Option type out of range.	An invalid <i>fOption</i> value is specified.
S1C00	Driver not capable.	The <i>fOption</i> is recognized, but is not supported.

### Restrictions

None.

### Example

```
/* ... */
rc = SQLGetConnectOption(hdbc, SQL_AUTOCOMMIT, &autocommit);
printf("Autocommit is: ");
if (autocommit == SQL_AUTOCOMMIT_ON)
    printf("ON\n");
else
    printf("OFF\n");
/* ... */
```

### References

- “SQLSetConnectOption - Set connection option” on page 316
- “SQLSetStmtOption - Set statement option” on page 334
- “SQLGetStmtOption - Returns current setting of a statement option” on page 249

## SQLGetCursorName - Get cursor name

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLGetCursorName() returns the cursor name associated with the input statement handle. If a cursor name is explicitly set by calling SQLSetCursorName(), this name is returned; otherwise, an implicitly generated name is returned.

### Syntax

```
SQLRETURN SQLGetCursorName (SQLHSTMT hstmt,
                             SQLCHAR FAR *szCursor,
                             SQLSMALLINT cbCursorMax,
                             SQLSMALLINT FAR *pcbCursor);
```

### Function arguments

Table 64. SQLGetCursorname arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle
SQLCHAR *	<i>szCursor</i>	output	Cursor name
SQLSMALLINT	<i>cbCursorMax</i>	input	Length of buffer <i>szCursor</i>
SQLSMALLINT *	<i>pcbCursor</i>	output	Number of bytes available to return for <i>szCursor</i>

### Usage

SQLGetCursorName() returns the cursor name set explicitly with SQLSetCursorName(), or if no name is set, it returns the cursor name internally generated by DB2 ODBC.

If a name is set explicitly using SQLSetCursorName(), this name is returned until the statement is dropped, or until another explicit name is set.

Internally generated cursor names always begin with SQLCUR or SQL\_CUR. For query result sets, DB2 ODBC also reserves SQLCURQRS as a cursor name prefix. Cursor names are always 18 characters or less, and are always unique within a connection.

### Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 65. SQLGetCursorName SQLSTATEs

SQLSTATE	Description	Explanation
01004	Data truncated.	The cursor name returned in <i>szCursor</i> is longer than the value in <i>cbCursorMax</i> , and is truncated to <i>cbCursorMax</i> - 1 bytes. The argument <i>pcbCursor</i> contains the length of the full cursor name available for return. The function returns SQL_SUCCESS_WITH_INFO.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1092	Option type out of range.	The value specified for the argument <i>hstmt</i> is not valid.
S1015	No cursor name available.	There is no open cursor on the statement handle specified by <i>hstmt</i> and no cursor name is set with SQLSetCursorName().
S1090	Invalid string or buffer length.	The value specified for the argument <i>cbCursorMax</i> is less than 0.

## Restrictions

ODBC generated cursor names begin with SQL\_CUR. X/Open CLI generated cursor names begin with either SQLCUR or SQL\_CUR. DB2 ODBC also generates a cursor name that begins with SQLCUR or SQL\_CUR.

## Example

```

/*****
/* DB2 for OS/390 example: */
/* Performs a positioned update on a column of a cursor. */
*****/

#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <sqlca.h>
#include "sqlcli1.h"

int main( )
{
    SQLHENV          hEnv      = SQL_NULL_HENV;
    SQLHDBC          hDbc      = SQL_NULL_HDBC;
    SQLHSTMT         hStmt     = SQL_NULL_HSTMT, hStmt2 = SQL_NULL_HSTMT;
    SQLRETURN        rc        = SQL_SUCCESS, rc2 = SQL_SUCCESS;
    SQLINTEGER       RETCODE = 0;
    char             *pDSN     = "STLEC1";

```



```

SWORD          cbCursor;
SDWORD         cbValue1;
SDWORD         cbValue2;
char           employee [30];
int            salary = 0;
char           cursor_name [20];
char           update [200];

char           *stmt = "SELECT NAME, SALARY FROM EMPLOYEE WHERE
                      SALARY > 100000 FOR UPDATE OF SALARY";

(void) printf ("**** Entering CLIP04.\n\n");

/*****
/* Allocate Environment Handle */
*****/

RETCODE = SQLAllocEnv(&hEnv);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Allocate Connection Handle to DSN */
*****/

RETCODE = SQLAllocConnect(hEnv,
                          &hDbc);

if( RETCODE != SQL_SUCCESS )    // Could not get a Connect Handle
    goto dberror;

/*****
/* CONNECT TO data source (STLEC1) */
*****/

RETCODE = SQLConnect(hDbc,    // Connect handle
                    (SQLCHAR *) pDSN, // DSN
                    SQL_NTS,    // DSN is nul-terminated
                    NULL,       // Null UID
                    0,          // Null Auth string
                    NULL,       // Null Auth string
                    0);

if( RETCODE != SQL_SUCCESS )    // Connect failed
    goto dberror;

/*****
/* Allocate Statement Handles */
*****/

rc = SQLAllocStmt (hDbc,
                  &hStmt);

if (rc != SQL_SUCCESS)
    goto exit;

```

## SQLGetCursorName

```
rc = SQLAllocStmt (hDbc,
                  &hStmt2);

if (rc != SQL_SUCCESS)
    goto exit;

/*****
/* Execute query to retrieve employee names */
*****/

rc = SQLExecDirect (hStmt,
                   (SQLCHAR *) stmt,
                   strlen(stmt));

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** EMPLOYEE QUERY FAILED.\n");
    goto dberror;
}

/*****
/* Extract cursor name -- required to build UPDATE statement. */
*****/

rc = SQLGetCursorName (hStmt,
                      (SQLCHAR *) cursor_name,
                      sizeof(cursor_name),
                      &cbCursor);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** GET CURSOR NAME FAILED.\n");
    goto dberror;
}

(void) printf ("**** Cursor Name is %s.\n");

rc = SQLBindCol (hStmt,          // bind employee name
                1,
                SQL_C_CHAR,
                employee,
                sizeof(employee),
                &cbValue1);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** BIND OF NAME FAILED.\n");
    goto dberror;
}

rc = SQLBindCol (hStmt,          // bind employee salary
                2,
                SQL_C_LONG,
                &salary,
                0,
                &cbValue2);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** BIND OF SALARY FAILED.\n");
    goto dberror;
}
```

```

/*****
/* Answer Set is available -- Fetch rows and update salary */
*****/

while ((rc = SQLFetch (hStmt)) == SQL_SUCCESS) &&
      (rc2 == SQL_SUCCESS)
{
    int new_salary = salary*1.1;

    (void) printf ("**** Employee Name %s with salary %d. New salary = %d.\n",
                  employee,
                  salary,
                  new_salary);

    sprintf (update,
             "UPDATE EMPLOYEE SET SALARY = %d WHERE CURRENT OF %s",
             new_salary,
             cursor_name);

    (void) printf ("***** Update statement is : %s\n", update);

    rc2 = SQLExecDirect (hStmt2,
                        (SQLCHAR *) update,
                        SQL_NTS);
}

if (rc2 != SQL_SUCCESS)
{
    (void) printf ("**** EMPLOYEE UPDATE FAILED.\n");
    goto dberror;
}

/*****
/* Reexecute query to validate that salary was updated */
*****/

rc = SQLFreeStmt (hStmt,
                  SQL_CLOSE);

rc = SQLExecDirect (hStmt,
                  (SQLCHAR *) stmt,
                  strlen(stmt));

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** EMPLOYEE QUERY FAILED.\n");
    goto dberror;
}

while ((rc = SQLFetch (hStmt)) == SQL_SUCCESS)
{
    (void) printf ("**** Employee Name %s has salary %d.\n",
                  employee,
                  salary);
}

```

## SQLGetCursorName

```

/*****
/* Deallocate Statement Handles */
*****/

rc = SQLFreeStmt (hStmt,
                  SQL_DROP);

rc = SQLFreeStmt (hStmt2,
                  SQL_DROP);

/*****
/* DISCONNECT from data source */
*****/

RETCODE = SQLDisconnect(hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Deallocate Connection Handle */
*****/

RETCODE = SQLFreeConnect (hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Free Environment Handle */
*****/

RETCODE = SQLFreeEnv (hEnv);

if (RETCODE == SQL_SUCCESS)
    goto exit;

dberror:
RETCODE=12;

exit:

(void) printf ("**** Exiting CLIP04.\n\n");

return RETCODE;
}

```

## References

- “SQLExecute - Execute a statement” on page 159
- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLPrepare - Prepare a statement” on page 277
- “SQLSetCursorName - Set cursor name” on page 322

## SQLGetData - Get data from a column

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLGetData() retrieves data for a single column in the current row of the result set. This is an alternative to SQLBindCol(), which is used to transfer data directly into application variables or LOB locators on each SQLFetch() or SQLExtendedFetch() call. SQLGetData() can also be used to retrieve large data values in pieces.

SQLFetch() must be called before SQLGetData().

After calling SQLGetData() for each column, SQLFetch() or SQLExtendedFetch() is called to retrieve the next row.

### Syntax

```
SQLRETURN SQLGetData(
    (SQLHSTMT          hstmt,
    SQLUSMALLINT      icol,
    SQLSMALLINT       fCType,
    SQLPOINTER        rgbValue,
    SQLINTEGER        cbValueMax,
    SQLINTEGER FAR    *pcbValue);
```

### Function arguments

Table 66 (Page 1 of 2). SQLGetData arguments

Data Type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle.
SQLUSMALLINT	<i>icol</i>	input	Column number for which the data retrieval is requested.
SQLSMALLINT	<i>fCType</i>	input	The C data type of the column identifier by <i>icol</i> . The following types are supported: <ul style="list-style-type: none"> <li>• SQL_C_BINARY</li> <li>• SQL_C_BIT</li> <li>• SQL_C_BLOB_LOCATOR</li> <li>• SQL_C_CHAR</li> <li>• SQL_C_CLOB_LOCATOR</li> <li>• SQL_C_DATE</li> <li>• SQL_C_DBCHAR</li> <li>• SQL_C_DBCLOB_LOCATOR</li> <li>• SQL_C_DOUBLE</li> <li>• SQL_C_FLOAT</li> <li>• SQL_C_LONG</li> <li>• SQL_C_SHORT</li> <li>• SQL_C_TIME</li> <li>• SQL_C_TIMESTAMP</li> <li>• SQL_C_TINYINT</li> </ul> Specifying SQL_C_DEFAULT results in the data being converted to its default C data type, see Table 4 on page 39 for more information.
SQLPOINTER	<i>rgbValue</i>	output	Pointer to buffer where the retrieved column data is to be stored.

Table 66 (Page 2 of 2). SQLGetData arguments

Data Type	Argument	Use	Description
SQLINTEGER	<i>cbValueMax</i>	input	Maximum size of the buffer pointed to by <i>rgbValue</i> .
SQLINTEGER *	<i>pcbValue</i>	output	<p>Pointer to value which indicates the number of bytes DB2 ODBC has available to return in the <i>rgbValue</i> buffer. If the data is being retrieved in pieces, this contains the number of bytes still remaining.</p> <p>The value is SQL_NULL_DATA if the data value of the column is null. If this pointer is NULL and SQLFetch() has obtained a column containing null data, then this function fails because it has no means of reporting this.</p> <p>If SQLFetch() has fetched a column containing binary data, then the pointer to <i>pcbValue</i> must not be NULL or this function fails because it has no other means of informing the application about the length of the data retrieved in the <i>rgbValue</i> buffer.</p>

**Note:** DB2 ODBC provides some performance enhancement if *rgbValue* is placed consecutively in memory after *pcbValue*.

## Usage

SQLGetData() can be used with SQLBindCol() for the same result set, as long as SQLFetch() and not SQLExtendedFetch() is used. The general steps are:

1. SQLFetch() - advances cursor to first row, retrieves first row, transfers data for bound columns.
2. SQLGetData() - transfers data for the specified column.
3. Repeat step 2 for each column needed.
4. SQLFetch() - advances cursor to next row, retrieves next row, transfers data for bound columns.
5. Repeat steps 2, 3 and 4 for each row in the result set, or until the result set is no longer needed.

SQLGetData() can also be used to retrieve long columns if the C data type (*fCType*) is SQL\_C\_CHAR, SQL\_C\_BINARY, SQL\_C\_DBCHAR, or if *fCType* is SQL\_C\_DEFAULT and the column type denotes a binary or character string.

Upon each SQLGetData() call, if the data available for return is greater than or equal to *cbValueMax*, truncation occurs. Truncation is indicated by a function return code of SQL\_SUCCESS\_WITH\_INFO coupled with a SQLSTATE denoting data truncation. The application can call SQLGetData() again, with the same *icol* value, to get subsequent data from the same unbound column starting at the point of truncation. To obtain the entire column, the application repeats such calls until the function returns SQL\_SUCCESS. The next call to SQLGetData() returns SQL\_NO\_DATA\_FOUND.

Although SQLGetData() can be used for the sequential retrieval of LOB column data, use the DB2 ODBC LOB functions if only a portion of the LOB data or a few sections of the LOB column data are needed:

1. Bind the column to a LOB locator.
2. Fetch the row.

3. Use the locator in a `SQLGetSubString()` call to retrieve the data in pieces. (`SQLGetLength()` and `SQLGetPosition()` might also be required for determining the values of some of the arguments).
4. Repeat step 2.

Truncation is also affected by the `SQL_MAX_LENGTH` statement option. The application can specify that truncation is not to be reported by calling `SQLSetStmtOption()` with `SQL_MAX_LENGTH` and a value for the maximum length to return for any one column, and by allocating a *rgbValue* buffer of the same size (plus the null-terminator). If the column data is larger than the set maximum length, `SQL_SUCCESS` is returned and the maximum length, not the actual length is returned in *pcbValue*.

To discard the column data part way through the retrieval, the application can call `SQLGetData()` with *icol* set to the next column position of interest. To discard data that has not been retrieved for the entire row, the application should call `SQLFetch()` to advance the cursor to the next row; or, if it is not interested in any more data from the result set, call `SQLFreeStmt()` to close the cursor.

The *fCType* input argument determines the type of data conversion (if any) needed before the column data is placed into the storage area pointed to by *rgbValue*.

For SQL graphic column data:

- The length of the *rgbValue* buffer (*cbValueMax*) should be a multiple of 2. The application can determine the SQL data type of the column by first calling `SQLDescribeCol()` or `SQLColAttributes()`.
- The pointer to *pcbValue* must not be NULL since DB2 ODBC stores the number of octets stored in *rgbValue*.
- If the data is retrieved in piecewise fashion, DB2 ODBC attempts to fill *rgbValue* to the nearest multiple of two octets that is still less than or equal to *cbValueMax*. This means if *cbValueMax* is not a multiple of two, the last byte in that buffer is untouched; DB2 ODBC does not split a double-byte character.

The contents returned in *rgbValue* are always null-terminated unless the column data to be retrieved is binary, or if the SQL data type of the column is graphic (DBCS) and the C buffer type is `SQL_C_CHAR`. If the application is retrieving the data in multiple chunks, it should make the proper adjustments (for example, strip off the null-terminator before concatenating the pieces back together assuming the null termination environment attribute is in effect).

Truncation of numeric data types is reported as a warning if the truncation involves digits to the right of the decimal point. If truncation occurs to the left of the decimal point, an error is returned (see the 'Diagnostics' section).

Applications that use `SQLExtendedFetch()` to retrieve data should call `SQLGetData()` only when the rowset size is 1 (equivalent to issuing `SQLFetch()`). `SQLGetData()` can only retrieve column data for a row where the cursor is currently positioned.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE
- SQL\_NO\_DATA\_FOUND

SQL\_NO\_DATA\_FOUND is returned when the preceding SQLGetData() call has retrieved all of the data for this column.

SQL\_SUCCESS is returned if a zero-length string is retrieved by SQLGetData(). If this is the case, *pcbValue* contains 0, and *rgbValue* contains a null terminator.

If the preceding call to SQLFetch() failed, SQLGetData() should not be called since the result is undefined.

## Diagnostics

Table 67 (Page 1 of 2). SQLGetData SQLSTATES

SQLSTATE	Description	Explanation
01004	Data truncated.	Data returned for the specified column ( <i>icol</i> ) is truncated. String or numeric values are right truncated. SQL_SUCCESS_WITH_INFO is returned.
07006	Invalid conversion.	The data value cannot be converted to the C data type specified by the argument <i>fCType</i> . The function has been called before for the same <i>icol</i> value but with a different <i>fCType</i> value.
22002	Invalid output or indicator buffer specified.	The pointer value specified for the argument <i>pcbValue</i> is a null pointer and the value of the column is null. There is no means to report SQL_NULL_DATA.
22003	Numeric value out of range.	Returning the numeric value (as numeric or string) for the column causes the whole part of the number to be truncated.
22005	Error in assignment.	A returned value is incompatible with the data type denoted by the argument <i>fCType</i> .
22008	Datetime field overflow.	Datetime field overflow occurred; for example, an arithmetic operation on a date or timestamp has a result that is not within the valid range of dates, or a datetime value cannot be assigned to a bound variable because it is too small.
24000	Invalid cursor state.	The previous SQLFetch() resulted in SQL_ERROR or SQL_NO_DATA found; as a result, the cursor is not positioned on a row.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.



Table 67 (Page 2 of 2). SQLGetData SQLSTATES

SQLSTATE	Description	Explanation
S1002	Invalid column number.	The specified column is less than 0 or greater than the number of result columns.  The specified column is 0, but DB2 ODBC does not support ODBC bookmarks ( <i>icol</i> = 0).  SQLExtendedFetch() is called for this result set.
S1003	Program type out of range.	<i>fcType</i> is not a valid data type or SQL_C_DEFAULT.
S1009	Invalid argument value.	The argument <i>rgbValue</i> is a null pointer.  The argument <i>pcbValue</i> is a null pointer; the column SQL data type is graphic (DBCS); and <i>fcType</i> is set to SQL_C_CHAR.
S1010	Function sequence error.	The specified <i>hstmt</i> is not in a cursor positioned state. The function is called without first calling SQLFetch().  The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1090	Invalid string or buffer length.	The value of the argument <i>cbValueMax</i> is less than 0 and the argument <i>fcType</i> is SQL_C_CHAR, SQL_C_BINARY, SQL_C_DBCHAR or (SQL_C_DEFAULT and the default type is one of SQL_C_CHAR, SQL_C_BINARY, or SQL_C_DBCHAR).
S1C00	Driver not capable.	The SQL data type for the specified data type is recognized but not supported by DB2 ODBC.  The requested conversion from the SQL data type to the application data <i>fcType</i> cannot be performed by DB2 ODBC or the data source.  SQLExtendedFetch() is called for the specified <i>hstmt</i> .

## Restrictions

ODBC has defined column 0 for bookmarks. DB2 ODBC does not support bookmarks.

## Example

See “Example” on page 173 for a comparison between using bound columns and using SQLGetData().

```

/*****
/* DB2 for OS/390 example:
/*     Populates BIOGRAPHY Table from flat file text. Inserts
/*     VITAE in 80-byte pieces via SQLPutData. Also retrieve
/*     NAME, UNIT and VITAE for all members. VITAE is retrieved
/*     via SQLGetData.
*****/

#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <sqlca.h>
#include "sqlcli1.h"

#define TEXT_SIZE 80

int insert_bio (SQLHSTMT hStmt,      // insert_bio prototype
               char      *bio,
               int       bcount);

int main( )
{
    SQLHENV      hEnv   = SQL_NULL_HENV;
    SQLHDBC      hDbc   = SQL_NULL_HDBC;
    SQLHSTMT     hStmt  = SQL_NULL_HSTMT, hStmt2 = SQL_NULL_HSTMT;
    SQLRETURN    rc     = SQL_SUCCESS;
    FILE         *fp;
    SQLINTEGER   RETCODE = 0;
    char         pTable [200];
    char         *pDSN = "STLEC1";
    UDWORD      pirow;
    SDWORD      cbValue;

```

```

char          *i_stmt = "INSERT INTO BIOGRAPHY VALUES (?, ?, ?)";
char          *query  = "SELECT NAME, UNIT, VITAE FROM BIOGRAPHY";
char          text [TEXT_SIZE]; // file text
char          vitae [3200];    // biography text
char          Narrative [TEXT_SIZE];
SQLINTEGER    vitae_ind = SQL_DATA_AT_EXEC; // bio data is
                                     // passed at execute time
                                     // via SQLPutData

SQLINTEGER    vitae_cbValue = TEXT_SIZE;
char          *t = NULL;
char          *c = NULL;
char          name [21];
SQLINTEGER    name_ind = SQL_NTS;
SQLINTEGER    name_cbValue = sizeof(name);
char          unit [31];
SQLINTEGER    unit_ind = SQL_NTS;
SQLINTEGER    unit_cbValue = sizeof(unit);
char          tmp [80];
char          *token = NULL, *pbio = vitae;
char          insert = SQL_FALSE;
int           i, bcount = 0;

(void) printf ("**** Entering CLIP09.\n\n");

/*****
/* Allocate Environment Handle
*****/

RETCODE = SQLAllocEnv(&hEnv);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Allocate Connection Handle to DSN
*****/

RETCODE = SQLAllocConnect(hEnv,
                          &hDbc);

if( RETCODE != SQL_SUCCESS )    // Could not get a Connect Handle
    goto dberror;

/*****
/* CONNECT TO data source (STLEC1)
*****/

RETCODE = SQLConnect(hDbc,          // Connect handle
                    (SQLCHAR *) pDSN, // DSN
                    SQL_NTS,        // DSN is nul-terminated
                    NULL,           // Null UID
                    0,               //
                    NULL,            // Null Auth string
                    0);

if( RETCODE != SQL_SUCCESS )    // Connect failed
    goto dberror;

```

```

/*****
/* Allocate Statement Handles */
/*****

rc = SQLAllocStmt (hDbc,
                  &hStmt);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** Allocate Statement Handle Failed.\n");
    goto dberror;
}

rc = SQLAllocStmt (hDbc,
                  &hStmt2);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** Allocate Statement Handle Failed.\n");
    goto dberror;
}

/*****
/* Prepare INSERT statement. */
/*****

rc = SQLPrepare (hStmt,
                (SQLCHAR *) i_stmt,
                SQL_NTS);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** Prepare of INSERT Failed.\n");
    goto dberror;
}

/*****
/* Bind NAME and UNIT. Bind VITAE so that data can be passed */
/* via SQLPutData. */
/*****

rc = SQLBindParameter (hStmt,          // bind NAME
                      1,
                      SQL_PARAM_INPUT,
                      SQL_C_CHAR,
                      SQL_CHAR,
                      sizeof(name),
                      0,
                      name,
                      sizeof(name),
                      &name_ind);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** Bind of NAME Failed.\n");
    goto dberror;
}

```

```

rc = SQLBindParameter (hStmt,          // bind Branch
                      2,
                      SQL_PARAM_INPUT,
                      SQL_C_CHAR,
                      SQL_CHAR,
                      sizeof(unit),
                      0,
                      unit,
                      sizeof(unit),
                      &unit_ind);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** Bind of UNIT Failed.\n");
    goto dberror;
}

rc = SQLBindParameter (hStmt,          // bind Rank
                      3,
                      SQL_PARAM_INPUT,
                      SQL_C_CHAR,
                      SQL_LONGVARCHAR,
                      3200,
                      0,
                      (SQLPOINTER) 3,
                      0,
                      &vitae_ind);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** Bind of VITAE Failed.\n");
    goto dberror;
}

/*****
/* Read Biographical text from flat file
*****/

if ((fp = fopen ("DD: BIOGRAF", "r")) == NULL) // open command file
{
    rc = SQL_ERROR;          // open failed
    goto exit;
}

/*****
/* Process file and insert Biographical text
*****/

while (((t = fgets (text, sizeof(text), fp)) != NULL) &&
      (rc == SQL_SUCCESS))
{
    if (text[0] == '#')      // if commander data
    {
        if (insert)          // if BIO data to be inserted
        {
            rc = insert_bio (hStmt,
                             vitae,
                             bcount);      // insert row into BIOGRAPHY Table
            bcount = 0;          // reset text line count
            pbio = vitae;       // reset text pointer
        }
    }
}

```

```

        token = strtok (text+1, ","); // get member NAME
        (void) strcpy (name, token);
        token = strtok (NULL, "#"); // extract UNIT
        (void) strcpy (unit, token); // copy to local variable
                                        // SQLPutData
        insert = SQL_TRUE; // have row to insert
    }
    else
    {
        memset (pbio, ' ', sizeof(text));
        strcpy (pbio, text); // populate text
        i = strlen (pbio); // remove '\n' and '\0'
        pbio [i--] = ' ';
        pbio [i] = ' ';
        pbio += sizeof (text); // advance pbio
        bcount++; // one more text line
    }
}

if (insert) // if BIO data to be inserted
{
    rc = insert_bio (hStmt,
                    vitae,
                    bcount); // insert row into BIOGRAPHY Table
}

fclose (fp); // close text flat file

/*****
/* Commit Insert of rows */
*****/

rc = SQLTransact (hEnv,
                  hDbc,
                  SQL_COMMIT);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** COMMIT FAILED.\n");
    goto dberror;
}

/*****
/* Open query to retrieve NAME, UNIT and VITAE. Bind NAME and */
/* UNIT but leave VITAE unbound. Retrieved using SQLGetData. */
*****/

RETCODE = SQLPrepare (hStmt2,
                      (SQLCHAR *)query,
                      strlen(query));

if (RETCODE != SQL_SUCCESS)
{
    (void) printf ("**** Prepare of Query Failed.\n");
    goto dberror;
}

```

```

RETCODE = SQLExecute (hStmt2);

if (RETCODE != SQL_SUCCESS)
{
    (void) printf ("**** Query Failed.\n");
    goto dberror;
}

RETCODE = SQLBindCol (hStmt2,          // bind NAME
                      1,
                      SQL_C_DEFAULT,
                      name,
                      sizeof(name),
                      &name_cbValue);

if (RETCODE != SQL_SUCCESS)
{
    (void) printf ("**** Bind of NAME Failed.\n");
    goto dberror;
}

RETCODE = SQLBindCol (hStmt2,          // bind UNIT
                      2,
                      SQL_C_DEFAULT,
                      unit,
                      sizeof(unit),
                      &unit_cbValue);

if (RETCODE != SQL_SUCCESS)
{
    (void) printf ("**** Bind of UNIT Failed.\n");
    goto dberror;
}

while ((RETCODE = SQLFetch (hStmt2)) != SQL_NO_DATA_FOUND)
{
    (void) printf ("**** Name is %s. Unit is %s.\n\n", name, unit);
    (void) printf ("**** Vitae follows:\n\n");

    for (i = 0; (i < 3200 && RETCODE != SQL_NO_DATA_FOUND); i += TEXT_SIZE)
    {
        RETCODE = SQLGetData (hStmt2,
                              3,
                              SQL_C_CHAR,
                              Narrative,
                              sizeof(Narrative) + 1,
                              &vitae_cbValue);

        if (RETCODE != SQL_NO_DATA_FOUND)
            (void) printf ("%s\n", Narrative);
    }
}

/*****
/* Deallocate Statement Handles */
*****/

rc = SQLFreeStmt (hStmt,
                  SQL_DROP);

rc = SQLFreeStmt (hStmt2,
                  SQL_DROP);

```

```

/*****
/* DISCONNECT from data source */
*****/

RETCODE = SQLDisconnect(hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Deallocate Connection Handle */
*****/

RETCODE = SQLFreeConnect (hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Free Environment Handle */
*****/

RETCODE = SQLFreeEnv (hEnv);

if (RETCODE == SQL_SUCCESS)
    goto exit;

dberror:
RETCODE=12;

exit:

(void) printf ("**** Exiting CLIP09.\n\n");

return RETCODE;
}

/*****
/* function insert_bio is invoked to insert one row into the */
/* BIOGRAPHY Table. The biography text is inserted in sets of */
/* 80 bytes via SQLPutData. */
*****/

int insert_bio (SQLHSTMT hStmt,
               char      *vitae,
               int       bcount)
{
    SQLINTEGER      rc = SQL_SUCCESS;
    SQLPOINTER      prgbValue;
    int             i;
    char            *text;

```



```

/*****
/* NAME and UNIT are bound... VITAE is provided after execution */
/* of the INSERT using SQLPutData. */
*****/

rc = SQLExecute (hStmt);

if (rc != SQL_NEED_DATA) // expect SQL_NEED_DATA
{
    rc = 12;
    (void) printf ("**** NEED DATA not returned.\n");
    goto exit;
}
/*****
/* Invoke SQLParamData to position ODBC driver on input parameter*/
*****/

if ((rc = SQLParamData (hStmt,
                        &prgbValue)) != SQL_NEED_DATA)
{
    rc = 12;
    (void) printf ("**** NEED DATA not returned.\n");
    goto exit;
}
/*****
/* Iterate through VITAE in 80 byte increments... pass to */
/* ODBC Driver via SQLPutData. */
*****/

for (i = 0, text = vitae, rc = SQL_SUCCESS;
     (i < bcount) && (rc == SQL_SUCCESS);
     i++, text += TEXT_SIZE)
{
    rc = SQLPutData (hStmt,
                    text,
                    TEXT_SIZE);
}
/*****
/* Invoke SQLParamData to trigger ODBC driver to execute the */
/* statement. */
*****/

if ((rc = SQLParamData (hStmt,
                        &prgbValue)) != SQL_SUCCESS)
{
    rc = 12;
    (void) printf ("**** INSERT Failed.\n");
}
exit:
return (rc);
}

```

## References

- “SQLExtendedFetch - Extended fetch (fetch array of rows)” on page 162
- “SQLFetch - Fetch next row” on page 169

## SQLGetEnvAttr - Returns current setting of an environment attribute

### Purpose

Specification:		X/OPEN CLI	ISO CLI
----------------	--	------------	---------

SQLGetEnvAttr() returns the current setting for the specified environment attribute.

These options are set using the SQLSetEnvAttr() function.

### Syntax

```
SQLRETURN SQLGetEnvAttr (SQLHENV henv,
                          SQLINTEGER Attribute,
                          SQLPOINTER Value,
                          SQLINTEGER BufferLength,
                          SQLINTEGER FAR *StringLength);
```

### Function arguments

Table 68. SQLGetEnvAttr arguments

Data type	Argument	Use	Description
SQLHENV	henv	input	Environment handle.
SQLINTEGER	Attribute	input	Attribute to get. See Table 125 on page 325 for the list of environment attributes and their descriptions.
SQLPOINTER	Value	output	The current value associated with <i>Attribute</i> . The type of the value returned depends on <i>Attribute</i> .
SQLINTEGER	BufferLength	input	Maximum size of buffer pointed to by <i>Value</i> , if the attribute value is a character string; otherwise, ignored.
SQLINTEGER *	StringLength	output	Length in bytes of the output data if the attribute value is a character string; otherwise, ignored.

If *Attribute* does not denote a string, then DB2 ODBC ignores *BufferLength* and does not set *StringLength*.

### Usage

SQLGetEnvAttr() can be called at any time between the allocation and freeing of the environment handle. It obtains the current value of the environment attribute.

For a list of valid environment attributes, see Table 125 on page 325.

### Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 69. SQLGetEnvAttr SQLSTATEs

SQLSTATE	Description	Explanation
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1092	Option type out of range.	An invalid <i>Attribute</i> value was specified.

## Restrictions

None.

## Example

```

/* ... */
rc = SQLGetEnvAttr(henv, SQL_ATTR_OUTPUT_NTS, &output_nts, 0, 0);
printf("Null Termination of Output strings is: ");
if (output_nts == SQL_TRUE)
    printf("True\n");
else
    printf("False\n");
/* ... */

```

## References

- “SQLSetEnvAttr - Set environment attribute” on page 325

## SQLGetFunctions - Get functions

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLGetFunctions() to query whether a specific function is supported. This allows applications to adapt to varying levels of support when connecting to different database servers.

A connection to a database server must exist before calling this function.

### Syntax

```
SQLRETURN SQLGetFunctions (SQLHDBC          hdbc,
                            SQLUSMALLINT    fFunction,
                            SQLUSMALLINT FAR *pfExists);
```

### Function arguments

Table 70. SQLGetFunctions arguments

Data type	Argument	Use	Description
SQLHDBC	<i>hdbc</i>	input	Database connection handle.
SQLUSMALLINT	<i>fFunction</i>	input	The function being queried. Valid <i>fFunction</i> values are shown in Figure 6 on page 215
SQLUSMALLINT *	<i>pfExists</i>	output	Pointer to location where this function returns SQL_TRUE or SQL_FALSE depending on whether the function being queried is supported.

### Usage

Figure 6 on page 215 shows the valid values for the *fFunction* argument and whether the corresponding function is supported.

If *fFunction* is set to SQL\_API\_ALL\_FUNCTIONS, then *pfExists* must point to an SQLSMALLINT array of 100 elements. The array is indexed by the *fFunction* values used to identify many of the functions. Some elements of the array are unused and reserved. Since some *fFunction* values are greater than 100, the array method can not be used to obtain a list of functions. The SQLGetFunction() call must be explicitly issued for all *fFunction* values equal to or above 100. The complete set of *fFunction* values is defined in sqlcli1.h.

### Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

SQL_API_SQLALLOCONNECT	= TRUE	SQL_API_SQLALLOCENV	= TRUE
SQL_API_SQLALLOCSMT	= TRUE	SQL_API_SQLBINDCOL	= TRUE
SQL_API_SQLBINDFILETOCOL	= FALSE	SQL_API_SQLBINDFILETOPARAM	= FALSE
SQL_API_SQLBINDPARAMETER	= TRUE	SQL_API_SQLBROWSECONNECT	= FALSE
SQL_API_SQLCANCEL	= TRUE	SQL_API_SQLCOLATTRIBUTES	= TRUE
SQL_API_SQLCOLUMNPRIVILEGES	= TRUE	SQL_API_SQLCOLUMNS	= TRUE
SQL_API_SQLCONNECT	= TRUE	SQL_API_SQLDATASOURCES	= TRUE
SQL_API_SQLDESCRIBECOL	= TRUE	SQL_API_SQLDESCRIBEPARAM	= TRUE
SQL_API_SQLDISCONNECT	= TRUE	SQL_API_SQLDRIVERCONNECT	= TRUE
SQL_API_SQLERROR	= TRUE	SQL_API_SQLEXECDIRECT	= TRUE
SQL_API_SQLEXECUTE	= TRUE	SQL_API_SQLEXTENDEDFETCH	= TRUE
SQL_API_SQLFETCH	= TRUE	SQL_API_SQLFOREIGNKEYS	= TRUE
SQL_API_SQLFREECONNECT	= TRUE	SQL_API_SQLFREEENV	= TRUE
SQL_API_SQLFREESTMT	= TRUE	SQL_API_SQLGETCONNECTOPTION	= TRUE
SQL_API_SQLGETCURSORNAME	= TRUE	SQL_API_SQLGETDATA	= TRUE
SQL_API_SQLGETENVATTR	= TRUE	SQL_API_SQLGETFUNCTIONS	= TRUE
SQL_API_SQLGETINFO	= TRUE	SQL_API_SQLGETLENGTH	= TRUE
SQL_API_SQLGETPOSITION	= TRUE	SQL_API_SQSQLGETSQLCA	= TRUE
SQL_API_SQLGETSTMTOPTION	= TRUE	SQL_API_SQLGETSUBSTRING	= TRUE
SQL_API_SQLGETTYPEINFO	= TRUE	SQL_API_SQLMORERESULTS	= TRUE
SQL_API_SQLNATIVESQL	= TRUE	SQL_API_SQLNUMPARAMS	= TRUE
SQL_API_SQLNUMRESULTCOLS	= TRUE	SQL_API_SQLPARAMDATA	= TRUE
SQL_API_SQLPARAMOPTIONS	= TRUE	SQL_API_SQLPREPARE	= TRUE
SQL_API_SQLPRIMARYKEYS	= TRUE	SQL_API_SQLPROCEDURECOLUMNS	= TRUE
SQL_API_SQLPROCEDURES	= TRUE	SQL_API_SQLPUTDATA	= TRUE
SQL_API_SQLROWCOUNT	= TRUE	SQL_API_SQLSETCOLATTRIBUTES	= TRUE
SQL_API_SQLSETCONNECTION	= TRUE	SQL_API_SQLSETCONNECTOPTION	= TRUE
SQL_API_SQLSETCURSORNAME	= TRUE	SQL_API_SQLSETENVATTR	= TRUE
SQL_API_SQLSETPARAM	= TRUE	SQL_API_SQLSETPOS	= FALSE
SQL_API_SQLSETSCROLLOPTIONS	= FALSE	SQL_API_SQLSETSTMTOPTION	= TRUE
SQL_API_SQLSPECIALCOLUMNS	= TRUE	SQL_API_SQLSTATISTICS	= TRUE
SQL_API_SQLTABLEPRIVILEGES	= TRUE	SQL_API_SQLTABLES	= TRUE
SQL_API_TRANSACT	= TRUE		

Figure 6. Supported functions list

## Diagnostics

Table 71. SQLGetFunctions SQLSTATEs

SQLSTATE	Description	Explanation
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	The argument pfExists was a null pointer.
S1010	Function sequence error.	SQLGetFunctions() was called before a database connection was established.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.

## Restrictions

None.

## Example

```

/*****
/* DB2 for OS/390 example: */
/* Executes SQLGetFunctions to verify that APIs required */
/* by application are supported. */
*****/

#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <sqlca.h>
#include "sqlcli1.h"

typedef struct odbc_api
{
    SQLUSMALLINT api;
    char api_name [40];
} ODBC_API;

/*****
/* CLI APIs required by application */
*****/

ODBC_API o_api [7] = {
    { SQL_API_SQLBINDPARAMETER, "SQLBindParameter" },
    { SQL_API_SQLDISCONNECT , "SQLDisconnect" },
    { SQL_API_SQLGETTYPEINFO , "SQLGetTypeInfo" },
    { SQL_API_SQLFETCH , "SQLFetch" },
    { SQL_API_SQLTRANSACT , "SQLTransact" },
    { SQL_API_SQLBINDCOL , "SQLBindCol" },
    { SQL_API_SQLEXCEDIRECT , "SQLExecDirect" }
    };

/*****
/* Validate that required APIs are supported. */
*****/

int main( )
{
    SQLHENV hEnv = SQL_NULL_HENV;
    SQLHDBC hDbc = SQL_NULL_HDBC;
    SQLRETURN rc = SQL_SUCCESS;
    SQLINTEGER RETCODE = 0;
    int i;

    // SQLGetFunctions parameters

    SQLUSMALLINT fExists = SQL_TRUE;
    SQLUSMALLINT *pfExists = &fExists;

    (void) printf ("**** Entering CLIP05.\n\n");

```

```

/*****
/* Allocate Environment Handle */
*****/

RETCODE = SQLAllocEnv(&hEnv);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Allocate Connection Handle to DSN */
*****/

RETCODE = SQLAllocConnect(hEnv,
                          &hDbc);

if( RETCODE != SQL_SUCCESS )    // Could not get a Connect Handle
    goto dberror;

/*****
/* CONNECT TO data source (STLEC1) */
*****/

RETCODE = SQLConnect(hDbc,      // Connect handle
                    (SQLCHAR *) "STLEC1", // DSN
                    SQL_NTS,    // DSN is nul-terminated
                    NULL,       // Null UID
                    0,          //
                    NULL,       // Null Auth string
                    0);

if( RETCODE != SQL_SUCCESS )    // Connect failed
    goto dberror;

/*****
/* See if DSN supports required ODBC APIs */
*****/

for (i = 0, (*pfExists = SQL_TRUE);
     (i < (sizeof(o_api)/sizeof(ODBC_API)) && (*pfExists) == SQL_TRUE);
     i++)
{
    RETCODE = SQLGetFunctions (hDbc,
                              o_api[i].api,
                              pfExists);

    if (*pfExists == SQL_TRUE)    // if api is supported then print
    {
        (void) printf ("**** ODBC api %s IS supported.\n",
                       o_api[i].api_name);
    }
}

if (*pfExists == SQL_FALSE)    // a required api is not supported
{
    (void) printf ("**** ODBC api %s not supported.\n",
                  o_api[i].api_name);
}

```

## SQLGetFunctions

```

/*****
/* DISCONNECT from data source
*****/

RETCODE = SQLDisconnect(hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Deallocate Connection Handle
*****/

RETCODE = SQLFreeConnect (hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Free Environment Handle
*****/

RETCODE = SQLFreeEnv (hEnv);

if (RETCODE == SQL_SUCCESS)
    goto exit;

dberror:
RETCODE=12;

exit:

(void) printf("\n\n**** Exiting CLIP05.\n\n ");

return(RETCODE);
}

```

## References

None.



## SQLGetInfo - Get general information

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLGetInfo() returns general information, (including supported data conversions) about the DBMS that the application is currently connected to.

### Syntax

```
SQLRETURN SQLGetInfo (SQLHDBC hdbc,
                      SQLUSMALLINT fInfoType,
                      SQLPOINTER rgbInfoValue,
                      SQLSMALLINT cbInfoValueMax,
                      SQLSMALLINT FAR *pcbInfoValue);
```

### Function arguments

Table 72. SQLGetInfo arguments

Data type	Argument	Use	Description
SQLHDBC	<i>hdbc</i>	input	Database connection handle
SQLUSMALLINT	<i>fInfoType</i>	input	The type of information desired.
SQLPOINTER	<i>rgbInfoValue</i>	output (also input)	Pointer to buffer where this function stores the desired information. Depending on the type of information being retrieved, 5 types of information can be returned: <ul style="list-style-type: none"> <li>• 16-bit integer value</li> <li>• 32-bit integer value</li> <li>• 32-bit binary value</li> <li>• 32-bit mask</li> <li>• Null-terminated character string</li> </ul>
SQLSMALLINT	<i>cbInfoValueMax</i>	input	Maximum length of the buffer pointed to by <i>rgbInfoValue</i> pointer.
SQLSMALLINT *	<i>pcbInfoValue</i>	output	Pointer to location where this function returns the total number of bytes available to return the desired information. In the case of string output, this size does not include the null terminating character. <p>If the value in the location pointed to by <i>pcbInfoValue</i> is greater than the size of the <i>rgbInfoValue</i> buffer as specified in <i>cbInfoValueMax</i>, then the string output information is truncated to <i>cbInfoValueMax</i> - 1 bytes and the function returns with SQL_SUCCESS_WITH_INFO.</p>

### Usage

See Table 73 on page 220 for a list of the possible values of *fInfoType* and a description of the information that SQLGetInfo() would return for that value.

## SQLGetInfo

Table 73 (Page 1 of 14). Information returned by SQLGetInfo

<i>InfoType</i>	Format	Description and notes
<p><b>Note:</b> DB2 ODBC returns a value for each <i>InfoType</i> in this table. If the <i>InfoType</i> does not apply or is not supported, the result is dependent on the return type. If the return type is a:</p> <ul style="list-style-type: none"> <li>• Character string containing 'Y' or 'N', "N" is returned.</li> <li>• Character string containing a value other than just 'Y' or 'N', an empty string is returned.</li> <li>• 16-bit integer, 0 (zero).</li> <li>• 32-bit integer, 0 (zero).</li> <li>• 32-bit mask, 0 (zero).</li> </ul>		
SQL_ACCESSIBLE_PROCEDURES [20]	string	A character string of "Y" indicates that the user can execute all procedures returned by the function <code>SQLProcedures()</code> . "N" indicates that procedures can be returned that the user cannot execute.
SQL_ACCESSIBLE_TABLES [19]	string	A character string of "Y" indicates that the user is guaranteed SELECT privilege to all tables returned by the function <code>SQLTables()</code> . "N" indicates that tables can be returned that the user cannot access.
SQL_ACTIVE_CONNECTIONS [0]	16-bit integer	The maximum number of active connections supported per application.  Zero is returned, indicating that the limit is dependent on system resources.  The MAXCONN keyword in the initialization file or the SQL_MAX_CONNECTIONS environment/connection option can be used to impose a limit on the number of connections. This limit is returned if it is set to any value other than zero.
SQL_ACTIVE_STATEMENTS [1]	16-bit integer	The maximum number of active statements per connection.  Zero is returned, indicating that the limit is dependent on database system and DB2 ODBC resources, and limits.
SQL_ALTER_TABLE [86]	32-bit mask	Indicates which clauses in ALTER TABLE are supported by the DBMS. <ul style="list-style-type: none"> <li>• SQL_AT_ADD_COLUMN</li> <li>• SQL_AT_DROP_COLUMN</li> </ul>
SQL_BOOKMARK_PERSISTENCE [82]	32-bit mask	Reserved option, zero is returned for the bit-mask.
SQL_COLUMN_ALIAS [87]	string	Returns "Y" if column aliases are supported, or "N" if they are not.
SQL_CONCAT_NULL_BEHAVIOR [22]	16-bit integer	Indicates how the concatenation of NULL valued character data type columns with non-NULL valued character data type columns is handled. <ul style="list-style-type: none"> <li>• SQL_CB_NULL - indicates the result is a NULL value (this is the case for IBM RDBMs).</li> <li>• SQL_CB_NON_NULL - indicates the result is a concatenation of non-NULL column values.</li> </ul>

Table 73 (Page 2 of 14). Information returned by SQLGetInfo

<i>InfoType</i>	Format	Description and notes
SQL_CONVERT_BIGINT [53] SQL_CONVERT_BINARY [54] SQL_CONVERT_BIT [55] SQL_CONVERT_CHAR [56] SQL_CONVERT_DATE [57] SQL_CONVERT_DECIMAL [58] SQL_CONVERT_DOUBLE [59] SQL_CONVERT_FLOAT [60] SQL_CONVERT_INTEGER [61] SQL_CONVERT_LONGVARBINARY [71] SQL_CONVERT_LONGVARCHAR [62] SQL_CONVERT_NUMERIC [63] SQL_CONVERT_REAL [64] SQL_CONVERT_ROWID [121] SQL_CONVERT_SMALLINT [65] SQL_CONVERT_TIME [66] SQL_CONVERT_TIMESTAMP [67] SQL_CONVERT_TINYINT [68] SQL_CONVERT_VARBINARY [69] SQL_CONVERT_VARCHAR [70]	32-bit mask	<p>Indicates the conversions supported by the data source with the CONVERT scalar function for data of the type named in the <i>InfoType</i>. If the bitmask equals zero, the data source does not support any conversions for the data of the named type, including conversions to the same data type.</p> <p>For example, to find out if a data source supports the conversion of SQL_INTEGER data to the SQL_DECIMAL data type, an application calls SQLGetInfo() with <i>InfoType</i> of SQL_CONVERT_INTEGER. The application then ANDs the returned bitmask with SQL_CVT_DECIMAL. If the resulting value is nonzero then the conversion is supported.</p> <p>The following bitmasks are used to determine which conversions are supported:</p> <ul style="list-style-type: none"> <li>• SQL_CVT_BIGINT</li> <li>• SQL_CVT_BINARY</li> <li>• SQL_CVT_BIT</li> <li>• SQL_CVT_CHAR</li> <li>• SQL_CVT_DATE</li> <li>• SQL_CVT_DECIMAL</li> <li>• SQL_CVT_DOUBLE</li> <li>• SQL_CVT_FLOAT</li> <li>• SQL_CVT_INTEGER</li> <li>• SQL_CVT_LONGVARBINARY</li> <li>• SQL_CVT_LONGVARCHAR</li> <li>• SQL_CVT_NUMERIC</li> <li>• SQL_CVT_REAL</li> <li>• SQL_CVT_ROWID</li> <li>• SQL_CVT_SMALLINT</li> <li>• SQL_CVT_TIME</li> <li>• SQL_CVT_TIMESTAMP</li> <li>• SQL_CVT_TINYINT</li> <li>• SQL_CVT_VARBINARY</li> <li>• SQL_CVT_VARCHAR</li> </ul>
SQL_CONVERT_FUNCTIONS [48]	32-bit mask	Indicates the scalar conversion functions supported by the driver and associated data source.
SQL_CORRELATION_NAME [74]	16-bit integer	<p>Indicates the degree of correlation name support by the server:</p> <ul style="list-style-type: none"> <li>• SQL_CN_ANY, supported and can be any valid user-defined name.</li> <li>• SQL_CN_NONE, correlation name not supported.</li> <li>• SQL_CN_DIFFERENT, correlation name supported but it must be different than the name of the table that it represents.</li> </ul>
SQL_CLOSE_BEHAVIOR	32-bit integer	<p>Indicates whether or not locks are released when the cursor is closed. The possible values are:</p> <ul style="list-style-type: none"> <li>• SQL_CC_NO_RELEASE: locks are not released when the cursor on this statement handle is closed. This is the default.</li> <li>• SQL_CC_RELEASE: locks are released when the cursor on this statement handle is closed.</li> </ul> <p>Typically cursors are explicitly closed when the function SQLFreeStmt() is called with the SQL_CLOSE or SQL_DROP option. In addition, the end of the transaction (when a commit or rollback is issued) can also cause the closing of the cursor (depending on the WITH HOLD attribute currently in use).</p>

## SQLGetInfo

Table 73 (Page 3 of 14). Information returned by SQLGetInfo

<i>InfoType</i>	<b>Format</b>	<b>Description and notes</b>
SQL_CURSOR_COMMIT_BEHAVIOR	16-bit integer	<p>Indicates how a COMMIT operation affects cursors. A value of:</p> <ul style="list-style-type: none"> <li>SQL_CB_DELETE, destroys cursors and drops access plans for dynamic SQL statements.</li> <li>SQL_CB_CLOSE, destroys cursors, but retains access plans for dynamic SQL statements (including non-query statements)</li> <li>SQL_CB_PRESERVE, retains cursors and access plans for dynamic statements (including non-query statements). Applications can continue to fetch data, or close the cursor and re-execute the query without re-preparing the statement.</li> </ul> <p>After COMMIT, a FETCH must be issued to reposition the cursor before actions such as positioned updates or deletes can be taken.</p>
SQL_CURSOR_ROLLBACK_BEHAVIOR [24]	16-bit integer	<p>Indicates how a ROLLBACK operation affects cursors. A value of:</p> <ul style="list-style-type: none"> <li>SQL_CB_DELETE, destroys cursors and drops access plans for dynamic SQL statements.</li> <li>SQL_CB_CLOSE, destroys cursors, but retains access plans for dynamic SQL statements (including non-query statements)</li> <li>SQL_CB_PRESERVE, retains cursors and access plans for dynamic statements (including non-query statements). Applications can continue to fetch data, or close the cursor and re-execute the query without re-preparing the statement.</li> </ul> <p>DB2 servers do not have the SQL_CB_PRESERVE property.</p>
SQL_DATA_SOURCE_NAME [2]	string	The name used as data source on the input to SQLConnect(), or the DSN keyword value in the SQLDriverConnect() connection string.
SQL_DATA_SOURCE_READ_ONLY [25]	string	A character string of "Y" indicates that the database is set to READ ONLY mode; an "N" indicates that it is not set to READ ONLY mode.
SQL_DATABASE_NAME [16]	string	<p>The name of the current database in use.</p> <p><b>Note:</b> Also returned by SELECT CURRENT SERVER on IBM DBMS's.</p>
SQL_DBMS_NAME [17]	string	<p>The name of the DBMS product being accessed. For example:</p> <ul style="list-style-type: none"> <li>"DB2/6000"</li> <li>"DB2/2"</li> </ul>
SQL_DBMS_VER [18]	string	The Version of the DBMS product accessed. A string of the form 'mm.vv.rrrr' where mm is the major version, vv is the minor version and rrrr is the release. For example, "02.01.0000" translates to major version 2, minor version 1, release 0.

Table 73 (Page 4 of 14). Information returned by SQLGetInfo

<i>InfoType</i>	<b>Format</b>	<b>Description and notes</b>
SQL_DEFAULT_TXN_ISOLATION [26]	32-bit mask	<p>The default transaction isolation level supported.</p> <p>One of the following masks are returned:</p> <ul style="list-style-type: none"> <li>• SQL_TXN_READ_UNCOMMITTED = Changes are immediately perceived by all transactions (dirty read, non-repeatable read, and phantoms are possible). This is equivalent to IBM's UR level.</li> <li>• SQL_TXN_READ_COMMITTED = Row read by transaction 1 can be altered and committed by transaction 2 (non-repeatable read and phantoms are possible) This is equivalent to IBM's CS level.</li> <li>• SQL_TXN_REPEATABLE_READ = A transaction can add or remove rows matching the search condition or a pending transaction (repeatable read, but phantoms are possible) This is equivalent to IBM's RS level.</li> <li>• SQL_TXN_SERIALIZABLE = Data affected by pending transaction is not available to other transactions (repeatable read, phantoms are not possible) This is equivalent to IBM's RR level.</li> <li>• SQL_TXN_VERSIONING = Not applicable to IBM DBMSs.</li> <li>• SQL_TXN_NOCOMMIT = Any changes are effectively committed at the end of a successful operation; no explicit commit or rollback is allowed. This is a DB2 for AS/400 isolation level.</li> </ul> <p>In IBM terminology,</p> <ul style="list-style-type: none"> <li>• SQL_TXN_READ_UNCOMMITTED is uncommitted read;</li> <li>• SQL_TXN_READ_COMMITTED is cursor stability;</li> <li>• SQL_TXN_REPEATABLE_READ is read stability;</li> <li>• SQL_TXN_SERIALIZABLE is repeatable read.</li> </ul>
SQL_DRIVER_HDBC [3]	32 bits	DB2 ODBC's current database handle.
SQL_DRIVER_HENV [4]	32 bits	DB2 ODBC's environment handle.
SQL_DRIVER_HLIB [76]	32 bits	Reserved.
SQL_DRIVER_HSTMT [5]	32 bits	DB2 ODBC's current statement handle for the current connection.
SQL_DRIVER_NAME [6]	string	The file name of the DB2 ODBC implementation. DB2 ODBC returns NULL.
SQL_DRIVER_ODBC_VER [77]	string	The version number of ODBC that the Driver supports. DB2 ODBC returns "2.1".
SQL_DRIVER_VER [7]	string	The version of the CLI driver. A string of the form 'mm.vv.rrrr' where mm is the major version, vv is the minor version and rrrr is the release. For example, "02.01.0000" translates to major version 2, minor version 1, release 0.
SQL_EXPRESSIONS_IN_ORDERBY [27]	string	The character string "Y" indicates the database server supports the DIRECT specification of expressions in the ORDER BY list, "N" indicates that it does not.

Table 73 (Page 5 of 14). Information returned by SQLGetInfo

<i>InfoType</i>	<b>Format</b>	<b>Description and notes</b>
SQL_FETCH_DIRECTION [8]	32-bit mask	<p>The supported fetch directions.</p> <p>The following bit-masks are used in conjunction with the flag to determine which options are supported.</p> <ul style="list-style-type: none"> <li>• SQL_FD_FETCH_NEXT</li> <li>• SQL_FD_FETCH_FIRST</li> <li>• SQL_FD_FETCH_LAST</li> <li>• SQL_FD_FETCH_PREV</li> <li>• SQL_FD_FETCH_ABSOLUTE</li> <li>• SQL_FD_FETCH_RELATIVE</li> <li>• SQL_FD_FETCH_RESUME</li> </ul>
SQL_FILE_USAGE [84]	16-bit integer	Reserved. Zero is returned.
SQL_GETDATA_EXTENSIONS [81]	32-bit mask	<p>Indicates whether extensions to the SQLGetData() function are supported. The following extensions are currently identified and supported by DB2 ODBC:</p> <ul style="list-style-type: none"> <li>• SQL_GD_ANY_COLUMN, SQLGetData() can be called for unbound columns that precede the last bound column.</li> <li>• SQL_GD_ANY_ORDER, SQLGetData() can be called for columns in any order.</li> </ul> <p>ODBC also defines SQL_GD_BLOCK and SQL_GD_BOUND; these bits are not returned by DB2 ODBC.</p>
SQL_GROUP_BY [88]	16-bit integer	<p>Indicates the degree of support for the GROUP BY clause by the server:</p> <ul style="list-style-type: none"> <li>• SQL_GB_NO_RELATION, there is no relationship between the columns in the GROUP BY and in the SELECT list</li> <li>• SQL_GB_NOT_SUPPORTED, GROUP BY not supported</li> <li>• SQL_GB_GROUP_BY_EQUALS_SELECT, GROUP BY must include all non-aggregated columns in the select list.</li> <li>• SQL_GB_GROUP_BY_CONTAINS_SELECT, the GROUP BY clause must contain all non-aggregated columns in the SELECT list.</li> </ul>
SQL_IDENTIFIER_CASE [28]	16-bit integer	<p>Indicates case sensitivity of object names (such as table-name).</p> <p>A value of:</p> <ul style="list-style-type: none"> <li>• SQL_IC_UPPER = identifier names are stored in upper case in the system catalog.</li> <li>• SQL_IC_LOWER = identifier names are stored in lower case in the system catalog.</li> <li>• SQL_IC_SENSITIVE = identifier names are case sensitive, and are stored in mixed case in the system catalog.</li> <li>• SQL_IC_MIXED = identifier names are not case sensitive, and are stored in mixed case in the system catalog.</li> </ul> <p><b>Note:</b> Identifier names in IBM DBMSs are not case sensitive.</p>
SQL_IDENTIFIER_QUOTE_CHAR [29]	string	Indicates the character used to surround a delimited identifier.
SQL_KEYWORDS [89]	string	This is a string of all the <i>keywords</i> at the DBMS that are not in the ODBC's list of reserved words.
SQL_LIKE_ESCAPE_CLAUSE [113]	string	A character string that indicates if an escape character is supported for the metacharacters percent and underscore in a LIKE predicate.

Table 73 (Page 6 of 14). Information returned by SQLGetInfo

<i>InfoType</i>	<b>Format</b>	<b>Description and notes</b>
SQL_LOCK_TYPES [78]	32-bit mask	Reserved option, zero is returned for the bit-mask.
SQL_MAX_BINARY_LITERAL_LEN [112]	32-bit integer	A 32-bit integer value specifying the maximum length of a hexadecimal literal in a SQL statement.
SQL_MAX_CHAR_LITERAL_LEN [108]	32-bit integer	The maximum length of a character literal in an SQL statement (in bytes).
SQL_MAX_COLUMN_NAME_LEN [30]	16-bit integer	The maximum length of a column name (in bytes).
SQL_MAX_COLUMNS_IN_GROUP_BY [97]	16-bit integer	Indicates the maximum number of columns that the server supports in a GROUP BY clause. Zero if no limit.
SQL_MAX_COLUMNS_IN_INDEX [98]	16-bit integer	Indicates the maximum number of columns that the server supports in an index. Zero if no limit.
SQL_MAX_COLUMNS_IN_ORDER_BY [99]	16-bit integer	Indicates the maximum number of columns that the server supports in an ORDER BY clause. Zero if no limit.
SQL_MAX_COLUMNS_IN_SELECT [100]	16-bit integer	Indicates the maximum number of columns that the server supports in a select list. Zero if no limit.
SQL_MAX_COLUMNS_IN_TABLE [101]	16-bit integer	Indicates the maximum number of columns that the server supports in a base table. Zero if no limit.
SQL_MAX_CURSOR_NAME_LEN [31]	16-bit integer	The maximum length of a cursor name (in bytes).
SQL_MAX_INDEX_SIZE [102]	32-bit integer	Indicates the maximum size in bytes that the server supports for the combined columns in an index. Zero if no limit.
SQL_MAX_OWNER_NAME_LEN [32] SQL_MAX_SCHEMA_NAME_LEN	16-bit integer	The maximum length of a schema qualifier name (in bytes).
SQL_MAX_PROCEDURE_NAME_LEN [33]	16-bit integer	The maximum length of a procedure name (in bytes).
SQL_MAX_QUALIFIER_NAME_LEN [34] SQL_MAX_CATALOG_NAME_LEN	16-bit integer	The maximum length of a catalog qualifier name; first part of a 3 part table name (in bytes).
SQL_MAX_ROW_SIZE [104]	32-bit integer	Specifies the maximum length in bytes that the server supports in single row of a base table. Zero if no limit.
SQL_MAX_ROW_SIZE_INCLUDES_LONG [103]	string	Set to "Y" to indicate that the value returned by SQL_MAX_ROW_SIZE <i>InfoType</i> includes the length of product-specific <i>long string</i> data types. Otherwise, set to "N".
SQL_MAX_STATEMENT_LEN [105]	32-bit integer	Indicates the maximum length of an SQL statement string in bytes, including the number of white spaces in the statement.
SQL_MAX_TABLE_NAME_LEN [35]	16-bit integer	The maximum length of a table name (in bytes).
SQL_MAX_TABLES_IN_SELECT [106]	16-bit integer	Indicates the maximum number of table names allowed in a FROM clause in a <query specification>.
SQL_MAX_USER_NAME_LEN [107]	16-bit integer	Indicates the maximum size allowed for a <user identifier> (in bytes).
SQL_MULT_RESULT_SETS [36]	string	The character string "Y" indicates that the database supports multiple result sets, "N" indicates that it does not.
SQL_MULTIPLE_ACTIVE_TXN [37]	string	The character string "Y" indicates that active transactions on multiple connections are allowed. "N" indicates that only one connection at a time can have an active transaction.
SQL_NEED_LONG_DATA_LEN [111]	string	A character string reserved for the use of ODBC. "N" is always returned.

## SQLGetInfo

Table 73 (Page 7 of 14). Information returned by SQLGetInfo

<i>InfoType</i>	<b>Format</b>	<b>Description and notes</b>
SQL_NON_NULLABLE_COLUMNS [75]	16-bit integer	Indicates whether non-nullable columns are supported: <ul style="list-style-type: none"> <li>• SQL_NNC_NON_NULL, columns can be defined as NOT NULL.</li> <li>• SQL_NNC_NULL, columns can not be defined as NOT NULL.</li> </ul>
SQL_NULL_COLLATION [85]	16-bit integer	Indicates where NULLs are sorted in a list: <ul style="list-style-type: none"> <li>• SQL_NC_HIGH, null values sort high</li> <li>• SQL_NC_LOW, to indicate that null values sort low</li> </ul>
SQL_NUMERIC_FUNCTIONS [49]	32-bit mask	Indicates the ODBC scalar numeric functions supported. These functions are intended to be used with the ODBC vendor escape sequence described in “Using vendor escape clauses” on page 404.  The following bit-masks are used to determine which numeric functions are supported: <ul style="list-style-type: none"> <li>• SQL_FN_NUM_ABS</li> <li>• SQL_FN_NUM_ACOS</li> <li>• SQL_FN_NUM_ASIN</li> <li>• SQL_FN_NUM_ATAN</li> <li>• SQL_FN_NUM_ATAN2</li> <li>• SQL_FN_NUM_CEILING</li> <li>• SQL_FN_NUM_COS</li> <li>• SQL_FN_NUM_COT</li> <li>• SQL_FN_NUM_DEGREES</li> <li>• SQL_FN_NUM_EXP</li> <li>• SQL_FN_NUM_FLOOR</li> <li>• SQL_FN_NUM_LOG</li> <li>• SQL_FN_NUM_LOG10</li> <li>• SQL_FN_NUM_MOD</li> <li>• SQL_FN_NUM_PI</li> <li>• SQL_FN_NUM_POWER</li> <li>• SQL_FN_NUM_RADIANS</li> <li>• SQL_FN_NUM_RAND</li> <li>• SQL_FN_NUM_ROUND</li> <li>• SQL_FN_NUM_SIGN</li> <li>• SQL_FN_NUM_SIN</li> <li>• SQL_FN_NUM_SQRT</li> <li>• SQL_FN_NUM_TAN</li> <li>• SQL_FN_NUM_TRUNCATE</li> </ul>
SQL_ODBC_API_CONFORMANCE [9]	16-bit integer	The level of ODBC conformance. <ul style="list-style-type: none"> <li>• SQL_OAC_NONE</li> <li>• SQL_OAC_LEVEL1</li> <li>• SQL_OAC_LEVEL2</li> </ul>
SQL_ODBC_SAG_CLI_CONFORMANCE [12]	16-bit integer	The compliance to the functions of the SQL Access Group (SAG) CLI specification.  A value of: <ul style="list-style-type: none"> <li>• SQL_OSCC_NOT_COMPLIANT - the driver is not SAG-compliant.</li> <li>• SQL_OSCC_COMPLIANT - the driver is SAG-compliant.</li> </ul>



Table 73 (Page 8 of 14). Information returned by SQLGetInfo

<i>InfoType</i>	<b>Format</b>	<b>Description and notes</b>
SQL_ODBC_SQL_CONFORMANCE [15]	16-bit integer	<p>A value of:</p> <ul style="list-style-type: none"> <li>SQL_OSC_MINIMUM - means minimum ODBC SQL grammar supported</li> <li>SQL_OSC_CORE - means core ODBC SQL Grammar supported</li> <li>SQL_OSC_EXTENDED - means extended ODBC SQL Grammar supported</li> </ul> <p>For the definition of the above 3 types of ODBC SQL grammar, see <i>Microsoft ODBC 3.0 Programmer's Reference and SDK Guide</i>.</p>
SQL_ODBC_SQL_OPT_IEF [73]	string	The "Y" character string indicates that the data source supports Integrity Enhanced Facility (IEF) in SQL89 and in X/Open XPG4 Embedded SQL; an "N" indicates it does not.
SQL_ODBC_VER [10]	string	<p>The version number of ODBC that the driver manager supports.</p> <p>DB2 ODBC returns the string "02.10".</p>
SQL_OJ_CAPABILITIES [65003]	32-bit mask	<p>A 32-bit bit-mask enumerating the types of outer join supported.</p> <p>The bitmasks are:</p> <ul style="list-style-type: none"> <li>SQL_OJ_LEFT : Left outer join is supported.</li> <li>SQL_OJ_RIGHT : Right outer join is supported.</li> <li>SQL_OJ_FULL : Full outer join is supported.</li> <li>SQL_OJ_NESTED : Nested outer join is supported.</li> <li>SQL_OJ_NOT_ORDERED : The order of the tables underlying the columns in the outer join ON clause need not be in the same order as the tables in the JOIN clause.</li> <li>SQL_OJ_INNER : The inner table of an outer join can also be an inner join.</li> <li>SQL_OJ_ALL_COMPARISONS_OPS : Any predicate can be used in the outer join ON clause. If this bit is not set, the equality (=) operator is the only valid comparison operator in the ON clause.</li> </ul>
SQL_ORDER_BY_COLUMNS_IN_SELECT [90]	string	Set to "Y" if columns in the ORDER BY clauses must be in the select list; otherwise set to "N".
SQL_OUTER_JOINS [38]	string	<p>The character string:</p> <ul style="list-style-type: none"> <li>"Y" indicates that outer joins are supported, and DB2 ODBC supports the ODBC outer join request syntax.</li> <li>"N" indicates that it is not supported.</li> </ul> <p>(See "Using vendor escape clauses" on page 404)</p>
SQL_OWNER_TERM [39] SQL_SCHEMA_TERM	string	The database vendor's terminology for a schema (owner)

## SQLGetInfo

Table 73 (Page 9 of 14). Information returned by SQLGetInfo

<i>InfoType</i>	<b>Format</b>	<b>Description and notes</b>
SQL_OWNER_USAGE [91]	32-bit mask	Indicates the type of SQL statements that have schema (owners) associated with them when these statements are executed. Schema qualifiers (owners) are: <ul style="list-style-type: none"> <li>• SQL_OU_DML_STATEMENTS - supported in all DML statements.</li> <li>• SQL_OU_PROCEDURE_INVOCATION - supported in the procedure invocation statement.</li> <li>• SQL_OU_TABLE_DEFINITION - supported in all table definition statements.</li> <li>• SQL_OU_INDEX_DEFINITION - supported in all index definition statements.</li> <li>• SQL_OU_PRIVILEGE_DEFINITION - supported in all privilege definition statements (i.e. grant and revoke statements).</li> </ul>
SQL_POS_OPERATIONS [79]	32-bit mask	Reserved option, zero is returned for the bit-mask.
SQL_POSITIONED_STATEMENTS [80]	32-bit mask	Indicates the degree of support for positioned UPDATE and positioned DELETE statements: <ul style="list-style-type: none"> <li>• SQL_PS_POSITIONED_DELETE</li> <li>• SQL_PS_POSITIONED_UPDATE</li> <li>• SQL_PS_SELECT_FOR_UPDATE, indicates whether or not the server requires the FOR UPDATE clause to be specified on a &lt;query expression&gt; in order for a column to be updateable using the cursor.</li> </ul>
SQL_PROCEDURE_TERM [40]	string	The name a database vendor uses for a procedure
SQL_PROCEDURES [21]	string	A character string of "Y" indicates that the data source supports procedures and DB2 ODBC supports the ODBC procedure invocation syntax specified in "Using stored procedures" on page 389. "N" indicates that it does not.
SQL_QUALIFIER_LOCATION [114]	16-bit integer	A 16-bit integer value indicated the position of the qualifier in a qualified table name. Zero indicates that qualified names are not supported.
SQL_QUALIFIER_NAME_SEPARATOR [41] SQL_CATALOG_NAME_SEPARATOR	string	The characters used as a separator between a catalog name and the qualified name element that follows it.
SQL_QUALIFIER_TERM [42] SQL_CATALOG_TERM	string	The database vendor's terminology for a qualifier.  The name that the vendor uses for the high order part of a three part name.  Since DB2 ODBC does not support three part names, a zero-length string is returned.  For non-ODBC applications, the SQL_CATALOG_TERM symbolic name should be used instead of SQL_QUALIFIER_NAME.
SQL_QUALIFIER_USAGE [92] SQL_CATALOG_USAGE	32-bit mask	This is similar to SQL_OWNER_USAGE except that this is used for catalog.

Table 73 (Page 10 of 14). Information returned by SQLGetInfo

<i>flnfoType</i>	Format	Description and notes
SQL_QUOTED_IDENTIFIER_CASE [93]	16-bit integer	<p>Returns:</p> <ul style="list-style-type: none"> <li>SQL_IC_UPPER - quoted identifiers in SQL are case insensitive and stored in upper case in the system catalog.</li> <li>SQL_IC_LOWER - quoted identifiers in SQL are case insensitive and are stored in lower case in the system catalog.</li> <li>SQL_IC_SENSITIVE - quoted identifiers (delimited identifiers) in SQL are case sensitive and are stored in mixed case in the system catalog.</li> <li>SQL_IC_MIXED - quoted identifiers in SQL are case insensitive and are stored in mixed case in the system catalog.</li> </ul> <p>This should be contrasted with the SQL_IDENTIFIER_CASE <i>flnfoType</i> which is used to determine how (unquoted) identifiers are stored in the system catalog.</p>
SQL_ROW_UPDATES [11]	string	A character string of "Y" indicates changes are detected in rows between multiple fetches of the same rows, "N" indicates that changes are not detected.
SQL_SCROLL_CONCURRENCY [43]	32-bit mask	<p>Indicates the concurrency options supported for the cursor.</p> <p>The following bit-masks are used in conjunction with the flag to determine which options are supported:</p> <ul style="list-style-type: none"> <li>SQL_SCCO_READ_ONLY</li> <li>SQL_SCCO_LOCK</li> <li>SQL_SCCO_OPT_TIMESTAMP</li> <li>SQL_SCCO_OPT_VALUES</li> </ul> <p>DB2 ODBC returns SQL_SCCO_LOCK indicating that the lowest level of locking that is sufficient to ensure the row can be updated is used.</p>
SQL_SCROLL_OPTIONS [44]	32-bit mask	<p>The scroll options supported for scrollable cursors.</p> <p>The following bit-masks are used in conjunction with the flag to determine which options are supported:</p> <ul style="list-style-type: none"> <li>SQL_SO_FORWARD_ONLY</li> <li>SQL_SO_KEYSET_DRIVEN</li> <li>SQL_SO_STATIC</li> <li>SQL_SO_DYNAMIC</li> <li>SQL_SO_MIXED</li> </ul> <p>DB2 ODBC returns SQL_SO_FORWARD_ONLY, indicating that the cursor scrolls forward only.</p>
SQL_SEARCH_PATTERN_ESCAPE [14]	string	Used to specify what the driver supports as an escape character for catalog functions such as (SQLTables(), SQLColumns()).
SQL_SERVER_NAME [13]	string	The name of DB2 subsystem to which the application is connected.
SQL_SPECIAL_CHARACTERS [94]	string	Contains all the characters in addition to a...z, A...Z, 0...9, and _ that the server allows in non-delimited identifiers.

Table 73 (Page 11 of 14). Information returned by SQLGetInfo

<i>InfoType</i>	<b>Format</b>	<b>Description and notes</b>
SQL_STATIC_SENSITIVITY [83]	32-bit mask	<p>Indicates whether changes made by an application with a positioned UPDATE or DELETE statement can be detected by that application:</p> <ul style="list-style-type: none"> <li>• SQL_SS_ADDITIONS: Added rows are visible to the cursor; the cursor can scroll to these rows. All DB2 servers see added rows.</li> <li>• SQL_SS_DELETIONS: Deleted rows are no longer available to the cursor and do not leave a hole in the result set; after the cursor scrolls from a deleted row, it cannot return to that row.</li> <li>• SQL_SS_UPDATES: Updates to rows are visible to the cursor; if the cursor scrolls from and returns to an updated row, the data returned by the cursor is the updated data, not the original data.</li> </ul>
SQL_STRING_FUNCTIONS [50]	32-bit mask	<p>Indicates which string functions are supported.</p> <p>The following bit-masks are used to determine which string functions are supported:</p> <ul style="list-style-type: none"> <li>• SQL_FN_STR_ASCII</li> <li>• SQL_FN_STR_CHAR</li> <li>• SQL_FN_STR_CONCAT</li> <li>• SQL_FN_STR_DIFFERENCE</li> <li>• SQL_FN_STR_INSERT</li> <li>• SQL_FN_STR_LCASE</li> <li>• SQL_FN_STR_LEFT</li> <li>• SQL_FN_STR_LENGTH</li> <li>• SQL_FN_STR_LOCATE</li> <li>• SQL_FN_STR_LOCATE_2</li> <li>• SQL_FN_STR_LTRIM</li> <li>• SQL_FN_STR_REPEAT</li> <li>• SQL_FN_STR_REPLACE</li> <li>• SQL_FN_STR_RIGHT</li> <li>• SQL_FN_STR_RTRIM</li> <li>• SQL_FN_STR_SOUNDEX</li> <li>• SQL_FN_STR_SPACE</li> <li>• SQL_FN_STR_SUBSTRING</li> <li>• SQL_FN_STR_UCASE</li> </ul> <p>If an application can call the LOCATE scalar function with the <i>string1</i>, <i>string2</i>, and <i>start</i> arguments, the SQL_FN_STR_LOCATE bitmask is returned. If an application can only call the LOCATE scalar function with the <i>string1</i> and <i>string2</i>, the SQL_FN_STR_LOCATE_2 bitmask is returned. If the LOCATE scalar function is fully supported, both bitmasks are returned.</p>
SQL_SUBQUERIES [95]	32-bit mask	<p>Indicates which predicates support subqueries:</p> <ul style="list-style-type: none"> <li>• SQL_SQ_COMPARISON - the <i>comparison</i> predicate</li> <li>• SQL_SQ_CORRELATE_SUBQUERIES - all predicates</li> <li>• SQL_SQ_EXISTS - the <i>exists</i> predicate</li> <li>• SQL_SQ_IN - the <i>in</i> predicate</li> <li>• SQL_SQ_QUANTIFIED - the predicates containing a quantification scalar function.</li> </ul>

Table 73 (Page 12 of 14). Information returned by SQLGetInfo

<i>InfoType</i>	<b>Format</b>	<b>Description and notes</b>
SQL_SYSTEM_FUNCTIONS [51]	32-bit mask	<p>Indicates which scalar system functions are supported.</p> <p>The following bit-masks are used to determine which scalar system functions are supported:</p> <ul style="list-style-type: none"> <li>• SQL_FN_SYS_DBNAME</li> <li>• SQL_FN_SYS_IFNULL</li> <li>• SQL_FN_SYS_USERNAME</li> </ul> <p><b>Note:</b> These functions are intended to be used with the escape sequence in ODBC.</p>
SQL_TABLE_TERM [45]	string	The database vendor's terminology for a table.
SQL_TIMEDATE_ADD_INTERVALS [109]	32-bit mask	<p>Indicates whether or not the special ODBC system function <code>TIMESTAMPADD</code> is supported, and, if it is, which intervals are supported.</p> <p>The following bitmasks are used to determine which intervals are supported:</p> <ul style="list-style-type: none"> <li>• SQL_FN_TSI_FRAC_SECOND</li> <li>• SQL_FN_TSI_SECOND</li> <li>• SQL_FN_TSI_MINUTE</li> <li>• SQL_FN_TSI_HOUR</li> <li>• SQL_FN_TSI_DAY</li> <li>• SQL_FN_TSI_WEEK</li> <li>• SQL_FN_TSI_MONTH</li> <li>• SQL_FN_TSI_QUARTER</li> <li>• SQL_FN_TSI_YEAR</li> </ul>
SQL_TIMEDATE_DIFF_INTERVALS [110]	32-bit mask	<p>Indicates whether or not the special ODBC system function <code>TIMESTAMPDIFF</code> is supported, and, if it is, which intervals are supported.</p> <p>The following bitmasks are used to determine which intervals are supported:</p> <ul style="list-style-type: none"> <li>• SQL_FN_TSI_FRAC_SECOND</li> <li>• SQL_FN_TSI_SECOND</li> <li>• SQL_FN_TSI_MINUTE</li> <li>• SQL_FN_TSI_HOUR</li> <li>• SQL_FN_TSI_DAY</li> <li>• SQL_FN_TSI_WEEK</li> <li>• SQL_FN_TSI_MONTH</li> <li>• SQL_FN_TSI_QUARTER</li> <li>• SQL_FN_TSI_YEAR</li> </ul>

Table 73 (Page 13 of 14). Information returned by SQLGetInfo

<i>InfoType</i>	<i>Format</i>	<i>Description and notes</i>
SQL_TIMEDATE_FUNCTIONS [52]	32-bit mask	<p>Indicates which time and date functions are supported.</p> <p>The following bit-masks are used to determine which date functions are supported:</p> <ul style="list-style-type: none"> <li>SQL_FN_TD_CURDATE</li> <li>SQL_FN_TD_CURTIME</li> <li>SQL_FN_TD_DAYNAME</li> <li>SQL_FN_TD_DAYOFMONTH</li> <li>SQL_FN_TD_DAYOFWEEK</li> <li>SQL_FN_TD_DAYOFYEAR</li> <li>SQL_FN_TD_HOUR</li> <li>SQL_FN_TD_JULIAN_DAY</li> <li>SQL_FN_TD_MINUTE</li> <li>SQL_FN_TD_MONTH</li> <li>SQL_FN_TD_MONTHNAME</li> <li>SQL_FN_TD_NOW</li> <li>SQL_FN_TD_QUARTER</li> <li>SQL_FN_TD_SECOND</li> <li>SQL_FN_TD_SECONDS_SINCE_MIDNIGHT</li> <li>SQL_FN_TD_TIMESTAMPADD</li> <li>SQL_FN_TD_TIMESTAMPDIFF</li> <li>SQL_FN_TD_WEEK</li> <li>SQL_FN_TD_YEAR</li> </ul> <p><b>Note:</b> These functions are intended to be used with the escape sequence in ODBC.</p>
SQL_TXN_CAPABLE [46]	16-bit integer	<p>Indicates whether transactions can contain DDL or DML or both.</p> <ul style="list-style-type: none"> <li>SQL_TC_NONE = transactions not supported.</li> <li>SQL_TC_DML = transactions can only contain DML statements (SELECT, INSERT, UPDATE, DELETE, etc.) DDL statements (CREATE TABLE, DROP INDEX, etc.) encountered in a transaction cause an error.</li> <li>SQL_TC_DDL_COMMIT = transactions can only contain DML statements. DDL statements encountered in a transaction cause the transaction to be committed.</li> <li>SQL_TC_DDL_IGNORE = transactions can only contain DML statements. DDL statements encountered in a transaction are ignored.</li> <li>SQL_TC_ALL = transactions can contain DDL and DML statements in any order.</li> </ul>
SQL_TXN_ISOLATION_OPTION [72]	32-bit mask	<p>The transaction isolation levels available at the currently connected database server.</p> <p>The following masks are used in conjunction with the flag to determine which options are supported:</p> <ul style="list-style-type: none"> <li>SQL_TXN_READ_UNCOMMITTED</li> <li>SQL_TXN_READ_COMMITTED</li> <li>SQL_TXN_REPEATABLE_READ</li> <li>SQL_TXN_SERIALIZABLE</li> <li>SQL_TXN_NOCOMMIT</li> <li>SQL_TXN_VERSIONING</li> </ul> <p>For descriptions of each level, see SQL_DEFAULT_TXN_ISOLATION.</p>
SQL_UNION [96]	32-bit mask	<p>Indicates if the server supports the UNION operator:</p> <ul style="list-style-type: none"> <li>SQL_U_UNION - supports the UNION clause</li> <li>SQL_U_UNION_ALL - supports the ALL keyword in the UNION clause</li> </ul> <p>If SQL_U_UNION_ALL is set, so is SQL_U_UNION.</p>

Table 73 (Page 14 of 14). Information returned by SQLGetInfo

<i>fInfoType</i>	Format	Description and notes
SQL_USER_NAME [47]	string	The user name used in a particular database. This is the identifier specified on the SQLConnect() call.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 74. SQLGetInfo SQLSTATEs

SQLSTATE	Description	Explanation
01004	Data truncated.	The requested information is returned as a string and its length exceeds the length of the application buffer as specified in <i>cbInfoValueMax</i> . The argument <i>pcbInfoValue</i> contains the actual (not truncated) length of the requested information. (Function returns SQL_SUCCESS_WITH_INFO.)
08003	Connection is closed.	The type of information requested in <i>fInfoType</i> requires an open connection. Only SQL_ODBC_VER does not require an open connection.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	The argument <i>rgbInfoValue</i> is a null pointer.  The <i>fInfoType</i> is SQL_DRIVER_HSTMT and the value pointed to by <i>rgbInfoValue</i> is not a valid handle.
S1090	Invalid string or buffer length.	The value specified for argument <i>cbInfoValueMax</i> is less than 0.
S1096	Information type out of range.	An invalid <i>fInfoType</i> is specified.
S1C00	Driver not capable.	The value specified in the argument <i>fInfoType</i> is not supported by either DB2 ODBC or the data source.

## Restrictions

None.

### Example

```
/* ... */
*/
/* Check to see if SQLGetInfo() is supported */
rc = SQLGetFunctions(hdbc, SQL_API_SQLGETINFO, &supported);

if (supported == SQL_TRUE) { /* get information about current connection */

    rc = SQLGetInfo(hdbc, SQL_DATA_SOURCE_NAME, buffer, 255, &outlen);
    printf("    Server Name: %s\n", buffer);

    rc = SQLGetInfo(hdbc, SQL_DATABASE_NAME,    buffer, 255, &outlen);
    printf("    Database Name: %s\n", buffer);

    rc = SQLGetInfo(hdbc, SQL_SERVER_NAME,      buffer, 255, &outlen);
    printf("    Instance Name: %s\n", buffer);

    rc = SQLGetInfo(hdbc, SQL_DBMS_NAME,       buffer, 255, &outlen);
    printf("    DBMS Name: %s\n", buffer);

    rc = SQLGetInfo(hdbc, SQL_DBMS_VER,        buffer, 255, &outlen);
    printf("    DBMS Version: %s\n", buffer);

    rc = SQLGetInfo(hdbc, SQL_DRIVER_NAME,     buffer, 255, &outlen);
    printf("    CLI Driver Name: %s\n", buffer);

    rc = SQLGetInfo(hdbc, SQL_DRIVER_VER,      buffer, 255, &outlen);
    printf("CLI Driver Version: %s\n", buffer);

    rc = SQLGetInfo(hdbc, SQL_ODBC_SQL_CONFORMANCE, &output,
                    sizeof(output), &outlen);
    switch (output) {
    case 0:
        strcpy(buffer, "Minimum Grammar");
        break;
    case 1:
        strcpy(buffer, "Core Grammar");
        break;
    case 2:
        strcpy(buffer, "Extended Grammar");
        break;
    default:
        printf("Error calling getinfo!");
        return (SQL_ERROR);
    }
    printf("ODBC SQL Conformance Level: %s\n", buffer);
} else {
    printf("SQLGetInfo is not supported!\n");
}
/* ... */
```

### References

- “SQLGetTypeInfo - Get data type information” on page 254



## SQLGetLength - Retrieve length of a string value

### Purpose

Specification:			
----------------	--	--	--

SQLGetLength() retrieves the length of a large object value, referenced by a large object locator that is returned from the server (as a result of a fetch, or an SQLGetSubString() call) during the current transaction.

### Syntax

```
SQLRETURN SQLGetLength (SQLHSTMT          hstmt,
                        SQLSMALLINT       LocatorCType,
                        SQLINTEGER        Locator,
                        SQLINTEGER FAR    *StringLength,
                        SQLINTEGER FAR    *IndicatorValue);
```

### Function arguments

Table 75. SQLGetLength arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle. This can be any statement handle that is allocated but does not currently have a prepared statement assigned to it.
SQLSMALLINT	LocatorCType	input	The C type of the source LOB locator. This can be: <ul style="list-style-type: none"> <li>• SQL_C_BLOB_LOCATOR</li> <li>• SQL_C_CLOB_LOCATOR</li> <li>• SQL_C_DBCLOB_LOCATOR</li> </ul>
SQLINTEGER	Locator	input	Must be set to the LOB locator value.
SQLINTEGER *	StringLength	output	The length of the returned information in <i>rgbValue</i> in bytes <sup>a</sup> if the target C buffer type is intended for a binary or character string variable and not a locator value. If the pointer is set to NULL, nothing is returned.
SQLINTEGER *	IndicatorValue	output	Always set to zero.

#### Note:

<sup>a</sup> This is in bytes even for DBCLOB data.

### Usage

SQLGetLength() can determine the length of the data value represented by a LOB locator. Applications use it to determine the overall length of the referenced LOB value so that the appropriate strategy for obtaining some or all of that value can be chosen.

The *Locator* argument can contain any valid LOB locator that is not explicitly freed using a FREE LOCATOR statement or implicitly freed because the transaction during which it was created has terminated.

## SQLGetLength

The statement handle must not be associated with any prepared statements or catalog function calls.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 76. SQLGetLength SQLSTATEs

SQLSTATE	Description	Explanation
07006	Invalid conversion.	The combination of <i>LocatorCType</i> and <i>Locator</i> is not valid.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1003	Program type out of range.	<i>LocatorCType</i> is not one of the following: <ul style="list-style-type: none"><li>• SQL_C_CLOB_LOCATOR</li><li>• SQL_C_BLOB_LOCATOR</li><li>• SQL_C_DBCLOB_LOCATOR</li></ul>
S1009	Invalid argument value.	Pointer to <i>StringLength</i> is NULL.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1C00	Driver not capable.	The application is currently connected to a data source that does not support large objects.
0F001	The LOB token variable does not currently represent any value.	The value specified for <i>Locator</i> is not associated with a LOB locator.

## Restrictions

This function is not available when connected to a DB2 server that does not support large objects. Call `SQLGetFunctions()` with the function type set to `SQL_API_SQLGETLENGTH` and check the *fExists* output argument to determine if the function is supported for the current connection.

## Example

See “Example” on page 240.

## References

- “SQLBindCol - Bind a column to an application variable” on page 87
- “SQLExtendedFetch - Extended fetch (fetch array of rows)” on page 162
- “SQLFetch - Fetch next row” on page 169
- “SQLGetPosition - Return starting position of string” on page 238
- “SQLGetSubString - Retrieve portion of a string value” on page 251

---

## SQLGetPosition - Return starting position of string

### Purpose

<b>Specification:</b>			
-----------------------	--	--	--

SQLGetPosition() returns the starting position of one string within a LOB value (the source). The source value must be a LOB locator; the search string can be a LOB locator or a literal string.

The source and search LOB locators can be any value that is returned from the database from a fetch or a SQLGetSubString() call during the current transaction.

### Syntax

```
SQLRETURN SQLGetPosition (SQLHSTMT          hstmt,
                          SQLSMALLINT       LocatorCType,
                          SQLINTEGER         SourceLocator,
                          SQLINTEGER         SearchLocator,
                          SQLCHAR          FAR *SearchLiteral,
                          SQLINTEGER         SearchLiteralLength,
                          SQLUINTEGER        FromPosition,
                          SQLUINTEGER FAR   *LocatedAt,
                          SQLUINTEGER FAR   *IndicatorValue);
```

### Function arguments

Table 77 (Page 1 of 2). SQLGetPosition arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle. This can be any statement handle that is allocated but does not currently have a prepared statement assigned to it.
SQLSMALLINT	LocatorCType	input	The C type of the source LOB locator. This can be: <ul style="list-style-type: none"> <li>• SQL_C_BLOB_LOCATOR</li> <li>• SQL_C_CLOB_LOCATOR</li> <li>• SQL_C_DBCLOB_LOCATOR</li> </ul>
SQLINTEGER	Locator	input	<i>Locator</i> must be set to the source LOB locator.
SQLINTEGER	SearchLocator	input	If the <i>SearchLiteral</i> pointer is NULL and if <i>SearchLiteralLength</i> is set to 0, then <i>SearchLocator</i> must be set to the LOB locator associated with the search string. Otherwise, this argument is ignored.
SQLCHAR *	SearchLiteral	input	This argument points to the area of storage that contains the search string literal. If <i>SearchLiteralLength</i> is 0, this pointer must be NULL.
SQLINTEGER	SearchLiteralLength	input	The length of the string in <i>SearchLiteral</i> (in bytes). <sup>a</sup> If this argument value is 0, then the argument <i>SearchLocator</i> is meaningful.
SQLUINTEGER	FromPosition	input	For BLOBs and CLOBs, this is the position of the first byte within the source string at which the search is to start. For DBCLOBs, this is the first character. The start byte or character is numbered 1.

Table 77 (Page 2 of 2). SQLGetPosition arguments

Data type	Argument	Use	Description
SQLUINTEGER *	LocatedAt	output	For BLOBs and CLOBs, this is the byte position at which the string was located or, if not located, the value zero. For DBCLOBs, this is the character position.  If the length of the source string is zero, the value 1 is returned.
SQLINTEGER *	IndicatorValue	output	Always set to zero.

**Note:**

<sup>a</sup> This is in bytes even for DBCLOB data.

## Usage

SQLGetPosition() is used in conjunction with SQLGetSubString() to obtain any portion of a string in a random manner. To use SQLGetSubString(), the location of the substring within the overall string must be known in advance. In situations where a search string finds the start of that substring, SQLGetPosition() can be used to obtain the starting position of the substring.

The *Locator* and *SearchLocator* arguments (if used) can contain any valid LOB locator that is not explicitly freed using a FREE LOCATOR statement or implicitly freed because the transaction during which it was created has terminated.

The *Locator* and *SearchLocator* must have the same LOB locator type.

The statement handle must not be associated with any prepared statements or catalog function calls.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 78 (Page 1 of 2). SQLGetPosition SQLSTATES

SQLSTATE	Description	Explanation
07006	Invalid conversion.	The combination of <i>LocatorCType</i> and either of the LOB locator values is not valid.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
42818	The operands of an operator or function are not compatible.	The length of the pattern is longer than 4000 bytes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.

## SQLGetPosition

Table 78 (Page 2 of 2). SQLGetPosition SQLSTATES

SQLSTATE	Description	Explanation
S1009	Invalid argument value.	The pointer to the <i>LocatedAt</i> argument is NULL. The argument value for <i>FromPosition</i> is not greater than 0. <i>LocatorCType</i> is not one of SQL_C_CLOB_LOCATOR, SQL_C_BLOB_LOCATOR, or SQL_C_DBCLOB_LOCATOR.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1090	Invalid string or buffer length.	The value of <i>SearchLiteralLength</i> is less than 1, and not SQL_NTS.
S1C00	Driver not capable.	The application is currently connected to a data source that does not support large objects.
0F001	The LOB token variable does not currently represent any value.	The value specified for <i>Locator</i> or <i>SearchLocator</i> is not currently a LOB locator.

## Restrictions

This function is not available when connected to a DB2 server that does not support large objects. Call `SQLGetFunctions()` with the function type set to `SQL_API_SQLGETPOSITION` and check the *fExists* output argument to determine if the function is supported for the current connection.

## Example

```
/* ... */
SQLCHAR      stmt2[] =
              "SELECT resume FROM emp_resume "
              "WHERE empno = ? AND resume_format = 'ascii'";
/* ... */
/*****
** Get CLOB locator to selected Resume **
*****/
rc = SQLSetParam(hstmt, 1, SQL_C_CHAR, SQL_CHAR, 7,
                0, Empno.s, &Empno.ind);

printf("\n>Enter an employee number:\n");
gets(Empno.s);

rc = SQLExecDirect(hstmt, stmt2, SQL_NTS);
rc = SQLBindCol(hstmt, 1, SQL_C_CLOB_LOCATOR, &ClobLoc1, 0,
                &pcbValue);
rc = SQLFetch(hstmt);
```

```

/*****
  Search CLOB locator to find "Interests"
  Get substring of resume (from position of interests to end)
*****/

rc = SQLAllocStmt(hdbc, &lhstmt);

/* Get total length */
rc = SQLGetLength(lhstmt, SQL_C_CLOB_LOCATOR, ClobLoc1, &SLength, &Ind);

/* Get Starting position */
rc = SQLGetPosition(lhstmt, SQL_C_CLOB_LOCATOR, ClobLoc1, 0,
                  "Interests", 9, 1, &Pos1, &Ind);

buffer = (SQLCHAR *)malloc(SLength - Pos1 + 1);

/* Get just the "Interests" section of the Resume CLOB */
/* (From Pos1 to end of CLOB) */
rc = SQLGetSubString(lhstmt, SQL_C_CLOB_LOCATOR, ClobLoc1, Pos1,
                    SLength - Pos1, SQL_C_CHAR, buffer, SLength - Pos1 + 1,
                    &OutLength, &Ind);

/* Print Interest section of Employee's resume */
printf("\nEmployee #: %s\n %s\n", Empno.s, buffer);
/* ... */

```

## References

- “SQLBindCol - Bind a column to an application variable” on page 87
- “SQLExtendedFetch - Extended fetch (fetch array of rows)” on page 162
- “SQLFetch - Fetch next row” on page 169
- “SQLGetFunctions - Get functions” on page 214
- “SQLGetLength - Retrieve length of a string value” on page 235
- “SQLGetSubString - Retrieve portion of a string value” on page 251

---

## SQLGetSQLCA - Get SQLCA data structure

### Purpose

<b>Specification:</b>			
-----------------------	--	--	--

SQLGetSQLCA() is used to return the SQLCA associated with preparing and executing an SQL statement, fetching data, or closing a cursor. The SQLCA can return information that supplements the information obtained by using SQLError().

For a detailed description of the SQLCA structure, see Appendix C of *DB2 SQL Reference*.

An SQLCA is not available if a function is processed strictly on the application side, such as allocating a statement handle. In this case, an empty SQLCA is returned with all values set to zero.

### Syntax

```
SQLRETURN SQLGetSQLCA (SQLHENV      henv,
                       SQLHDBC      hdbc,
                       SQLHSTMT     hstmt,
                       struct sqlca FAR *pSqlca );
```

### Function arguments

Table 79. SQLGetSQLCA arguments

Data type	Argument	Use	Description
SQLHENV	henv	input	Environment Handle
SQLHDBC	hdbc	input	Connection Handle
SQLHSTMT	hstmt	input	Statement Handle
SQLCA *	pqlca	output	SQL Communication Area

### Usage

The handles are used in the same way as for the SQLError() function. To obtain the SQLCA associated with:

- An environment, pass a valid environment handle. Set *hdbc* and *hstmt* to SQL\_NULL\_HDBC and SQL\_NULL\_HSTMT respectively.
- A connection, pass a valid database connection handle, and set *hstmt* to SQL\_NULL\_HSTMT. The *henv* argument is ignored.
- A statement, pass a valid statement handle. The *henv* and *hdbc* arguments are ignored.

If diagnostic information generated by one DB2 ODBC function is not retrieved before a function other than SQLError() is called with the same handle, the information for the previous function call is lost. This is true whether or not diagnostic information is generated for the second DB2 ODBC function call.



If a DB2 ODBC function is called that does not result in interaction with the DBMS, then the SQLCA contains all zeroes. Meaningful information is returned for the following functions:

- SQLCancel()
- SQLConnect(), SQLDisconnect()
- SQLExecDirect(), SQLExecute()
- SQLFetch()
- SQLPrepare()
- SQLTransact()
- SQLColumns()
- SQLConnect()
- SQLGetData (if LOB column is involved)
- SQLSetConnectOption() (for SQL\_AUTOCOMMIT)
- SQLStatistics()
- SQLTables()
- SQLColumnPrivileges()
- SQLExtendedFetch()
- SQLForeignKeys()
- SQLMoreResults()
- SQLPrimaryKeys()
- SQLProcedureColumns()
- SQLProcedures()
- SQLTablePrivileges()

## Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

None.

## Restrictions

None.

## Example

```

/*****
/* DB2 for OS/390 example: */
/* Prepares a query and executes that query against a non- */
/* existent table. Then invoke SQLGetSQLCA to extract */
/* native SQLCA data structure. Note that this API is NOT */
/* defined within ODBC, i.e. this is unique to IBM CLI. */
*****/

#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <sqlca.h>
#include "sqlcli1.h"

void print_sqlca (SQLHENV,          // prototype for print_sqlca
                 SQLHDBC,
                 SQLHSTMT);

int main( )
{
    SQLHENV      hEnv   = SQL_NULL_HENV;
    SQLHDBC      hDbc   = SQL_NULL_HDBC;
    SQLHSTMT     hStmt  = SQL_NULL_HSTMT;
    SQLRETURN    rc     = SQL_SUCCESS;
    SQLINTEGER   RETCODE = 0;
    char         *pDSN  = "STLEC1";
    SWORD        cbCursor;
    SDWORD       cbValue1;
    SDWORD       cbValue2;
    char         employee [30];
    int          salary = 0;
    int          param_salary = 30000;

    char         *stmt = "SELECT NAME, SALARY FROM EMPLOYEES WHERE SALARY > ?";

    (void) printf ("**** Entering CLIP11.\n\n");

    /*****
    /* Allocate Environment Handle */
    *****/

    RETCODE = SQLAllocEnv(&hEnv);

    if (RETCODE != SQL_SUCCESS)
        goto dberror;

    /*****
    /* Allocate Connection Handle to DSN */
    *****/

    RETCODE = SQLAllocConnect(hEnv,
                              &hDbc);

    if( RETCODE != SQL_SUCCESS ) // Could not get a Connect Handle
        goto dberror;

```

```

/*****
/* CONNECT TO data source (STLEC1) */
*****/

RETCODE = SQLConnect(hDbc,      // Connect handle
                    (SQLCHAR *) pDSN, // DSN
                    SQL_NTS,    // DSN is nul-terminated
                    NULL,       // Null UID
                    0,          // Null Auth string
                    NULL,
                    0);

if( RETCODE != SQL_SUCCESS )    // Connect failed
    goto dberror;

/*****
/* Allocate Statement Handles */
*****/

rc = SQLAllocStmt (hDbc,
                  &hStmt);

if (rc != SQL_SUCCESS)
    goto exit;

/*****
/* Prepare the query for multiple execution within current */
/* transaction. Note that query is collapsed when transaction */
/* is committed or rolled back. */
*****/

rc = SQLPrepare (hStmt,
                (SQLCHAR *) stmt,
                strlen(stmt));

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** PREPARE OF QUERY FAILED.\n");
    (void) print_sqlca (hStmt,
                      hDbc,
                      hEnv);
    goto dberror;
}

rc = SQLBindCol (hStmt,      // bind employee name
                1,
                SQL_C_CHAR,
                employee,
                sizeof(employee),
                &cbValue1);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** BIND OF NAME FAILED.\n");
    goto dberror;
}

```

```

rc = SQLBindCol (hStmt,          // bind employee salary
                2,
                SQL_C_LONG,
                &salary,
                0,
                &cbValue2);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** BIND OF SALARY FAILED.\n");
    goto dberror;
}

/*****
/* Bind parameter to replace '?' in query. This has an initial
/* value of 30000.
*****/

rc = SQLBindParameter (hStmt,
                      1,
                      SQL_PARAM_INPUT,
                      SQL_C_LONG,
                      SQL_INTEGER,
                      0,
                      0,
                      &param_salary,
                      0,
                      NULL);

/*****
/* Execute prepared statement to generate answer set.
*****/

rc = SQLExecute (hStmt);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** EXECUTE OF QUERY FAILED.\n");
    (void) print_sqlca (hStmt,
                      hDbc,
                      hEnv);
    goto dberror;
}

/*****
/* Answer Set is available -- Fetch rows and print employees
/* and salary.
*****/

(void) printf ("**** Employees whose salary exceeds %d follow.\n\n",
              param_salary);

while ((rc = SQLFetch (hStmt)) == SQL_SUCCESS)
{
    (void) printf ("**** Employee Name %s with salary %d.\n",
                  employee,
                  salary);
}

```

```

/*****
/* Deallocate Statement Handles -- statement is no longer in a
/* Prepared state.
*****/

rc = SQLFreeStmt (hStmt,
                  SQL_DROP);

/*****
/* DISCONNECT from data source
*****/

RETCODE = SQLDisconnect(hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Deallocate Connection Handle
*****/

RETCODE = SQLFreeConnect (hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Free Environment Handle
*****/

RETCODE = SQLFreeEnv (hEnv);

if (RETCODE == SQL_SUCCESS)
    goto exit;

dberror:
RETCODE=12;

exit:

(void) printf ("**** Exiting CLIP11.\n\n");

return RETCODE;
}

/*****
/* print_sqlca invokes SQLGetSQLCA and prints the native SQLCA.
*****/

void print_sqlca (SQLHENV hEnv ,
                 SQLHDBC hDbc ,
                 SQLHSTMT hStmt)
{
    SQLRETURN rc = SQL_SUCCESS;
    struct sqlca sqlca;
    struct sqlca *pSQLCA = &sqlca;
    int code ;
    char state [6];
    char errp [9];
    char tok [40];
    int count, len, start, end, i;

```

## SQLGetSQLCA

```
if ((rc = SQLGetSQLCA (hEnv ,
                      hDbc ,
                      hStmt,
                      pSQLCA)) != SQL_SUCCESS)
{
    (void) printf ("**** SQLGetSQLCA failed Return Code = %d.\n", rc);
    goto exit;
}

code = (int) pSQLCA->sqlcode;
memcpy (state, pSQLCA->sqlstate, 5);
state [5] = '\0';

(void) printf ("**** sqlcode = %d, sqlstate = %s.\n", code, state);

memcpy (errp, pSQLCA->sqlerrp, 8);
errp [8] = '\0';
(void) printf ("**** sqlerrp = %s.\n", errp);

if (pSQLCA->sqlerrml == 0)
    (void) printf ("**** No tokens.\n");
else
{
    for (len = 0, count = 0; len < pSQLCA->sqlerrml; len = ++end)
    {
        start = end = len;
        while ((pSQLCA->sqlerrmc [end] != 0xFF) &&
              (end < pSQLCA->sqlerrml))

            end++;
        if (start != end)
        {
            memcpy (tok, &pSQLCA->sqlerrmc[start],
                  (end-start));

            tok [end-start+1] = '\0';
            (void) printf ("**** Token # %d = %s.\n", count++, tok);
        }
    }
}

for (i = 0; i <= 5; i++)
    (void) printf ("**** sqlerrd # %d = %d.\n", i+1, pSQLCA->sqlerrd_i];

for (i = 0; i <= 10; i++)
    (void) printf ("**** sqwarn # %d = %c.\n", i+1, pSQLCA->sqlwarn_i]);

exit:
return;
}
```

## References

- “SQLError - Retrieve error information” on page 148

## SQLGetStmtOption - Returns current setting of a statement option

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	
----------------	----------	------------	--

SQLGetStmtOption() returns the current settings of the specified statement option.

These options are set using the SQLSetStmtOption() function.

### Syntax

```
SQLRETURN SQLGetStmtOption (SQLHSTMT          hstmt,
                             SQLUSMALLINT      fOption,
                             SQLPOINTER       pvParam);
```

### Function arguments

Table 80. SQLStmtOption arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle.
SQLUSMALLINT	fOption	input	Option to set. See Table 130 on page 335 for the list of statement options and their descriptions.
SQLPOINTER	pvParam	output	Value of the option. Depending on the value of <i>fOption</i> this can be a 32-bit integer value, or a pointer to a null terminated character string. The maximum length of any character string returned is SQL_MAX_OPTION_STRING_LENGTH bytes (excluding the null-terminator).

### Usage

See Table 130 on page 335 in the function description of SQLSetStmtOption() for a list of statement options. The following table lists the statement options that are read-only (can be read but not set).

Table 81. Statement options

<i>fOption</i>	Contents
SQL_ROW_NUMBER	A 32-bit integer value that specifies the number of the current row in the entire result set. If the number of the current row cannot be determined or there is no current row, 0 is returned.

**Note:** ODBC architecture also defines the read-only statement option SQL\_GET\_BOOKMARK. This option is not supported by DB2 ODBC. If it is specified, this function returns SQL\_ERROR (SQLSTATE S1011 -- Operation invalid at this time.)

### Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

### Diagnostics

Table 82. SQLGetStmtOption SQLSTATES

SQLSTATE	Description	Explanation
24000	Invalid cursor state.	There is no open cursor on the statement handle.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	<i>pvParam</i> was null.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1092	Option type out of range.	An invalid <i>fOption</i> value was specified.
S1C00	Driver not capable.	DB2 ODBC recognizes the option but does not support it.

### Restrictions

None.

### Example

```
/* ... */
rc = SQLGetStmtOption(hstmt, SQL_CURSOR_HOLD, &cursor_hold);
printf("Cursor With Hold is: ");
if (cursor_hold == SQL_CURSOR_HOLD_ON )
    printf("ON\n");
else
    printf("OFF\n");
/* ... */
```

### References

- “SQLSetConnectOption - Set connection option” on page 316
- “SQLSetStmtOption - Set statement option” on page 334



## SQLGetSubString - Retrieve portion of a string value

### Purpose

Specification:			
----------------	--	--	--

SQLGetSubString() retrieves a portion of a large object value, referenced by a LOB locator that the server returns (returned by a fetch or a previous SQLGetSubString() call) during the current transaction.

### Syntax

```
SQLRETURN SQLGetSubString (SQLHSTMT          hstmt,
                           SQLSMALLINT       LocatorCType,
                           SQLINTEGER         SourceLocator,
                           SQLUINTEGER       FromPosition,
                           SQLUINTEGER       ForLength,
                           SQLSMALLINT       TargetCType,
                           SQLPOINTER        rgbValue,
                           SQLINTEGER        cbValueMax,
                           SQLINTEGER FAR    *StringLength,
                           SQLINTEGER FAR    *IndicatorValue);
```

### Function arguments

Table 83 (Page 1 of 2). SQLGetSubString arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle. This can be any statement handle that is allocated but does not currently have a prepared statement assigned to it.
SQLSMALLINT	LocatorCType	input	The C type of the source LOB locator. This can be: <ul style="list-style-type: none"> <li>• SQL_C_BLOB_LOCATOR</li> <li>• SQL_C_CLOB_LOCATOR</li> <li>• SQL_C_DBCLOB_LOCATOR</li> </ul>
SQLINTEGER	Locator	input	<i>Locator</i> must be set to the source LOB locator value.
SQLUINTEGER	FromPosition	input	For BLOBs and CLOBs, this is the position of the first byte the function returns. For DBCLOBs, this is the first character. The start byte or character is numbered 1.
SQLUINTEGER	ForLength	input	This is the length of the string to be returned by the function. For BLOBs and CLOBs, this is the length in bytes. For DBCLOBs, this is the length in characters.  If <i>FromPosition</i> is less than the length of the source string but <i>FromPosition</i> + <i>ForLength</i> - 1 extends beyond the end of the source string, the result is padded on the right with the necessary number of characters (X'00' for BLOBs, single-byte blank character for CLOBs, and double-byte blank character for DBCLOBs).
SQLSMALLINT	TargetCType	input	The C data type of the <i>rgbValue</i> . The target must always be either a LOB locator C buffer type (SQL_C_CLOB_LOCATOR, SQL_C_BLOB_LOCATOR, SQL_C_DBCLOB_LOCATOR) or a C string variable (SQL_C_CHAR for CLOB, SQL_C_BINARY for BLOB, and SQL_C_DBCHAR for DBCLOB).

## SQLGetSubString

Table 83 (Page 2 of 2). SQLGetSubstring arguments

Data type	Argument	Use	Description
SQLPOINTER	rgbValue	output	Pointer to the buffer where the retrieved string value or a LOB locator is stored.
SQLINTEGER	cbValueMax	input	Maximum size of the buffer pointed to by <i>rgbValue</i> in bytes.
SQLINTEGER *	StringLength	output	The length of the returned information in <i>rgbValue</i> in bytes <sup>a</sup> if the target C buffer type is intended for a binary or character string variable and not a locator value. If the pointer is set to NULL, nothing is returned.
SQLINTEGER *	IndicatorValue	output	Always set to zero.

**Note:**

<sup>a</sup> This is in bytes even for DBCLOB data.

## Usage

SQLGetSubString() is used to obtain any portion of the string that is represented by the LOB locator. The target can be one of the following:

- An appropriate C string variable.
- A new LOB value that is created on the server. The LOB locator for this value can be assigned to a target application variable on the client.

SQLGetSubString() can be used as an alternative to SQLGetData() for getting data in pieces. In this case, a column is first bound to a LOB locator, which is then used to fetch the LOB as a whole or in pieces.

The *Locator* argument can contain any valid LOB locator that is not explicitly freed using a FREE LOCATOR statement or implicitly freed because the transaction during which it was created has terminated.

The statement handle must not be associated with any prepared statements or catalog function calls.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 84 (Page 1 of 2). SQLGetSubstring SQLSTATEs

SQLSTATE	Description	Explanation
01004	Data truncated.	The amount of returned data is longer than <i>cbValueMax</i> . Actual length available for return is stored in <i>StringLength</i> .

Table 84 (Page 2 of 2). SQLGetSubstring SQLSTATEs

SQLSTATE	Description	Explanation
07006	Invalid conversion.	The value specified for <i>TargetCType</i> is not SQL_C_CHAR, SQL_C_BINARY, SQL_C_DBCHAR or a LOB locator.  The value specified for <i>TargetCType</i> is inappropriate for the source (for example SQL_C_DBCHAR for a BLOB column).
22011	A substring error occurred.	<i>FromPosition</i> is greater than the length of the source string.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1003	Program type out of range.	<i>LocatorCType</i> is not one of the following: <ul style="list-style-type: none"> <li>• SQL_C_CLOB_LOCATOR</li> <li>• SQL_C_BLOB_LOCATOR</li> <li>• SQL_C_DBCLOB_LOCATOR</li> </ul>
S1009	Invalid argument value.	The value specified for <i>FromPosition</i> or for <i>ForLength</i> is not a positive integer.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1090	Invalid string or buffer length.	The value of <i>cbValueMax</i> is less than 0.
S1C00	Driver not capable.	The application is currently connected to a data source that does not support large objects.
0F001	The LOB token variable does not currently represent any value.	The value specified for <i>Locator</i> or <i>SearchLocator</i> is not currently a LOB locator.

## Restrictions

This function is not available when connected to a DB2 server that does not support large objects. Call `SQLGetFunctions()` with the function type set to `SQL_API_SQLGETSUBSTRING`, and check the *fExists* output argument to determine if the function is supported for the current connection.

## Example

See “Example” on page 240.

## References

- “SQLBindCol - Bind a column to an application variable” on page 87
- “SQLExtendedFetch - Extended fetch (fetch array of rows)” on page 162
- “SQLFetch - Fetch next row” on page 169
- “SQLGetLength - Retrieve length of a string value” on page 235
- “SQLGetSubString - Retrieve portion of a string value” on page 251

## SQLGetTypeInfo - Get data type information

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLGetTypeInfo() returns information about the data types that are supported by the DBMSs associated with DB2 ODBC. The information is returned in an SQL result set. The columns can be received using the same functions that are used to process a query.

### Syntax

```
SQLRETURN SQLGetTypeInfo (SQLHSTMT hstmt,
                          SQLSMALLINT fSqlType);
```

### Function arguments

Table 85. SQLGetTypeInfo arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle.
SQLSMALLINT	<i>fSqlType</i>	input	<p>The SQL data type being queried. The supported types are:</p> <ul style="list-style-type: none"> <li>• SQL_ALL_TYPES</li> <li>• SQL_BINARY</li> <li>• SQL_BLOB</li> <li>• SQL_CHAR</li> <li>• SQL_CLOB</li> <li>• SQL_DATE</li> <li>• SQL_DBCLOB</li> <li>• SQL_DECIMAL</li> <li>• SQL_DOUBLE</li> <li>• SQL_FLOAT</li> <li>• SQL_GRAPHIC</li> <li>• SQL_INTEGER</li> <li>• SQL_LONGVARBINARY</li> <li>• SQL_LONGVARCHAR</li> <li>• SQL_LONGVARGRAPHIC</li> <li>• SQL_NUMERIC</li> <li>• SQL_REAL</li> <li>• SQL_ROWID</li> <li>• SQL_SMALLINT</li> <li>• SQL_TIME</li> <li>• SQL_TIMESTAMP</li> <li>• SQL_VARBINARY</li> <li>• SQL_VARCHAR</li> <li>• SQL_VARGRAPHIC</li> </ul> <p>If SQL_ALL_TYPES is specified, information about all supported data types is returned in ascending order by TYPE_NAME. All unsupported data types are absent from the result set.</p>

## Usage

Since `SQLGetTypeInfo()` generates a result set and is equivalent to executing a query, it generates a cursor and begins a transaction. To prepare and execute another statement on this statement handle, the cursor must be closed.

If `SQLGetTypeInfo()` is called with an invalid *fSql/Type*, an empty result set is returned.

The columns of the result set generated by this function are described below.

Although new columns might be added and the names of the existing columns changed in future releases, the position of the current columns does not change. The data types returned are those that can be used in a `CREATE TABLE`, `ALTER TABLE`, `DDL` statement. Non-persistent data types such as the locator data types are not part of the returned result set. User defined data types are not returned either.

Table 86 (Page 1 of 2). Columns returned by `SQLGetTypeInfo`

Column number/name	Data type	Description
1 TYPE_NAME	VARCHAR(128) NOT NULL	Character representation of the SQL data type name. For example, VARCHAR, BLOB, DATE, INTEGER.
2 DATA_TYPE	SMALLINT NOT NULL	SQL data type define values. For example, SQL_VARCHAR, SQL_BLOB, SQL_DATE, SQL_INTEGER.
3 COLUMN_SIZE	INTEGER	If the data type is a character or binary string, then this column contains the maximum length in bytes; if it is a graphic (DBCS) string, this is the number of double byte characters for the column.  For date, time, timestamp data types, this is the total number of characters required to display the value when converted to character.  For numeric data types, this is the total number of digits.
4 LITERAL_PREFIX	VARCHAR(128)	Character that DB2 recognizes as a prefix for a literal of this data type. This column is null for data types where a literal prefix is not applicable.
5 LITERAL_SUFFIX	VARCHAR(128)	Character that DB2 recognizes as a suffix for a literal of this data type. This column is null for data types where a literal prefix is not applicable.
6 CREATE_PARAMS	VARCHAR(128)	The text of this column contains a list of keywords, separated by commas, that correspond to each parameter the application can specify in parenthesis when using the name in the TYPE_NAME column as a data type in SQL. The keywords in the list can be any of the following: LENGTH, PRECISION, SCALE. They appear in the order that the SQL syntax requires that they be used.  A NULL indicator is returned if there are no parameters for the data type definition, (such as INTEGER).  <b>Note:</b> The intent of CREATE_PARAMS is to enable an application to customize the interface for a <i>DDL builder</i> . An application should expect, using this, only to be able to determine the number of arguments required to define the data type and to have localized text that could be used to label an edit control.

## SQLGetTypeInfo

Table 86 (Page 2 of 2). Columns returned by SQLGetTypeInfo

Column number/name	Data type	Description
7 NULLABLE	SMALLINT NOT NULL	Indicates whether the data type accepts a NULL value <ul style="list-style-type: none"> <li>• Set to SQL_NO_NULLS if NULL values are disallowed.</li> <li>• Set to SQL_NULLABLE if NULL values are allowed.</li> </ul>
8 CASE_SENSITIVE	SMALLINT NOT NULL	Indicates whether the data type can be treated as case sensitive for collation purposes; valid values are SQL_TRUE and SQL_FALSE.
9 SEARCHABLE	SMALLINT NOT NULL	Indicates how the data type is used in a WHERE clause. Valid values are: <ul style="list-style-type: none"> <li>• SQL_UNSEARCHABLE : if the data type cannot be used in a WHERE clause.</li> <li>• SQL_LIKE_ONLY : if the data type can be used in a WHERE clause only with the LIKE predicate.</li> <li>• SQL_ALL_EXCEPT_LIKE : if the data type can be used in a WHERE clause with all comparison operators except LIKE.</li> <li>• SQL_SEARCHABLE : if the data type can be used in a WHERE clause with any comparison operator.</li> </ul>
10 UNSIGNED_ATTRIBUTE	SMALLINT	Indicates whether the data type is unsigned. The valid values are: SQL_TRUE, SQL_FALSE or NULL. A NULL indicator is returned if this attribute is not applicable to the data type.
11 FIXED_PREC_SCALE	SMALLINT NOT NULL	Contains the value SQL_TRUE if the data type is exact numeric and always has the same precision and scale; otherwise, it contains SQL_FALSE.
12 AUTO_INCREMENT	SMALLINT	Contains SQL_TRUE if a column of this data type is automatically set to a unique value when a row is inserted; otherwise, contains SQL_FALSE.
13 LOCAL_TYPE_NAME	VARCHAR(128)	This column contains any localized (native language) name for the data type that is different from the regular name of the data type. If there is no localized name, this column is NULL.  This column is intended for display only. The character set of the string is locale-dependent and is typically the default character set of the database.
14 MINIMUM_SCALE	SMALLINT	The minimum scale of the SQL data type. If a data type has a fixed scale, the MINIMUM_SCALE and MAXIMUM_SCALE columns both contain the same value. NULL is returned where scale is not applicable.
15 MAXIMUM_SCALE	SMALLINT	The maximum scale of the SQL data type. NULL is returned where scale is not applicable. If the maximum scale is not defined separately in the DBMS, but is defined instead to be the same as the maximum length of the column, then this column contains the same value as the COLUMN_SIZE column.

## Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 87. SQLGetTypeInfo SQLSTATEs

SQLSTATE	Description	Explanation
24000	Invalid cursor state.	A cursor is already opened on the statement handle. <i>hstmt</i> is not closed.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1004	SQL data type out of range.	An invalid <i>fSqlType</i> is specified.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.

## Restrictions

The following ODBC specified SQL data types (and their corresponding *fSqlType* define values) are not supported by any IBM RDBMS:

<b>Data type</b>	<b>fSqlType</b>
TINY INT	SQL_TINYINT
BIG INT	SQL_BIGINT
BIT	SQL_BIT

## Example

```

/*****
/* DB2 for OS/390 example:
/*     Invokes SQLGetTypeInfo to retrieve SQL data types sup-
/*     ported.
*****/

#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <sqlca.h>
#include "sqlcli1.h"

/*****
/* Invoke SQLGetTypeInfo to retrieve all SQL data types supported
/* by Data Source.
*****/

int main( )
{
    SQLHENV      hEnv   = SQL_NULL_HENV;
    SQLHDBC      hDbc   = SQL_NULL_HDBC;
    SQLHSTMT     hStmt  = SQL_NULL_HSTMT;

```

## SQLGetTypeInfo

```
SQLRETURN      rc      = SQL_SUCCESS;
SQLINTEGER     RETCODE = 0;

(void) printf ("**** Entering CLIP06.\n\n");

/*****
/* Allocate Environment Handle */
*****/

RETCODE = SQLAllocEnv(&hEnv);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Allocate Connection Handle to DSN */
*****/

RETCODE = SQLAllocConnect(hEnv,
                          &hDbc);

if( RETCODE != SQL_SUCCESS )    // Could not get a Connect Handle
    goto dberror;

/*****
/* CONNECT TO data source (STLEC1) */
*****/

RETCODE = SQLConnect(hDbc,      // Connect handle
                    (SQLCHAR *) "STLEC1", // DSN
                    SQL_NTS,    // DSN is null-terminated
                    NULL,       // Null UID
                    0,          // Null Auth string
                    NULL,       // Null Auth string
                    0);

if( RETCODE != SQL_SUCCESS )    // Connect failed
    goto dberror;

/*****
/* Retrieve SQL data types from DSN */
*****/
// local variables to Bind to retrieve TYPE_NAME, DATA_TYPE,
// COLUMN_SIZE and NULLABLE

struct                // TYPE_NAME is VARCHAR(128)
{
    SQLSMALLINT length;
    SQLCHAR     name [128];
    SQLINTEGER  ind;
} typename;

SQLSMALLINT data_type;    // DATA_TYPE is SMALLINT
SQLINTEGER  data_type_ind;
SQLINTEGER  column_size; // COLUMN_SIZE is integer
SQLINTEGER  column_size_ind;
SQLSMALLINT nullable;    // NULLABLE is SMALLINT
SQLINTEGER  nullable_ind;
```



```

/*****
/* Allocate Statement Handle */
/*****

rc = SQLAllocStmt (hDbc,
                  &hStmt);

if (rc != SQL_SUCCESS)
    goto exit;

/*****
/*
/* Retrieve native SQL types from DSN -----> */
/*
/* The result set consists of 15 columns. We only bind */
/* TYPE_NAME, DATA_TYPE, COLUMN_SIZE and NULLABLE. Note: Need */
/* not bind all columns of result set -- only those required. */
/*
/*****

rc = SQLGetTypeInfo (hStmt,
                   SQL_ALL_TYPES);

if (rc != SQL_SUCCESS)
    goto exit;

rc = SQLBindCol (hStmt,          // bind TYPE_NAME
                1,
                SQL_CHAR,
                (SQLPOINTER) typename.name,
                128,
                &typename.ind);

if (rc != SQL_SUCCESS)
    goto exit;

rc = SQLBindCol (hStmt,          // bind DATA_NAME
                2,
                SQL_C_DEFAULT,
                (SQLPOINTER) &data_type,
                sizeof(data_type),
                &data_type_ind);

if (rc != SQL_SUCCESS)
    goto exit;

rc = SQLBindCol (hStmt,          // bind COLUMN_SIZE
                3,
                SQL_C_DEFAULT,
                (SQLPOINTER) &column_size,
                sizeof(column_size),
                &column_size_ind);

if (rc != SQL_SUCCESS)
    goto exit;

rc = SQLBindCol (hStmt,          // bind NULLABLE
                7,
                SQL_C_DEFAULT,
                (SQLPOINTER) &nullable,
                sizeof(nullable),
                &nullable_ind);

```

## SQLGetTypeInfo

```
if (rc != SQL_SUCCESS)
    goto exit;

/*****
/* Fetch all native DSN SQL Types and print Type Name, Type,
/* Precision and nullability.
*****/

while ((rc = SQLFetch (hStmt)) == SQL_SUCCESS)
{
    (void) printf ("**** Type Name is %s. Type is %d. Precision is %d.",
                  typename.name,
                  data_type,
                  column_size);
    if (nullable == SQL_NULLABLE)
        (void) printf (" Type is nullable.\n");
    else
        (void) printf (" Type is not nullable.\n");
}

if (rc == SQL_NO_DATA_FOUND) // if result set exhausted reset
    rc = SQL_SUCCESS;       // rc to OK

/*****
/* Free Statement handle.
*****/

rc = SQLFreeStmt (hStmt,
                  SQL_DROP);

if (RETCODE != SQL_SUCCESS) // An advertised API failed
    goto dberror;

/*****
/* DISCONNECT from data source
*****/

RETCODE = SQLDisconnect(hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Deallocate Connection Handle
*****/

RETCODE = SQLFreeConnect (hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Free Environment Handle
*****/

RETCODE = SQLFreeEnv (hEnv);

if (RETCODE == SQL_SUCCESS)
    goto exit;
```

```
dberror:  
RETCODE=12;  
  
exit:  
  
(void) printf ("**** Exiting CLIP06.\n\n");  
  
return(RETCODE);  
}
```

## References

- “SQLColAttributes - Get column attributes” on page 106
- “SQLExtendedFetch - Extended fetch (fetch array of rows)” on page 162
- “SQLGetInfo - Get general information” on page 219

---

## SQLMoreResults - Determine if there are more result sets

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>		
-----------------------	-----------------	--	--

SQLMoreResults() determines whether there is more information available on the statement handle which has been associated with:

- Array input of parameter values for a query, or
- A stored procedure that is returning result sets.

### Syntax

```
SQLRETURN SQLMoreResults (SQLHSTMT hstmt);
```

### Function arguments

Table 88. SQLMoreResults arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle.

### Usage

This function is used to return multiple results set in a sequential manner upon the execution of:

- A parameterized query with an array of input parameter values specified with SQLParamOptions() and SQLBindParameter(), or
- A stored procedure containing SQL queries, the cursors of which have been left open so that the result sets remain accessible when the stored procedure has finished execution.

See “Using arrays to input parameter values” on page 375 and “Returning result sets from stored procedures” on page 392 for more information.

After completely processing the first result set, the application can call SQLMoreResults() to determine if another result set is available. If the current result set has unfetched rows, SQLMoreResults() discards them by closing the cursor and, if another result set is available, returns SQL\_SUCCESS.

If all the result sets have been processed, SQLMoreResults() returns SQL\_NO\_DATA\_FOUND.

If SQLFreeStmt() is called with the SQL\_CLOSE or SQL\_DROP option, all pending result sets on this statement handle are discarded.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE
- SQL\_NO\_DATA\_FOUND

## Diagnostics

Table 89. SQLMoreResults SQLSTATES

SQLSTATE	Description	Explanation
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.

In addition SQLMoreResults() can return the SQLSTATES associated with SQLExecute().

## Restrictions

The ODBC specification of SQLMoreResults() also allows counts associated with the execution of parameterized INSERT, UPDATE, and DELETE statements with arrays of input parameter values to be returned. However, DB2 ODBC does not support the return of such count information.

## Example

```

/* ... */
#define NUM_CUSTOMERS 25
SQLCHAR      stmt[] =
{ "WITH " /* Common Table expression (or Define Inline View) */
  "order (ord_num, cust_num, prod_num, quantity, amount) AS "
  "( "
  "SELECT c.ord_num, c.cust_num, l.prod_num, l.quantity, "
  "       price(char(p.price, '.'), p.units, char(l.quantity, '.')) "
  "FROM ord_cust c, ord_line l, product p "
  "WHERE c.ord_num = l.ord_num AND l.prod_num = p.prod_num "
  "AND cust_num = CNUM(cast (? as integer)) "
  ")", "
  "totals (ord_num, total) AS "
  "( "
  "SELECT ord_num, sum(decimal(amount, 10, 2)) "
  "FROM order GROUP BY ord_num "
  ")", "

```

## SQLMoreResults

```
/* The 'actual' SELECT from the inline view */
"SELECT order.ord_num, cust_num, prod_num, quantity, "
      "DECIMAL(amount,10,2) amount, total "
"FROM order, totals "
"WHERE order.ord_num = totals.ord_num "
};
/* Array of customers to get list of all orders for */
SQLINTEGER    Cust[]=
{
    10, 20, 30, 40, 50, 60, 70, 80, 90, 100,
    110, 120, 130, 140, 150, 160, 170, 180, 190, 200,
    210, 220, 230, 240, 250
};

#define NUM_CUSTOMERS sizeof(Cust)/sizeof(SQLINTEGER)

/* Row-Wise (Includes buffer for both column data and length) */
struct {
    SQLINTEGER    Ord_Num_L;
    SQLINTEGER    Ord_Num;
    SQLINTEGER    Cust_Num_L;
    SQLINTEGER    Cust_Num;
    SQLINTEGER    Prod_Num_L;
    SQLINTEGER    Prod_Num;
    SQLINTEGER    Quant_L;
    SQLDOUBLE     Quant;
    SQLINTEGER    Amount_L;
    SQLDOUBLE     Amount;
    SQLINTEGER    Total_L;
    SQLDOUBLE     Total;
} Ord[ROWSET_SIZE];

SQLINTEGER    pirow = 0;
SQLINTEGER    pcrow;
SQLINTEGER    i;
SQLINTEGER    j;
/* ... */
/* Get details and total for each order Row-Wise */
rc = SQLAllocStmt(hdbc, &hstmt);

rc = SQLParamOptions(hstmt, NUM_CUSTOMERS, &pirow);

rc = SQLBindParameter(hstmt, 1, SQL_PARAM_INPUT, SQL_C_LONG, SQL_INTEGER,
                     0, 0, Cust, 0, NULL);

rc = SQLExecDirect(hstmt, stmt, SQL_NTS);

/* SQL_ROWSET_SIZE sets the max number of result rows to fetch each time */
rc = SQLSetStmtOption(hstmt, SQL_ROWSET_SIZE, ROWSET_SIZE);

/* Set Size of One row, Used for Row-Wise Binding Only */
rc = SQLSetStmtOption(hstmt, SQL_BIND_TYPE, sizeof(Ord) / ROWSET_SIZE);

/* Bind column 1 to the Ord_num Field of the first row in the array*/
rc = SQLBindCol(hstmt, 1, SQL_C_LONG, (SQLPOINTER) &Ord[0].Ord_Num, 0,
                &Ord[0].Ord_Num_L);

/* Bind remaining columns ... */
/* ... */
```

```

/* NOTE: This sample assumes that an order never has more
        rows than ROWSET_SIZE. A check should be added below to call
        SQLExtendedFetch multiple times for each result set.
*/
do /* for each result set .... */
{ rc = SQLExtendedFetch(hstmt, SQL_FETCH_NEXT, 0, &pcrow, NULL);

  if (pcrow > 0) /* if 1 or more rows in the result set */
  {
    i = j = 0;

    printf("*****\n");
    printf("Orders for Customer: %ld\n", Ord[0].Cust_Num);
    printf("*****\n");

    while (i < pcrow)
    { printf("\nOrder #: %ld\n", Ord[i].Ord_Num);
      printf("      Product  Quantity      Price\n");
      printf("      -----  -----  ----- \n");
      j = i;
      while (Ord[j].Ord_Num == Ord[i].Ord_Num)
      { printf("      %8ld %16.7lf %12.2lf\n",
                Ord[i].Prod_Num, Ord[i].Quant, Ord[i].Amount);
        i++;
      }
      printf("                                =====\n");
      printf("                                %12.2lf\n", Ord[j].Total);
    } /* end while */
  } /* end if */
}
while ( SQLMoreResults(hstmt) == SQL_SUCCESS);
/* ... */

```

## References

- “SQLParamOptions - Specify an input array for a parameter” on page 275

## SQLNativeSql - Get native SQL text

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>		
-----------------------	-----------------	--	--

SQLNativeSql() is used to show how DB2 ODBC interprets vendor escape clauses. If the original SQL string passed in by the application contains vendor escape clause sequences, then DB2 ODBC returns the transformed SQL string that the data source sees (with vendor escape clauses either converted or discarded, as appropriate).

### Syntax

```
SQLRETURN SQLNativeSql (SQLHDBC hdbc,
                        SQLCHAR FAR *szSqlStrIn,
                        SQLINTEGER cbSqlStrIn,
                        SQLCHAR FAR *szSqlStr,
                        SQLINTEGER cbSqlStrMax,
                        SQLINTEGER FAR *pcbSqlStr);
```

### Function arguments

Table 90. SQLNativeSQL arguments

Data type	Argument	Use	Description
SQLHDBC	hdbc	input	Connection handle.
SQLCHAR *	szSqlStrIn	input	Input SQL string.
SQLINTEGER	cbSqlStrIn	input	Length of <i>szSqlStrIn</i> .
SQLCHAR *	szSqlStr	output	Pointer to buffer for the transformed output string.
SQLINTEGER	cbSqlStrMax	input	Size of buffer pointed by <i>szSqlStr</i> .
SQLINTEGER *	pcbSqlStr	output	The total number of bytes (excluding the null-terminator) available to return in <i>szSqlStr</i> . If the number of bytes available to return is greater than or equal to <i>cbSqlStrMax</i> , the output SQL string in <i>szSqlStr</i> is truncated to <i>cbSqlStrMax - 1</i> bytes.

### Usage

This function is called when the application wishes to examine or display the transformed SQL string that is passed to the data source by DB2 ODBC. Translation (mapping) only occurs if the input SQL statement string contains vendor escape clause sequences. For more information on vendor escape clause sequences, see “Using vendor escape clauses” on page 404.

DB2 ODBC can only detect vendor escape clause syntax errors; since DB2 ODBC does not pass the transformed SQL string to the data source for preparation, syntax errors that are detected by the DBMS are not generated at this time. (The statement is not passed to the data source for preparation because the preparation can potentially cause the initiation of a transaction.)



## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 91. SQLNativeSQL SQLSTATEs

SQLSTATE	Description	Explanation
01004	Data truncated.	The buffer <i>szSqlStr</i> is not large enough to contain the entire SQL string, so truncation occurs. The argument <i>pcbSqlStr</i> contains the total length of the untruncated SQL string. (Function returns with SQL_SUCCESS_WITH_INFO)
08003	Connection is closed.	The <i>hdbc</i> does not reference an open database connection.
37000	Invalid SQL syntax.	The input SQL string in <i>szSqlStrIn</i> contained a syntax error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	The argument <i>szSqlStrIn</i> is a NULL pointer. The argument <i>szSqlStr</i> is a NULL pointer.
S1090	Invalid string or buffer length.	The argument <i>cbSqlStrIn</i> was less than 0, but not equal to SQL_NTS. The argument <i>cbSqlStrMax</i> was less than 0.

## Restrictions

None.

### Example

```
/* ... */
SQLCHAR      in_stmt[1024];
SQLCHAR      out_stmt[1024];
SQLSMALLINT  pcPar;
SQLINTEGER   indicator;
/* ... */
/* Prompt for a statement to prepare */
printf("Enter an SQL statement: \n");
gets(in_stmt);

/* prepare the statement */
rc = SQLPrepare(hstmt, in_stmt, SQL_NTS);

SQLNumParams(hstmt, &pcPar);

SQLNativeSql(hstmt, in_stmt, SQL_NTS, out_stmt, 1024, &indicator);

if (indicator == SQL_NULL_DATA)
{ printf("Invalid statement\n"); }
else
{ printf(" Input Statement: \n %s \n", in_stmt);
  printf("Output Statement: \n %s \n", out_stmt);
  printf("Number of Parameter Markers = %ld\n", pcPar);
}
rc = SQLFreeStmt(hstmt, SQL_DROP);
/* ... */
```

### References

- “Using vendor escape clauses” on page 404

## SQLNumParams - Get number of parameters in a SQL statement

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>		
-----------------------	-----------------	--	--

SQLNumParams() returns the number of parameter markers in a SQL statement.

### Syntax

```
SQLRETURN SQLNumParams (SQLHSTMT hstmt,
                        SQLSMALLINT FAR *pcpar);
```

### Function arguments

Table 92. SQLNumParams arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	Input	Statement handle.
SQLSMALLINT *	pcpar	Output	Number of parameters in the statement.

### Usage

This function can only be called after the statement associated with *hstmt* has been prepared. If the statement does not contain any parameter markers, *pcpar* is set to 0.

An application can call this function to determine how many SQLBindParameter() calls are necessary for the SQL statement associated with the statement handle.

### Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

### Diagnostics

Table 93 (Page 1 of 2). SQLNumParams SQLSTATEs

SQLSTATE	Description	Explanation
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	<i>pcpar</i> is null.
S1010	Function sequence error.	This function is called before SQLPrepare() is called for the specified <i>hstmt</i> .  The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.

## SQLNumParams

Table 93 (Page 2 of 2). SQLNumParams SQLSTATES

SQLSTATE	Description	Explanation
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.

### Restrictions

None.

### Example

See "Example" on page 268.

### References

- "SQLBindParameter - Binds a parameter marker to a buffer or LOB locator" on page 93
- "SQLPrepare - Prepare a statement" on page 277

## SQLNumResultCols - Get number of result columns

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLNumResultCols() returns the number of columns in the result set associated with the input statement handle.

SQLPrepare() or SQLExecDirect() must be called before calling this function.

After calling this function, you can call SQLColAttributes(), or one of the bind column functions.

### Syntax

```
SQLRETURN SQLNumResultCols (SQLHSTMT hstmt,
                             SQLSMALLINT FAR *pccol);
```

### Function arguments

Table 94. SQLNumResultCols arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle
SQLSMALLINT *	<i>pccol</i>	output	Number of columns in the result set

### Usage

The function sets the output argument to zero if the last statement or function executed on the input statement handle did not generate a result set.

### Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

### Diagnostics

Table 95 (Page 1 of 2). SQLNumResultCols SQLSTATEs

SQLSTATE	Description	Explanation
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	<i>pccol</i> is a null pointer.

## SQLNumResultCols

Table 95 (Page 2 of 2). SQLNumResultCols SQLSTATES

SQLSTATE	Description	Explanation
S1010	Function sequence error.	The function is called prior to calling SQLPrepare() or SQLExecDirect() for the <i>hstmt</i> .  The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.

## Restrictions

None.

## Example

See "Example" on page 135

## References

- "SQLColAttributes - Get column attributes" on page 106
- "SQLDescribeCol - Describe column attributes" on page 132
- "SQLExecDirect - Execute a statement directly" on page 154
- "SQLGetData - Get data from a column" on page 199
- "SQLPrepare - Prepare a statement" on page 277

## SQLParamData - Get next parameter for which a data value is needed

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLParamData() is used in conjunction with SQLPutData() to send long data in pieces. It can also be used to send fixed length data as well. For a description of the exact sequence of this input method, see “Sending/retrieving long data in pieces” on page 373.

### Syntax

```
SQLRETURN SQLParamData (SQLHSTMT hstmt,
                        SQLPOINTER FAR *prgbValue);
```

### Function arguments

Table 96. SQLParamData arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle.
SQLPOINTER *	prgbValue	output	Pointer to the value of the <i>prgbValue</i> argument specified on the SQLBindParameter() or SQLSetParam() call.

### Usage

SQLParamData() returns SQL\_NEED\_DATA if there is at least one SQL\_DATA\_AT\_EXEC parameter for which data is not assigned. This function returns an application provided value in *prgbValue* supplied by the application during the previous SQLBindParameter() call. SQLPutData() is called one or more times (in the case of long data) to send the parameter data. SQLParamData() is called to signal that all the data has been sent for the current parameter and to advance to the next SQL\_DATA\_AT\_EXEC parameter. SQL\_SUCCESS is returned when all the parameters have been assigned data values and the associated statement has been executed successfully. If any errors occur during or before actual statement execution, SQL\_ERROR is returned.

If SQLParamData() returns SQL\_NEED\_DATA, then only SQLPutData() or SQLCancel() calls can be made. All other function calls using this statement handle fails. In addition, all function calls referencing the parent *hdbc* of *hstmt* fail if they involve changing any attribute or state of that connection; that is, the following function calls on the parent *hdbc* are also not permitted:

- SQLAllocConnect()
- SQLAllocStmt()
- SQLSetConnectOption()
- SQLNativeSql()
- SQLTransact()

Should they be invoked during an SQL\_NEED\_DATA sequence, these functions return SQL\_ERROR with SQLSTATE of S1010 and the processing of the SQL\_DATA\_AT\_EXEC parameters is not affected.

### Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE
- SQL\_NEED\_DATA

### Diagnostics

SQLParamData() can return any SQLSTATE returned by the SQLExecDirect() and SQLExecute() functions. In addition, the following diagnostics can also be generated:

Table 97. SQLParamData SQLSTATES

SQLSTATE	Description	Explanation
40001	Transaction rollback.	The transaction to which this SQL statement belonged is rolled back due to a deadlock or timeout.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	SQLParamData() is called out of sequence. This call is only valid after an SQLExecDirect() or an SQLExecute(), or after an SQLPutData() call.  Even though this function is called after an SQLExecDirect() or an SQLExecute() call, there are no SQL_DATA_AT_EXEC parameters (left) to process.

### Restrictions

None.

### Example

See “Example” on page 306.

### References

- “SQLBindParameter - Binds a parameter marker to a buffer or LOB locator” on page 93
- “SQLCancel - Cancel statement” on page 104
- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLPutData - Passing data value for a parameter” on page 304
- “SQLSetParam - Binds a parameter marker to a buffer” on page 328



## SQLParamOptions - Specify an input array for a parameter

### Purpose

Specification:	ODBC 1.0		
----------------	----------	--	--

SQLParamOptions() provides the ability to set multiple values for each parameter set by SQLBindParameter(). This allows the application to perform batched processing of the same SQL statement with one set of prepare, execute and SQLBindParameter() calls.

### Syntax

```
SQLRETURN SQLParamOptions (SQLHSTMT hstmt,
                            SQLINTEGER crow,
                            SQLINTEGER FAR *pirow);
```

### Function arguments

Table 98. SQLParamOptions arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	Input	Statement handle.
SQLINTEGER	crow	Input	Number of values for each parameter. If this is greater than 1, then the <i>rgbValue</i> argument in SQLBindParameter() points to an array of parameter values, and <i>pcbValue</i> points to an array of lengths.
SQLINTEGER *	pirow	Output (deferred)	Pointer to the buffer for the current parameter array index. As each set of parameter values is processed, <i>pirow</i> is set to the array index of that set. If a statement fails, <i>pirow</i> can be used to determine how many statements were successfully processed. Nothing is returned if the <i>pirow</i> pointer is NULL.

### Usage

DB2 ODBC prepares the statement, and executes it repeatedly for the array of parameter markers.

As a statement executes, *pirow* is set to the index of the current array of parameter values. If an error occurs during execution for a particular element in the array, execution halts and SQLExecute(), SQLExecDirect() or SQLParamData() returns SQL\_ERROR.

The contents of *pirow* have the following uses:

- When SQLParamData() returns SQL\_NEED\_DATA, the application can access the value in *pirow* to determine which set of parameters is being assigned values.
- When SQLExecute() or SQLExecDirect() returns an error, the application can access the value in *pirow* to find out which element in the parameter value array failed.

## SQLParamOptions

- When `SQLExecute()`, `SQLExecDirect()`, `SQLParamData()`, or `SQLPutData()` succeeds, the value in *pirow* is set to the input value in *crow* to indicate that all elements of the array have been processed successfully.

The output argument *pirow* indicates how many sets of parameters were successfully processed. If the statement processed is a query, *pirow* indicates the array index associated with the current result set returned by `SQLMoreResults()` and is incremented each time `SQLMoreResults()` is called.

## Return codes

- `SQL_SUCCESS`
- `SQL_SUCCESS_WITH_INFO`
- `SQL_ERROR`
- `SQL_INVALID_HANDLE`

## Diagnostics

Table 99. *SQLParamOptions* SQLSTATES

SQLSTATE	Description	Explanation
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute ( <code>SQLParamData()</code> , <code>SQLPutData()</code> ) operation.
S1107	Row value out of range.	The value in the argument <i>crow</i> is less than 1.

## Restrictions

None.

## Example

See “Array input example” on page 378.

## References

- “`SQLBindParameter` - Binds a parameter marker to a buffer or LOB locator” on page 93
- “`SQLMoreResults` - Determine if there are more result sets” on page 262
- “`SQLSetStmtOption` - Set statement option” on page 334

## SQLPrepare - Prepare a statement

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLPrepare() associates an SQL statement with the input statement handle and sends the statement to the DBMS to be prepared. The application can reference this prepared statement by passing the statement handle to other functions.

If the statement handle has been previously used with a query statement (or any function that returns a result set), SQLFreeStmt() must be called to close the cursor, before calling SQLPrepare().

### Syntax

```
SQLRETURN SQLPrepare(
    (SQLHSTMT
     SQLCHAR FAR *szSqlStr,
     SQLINTEGER
     hstmt,
     *szSqlStr,
     cbSqlStr);
```

### Function arguments

Table 100. SQLPrepare arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle. There must not be an open cursor associated with <i>hstmt</i> .
SQLCHAR *	<i>szSqlStr</i>	input	SQL statement string.
SQLINTEGER	<i>cbSqlStr</i>	input	Length of contents of <i>szSqlStr</i> argument.  This must be set to either the exact length of the SQL statement in <i>szSqlStr</i> , or to SQL_NTS if the statement text is null-terminated.

### Usage

If the SQL statement text contains vendor escape clause sequences, DB2 ODBC first modifies the SQL statement text to the appropriate DB2 specific format before submitting it to the database for preparation. If the application does not generate SQL statements that contain vendor escape clause sequences (see “Using vendor escape clauses” on page 404); then the SQL\_NOSCAN statement option should be set to SQL\_NOSCAN\_ON at the statement level so that DB2 ODBC does not perform a scan for any vendor escape clauses.

When a statement is prepared using SQLPrepare(), the application can request information about the format of the result set (if the statement was a query) by calling:

- SQLNumResultCols()
- SQLDescribeCol()
- SQLColAttributes()

The SQL statement string can contain parameter markers and SQLNumParams() can be called to determine the number of parameter markers in the statement. A

## SQLPrepare

parameter marker is represented by a "?" character that indicates a position in the statement where an application supplied value is to be substituted when `SQLExecute()` is called. The bind parameter functions, `SQLBindParameter()` and `SQLSetParam()` are used to bind (associate) application values with each parameter marker and to indicate if any data conversion should be performed at the time the data is transferred.

All parameters must be bound before calling `SQLExecute()`. For more information see "SQLExecute - Execute a statement" on page 159.

After the application processes the results from the `SQLExecute()` call, it can execute the statement again with new (or the same) parameter values.

The SQL statement cannot be a COMMIT or ROLLBACK. `SQLTransact()` must be called to issue COMMIT or ROLLBACK. For more information about SQL statements, that DB2 for OS/390 supports, see Table 1 on page 16.

If the SQL statement is a positioned DELETE or a positioned UPDATE, the cursor referenced by the statement must be defined on a separate statement handle under the same connection handle and same isolation level.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 101 (Page 1 of 2). SQLPrepare SQLSTATEs

SQLSTATE	Description	Explanation
01504	The UPDATE or DELETE statement does not include a WHERE clause.	<i>szSqlStr</i> contains an UPDATE or DELETE statement which did not contain a WHERE clause.
21S01	Insert value list does not match column list.	<i>szSqlStr</i> contains an INSERT statement and the number of values to be inserted did not match the degree of the derived table.
21S02	Degrees of derived table does not match column list.	<i>szSqlStr</i> contains a CREATE VIEW statement and the number of names specified is not the same degree as the derived table defined by the query specification.
24000	Invalid cursor state.	A cursor is already opened on the statement handle.
34000	Invalid cursor name.	<i>szSqlStr</i> contains a positioned DELETE or a positioned UPDATE and the cursor referenced by the statement being executed is not open.
37xxx a	Invalid SQL syntax.	<i>szSqlStr</i> contains one or more of the following: <ul style="list-style-type: none"><li>• A COMMIT</li><li>• A ROLLBACK</li><li>• An SQL statement that the connected database server cannot prepare</li><li>• A statement containing a syntax error</li></ul>
40001	Transaction rollback.	The transaction to which this SQL statement belongs is rolled back due to deadlock or timeout.

Table 101 (Page 2 of 2). SQLPrepare SQLSTATES

SQLSTATE	Description	Explanation
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
42xxx <sup>a</sup>	Syntax error or access rule violation	<b>425xx</b> indicates the authorization ID does not have permission to execute the SQL statement contained in <i>szSqlStr</i> .  Other <b>42xxx</b> SQLSTATES indicate a variety of syntax or access problems with the statement.
58004	Unexpected system failure.	Unrecoverable system error.
S0001	Database object already exists.	<i>szSqlStr</i> contains a CREATE TABLE or CREATE VIEW statement and the table name or view name specified already exists.
S0002	Database object does not exist.	<i>szSqlStr</i> contains an SQL statement that references a table name or a view name that does not exist.
S0011	Index already exists.	<i>szSqlStr</i> contains a CREATE INDEX statement and the specified index name already exists.
S0012	Index not found.	<i>szSqlStr</i> contains a DROP INDEX statement and the specified index name does not exist.
S0021	Column already exists.	<i>szSqlStr</i> contains an ALTER TABLE statement and the column specified in the ADD clause is not unique or identifies an existing column in the base table.
S0022	Column not found.	<i>szSqlStr</i> contains an SQL statement that references a column name that does not exist.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	<i>szSqlStr</i> is a null pointer.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1014	No more handles.	DB2 ODBC is not able to allocate a handle due to internal resources.
S1090	Invalid string or buffer length.	The argument <i>cbSqlStr</i> is less than 1, but not equal to SQL_NTS.
<b>Note:</b>		
<sup>a</sup>	xxx refers to any SQLSTATE with that class code. Example, <b>37xxx</b> refers to any SQLSTATE in the <b>37</b> class.	

Not all DBMSs report all of the above diagnostic messages at prepare time. Therefore, an application must also be able to handle these conditions when calling SQLExecute().

## Restrictions

None.

## Example

```

/*****
/* DB2 for OS/390 example: */
/*     Prepares a query and executes that query twice */
/*     specifying a unique value for the parameter marker. */
*****/

#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <sqlca.h>
#include "sqlcli1.h"

int main( )
{
    SQLHENV      hEnv   = SQL_NULL_HENV;
    SQLHDBC      hDbc   = SQL_NULL_HDBC;
    SQLHSTMT     hStmt  = SQL_NULL_HSTMT;
    SQLRETURN    rc     = SQL_SUCCESS;
    SQLINTEGER   RETCODE = 0;
    char         *pDSN = "STLEC1";
    SWORD        cbCursor;
    SDWORD       cbValue1;
    SDWORD       cbValue2;
    char         employee [30];
    int          salary = 0;
    int          param_salary = 30000;

    char         *stmt = "SELECT NAME, SALARY FROM EMPLOYEE WHERE SALARY > ?";

    (void) printf ("**** Entering CLIP07.\n\n");

    /*****
    /* Allocate Environment Handle */
    *****/

    RETCODE = SQLAllocEnv(&hEnv);

    if (RETCODE != SQL_SUCCESS)
        goto dberror;

    /*****
    /* Allocate Connection Handle to DSN */
    *****/

    RETCODE = SQLAllocConnect(hEnv,
                              &hDbc);

    if( RETCODE != SQL_SUCCESS )    // Could not get a Connect Handle
        goto dberror;

```

```

/*****
/* CONNECT TO data source (STLEC1) */
*****/

    RETCODE = SQLConnect(hDbc,          // Connect handle
                        (SQLCHAR *) pDSN, // DSN
                        SQL_NTS,        // DSN is nul-terminated
                        NULL,           // Null UID
                        0,
                        NULL,          // Null Auth string
                        0);
    if( RETCODE != SQL_SUCCESS )      // Connect failed
        goto dberror;

/*****
/* Allocate Statement Handles */
*****/

    rc = SQLAllocStmt (hDbc,
                      &hStmt);

    if (rc != SQL_SUCCESS)
        goto exit;

/*****
/* Prepare the query for multiple execution within current */
/* transaction. Note that query is collapsed when transaction */
/* is committed or rolled back. */
*****/

    rc = SQLPrepare (hStmt,
                    (SQLCHAR *) stmt,
                    strlen(stmt));

    if (rc != SQL_SUCCESS)
    {
        (void) printf ("**** PREPARE OF QUERY FAILED.\n");
        goto dberror;
    }

    rc = SQLBindCol (hStmt,          // bind employee name
                    1,
                    SQL_C_CHAR,
                    employee,
                    sizeof(employee),
                    &cbValue1);

    if (rc != SQL_SUCCESS)
    {
        (void) printf ("**** BIND OF NAME FAILED.\n");
        goto dberror;
    }

    rc = SQLBindCol (hStmt,          // bind employee salary
                    2,
                    SQL_C_LONG,
                    &salary,
                    0,
                    &cbValue2);
    if (rc != SQL_SUCCESS)

```

```

{
    (void) printf ("**** BIND OF SALARY FAILED.\n");
    goto dberror;
}

/*****
/* Bind parameter to replace '?' in query. This has an initial */
/* value of 30000. */
*****/

rc = SQLBindParameter (hStmt,
                      1,
                      SQL_PARAM_INPUT,
                      SQL_C_LONG,
                      SQL_INTEGER,
                      0,
                      0,
                      &param_salary,
                      0,
                      NULL);

/*****
/* Execute prepared statement to generate answer set. */
*****/

rc = SQLExecute (hStmt);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** EXECUTE OF QUERY FAILED.\n");
    goto dberror;
}

/*****
/* Answer Set is available -- Fetch rows and print employees */
/* and salary. */
*****/

(void) printf ("**** Employees whose salary exceeds %d follow.\n\n",
              param_salary);

while ((rc = SQLFetch (hStmt)) == SQL_SUCCESS)
{
    (void) printf ("**** Employee Name %s with salary %d.\n",
                  employee,
                  salary);
}

/*****
/* Close query --- note that query is still prepared. Then change*/
/* bound parameter value to 100000. Then re-execute query. */
*****/

rc = SQLFreeStmt (hStmt,
                  SQL_CLOSE);

param_salary = 100000;

rc = SQLExecute (hStmt);
if (rc != SQL_SUCCESS)

```



```

{
    (void) printf ("**** EXECUTE OF QUERY FAILED.\n");
    goto dberror;
}

/*****
/* Answer Set is available -- Fetch rows and print employees    */
/* and salary.                                                */
*****/

(void) printf ("**** Employees whose salary exceeds %d follow.\n\n",
              param_salary);

while ((rc = SQLFetch (hStmt)) == SQL_SUCCESS)
{
    (void) printf ("**** Employee Name %s with salary %d.\n",
                  employee,
                  salary);
}

/*****
/* Deallocate Statement Handles -- statement is no longer in a  */
/* Prepared state.                                             */
*****/

rc = SQLFreeStmt (hStmt,
                  SQL_DROP);

/*****
/* DISCONNECT from data source                                */
*****/

RETCODE = SQLDisconnect(hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Deallocate Connection Handle                                */
*****/

RETCODE = SQLFreeConnect (hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

```



## SQLPrimaryKeys - Get primary key columns of a table

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>		
-----------------------	-----------------	--	--

SQLPrimaryKeys() returns a list of column names that comprise the primary key for a table. The information is returned in an SQL result set, which can be retrieved using the same functions that are used to process a result set generated by a query.

### Syntax

```
SQLRETURN SQLPrimaryKeys (SQLHSTMT hstmt,
                          SQLCHAR FAR *szCatalogName,
                          SQLSMALLINT cbCatalogName,
                          SQLCHAR FAR *szSchemaName,
                          SQLSMALLINT cbSchemaName,
                          SQLCHAR FAR *szTableName,
                          SQLSMALLINT cbTableName);
```

### Function arguments

Table 102. SQLPrimaryKeys arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle.
SQLCHAR *	szCatalogName	input	Catalog qualifier of a 3 part table name. This must be a NULL pointer or a zero length string.
SQLSMALLINT	cbCatalogName	input	Length of <i>szCatalogName</i> .
SQLCHAR *	szSchemaName	input	Schema qualifier of table name.
SQLSMALLINT	cbSchemaName	input	Length of <i>szSchemaName</i> .
SQLCHAR *	szTableName	input	Table name.
SQLSMALLINT	cbTableName	input	Length of <i>szTableName</i> .

### Usage

SQLPrimaryKeys() returns the primary key columns from a single table. Search patterns cannot be used to specify the schema qualifier or the table name.

The result set contains the columns listed in Table 103 on page 286, ordered by TABLE\_CAT, TABLE\_SCHEM, TABLE\_NAME and ORDINAL\_POSITION.

Since calls to SQLPrimaryKeys() in many cases map to a complex and, thus, expensive query against the system catalog, they should be used sparingly, and the results saved rather than repeating calls.

The VARCHAR columns of the catalog functions result set have been declared with a maximum length attribute of 128 to be consistent with SQL92 limits. Since DB2 names are less than 128, the application can choose to always set aside 128 characters (plus the null-terminator) for the output buffer, or alternatively, call

## SQLPrimaryKeys

SQLGetInfo() with the SQL\_MAX\_CATALOG\_NAME\_LEN, SQL\_MAX\_SCHEMA\_NAME\_LEN, SQL\_MAX\_TABLE\_NAME\_LEN, and SQL\_MAX\_COLUMN\_NAME\_LEN to determine respectively the actual lengths of the TABLE\_CAT, TABLE\_SCHEM, TABLE\_NAME, and COLUMN\_NAME columns supported by the connected DBMS.

Although new columns might be added and the names of the existing columns changed in future releases, the position of the current columns does not change.

Table 103. Columns returned by SQLPrimaryKeys

Column number/name	Data type	Description
1 TABLE_CAT	VARCHAR(128)	This is always null.
2 TABLE_SCHEM	VARCHAR(128)	The name of the schema containing TABLE_NAME.
3 TABLE_NAME	VARCHAR(128) NOT NULL	Name of the specified table.
4 COLUMN_NAME	VARCHAR(128) NOT NULL	Primary key column name.
5 KEY_SEQ	SMALLINT NOT NULL	Column sequence number in the primary key, starting with 1.
6 PK_NAME	VARCHAR(128)	Primary key identifier. NULL if not applicable to the data source.

**Note:** The column names used by DB2 ODBC follow the X/Open CLI CAE specification style. The column types, contents and order are identical to those defined for the SQLPrimaryKeys() result set in ODBC.

If the specified table does not contain a primary key, an empty result set is returned.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 104 (Page 1 of 2). SQLPrimaryKeys SQLSTATEs

SQLSTATE	Description	Explanation
24000	Invalid cursor state.	A cursor was already opened on the statement handle.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1014	No more handles.	DB2 ODBC is not able to allocate a handle due to internal resources.
S1090	Invalid string or buffer length.	The value of one of the name length arguments is less than 0, but not equal SQL_NTS.

Table 104 (Page 2 of 2). SQLPrimaryKeys SQLSTATES

SQLSTATE	Description	Explanation
S1C00	Driver not capable.	DB2 ODBC does not support <i>catalog</i> as a qualifier for table name.

## Restrictions

None.

## Example

The following example uses SQLPrimaryKeys to locate a primary key for a table, and calls SQLColAttributes to find its data type.

```

/* ... */
#include <sqlcli1.h>

void main()
{
    SQLCHAR    rgbDesc_20];
    SQLCHAR    szTableName_20];
    SQLCHAR    szSchemaName_20];
    SQLCHAR    rgbValue_20];
    SQLINTEGER pcbValue;
    SQLHENV    henv;
    SQLHDBC    hdbc;
    SQLHSTMT   hstmt;
    SQLSMALLINT pscDesc;
    SQLINTEGER pdDesc;
    SQLRETURN  rc;

    /*****
    /* Initialization...
    *****/

    if( SQLAllocEnv( &henv ) != SQL_SUCCESS )
    {
        fprintf( stdout, "Error in SQLAllocEnv\n" );
        exit(1);
    }
    if( SQLAllocConnect( henv, &hdbc ) != SQL_SUCCESS )
    {
        fprintf( stdout, "Error in SQLAllocConnect\n" );
        exit(1);
    }
    if( SQLConnect( hdbc,
                   NULL, SQL_NTS,
                   NULL, SQL_NTS,
                   NULL, SQL_NTS ) != SQL_SUCCESS )
    {
        fprintf( stdout, "Error in SQLConnect\n" );
        exit(1);
    }
    if( SQLAllocStmt( hdbc, &hstmt ) != SQL_SUCCESS )
    {
        fprintf( stdout, "Error in SQLAllocStmt\n" );
        exit(1);
    }
}

```

## SQLPrimaryKeys

```
/* Get primary key for table 'myTable' by using SQLPrimaryKeys */
rc = SQLPrimaryKeys( hstmt,
                    NULL, SQL_NTS,
                    (SQLCHAR*)szSchemaName, SQL_NTS,
                    (SQLCHAR*)szTableName, SQL_NTS );
if( rc != SQL_SUCCESS )
{
    goto exit;
}

/*
 * Since all we need is the ordinal position, we'll bind column 5 from
 * the result set.
 */
rc = SQLBindCol( hstmt,
                5,
                SQL_C_CHAR,
                (SQLPOINTER)rgbValue,
                20,
                &pcbValue );
if( rc != SQL_SUCCESS )
{
    goto exit;
}

/*
 * Fetch data...
 */
if( SQLFetch( hstmt ) != SQL_SUCCESS )
{
    goto exit;
}

/* Get data type for that column by calling SQLColAttributes(). */
rc = SQLColAttributes( hstmt,
                      pcbValue,
                      SQL_COLUMN_TYPE,
                      rgbDesc,
                      20,
                      &pcbDesc,
                      &pfDesc );
if( rc != SQL_SUCCESS )
{
    goto exit;
}

/*
 * Display the data type.
 */
fprintf( stdout, "Data type ==> %s\n", rgbDesc );
```

```
exit:
/*****
/* Clean up the environment...
*****/

SQLTransact( henv,
             hdbc,
             SQL_ROLLBACK );

SQLDisconnect( hdbc );

SQLFreeConnect( hdbc );

SQLFreeEnv( henv );
}
```

## References

- “SQLForeignKeys - Get the list of foreign key columns” on page 175
- “SQLStatistics - Get index and statistics information for a base table” on page 344

## SQLProcedureColumns - Get input/output parameter information for a procedure

### Purpose

Specification:	ODBC 1.0		
----------------	----------	--	--

SQLProcedureColumns() returns a list of input and output parameters associated with a procedure. The information is returned in an SQL result set, which can be retrieved using the same functions that are used to process a result set generated by a query.

### Syntax

```
SQLRETURN SQLProcedureColumns (
    SQLHSTMT hstmt,
    SQLCHAR FAR *szProcCatalog,
    SQLSMALLINT cbProcCatalog,
    SQLCHAR FAR *szProcSchema,
    SQLSMALLINT cbProcSchema,
    SQLCHAR FAR *szProcName,
    SQLSMALLINT cbProcName,
    SQLCHAR FAR *szColumnName,
    SQLSMALLINT cbColumnName);
```

### Function arguments

Table 105. SQLProcedureColumns arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle.
SQLCHAR *	szProcCatalog	input	Catalog qualifier of a 3 part procedure name. This must be a NULL pointer or a zero length string.
SQLSMALLINT	cbProcCatalog	input	Length of <i>szProcCatalog</i> . This must be set to 0.
SQLCHAR *	szProcSchema	input	Buffer that can contain a <i>pattern-value</i> to qualify the result set by schema name.  For Version 4 and Version 5 of DB2 for OS/390, all the stored procedures are in one schema; the only acceptable value for the <i>szProcSchema</i> argument is a null pointer. For DB2 UDB, <i>szProcSchema</i> can contain a valid pattern value. For more information about valid search patterns, see "Querying system catalog information" on page 368.
SQLSMALLINT	cbProcSchema	input	Length of <i>szProcSchema</i> .
SQLCHAR *	szProcName	input	Buffer that can contain a <i>pattern-value</i> to qualify the result set by procedure name.
SQLSMALLINT	cbProcName	input	Length of <i>szProcName</i> .
SQLCHAR *	szColumnName	input	Buffer that can contain a <i>pattern-value</i> to qualify the result set by parameter name. This argument is to be used to further qualify the result set already restricted by specifying a non-empty value for <i>szProcName</i> and/or <i>szProcSchema</i> .
SQLSMALLINT	cbColumnName	input	Length of <i>szColumnName</i> .



## Usage

If the stored procedure is at a DB2 for MVS/ESA Version 4 server or a DB2 for OS/390 Version 5 server, the name of the stored procedures must be registered in the server's SYSIBM.SYSPROCEDURES catalog table. If the stored procedure is at a DB2 for OS/390 Version 6 server, the name of the stored procedures must be registered in the server's SYSIBM.SYSROUTINES catalog table.

For versions of other DB2 servers that do not provide facilities for a stored procedure catalog, an empty result set is returned.

DB2 ODBC returns information on the input, input/output, and output parameters associated with the stored procedure, but cannot return information on the descriptor information for any result sets returned.

SQLProcedureColumns() returns the information in a result set, ordered by PROCEDURE\_CAT, PROCEDURE\_SCHEM, PROCEDURE\_NAME, and COLUMN\_TYPE. Table 106 lists the columns in the result set.

Since calls to SQLProcedureColumns() in many cases map to a complex and thus expensive query against the system catalog, they should be used sparingly, and the results saved rather than repeating calls.

The VARCHAR columns of the catalog functions result set have been declared with a maximum length attribute of 128 to be consistent with SQL92 limits. Since DB2 names are less than 128, the application can choose to always set aside 128 characters (plus the null-terminator) for the output buffer, or alternatively, call SQLGetInfo() with the SQL\_MAX\_CATALOG\_NAME\_LEN, SQL\_MAX\_SCHEMA\_NAME\_LEN, SQL\_MAX\_TABLE\_NAME\_LEN, and SQL\_MAX\_COLUMN\_NAME\_LEN to determine respectively the actual lengths of the TABLE\_CAT, TABLE\_SCHEM, TABLE\_NAME, and COLUMN\_NAME columns supported by the connected DBMS.

Applications should be aware that columns beyond the last column might be defined in future releases. Although new columns might be added and the names of the existing columns changed in future releases, the position of the current columns does not change.

Table 106 (Page 1 of 4). Columns returned by SQLProcedureColumns

Column number/name	Data type	Description
1 PROCEDURE_CAT	VARCHAR(128)	The is always null.
2 PROCEDURE_SCHEM	VARCHAR(128)	The name of the schema containing PROCEDURE_NAME. (This is also NULL for DB2 for OS/390 SQLProcedureColumns() result sets.)
3 PROCEDURE_NAME	VARCHAR(128)	Name of the procedure.
4 COLUMN_NAME	VARCHAR(128)	Name of the parameter.

## SQLProcedureColumns

Table 106 (Page 2 of 4). Columns returned by SQLProcedureColumns

Column number/name	Data type	Description
5 COLUMN_TYPE	SMALLINT NOT NULL	<p>Identifies the type information associated with this row. The values can be:</p> <ul style="list-style-type: none"> <li>SQL_PARAM_TYPE_UNKNOWN: the parameter type is unknown.</li> </ul> <p><b>Note:</b> This is not returned.</p> <ul style="list-style-type: none"> <li>SQL_PARAM_INPUT: this parameter is an input parameter.</li> <li>SQL_PARAM_INPUT_OUTPUT: this parameter is an input / output parameter.</li> <li>SQL_PARAM_OUTPUT: this parameter is an output parameter.</li> <li>SQL_RETURN_VALUE: the procedure column is the return value of the procedure.</li> </ul> <p><b>Note:</b> This is not returned.</p> <ul style="list-style-type: none"> <li>SQL_RESULT_COL: this parameter is actually a column in the result set.</li> </ul> <p><b>Note:</b> This is not returned.</p> <p><b>Note:</b> SQL_PARAM_OUTPUT and SQL_RETURN_VALUE are supported only on ODBC 2.0 or higher.</p>
6 DATA_TYPE	SMALLINT NOT NULL	SQL data type.
7 TYPE_NAME	VARCHAR(128) NOT NULL	Character string representing the name of the data type corresponding to DATA_TYPE.
8 COLUMN_SIZE	INTEGER	<p>If the DATA_TYPE column value denotes a character or binary string, then this column contains the maximum length in bytes; if it is a graphic (DBCS) string, this is the number of double byte characters for the parameter.</p> <p>For date, time, timestamp data types, this is the total number of bytes required to display the value when converted to character.</p> <p>For numeric data types, this is either the total number of digits, or the total number of bits allowed in the column, depending on the value in the NUM_PREC_RADIX column in the result set.</p> <p>See Table 151 on page 445.</p>
9 BUFFER_LENGTH	INTEGER	<p>The maximum number of bytes for the associated C buffer to store data from this parameter if SQL_C_DEFAULT is specified on the SQLBindCol(), SQLGetData() and SQLBindParameter() calls. This length excludes any null-terminator. For exact numeric data types, the length accounts for the decimal and the sign.</p> <p>See Table 153 on page 447.</p>
10 DECIMAL_DIGITS	SMALLINT	<p>The scale of the parameter. NULL is returned for data types where scale is not applicable.</p> <p>See Table 152 on page 446.</p>

Table 106 (Page 3 of 4). Columns returned by SQLProcedureColumns

Column number/name	Data type	Description
11 NUM_PREC_RADIX	SMALLINT	<p>Either 10 or 2 or NULL. If DATA_TYPE is an approximate numeric data type, this column contains the value 2, then the COLUMN_SIZE column contains the number of bits allowed in the parameter.</p> <p>If DATA_TYPE is an exact numeric data type, this column contains the value 10 and the COLUMN_SIZE and DECIMAL_DIGITS columns contain the number of decimal digits allowed for the parameter.</p> <p>For numeric data types, the DBMS can return a NUM_PREC_RADIX of either 10 or 2.</p> <p>NULL is returned for data types where radix is not applicable.</p>
12 NULLABLE	SMALLINT NOT NULL	<p>SQL_NO_NULLS if the parameter does not accept NULL values.</p> <p>SQL_NULLABLE if the parameter accepts NULL values.</p>
13 REMARKS	VARCHAR(254)	Might contain descriptive information about the parameter.
14 COLUMN_DEF	VARCHAR(254)	<p>The column's default value. If the default value is:</p> <ul style="list-style-type: none"> <li>• A numeric literal, this column contains the character representation of the numeric literal with no enclosing single quotes.</li> <li>• A character string, this column is that string enclosed in single quotes.</li> <li>• A pseudo-literal, such as for DATE, TIME, and TIMESTAMP columns, this column contains the keyword of the pseudo-literal (for example, CURRENT DATE) with no enclosing single quotes.</li> <li>• NULL, this column returns the word NULL, with no enclosing single quotes.</li> </ul> <p>If the default value cannot be represented without truncation, this column contains TRUNCATED with no enclosing single quotes. If no default value is specified, this column is NULL.</p>
15 SQL_DATA_TYPE	SMALLINT NOT NULL	The SQL data type. This column is the same as the DATA_TYPE column. For datetime data types, the SQL_DATA_TYPE field in the result set is SQL_DATETIME, and the SQL_DATETIME_SUB field returns the subcode for the specific datetime data type (SQL_CODE_DATE, SQL_CODE_TIME or SQL_CODE_TIMESTAMP).
16 SQL_DATETIME_SUB	SMALLINT	<p>The subtype code for datetime data types:</p> <ul style="list-style-type: none"> <li>• SQL_CODE_DATE</li> <li>• SQL_CODE_TIME</li> <li>• SQL_CODE_TIMESTAMP</li> </ul> <p>For all other data types, this column returns a NULL.</p>
17 CHAR_OCTET_LENGTH	INTEGER	The maximum length in bytes of a character data type column. For all other data types, this column returns a NULL.
18 ORDINAL_POSITION	INTEGER NOT NULL	Contains the ordinal position of the parameter given by COLUMN_NAME in this result set. This is the ordinal position of the argument provided on the CALL statement. The leftmost argument has an ordinal position of 1.

## SQLProcedureColumns

Table 106 (Page 4 of 4). Columns returned by SQLProcedureColumns

Column number/name	Data type	Description
19 IS_NULLABLE	VARCHAR(128)	One of the following: <ul style="list-style-type: none"><li>• "NO", if the column does not include NULLs</li><li>• "YES", if the column can include NULLs</li><li>• Zero-length string if nullability is unknown.</li></ul> The value returned for this column is different than the value returned for the NULLABLE column. (See the description of the NULLABLE column.)

**Note:** The column names used by DB2 ODBC follow the X/Open CLI CAE specification style. The column types, contents and order are identical to those defined for the SQLProcedureColumns() result set in ODBC.

# DB2 ODBC applications that issue SQLProcedureColumns() against a DB2 for OS/390 server, Version 5 or  
# later, should expect the result set columns listed in the table above. Revision bars identify the new and  
# changed columns.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 107. SQLProcedureColumns SQLSTATEs

SQLSTATE	Description	Explanation
24000	Invalid cursor state.	A cursor is already opened on the statement handle.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
42601	PARMLIST syntax error.	The PARMLIST value in the stored procedures catalog table contains a syntax error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1014	No more handles.	DB2 ODBC is not able to allocate a handle due to internal resources.
S1090	Invalid string or buffer length	The value of one of the name length arguments is less than 0, but not equal SQL_NTS.
S1C00	Driver not capable.	DB2 ODBC does not support <i>catalog</i> as a qualifier for procedure name.  The connected server does not support <i>schema</i> as a qualifier for procedure name.

## Restrictions

SQLProcedureColumns() does not return information about the attributes of result sets that stored procedures can return.

If an application is connected to a DB2 server that does not provide support for stored procedures, or for a stored procedure catalog, SQLProcedureColumns() returns an empty result set.

## Example

```

/*****
/* DB2 for OS/390 example: */
/* Invokes SQLProcedureColumns and enumerates all rows */
/* retrieved. */
*****/

#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <sqlca.h>
#include "sqlcli1.h"

int main( )
{
    SQLHENV          hEnv    = SQL_NULL_HENV;
    SQLHDBC          hDbc    = SQL_NULL_HDBC;
    SQLHSTMT         hStmt   = SQL_NULL_HSTMT;
    SQLRETURN        rc      = SQL_SUCCESS;
    SQLINTEGER       RETCODE = 0;
    char             *pDSN   = "STLEC1";
    char             procedure_name [20];
    char             parameter_name [20];
    char             ptype        [20];
    SQLSMALLINT      parameter_type = 0;
    SQLSMALLINT      data_type    = 0;
    char             type_name     [20];
    SWORD            cbCursor;
    SDWORD           cbValue3;
    SDWORD           cbValue4;
    SDWORD           cbValue5;
    SDWORD           cbValue6;
    SDWORD           cbValue7;
    char             ProcCatalog [20] = {0};
    char             ProcSchema  [20] = {0};
    char             ProcName    [20] = {"DOIT%"};
    char             ColumnName  [20] = {"P%"};
    SQLSMALLINT      cbProcCatalog = 0;
    SQLSMALLINT      cbProcSchema  = 0;
    SQLSMALLINT      cbProcName    = strlen(ProcName);
    SQLSMALLINT      cbColumnName  = strlen(ColumnName);

```

## SQLProcedureColumns

```
(void) printf ("**** Entering CLIP12.\n\n");

/*****
/* Allocate Environment Handle */
*****/

RETCODE = SQLAllocEnv(&hEnv);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Allocate Connection Handle to DSN */
*****/

RETCODE = SQLAllocConnect(hEnv,
                          &hDbc);

if( RETCODE != SQL_SUCCESS )    // Could not get a Connect Handle
    goto dberror;

/*****
/* CONNECT TO data source (STLEC1) */
*****/

RETCODE = SQLConnect(hDbc,      // Connect handle
                    (SQLCHAR *) pDSN, // DSN
                    SQL_NTS,    // DSN is nul-terminated
                    NULL,       // Null UID
                    0,          //
                    NULL,       // Null Auth string
                    0);

if( RETCODE != SQL_SUCCESS )    // Connect failed
    goto dberror;

/*****
/* Allocate Statement Handles */
*****/

rc = SQLAllocStmt (hDbc,
                  &hStmt);

if (rc != SQL_SUCCESS)
    goto exit;

/*****
/* Invoke SQLProcedureColumns and retrieve all rows within
/* answer set. */
*****/

rc = SQLProcedureColumns (hStmt
                        ,
                        (SQLCHAR *) ProcCatalog,
                        cbProcCatalog
                        ,
                        (SQLCHAR *) ProcSchema ,
                        cbProcSchema
                        ,
                        (SQLCHAR *) ProcName   ,
                        cbProcName
                        ,
                        (SQLCHAR *) ColumnName ,
                        cbColumnName);
```

```

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** SQLProcedureColumns Failed.\n");
    goto dberror;
}

rc = SQLBindCol (hStmt,          // bind procedure_name
                3,
                SQL_C_CHAR,
                procedure_name,
                sizeof(procedure_name),
                &cbValue3);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** Bind of procedure_name Failed.\n");
    goto dberror;
}

rc = SQLBindCol (hStmt,          // bind parameter_name
                4,
                SQL_C_CHAR,
                parameter_name,
                sizeof(parameter_name),
                &cbValue4);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** Bind of parameter_name Failed.\n");
    goto dberror;
}

rc = SQLBindCol (hStmt,          // bind parameter_type
                5,
                SQL_C_SHORT,
                &parameter_type,
                0,
                &cbValue5);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** Bind of parameter_type Failed.\n");
    goto dberror;
}

rc = SQLBindCol (hStmt,          // bind SQL data type
                6,
                SQL_C_SHORT,
                &data_type,
                0,
                &cbValue6);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** Bind of data_type Failed.\n");
    goto dberror;
}

```

## SQLProcedureColumns

```
rc = SQLBindCol (hStmt,          // bind type_name
                7,
                SQL_C_CHAR,
                type_name,
                sizeof(type_name),
                &cbValue7);

if (rc != SQL_SUCCESS)
{
    (void) printf ("**** Bind of type_name Failed.\n");
    goto dberror;
}

/*****
/* Answer Set is available - Fetch rows and print parameters for */
/* all procedures. */
*****/

while ((rc = SQLFetch (hStmt)) == SQL_SUCCESS)
{
    (void) printf ("**** Procedure Name = %s. Parameter %s",
                  procedure_name,
                  parameter_name);

    switch (parameter_type)
    {
        case SQL_PARAM_INPUT      :
            (void) strcpy (ptype, "INPUT");
            break;
        case SQL_PARAM_OUTPUT    :
            (void) strcpy (ptype, "OUTPUT");
            break;
        case SQL_PARAM_INPUT_OUTPUT :
            (void) strcpy (ptype, "INPUT/OUTPUT");
            break;
        default                    :
            (void) strcpy (ptype, "UNKNOWN");
            break;
    }

    (void) printf (" is %s. Data Type is %d. Type Name is %s.\n",
                  ptype,
                  data_type,
                  type_name);
}

/*****
/* Deallocate Statement Handles -- statement is no longer in a */
/* Prepared state. */
*****/

rc = SQLFreeStmt (hStmt,
                  SQL_DROP);

/*****
/* DISCONNECT from data source */
*****/

RETCODE = SQLDisconnect(hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;
```



```

/*****
/* Deallocate Connection Handle */
*****/

RETCODE = SQLFreeConnect (hDbc);

if (RETCODE != SQL_SUCCESS)
    goto dberror;

/*****
/* Free Environment Handle */
*****/

RETCODE = SQLFreeEnv (hEnv);

if (RETCODE == SQL_SUCCESS)
    goto exit;

dberror:
RETCODE=12;

exit:

(void) printf ("**** Exiting CLIP12.\n\n");

return RETCODE;
}

```

## References

- “SQLProcedures - Get list of procedure names” on page 300

## SQLProcedures - Get list of procedure names

### Purpose

Specification:	ODBC 1.0		
----------------	----------	--	--

SQLProcedures() returns a list of procedure names that have been registered at the server, and which match the specified search pattern.

The information is returned in an SQL result set, which can be retrieved using the same functions that are used to process a result set generated by a query.

### Syntax

```
SQLRETURN SQLProcedures (SQLHSTMT hstmt,
SQLCHAR FAR *szProcCatalog,
SQLSMALLINT cbProcCatalog,
SQLCHAR FAR *szProcSchema,
SQLSMALLINT cbProcSchema,
SQLCHAR FAR *szProcName,
SQLSMALLINT cbProcName);
```

### Function arguments

Table 108. SQLProcedures arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	Input	Statement handle.
SQLCHAR *	szProcCatalog	Input	Catalog qualifier of a 3 part procedure name. This must be a NULL pointer or a zero length string.
SQLSMALLINT	cbProcCatalog	Input	Length of <i>szProcCatalog</i> . This must be set to 0.
SQLCHAR *	szProcSchema	Input	Buffer that can contain a <i>pattern-value</i> to qualify the result set by schema name.  For Version 4 and Version 5 of DB2 for OS/390, all the stored procedures are in one schema; the only acceptable value for the <i>szProcSchema</i> argument is a null pointer. For DB2 UDB, <i>szProcSchema</i> can contain a valid pattern value. For more information about valid search patterns, see "Querying system catalog information" on page 368.
SQLSMALLINT	cbProcSchema	Input	Length of <i>szProcSchema</i> .
SQLCHAR *	szProcName	Input	Buffer that can contain a <i>pattern-value</i> to qualify the result set by table name.
SQLSMALLINT	cbProcName	Input	Length of <i>szProcName</i> .

### Usage

If the stored procedure is at a DB2 for MVS/ESA Version 4 server or a DB2 for OS/390 Version 5 server, the name of the stored procedures must be registered in the servers SYSIBM.SYSPROCEDURES catalog table. If the stored procedure is at a DB2 for OS/390 Version 6 server, the name of the stored procedure must be registered in the servers SYSIBM.SYSROUTINES catalog table.

For other versions of DB2 servers that do not provide facilities for a stored procedure catalog, an empty result set is returned.

The result set returned by `SQLProcedures()` contains the columns listed in Table 109 in the order given. The rows are ordered by `PROCEDURE_CAT`, `PROCEDURE_SCHEMA`, and `PROCEDURE_NAME`.

Since calls to `SQLProcedures()` in many cases map to a complex and thus expensive query against the system catalog, they should be used sparingly, and the results saved rather than repeating calls.

The `VARCHAR` columns of the catalog functions result set have been declared with a maximum length attribute of 128 to be consistent with SQL92 limits. Since DB2 names are less than 128, the application can choose to always set aside 128 characters (plus the null-terminator) for the output buffer, or alternatively, call `SQLGetInfo()` with the `SQL_MAX_CATALOG_NAME_LEN`, `SQL_MAX_SCHEMA_NAME_LEN`, `SQL_MAX_TABLE_NAME_LEN`, and `SQL_MAX_COLUMN_NAME_LEN` to determine respectively the actual lengths of the `TABLE_CAT`, `TABLE_SCHEM`, `TABLE_NAME`, and `COLUMN_NAME` columns supported by the connected DBMS.

Although new columns might be added and the names of the existing columns changed in future releases, the position of the current columns does not change.

Table 109. Columns returned by `SQLProcedures`

Column number/name	Data type	Description
1 PROCEDURE_CAT	VARCHAR(128)	This is always null.
2 PROCEDURE_SCHEM	VARCHAR(128)	The name of the schema containing PROCEDURE_NAME.
3 PROCEDURE_NAME	VARCHAR(128) NOT NULL	The name of the procedure.
4 NUM_INPUT_PARAMS	INTEGER not NULL	Number of input parameters.
5 NUM_OUTPUT_PARAMS	INTEGER not NULL	Number of output parameters.
6 NUM_RESULT_SETS	INTEGER not NULL	Number of result sets returned by the procedure.
7 REMARKS	VARCHAR(254)	Contains the descriptive information about the procedure.
8 PROCEDURE_TYPE	SMALLINT	Defines the procedure type: <ul style="list-style-type: none"> <li>• <code>SQL_PT_UNKNOWN</code>: It cannot be determined whether the procedure returns a value.</li> <li>• <code>SQL_PT_PROCEDURE</code>: The returned object is a procedure; that is, it does not have a return value.</li> <li>• <code>SQL_PT_FUNCTION</code>: The returned object is a function; that is, it has a return value.</li> </ul> DB2 ODBC always returns <code>SQL_PT_PROCEDURE</code> .

**Note:** The column names used by DB2 ODBC follow the X/Open CLI CAE specification style. The column types, contents and order are identical to those defined for the `SQLProcedures()` result set in ODBC.

### Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

### Diagnostics

Table 110. SQLProcedures SQLSTATEs

SQLSTATE	Description	Explanation
24000	Invalid cursor state.	A cursor is already opened on the statement handle.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1014	No more handles.	DB2 ODBC is not able to allocate a handle due to internal resources.
S1090	Invalid string or buffer length.	The value of one of the name length arguments is less than 0, but not equal to SQL_NTS.
S1C00	Driver not capable.	DB2 ODBC does not support <i>catalog</i> as a qualifier for procedure name.  The connected server does not supported schema as a qualifier for procedure name.

### Restrictions

If an application is connected to a DB2 server that does not provide support for stored procedures, or for a stored procedure catalog, `SQLProcedureColumns()` returns an empty result set.

## Example

```

/* ... */

printf("Enter Procedure Schema Name Search Pattern:\n");
gets(proc_schem.s);

rc = SQLProcedures(hstmt, NULL, 0, proc_schem.s, SQL_NTS, "%", SQL_NTS);

rc = SQLBindCol(hstmt, 2, SQL_C_CHAR, (SQLPOINTER) proc_schem.s, 129,
                &proc_schem.ind);

rc = SQLBindCol(hstmt, 3, SQL_C_CHAR, (SQLPOINTER) proc_name.s, 129,
                &proc_name.ind);

rc = SQLBindCol(hstmt, 7, SQL_C_CHAR, (SQLPOINTER) remarks.s, 255,
                &remarks.ind);

printf("PROCEDURE SCHEMA          PROCEDURE NAME          \n");
printf("----- \n");
/* Fetch each row, and display */
while ((rc = SQLFetch(hstmt)) == SQL_SUCCESS) {
    printf("%-25s %-25s\n", proc_schem.s, proc_name.s);
    if (remarks.ind != SQL_NULL_DATA) {
        printf(" (Remarks) %s\n", remarks.s);
    }
}
/* endwhile */
/* ... */

```

## References

- “SQLProcedureColumns - Get input/output parameter information for a procedure” on page 290

## SQLPutData - Passing data value for a parameter

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLPutData() is called following an SQLParamData() call returning SQL\_NEED\_DATA to supply parameter data values. This function can be used to send large parameter values in pieces.

The information is returned in an SQL result set, which can be retrieved using the same functions that are used to process a result set generated by a query.

### Syntax

```
SQLRETURN SQLPutData(
    (SQLHSTMT
    SQLPOINTER
    SQLINTEGER);
    hstmt,
    rgbValue,
    cbValue);
```

### Function arguments

Table 111. SQLPutData arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	Input	Statement handle.
SQLPOINTER	rgbValue	Input	Pointer to the actual data, or portion of data, for a parameter. The data must be in the form specified in the SQLBindParameter() call that the application used when specifying the parameter.
SQLINTEGER	cbValue	Input	Length of <i>rgbValue</i> . Specifies the amount of data sent in a call to SQLPutData() .  The amount of data can vary with each call for a given parameter. The application can also specify SQL_NTS or SQL_NULL_DATA for <i>cbValue</i> .  <i>cbValue</i> is ignored for all fixed length C buffer types, such as date, time, timestamp, and all numeric C buffer types.  For cases where the C buffer type is SQL_C_CHAR or SQL_C_BINARY, or if SQL_C_DEFAULT is specified as the C buffer type and the C buffer type default is SQL_C_CHAR or SQL_C_BINARY, this is the number of bytes of data in the <i>rgbValue</i> buffer.

### Usage

For a description on the SQLParamData() and SQLPutData() sequence, see "Sending/retrieving long data in pieces" on page 373.

The application calls SQLPutData() after calling SQLParamData() on a statement in the SQL\_NEED\_DATA state to supply the data values for an SQL\_DATA\_AT\_EXEC parameter. Long data can be sent in pieces using repeated calls to SQLPutData(). After all the pieces of data for the parameter have been sent, the application calls SQLParamData() again to proceed to the next

SQL\_DATA\_AT\_EXEC parameter, or, if all parameters have data values, to execute the statement.

SQLPutData() cannot be called more than once for a fixed length C buffer type, such as SQL\_C\_LONG.

After an SQLPutData() call, the only legal function calls are SQLParamData(), SQLCancel(), or another SQLPutData() if the input data is character or binary data. As with SQLParamData(), all other function calls using this statement handle fail. In addition, all function calls referencing the parent *hdbc* of *hstmt* fail if they involve changing any attribute or state of that connection; that is, the following function calls on the parent *hdbc* are also not permitted:

- SQLAllocConnect()
- SQLAllocStmt()
- SQLSetConnectOption()
- SQLNativeSql()
- SQLTransact()

Should they be invoked during an SQL\_NEED\_DATA sequence, these functions return SQL\_ERROR with SQLSTATE of **S1010** and the processing of the SQL\_DATA\_AT\_EXEC parameters is not affected.

If one or more calls to SQLPutData() for a single parameter results in SQL\_SUCCESS, attempting to call SQLPutData() with *cbValue* set to SQL\_NULL\_DATA for the same parameter results in an error with SQLSTATE of **22005**. This error does not result in a change of state; the statement handle is still in a *Need Data* state and the application can continue sending parameter data.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Some of the following diagnostic conditions are also reported on the final SQLParamData() call rather than at the time the SQLPutData() is called.

Table 112 (Page 1 of 2). SQLPutData SQLSTATEs

SQLSTATE	Description	Explanation
01004	Data truncated.	The data sent for a numeric parameter is truncated without the loss of significant digits.  Timestamp data sent for a date or time column is truncated.  Function returns with SQL_SUCCESS_WITH_INFO.
22001	String data right truncation.	More data is sent for a binary or char data than the data source can support for that column.
22003	Numeric value out of range.	The data sent for a numeric parameter cause the whole part of the number to be truncated when assigned to the associated column.  SQLPutData() was called more than once for a fixed length parameter.

Table 112 (Page 2 of 2). SQLPutData SQLSTATES

SQLSTATE	Description	Explanation
22005	Error in assignment.	The data sent for a parameter is incompatible with the data type of the associated table column.
22007	Invalid datetime format.	The data value sent for a date, time, or timestamp parameters is invalid.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	The argument <i>rgbValue</i> is a NULL pointer, and the argument <i>cbValue</i> is neither 0 nor SQL_NULL_DATA.
S1010	Function sequence error.	The statement handle <i>hstmt</i> must be in a need data state and must have been positioned on an SQL_DATA_AT_EXEC parameter using a previous SQLParamData() call.
S1090	Invalid string or buffer length.	The argument <i>rgbValue</i> is not a NULL pointer, and the argument <i>cbValue</i> is less than 0, but not equal to SQL_NTS or SQL_NULL_DATA.

## Restrictions

A new value for *pcbValue*, SQL\_DEFAULT\_PARAM, was introduced in ODBC 2.0, to indicate that the procedure is to use the default value of a parameter, rather than a value sent from the application. Since the concept of default values does not apply to DB2 stored procedure arguments, specification of this value for the *pcbValue* argument results in an error when the CALL statement is executed because the SQL\_DEFAULT\_PARAM value is considered an invalid length.

ODBC 2.0 also introduced the SQL\_LEN\_DATA\_AT\_EXEC(*length*) macro to be used with the *pcbValue* argument. The macro is used to specify the sum total length of the entire data that would be sent for character or binary C data using the subsequent SQLPutData() calls. Since the DB2 ODBC driver does not need this information, the macro is not needed. An ODBC application calls SQLGetInfo() with the SQL\_NEED\_LONG\_DATA\_LEN option to check if the driver needs this information. The DB2 ODBC driver returns 'N' to indicate that this information is not needed by SQLPutData().

## Example

See “Example” on page 204.

## References

- “SQLBindParameter - Binds a parameter marker to a buffer or LOB locator” on page 93
- “SQLExecute - Execute a statement” on page 159
- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLParamData - Get next parameter for which a data value is needed” on page 273
- “SQLCancel - Cancel statement” on page 104



## SQLRowCount - Get row count

### Purpose

<b>Specification:</b>	<b>ODBC 1.0</b>	<b>X/OPEN CLI</b>	<b>ISO CLI</b>
-----------------------	-----------------	-------------------	----------------

SQLRowCount() returns the number of rows in a table that were affected by an UPDATE, INSERT, or DELETE statement executed against the table, or a view based on the table.

SQLExecute() or SQLExecDirect() must be called before calling this function.

### Syntax

```
SQLRETURN SQLRowCount (SQLHSTMT hstmt,
                        SQLINTEGER FAR *pcrow);
```

### Function arguments

Table 113. SQLrowcount arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle
SQLINTEGER *	<i>pcrow</i>	output	Pointer to location where the number of rows affected is stored.

### Usage

If the last executed statement referenced by the input statement handle was not an UPDATE, INSERT, or DELETE statement, or if it did not execute successfully, then the function sets the contents of *pcrow* to -1.

If SQLRowCount() is executed after the SQLExecDirect() or SQLExecute() of an SQL statement other than INSERT, UPDATE, or DELETE, it results in return code 0 and *pcrow* is set to -1.

Any rows in other tables that might be affected by the statement (for example, cascading deletes) are not included in the count.

If SQLRowCount() is executed after a built-in function (for example, SQLTables()), it results in return code -1 and SQLSTATE S1010.

### Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

### Diagnostics

Table 114. SQLRowCount SQLSTATEs

SQLSTATE	Description	Explanation
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called prior to calling SQLExecute() or SQLExecDirect() for the <i>hstmt</i> .
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.

### Restrictions

None.

### Example

See “Example” on page 135.

### References

- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLExecute - Execute a statement” on page 159
- “SQLNumResultCols - Get number of result columns” on page 271

---

## SQLSetColAttributes - Set column attributes

### Purpose

<b>Specification:</b>			
-----------------------	--	--	--

SQLSetColAttributes() sets the data source result descriptor (column name, type, precision, scale and nullability) for one column in the result set so that the DB2 ODBC implementation does not have to obtain the descriptor information from the DBMS server.

### Syntax

```
SQLRETURN SQLSetColAttributes (SQLHSTMT      hstmt,
                               SQLUSMALLINT   icol,
                               SQLCHAR FAR    *pszColName,
                               SQLSMALLINT    cbColName,
                               SQLSMALLINT    fSQLType,
                               SQLINTEGER     cbColDef,
                               SQLSMALLINT    ibScale,
                               SQLSMALLINT    fNullable);
```

### Function arguments

Table 115 (Page 1 of 2). SQLSetColAttributes arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle.
SQLUSMALLINT	icol	input	Column number of result data, ordered sequentially left to right, starting at 1.
SQLCHAR *	szColName	input	Pointer to the column name. If the column is unnamed or is an expression, this pointer can be set to NULL, or an empty string can be used.
SQLSMALLINT	cbColName	input	Length of szColName buffer.

## SQLSetColAttributes

Table 115 (Page 2 of 2). SQLSetColAttributes arguments

Data type	Argument	Use	Description
SQLSMALLINT	fSqlType	input	The SQL data type of the column. The following values are recognized: <ul style="list-style-type: none"><li>• SQL_BINARY</li><li>• SQL_BLOB</li><li>• SQL_CHAR</li><li>• SQL_CLOB</li><li>• SQL_DATE</li><li>• SQL_DBCLOB</li><li>• SQL_DECIMAL</li><li>• SQL_DOUBLE</li><li>• SQL_FLOAT</li><li>• SQL_GRAPHIC</li><li>• SQL_INTEGER</li><li>• SQL_LONGVARIABLE</li><li>• SQL_LONGVARCHAR</li><li>• SQL_LONGVARGRAPHIC</li><li>• SQL_NUMERIC</li><li>• SQL_REAL</li><li>• SQL_ROWID</li><li>• SQL_SMALLINT</li><li>• SQL_TIME</li><li>• SQL_TIMESTAMP</li><li>• SQL_VARBINARY</li><li>• SQL_VARCHAR</li><li>• SQL_VARGRAPHIC</li></ul>
SQLINTEGER	cbColDef	input	The precision of the column on the data source.
SQLSMALLINT	ibScale	input	The scale of the column on the data source. This is ignored for all data types except SQL_DECIMAL, SQL_NUMERIC, SQL_TIMESTAMP.
SQLSMALLINT	fNullable	input	Indicates whether the column allows NULL value. This must be one of the following values: <ul style="list-style-type: none"><li>• SQL_NO_NULLS - the column does not allow NULL values.</li><li>• SQL_NULLABLE - the column allows NULL values.</li></ul>

## Usage

This function is designed to help reduce the amount of network traffic that can result when an application is fetching result sets that contain an extremely large number of columns. If the application has advanced knowledge of the characteristics of the descriptor information of a result set (that is, the exact number of columns, column name, data type, nullability, precision, or scale), then it can inform DB2 ODBC rather than having DB2 ODBC obtain this information from the database, thus reducing the quantity of network traffic.

An application typically calls `SQLSetColAttributes()` after a call to `SQLPrepare()` and before the associated call to `SQLExecute()`. An application can also call `SQLSetColAttributes()` before a call to `SQLExecDirect()`. This function is valid only after the statement option `SQL_NODESCRIBE` has been set to `SQL_NODESCRIBE_ON` for this statement handle.

SQLSetColAttributes() informs DB2 ODBC of the column name, type, and length that would be generated by the subsequent execution of the query. This information allows DB2 ODBC to determine whether any data conversion is necessary when the result is returned to the application.

**Recommendation:** The application should only use this function if it has prior knowledge of the exact nature of the result set.

The application must provide the result descriptor information for every column in the result set or an error occurs on the subsequent fetch (SQLSTATE 07002). Using this function only benefits those applications that handle an extremely large number (hundreds) of columns in a result set, otherwise the effect is minimal.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 116 (Page 1 of 2). SQLSetColAttributes SQLSTATEs

SQLSTATE	Description	Explanation
01004	Data truncated.	szColName contains a column name that is too long. To obtain the maximum length of the column name, call SQLGetInfo with the flInfoType SQL_MAX_COLUMN_NAME_LEN.
24000	Invalid cursor state.	A cursor is already open on the statement handle.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1000	General error.	An error occurred for which there is no specific SQLSTATE and for which no implementation defined SQLSTATE is defined. The error message returned by SQLError() in the argument szErrorMsg describes the error and its cause.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1002	Invalid column number.	The value specified for the argument icol is less than 1 or greater than the maximum number of columns supported by the server.
S1004	SQL data type out of range.	The value specified for the argument fSqlType is not a valid SQL data type.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1090	Invalid string or buffer length.	The value specified for the argument cbColName is less than 0 and not equal to SQL_NTS.
S1094	Invalid scale value.	The value specified for fSqlType is either SQL_DECIMAL or SQL_NUMERIC and the value specified for ibScale is less than 0 or greater than the value for the argument cbColDef (precision). The value specified for fSqlType is SQL_TIMESTAMP and the value for ibScale is less than 0 or greater than 6.

## SQLSetColAttributes

Table 116 (Page 2 of 2). SQLSetColAttributes SQLSTATES

SQLSTATE	Description	Explanation
S1099	Nullable type out of range.	The value specified for <i>fNullable</i> is invalid.
S1104	Invalid precision value.	The value specified for <i>fSqlType</i> is either SQL_DECIMAL or SQL_NUMERIC and the value specified for <i>cbColDef</i> is less than 1.

## Restrictions

None.

## Example

```
/* ... */
SQLCHAR      stmt[] =
    { "Select id, name from staff" };
/* ... */

/* Tell DB2 ODBC not to get Column Attribute from the server for this hstmt */
rc = SQLSetStmtOption(hstmt, SQL_NODESCRIBE, SQL_NODESCRIBE_ON);

rc = SQLPrepare(hstmt, stmt, SQL_NTS);

/* Provide the columns attributes to DB2 ODBC for this hstmt */
rc = SQLSetColAttributes(hstmt, 1, "-ID-", SQL_NTS, SQL_SMALLINT,
                        5, 0, SQL_NO_NULLS);
rc = SQLSetColAttributes(hstmt, 2, "-NAME-", SQL_NTS, SQL_CHAR,
                        9, 0, SQL_NULLABLE);
rc = SQLExecute(hstmt);

print_results(hstmt); /* Call sample function to print column attributes
                        and fetch and print rows. */

rc = SQLFreeStmt(hstmt, SQL_DROP);

rc = SQLTransact(henv, hdbc, SQL_COMMIT);

printf("Disconnecting ....\n");
rc = SQLDisconnect(hdbc);

rc = SQLFreeConnect(hdbc);
if (rc != SQL_SUCCESS)
    return (terminate(henv, rc));

rc = SQLFreeEnv(henv);
if (rc != SQL_SUCCESS)
    return (terminate(henv, rc));

return (SQL_SUCCESS);
} /* end main */
```

## References

- “SQLColAttributes - Get column attributes” on page 106
- “SQLDescribeCol - Describe column attributes” on page 132
- “SQLExecute - Execute a statement” on page 159
- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLPrepare - Prepare a statement” on page 277

---

## SQLSetConnection - Set connection handle

### Purpose

<b>Specification:</b>			
-----------------------	--	--	--

This function is needed if the application needs to deterministically switch to a particular connection before continuing execution. It should only be used when the application is mixing DB2 ODBC function calls with embedded SQL function calls and multiple connections are involved.

### Syntax

```
SQLRETURN SQLSetConnection (SQLHDBC          hdbc);
```

### Function arguments

Table 117. SQLSetConnection arguments

Data type	Argument	Use	Description
SQLHDBC	hdbc	input	The connection handle associated with the connection to which the application wishes to switch.

### Usage

ODBC allows multiple concurrent connections. It is not clear which connection an embedded SQL routine uses when invoked. In practice, the embedded routine uses the connection associated with the most recent network activity. However, from the application's perspective, this is not always easy to determine and it is difficult to keep track of this information. `SQLSetConnection()` is used to allow the application to *explicitly* specify which connection is active. The application can then call the embedded SQL routine.

`SQLSetConnection()` is not needed at all if the application makes purely DB2 ODBC calls. This is because each statement handle is implicitly associated with a connection handle and there is never any confusion as to which connection a particular DB2 ODBC function applies.

For more information on using embedded SQL within DB2 ODBC applications see "Mixing embedded SQL and DB2 ODBC" on page 401.

### Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE



## Diagnostics

Table 118. SQLSetConnection SQLSTATES

SQLSTATE	Description	Explanation
08003	Connection is closed.	The connection handle provided is not currently associated with an open connection to a database server.
S1000	General error.	An error occurred for which there is no specific SQLSTATE and for which the implementation does not define an SQLSTATE. SQLError() returns an error message in the argument <i>szErrorMsg</i> that describes the error and its cause.

## Restrictions

None.

## Example

See “Mixed embedded SQL and DB2 ODBC example” on page 401.

## References

- “SQLConnect - Connect to a data source” on page 123
- “SQLDriverConnect - (Expanded) connect to a data source” on page 142

## SQLSetConnectOption - Set connection option

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	
----------------	----------	------------	--

SQLSetConnectOption() sets connection attributes for a particular connection.

### Syntax

```
SQLRETURN SQLSetConnectOption(
    SQLHDBC          hdbc,
    SQLUSMALLINT     fOption,
    SQLINTEGER        vParam);
```

### Function arguments

Table 119. SQLSetConnectOption arguments

Data Type	Argument	Use	Description
HDBC	hdbc	Input	Connection handle.
SQLUSMALLINT	fOption	Input	Connect option to set, see Table 120 on page 317 for the complete list of connect options and their description.
SQLINTEGER	vParam	Input	Value associated with <i>fOption</i> . Depending on the option, this can be a 32-bit integer value, or a pointer to a null-terminated string.

### Usage

The SQLSetConnectOption() can be used to specify statement options for *all* statement handles on this connection, as well as for all future statement handles on this connection. For a list of statement options, see “SQLSetStmtOption - Set statement option” on page 334.

All connection and statement options set using the SQLSetConnectOption() persist until SQLFreeConnect() is called or the next SQLSetConnectOption() call.

It is illegal to call SQLSetConnectOption() (SQLSTATE **S1010**) if any of the statement handles associated with this connection is in a need data state (that is, in the middle of an SQLParamData() -- SQLPutData() sequence to process SQL\_DATA\_AT\_EXEC parameters). This sequence is described in “Sending/retrieving long data in pieces” on page 373.

The format of information set with *vParam* depends on the specified *fOption*. The option information can be either a 32-bit integer or a pointer to a null-terminated character string. In the case of the null-terminated character string, the maximum length of the string can be SQL\_MAX\_OPTION\_STRING\_LENGTH bytes (excluding the null-terminator).

Table 120 (Page 1 of 3). Connect options

Option	Contents
SQL_ACCESS_MODE	<p>A 32-bit integer value which can be either:</p> <ul style="list-style-type: none"> <li>• <b>SQL_MODE_READ_ONLY</b>: Indicates that the application is not performing any updates on data from this point on. Therefore, a less restrictive isolation level and locking can be used on transactions; that is, uncommitted read (SQL_TXN_READ_UNCOMMITTED).</li> </ul> <p>DB2 ODBC does not ensure that requests to the database are <i>read-only</i>. If an update request is issued, DB2 ODBC processes it using the transaction isolation level it selected as a result of the SQL_MODE_READ_ONLY setting.</p> <ul style="list-style-type: none"> <li>• <b>SQL_MODE_READ_WRITE</b>: Indicates that the application is making updates on data from this point on. DB2 ODBC goes back to using the default transaction isolation level for this connection.</li> </ul> <p>SQL_MODE_READ_WRITE is the default.</p> <p>There must not be any outstanding transactions on this connection.</p>
SQL_AUTOCOMMIT	<p>A 32-bit integer value that specifies whether to use auto-commit or manual commit mode:</p> <ul style="list-style-type: none"> <li>• <b>SQL_AUTOCOMMIT_OFF</b>: The application must manually, explicitly commit or rollback transactions with SQLTransact() calls.</li> <li>• <b>SQL_AUTOCOMMIT_ON</b>: DB2 ODBC operates in auto-commit mode. Each statement is implicitly committed. Each statement, that is not a query, is committed immediately after it has been executed. Each query is committed immediately after the associated cursor is closed.</li> </ul> <p>SQL_AUTOCOMMIT_ON is the default.</p> <p><b>Note:</b> If this is a coordinated distributed unit of work connection, then the default is <b>SQL_AUTOCOMMIT_OFF</b></p> <p>When specifying auto-commit, the application can have only one outstanding statement per connection. For example, there must not be two open cursors, or unpredictable results can occur. An open cursor must be closed before another query is executed.</p> <p>Since in many DB2 environments, the execution of the SQL statements and the commit can be flowed separately to the database server, autocommit can be expensive. The application developer should take this into consideration when selecting the auto-commit mode.</p> <p>Changing from manual-commit to auto-commit mode commits any open transaction on the connection.</p>

## SQLSetConnectOption

Table 120 (Page 2 of 3). Connect options

Option	Contents
SQL_CONNECTTYPE	<p data-bbox="532 268 1433 384">A 32-bit integer value that specifies whether this application is to operate in a coordinated or uncoordinated distributed environment. If the processing needs to be coordinated, then this option must be considered in conjunction with the SQL_SYNC_POINT connection option. The possible values are:</p> <ul data-bbox="548 405 1433 548" style="list-style-type: none"><li>• <b>SQL_CONCURRENT_TRANS:</b> The application can have concurrent multiple connections to any one database or to multiple databases. This option setting corresponds to the specification of the type 1 CONNECT in embedded SQL. Each connection has its own commit scope. No effort is made to enforce coordination of transaction.</li></ul> <p data-bbox="573 569 1260 596">The current setting of the SQL_SYNC_POINT option is ignored.</p> <p data-bbox="573 617 773 644">This is the default.</p> <ul data-bbox="548 665 1433 831" style="list-style-type: none"><li>• <b>SQL_COORDINATED_TRANS:</b> The application wishes to have commit and rollbacks coordinated among multiple database connections. This option setting corresponds to the specification of the type 2 CONNECT in embedded SQL and must be considered in conjunction with the SQL_SYNC_POINT connection option. In contrast to the SQL_CONCURRENT_TRANS setting described above, the application is permitted only one open connection per database.</li></ul> <p data-bbox="573 852 1373 905"><b>Note:</b> This connection type results in the default for SQL_AUTOCOMMIT connection option to be SQL_AUTOCOMMIT_OFF.</p> <p data-bbox="532 926 1312 978">This option must be set before making a connect request; otherwise, the SQLSetConnectOption() call is rejected.</p> <p data-bbox="532 999 1349 1079">All the connections within an application must have the same SQL_CONNECTTYPE and SQL_SYNC_POINT values. The first connection determines the acceptable attributes for the subsequent connections.</p> <p data-bbox="532 1100 1433 1245"><b>Recommendation:</b> Have the application set the SQL_CONNECTTYPE attribute at the environment level rather than on a per connection basis. ODBC applications written to take advantage of coordinated DB2 transactions must set these attributes at the connection level for each connection as SQLSetEnvAttr() is not supported in ODBC.</p> <p data-bbox="532 1266 972 1293"><b>Note:</b> This is an IBM-defined extension.</p>
SQL_CURRENT_SCHEMA	<p data-bbox="532 1314 1433 1367">A null-terminated character string containing the name of the schema to be used by DB2 ODBC for the SQLColumns() call if the szSchemaName pointer is set to null.</p> <p data-bbox="532 1388 1401 1440">To reset this option, specify this option with a zero length or a null pointer for the vParam argument.</p> <p data-bbox="532 1461 1393 1541">This option is useful when the application developer has coded a generic call to SQLColumns() that does not restrict the result set by schema name, but needs to constrain the result set at isolated places in the code.</p> <p data-bbox="532 1562 1409 1614">This option can be set at any time and is effective on the next SQLColumns() call where the szSchemaName pointer is null.</p> <p data-bbox="532 1635 972 1663"><b>Note:</b> This is an IBM-defined extension.</p>

Table 120 (Page 3 of 3). Connect options

fOption	Contents
SQL_MAXCONN	<p>A 32-bit integer value corresponding to the number of maximum concurrent connections that an application wants to set up. The default value is <b>0</b>, which means no maximum - the application is allowed to set up as many connections as the system resources permit. The integer value must be 0 or a positive number.</p> <p>This can be used as a governor for the maximum number of connections on a per application basis.</p> <p>The value that is in effect when the first connection is established is the value that is used. When the first connection is established, attempts to change this value are rejected.</p> <p><b>Recommendation:</b> Have the application set SQL_MAXCONN at the environment level rather than on a connection basis. ODBC applications must set this attribute at the connection level since SQLSetEnvAttr() is not supported in ODBC.</p> <p><b>Note:</b> This is an IBM-defined extension.</p>
SQL_PARAMOPT_ATOMIC	<p>If specified, DB2 ODBC returns S1C00 on SQLSetConnectOption and S1011 on SQLGetConnectOption.</p>
SQL_TXN_ISOLATION	<p>A 32-bit bitmask that sets the transaction isolation level for the current connection referenced by <i>hdbc</i>. The valid values for <i>vParam</i> can be determined at runtime by calling SQLGetInfo() with <i>fInfoType</i> set to SQL_TXN_ISOLATION_OPTIONS. The following values are accepted by DB2 ODBC, but each server might only support a subset of these isolation levels:</p> <ul style="list-style-type: none"> <li>• SQL_TXN_READ_UNCOMMITTED - Dirty reads, reads that cannot be repeated, and phantoms are possible.</li> <li>• <b>SQL_TXN_READ_COMMITTED</b> - Dirty reads are not possible. Reads that cannot be repeated, and phantoms are possible.</li> </ul> <p>This is the default.</p> <ul style="list-style-type: none"> <li>• SQL_TXN_REPEATABLE_READ - Dirty reads and reads that cannot be repeated are not possible. Phantoms are possible.</li> <li>• SQL_TXN_SERIALIZABLE - Transactions can be serialized. Dirty reads, non-repeatable reads, and phantoms are not possible.</li> <li>• SQL_TXN_NOCOMMIT - Any changes are effectively committed at the end of a successful operation; no explicit commit or rollback is allowed. This is analogous to autocommit. This is not an SQL92 isolation level, but an IBM defined extension, supported only by DB2 for AS/400.</li> </ul> <p>In IBM terminology,</p> <ul style="list-style-type: none"> <li>• SQL_TXN_READ_UNCOMMITTED is uncommitted read;</li> <li>• SQL_TXN_READ_COMMITTED is cursor stability;</li> <li>• SQL_TXN_REPEATABLE_READ is read stability;</li> <li>• SQL_TXN_SERIALIZABLE is repeatable read.</li> </ul> <p>For a detailed explanation of isolation levels, see <i>IBM SQL Reference</i></p> <p>This option cannot be specified while there is an open cursor on any hstmt, or an outstanding transaction for this connection; otherwise, SQL_ERROR is returned on the function call (SQLSTATE <b>S1011</b>).</p> <p><b>Note:</b> There is an IBM extension that permits the setting of transaction isolation levels on a per statement handle basis. See the SQL_STMTTXN_ISOLATION option in the function description for SQLSetStmtOption().</p>

### Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

### Diagnostics

Table 121. SQLSetConnectOption SQLSTATES

SQLSTATE	Description	Explanation
<b>Note:</b> Since SQLSetConnectOption() can also be used to set statement options, SQLSTATES for SQLSetConnectOption() can also include those listed under "Diagnostics" for the SQLSetStmtOption() API.		
01000	Warning.	Informational message indicating an internal commit has been issued on behalf of the application as part of the processing to set the specified connection option.
01S02	Option value changed.	SQL_CONNECTTYPE changed to SQL_CONCURRENT_TRANS when MULTICONTEXT=1 in use.
08003	Connection is closed.	An <i>fOption</i> is specified that required an open connection.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	Given the <i>fOption</i> value, an invalid value is specified for the argument <i>vParam</i> .
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1011	Operation invalid at this time.	The <i>fOption</i> specified option cannot be set at this time: <ul style="list-style-type: none"> <li>• SQL_CONNECTTYPE: attempt is made to change the value of this option from its current value but the connection is open.</li> <li>• SQL_TXN_ISOLATION, SQL_ACCESS_MODE: a transaction is outstanding.</li> </ul>
S1092	Option type out of range.	An invalid <i>fOption</i> value is specified.
S1C00	Driver not capable.	The specified <i>fOption</i> is not supported.  Given specified <i>fOption</i> value, the value specified for the argument <i>vParam</i> is not supported.

### Restrictions

For compatibility with ODBC applications, *fOption* values of SQL\_CURRENT\_QUALIFIER and SQL\_PACKET\_SIZE are also recognized, but not supported. If either of these two options are specified, SQL\_ERROR is returned on the function call (SQLSTATE S1C00).

ODBC *fOption* values of SQL\_TRANSLATE\_DLL and SQL\_TRANSLATE\_OPTION are not supported since DB2 handles codepage conversion at the server, not the client.

## Example

See “Example” on page 125.

## References

- “SQLGetConnectOption - Returns current setting of a connect option” on page 191
- “SQLGetStmtOption - Returns current setting of a statement option” on page 249
- “SQLSetStmtOption - Set statement option” on page 334

---

## SQLSetCursorName - Set cursor name

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLSetCursorName() associates a cursor name with the statement handle. This function is optional since DB2 ODBC implicitly generates a cursor name when each statement handle is allocated.

### Syntax

```
SQLRETURN SQLSetCursorName (SQLHSTMT          hstmt,
                             SQLCHAR FAR       *szCursor,
                             SQLSMALLINT      cbCursor);
```

### Function arguments

Table 122. SQLSetCursorName arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle
SQLCHAR *	<i>szCursor</i>	input	Cursor name
SQLSMALLINT	<i>cbCursor</i>	input	Length of contents of <i>szCursor</i> argument

### Usage

DB2 ODBC always generates and uses an internally generated cursor name when a query is prepared or executed directly. SQLSetCursorName() allows an application defined cursor name to be used in an SQL statement (a positioned UPDATE or DELETE). DB2 ODBC maps this name to the internal name. The name remains associated with the statement handle, until the handle is dropped, or another SQLSetCursorName() is called on this statement handle.

Although SQLGetCursorName() returns the name set by the application (if one is set), error messages associated with positioned UPDATE and DELETE statements refer to the internal name.

**Recommendation:** Do not use SQLSetCursorName(). Instead, use the internal name which can be obtained by calling SQLGetCursorName().

Cursor names must follow these rules:

- All cursor names within the connection must be unique.
- Each cursor name must be less than or equal to 18 bytes in length. Any attempt to set a cursor name longer than 18 bytes results in truncation of that cursor name to 18 bytes. (No warning is generated.)
- Since internally generated names begin with SQLCUR, SQL\_CUR, or SQLCURQRS, the application must not input a cursor name starting with either SQLCUR or SQL\_CUR in order to avoid conflicts with internal names.
- Since a cursor name is considered an identifier in SQL, it must begin with an English letter (a-z, A-Z) followed by any combination of digits (0-9), English letters or the underscore character (\_).



- To permit cursor names containing characters other than those listed above (such as National Language Set or Double Bytes Character Set characters), the application must enclose the cursor name in double quotes ("").
- Unless the input cursor name is enclosed in double quotes, all leading and trailing blanks from the input cursor name string are removed.

For efficient processing, applications should not include any leading or trailing spaces in the *szCursor* buffer. If the *szCursor* buffer contains a delimited identifier, applications should position the first double quote as the first character in the *szCursor* buffer.

## Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 123. SQLSetCursorName SQLSTATEs

SQLSTATE	Description	Explanation
34000	Invalid cursor name.	The cursor name specified by the argument <i>szCursor</i> is invalid. The cursor name either begins with SQLCUR, SQL_CUR, or SQLCURQRS or violates the cursor naming rules (Must begin with a-z or A-Z followed by any combination of English letters, digits, or the '_' character.  The cursor name specified by the argument <i>szCursor</i> already exists.  The cursor name length is greater than the value returned by SQLGetInfo() with the SQL_MAX_CURSOR_NAME_LEN argument.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	<i>szCursor</i> is a null pointer.
S1010	Function sequence error.	There is an open or positioned cursor on the statement handle.  The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation called prior to SQLSetCursorName().
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1090	Invalid string or buffer length.	The argument <i>cbCursor</i> is less than 0, but not equal to SQL_NTS.

### Restrictions

None.

### Example

```
/* ... */
SQLCHAR      sqlstmt[] =
              "SELECT name, job FROM staff "
              "WHERE job='Clerk'  FOR UPDATE OF job";
/* ... */
/* allocate second statement handle for update statement */
rc2 = SQLAllocStmt(hdbc, &hstmt2);

/* Set Cursor for the SELECT statement's handle */
rc = SQLSetCursorName(hstmt1, "JOB CURS", SQL_NTS);

rc = SQLExecDirect(hstmt1, sqlstmt, SQL_NTS);

/* bind name to first column in the result set */
rc = SQLBindCol(hstmt1, 1, SQL_C_CHAR, (SQLPOINTER) name.s, 10,
               &name.ind);

/* bind job to second column in the result set */
rc = SQLBindCol(hstmt1, 2, SQL_C_CHAR, (SQLPOINTER) job.s, 6,
               &job.ind);

printf("Job Change for all clerks\n");

while ((rc = SQLFetch(hstmt1)) == SQL_SUCCESS) {
    printf("Name: %-9.9s Job: %-5.5s \n", name.s, job.s);
    printf("Enter new job or return to continue\n");
    gets(newjob);
    if (newjob[0] != '\0') {
        sprintf(updstmt,
              "UPDATE staff set job = '%s' where current of JOB CURS",
              newjob);
        rc2 = SQLExecDirect(hstmt2, updstmt, SQL_NTS);
    }
}
if (rc != SQL_NO_DATA_FOUND)
    check_error(henv, hdbc, hstmt1, rc, __LINE__, __FILE__);
/* ... */
```

### References

- "SQLGetCursorName - Get cursor name" on page 193

## SQLSetEnvAttr - Set environment attribute

### Purpose

Specification:		X/OPEN CLI	ISO CLI
----------------	--	------------	---------

SQLSetEnvAttr() sets an environment attribute for the current environment.

### Syntax

```
SQLRETURN SQLSetEnvAttr (SQLHENV          henv,
                          SQLINTEGER       Attribute,
                          SQLPOINTER      Value,
                          SQLINTEGER       StringLength);
```

### Function arguments

Table 124. SQLSetEnvAttr arguments

Data type	Argument	Use	Description
SQLHENV	henv	Input	Environment handle.
SQLINTEGER	Attribute	Input	Environment attribute to set. See Table 125 for the list of attributes and their descriptions.
SQLPOINTER	Value	Input	The desired value for <i>Attribute</i> .
SQLINTEGER	StringLength	Input	Length of <i>Value</i> in bytes if the attribute value is a character string; if <i>Attribute</i> does not denote a string, then DB2 ODBC ignores <i>StringLength</i> .

### Usage

When set, the attribute's value affects all connections in this environment.

The application can obtain the current attribute value by calling SQLGetEnvAttr().

Table 125 (Page 1 of 2). Environment attributes

Attribute	Contents
SQL_ATTR_OUTPUT_NTS	<p>A 32-bit integer value which controls the use of null-termination in output arguments. The possible values are:</p> <ul style="list-style-type: none"> <li>• <b>SQL_TRUE</b>: DB2 ODBC uses null termination to indicate the length of output character strings. This is the default.</li> <li>• <b>SQL_FALSE</b>: DB2 ODBC does not use null termination in output character strings.</li> </ul> <p>The CLI functions affected by this attribute are all functions called for the environment (and for any connections and statements allocated under the environment) that have character string parameters.</p> <p>This attribute can only be set when there are no connection handles allocated under this environment.</p>

Table 125 (Page 2 of 2). Environment attributes

Attribute	Contents
SQL_CONNECTTYPE	<p>A 32-bit integer value that specifies whether this application is to operate in a coordinated or uncoordinated distributed environment. The possible values are:</p> <ul style="list-style-type: none"> <li>• <b>SQL_CONCURRENT_TRANS:</b> Each connection has its own commit scope. No effort is made to enforce coordination of transaction. If an application issues a commit using the environment handle on <code>SQLTransact()</code> and not all of the connections commit successfully, the application is responsible for recovery. This corresponds to <code>CONNECT</code> (type 1) semantics subject to the restrictions described in “DB2 ODBC restrictions on the ODBC connection model” on page 23. This is the default.</li> <li>• <b>SQL_COORDINATED_TRANS:</b> The application wishes to have commit and rollbacks coordinated among multiple database connections. In contrast to the <code>SQL_CONCURRENT_TRANS</code> setting described above, the application is permitted only one open connection per database.</li> </ul> <p>This attribute must be set before allocating any connection handles, otherwise, the <code>SQLSetEnvAttr()</code> call is rejected.</p> <p>All the connections within an application must have the same <code>SQL_CONNECTTYPE</code> and <code>SQL_SYNCPOINT</code> values. This attribute can also be set using the <code>SQLSetConnectOption</code> function.</p> <p><b>Recommendation:</b> Have the application set the <code>SQL_CONNECTTYPE</code> attribute at the environment level rather than on a per connection basis. ODBC applications written to take advantage of coordinated DB2 transactions must set these attributes at the connection level for each connection using <code>SQLSetConnectOption()</code> as <code>SQLSetEnvAttr()</code> is not supported in ODBC.</p> <p><b>Note:</b> This is an IBM-defined extension.</p>
SQL_MAXCONN	<p>A 32-bit integer value corresponding to the number that maximum concurrent connections that an application wants to set up. The default value is <b>0</b>, which means no maximum - the application is allowed to set up as many connections as the system resources permit. The integer value must be 0 or a positive number.</p> <p>This can be used as a governor for the maximum number of connections on a per application basis.</p> <p>The value that is in effect when the first connection is established is the value that is used. When the first connection is established, attempts to change this value are rejected.</p> <p><b>Recommendation:</b> Have the application set <code>SQL_MAXCONN</code> at the environment level rather than on a connection basis. ODBC applications must set this attribute at the connection level since <code>SQLSetEnvAttr()</code> is not supported in ODBC.</p> <p><b>Note:</b> This is an IBM-defined extension.</p>

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 126. SQLSetEnvAttr SQLSTATEs

SQLSTATE	Description	Explanation
01S02	Option value changed.	SQL_CONNECTTYPE changed to SQL_CONCURRENT_TRANS when MULTICONTEXT=1 in use.
S1009	Invalid argument value.	Given the <i>fOption</i> value, an invalid value is specified for the argument <i>vParam</i> .
S1011	Operation invalid at this time.	Applications cannot set environment attributes while connection handles are allocated on the environment handle.
S1092	Option type out of range.	An invalid <i>Attribute</i> value is specified.
S1C00	Driver not capable.	The specified <i>Attribute</i> is not supported by DB2 ODBC. Given specified <i>Attribute</i> value, the value specified for the argument <i>Value</i> is not supported.

## Restrictions

None.

## Example

See also, “Distributed unit of work example” on page 367.

```

/* ... */
int
main()
{
    SQLHENV        henv;
    SQLRETURN      rc;
    SQLINTEGER     output_nts = SQL_TRUE;

    rc = SQLAllocEnv(&henv); /* allocate an environment handle */
    if (rc == SQL_SUCCESS)
    { rc = SQLSetEnvAttr(henv, SQL_ATTR_OUTPUT_NTS, output_nts,
                        0);
    }

    rc = SQLFreeEnv(henv);

}
/* ... */

```

## References

- “SQLGetEnvAttr - Returns current setting of an environment attribute” on page 212

---

## SQLSetParam - Binds a parameter marker to a buffer

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	
----------------	----------	------------	--

**Note:** In ODBC 2.0, this function is replaced by `SQLBindParameter()`. See the restrictions section below for details.

`SQLSetParam()` is used to associate (bind) parameter markers in an SQL statement to one of the following:

- Application variables (storage buffers), for all data types. In this case data is transferred from the application to the DBMS when `SQLExecute()` or `SQLExecDirect()` is called. Data conversion can occur as the data is transferred.
- A LOB locator, for SQL LOB data types. In this case, the application transfers a LOB locator value (not the LOB data itself) to the server when the SQL statement is executed.

### Syntax

```
SQLRETURN SQLSetParam (SQLHSTMT hstmt,
                        SQLUSMALLINT ipar,
                        SQLSMALLINT fCType,
                        SQLSMALLINT fSqlType,
                        SQLINTEGER cbParamDef,
                        SQLSMALLINT ibScale,
                        SQLPOINTER rgbValue,
                        SQLINTEGER FAR *pcbValue);
```

### Function arguments

Table 127 (Page 1 of 3). *SQLSetParam* arguments

Data type	Argument	Use	Description
SQLHSTMT	<i>hstmt</i>	input	Statement handle.
SQLUSMALLINT	<i>ipar</i>	input	Parameter marker number, ordered sequentially left to right, starting at 1.

Table 127 (Page 2 of 3). SQLSetParam arguments

Data type	Argument	Use	Description
SQLSMALLINT	<i>fCType</i>	input	<p>C data type of argument. The following types are supported:</p> <ul style="list-style-type: none"> <li>• SQL_C_BINARY</li> <li>• SQL_C_BIT</li> <li>• SQL_C_BLOB_LOCATOR</li> <li>• SQL_C_CHAR</li> <li>• SQL_C_CLOB_LOCATOR</li> <li>• SQL_C_DATE</li> <li>• SQL_C_DBCHAR</li> <li>• SQL_C_DBCLOB_LOCATOR</li> <li>• SQL_C_DOUBLE</li> <li>• SQL_C_FLOAT</li> <li>• SQL_C_LONG</li> <li>• SQL_C_SHORT</li> <li>• SQL_C_TIME</li> <li>• SQL_C_TIMESTAMP</li> <li>• SQL_C_TINYINT</li> </ul> <p>Specifying SQL_C_DEFAULT causes data to be transferred from its default C data type for the type indicated in <i>fSqlType</i>. See Table 4 on page 39 for more information.</p>
SQLSMALLINT	<i>fSqlType</i>	input	<p>SQL data type of column. The supported types are:</p> <ul style="list-style-type: none"> <li>• SQL_BINARY</li> <li>• SQL_BLOB</li> <li>• SQL_BLOB_LOCATOR</li> <li>• SQL_CHAR</li> <li>• SQL_CLOB</li> <li>• SQL_CLOB_LOCATOR</li> <li>• SQL_DATE</li> <li>• SQL_DBCLOB</li> <li>• SQL_DBCLOB_LOCATOR</li> <li>• SQL_DECIMAL</li> <li>• SQL_DOUBLE</li> <li>• SQL_FLOAT</li> <li>• SQL_GRAPHIC</li> <li>• SQL_INTEGER</li> <li>• SQL_LONGVARIABLE</li> <li>• SQL_LONGVARCHAR</li> <li>• SQL_LONGVARGRAPHIC</li> <li>• SQL_NUMERIC</li> <li>• SQL_REAL</li> <li>• SQL_ROWID</li> <li>• SQL_SMALLINT</li> <li>• SQL_TIME</li> <li>• SQL_TIMESTAMP</li> <li>• SQL_VARBINARY</li> <li>• SQL_VARCHAR</li> <li>• SQL_VARGRAPHIC</li> </ul> <p><b>Note:</b> SQL_BLOB_LOCATOR, SQL_CLOB_LOCATOR, and SQL_DBCLOB_LOCATOR are application related concepts and do not map to a data type for column definition during a CREATE TABLE.</p>

Table 127 (Page 3 of 3). SQLSetParam arguments

Data type	Argument	Use	Description
SQLINTEGER	<i>cbParamDef</i>	input	<p>Precision of the corresponding parameter marker. If <i>fSqlType</i> denotes:</p> <ul style="list-style-type: none"> <li>• A binary or single byte character string (for example, SQL_CHAR, SQL_BINARY), this is the maximum length in bytes for this parameter marker.</li> <li>• A double byte character string (for example, SQL_GRAPHIC), this is the maximum length in double-byte characters for this parameter.</li> <li>• SQL_DECIMAL, SQL_NUMERIC, this is the maximum decimal precision.</li> <li>• SQL_ROWID, this must be set to 40, the maximum length in bytes for this data type. Otherwise, an error is returned.</li> <li>• Otherwise, this argument is ignored.</li> </ul>
SQLSMALLINT	<i>ibScale</i>	input	<p>Scale of the corresponding parameter marker if <i>fSqlType</i> is SQL_DECIMAL or SQL_NUMERIC. If <i>fSqlType</i> is SQL_TIMESTAMP, this is the number of digits to the right of the decimal point in the character representation of a timestamp (for example, the scale of yyyy-mm-dd hh:mm:ss.fff is 3).</p> <p>Other than for the <i>fSqlType</i> values mentioned here, <i>ibScale</i> is ignored.</p>
SQLPOINTER	<i>rgbValue</i>	input (deferred)	<p>Pointer to the location which contains (when the statement is executed) the actual values for the associated parameter marker.</p>
SQLINTEGER *	<i>pcbValue</i>	input (deferred)	<p>Pointer to the location which contains (when the statement is executed) the length of the parameter marker value stored at <i>rgbValue</i>.</p> <p>To specify a null value for a parameter marker, this storage location must contain SQL_NULL_DATA.</p> <p>If <i>fCType</i> is SQL_C_CHAR, this storage location must contain either the exact length of the data stored at <i>rgbValue</i>, or SQL_NTS if the contents at <i>rgbValue</i> are null-terminated.</p> <p>If <i>fCType</i> indicates character data (explicitly, or implicitly using SQL_C_DEFAULT), and this pointer is set to NULL, it is assumed that the application always provides a null-terminated string in <i>rgbValue</i>. This also implies that this parameter marker never contains a null value.</p> <p>If <i>fSqlType</i> indicates a graphic data type, and the <i>fCType</i> is SQL_C_CHAR, the pointer to <i>pcbValue</i> can never be NULL and the contents of <i>pcbValue</i> can never hold SQL_NTS. In general for graphic data types, this length should be the number of octets that the double byte data occupies; therefore, the length should always be a multiple of 2. In fact, if the length is odd, then an error occurs when the statement is executed.</p>



## Usage

A parameter marker is represented by a "?" character in an SQL statement and is used to indicate a position in the statement where an application supplied value is to be substituted when the statement is executed. This value can be obtained from:

- An application variable.

SQLSetParam() (or SQLBindParameter()) is used to bind the application storage area to the parameter marker.

- A LOB value from the database server (by specifying a LOB locator).

SQLSetParam() (or SQLBindParameter()) is used to bind a LOB locator to the parameter marker. The LOB value itself is supplied by the database server, so only the LOB locator is transferred between the database server and the application.

An application can use a locator with SQLGetSubString(), SQLGetSubString() or SQLGetLength(). SQLGetSubString() can either return another locator, or the data itself. All locators remain valid until the end of the transaction in which they are created (even when the cursor moves to another row, or until it is freed using the FREE LOCATOR statement).

The application must bind a variable to each parameter marker in the SQL statement before executing the SQL statement. For this function, *rgbValue* and *pcbValue* are deferred arguments. The storage locations must be valid and contain input data values when the statement is executed. This means either keeping the SQLExecDirect() or SQLExecute() call in the same procedure scope as the SQLBindParameter() calls, or, these storage locations must be dynamically allocated or declared statically or globally.

SQLSetParam() can be called before SQLPrepare() if the columns in the result set are known, otherwise the attributes of the result set can be obtained after the statement is prepared.

Parameter markers are referenced by number (*icol*) and are numbered sequentially from left to right, starting at 1.

All parameters bound by this function remain in effect until SQLFreeStmt() is called with either the SQL\_DROP or SQL\_RESET\_PARAMS option, or until SQLSetParam() is called again for the same parameter *ipar* number.

After the SQL statement is executed, and the results processed, the application can reuse the statement handle to execute a different SQL statement. If the parameter marker specifications are different (number of parameters, length or type), then SQLFreeStmt() should be called with SQL\_RESET\_PARAMS to reset or clear the parameter bindings.

The C buffer data type given by *fCType* must be compatible with the SQL data type indicated by *fSqlType*, or an error occurs.

If *fSqlType* is SQL\_ROWID, the *cbParamDef* value must be set to 40, the maximum length in bytes for a ROWID data type. If *cbParamDef* is not set to 40, the application receives SQLSTATE=22001 when *cbParamDef* is less than 40 and SQLSTATE=S1104 when *cbParamDef* is greater than 40.

## SQLSetParam

An application can pass the value for a parameter either in the *rgbValue* buffer or with one or more calls to `SQLPutData()`. In the latter case, these parameters are data-at-execution parameters. The application informs DB2 ODBC of a data-at-execution parameter by placing the `SQL_DATA_AT_EXEC` value in the *pcbValue* buffer. It sets the *rgbValue* input argument to a 32-bit value which is returned on a subsequent `SQLParamData()` call and can be used to identify the parameter position.

Since the data in the variables referenced by *rgbValue* and *pcbValue* is not verified until the statement is executed, data content or format errors are not detected or reported until `SQLExecute()` or `SQLExecDirect()` is called.

## Return codes

- `SQL_SUCCESS`
- `SQL_ERROR`
- `SQL_INVALID_HANDLE`

## Diagnostics

Table 128 (Page 1 of 2). *SQLSetParam* SQLSTATEs

SQLSTATE	Description	Explanation
07006	Invalid conversion.	The conversion from the data value identified by the <i>fCType</i> argument to the data type identified by the <i>fSqlType</i> argument is not a meaningful conversion. (For example, conversion from <code>SQL_C_DATE</code> to <code>SQL_DOUBLE</code> .)
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1003	Program type out of range.	The value specified by the argument <i>fCType</i> is not a valid data type or <code>SQL_C_DEFAULT</code> .
S1004	SQL data type out of range.	The value specified for the argument <i>fSqlType</i> is not a valid SQL data type.
S1009	Invalid argument value.	The argument <i>rgbValue</i> is a null pointer.
S1010	Function sequence error.	There is an open or positioned cursor on the statement handle. The function is called while in a data-at-execute ( <code>SQLParamData()</code> , <code>SQLPutData()</code> ) operation, called prior to <code>SQLSetCursorName()</code> .
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.
S1093	Invalid parameter number.	The value specified for the argument <i>ipar</i> is less than 1 or greater than the maximum number of parameters supported by the server.

Table 128 (Page 2 of 2). SQLSetParam SQLSTATES

SQLSTATE	Description	Explanation
S1094	Invalid scale value.	The value specified for <i>fSqlType</i> is either SQL_DECIMAL or SQL_NUMERIC and the value specified for <i>ibScale</i> is less than 0 or greater than the value for the argument <i>cbParamDef</i> (precision).  The value specified for <i>fCType</i> is SQL_C_TIMESTAMP and the value for <i>fSqlType</i> is either SQL_CHAR or SQL_VARCHAR and the value for <i>ibScale</i> is less than 0 or greater than 6.
S1104	Invalid precision value.	The value specified for <i>cbParamDef</i> is either less than 0 or greater than the permissible range for the <i>fSQLType</i> .
S1C00	Driver not capable.	DB2 ODBC or the data source does not support the conversion specified by the combination of the value specified for the argument <i>fCType</i> and the value specified for the argument <i>fSqlType</i> .  The value specified for the argument <i>fCType</i> or <i>fSqlType</i> is not supported by either DB2 ODBC or the data source.

## Restrictions

In ODBC 2.0, SQLSetParam() has replaced by SQLBindParameter().

SQLSetParam() cannot be used to:

- Bind application variables to parameter markers in a stored procedure CALL statement.
- Bind arrays of application variables when *SQLParamOptions()* has been used to specify multiple input parameter values.

*SQLBindParameter()* should be used instead in both of the above situations.

## Example

See “Example” on page 280.

## References

- “SQLBindParameter - Binds a parameter marker to a buffer or LOB locator” on page 93
- “SQLExecDirect - Execute a statement directly” on page 154
- “SQLExecute - Execute a statement” on page 159
- “SQLPrepare - Prepare a statement” on page 277

---

## SQLSetStmtOption - Set statement option

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	
----------------	----------	------------	--

SQLSetStmtOption() sets an attribute of a specific statement handle. To set an option for all statement handles associated with a connection handle, the application can call SQLSetConnectOption() (see “SQLSetConnectOption - Set connection option” on page 316).

### Syntax

```
SQLRETURN SQLSetStmtOption (SQLHSTMT          hstmt,
                             SQLUSMALLINT      fOption,
                             SQLINTEGER        vParam);
```

### Function arguments

Table 129. SQLSetStmtOption arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	input	Statement handle.
SQLUSMALLINT	fOption	input	Option to set. See Table 130 on page 335 for the list of statement options that can be set and their descriptions.
SQLINTEGER	vParam	input	Value associated with <i>fOption</i> . <i>vParam</i> can be a 32-bit integer value or a pointer to a null-terminated string.

### Usage

Statement options for an *hstmt* remain in effect until they are changed by another call to SQLSetStmtOption() or SQLSetConnectOption(), or the *hstmt* is dropped by calling SQLFreeStmt() with the SQL\_DROP option. Calling SQLFreeStmt() with the SQL\_CLOSE, SQL\_UNBIND, or SQL\_RESET\_PARAMS options does not reset statement options.

The format of *vParam* depends on the value specified *fOption*. The format of each is noted in Table 130 on page 335. If the format denotes a pointer to a null-terminated character string the maximum length is SQL\_MAX\_OPTION\_STRING\_LENGTH (excluding the null terminator).

**Note:** Currently no statement option requires a string.

Table 130 (Page 1 of 3). Statement options

Option	Contents
<b>Note:</b> Values shown in <b>bold</b> are default values.	
SQL_BIND_TYPE	<p>A 32-bit integer value that sets the binding orientation to be used when <code>SQLExtendedFetch()</code> is called with this statement handle. <i>Column-wise binding</i> is selected by supplying the value <b>SQL_BIND_BY_COLUMN</b> for the argument <i>vParam</i>. <i>Row-wise binding</i> is selected by supplying a value for <i>vParam</i> specifying the length of the structure or an instance of a buffer into which result columns are bound.</p> <p>For row-wise binding, the length specified in <i>vParam</i> must include space for all of the bound columns and any padding of the structure or buffer to ensure that when the address of a bound column is incremented with the specified length, the result points to the beginning of the same column in the next row. (When using the <code>sizeof</code> operator with structures or unions in ANSI C, this behavior is guaranteed.)</p>
SQL_CLOSE_BEHAVIOR	<p>A 32-bit integer that forces the release of locks upon an underlying <code>CLOSE CURSOR</code> operation. The possible values are:</p> <ul style="list-style-type: none"> <li>• <b>SQL_CC_NO_RELEASE:</b> locks are not released when the cursor on this statement handle is closed.</li> <li>• <b>SQL_CC_RELEASE:</b> locks are released when the cursor on this statement handle is closed.</li> </ul> <p>Typically cursors are explicitly closed when the function <code>SQLFreeStmt()</code> is called with the <code>SQL_CLOSE</code> or <code>SQL_DROP</code> option. In addition, the end of the transaction (when a commit or rollback is issued) can also close the cursor (depending on the <code>WITH HOLD</code> attribute currently in use).</p>
SQL_CONCURRENCY	<p>If specified, DB2 ODBC returns <code>S1C00</code> on <code>SQLSetConnectOption</code> and <code>S1011</code> on <code>SQLGetConnectOption</code>.</p>
SQL_CURSOR_HOLD	<p>A 32-bit integer which specifies whether the cursor associated with this <code>hstmt</code> is preserved in the same position as before the <code>COMMIT</code> operation, and whether the application can fetch without executing the statement again.</p> <ul style="list-style-type: none"> <li>• <b>SQL_CURSOR_HOLD_ON</b></li> <li>• <b>SQL_CURSOR_HOLD_OFF</b></li> </ul> <p>The default value when an <code>hstmt</code> is first allocated is <code>SQL_CURSOR_HOLD_ON</code>. This option cannot be specified while there is an open cursor on this <code>hstmt</code>.</p>
SQL_CURSOR_TYPE	<p>A 32-bit integer value that specifies the cursor type. The currently supported value is:</p> <ul style="list-style-type: none"> <li>• <b>SQL_CURSOR_FORWARD_ONLY</b> - Cursor behaves as a forward only scrolling cursor.</li> </ul> <p>This option cannot be set if there is an open cursor on the associated <code>hstmt</code>.</p> <p><b>Note:</b> ODBC architecture has also defined the following values, which are not supported by DB2 ODBC:</p> <ul style="list-style-type: none"> <li>• <code>SQL_CURSOR_STATIC</code> - The data in the result set appears to be static.</li> <li>• <code>SQL_CURSOR_KEYSET_DRIVEN</code> - The keys for the number of rows specified in the <code>SQL_KEYSET_SIZE</code> option is stored. DB2 ODBC does not support this option value.</li> <li>• <code>SQL_CURSOR_DYNAMIC</code> - The keys for the rows in the rowset are saved. DB2 ODBC does not support this option value.</li> </ul> <p>If one of these values is used, <code>SQL_SUCCESS_WITH_INFO</code> (<code>SQLSTATE 01S02</code>) is returned and the value remains unchanged.</p>

## SQLSetStmtOption

Table 130 (Page 2 of 3). Statement options

<i>fOption</i>	<b>Contents</b>
SQL_MAX_LENGTH	<p>A 32-bit integer value corresponding to the maximum amount of data that can be retrieved from a single character or binary column. If data is truncated because the value specified for SQL_MAX_LENGTH is less than the amount of data available, an SQLGetData() call or fetch returns SQL_SUCCESS instead of returning SQL_SUCCESS_WITH_INFO and SQLSTATE 01004 (data truncated). The default value for <i>vParam</i> is 0; 0 means that DB2 ODBC attempts to return all available data for character or binary type data.</p>
SQL_MAX_ROWS	<p>A 32-bit integer value corresponding to the maximum number of rows to return to the application from a query. The default value for <i>vParam</i> is 0; 0 means all rows are returned.</p>
SQL_NODESCRIBE	<p>A 32-bit integer which specifies whether DB2 ODBC should automatically describe the column attributes of the result set or wait to be informed by the application using SQLSetColAttributes().</p> <ul style="list-style-type: none"><li>• <b>SQL_NODESCRIBE_OFF</b></li><li>• <b>SQL_NODESCRIBE_ON</b></li></ul> <p>This option cannot be specified while there is an open cursor on this hstmt.</p> <p>This option is used in conjunction with the function SQLSetColAttributes() by an application which has prior knowledge of the exact nature of the result set to be returned and which does not wish to incur the extra network traffic associated with the descriptor information needed by DB2 ODBC to provide client side processing.</p> <p><b>Note:</b> This option is an IBM-defined extension.</p>
SQL_NOSCAN	<p>A 32-bit integer value that specifies whether DB2 ODBC will scan SQL strings for escape clauses. The two permitted values are:</p> <ul style="list-style-type: none"><li>• <b>SQL_NOSCAN_OFF</b> - SQL strings are scanned for escape clause sequences.</li><li>• <b>SQL_NOSCAN_ON</b> - SQL strings are not scanned for escape clauses. Everything is sent directly to the server for processing.</li></ul> <p>This application can choose to turn off the scanning if it never uses vendor escape sequences in the SQL strings that it sends. This eliminates some of the overhead processing associated with scanning.</p>
SQL_RETRIEVE_DATA	<p>A 32-bit integer value indicating whether DB2 ODBC should actually retrieve data from the database when SQLExtendedFetch() is called. The possible values are:</p> <ul style="list-style-type: none"><li>• <b>SQL_RD_ON:</b> SQLExtendedFetch() does retrieve data.</li><li>• <b>SQL_RD_OFF:</b> SQLExtendedFetch() does not retrieve data. This is useful for verifying whether rows exist without incurring the overhead of sending long data from the database server. DB2 ODBC internally retrieves all the fixed length columns, such as INTEGER and SMALLINT; so there is still some overhead.</li></ul> <p>This option cannot be set if the cursor is open.</p>
SQL_ROWSET_SIZE	<p>A 32-bit integer value that specifies the number of rows in the rowset. A rowset is the array of rows returned by each call to SQLExtendedFetch(). The default value is 1, which is equivalent to making a single SQLFetch(). This option can be specified even when the cursor is open and becomes effective on the next SQLExtendedFetch() call.</p>

Table 130 (Page 3 of 3). Statement options

<i>fOption</i>	Contents
SQL_STMTTXN_ISOLATION SQL_TXN_ISOLATION	<p>A 32-bit integer value that sets the transaction isolation level for the current hstmt. This overrides the default value set at the connection level (refer also to “SQLSetConnectOption - Set connection option” on page 316 for the permitted values).</p> <p>This option cannot be set if there is an open cursor on this statement handle (SQLSTATE 24000).</p> <p>The value SQL_STMTTXN_ISOLATION is synonymous with SQL_TXN_ISOLATION.</p> <p><b>Note:</b> It is an IBM extension to allow setting this option at the statement level.</p>

## Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 131. SQLSetStmtOption SQLSTATEs

SQLSTATE	Description	Explanation
01S02	Option value changed.	A recognized concurrency value is specified for the SQL_CONCURRENCY option, but is not supported by DB2 ODBC.
24000	Invalid cursor state.	<i>fOption</i> is set to SQL_CONCURRENCY, SQL_CURSOR_TYPE, SQL_STMTTXN_ISOLATION, or SQL_TXN_ISOLATION and a cursor is already opened on the statement handle.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1000	General error.	An error occurred for which there is no specific SQLSTATE and for which no implementation defined SQLSTATE is defined. The error message returned by <code>SQLError()</code> in the argument <code>szErrorMsg</code> describes the error and its cause.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value.	Given the specified <i>fOption</i> value, an invalid value is specified for the argument <i>vParam</i> .
S1010	Function sequence error.	The function is called while in a data-at-execute ( <code>SQLParamData()</code> , <code>SQLPutData()</code> ) operation; called prior to <code>SQLSetCursorName()</code> .
S1011	Operation invalid at this time.	The <i>fOption</i> is SQL_CONCURRENCY, SQL_CURSOR_HOLD, SQL_NODESCRIBE, SQL_RETRIEVE_DATA, SQL_(STMT)TXN_ISOLATION, or SQL_CURSOR_TYPE and the statement is prepared.
S1092	Option type out of range.	An invalid <i>fOption</i> value is specified.
S1C00	Driver not capable.	The option or option value is not supported.

### Restrictions

ODBC also defines statement options SQL\_KEYSET\_SIZE, SQL\_BOOKMARKS and SQL\_SIMULATE\_CURSOR. These options are not supported by DB2 ODBC. If either one is specified, SQL\_ERROR (SQLSTATE S1C00) is returned.

### Example

See “Example” on page 312.

### References

- “SQLColAttributes - Get column attributes” on page 106
- “SQLExtendedFetch - Extended fetch (fetch array of rows)” on page 162
- “SQLFetch - Fetch next row” on page 169
- “SQLGetConnectOption - Returns current setting of a connect option” on page 191
- “SQLGetData - Get data from a column” on page 199
- “SQLGetStmtOption - Returns current setting of a statement option” on page 249
- “SQLParamOptions - Specify an input array for a parameter” on page 275
- “SQLSetConnectOption - Set connection option” on page 316



## SQLSpecialColumns - Get special (row identifier) columns

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	
----------------	----------	------------	--

SQLSpecialColumns() returns unique row identifier information (primary key or unique index) for a table. The information is returned in an SQL result set, which can be retrieved using the same functions that are used to process a result set generated by a query.

### Syntax

```
SQLRETURN SQLSpecialColumns(SQLHSTMT hstmt,
                             SQLUSMALLINT fColType,
                             SQLCHAR FAR *szCatalogName,
                             SQLSMALLINT cbCatalogName,
                             SQLCHAR FAR *szSchemaName,
                             SQLSMALLINT cbSchemaName,
                             SQLCHAR FAR *szTableName,
                             SQLSMALLINT cbTableName,
                             SQLUSMALLINT fScope,
                             SQLUSMALLINT fNullable);
```

### Function arguments

Table 132 (Page 1 of 2). SQLSpecialColumns arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	Input	Statement handle.
SQLUSMALLINT	fColType	Input	Type of unique row identifier to return. Only the following type is supported: <ul style="list-style-type: none"> <li>SQL_BEST_ROWID</li> </ul> Returns the optimal set of columns that can uniquely identify any row in the specified table. <p><b>Note:</b> For compatibility with ODBC applications, SQL_ROWVER is also recognized, but not supported; therefore, if SQL_ROWVER is specified, an empty result is returned.</p>
SQLCHAR *	szCatalogName	Input	Catalog qualifier of a 3 part table name. This must be a null pointer or a zero length string.
SQLSMALLINT	cbCatalogName	Input	Length of <i>szCatalogName</i> . This must be a set to 0.
SQLCHAR *	szSchemaName	Input	Schema qualifier of the specified table.
SQLSMALLINT	cbSchemaName	Input	Length of <i>szSchemaName</i> .
SQLCHAR *	szTableName	Input	Table name.
SQLSMALLINT	cbTableName	Input	Length of <i>cbTableName</i> .

Table 132 (Page 2 of 2). SQLSpecialColumns arguments

Data type	Argument	Use	Description
SQLUSMALLINT	fScope	Input	<p>Minimum required duration for which the unique row identifier is valid.</p> <p><i>fScope</i> must be one of the following:</p> <ul style="list-style-type: none"> <li>SQL_SCOPE_CURROW: The row identifier is guaranteed to be valid only while positioned on that row. A later re-select using the same row identifier values might not return a row if the row was updated or deleted by another transaction.</li> <li>SQL_SCOPE_TRANSACTION: The row identifier is guaranteed to be valid for the duration of the current transaction.</li> </ul> <p><b>Note:</b> This option is only valid if SQL_TXN_SERIALIZABLE and SQL_TXN_REPEATABLE_READ isolation options are set.</p> <ul style="list-style-type: none"> <li>SQL_SCOPE_SESSION: The row identifier is guaranteed to be valid for the duration of the connection.</li> </ul> <p><b>Note:</b> This option is not supported by DB2 for OS/390.</p> <p>The duration over which a row identifier value is guaranteed to be valid depends on the current transaction isolation level. For information and scenarios involving isolation levels, see <i>DB2 SQL Reference</i>.</p>
SQLUSMALLINT	fNullable	Input	<p>Determines whether to return special columns that can have a NULL value.</p> <p>Must be one of the following:</p> <ul style="list-style-type: none"> <li>SQL_NO_NULLS - The row identifier column set returned cannot have any NULL values.</li> <li>SQL_NULLABLE - The row identifier column set returned can include columns where NULL values are permitted.</li> </ul>

## Usage

If multiple ways exist to uniquely identify any row in a table (that is, if there are multiple unique indexes on the specified table), then DB2 ODBC returns the *best* set of row identifier column sets based on its internal criterion.

If there is no column set that allows any row in the table to be uniquely identified, an empty result set is returned.

The unique row identifier information is returned in the form of a result set where each column of the row identifier is represented by one row in the result set. Table 133 on page 341 shows the order of the columns in the result set returned by `SQLSpecialColumns()`, sorted by SCOPE.

Since calls to `SQLSpecialColumns()` in many cases map to a complex and thus expensive query against the system catalog, they should be used sparingly, and the results saved rather than repeating calls.

The VARCHAR columns of the catalog functions result set are declared with a maximum length attribute of 128 to be consistent with SQL92 limits. Since DB2 names are less than 128, the application can choose to always set aside 128 characters (plus the null-terminator) for the output buffer, or alternatively, call `SQLGetInfo()` with the `SQL_MAX_COLUMN_NAME_LEN` to determine the actual length of the `COLUMN_NAME` column supported by the connected DBMS.

Although new columns might be added and the names of the columns changed in future releases, the position of the current columns does not change.

Table 133 (Page 1 of 2). Columns returned by `SQLSpecialColumns`

Column number/name	Data type	Description
1 SCOPE	SMALLINT	The duration for which the name in <code>COLUMN_NAME</code> is guaranteed to point to the same row. Valid values are the same as for the <i>fScope</i> argument: Actual scope of the row identifier. Contains one of the following values: <ul style="list-style-type: none"> <li>• <code>SQL_SCOPE_CURROW</code></li> <li>• <code>SQL_SCOPE_TRANSACTION</code></li> <li>• <code>SQL_SCOPE_SESSION</code></li> </ul> See <i>fScope</i> in Table 132 on page 339 for a description of each value.
2 COLUMN_NAME	VARCHAR(128) NOT NULL	Name of the column that is (or part of) the table's primary key.
3 DATA_TYPE	SMALLINT NOT NULL	SQL data type of the column. One of the values in the Symbolic SQL Data Type column in Table 4 on page 39.
4 TYPE_NAME	VARCHAR(128) NOT NULL	DBMS character string represented of the name associated with <code>DATA_TYPE</code> column value.
5 COLUMN_SIZE	INTEGER	If the <code>DATA_TYPE</code> column value denotes a character or binary string, then this column contains the maximum length in bytes; if it is a graphic (DBCS) string, this is the number of double byte characters for the parameter.  For date, time, timestamp data types, this is the total number of bytes required to display the value when converted to character.  For numeric data types, this is either the total number of digits, or the total number of bits allowed in the column, depending on the value in the <code>NUM_PREC_RADIX</code> column in the result set.  See Table 151 on page 445.
6 BUFFER_LENGTH	INTEGER	The maximum number of bytes for the associated C buffer to store data from this column if <code>SQL_C_DEFAULT</code> is specified on the <code>SQLBindCol()</code> , <code>SQLGetData()</code> and <code>SQLBindParameter()</code> calls. This length does not include any null-terminator. For exact numeric data types, the length accounts for the decimal and the sign.  See Table 153 on page 447.
7 DECIMAL_DIGITS	SMALLINT	The scale of the column. NULL is returned for data types where scale is not applicable. See Table 152 on page 446.

## SQLSpecialColumns

Table 133 (Page 2 of 2). Columns returned by SQLSpecialColumns

Column number/name	Data type	Description
8 PSEUDO_COLUMN	SMALLINT	Indicates whether or not the column is a pseudo-column. DB2 ODBC only returns: <ul style="list-style-type: none"><li>• SQL_PC_NOT_PSEUDO</li></ul> DB2 DBMSs do not support pseudo columns. ODBC applications can receive the following values from other non-IBM RDBMS servers: <ul style="list-style-type: none"><li>• SQL_PC_UNKNOWN</li><li>• SQL_PC_PSEUDO</li></ul>

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 134. SQLSpecialColumns SQLSTATEs

SQLSTATE	Description	Explanation
24000	Invalid cursor state.	A cursor is already opened on the statement handle.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1014	No more handles.	DB2 ODBC is not able to allocate a handle due to internal resources.
S1090	Invalid string or buffer length.	The value of one of the length arguments is less than 0, but not equal to SQL_NTS. The value of one of the length arguments exceeded the maximum length supported by the DBMS for that qualifier or name.
S1097	Column type out of range.	An invalid <i>fColType</i> value is specified.
S1098	Scope type out of range.	An invalid <i>fScope</i> value is specified.
S1099	Nullable type out of range.	An invalid <i>fNullable</i> values is specified.
S1C00	Driver not capable.	DB2 ODBC does not support <i>catalog</i> as a qualifier for table name.

## Restrictions

None.

## Example

```

/* ... */
SQLRETURN
list_index_columns(SQLHDBC hdbc, SQLCHAR *schema, SQLCHAR *tablename )
{
/* ... */
    rc = SQLSpecialColumns(hstmt, SQL_BEST_ROWID, NULL, 0, schema, SQL_NTS,
                          tablename, SQL_NTS, SQL_SCOPE_CURROW, SQL_NULLABLE);

    rc = SQLBindCol(hstmt, 2, SQL_C_CHAR, (SQLPOINTER) column_name.s, 129,
                   &column_name.ind);

    rc = SQLBindCol(hstmt, 4, SQL_C_CHAR, (SQLPOINTER) type_name.s, 129,
                   &type_name.ind);

    rc = SQLBindCol(hstmt, 5, SQL_C_LONG, (SQLPOINTER) & precision,
                   sizeof(precision), &precision_ind);

    rc = SQLBindCol(hstmt, 7, SQL_C_SHORT, (SQLPOINTER) & scale,
                   sizeof(scale), &scale_ind);

    printf("Primary Key or Unique Index for %s.%s\n", schema, tablename);
    /* Fetch each row, and display */
    while ((rc = SQLFetch(hstmt)) == SQL_SUCCESS) {
        printf("  %s, %s ", column_name.s, type_name.s);
        if (precision_ind != SQL_NULL_DATA) {
            printf(" (%ld", precision);
        } else {
            printf("\n");
        }
        if (scale_ind != SQL_NULL_DATA) {
            printf(", %d)\n", scale);
        } else {
            printf(")\n");
        }
    }
}
/* ... */

```

## References

- “SQLColumns - Get column information for a table” on page 117
- “SQLStatistics - Get index and statistics information for a base table” on page 344
- “SQLTables - Get table information” on page 353

## SQLStatistics - Get index and statistics information for a base table

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	
----------------	----------	------------	--

SQLStatistics() retrieves index information for a given table. It also returns the cardinality and the number of pages associated with the table and the indexes on the table. The information is returned in a result set, which can be retrieved using the same functions that are used to process a result set generated by a query.

### Syntax

```
SQLRETURN SQLStatistics (SQLHSTMT hstmt,
                        SQLCHAR FAR *szCatalogName,
                        SQLSMALLINT cbCatalogName,
                        SQLCHAR FAR *szSchemaName,
                        SQLSMALLINT cbSchemaName,
                        SQLCHAR FAR *szTableName,
                        SQLSMALLINT cbTableName,
                        SQLUSMALLINT fUnique,
                        SQLUSMALLINT fAccuracy);
```

### Function arguments

Table 135 (Page 1 of 2). SQLStatistics arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	Input	Statement handle.
SQLCHAR *	szCatalogName	Input	Catalog qualifier of a 3 part table name. This must be a null pointer or a zero length string.
SQLSMALLINT	cbCatalogName	Input	Length of <i>cbCatalogName</i> . This must be set to 0.
SQLCHAR *	szSchemaName	Input	Schema qualifier of the specified table.
SQLSMALLINT	cbSchemaName	Input	Length of <i>szSchemaName</i> .
SQLCHAR *	szTableName	Input	Table name.
SQLSMALLINT	cbTableName	Input	Length of <i>cbTableName</i> .
SQLUSMALLINT	fUnique	Input	Type of index information to return: <ul style="list-style-type: none"> <li>• SQL_INDEX_UNIQUE</li> </ul> Only unique indexes are returned. <ul style="list-style-type: none"> <li>• SQL_INDEX_ALL</li> </ul> All indexes are returned.

Table 135 (Page 2 of 2). SQLStatistics arguments

Data type	Argument	Use	Description
SQLUSMALLINT	fAccuracy	Input	<p>Indicate whether the CARDINALITY and PAGES columns in the result set contain the most current information:</p> <ul style="list-style-type: none"> <li>SQL_ENSURE : This value is reserved for future use, when the application requests the most up to date statistics information. Existing applications that specify this value receive the same results as SQL_QUICK. <b>Recommendation:</b> New applications should not use this value.</li> <li>SQL_QUICK : Statistics which are readily available at the server are returned. The values might not be current, and no attempt is made to ensure that they be up to date.</li> </ul>

## Usage

SQLStatistics() returns two types of information:

- Statistics information for the table (if it is available):
  - when the TYPE column in the table below is set to SQL\_TABLE\_STAT, the number of rows in the table and the number of pages used to store the table.
  - when the TYPE column indicates an index, the number of unique values in the index, and the number of pages used to store the indexes.
- Information about each index, where each index column is represented by one row of the result set. The result set columns are given in Table 136 in the order shown; the rows in the result set are ordered by NON\_UNIQUE, TYPE, INDEX\_QUALIFIER, INDEX\_NAME and ORDINAL\_POSITION.

Since calls to SQLStatistics() in many cases map to a complex and thus expensive query against the system catalog, they should be used sparingly, and the results saved rather than repeating calls.

The VARCHAR columns of the catalog functions result set are declared with a maximum length attribute of 128 to be consistent with SQL92 limits. Since DB2 names are less than 128, the application can choose to always set aside 128 characters (plus the null-terminator) for the output buffer, or alternatively, call SQLGetInfo() with the SQL\_MAX\_CATALOG\_NAME\_LEN, SQL\_MAX\_OWNER\_SCHEMA\_LEN, SQL\_MAX\_TABLE\_NAME\_LEN, and SQL\_MAX\_COLUMN\_NAME\_LEN to determine respectively the actual lengths of the TABLE\_CAT, TABLE\_SCHEM, TABLE\_NAME, and COLUMN\_NAME columns supported by the connected DBMS.

Although new columns might be added and the names of the existing columns changed in future releases, the position of the current columns does not change.

Table 136 (Page 1 of 2). Columns returned by SQLStatistics

Column number/name	Data type	Description
1 TABLE_CAT	VARCHAR(128)	The is always null.
2 TABLE_SCHEM	VARCHAR(128)	The name of the schema containing TABLE_NAME.

## SQLStatistics

Table 136 (Page 2 of 2). Columns returned by SQLStatistics

Column number/name	Data type	Description
3 TABLE_NAME	VARCHAR(128) NOT NULL	Name of the table.
4 NON_UNIQUE	SMALLINT	Indicates whether the index prohibits duplicate values: <ul style="list-style-type: none"> <li>• SQL_TRUE if the index allows duplicate values.</li> <li>• SQL_FALSE if the index values must be unique.</li> <li>• NULL is returned if the TYPE column indicates that this row is SQL_TABLE_STAT (statistics information on the table itself).</li> </ul>
5 INDEX_QUALIFIER	VARCHAR(128)	The string is used to qualify the index name in the DROP INDEX statement. Appending a period (.) plus the INDEX_NAME results in a full specification of the index.
6 INDEX_NAME	VARCHAR(128)	The name of the index. If the TYPE column has the value SQL_TABLE_STAT, this column has the value NULL.
7 TYPE	SMALLINT NOT NULL	Indicates the type of information contained in this row of the result set: <ul style="list-style-type: none"> <li>• SQL_TABLE_STAT - Indicates this row contains statistics information on the table itself.</li> <li>• SQL_INDEX_CLUSTERED - Indicates this row contains information on an index, and the index type is a clustered index.</li> <li>• SQL_INDEX_HASHED - Indicates this row contains information on an index, and the index type is a hashed index.</li> <li>• SQL_INDEX_OTHER - Indicates this row contains information on an index, and the index type is other than clustered or hashed.</li> </ul>
8 ORDINAL_POSITION	SMALLINT	Ordinal position of the column within the index whose name is given in the INDEX_NAME column. A NULL value is returned for this column if the TYPE column has the value of SQL_TABLE_STAT.
9 COLUMN_NAME	VARCHAR(128)	Name of the column in the index. A NULL value is returned for this column if the TYPE column has the value of SQL_TABLE_STAT.
10 ASC_OR_DESC	CHAR(1)	Sort sequence for the column; A for ascending, D for descending. NULL value is returned if the value in the TYPE column is SQL_TABLE_STAT.
11 CARDINALITY	INTEGER	<ul style="list-style-type: none"> <li>• If the TYPE column contains the value SQL_TABLE_STAT, this column contains the number of rows in the table.</li> <li>• If the TYPE column value is not SQL_TABLE_STAT, this column contains the number of unique values in the index.</li> <li>• A NULL value is returned if information is not available from the DBMS.</li> </ul>
12 PAGES	INTEGER	<ul style="list-style-type: none"> <li>• If the TYPE column contains the value SQL_TABLE_STAT, this column contains the number of pages used to store the table.</li> <li>• If the TYPE column value is not SQL_TABLE_STAT, this column contains the number of pages used to store the indexes.</li> <li>• A NULL value is returned if information is not available from the DBMS.</li> </ul>
13 FILTER_CONDITION	VARCHAR(128)	If the index is a filtered index, this is the filter condition. Since DATABASE 2 servers do not support filtered indexes, NULL is always returned. NULL is also returned if TYPE is SQL_TABLE_STAT.

For the row in the result set that contains table statistics (TYPE is set to SQL\_TABLE\_STAT), the columns values of NON\_UNIQUE, INDEX\_QUALIFIER,



INDEX\_NAME, ORDINAL\_POSITION, COLUMN\_NAME, and ASC\_OR\_DESC are set to NULL. If the CARDINALITY or PAGES information cannot be determined, then NULL is returned for those columns.

**Note:** The accuracy of the information returned in the SQLERRD(3) and SQLERRD(4) fields is dependent on many factors such as the use of parameter markers and expressions within the statement. The main factor which can be controlled is the accuracy of the database statistics. That is, when the statistics were last updated, (for example, for DB2 for OS/390, the last time the RUNSTATS utility was run.)

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 137. SQLStatistics SQLSTATEs

SQLSTATE	Description	Explanation
24000	Invalid cursor state.	A cursor is already opened on the statement handle.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1014	No more handles.	DB2 ODBC is not able to allocate a handle due to internal resources.
S1090	Invalid string or buffer length.	The value of one of the name length arguments was less than 0, but not equal to SQL_NTS.  The valid of one of the name length arguments exceeded the maximum value supported for that data source. The maximum supported value can be obtained by calling the SQLGetInfo() function.
S1100	Uniqueness option type out of range.	An invalid <i>fUnique</i> value was specified.
S1101	Accuracy option type out of range.	An invalid <i>fAccuracy</i> value was specified.
S1C00	Driver not capable.	DB2 ODBC does not support <i>catalog</i> as a qualifier for table name.

## Restrictions

None.

## Example

```
/* ... */
SQLRETURN
list_stats(SQLHDBC hdbc, SQLCHAR *schema, SQLCHAR *tablename )
{
/* ... */
    rc = SQLStatistics(hstmt, NULL, 0, schema, SQL_NTS,
                      tablename, SQL_NTS, SQL_INDEX_UNIQUE, SQL_QUICK);
    rc = SQLBindCol(hstmt, 4, SQL_C_SHORT,
                    &non_unique, 2, &non_unique_ind);
    rc = SQLBindCol(hstmt, 6, SQL_C_CHAR,
                    index_name.s, 129, &index_name.ind);
    rc = SQLBindCol(hstmt, 7, SQL_C_SHORT,
                    &type, 2, &type_ind);
    rc = SQLBindCol(hstmt, 9, SQL_C_CHAR,
                    column_name.s, 129, &column_name.ind);
    rc = SQLBindCol(hstmt, 11, SQL_C_LONG,
                    &cardinality, 4, &card_ind);
    rc = SQLBindCol(hstmt, 12, SQL_C_LONG,
                    &pages, 4, &pages_ind);

    printf("Statistics for %s.%s\n", schema, tablename);

    while ((rc = SQLFetch(hstmt)) == SQL_SUCCESS)
    { if (type != SQL_TABLE_STAT)
      { printf(" Column: %-18s Index Name: %-18s\n",
              column_name.s, index_name.s);
        }
      else
      { printf(" Table Statistics:\n");
        }
      if (card_ind != SQL_NULL_DATA)
        printf(" Cardinality = %13ld", cardinality);
      else
        printf(" Cardinality = (Unavailable)");

      if (pages_ind != SQL_NULL_DATA)
        printf(" Pages = %13ld\n", pages);
      else
        printf(" Pages = (Unavailable)\n");
    }
}
/* ... */
```

## References

- “SQLColumns - Get column information for a table” on page 117
- “SQLSpecialColumns - Get special (row identifier) columns” on page 339

## SQLTablePrivileges - Get privileges associated with a table

### Purpose

Specification:	ODBC 1.0		
----------------	----------	--	--

SQLTablePrivileges() returns a list of tables and associated privileges for each table. The information is returned in an SQL result set, which can be retrieved using the same functions that are used to process a result set generated by a query.

### Syntax

```
SQLRETURN SQLTablePrivileges (SQLHSTMT      hstmt,
                              SQLCHAR FAR   *szCatalogName,
                              SQLSMALLINT   cbCatalogName,
                              SQLCHAR FAR   *szSchemaName,
                              SQLSMALLINT   cbSchemaName,
                              SQLCHAR FAR   *szTableName,
                              SQLSMALLINT   cbTableName);
```

### Function arguments

Table 138. SQLTablePrivileges arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	Input	Statement handle.
SQLCHAR *	szTableQualifier	Input	Catalog qualifier of a 3 part table name. This must be a null pointer or a zero length string.
SQLSMALLINT	cbTableQualifier	Input	Length of <i>szCatalogName</i> . This must be set to 0.
SQLCHAR *	szSchemaName	Input	Buffer that can contain a <i>pattern-value</i> to qualify the result set by schema name.
SQLSMALLINT	cbSchemaName	Input	Length of <i>szSchemaName</i> .
SQLCHAR *	szTableName	Input	Buffer that can contain a <i>pattern-value</i> to qualify the result set by table name.
SQLSMALLINT	cbTableName	Input	Length of <i>szTableName</i> .

The *szSchemaName* and *szTableName* arguments accept search pattern. For more information about valid search patterns, see “Input arguments on catalog functions” on page 369.

### Usage

The results are returned as a standard result set containing the columns listed in the following table. The result set is ordered by TABLE\_CAT, TABLE\_SCHEM, TABLE\_NAME, and PRIVILEGE. If multiple privileges are associated with any given table, each privilege is returned as a separate row.

Since calls to SQLTablePrivileges() in many cases map to a complex and thus expensive query against the system catalog, they should be used sparingly, and the results saved rather than repeating calls.

## SQLTablePrivileges

The VARCHAR columns of the catalog functions result set are declared with a maximum length attribute of 128 to be consistent with SQL92 limits. Since DB2 names are less than 128, the application can choose to always set aside 128 characters (plus the null-terminator) for the output buffer, or alternatively, call `SQLGetInfo()` with the `SQL_MAX_CATALOG_NAME_LEN`, `SQL_MAX_OWNER_SCHEMA_LEN`, `SQL_MAX_TABLE_NAME_LEN`, and `SQL_MAX_COLUMN_NAME_LEN` to determine respectively the actual lengths of the `TABLE_CAT`, `TABLE_SCHEM`, `TABLE_NAME`, and `COLUMN_NAME` columns supported by the connected DBMS.

Although new columns might be added and the names of the existing columns changed in future releases, the position of the current columns does not change.

Table 139. Columns returned by `SQLTablePrivileges`

Column Number/Name	Data Type	Description
1 TABLE_CAT	VARCHAR(128)	The is always null.
2 TABLE_SCHEM	VARCHAR(128)	The name of the schema contain TABLE_NAME.
3 TABLE_NAME	VARCHAR(128) NOT NULL	The name of the table.
4 GRANTOR	VARCHAR(128)	Authorization ID of the user who granted the privilege.
5 GRANTEE	VARCHAR(128)	Authorization ID of the user to whom the privilege is granted.
6 PRIVILEGE	VARCHAR(128)	The table privilege. This can be one of the following strings: <ul style="list-style-type: none"><li>• ALTER</li><li>• CONTROL</li><li>• DELETE</li><li>• INDEX</li><li>• INSERT</li><li>• REFERENCES</li><li>• SELECT</li><li>• UPDATE</li></ul>
7 IS_GRANTABLE	VARCHAR(3)	Indicates whether the grantee is permitted to grant the privilege to other users.  This can be "YES", "NO" or NULL.

**Note:** The column names used by DB2 ODBC follow the X/Open CLI CAE specification style. The column types, contents and order are identical to those defined for the `SQLProcedures()` result set in ODBC.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 140 (Page 1 of 2). `SQLTablePrivileges SQLSTATEs`

SQLSTATE	Description	Explanation
24000	Invalid cursor state.	A cursor is already opened on the statement handle.

Table 140 (Page 2 of 2). SQLTablePrivileges SQLSTATEs

SQLSTATE	Description	Explanation
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1014	No more handles.	DB2 ODBC is not able to allocate a handle due to internal resources.
S1090	Invalid string or buffer length.	The value of one of the name length arguments is less than 0, but not equal to SQL_NTS.  The value of one of the name length arguments exceeded the maximum value supported for that data source. The maximum supported value can be obtained by calling the SQLGetInfo() function.
S1C00	Driver not capable.	DB2 ODBC does not support <i>catalog</i> as a qualifier for table name.

## Restrictions

None.

### Example

```
/* ... */
SQLRETURN
list_table_privileges(SQLHDBC hdbc, SQLCHAR *schema,
                      SQLCHAR *tablename )
{
    SQLHSTMT      hstmt;
    SQLRETURN      rc;
    struct { SQLINTEGER ind; /* Length & Indicator variable */
            SQLCHAR  s[129]; /* String variable */
    } grantor, grantee, privilege;

    struct { SQLINTEGER ind;
            SQLCHAR  s[4];
    } is_grantable;

    SQLCHAR      cur_name[512] = ""; /* Used when printing the */
    SQLCHAR      pre_name[512] = ""; /* Result set */

    /* Allocate a statment handle to reference the result set */
    rc = SQLAllocStmt(hdbc, &hstmt);

    /* Create Table Privilges result set */
    rc = SQLTablePrivileges(hstmt, NULL, 0, schema, SQL_NTS,
                            tablename, SQL_NTS);

    rc = SQLBindCol(hstmt, 4, SQL_C_CHAR, (SQLPOINTER) grantor.s, 129,
                    &grantor.ind);

    /* Continue Binding, then fetch and display result set */
    /* ... */
}
```

### References

- “SQLTables - Get table information” on page 353

## SQLTables - Get table information

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	
----------------	----------	------------	--

SQLTables() returns a list of table names and associated information stored in the system catalog of the connected data source. The list of table names is returned as a result set, which can be retrieved using the same functions that are used to process a result set generated by a query.

### Syntax

```
SQLRETURN SQLTables(
    (SQLHSTMT hstmt,
    SQLCHAR FAR *szCatalogName,
    SQLSMALLINT cbCatalogName,
    SQLCHAR FAR *szSchemaName,
    SQLSMALLINT cbSchemaName,
    SQLCHAR FAR *szTableName,
    SQLSMALLINT cbTableName,
    SQLCHAR FAR *szTableType,
    SQLSMALLINT cbTableType);
```

### Function arguments

Table 141. SQLTables arguments

Data type	Argument	Use	Description
SQLHSTMT	hstmt	Input	Statement handle.
SQLCHAR *	szCatalogName	Input	Buffer that can contain a <i>pattern-value</i> to qualify the result set. <i>Catalog</i> is the first part of a 3 part table name. This must be a NULL pointer or a zero length string.
SQLSMALLINT	cbCatalogName	Input	Length of <i>szCatalogName</i> . This must be set to 0.
SQLCHAR *	szSchemaName	Input	Buffer that can contain a <i>pattern-value</i> to qualify the result set by schema name.
SQLSMALLINT	cbSchemaName	Input	Length of <i>szSchemaName</i> .
SQLCHAR *	szTableName	Input	Buffer that can contain a <i>pattern-value</i> to qualify the result set by table name.
SQLSMALLINT	cbTableName	Input	Length of <i>szTableName</i> .
SQLCHAR *	szTableType	Input	Buffer that can contain a <i>value list</i> to qualify the result set by table type.  The value list is a list of upper-case comma-separated single quoted values for the table types of interest. Valid table type identifiers can include: TABLE, VIEW, SYSTEM TABLE, ALIAS, SYNONYM. If <i>szTableType</i> argument is a NULL pointer or a zero length string, then this is equivalent to specifying all of the possibilities for the table type identifier.  If SYSTEM TABLE is specified, then both system tables and system views (if there are any) are returned.
SQLSMALLINT	cbTableType	Input	Size of <i>cbTableType</i>

Note that the *szCatalogName*, *szSchemaName*, and *szTableName* arguments accept search patterns. For more information about valid search patterns, see “Input arguments on catalog functions” on page 369.

## Usage

Table information is returned in a result set where each table is represented by one row of the result set. To determine the type of access permitted on any given table in the list, the application can call `SQLTablePrivileges()`. Otherwise, the application must be able to handle a situation where the user selects a table for which `SELECT` privileges are not granted.

To support obtaining just a list of schemas, the following special semantics for the *szSchemaName* argument can be applied: if *szSchemaName* is a string containing a single percent (%) character, and *szCatalogName* and *szTableName* are empty strings, then the result set contains a list of valid schemas in the data source.

If *szTableType* is a single percent character (%) and *szCatalogName*, *szSchemaName*, and *szTableName* are empty strings, then the result set contains a list of valid table types for the data source. (All columns except the `TABLE_TYPE` column contain NULLs.)

If *szTableType* is not an empty string, it must contain a list of upper-case, comma-separated values for the types of interest; each value can be enclosed in single quotes or unquoted. For example, `"TABLE','VIEW"` or `"TABLE,VIEW"`. If the data source does not support or does not recognize a specified table type, nothing is returned for that type.

Sometimes, an application calls `SQLTables()` with null pointers for some or all of the *szSchemaName*, *szTableName*, and *szTableType* arguments so that no attempt is made to restrict the result set returned. For some data sources that contain a large quantity of tables, views, or aliases, this scenario maps to an extremely large result set and very long retrieval times. Three mechanisms are introduced to help the end user reduce the long retrieval times: three keywords (`SCHEMALIST`, `SYSCHEMA`, `TABLETYPE`) can be specified in the DB2 ODBC initialization file to help restrict the result set when the application has supplied null pointers for either or both of *szSchemaName* and *szTableType*. These keywords and their usage are discussed in detail in “Initialization keywords” on page 63. If the application did not specify a null pointer for *szSchemaName* or *szTableType* then the associated keyword specification in the DB2 ODBC initialization file is ignored.

The result set returned by `SQLTables()` contains the columns listed in Table 142 on page 355 in the order given. The rows are ordered by `TABLE_TYPE`, `TABLE_CAT`, `TABLE_SCHEM`, and `TABLE_NAME`.

Since calls to `SQLTables()` in many cases map to a complex and thus expensive query against the system catalog, they should be used sparingly, and the results saved rather than repeating calls.

The `VARCHAR` columns of the catalog functions result set are declared with a maximum length attribute of 128 to be consistent with SQL92 limits. Since DB2 names are less than 128, the application can choose to always set aside 128 characters (plus the null-terminator) for the output buffer, or alternatively, call `SQLGetInfo()` with the `SQL_MAX_CATALOG_NAME_LEN`, `SQL_MAX_OWNER_SCHEMA_LEN`, `SQL_MAX_TABLE_NAME_LEN`, and



SQL\_MAX\_COLUMN\_NAME\_LEN to determine respectively the actual lengths of the TABLE\_CAT, TABLE\_SCHEM, TABLE\_NAME, and COLUMN\_NAME columns supported by the connected DBMS.

Although new columns might be added and the names of the existing columns changed in future releases, the position of the current columns does not change.

Table 142. Columns returned by SQLTables

Column Name	Data type	Description
TABLE_CAT	VARCHAR(128)	The name of the catalog containing TABLE_SCHEM. This column contains a NULL value.
TABLE_SCHEM	VARCHAR(128)	The name of the schema containing TABLE_NAME.
TABLE_NAME	VARCHAR(128)	The name of the table, or view, or alias, or synonym.
TABLE_TYPE	VARCHAR(128)	Identifies the type given by the name in the TABLE_NAME column. It can have the string values 'TABLE', 'VIEW', 'INOPERATIVE VIEW', 'SYSTEM TABLE', 'ALIAS', or 'SYNONYM'.
REMARKS	VARCHAR(254)	Contains the descriptive information about the table.

## Return codes

- SQL\_SUCCESS
- SQL\_SUCCESS\_WITH\_INFO
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 143. SQLTables SQLSTATEs

SQLSTATE	Description	Explanation
24000	Invalid cursor state.	A cursor is already opened on the statement handle.
40003 08S01	Communication link failure.	The communication link between the application and data source fails before the function completes.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1010	Function sequence error.	The function is called while in a data-at-execute (SQLParamData(), SQLPutData()) operation.
S1014	No more handles.	DB2 ODBC is not able to allocate a handle due to internal resources.
S1090	Invalid string or buffer length.	The value of one of the name length arguments is less than 0, but not equal to SQL_NTS.  The value of one of the name length arguments exceeded the maximum value supported for that data source. The maximum supported value can be obtained by calling the SQLGetInfo() function.
S1C00	Driver not capable.	DB2 ODBC does not support <i>catalog</i> as a qualifier for table name.

## SQLTables

## Restrictions

None.

## Example

Also, see “Querying environment information example” on page 46.

```
/* ... */
SQLRETURN init_tables(SQLHDBC hdbc )
{
    SQLHSTMT      hstmt;
    SQLRETURN      rc;

    SQLUSMALLINT   rowstat[MAX_TABLES];
    SQLUIINTEGER   pcrow;

    rc = SQLAllocStmt(hdbc, &hstmt);

    /* SQL_ROWSET_SIZE sets the max number of result rows to fetch each time */
    rc = SQLSetStmtOption(hstmt, SQL_ROWSET_SIZE, MAX_TABLES);

    /* Set Size of One row, Used for Row-Wise Binding Only */
    rc = SQLSetStmtOption(hstmt, SQL_BIND_TYPE, sizeof(table) / MAX_TABLES);

    printf("Enter Search Pattern for Table Schema Name:\n");
    gets(table->schem);
    printf("Enter Search Pattern for Table Name:\n");
    gets(table->name);
    rc = SQLTables(hstmt, NULL, 0, table->schem, SQL_NTS,
                  table->name, SQL_NTS, NULL, 0);

    rc = SQLBindCol(hstmt, 2, SQL_C_CHAR, (SQLPOINTER) &table->schem, 129,
                   &table->schem_1);

    rc = SQLBindCol(hstmt, 3, SQL_C_CHAR, (SQLPOINTER) &table->name, 129,
                   &table->name_1);

    rc = SQLBindCol(hstmt, 4, SQL_C_CHAR, (SQLPOINTER) &table->type, 129,
                   &table->type_1);

    rc = SQLBindCol(hstmt, 5, SQL_C_CHAR, (SQLPOINTER) &table->remarks, 255,
                   &table->remarks_1);

    /* Now fetch the result set */
    /* ... */
}
```

## References

- “SQLColumns - Get column information for a table” on page 117
- “SQLTablePrivileges - Get privileges associated with a table” on page 349

## SQLTransact - Transaction management

### Purpose

Specification:	ODBC 1.0	X/OPEN CLI	ISO CLI
----------------	----------	------------	---------

SQLTransact() commits or rolls back the current transaction in the specified connection. SQLTransact() can also be used to request that a commit or rollback be issued for each of the connections associated with the environment.

All changes to the database performed on the connection since connect time or the previous call to SQLTransact() (whichever is the most recent) are committed or rolled back.

If a transaction is active on a connection, the application must call SQLTransact() before it can disconnect from the database.

### Syntax

```
SQLRETURN SQLTransact (SQLHENV     henv,
                       SQLHDBC     hdbc,
                       SQLUSMALLINT fType);
```

### Function arguments

Table 144. SQLTransact arguments

Data type	Argument	Use	Description
SQLHENV	<i>henv</i>	input	Environment handle. If <i>hdbc</i> is a valid connection handle, <i>henv</i> is ignored.
SQLHDBC	<i>hdbc</i>	input	Database connection handle. If <i>hdbc</i> is set to SQL_NULL_HDBC, then <i>henv</i> must contain the environment handle that the connection is associated with.
SQLUSMALLINT	<i>fType</i>	input	The desired action for the transaction. The value for this argument must be one of: <ul style="list-style-type: none"> <li>• SQL_COMMIT</li> <li>• SQL_ROLLBACK</li> </ul>

### Usage

In DB2 ODBC, a transaction begins implicitly when an application that does not already have an active transaction, issues SQLPrepare(), SQLExecDirect(), SQLExecDirect(), SQLGetTypeInfo(), or one of the catalog functions. The transaction ends when the application calls SQLTransact() or disconnects from the data source.

If the input connection handle is SQL\_NULL\_HDBC and the environment handle is valid, then a commit or rollback is issued on each of the open connections in the environment. SQL\_SUCCESS is returned only if success is reported on all the connections. If the commit or rollback fails for one or more of the connections, SQLTransact() returns SQL\_ERROR. To determine which connections failed the

commit or rollback operation, the application needs to call `SQLError()` on each connection handle in the environment.

It is important to note that unless the connection option `SQL_CONNECTTYPE` is set to `SQL_COORDINATED_TRANS` (to indicate coordinated distributed transactions), there is no attempt to provide coordinated global transaction with one-phase or two-phase commit protocols.

Completing a transaction has the following effects:

- Prepared SQL statements (using `SQLPrepare()`) survive transactions; they can be executed again without first calling `SQLPrepare()`.
- Cursor positions are maintained after a commit unless one or more of the following is true:
  - The server is DB2 Server for VSE & VM.
  - The `SQL_CURSOR_HOLD` statement option for this handle is set to `SQL_CURSOR_HOLD_OFF`.
  - The `CURSORHOLD` keyword in the DB2 ODBC initialization file is set so that cursor with hold is not in effect and this has not been overridden by resetting the `SQL_CURSOR_HOLD` statement option.
  - The `CURSORHOLD` keyword is present in a the connection string on the `SQLDriverConnect()` call that set up this connection, and it indicates cursor with hold is not in effect, and this has not been overridden by resetting the `SQL_CURSOR_HOLD` statement option.

If the cursor position is not maintained due to any one of the above circumstances, the cursor is closed and all pending results are discarded.

If the cursor position is maintained after a commit, the application must issue a fetch to re-position the cursor (to the next row) before continuing with processing of the remaining result set.

To determine whether cursor position is maintained after a commit, call `SQLGetInfo()` with the `SQL_CURSOR_COMMIT_BEHAVIOR` information type.

- Cursors are closed after a rollback and all pending results are discarded.
- Statement handles are still valid after a call to `SQLTransact()`, and can be reused for subsequent SQL statements or de-allocated by calling `SQLFreeStmt()`.
- Cursor names, bound parameters, and column bindings survive transactions.

If no transaction is currently active on the connection, calling `SQLTransact()` has no effect on the database server and returns `SQL_SUCCESS`.

`SQLTransact()` can fail while executing the `COMMIT` or `ROLLBACK` due to a loss of connection. In this case the application might not be able to determine whether the `COMMIT` or `ROLLBACK` was processed, and a database administrator's help might be required. See the DBMS product information for more information on transaction logs and other transaction management tasks.

## Return codes

- SQL\_SUCCESS
- SQL\_ERROR
- SQL\_INVALID\_HANDLE

## Diagnostics

Table 145. SQLTransact SQLSTATEs

SQLSTATE	Description	Explanation
08003	Connection is closed.	The <i>hdbc</i> is not in a connected state.
08007	Connection failure during transaction.	The connection associated with the <i>hdbc</i> failed during the execution of the function and it cannot be determined whether the requested COMMIT or ROLLBACK occurred before the failure.
58004	Unexpected system failure.	Unrecoverable system error.
S1001	Memory allocation failure.	DB2 ODBC is not able to allocate memory required to support execution or completion of the function.
S1012	Invalid transaction code.	The value specified for the argument <i>fType</i> was neither SQL_COMMIT not SQL_ROLLBACK.
S1013	Unexpected memory handling error.	DB2 ODBC is not able to access memory required to support execution or completion of the function.

## Restrictions

SQLTransact() can not be issued if the application is executing as a stored procedure.

## Example

See “Example” on page 173.

## References

- “SQLSetStmtOption - Set statement option” on page 334
- “SQLGetInfo - Get general information” on page 219



---

## Chapter 6. Using advanced features

This section covers a series of advanced tasks.

- “Environment, connection, and statement options”
- “Distributed unit of work (coordinated distributed transactions)” on page 363
- “Querying system catalog information” on page 368
- “Sending/retrieving long data in pieces” on page 373
- “Using arrays to input parameter values” on page 375
- “Retrieving a result set into an array” on page 379
- “Using large objects” on page 383
- “Using distinct types” on page 386
- “Using stored procedures” on page 389
- “Writing multithreaded applications” on page 394
- “Mixing embedded SQL and DB2 ODBC” on page 401
- “Using vendor escape clauses” on page 404
- “Programming hints and tips” on page 407

---

### Environment, connection, and statement options

Environments, connections, and statements each have a defined set of options (or attributes). All attributes can be queried by the application, but only some attributes can be changed from their default values. By changing attribute values, the application can change the behavior of DB2 ODBC.

An environment handle has attributes which affect the behavior of DB2 ODBC functions under that environment. The application can specify the value of an attribute by calling `SQLSetEnvAttr()` and can obtain the current attribute value by calling `SQLGetEnvAttr()`. `SQLSetEnvAttr()` can only be called before connection handles have been allocated.

A connection handle has options which affect the behavior of DB2 ODBC functions under that connection. Of the options that can be changed:

- Some can be set any time after the connection handle is allocated.
- Some can be set only before the actual connection is established.
- Some can be set only after the connection is established.
- Some can be set after the connection is established, but only while there are no outstanding transactions or open cursors.

The application can change the value of connection options by calling `SQLSetConnectOption()` and can obtain the current value of an option by calling `SQLGetConnectOption()`. An example of a connection option which can be set any time after a handle is allocated is the auto-commit option introduced in “Commit or rollback” on page 32. For complete details on when each option can be set, see “`SQLSetConnectOption` - Set connection option” on page 316.

A statement handle has options which affect the behavior of ODBC functions executed using that statement handle. Of the statement options that can be changed:

- Some options can be set, but currently can be set to only one specific value.
- Some options can be set any time after the statement handle is allocated.

- Some options can only be set if there is no open cursor on that statement handle.

The application can specify the value of any settable statement option by calling `SQLSetStmtOption()`, and can obtain the current value of an option by calling `SQLGetStmtOption()`. For complete details on when each option can be set, see “`SQLSetStmtOption` - Set statement option” on page 334.

The `SQLSetConnectOption()` function can also be used to set statement options for all statement handles currently associated with the connection as well as for all future statement handles to be allocated under this connection. However, `SQLGetConnectOption()` can only be used to obtain connection options; `SQLGetStmtOption()` must be used to obtain the current value of a statement option.

Many applications use just the default option settings; however, there can be situations where some of these defaults are not suitable for a particular user of the application. DB2 ODBC provides end users with two methods to change some of these default values at run time. The first method is to specify the new default attribute values in the connection string input to the `SQLDriverConnect()` function. The second method involves the specification of the new default attribute values in a DB2 ODBC initialization file.

The DB2 ODBC initialization file can be used to change default values for all DB2 ODBC applications. This might be the end user's only means of changing the defaults if the application does not have a way for the user to provide default attribute values in the `SQLDriverConnect()` connection string. Default attribute values that are specified on `SQLDriverConnect()` override the values in the DB2 ODBC initialization file for that particular connection. For information on how the end user can use the DB2 ODBC initialization file as well as for a list of changeable defaults, see “DB2 ODBC initialization file” on page 60.

The mechanisms for changing defaults are intended for end user tuning; application developers must use the appropriate set-option function. If an application does call a set-option or attribute function with a value different from the initialization file or the connection string specification, then the initial default value is overridden and the new value takes effect.

The options that can be changed, are listed in the detailed function descriptions of the *set* option or attributes functions, see “Chapter 5. Functions” on page 73. The read-only options (if any exist) are listed with the detailed function descriptions of the *get* option or attribute functions.

For information on some commonly used options, see “Programming hints and tips” on page 407.

Figure 7 on page 363 shows the addition of the option or attribute functions to the basic connect scenario.



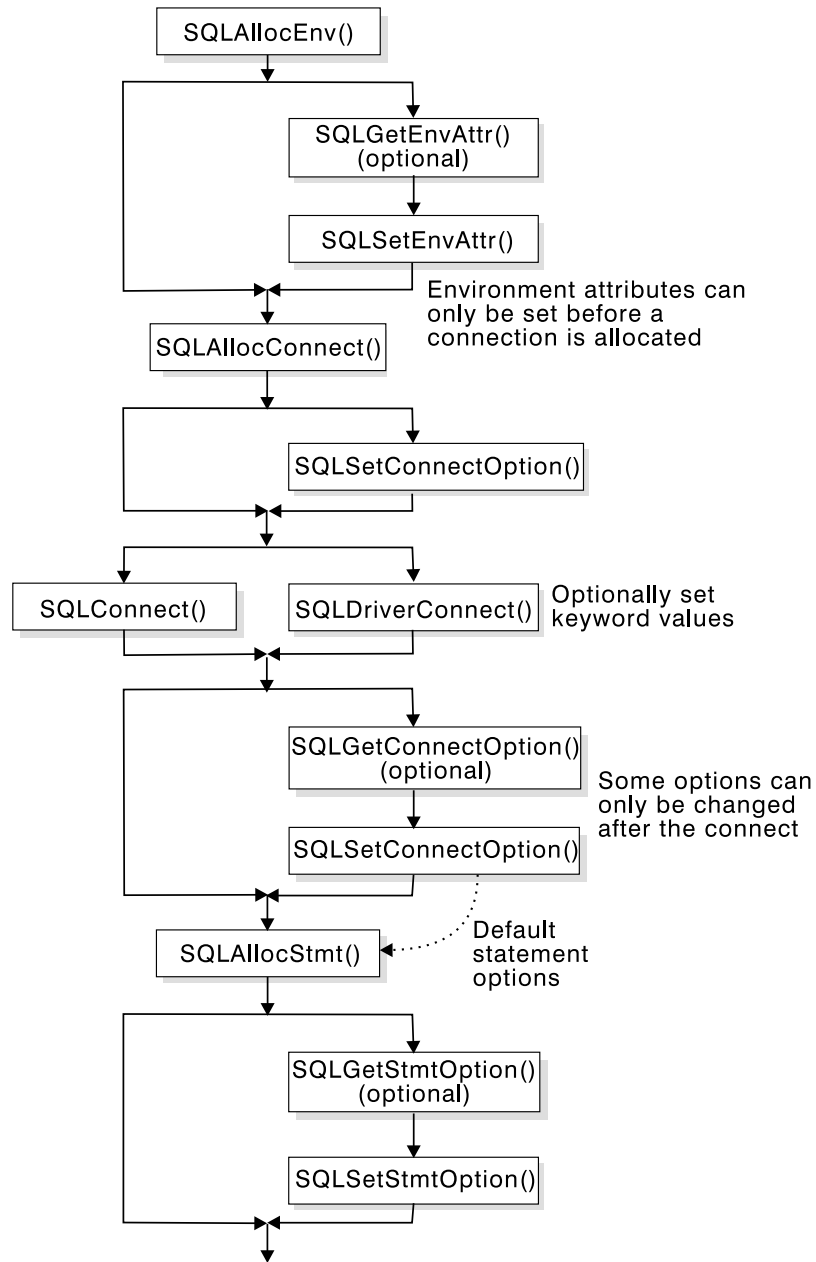


Figure 7. Setting and retrieving options (attributes)

## Distributed unit of work (coordinated distributed transactions)

The transaction scenario described in “Connecting to one or more data sources” on page 24 portrays an application which interacts with only one database server in a transaction. Further, only one transaction (that associated with the current server) existed at any given time.

With distributed unit of work (coordinated distributed transactions), the application, if executing CONNECT (type 2), is able to access multiple database servers from within the same coordinated transaction. This section describes how DB2 ODBC applications can be written to use coordinated distributed unit of work.

First, consider the environment attribute (SQL\_CONNECTTYPE) which controls whether the application is to operate in a coordinated or uncoordinated distributed environment. The two possible values for this attribute are:

- SQL\_CONCURRENT\_TRANS - supports the single data source per transaction semantics described in Chapter 2. Multiple (logical) concurrent connections to different data sources are permitted. This is the default.
- SQL\_COORDINATED\_TRANS - supports the multiple data sources per transaction semantics, as discussed below.

All connections within an application must have the same SQL\_CONNECTTYPE setting.

**Recommendation:** Have the application set this environment attribute, if necessary, as soon as `SQLAllocEnv()` is called successfully.

## Options that govern distributed unit of work semantics

A coordinated transaction means that commits or rollbacks among multiple data source connections are coordinated. The SQL\_COORDINATED\_TRANS setting of the SQL\_CONNECTTYPE attribute corresponds to the CONNECT (type 2) in IBM embedded SQL.

All the connections within an application must have the same SQL\_CONNECTTYPE setting. After the first connection is established, all subsequent connect types must be the same as the first. Coordinated connections default to manual-commit mode (for discussion on auto-commit mode, see “Commit or rollback” on page 32).

Figure 8 on page 365 shows the logical flow of an application executing statements on two SQL\_CONCURRENT\_TRANS connections ('A' and 'B'), and indicates the scope of the transactions.

Figure 9 on page 366 shows the same statements being executed on two SQL\_COORDINATED\_TRANS connections (A and B), and the scope of a coordinated distributed transaction.

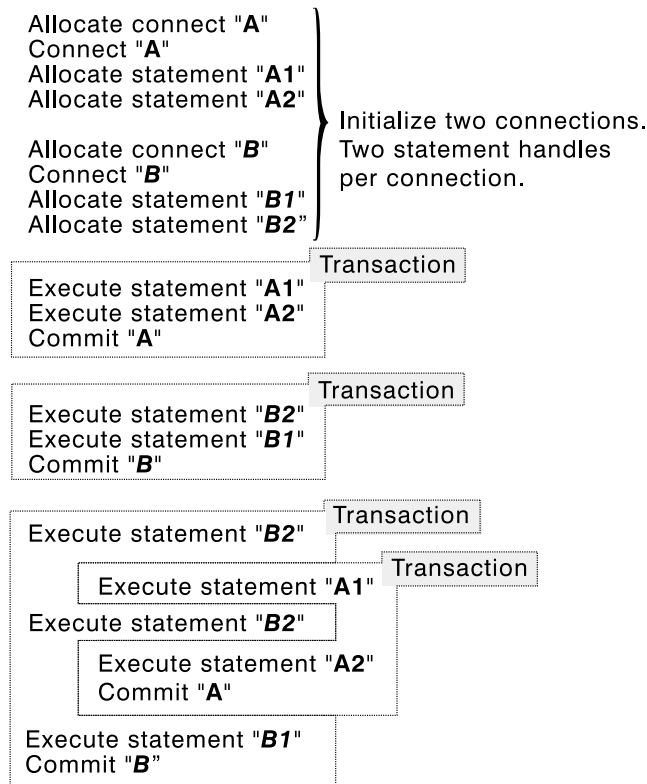


Figure 8. Multiple connections with concurrent transactions

In Figure 8, within the context of the ODBC connection model, the third and fourth transactions can be interleaved as shown. That is, if the application has specified `SQL_CONCURRENT_TRANS`, then the ODBC model supports one transaction for each active connection. The third transaction, consisting of the execution of statements B2, B2 again and B1 at data source B, can be managed and committed independent of the fourth transaction, consisting of the execution of statements A1 and A2 at data source A. That is, the transactions at A and B are independent and exist concurrently.

If the application specifies `SQL_CONCURRENT_TRANS` and is executing with `MULTICONTEXT=0` specified in the initialization file, then DB2 for OS/390 allows any number of concurrent connection handles to be allocated, subject to the restriction that only one physical connection can exist at any given time. This behavior precludes support for the ODBC connection model, and consequently the behavior of the application is substantially different (than that described for the ODBC execution model described above.)

In particular, the third transaction is executed as three transactions. Prior to executing statement 'B2' DB2 ODBC connects to B. This statement is executed and committed prior to reconnecting to data source A to execute A1. Similarly, this statement at data source A is committed prior to reconnecting to data source B to execute statement B2. This statement is then committed and a reconnection is made to A to execute A2. Next, another commit occurs and a reconnection to B to execute B1.

From an application point of view, the transaction at data source B, consisting of B2->B2->B1, is broken into three independent transactions: B2, B2 and B1. The

fourth transaction at data source A, consisting of A1->A2, is broken into two independent transactions: A1 and A2.

```

Allocate environment
Set environment attributes
(SQL_CONNECTTYPE)

Allocate connect "A"
Connect "A"
(SQL_CONCURRENT_TRANS)

Allocate statement "A1"
Allocate statement "A2"

Allocate connect "B"
Connect "B"
(SQL_CONCURRENT_TRANS)

Allocate statement "B1"
Allocate statement "B2"

```

} Initialize two connections.  
Two statement handles  
per connection.

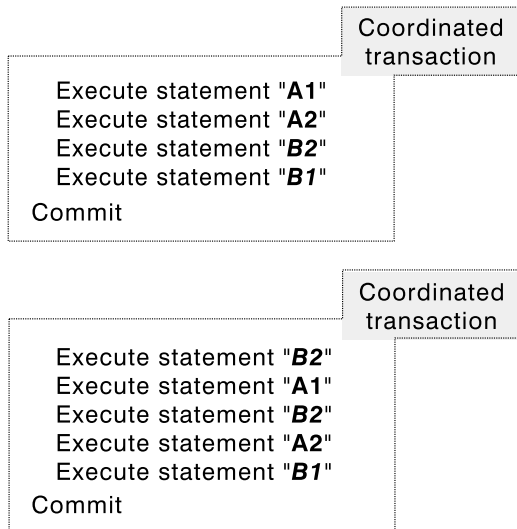


Figure 9. Multiple connections with coordinated transactions

For a discussion of multiple active transaction support, see "DB2 ODBC support of multiple contexts" on page 396.

## Establishing a coordinated transaction connection

An application can establish coordinated transaction connections by calling the `SQLSetEnvAttr()` or `SQLSetConnectOption()` functions, or by setting the `CONNECTTYPE` keyword in the DB2 ODBC initialization file or in the connection string for `SQLDriverConnect()`. The initialization file is intended for existing applications that do not use the `SQLSetConnectOption()` function. For information about the keywords, see "DB2 ODBC initialization file" on page 60.

An application cannot have a mixture of concurrent and coordinated connections; the type of the first connection determines the type of all subsequent connections. `SQLSetEnvAttr()` and `SQLSetConnectOption()` return an error if an application attempts to change the connect type while there is an active connection. When the connection type is established, it persists until `SQLFreeEnv` is called.

## Distributed unit of work example

The following example connects to two data sources using a SQL\_CONNECTTYPE set to SQL\_COORDINATED\_TRANS (CONNECT (type 2)).

```
/* ... */
#define MAX_CONNECTIONS 2

int
DBconnect(SQLHENV henv,
          SQLHDBC * hdbc,
          char * server);

int
main()
{
    SQLHENV      henv;
    SQLHDBC     hdbc[MAX_CONNECTIONS];
    SQLRETURN    rc;

    char *      svr[MAX_CONNECTIONS] =
        {
            "KARACHI" ,
            "DAMASCUS"
        }

    /* allocate an environment handle */
    SQLAllocEnv(&henv);

    /* Before allocating any connection handles, set Environment wide
       Connect Options */
    /* Set to Connect type 2 */
    rc = SQLSetEnvAttr(henv, SQL_CONNECTTYPE,
                      (SQLPOINTER) SQL_COORDINATED_TRANS, 0);
/* ... */
/* Connect to first data source */
/* allocate a connection handle */
if (SQLAllocConnect(henv, &hdbc[0]) != SQL_SUCCESS) {
    printf(">---ERROR while allocating a connection handle-----\n");
    return (SQL_ERROR);
}

/* Connect to first data source (Type-II) */

DBconnect (henv,
          &hdbc[0],
          svr[0]);

/* allocate a second connection handle */
if (SQLAllocConnect(henv, &hdbc[1]) != SQL_SUCCESS) {
    printf(">---ERROR while allocating a connection handle-----\n");
    return (SQL_ERROR);
}
/* Connect to second data source (Type-II) */

DBconnect (henv,
          &hdbc[1],
          svr[1]);
```

```

/***** Start Processing Step *****/
/* allocate statement handle, execute statement, etc. */
/* Note that both connections participate in the disposition*/
/* of the transaction. Note that a NULL connection handle */
/* is passed as all work is committed on all connections. */
/***** End Processing Step *****/

(void) SQLTransact (henv,
                  SQL_NULL_HDBC,
                  SQL_COMMIT);

/* Disconnect, free handles and exit */
}

/*****
** Server is passed as a parameter. Note that USERID and PASSWORD**
** are always NULL. **
*****/

int
DBconnect(SQLHENV henv,
          SQLHDBC * hdbc,
          char * server)
{
    SQLRETURN rc;
    SQLCHAR buffer[255];
    SQLSMALLINT outlen;

    SQLAllocConnect(henv, hdbc);/* allocate a connection handle */

    rc = SQLConnect(*hdbc, server, SQL_NTS, NULL, SQL_NTS, NULL, SQL_NTS);
    if (rc != SQL_SUCCESS) {
        printf(">--- Error while connecting to database: %s -----\n", server);
        return (SQL_ERROR);
    } else {
        printf(">Connected to %s\n", server);
        return (SQL_SUCCESS);
    }
}
/* ... */

```

---

## Querying system catalog information

Often, one of the first tasks an application performs is to display a list of tables from which the user selects one or more tables to work with. Although the application can issue its own queries against the database system catalog to get this type of catalog information, it is best for the application to either call the DB2 ODBC catalog query functions or direct the catalog query functions to the DB2 ODBC shadow catalog.

The following sections describe the methods you can use to query the system catalog.

- “Using the catalog query functions” on page 369
- “Directing catalog queries to the DB2 ODBC shadow catalog” on page 371

## Using the catalog query functions

Often, one of the first tasks an application performs is to display a list of tables to the user from which the user selects one or more tables to work with. Although the application can issue its own queries against the database system catalog to get this type of catalog information, it is best that the application calls the DB2 ODBC catalog functions instead. These catalog functions provide a generic interface to issue queries and return consistent result sets across the DB2 family of servers. In most cases, this allows the application to avoid server specific and release specific catalog queries.

The catalog functions operate by returning a result set to the application through a statement handle. Calling these functions is conceptually equivalent to using `SQLExecDirect()` to execute a `SELECT` against the system catalog tables. After calling these functions, the application can fetch individual rows of the result set as it would process column data from an ordinary `SQLFetch()`. The DB2 ODBC catalog functions are:

- “SQLColumnPrivileges - Get privileges associated with the columns of a table” on page 112
- “SQLColumns - Get column information for a table” on page 117
- “SQLForeignKeys - Get the list of foreign key columns” on page 175
- “SQLPrimaryKeys - Get primary key columns of a table” on page 285
- “SQLProcedureColumns - Get input/output parameter information for a procedure” on page 290
- “SQLProcedures - Get list of procedure names” on page 300
- “SQLSpecialColumns - Get special (row identifier) columns” on page 339
- “SQLStatistics - Get index and statistics information for a base table” on page 344
- “SQLTablePrivileges - Get privileges associated with a table” on page 349
- “SQLTables - Get table information” on page 353

The result sets returned by these functions are defined in the descriptions for each catalog function. The columns are defined in a specified order. The results sets for each API are fixed and cannot be changed. In future releases, other columns might be added to the end of each defined result set, therefore applications should be written in a way that would not be affected by such changes.

Some of the catalog functions result in execution of fairly complex queries, and for this reason should only be called when needed. It is best for the application to save the information returned rather than making repeated calls to get the same information.

### Input arguments on catalog functions

All of the catalog functions have *CatalogName* and *SchemaName* (and their associated lengths) on their input argument list. Other input arguments can also include *TableName*, *ProcedureName*, or *ColumnName* (and their associated lengths). These input arguments are used to either identify or constrain the amount of information to be returned. *CatalogName*, however, must always be a null pointer (with its length set to 0) as DB2 ODBC does not support three-part naming.

In the Function Arguments sections for these catalog functions in “Chapter 5. Functions” on page 73, each of the above input arguments are described either as a pattern-value or just as an ordinary argument. For example,

SQLColumnPrivileges() treats *SchemaName* and *TableName* as ordinary arguments and *ColumnName* as a pattern-value.

Inputs treated as ordinary arguments are taken literally and the case of letters is significant. The argument does not qualify a query but rather identifies the information desired. If the application passes a null pointer for this argument, the results can be unpredictable.

Inputs treated as pattern-values are used to constrain the size of the result set by including only matching rows as though the underlying query were qualified by a WHERE clause. If the application passes a null pointer for a pattern-value input, the argument is not used to restrict the result set (that is, there is no WHERE clause). If a catalog function has more than one pattern-value input argument, they are treated as though the WHERE clauses in the underlying query were joined by AND; a row appears in this result set only if it meets all the conditions of the WHERE clauses.

Each pattern-value argument can contain:

- The underscore ( `_` ) character which stands for any single character.
- The percent ( `%` ) character which stands for any sequence of zero or more characters.
- Characters which stand for themselves. The case of a letter is significant.

These argument values are used on conceptual LIKE predicates in the WHERE clause. To treat the metadata characters ( `_` , `%` ) as themselves, an escape character must immediately precede the `_` or `%`. The escape character itself can be specified as part of the pattern by including it twice in succession. An application can determine the escape character by calling SQLGetInfo() with SQL\_SEARCH\_PATTERN\_ESCAPE.

## Catalog functions example

In the sample application:

- A list of all tables are displayed for the specified schema (qualifier) name or search pattern.
- Column, special column, foreign key, and statistics information is returned for a selected table.

DB2 UDB is the data source in this example. Output from a sample is shown below. Relevant segments of the sample are listed for each of the catalog functions.



```

Enter Search Pattern for Table Schema Name:
STUDENT
Enter Search Pattern for Table Name:
%
### TABLE SCHEMA          TABLE_NAME          TABLE_TYPE
-----
1  STUDENT                  CUSTOMER             TABLE
2  STUDENT                  DEPARTMENT          TABLE
3  STUDENT                  EMP_ACT             TABLE
4  STUDENT                  EMP_PHOTO           TABLE
5  STUDENT                  EMP_RESUME          TABLE
6  STUDENT                  EMPLOYEE            TABLE
7  STUDENT                  NAMEID              TABLE
8  STUDENT                  ORD_CUST            TABLE
9  STUDENT                  ORD_LINE            TABLE
10 STUDENT                  ORG                 TABLE
11 STUDENT                  PROD_PARTS          TABLE
12 STUDENT                  PRODUCT             TABLE
13 STUDENT                  PROJECT             TABLE
14 STUDENT                  STAFF               TABLE
Enter a table Number and an action:(n [Q | C | P | I | F | T | O | L])
|Q=Quit   C=cols   P=Primary Key I=Index   F=Foreign Key |
|T=Tab Priv O=Col Priv S=Stats   L=List Tables   |
1c
Schema: STUDENT Table Name: CUSTOMER
  CUST_NUM, NOT NULLABLE, INTEger (10)
  FIRST_NAME, NOT NULLABLE, CHARacter (30)
  LAST_NAME, NOT NULLABLE, CHARacter (30)
  STREET, NULLABLE, CHARacter (128)
  CITY, NULLABLE, CHARacter (30)
  PROV_STATE, NULLABLE, CHARacter (30)
  PZ_CODE, NULLABLE, CHARacter (9)
  COUNTRY, NULLABLE, CHARacter (30)
  PHONE_NUM, NULLABLE, CHARacter (20)
>> Hit Enter to Continue<<

1p
Primary Keys for STUDENT.CUSTOMER
  1 Column: CUST_NUM          Primary Key Name: = NULL
>> Hit Enter to Continue<<

1f
Primary Key and Foreign Keys for STUDENT.CUSTOMER
  CUST_NUM STUDENT.ORD_CUST.CUST_NUM
  Update Rule SET NULL , Delete Rule: NO ACTION
>> Hit Enter to Continue<<

```

## Directing catalog queries to the DB2 ODBC shadow catalog

DB2 ODBC applications that frequently use ODBC catalog functions to access the DB2 system catalog can redirect catalog queries to the DB2 ODBC shadow catalog. The shadow catalog consists of a set of pseudo-catalog tables that contain rows representing objects defined in the DB2 system catalog. These tables are pre-joined and indexed to provide faster catalog access for ODBC applications. The table rows are limited to the columns necessary for supporting ODBC operations.

The tables that comprise the DB2 ODBC shadow catalog have a default schema name of CLISCHEM. To redirect catalog functions to access the base DB2 ODBC shadow catalog tables, you must add the entry CLISCHEMA=CLISCHEM to the data source section of the DB2 ODBC initialization file as follows:

```
[DATASOURCE]
MVSDEFAULTSSID=V61A
CLISCHEMA=CLISCHEM
```

Optionally, you can create views for the DB2 ODBC shadow catalog tables under your own schema name, and redirect the ODBC catalog functions to access these views instead of the base DB2 ODBC shadow catalog tables. To redirect the catalog functions to access your own set of views, add the entry `CLISCHEMA=myschema` to the data source section of the DB2 ODBC initialization file as follows:

```
[DATASOURCE]
MVSDEFAULTSSID=V61A
CLISCHEMA=PAYROLL
APPLTRACE=1
APPLTRACEFILENAME="DD:APPLTRC"
```

where `myschema` is the schema name of the set of views you create.

You can use the `CREATE VIEW` SQL statement to create views of the DB2 ODBC shadow catalog tables. To use your own set of views, you must create a view for each DB2 ODBC shadow catalog table. For example:

```
CREATE VIEW PAYROLL.table_name AS
  SELECT * FROM PAYROLL.table_name WHERE TABLE_SCHEM='USER01';
```

where `table_name` is the name of a DB2 ODBC shadow catalog table (for example, `COLUMNPRIVILEGES`).

Table 146 lists the DB2 ODBC shadow catalog tables and the catalog functions used to access the tables.

*Table 146. Shadow catalog tables and DB2 ODBC APIs*

Shadow catalog table	DB2 ODBC catalog function
CLISCHEM.COLUMNPRIVILEGES	SQLColumnPrivileges()
CLISCHEM.COLUMNS	SQLColumns()
CLISCHEM.FOREIGNKEYS	SQLForeignKeys()
CLISCHEM.PRIMARYKEYS	SQLPrimaryKeys()
CLISCHEM.PROCEDURECOLUMNS	SQLProcedureColumns()
CLISCHEM.PROCEDURES	SQLProcedures()
CLISCHEM.SPECIALCOLUMNS	SQLSpecialColumns()
CLISCHEM.TSTATISTICS	SQLStatistics()
CLISCHEM.TABLEPRIVILEGES	SQLTablePrivileges()
CLISCHEM.TABLE	SQLTables()

### **Creating and maintaining the DB2 ODBC shadow catalog**

DB2 DataPropagator for OS/390 populates and maintains the DB2 ODBC shadow catalog. DB2 for OS/390 supports the `DATA CAPTURE CHANGE` clause of the `ALTER TABLE` SQL statement for DB2 system catalog tables. This support allows DB2 to mark log records associated with any statements that change the DB2 catalog (for example, `CREATE` and `DROP`). In addition, the DB2 DataPropagator Capture and Apply process identifies and propagates the DB2 system catalog changes to the DB2 ODBC shadow based on marked log records.

For detailed information about setting up the DB2 ODBC shadow catalog and running the DB2 DataPropagator Capture and Apply programs, see *DB2 Replication Guide and Reference*.

### Shadow catalog example

If you specify `CLISHEMA=PAYROLL` in the data source section of the DB2 ODBC initialization file, the ODBC catalog functions that normally query the DB2 system catalog tables (SYSIBM schema) will reference the following set of views of the ODBC shadow catalog base tables:

- PAYROLL.COLUMNS
- PAYROLL.TABLES
- PAYROLL.COLUMNPRIVILIGES
- PAYROLL.TABLEPRIVILIGES
- PAYROLL.SPECIALCOLUMNS
- PAYROLL.PRIMARYKEYS
- PAYROLL.FOREIGNKEYS
- PAYROLL.TSTATISTICS
- PAYROLL.PROCEDURES
- PAYROLL.PROCEDURECOLUMNS

---

## Sending/retrieving long data in pieces

When manipulating long data, it might not be feasible for the application to load the entire parameter data value into storage at the time the statement is executed, or when the data is fetched from the database. A method is provided to allow the application to handle the data in pieces. The technique to send long data in pieces is called *Specifying Parameter Values at Execute Time* because it can also be used to specify values for fixed size non-character data types such as integers.

## Specifying parameter values at execute time

A bound parameter for which value is prompted at execution time instead of stored in memory before calling `SQLExecute()` or `SQLExecDirect()` is called a data-at-execute parameter. To indicate such a parameter on an `SQLBindParameter()` or `SQLSetParam()` call, the application:

- Sets the input data length pointer to point to a variable that, at execute time, contains the value `SQL_DATA_AT_EXEC`.
- If there is more than one data-at-execute parameter, sets each input data pointer argument to some value that it recognizes as uniquely identifying the field in question.

If there are any data-at-execute parameters when the application calls `SQLExecDirect()` or `SQLExecute()`, the call returns with `SQL_NEED_DATA` to prompt the application to supply values for these parameters. The application responds as follows:

1. It calls `SQLParamData()` to conceptually advance to the first such parameter. `SQLParamData()` returns `SQL_NEED_DATA` and provides the contents of the input data pointer argument specified on the associated `SQLBindParameter()` or `SQLSetParam()` call to help identify the information required.
2. It calls `SQLPutData()` to pass the actual data for the parameter. Long data can be sent in pieces by calling `SQLPutData()` repeatedly.

3. It calls `SQLParamData()` again after it has provided the entire data for this data-at-execute parameter. If more data-at-execute parameters exist, `SQLParamData()` again returns `SQL_NEED_DATA` and the application repeats steps 2 and 3 above.

When all data-at-execute parameters are assigned values, `SQLParamData()` completes execution of the SQL statement and produces a return value and diagnostics as the original `SQLExecDirect()` or `SQLExecute()` would have produced. The right side of Figure 10 on page 375 illustrates this flow.

While the data-at-execution flow is in progress, the only DB2 ODBC functions the application can call are:

- `SQLParamData()` and `SQLPutData()` as given in the sequence above.
- The `SQLCancel()` function which is used to cancel the flow and force an exit from the loops on the right side of Figure 10 on page 375 without executing the SQL statement.
- The `SQLError()` function. The application also must not end the transaction nor set any connection attributes.

## Fetching data in pieces

Typically, based on its knowledge of a column in the result set (using `SQLDescribeCol()` or prior knowledge), the application can choose to allocate the maximum memory the column value can occupy and bind it using `SQLBindCol()`. However, in the case of character and binary data, the column can be long. If the length of the column value exceeds the length of the buffer the application can allocate, or afford to allocate, a feature of `SQLGetData()` lets the application use repeated calls to obtain in sequence the value of a single column in more manageable pieces.

Basically, as shown on the left side of Figure 10 on page 375, a call to `SQLGetData()` returns `SQL_SUCCESS_WITH_INFO` (with `SQLSTATE 01004`) to indicate more data exists for this column. `SQLGetData()` is called repeatedly to get the remaining pieces of data until it returns `SQL_SUCCESS`, signifying that the entire data were retrieved for this column.

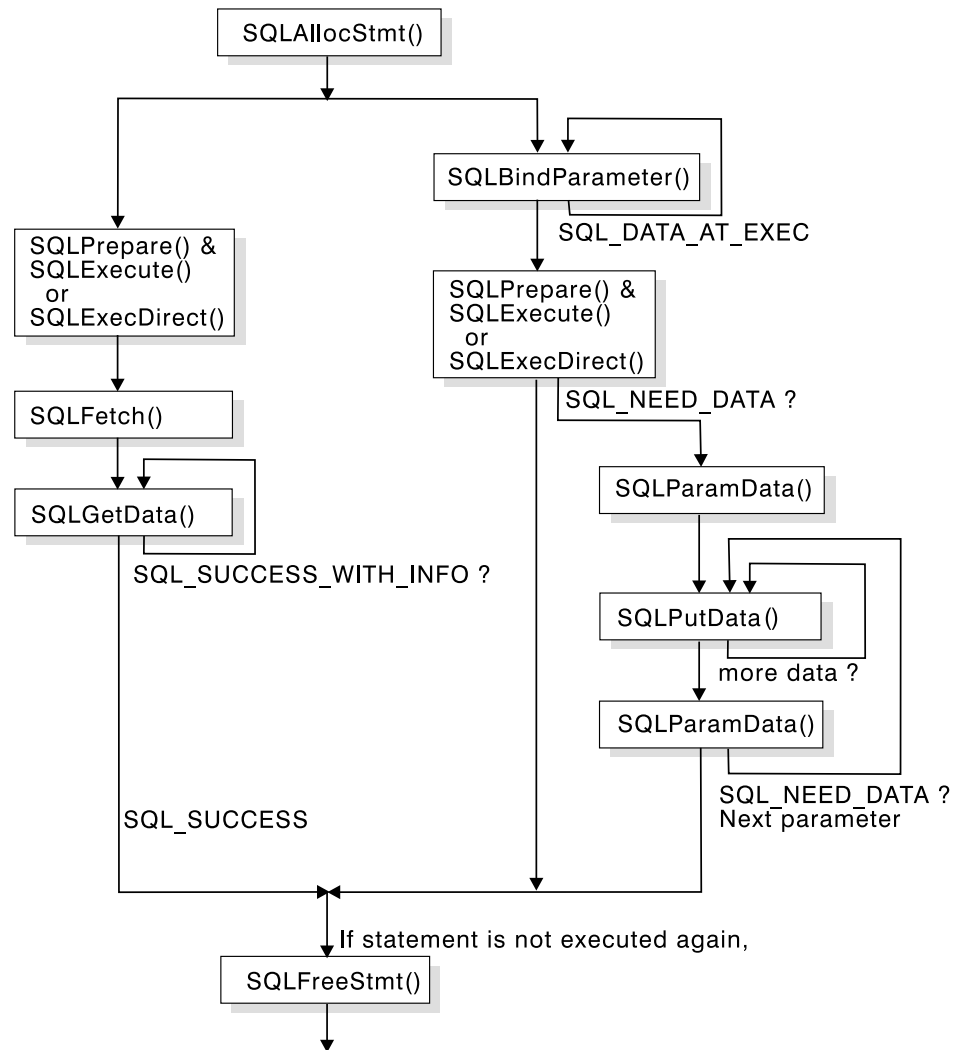


Figure 10. Input and retrieval of data in pieces

## Using arrays to input parameter values

For some data entry and update applications, users might often insert, delete, or change many cells in a data entry form and then ask for the data to be sent to the database. For these situations of bulk insert, delete, or update, DB2 ODBC provides an array input method to save the application from having to call `SQLExecute()` repeatedly on the same INSERT, DELETE, or UPDATE statement. In addition, there can be savings in network flows.

This method involves the binding of parameter markers to arrays of storage locations using the `SQLBindParameter()` call. For character and binary input data, the application uses the maximum input buffer size argument (`cbValueMax`) on `SQLBindParameter()` call to indicate to DB2 ODBC the location of values in the input array. For other input data types, the length of each element in the array is assumed to be the size of the C data type. `SQLParamOptions()` is called to specify how many elements are in the array before the execution of the SQL statement.

Suppose for Figure 11 on page 377 there is an application that allows the user to change values in the OVERTIME\_WORKED and OVERTIME\_PAID columns of a time sheet data entry form. Also suppose that the primary key of the underlying EMPLOYEE table is EMPLOY\_ID. The application can then request to prepare the following SQL statement:

```
UPDATE EMPLOYEE SET OVERTIME_WORKED= ? and OVERTIME_PAID= ?
WHERE EMPLOY_ID=?
```

When the user enters all the changes, the application counts that  $n$  rows are to change and allocates  $m=3$  arrays to store the changed data and the primary key. Then, the application makes calls to functions as follows:

1. `SQLBindParameter()` to bind the three parameter markers to the location of three arrays in memory.
2. `SQLParamOptions()` to specify the number of rows to change (the size of the array).
3. `SQLExecute()` once and all the updates are sent to the database.

This is the flow shown on the right side of Figure 11 on page 377.

The basic method is shown on the left side of Figure 11 on page 377 where `SQLBindParameter()` is called to bind the three parameter markers to the location of three variables in memory. `SQLExecute()` is called to send the first set of changes to the database. The variables are updated to reflect values for the next row of changes and again `SQLExecute()` is called. This method has  $n-1$  extra `SQLExecute()` calls. For insert and update, use `SQLRowCount` to verify the number of changed rows.

**Note:** `SQLSetParam()` must not be used to bind an array storage location to a parameter marker. In the case of character or binary input data, there is no method to specify the size of each element in the input array.

For queries with parameter markers on the WHERE clauses, an array of input values generate multiple sequential result sets. Each result set can be processed before moving onto the next one by calling `SQLMoreResults()`. See “SQLMoreResults - Determine if there are more result sets” on page 262 for more information and an example.

Figure 11 on page 377 shows the two methods of executing a statement with  $m$  parameters  $n$  times. The basic method shown on the left calls `SQLExecute()` once for each set of parameter values. The array method on the right, calls `SQLParamOptions()` to specify the number of rows ( $n$ ), then calls `SQLExecute()` once. Both methods must call `SQLBindParameter()` once for each parameter.

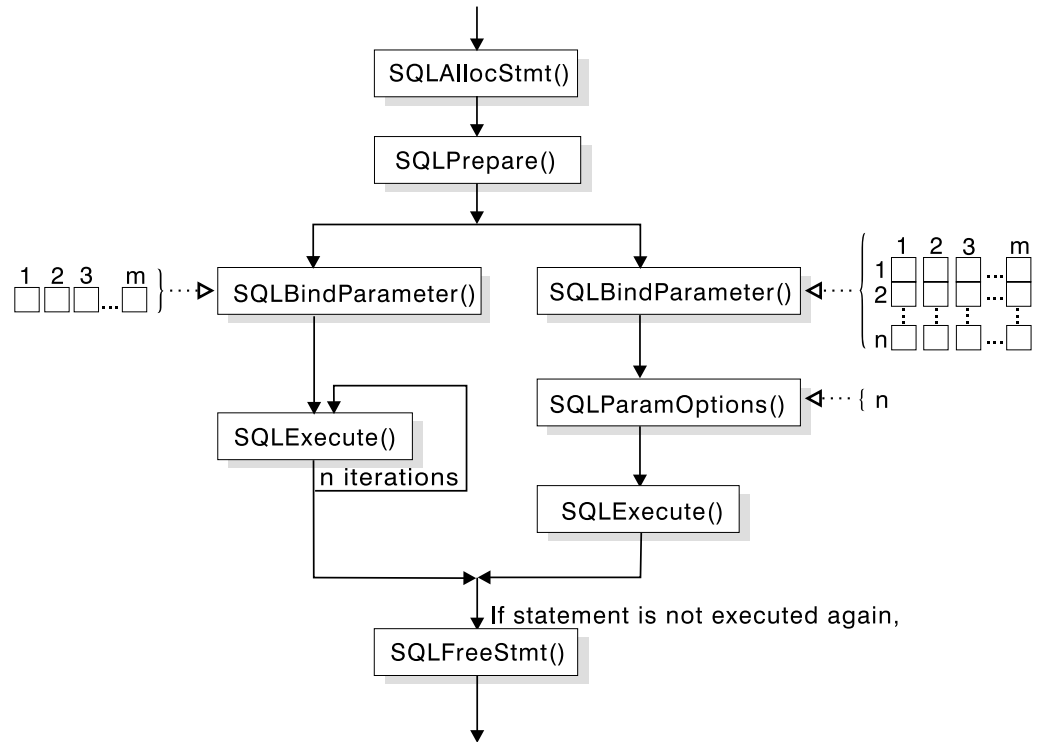


Figure 11. Array insert

## Array input example

This example shows an array INSERT statement. For an example of an array query statement, see “SQLMoreResults - Determine if there are more result sets” on page 262.

```
/* ... */
SQLINTEGER pirow = 0;
SQLCHAR      stmt[] =
    "INSERT INTO CUSTOMER ( Cust_Num, First_Name, Last_Name) "
    "VALUES (?, ?, ?)";

SQLINTEGER      Cust_Num[25] = {
    10, 20, 30, 40, 50, 60, 70, 80, 90, 100,
    110, 120, 130, 140, 150, 160, 170, 180, 190, 200,
    210, 220, 230, 240, 250
};

SQLCHAR      First_Name[25][31] = {
    "EVA",      "EILEEN",    "THEODORE", "VINCENZO", "SEAN",
    "DOLORES", "HEATHER",    "BRUCE",    "ELIZABETH", "MASATOSHI",
    "MARILYN", "JAMES",      "DAVID",    "WILLIAM",  "JENNIFER",
    "JAMES",   "SALVATORE", "DANIEL",   "SYBIL",    "MARIA",
    "ETHEL",   "JOHN",      "PHILIP",   "MAUDE",    "BILL"
};

SQLCHAR      Last_Name[25][31] = {
    "SPENSER", "LUCCHESI", "O'CONNELL", "QUINTANA",
    "NICHOLLS", "ADAMSON", "PIANKA", "YOSHIMURA",
    "SCOUTTEN", "WALKER", "BROWN", "JONES",
    "LUTZ", "JEFFERSON", "MARINO", "SMITH",
    "JOHNSON", "PEREZ", "SCHNEIDER", "PARKER",
    "SMITH", "SETRIGHT", "MEHTA", "LEE",
    "GOUNOT"
};
/* ... */
/* Prepare the statement */
rc = SQLPrepare(hstmt, stmt, SQL_NTS);

rc = SQLParamOptions(hstmt, 25, &pirow);

rc = SQLBindParameter(hstmt, 1, SQL_PARAM_INPUT, SQL_C_SLONG, SQL_INTEGER,
    0, 0, Cust_Num, 0, NULL);

rc = SQLBindParameter(hstmt, 2, SQL_PARAM_INPUT, SQL_C_CHAR, SQL_CHAR,
    31, 0, First_Name, 31, NULL);

rc = SQLBindParameter(hstmt, 3, SQL_PARAM_INPUT, SQL_C_CHAR, SQL_CHAR,
    31, 0, Last_Name, 31, NULL);

rc = SQLExecute(hstmt);
printf("Inserted %ld Rows\n", pirow);

/* ... */
```



## Retrieving a result set into an array

One of the most common tasks performed by an application is to issue a query statement, and then fetch each row of the result set into application variables that were bound using `SQLBindCol()`. If the application requires that each column or each row of the result set be stored in an array, each fetch must be followed by either a data copy operation or a new set of `SQLBindCol()` calls to assign new storage areas for the next fetch.

Alternatively, applications can eliminate the overhead of extra data copies or extra `SQLBindCol()` calls by retrieving multiple rows of data (called a rowset) at a time into an array. `SQLBindCol()` is also used to assign storage for application array variables. By default, the binding of rows is in column-wise fashion: this is symmetrical to using `SQLBindParameter()` to bind arrays of input parameter values as described in the previous section.

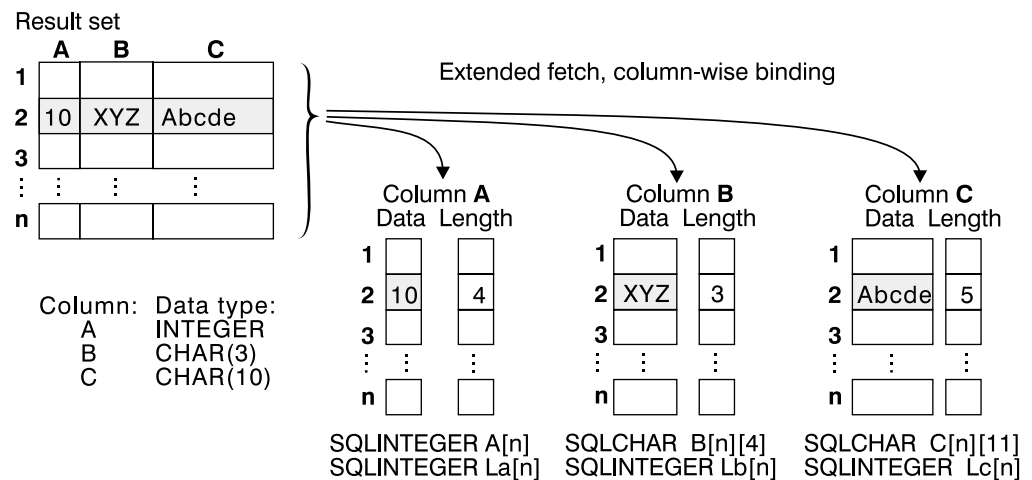


Figure 12. Column-wise binding

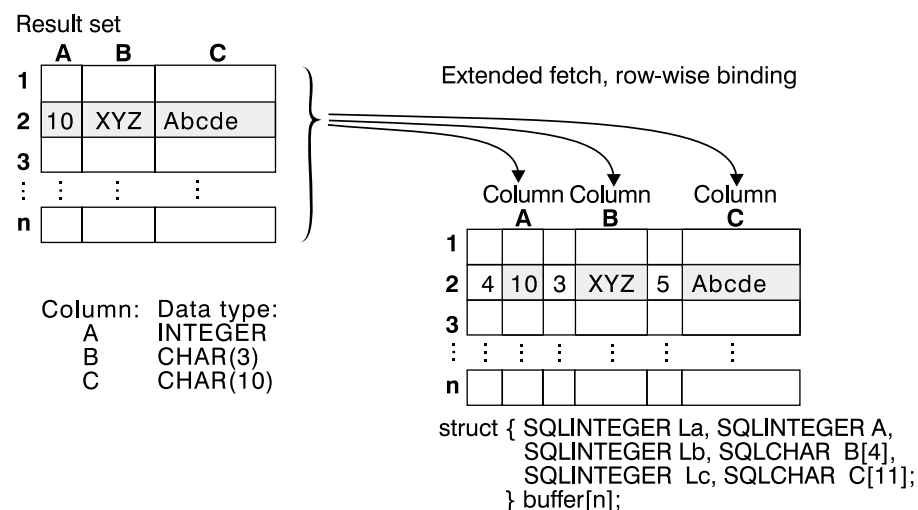


Figure 13. Row-wise binding

## Returning array data for column-wise bound data

Figure 12 on page 379 is a logical view of column-wise binding. The right side of Figure 14 on page 381 shows the function flows for column-wise retrieval.

To specify column-wise array retrieval, the application calls `SQLSetStmtOption()` with the `SQL_ROWSET_SIZE` attribute to indicate how many rows to retrieve at a time. When the value of the `SQL_ROWSET_SIZE` attribute is greater than 1, DB2 ODBC knows to treat the deferred output data pointer and length pointer as pointers to arrays of data and length rather than to one single element of data and length of a result set column.

The application then calls `SQLExtendedFetch()` to retrieve the data. When returning data, DB2 ODBC uses the maximum buffer size argument (*cbValueMax*) on `SQLBindCol()` to determine where to store successive rows of data in the array; the number of bytes available for return for each element is stored in the deferred length array. If the number of rows in the result set is greater than the `SQL_ROWSET_SIZE` attribute value, multiple calls to `SQLExtendedFetch()` are required to retrieve all the rows.

## Returning array data for row-wise bound data

The application can also do row-wise binding which associates an entire row of the result set with a structure. In this case the rowset is retrieved into an array of structures, each of which holds the data in one row and the associated length fields. Figure 13 on page 379 gives a pictorial view of row-wise binding.

To perform row-wise array retrieval, the application needs to call `SQLSetStmtOption()` with the `SQL_ROWSET_SIZE` attribute to indicate how many rows to retrieve at a time. In addition, it must call `SQLSetStmtOption()` with the `SQL_BIND_TYPE` attribute value set to the size of the structure to which the result columns are bound. DB2 ODBC treats the deferred output data pointer of `SQLBindCol()` as the address of the data field for the column in the first element of the array of these structures. It treats the deferred output length pointer as the address of the associated length field of the column.

The application then calls `SQLExtendedFetch()` to retrieve the data. When returning data, DB2 ODBC uses the structure size provided with the `SQL_BIND_TYPE` attribute to determine where to store successive rows in the array of structures.

Figure 14 on page 381 shows the required functions for each method. The left side shows  $n$  rows being selected, and retrieved one row at a time into  $m$  application variables. The right side shows the same  $n$  rows being selected, and retrieved directly into an array.

### Note:

- The diagram shows  $m$  columns bound, so  $m$  calls to `SQLBindCol()` are required in both cases.
- If arrays of less than  $n$  elements had been allocated, then multiple `SQLExtendedFetch()` calls would be required.

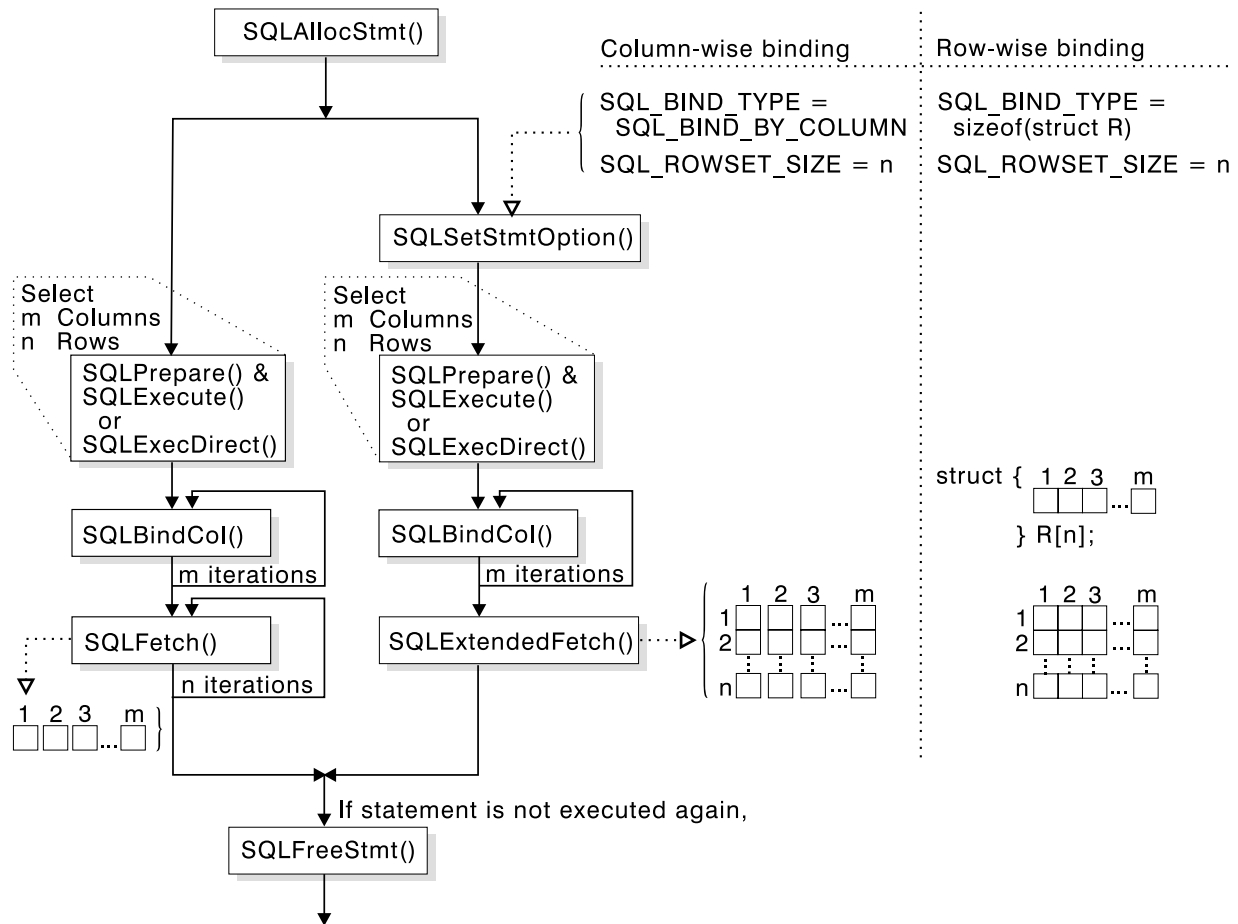


Figure 14. Array retrieval

## Column-wise, row-wise binding example

```

/* ... */
#define NUM_CUSTOMERS 25
SQLCHAR      stmt[] =
{ "WITH " /* Common Table expression (or Define Inline View) */
  "order (ord_num, cust_num, prod_num, quantity, amount) AS "
  "( "
  "SELECT c.ord_num, c.cust_num, l.prod_num, l.quantity, "
  "       price(char(p.price, '.'), p.units, char(l.quantity, '.')) "
  "FROM ord_cust c, ord_line l, product p "
  "WHERE c.ord_num = l.ord_num AND l.prod_num = p.prod_num "
  "AND cust_num = CNUM(cast (? as integer)) "
  ")", "
  "totals (ord_num, total) AS "
  "( "
  "SELECT ord_num, sum(decimal(amount, 10, 2)) "
  "FROM order GROUP BY ord_num "
  ")", "

```

```

        /* The 'actual' SELECT from the inline view */
        "SELECT order.ord_num, cust_num, prod_num, quantity, "
        "DECIMAL(amount,10,2) amount, total "
        "FROM order, totals "
        "WHERE order.ord_num = totals.ord_num "
    };
    /* Array of customers to get list of all orders for */
    SQLINTEGER    Cust[] =
    {
        10, 20, 30, 40, 50, 60, 70, 80, 90, 100,
        110, 120, 130, 140, 150, 160, 170, 180, 190, 200,
        210, 220, 230, 240, 250
    };

#define NUM_CUSTOMERS sizeof(Cust)/sizeof(SQLINTEGER)

    /* Row-Wise (Includes buffer for both column data and length) */
    struct {
        SQLINTEGER    Ord_Num_L;
        SQLINTEGER    Ord_Num;
        SQLINTEGER    Cust_Num_L;
        SQLINTEGER    Cust_Num;
        SQLINTEGER    Prod_Num_L;
        SQLINTEGER    Prod_Num;
        SQLINTEGER    Quant_L;
        SQLDOUBLE     Quant;
        SQLINTEGER    Amount_L;
        SQLDOUBLE     Amount;
        SQLINTEGER    Total_L;
        SQLDOUBLE     Total;
    } Ord[ROWSET_SIZE];

    SQLINTEGER    pirow = 0;
    SQLINTEGER    pcrow;
    SQLINTEGER    i;
    SQLINTEGER    j;
    /* ... */
    /* Get details and total for each order Row-Wise */
    rc = SQLAllocStmt(hdbc, &hstmt);

    rc = SQLParamOptions(hstmt, NUM_CUSTOMERS, &pirow);

    rc = SQLBindParameter(hstmt, 1, SQL_PARAM_INPUT, SQL_C_LONG, SQL_INTEGER,
        0, 0, Cust, 0, NULL);

    rc = SQLExecDirect(hstmt, stmt, SQL_NTS);

    /* SQL_ROWSET_SIZE sets the max number of result rows to fetch each time */
    rc = SQLSetStmtOption(hstmt, SQL_ROWSET_SIZE, ROWSET_SIZE);

    /* Set Size of One row, Used for Row-Wise Binding Only */
    rc = SQLSetStmtOption(hstmt, SQL_BIND_TYPE, sizeof(Ord) / ROWSET_SIZE);

    /* Bind column 1 to the Ord_num Field of the first row in the array*/
    rc = SQLBindCol(hstmt, 1, SQL_C_LONG, (SQLPOINTER) &Ord[0].Ord_Num, 0,
        &Ord[0].Ord_Num_L);

    /* Bind remaining columns ... */
    /* ... */

```

```

/* NOTE: This sample assumes that an order never has more
        rows than ROWSET_SIZE. A check should be added below to call
        SQLExtendedFetch multiple times for each result set.
*/
do /* for each result set .... */
{ rc = SQLExtendedFetch(hstmt, SQL_FETCH_NEXT, 0, &pcrow, NULL);

  if (pcrow > 0) /* if 1 or more rows in the result set */
  {
    i = j = 0;

    printf("*****\n");
    printf("Orders for Customer: %ld\n", Ord[0].Cust_Num);
    printf("*****\n");

    while (i < pcrow)
    { printf("\nOrder #: %ld\n", Ord[i].Ord_Num);
      printf("      Product  Quantity      Price\n");
      printf("      -----  -----  ----- \n");
      j = i;
      while (Ord[j].Ord_Num == Ord[i].Ord_Num)
      { printf("      %8ld %16.7lf %12.2lf\n",
                Ord[i].Prod_Num, Ord[i].Quant, Ord[i].Amount);
        i++;
      }
      printf("                                =====\n");
      printf("                                %12.2lf\n", Ord[j].Total);
    } /* end while */
  } /* end if */
}
while ( SQLMoreResults(hstmt) == SQL_SUCCESS);
/* ... */

```

---

## Using large objects

The term large object and the generic acronym LOB see any type of large object. There are three LOB data types:

- Binary large object (BLOB)
- Character large object (CLOB)
- Double-byte character large object (DBCLOB)

These LOB data types are represented symbolically as SQL\_BLOB, SQL\_CLOB, SQL\_DBCLOB respectively. The LOB symbolic constants can be specified or returned on any of the DB2 ODBC functions that take in or return an SQL data type argument (for example, SQLBindParameter() and SQLDescribeCol()). See Table 4 on page 39 for a complete list of symbolic and default C symbolic names for SQL data types.

Although LOB values can be very large, they can be retrieved and manipulated in the application address space. However, the application might not need or want the entire LOB transferred from the database server into application memory. In many cases an application needs to select a LOB value and operate on pieces of it. Transfer of LOB data using the piece-wise sequential method provided by SQLGetData() and SQLPutData(), following the ODBC model, can be very expensive and time consuming. The application can reference an individual LOB value using a LOB locator.

## Using LOB locators

A LOB locator is a mechanism that allows an application program to manipulate a large object value in an efficient, random access manner. The locator is a runtime concept: it is not a persistent type and is not stored in the database; it is a mechanism used to see a LOB value during a transaction and does not persist beyond the transaction in which it is created. Each LOB locator type has its own C data type (SQL\_C\_BLOB\_LOCATOR, SQL\_C\_CLOB\_LOCATOR, SQL\_C\_DBCLOB\_LOCATOR) which enables transfer of LOB locator values to and from the database server.

A LOB locator is a simple token value that represents a single LOB value. A locator is not a reference to a column in a row, rather it references a large object value. An operation performed on a locator does not affect the original LOB value stored in the row. An application can retrieve a LOB locator into an application variable (using the `SQLBindCol()` or `SQLGetData()` function) and can then apply the following DB2 ODBC functions to the associated LOB value by using the locator:

- `SQLGetLength()` gets the length of a string a LOB locator represents.
- `SQLGetPosition()` gets the position of a search string within a source string where a LOB locator represents the source string. A LOB locator can also represent the search string.

Locators are implicitly allocated by:

- Fetching a bound LOB column to the appropriate C locator type.
- Calling `SQLGetSubString()` and specifying that the substring be retrieved as a locator.
- Calling `SQLGetData()` on an unbound LOB column and specifying the appropriate C locator type. The C locator type must match the LOB column type or an error occurs.

LOB locators also provide an efficient method of moving data from one column of a table at the server to another column (of the same or different table) without having to first pull the data into application memory and then send it back to the server.

For example, the following INSERT statement inserts a LOB value that is a concatenation of two LOB values as represented by their locators:

```
INSERT INTO TABLE4A
VALUES(1,CAST(? AS CLOB(2K)) CONCAT CAST(? AS CLOB(3K)))
```

The locator can be explicitly freed before the end of a transaction by executing the `FREE LOCATOR` statement or retained beyond a unit of work by executing the `HOLD LOCATOR` statement. The syntax is shown below.

```
▶▶ FREE LOCATOR host_variable ◀◀
```

```
▶▶ HOLD LOCATOR host_variable ◀◀
```

Although this statement cannot be prepared dynamically, DB2 ODBC accepts it as a valid statement on `SQLPrepare()` and `SQLExecDirect()`. The application uses `SQLBindParameter()` with the SQL data type argument set to the appropriate SQL and C symbolic data types listed in Table 4 on page 39.

This function is not available when connected to a DB2 server that does not support large objects. Call `SQLGetFunctions()` with the function type set to `SQL_API_SQLGETSUBSTRING`, and check the *fExists* output argument to determine if the function is supported for the current connection.

## LOB example

This example extracts the 'Interests' section from the Resume CLOB column of the `EMP_RESUME` table. Only the substring is transferred to the application.

```

/* ... */
SQLCHAR          stmt2[] =
                  "SELECT resume FROM emp_resume "
                  "WHERE empno = ? AND resume_format = 'ascii'";
/* ... */
/*****
** Get CLOB locator to selected Resume **
*****/
rc = SQLSetParam(hstmt, 1, SQL_C_CHAR, SQL_CHAR, 7,
                0, Empno.s, &Empno.ind);

printf("\n>Enter an employee number:\n");
gets(Empno.s);

rc = SQLExecDirect(hstmt, stmt2, SQL_NTS);
rc = SQLBindCol(hstmt, 1, SQL_C_CLOB_LOCATOR, &ClobLoc1, 0,
                &pcbValue);
rc = SQLFetch(hstmt);

```

```

/*****
  Search CLOB locator to find "Interests"
  Get substring of resume (from position of interests to end)
*****/

rc = SQLAllocStmt(hdbc, &lhstmt);

/* Get total length */
rc = SQLGetLength(lhstmt, SQL_C_CLOB_LOCATOR, ClobLoc1, &SLength, &Ind);

/* Get Starting position */
rc = SQLGetPosition(lhstmt, SQL_C_CLOB_LOCATOR, ClobLoc1, 0,
                  "Interests", 9, 1, &Pos1, &Ind);

buffer = (SQLCHAR *)malloc(SLength - Pos1 + 1);

/* Get just the "Interests" section of the Resume CLOB */
/* (From Pos1 to end of CLOB) */
rc = SQLGetSubString(lhstmt, SQL_C_CLOB_LOCATOR, ClobLoc1, Pos1,
                    SLength - Pos1, SQL_C_CHAR, buffer, SLength - Pos1 + 1,
                    &OutLength, &Ind);

/* Print Interest section of Employee's resume */
printf("\nEmployee #: %s\n %s\n", Empno.s, buffer);
/* ... */

```

---

## Using distinct types

In addition to SQL data types (referred to as base SQL data types), you can define distinct types. A distinct type, or user-defined data type, shares its internal representation with a source type, but is considered to be a separate and incompatible type for most operations. Distinct types are created using the CREATE DISTINCT TYPE SQL statement. See *DB2 SQL Reference* for information about using this statement.

Distinct types help provide the strong typing control needed in object-oriented programming by ensuring that only those functions and operators explicitly defined on a distinct type can be applied to its instances. Applications continue to work with C data types for application variables, and only need to consider the distinct types when constructing SQL statements.

The following guidelines apply to distinct types:

- All SQL to C data type conversion rules that apply to the built-in type also apply to the distinct type.
- The distinct type has the same default C type as the built-in type.
- `SQLDescribeCol()` returns the built-in type information. The distinct type name can be obtained by calling `SQLColAttributes()` with the input descriptor type set to `SQL_COLUMN_DISTINCT_TYPE`.
- SQL predicates that involve parameter markers must be explicitly cast to the distinct type or base SQL data types. This is required since the application can only deal with the built-in types, so before any operation can be performed using the parameter, it must be cast from the C built-in type to the distinct type; otherwise an error occurs when the statement is prepared. For more



information about casting distinct types, see “Casting distinct types” on page 411.

## Distinct type example

This example shows the definition of distinct types and user-defined functions, and tables with distinct type columns. for an example that inserts rows into a table with distinct type columns, see “Array input example” on page 378.

```
/* ... */
/* Initialize SQL statement strings */
SQLCHAR      stmt[][MAX_STMT_LEN] = {
    "CREATE DISTINCT TYPE CNUM AS INTEGER WITH COMPARISONS",

    "CREATE DISTINCT TYPE PUNIT AS CHAR(2) WITH COMPARISONS",

    "CREATE DISTINCT TYPE UPRICE AS DECIMAL(10, 2) \
    WITH COMPARISONS",

    "CREATE DISTINCT TYPE PRICE AS DECIMAL(10, 2) \
    WITH COMPARISONS",

    "CREATE FUNCTION PRICE (CHAR(12), PUNIT, char(16) ) \
    returns char(12) \
    NOT FENCED EXTERNAL NAME 'order!price' \
    NOT VARIANT NO SQL LANGUAGE C PARAMETER STYLE DB2SQL \
    NO EXTERNAL ACTION",
```

```

"CREATE DISTINCT TYPE PNUM AS INTEGER WITH COMPARISONS",

"CREATE FUNCTION "+" (PNUM, INTEGER) RETURNS PNUM \
  source sysibm.\"+\"(integer, integer)",

"CREATE FUNCTION MAX (PNUM) RETURNS PNUM \
  source max(integer)",

"CREATE DISTINCT TYPE ONUM AS INTEGER WITH COMPARISONS",

"CREATE TABLE CUSTOMER ( \
Cust_Num      CNUM NOT NULL, \
First_Name    CHAR(30) NOT NULL, \
Last_Name     CHAR(30) NOT NULL, \
Street        CHAR(128) WITH DEFAULT, \
City          CHAR(30) WITH DEFAULT, \
Prov_State    CHAR(30) WITH DEFAULT, \
PZ_Code       CHAR(9) WITH DEFAULT, \
Country       CHAR(30) WITH DEFAULT, \
Phone_Num     CHAR(20) WITH DEFAULT, \
PRIMARY KEY (Cust_Num) )",

"CREATE TABLE PRODUCT ( \
Prod_Num      PNUM NOT NULL, \
Description   VARCHAR(256) NOT NULL, \
Price        DECIMAL(10,2) WITH DEFAULT , \
Units        PUNIT NOT NULL, \
Combo        CHAR(1) WITH DEFAULT, \
PRIMARY KEY (Prod_Num), \
CHECK (Units in (PUNIT('m'), PUNIT('l'), PUNIT('g'), PUNIT('kg'), \
PUNIT(' '))) )",

"CREATE TABLE PROD_PARTS ( \
Prod_Num      PNUM NOT NULL, \
Part_Num      PNUM NOT NULL, \
Quantity      DECIMAL(14,7), \
PRIMARY KEY (Prod_Num, Part_Num), \
FOREIGN KEY (Prod_Num) REFERENCES Product, \
FOREIGN KEY (Part_Num) REFERENCES Product, \
CHECK (Prod_Num <> Part_Num) )",

```

```

        "CREATE TABLE ORD_CUST( \
Ord_Num      ONUM NOT NULL, \
Cust_Num     CNUM NOT NULL, \
Ord_Date     DATE NOT NULL, \
PRIMARY KEY (Ord_Num), \
FOREIGN KEY (Cust_Num) REFERENCES Customer )" ,

        "CREATE TABLE ORD_LINE( \
Ord_Num      ONUM NOT NULL, \
Prod_Num     PNUM NOT NULL, \
Quantity     DECIMAL(14,7), \
PRIMARY KEY (Ord_Num, Prod_Num), \
FOREIGN KEY (Prod_Num) REFERENCES Product, \
FOREIGN KEY (Ord_Num) REFERENCES Ord_Cust )"
    };
/* ... */
    num_stmts = sizeof(stmt) / MAX_STMT_LEN;

    printf(">Executing %ld Statements\n", num_stmts);

    /* Execute Direct statements */
    for (i = 0; i < num_stmts; i++) {
        rc = SQLExecDirect(hstmt, stmt[i], SQL_NTS);
    }
/* ... */

```

---

## Using stored procedures

An application can be designed to run in two parts, one on the client and the other on the server. The stored procedure is a server application that runs at the database within the same transaction as the client application. Stored procedures can be written in either embedded SQL or using the DB2 ODBC functions (see "Writing a DB2 ODBC stored procedure" on page 391).

Both the main application that calls a stored procedure and a stored procedure itself can be either a DB2 ODBC application or a standard DB2 precompiled application. Any combination of embedded SQL and DB2 ODBC applications can be used. Figure 15 illustrates this concept.

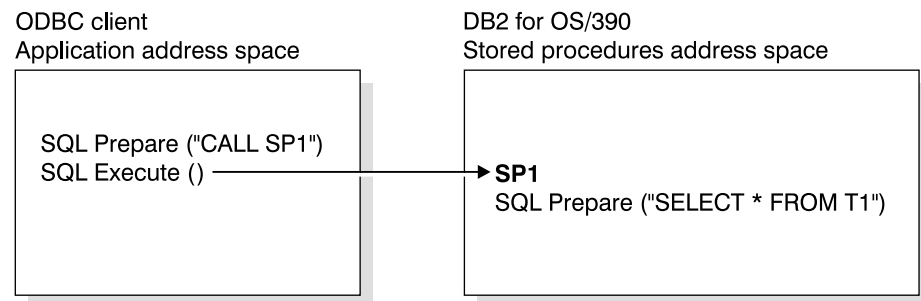


Figure 15. Running stored procedures

## Advantages of using stored procedures

In general, stored procedures have the following advantages:

- Avoid network transfer of large amounts of data obtained as part of intermediate results in a long sequence of queries.
- Deployment of client database applications into client/server pieces.

In addition, stored procedures written in embedded static SQL have the following advantages:

- Performance - static SQL is prepared at precompile time and has no run time overhead of access plan (package) generation.
- Encapsulation (information hiding) - users do not need to know the details about database objects in order to access them. Static SQL can help enforce this encapsulation.
- Security - users access privileges are encapsulated within the packages associated with the stored procedures, so there is no need to grant explicit access to each database object. For example, you can grant a user run access for a stored procedure that selects data from tables for which the user does not have select privilege.

## Catalog table for stored procedures


If the server is DB2 UDB Version 2 Release 1 or later, or DB2 for MVS/ESA Version 4 or later, an application can call `SQLProcedureColumns()` to determine the type of a parameter in a procedure call.

If the stored procedure resides on a DB2 for MVS/ESA Version 4 or a DB2 for OS/390 Version 5 server, the name of the stored procedure must be defined in the `SYSIBM.SYSPROCEDURES` catalog table. If the stored procedure resides on a DB2 for OS/390 Version 6 server, the name of the stored procedure must be defined in the `SYSIBM.SYSROUTINES` catalog table.

For more information, see “Catalog table for stored procedures” and “SQLProcedureColumns - Get input/output parameter information for a procedure” on page 290.

## Calling stored procedures from a DB2 ODBC application

Stored procedures are invoked from a DB2 ODBC application by passing the following `CALL` statement syntax to `SQLExecDirect()` or `SQLPrepare()` followed by `SQLExecute()`.



### **procedure-name**

The name of the stored procedure to execute.

Although the `CALL` statement cannot be prepared dynamically, DB2 ODBC accepts the `CALL` statement as if it can be dynamically prepared. Stored procedures can

also be called using the ODBC vendor escape sequence shown in “Stored procedure CALL syntax” on page 406.

For more information regarding the use of the CALL statement and stored procedures, see *DB2 SQL Reference* and *DB2 Application Programming and SQL Guide*.

If the server is DB2 UDB Version 2 Release 1 or later, or DB2 for MVS/ESA Version 4 or later, `SQLProcedures()` can be called to obtain a list of stored procedures available at the database.

The ? in the CALL statement syntax diagram denotes parameter markers corresponding to the arguments for a stored procedure. All arguments must be passed using parameter markers; literals, the NULL keyword, and special registers are not allowed. However, literals can be used if a vendor escape clause for the CALL statement is used. See “Using vendor escape clauses” on page 404.

The parameter markers in the CALL statement are bound to application variables using `SQLBindParameter()`. Although stored procedure arguments can be used both for input and output, in order to avoid sending unnecessary data between the client and the server, the application should specify on `SQLBindParameter()` the parameter type of an input argument to be `SQL_PARAM_INPUT` and the parameter type of an output argument to be `SQL_PARAM_OUTPUT`. Those arguments that are both input and output have a parameter type of `SQL_PARAM_INPUT_OUTPUT`. Literals are considered type `SQL_PARAM_INPUT` only.

## Writing a DB2 ODBC stored procedure

Although embedded SQL stored procedures provide the most advantages, application developers who have existing DB2 ODBC applications might wish to move components of the application to run on the server. In order to minimize the required changes to the code and logic of the application, these components can be implemented by writing stored procedures using DB2 ODBC.

Auto-commit must be off. This is achieved by using the `AUTOCOMMIT` keyword in the initialization file or by using the `SQLSetConnectOption` API with the `SQL_AUTOCOMMIT` connect option `SQL_AUTOCOMMIT_OFF`.

`SQLConnect()` should be a null connect. Since all the internal information related to a DB2 ODBC connection is referenced by the connection handle, and since a stored procedure runs under the same connection and transaction as the client application, a stored procedure written using DB2 ODBC must make a null `SQLConnect()` call to associate a connection handle with the underlying connection of the client application. In a null `SQLConnect()` call, the `szDSN`, `szUID`, and `szAuthStr` argument pointers are all set to NULL and their respective length arguments all set to 0.

For stored procedures written in DB2 ODBC, the `COMMIT_ON_RETURN` option has no effect on DB2 ODBC rules; set it to 'N'. However, be aware that setting this option to 'N' overrides the manual-commit mode set in the client application.

The `MVSATTACHTYPE` keyword should be set to `RRSAF` if the stored procedure contains any LOB data types or distinct types in its parameter list. DB2 for OS/390 requires that stored procedures containing any LOBs or distinct types must run in a

WLM-established stored procedure address space. For further information on setting up the stored procedures environment, see Section 6 of *DB2 Application Programming and SQL Guide*.

# For information about binding a stored procedure that runs under DB2 ODBC, see  
# "Bind stored procedures" on page 54.

## Returning result sets from stored procedures

DB2 ODBC provides the ability to retrieve one or more result sets from a stored procedure call, provided the stored procedure is coded such that one or more cursors, each associated with a query, is opened and left open when the stored procedure exits. If more than one cursor is left open, multiple result sets are returned.

Stored procedures written to return one or more result sets to a DB2 ODBC application should indicate the maximum number of result sets that can be returned in the RESULT\_SETS column of the SYSIBM.SYSPROCEDURES table (DB2 for MVS/ESA Version 4 or DB2 for OS/390 Version 5) or the RESULT\_SETS column of the SYSIBM.SYSROUTINES table (DB2 for OS/390 Version 6). A zero in this column indicates that open cursors returned no result sets.

### Programming stored procedures to return result sets

To return one or more result sets to a DB2 ODBC application the stored procedure must satisfy the following requirements:

- The stored procedure indicates that it wants a result set returned by declaring a cursor on the result set, opening a cursor on the result set (that is, executing the query), and leaving the cursor open when exiting the stored procedure.
- For every cursor that is left open, a result set is returned to the application.
- If more than one cursor is left open, the result sets are returned in the order in which their cursors were opened in the stored procedure.
- In a stored procedure, DB2 ODBC uses a cursor declared WITH RETURN. If the cursor is closed before the stored procedure exit, it is a local cursor. If the cursor remains open upon stored procedure exit, it returns a query result set (also called a multiple result set) to the client application.

# To leave the cursor open to return result sets, the application must follow these  
# guidelines:

- # – Issue SQLExecute() or SQLExecDirect().
- # – Optionally, SQLFetch() rows.
- # – Do not issue SQLFreeStmt() with either SQL\_DROP or SQL\_CLOSE.
- # – Issue SQLDisconnect(), SQLFreeConnect(), and SQLFreeEnv() to terminate  
# with the statement handle in a valid state.

# By avoiding SQLFreeStmt(), the cursor remains open to return result sets.  
# Appendix E, "Example code" on page 469 provides an example; see case 2 of  
# step 4.

- Only unread rows are passed back. For example, if the result set of a cursor has 500 rows, and 150 of those rows were already read by the stored procedure when it terminated, then rows 151 through 500 are returned to the

stored procedure. This can be useful if the stored procedure wishes to filter out some initial rows and not return them to the application.

### **Restrictions on stored procedures returning result sets**

In general, calling a stored procedure that returns a result set is equivalent to executing a query statement. The following restrictions apply:

- Column names are not returned by either `SQLDescribeCol()` or `SQLColAttributes()` for static query statements. In this case, the ordinal position of the column is returned instead.
- All result sets are read-only.
- Schema functions (such as `SQLTables()`) cannot be used to return a result set. If schema functions are used within a stored procedure, all of the cursors for the associated statement handles must be closed before returning, otherwise extraneous result sets might be returned.
- When a query is prepared, result set column information is available before the execute. When a stored procedure is prepared, the result set column information is not available until the `CALL` statement is executed.

### **Programming DB2 ODBC client applications to receive result sets**

DB2 ODBC applications can retrieve result sets after the execution of a stored procedure that leaves cursors open. The following guidelines explain the process and requirements.

- Before the stored procedure is called, ensure that there are no open cursors associated with the statement handle.
- Call the stored procedure.
- The execution of the stored procedure `CALL` statement effectively causes the cursors associated with the result sets to open.
- Examine any output parameters that are returned by the stored procedure. For example, the procedure might be designed so that there is an output parameter that indicates exactly how many result sets are generated.
- The DB2 ODBC application can then process a query as it normally does. If the application does not know the nature of the result set or the number of columns returned, it can call `SQLNumResultCols()`, `SQLDescribeCol()` or `SQLColAttributes()`. Next, the application can use any permitted combination of `SQLBindCol()`, `SQLFetch()`, and `SQLGetData()` to obtain the data in the result set.
- When `SQLFetch()` returns `SQL_NO_DATA_FOUND` or if the application is finished with the current result set, the application can call `SQLMoreResults()` to determine if there are more result sets to retrieve. Calling `SQLMoreResults()` closes the current cursor and advances processing to the next cursor that was left open by the stored procedure.
- If there is another result set, then `SQLMoreResults()` returns `SQL_SUCCESS`; otherwise, it returns an `SQL_NO_DATA_FOUND`.
- Result sets must be processed in serial fashion by the application, that is, one at a time in the order that they were opened in the stored procedure.

## Stored procedure example with query result set

A detailed stored procedure example is provided in Appendix E, “Example code” on page 469.

---

## Writing multithreaded applications

This section explains DB2 ODBC's support of multithreading and multiple contexts, and provides guidelines for programming techniques.

### DB2 ODBC support of multiple LE threads

All DB2 ODBC applications have at least one LE thread created automatically in the application's LE *enclave*. A multithreaded DB2 ODBC application creates additional LE threads using the *POSIX Pthread* function **pthread\_create()**. These additional threads share the same reentrant copy of DB2 ODBC code within the LE enclave.

DB2 ODBC code is *reentrant* but uses shared storage areas that must be protected if multiple LE threads are running concurrently in the same enclave. The quality of being reentrant and correctly handling shared storage areas is referred to as *threadsafe*. This quality is required by multithreaded applications.

DB2 ODBC supports concurrent execution of LE threads by making all of the DB2 ODBC function calls *threadsafe*. This *threadsafe* quality of function calls is only available if DB2 ODBC has access to OS/390 UNIX. DB2 ODBC uses the Pthread *mutex* functions of OS/390 UNIX to provide *threadsafe* function calls by serializing critical sections of DB2 ODBC code. See “Initialization keywords” on page 63 for a description of the THREADSAFE keyword.

Because OS/390 UNIX is present, *threadsafe* capability is available by default when executing a DB2 ODBC application in the following environments:

- The OS/390 UNIX shell
- TSO or batch for HFS-resident applications using the IBM supplied BPXBATCH program. (See *OS/390 UNIX System Services Command Reference* for more information about BPXBATCH).
- TSO or batch for applications that are not HFS-resident if the LE runtime option POSIX(ON) is specified when the application runs.

For example, to specify POSIX(ON) in TSO, you can invoke the DB2 ODBC application APP1 in the MVS dataset USER.RUNLIB.LOAD as follows:

```
CALL 'USER.RUNLIB.LOAD(APP1)' 'POSIX(ON)'
```

Using batch JCL, you can invoke the same application:

```
//STEP1 EXEC PGM=APP1,PARM='POSIX(ON)'  
//STEPLIB DD DSN=USER.RUNLIB.LOAD,DISP=SHR  
//          DD ...other libraries needed at runtime...
```

For more OS/390 UNIX information relating to DB2 ODBC, see “Special considerations for OS/390 UNIX” on page 57. Also, see *Language Environment for OS/390 & VM Programming Guide* for more information about running programs that use OS/390 UNIX.



Multithreaded applications allow threads to perform work in parallel on different connections as shown in Figure 16 on page 395.

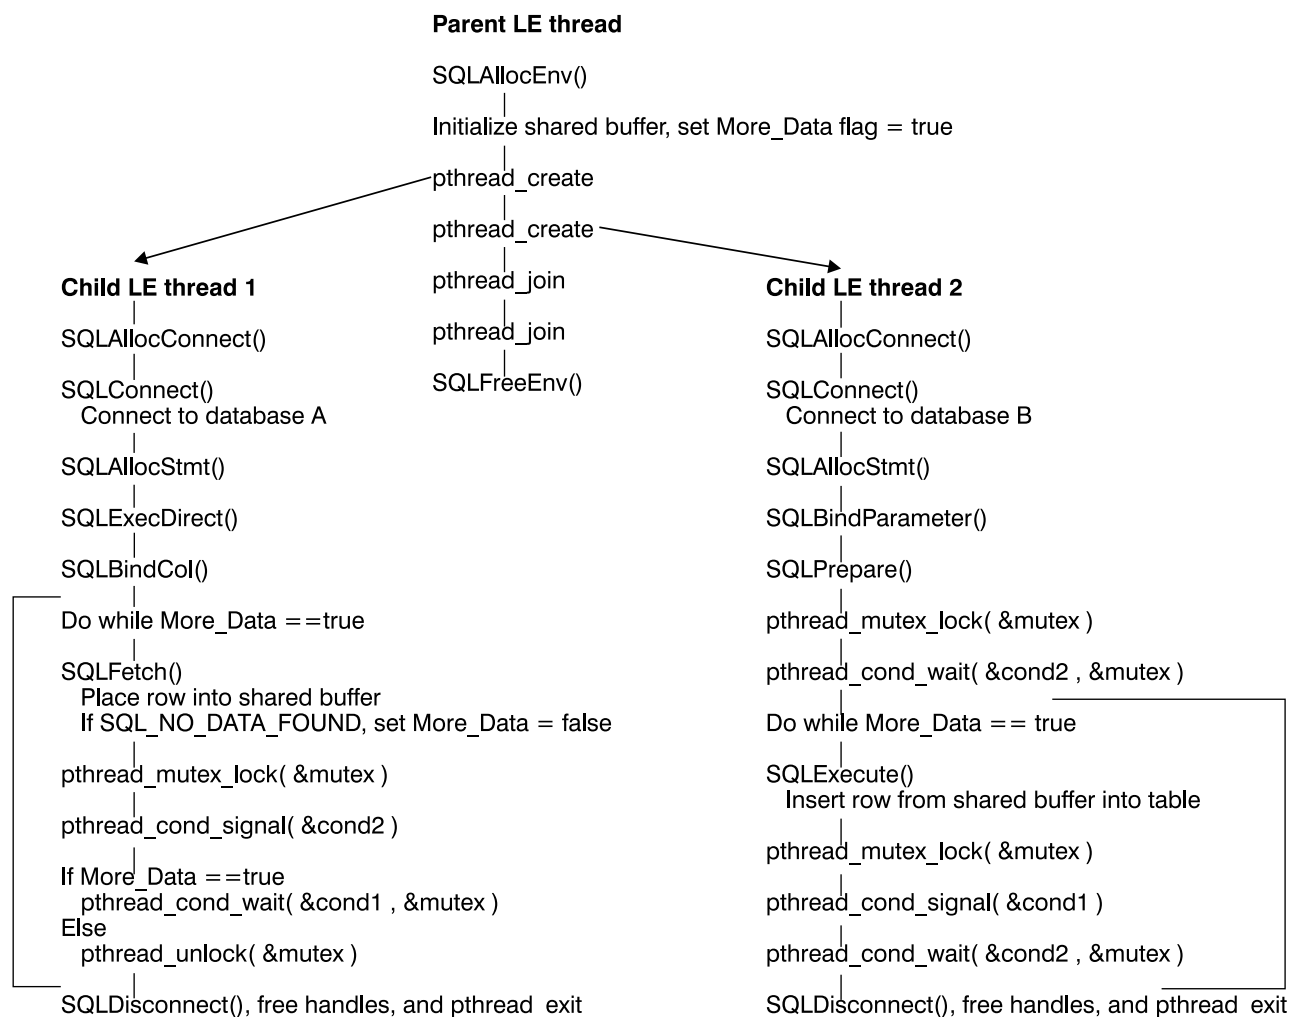


Figure 16. Multithreaded application

In Figure 16, an application implements a database-to-database copy as follows:

- A parent LE thread creates two child LE threads. The parent thread remains present for the duration of the activity of the child threads. DB2 ODBC requires that the thread which establishes the environment using `SQLAllocEnv()` must remain present for the duration of the application, so that DB2 language interface routines remain resident in the LE enclave.
- One child LE thread connects to database A and uses `SQLFetch()` calls to read data from one connection into a shared application buffer.
- Another child LE thread connects to database B and concurrently reads from the shared buffer, inserting the data into database B.
- Pthread functions are used to synchronize the use of the shared application buffer. See *OS/390 C/C++ Run-Time Library Reference* for a description of the Pthread functions.

## When to use multiple LE threads

Detailed discussion of evaluating application requirements and making decisions about whether or not to use multithreading is beyond the scope of this book. However, there are some general application types that are well-suited to multithreading. For example, applications that handle asynchronous work requests are candidates for multithreading.

An application that handles asynchronous work requests can take the form of a parent/child threading model in which the parent LE thread creates child LE threads to handle incoming work. The parent thread can then act as a dispatcher of these work requests as they arrive, directing them to child threads that are not currently busy servicing other work.

## DB2 ODBC support of multiple contexts

The context consists of the application's logical connection to the data source and associated internal DB2 ODBC connection information that allows the application to direct its operations to a data source. In DB2 ODBC, the context is established when `SQLAllocConnect()` is issued.

The context is the DB2 ODBC equivalent of a DB2 thread. DB2 ODBC always creates a context when a successful `SQLAllocConnect()` is first issued by an application LE thread. If DB2 ODBC support for multiple contexts is not enabled, only the first `SQLAllocConnect()` for a LE thread establishes a context. With support for multiple contexts, DB2 ODBC establishes a separate context (and DB2 thread) each time the `SQLAllocConnect()` function is invoked.

If the initialization file specifies `MULTICONTEXT=0` (see "Initialization keywords" on page 63), there can only be one context for each LE thread that the application creates. This single context per thread can provide only the simulated support of the ODBC connection model, explained in "DB2 ODBC restrictions on the ODBC connection model" on page 23.

When the initialization file specifies `MULTICONTEXT=1`, a distinct context is established for each connection handle that is allocated when `SQLAllocConnect()` is issued. Using `MULTICONTEXT=1` requires:

- The RRSF attachment facility, specified by `MVSATTACHTYPE=RRSAF` in the initialization file.
- OS/390 Unauthorized Context Services, available in OS/390 Version 2 Release 5 and higher releases.

Consistent with the ODBC connection model, the use of `MULTICONTEXT=1` implies `CONNECTTYPE=1`. The connections are independently handled by `SQLTransact` for both commit and rollback.

The creation of a context for each connection is consistent with, and provides full support for, the ODBC connection model.

The context is established with `SQLAllocConnect()` and deleted by `SQLFreeConnect()`. All `SQLConnect()` and `SQLDisconnect()` operations that use the same connection handle belong to the same context. Although there can be only one active connection to a data source at any given time for the duration of the

context, the target data source can be changed by `SQLDisconnect()` and `SQLConnect()`, subject to the rules of `CONNECTTYPE=1`.

With `MULTICONTEXT=1` specified, DB2 ODBC automatically uses OS/390 Unauthorized Context Services to create and manage contexts for the application. However, DB2 ODBC does not perform context management for the application if any of the following are true:

- The DB2 ODBC application created a DB2 thread before invoking DB2 ODBC. This is always the case for a stored procedure using DB2 ODBC.
- The DB2 ODBC application created and switched to a private context before invoking DB2 ODBC. For example, an application that is explicitly using OS/390 Context Services and issues `ctxswch` to switch to a private context prior to invoking DB2 ODBC cannot take advantage of `MULTICONTEXT=1`.
- The DB2 ODBC application started a unit of recovery with any RRS resource manager before invoking DB2 ODBC.
- `MVSATTACHTYPE=CAF` is specified in the initialization file.
- The OS/390 operating system level does not support Unauthorized Context Services.

The application can use the `SQLGetInfo()` function with `infoType=SQL_MULTIPLE_ACTIVE_TXN` to determine if `MULTICONTEXT=1` is active for the DB2 ODBC application. See “SQLGetInfo - Get general information” on page 219 for the description of `SQLGetInfo()`.

Table 147. Connection characteristics

Settings		Results		
MULTICONTEXT	CONNECTTYPE	Can LE thread have more than one ODBC connection with an outstanding unit of work?	Can LE thread commit/rollback ODBC connection independently?	Number of DB2 created by DB2 ODBC on behalf of application
0	2	Y	N	1 per LE thread
0	1	N	Y	1 per LE thread
1 1	1 or 2 2	Y	Y	1 per ODBC connection handle

**Note:**

1. `MULTICONTEXT=1` requires `MVSATTACHTYPE=RRSAF` and OS/390 Version 2 Release 5 or higher.
2. `MULTICONTEXT=1` implies `CONNECTTYPE=1` characteristics. With `MULTICONTEXT=1` and `CONNECTTYPE=2` specified in the initialization file, DB2 ODBC ignores `CONNECTTYPE=2`. With `MULTICONTEXT=1` specified, any attempts to set `CONNECTTYPE=2` using `SQLSetEnvAttr()`, `SQLSetConnectOptions()`, or `SQLDriverConnect()` are rejected with `SQLSTATE=01S02`.
  - All connections in a DB2 ODBC application have the same `CONNECTTYPE` and `MULTICONTEXT` characteristics. `CONNECTTYPE` is established at the first `SQLConnect()`. `MULTICONTEXT` is established at `SQLAllocEnv()`.
  - For `CONNECTTYPE=1` or `MULTICONTEXT=1`, the `AUTOCOMMIT` default is ON. For `CONNECTTYPE=2` or `MULTICONTEXT=0`, the `AUTOCOMMIT` default is OFF.

## Multiple contexts, one LE thread

When using the initialization file setting MULTICONTEXT=1, a DB2 ODBC application can create multiple independent connections for a LE thread. Figure 17 is an example of a multicontext, one LE thread application.

```
/* Get an environment handle (henv).          */
SQLAllocEnv( &henv );

/*
 * Get two connection handles, hdbc1 and hdbc2, which
 * represent two independent DB2 threads.
 */
SQLAllocConnect( henv , &hdbc1 );
SQLAllocConnect( henv , &hdbc2 );

/* Set autocommit off for both connections.    */
/* This is done only to emphasize the          */
/* independence of the connections for purposes */
/* of this example, and is not intended as     */
/* a general recommendation.                  */

SQLSetConnectOption( hdbc1 , SQL_AUTOCOMMIT, SQL_AUTOCOMMIT_OFF );
SQLSetConnectOption( hdbc2 , SQL_AUTOCOMMIT, SQL_AUTOCOMMIT_OFF );

/* Perform SQL under DB2 thread 1 at STLEC1.   */

SQLConnect( hdbc1, (SQLCHAR *) "STLEC1", ... );
SQLAllocStmt ...
SQLExecDirect ...
.
.
/* Perform SQL under DB2 thread 2 at STLEC1.   */

SQLConnect( hdbc2, (SQLCHAR *) "STLEC1", ... );
SQLAllocStmt ...
SQLExecDirect ...
.
.
/* Commit changes on connection 1.            */

SQLTransact( henv , hdbc1 , SQL_COMMIT );

/* Rollback changes on connection 2.          */

SQLTransact( henv , hdbc2 , SQL_ROLLBACK );
.
.
```

Figure 17. Example of independent connections on a single LE thread.

## Multiple contexts, multiple LE threads

Using the initialization file setting MULTICONTEXT=1, combined with the default THREADSAFE=1, the application can create multiple independent connections under multiple LE threads. This capability can support complex DB2 ODBC server applications that handle multiple incoming work requests by using a fixed number of threads.

The multiple context, multiple LE thread capability requires some special considerations for the application using it. The Pthread functions should be used by

the application for serialization of the use of connection handles and the associated statement handles. As an example of what can go wrong without proper serialization, see Figure 18 on page 399.

```

LE_Thread_1                LE_Thread_2
.                            .
.                            .
.                            .
rc = SQLExecDirect( hstmt1 , ... ); .
.                            .
.                            .
.                            .
.                            .
SQLFreeStmt( hstmt1 , SQL_DROP);
.                            .
.                            .
rc = SQLExecDirect( hstmt1 , ... ); .
.                            .

rc has the value SQL_INVALID_HANDLE because
LE_Thread_2 has freed statement handle hstmt1.

```

Figure 18. Example of improper serialization.

The suggested design is to map one LE thread per connection by establishing a pool of connections as shown in Figure 19.

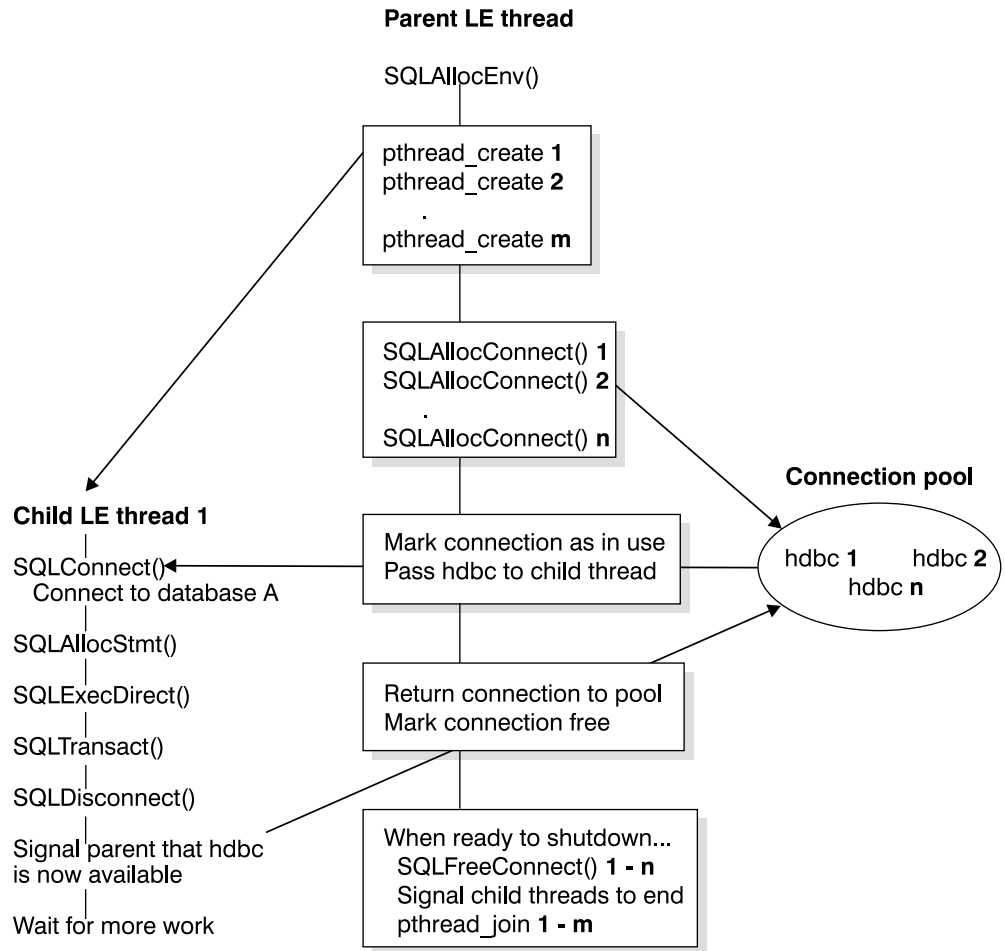


Figure 19. Model for multithreading with connection pooling (MULTICONTEXT=1)

In Figure 19 on page 399, a pool of connections is established as follows:

- Designate a parent LE thread which allocates:
  - $m$  child LE threads
  - $n$  connection handles
- Each task that requires a connection is executed by one of the child LE threads, and is given one of the  $n$  connections by the parent LE thread. The parent thread remains active, acting as a dispatcher of tasks.
- The parent LE thread marks a connection as being in use until the child thread returns it to the connection pool.
- The parent LE thread frees the connections using `SQLFreeConnect()` when the parent thread is ending.

DB2 ODBC requires that the LE thread which establishes the environment using `SQLAllocEnv()` must remain present for the duration of the application, so that DB2 language interface routines will remain resident in the LE enclave.

This suggested design allows the parent LE thread to create more LE threads than connections if the threads are also used to perform non-SQL related tasks, or more connections than threads if the application should maintain a pool of active connections but limit the number of active tasks.

Connections can move from one application LE thread to another as the connections in the pool are assigned to child threads, returned to the pool, and assigned again.

The use of this design prevents two LE threads from trying to use the same connection (or an associated statement handle) at the same time. Although DB2 ODBC controls access to its internal resources, the application resources such as bound columns, parameter buffers, and files are not controlled by DB2 ODBC. If it is necessary for two threads to share an application resource, the application must implement some form of synchronization mechanism. For example, the database copy scenario in Figure 16 on page 395 uses Pthread functions to synchronize use of the shared buffer.

## Application deadlocks

The possibility of timeouts, deadlocks, and general contention for database resources exists when multiple connections are used to access the same database resources concurrently.

An application that creates multiple connections by using multithreading or multiple context support can potentially create deadlocks with shared resources in the database.

A DB2 subsystem can detect deadlocks and rollback one or more connections to resolve them. An application can still deadlock if the following sequence occurs:

- Two LE threads connect to the same data source using two DB2 threads.
- One LE thread holds an internal application resource (such as a mutex) while its DB2 thread waits for access to a database resource.
- The other LE thread has a lock on a database resource while waiting for the internal application resource.

In this case the DB2 subsystem does not detect a deadlock since the application's internal resources cannot be monitored by a DB2 subsystem. However, the application is still subject to the DB2 subsystem detecting and handling any DB2 thread timeouts.

---

## Mixing embedded SQL and DB2 ODBC

It is possible, and sometimes desirable for an application to use DB2 ODBC in conjunction with embedded static SQL. Consider the scenario where the application developer wishes to take advantage of the ease of use provided by the DB2 ODBC catalog functions and maximize the portion of the application's processing where performance is critical. In order to mix the use of DB2 ODBC and embedded SQL, the application must comply to the following rules:

- All connection management and transaction management must be performed completely using either DB2 ODBC or embedded SQL. Either the DB2 ODBC application performs all the connects and commits/rollback and calls functions written using embedded SQL; or an embedded SQL application performs all the connects and commits/rollback and calls functions written in DB2 ODBC which use a null connection (see “Writing a DB2 ODBC stored procedure” on page 391 for details on null connections).
- Query statement processing must not and cannot straddle across DB2 ODBC and embedded SQL interfaces for the same statement; for example, the application cannot open a cursor in an embedded SQL routine, and then call the DB2 ODBC `SQLFetch()` function to retrieve row data.

Since DB2 ODBC permits multiple connections, the `SQLSetConnection()` function must be called prior to making a function call to a routine written in embedded SQL. This allows the application to explicitly specify the connection under which the embedded SQL routine should perform its processing. If the application only ever sets up one connection, or if the application is written entirely in DB2 ODBC, then calls to `SQLSetConnection()` are not needed.

## Mixed embedded SQL and DB2 ODBC example

The following example demonstrates an application that connects to two data sources, and executes both embedded SQL and dynamic SQL using DB2 ODBC.

```

/* ... */
/* allocate an environment handle */
SQLAllocEnv(&henv);

/* Connect to first data source */
DBconnect(henv, &hdbc[0]);

/* Connect to second data source */
DBconnect(henv, &hdbc[1]);

/***** Start Processing Step *****/
/* NOTE: at this point there are two active connections */

/* set current connection to the first database */
if ( (rc = SQLSetConnection(hdbc[0])) != SQL_SUCCESS )
    printf("Error setting connection 1\n");

/* call function that contains embedded SQL */
if ((rc = Create_Tab() ) != 0)
    printf("Error Creating Table on 1st connection, RC=%ld\n", rc);

/* Commit transaction on connection 1 */
SQLTransact(henv, hdbc[0], SQL_COMMIT);

/* set current connection to the second database */
if ( (rc = SQLSetConnection(hdbc[1])) != SQL_SUCCESS )
    printf("Error setting connection 2\n");

/* call function that contains embedded SQL */
if ((rc = Create_Tab() ) != 0)
    printf("Error Creating Table on 2nd connection, RC=%ld\n", rc);

/* Commit transaction on connection 2 */
SQLTransact(henv, hdbc[1], SQL_COMMIT);

/* Pause to allow the existence of the tables to be verified. */
printf("Tables created, hit Return to continue\n");
getchar();

SQLSetConnection(hdbc[0]);
if (( rc = Drop_Tab() ) != 0)
    printf("Error dropping Table on 1st connection, RC=%ld\n", rc);

```



```

/* Commit transaction on connection 1 */
SQLTransact(henv, hdbc[0], SQL_COMMIT);

SQLSetConnection(hdbc[1]);
if (( rc = Drop_Tab() ) != 0)
    printf("Error dropping Table on 2nd connection, RC=%ld\n", rc);

/* Commit transaction on connection 2 */
SQLTransact(henv, hdbc[1], SQL_COMMIT);

printf("Tables dropped\n");
/***** End Processing Step *****/

/* ... */
/***** Embedded SQL Functions *****/
** This would normally be a separate file to avoid having to **
** keep precompiling the embedded file in order to compile the DB2 CLI **
** section50 **
*****/

EXEC SQL INCLUDE SQLCA;

int
Create_Tab( )
{

    EXEC SQL CREATE TABLE mixedup
        (ID INTEGER, NAME CHAR(10));

    return( SQLCODE);
}

int
Drop_Tab( )
{
    EXEC SQL DROP TABLE mixedup;

    return( SQLCODE);
}
/* ... */

```

---

## Using vendor escape clauses

The X/Open SQL CAE specification defines an *escape clause* as: “a syntactic mechanism for vendor-specific SQL extensions to be implemented in the framework of standardized SQL”. Both DB2 ODBC and ODBC support vendor escape clauses as defined by X/Open.

**Note:** ODBC defines short forms of vendor escape clauses that are not defined by X/Open.

Currently, escape clauses are used extensively by ODBC to define SQL extensions. DB2 ODBC translates the ODBC extensions into the correct DB2 syntax. The `SQLNativeSql()` function can be used to display the resulting syntax.

If an application is only going to access DB2 data sources, then there is no reason to use the escape clauses. If an application is going to access other data sources that offer the same support, but uses different syntax, then the escape clauses increase the portability of the application.

## Escape clause syntax

The format of an X/Open SQL escape clause definition is:

```
--(*vendor(vendor-identifier),  
product(product-identifier) extended SQL text*)--
```

**vendor-identifier** Vendor identification that is consistent across all of that vendor's SQL products (for example, IBM).

**product-identifier** Identifies an SQL product (for example, DB2).

These two parts make up the *SQL-escape-identification*.

## Using ODBC defined SQL extensions

ODBC has used a vendor escape clause of:

```
--(* vendor(Microsoft), product(ODBC) extended SQL text*)--
```

to define the following SQL extensions (these extensions are not defined by X/Open):

- Extended date, time, timestamp data
- Outer join
- LIKE predicate
- Call stored procedure
- Extended scalar functions
  - Numeric functions
  - String functions
  - System functions

ODBC also defines a shorthand syntax for specifying these extensions:

```
{ extended SQL text }
```

X/Open does not support this shorthand syntax; however, it is widely used by ODBC applications.

## ODBC date, time, timestamp data

The ODBC escape clauses for date, time, and timestamp data are:

```
--(*vendor(Microsoft),product(ODBC) d 'value'*)--  
--(*vendor(Microsoft),product(ODBC) t 'value'*)--  
--(*vendor(Microsoft),product(ODBC) ts 'value'*)--
```

**d** indicates *value* is a date in the *yyyy-mm-dd* format,

**t** indicates *value* is a time in the *hh:mm:ss* format

**ts** indicates *value* is a timestamp in the *yyyy-mm-dd hh:mm:ss.ffffff* format.

The shorthand syntax for date, time, and timestamp data is:

```
{d 'value'}  
{t 'value'}  
{ts 'value'}
```

For example, each of the following statements can be used to issue a query against the **EMPLOYEE** table:

```
SELECT * FROM EMPLOYEE  
WHERE HIREDATE=--(*vendor(Microsoft),product(ODBC) d '1994-03-29' *)--
```

```
SELECT * FROM EMPLOYEE WHERE HIREDATE={d '1994-03-29'}
```

DB2 ODBC translates either of the above statements to a DB2 format.

SQLNativeSql() can be used to return the translated statement.

The ODBC escape clauses for date, time, and timestamp literals can be used in input parameters with a C data type of SQL\_C\_CHAR.

## ODBC outer join syntax

The ODBC escape clause for outer join is:

```
--(*vendor(Microsoft),product(ODBC) oj outer join*)--
```

where *outer join* is:

```
table-name {LEFT | RIGHT | FULL} OUTER JOIN  
           {table-name | outer-join}  
           ON search-condition
```

Or alternatively, the ODBC shorthand syntax is:

```
{oj outer-join}
```

For example, DB2 ODBC translates the following two statements:

```
SELECT * FROM  
--(*vendor(Microsoft),product(ODBC) oj  
T1 LEFT OUTER JOIN T2 ON T1.C1=T2.C3*)--  
WHERE T1.C2>20
```

```
SELECT * FROM {oj T1 LEFT OUTER JOIN T2 ON T1.C1=T2.C3}  
WHERE T1.C2>20
```

to IBM's format, which corresponds to the SQL92 outer join syntax.

```
SELECT * FROM T1 LEFT OUTER JOIN T2 ON T1.C1=T2.C3 WHERE T1.C2>20
```

**Note:** Not all servers support outer join. To determine if the current server supports outer joins, call `SQLGetInfo()` with the `SQL_OUTER_JOIN` and `SQL_OJ_CAPABILITIES` options.

## Like predicate escape clauses

In an SQL LIKE predicate, the metacharacter `%` matches zero or more of any character and the metacharacter `_` matches any one character. The ESCAPE clause allows the definition of patterns intended to match values that contain the actual percent and underscore characters by preceding them with an escape character. The escape clause ODBC uses to define the LIKE predicate escape character is:

```
--(*vendor(Microsoft),product(ODBC) escape 'escape-character'*)--
```

where *escape-character* is any character supported by the DB2 rules governing the use of the ESCAPE clause. The shorthand syntax for the LIKE predicate escape character is:

```
{escape 'escape-character'}
```

Applications that are not concerned about portability across different vendor DBMS products should pass the ESCAPE clause directly to the data source. To determine when LIKE predicate escape characters are supported by a particular data source, an application should call `SQLGetInfo()` with the `SQL_LIKE_ESCAPE_CLAUSE` information type.

## Stored procedure CALL syntax

The ODBC escape clause for calling a stored procedure is:

```
--(*vendor(Microsoft),product(ODBC),  
[?]=call procedure-name([[parameter][, [parameter]]...])*)--
```

### **procedure-name**

Specifies the name of a procedure stored at the data source.

### **parameter**

Specifies a procedure parameter.

A procedure can have zero or more parameters. The short form syntax is:

```
{[?]=call procedure-name([[parameter][, [parameter]]...])}
```

(The square brackets ([ ]) indicate optional arguments.)

ODBC specifies the optional parameter `?=` to represent the procedure's return value, which, if present, is stored in the location specified by the first parameter marker defined by `SQLBindParameter()`. DB2 ODBC returns the `SQLCODE` as the procedure's return value if `?=` is present in the escape clause. If `?=` is not present, then the application can retrieve the `SQLCA` by using the `SQLGetSQLCA()` function. Unlike ODBC, DB2 ODBC does not support literals as procedure arguments. Parameter markers must be used.

For more information about stored procedures, see "Using stored procedures" on page 389 or *DB2 Application Programming and SQL Guide*.

For example, DB2 ODBC translates the following two statements:

```
--(*vendor(Microsoft),product(ODBC) CALL NEBT94(?,?,?)*)--  
{CALL NETB94(?,?,?)}
```

To an internal CALL statement format:

```
CALL NEBT94(?, ?, ?)
```

## ODBC scalar functions

Scalar functions such as string length, substring, or trim can be used on columns of a result sets and on columns that restrict rows of a result set. The ODBC escape clauses for scalar functions and its shorthand are:

```
--(*vendor(Microsoft),product(ODBC) fn scalar-function*)--  
or,  
{fn scalar-function}
```

Where, *scalar-function* can be any function listed in Appendix B, “Extended scalar functions” on page 429.

For example, the ODBC driver translates both of the following statements:

```
SELECT --(*vendor(Microsoft),product(ODBC) fn CONCAT(FIRSTNAME, LASTNAME) *)--  
FROM EMPLOYEE
```

```
SELECT {fn CONCAT(FIRSTNAME, LASTNAME)} FROM EMPLOYEE
```

to:

```
SELECT FIRSTNAME CONCAT LASTNAME FROM EMPLOYEE
```

SQLNativeSql() can be called to obtain the translated SQL statement.

To determine which scalar functions are supported by the current server referenced by a specific connection handle, call SQLGetInfo() with the SQL\_NUMERIC\_FUNCTIONS, SQL\_STRING\_FUNCTIONS, SQL\_SYSTEM\_FUNCTIONS, and SQL\_TIMEDATE\_FUNCTIONS options.

---

## Programming hints and tips

This section provides some hints and tips to help improve DB2 ODBC and ODBC application performance and portability.

### Avoiding common initialization file problems

You can avoid two common problems when using the DB2 ODBC initialization file by ensuring that these contents are accurate.

#### Square brackets

The square brackets in the initialization file must consist of the correct EBCDIC characters. The open square bracket must use the hexadecimal characters X'AD'. The close square bracket must use the hexadecimal characters X'BD'. DB2 ODBC does not recognize brackets if coded differently.

### **Sequence numbers**

The initialization file cannot accept sequence numbers. All sequence numbers must be removed.

## **Setting common connection options**

The following connection options might need to be set (or considered) by DB2 ODBC applications.

### **SQL\_AUTOCOMMIT**

Generally this option should be set to `SQL_AUTOCOMMIT_OFF`, since each commit request can generate extra network flow. Only leave `SQL_AUTOCOMMIT_ON` if specifically needed.

**Note:** The default is `SQL_AUTOCOMMIT_ON`.

### **SQL\_TXN\_ISOLATION**

This connection (and statement) option determines the isolation level at which the connection or statement operate. The isolation level determines the level of concurrency possible, and the level of locking required to execute the statement. Applications need to choose an isolation level that maximizes concurrency, yet ensures data consistency.

See Section 5 of *DB2 Application Programming and SQL Guide* for a complete discussion of isolation levels and their effect.

## **Setting common statement options**

The following statement options might need to be set by DB2 ODBC applications.

### **SQL\_MAX\_ROWS**

Setting this option limits the number of rows returned to the application. This can be used to avoid an application from being overwhelmed with a very large result set generated inadvertently, especially for applications on clients with limited memory resources.

**Note:** The full result set is still generated at the server. DB2 ODBC only fetches up to `SQL_MAX_ROWS` rows.

### **SQL\_CURSOR\_HOLD**

This statement option determines if the cursor for this statement is defined with the equivalent of the `CURSOR WITH HOLD` clause.

Resources associated with statement handles can be better utilized by DB2 ODBC if the statements that do not require `CURSOR WITH HOLD` are set to `SQL_CURSOR_HOLD_OFF`.

**Note:** Many ODBC applications expect a default behavior where the cursor position is maintained after a commit.

## SQL\_STMTTXN\_ISOLATION

DB2 ODBC allows the isolation level to be set at the statement level (however, it is best to set the isolation level at the connection level). The isolation level determines the level of concurrency possible, and the level of locking required to execute the statement.

Resources associated with statement handles can be better utilized by DB2 ODBC if statements are set to the required isolation level, rather than leaving all statements at the default isolation level. This should only be attempted with a thorough understanding of the locking and isolation levels of the connected DBMS. See *DB2 SQL Reference* for a complete discussion of isolation levels and their effect.

Applications should use the minimum isolation level possible to maximize concurrency.

## Using SQLSetColAttributes to reduce network flow

Each time a query statement is prepared or executed directly, DB2 ODBC retrieves information about the SQL data type, and its size from the data source. If the application *knows* this information ahead of time, `SQLSetColAttributes()` can be used to provide DB2 ODBC with the information. This can significantly reduce the network flow from remote data sources if the result set coming back contains a very large number (hundreds) of columns.

**Note:** The application must provide DB2 ODBC with exact result descriptor information for ALL columns; otherwise, an error occurs when the data is fetched.

Queries that generate result sets that contain a large number of columns, but relatively small number of rows, have the most to gain from using `SQLSetColAttributes()`.

## Comparing binding and SQLGetData

Generally it is more efficient to bind application variables to result sets than to use `SQLGetData()`. Use `SQLGetData()` when the data value is large variable-length data that:

- Must be received in pieces, or
- Might not need to be retrieved (dependent on another application action.)

## Increasing transfer efficiency

The efficiency of transferring character data between bound application variables and DB2 ODBC can be increased if the *pcbValue* and *rgbValue* arguments are contiguous in memory. (This allows DB2 ODBC to fetch both values with one copy operation.)

For example:

```
struct { SQLINTEGER pcbValue;  
        SQLCHAR   rgbValue[MAX_BUFFER];  
        } column;
```

## Limiting use of catalog functions

In general, try to limit the number of times the catalog functions are called, and limit the number of rows returned.

The number of catalog function calls can be reduced by calling the function once, and storing the information at the application.

The number of rows returned can be limited by specifying a:

- Schema name or pattern for all catalog functions
- Table name or pattern for all catalog functions other than `SQLTables`
- Column name or pattern for catalog functions that return detailed column information.

Remember, although an application can be developed and tested against a data source with hundreds of tables, it can be run against a database with thousands of tables. Plan ahead.

Close any open cursors (call `SQLFreeStmt()` with `SQL_CLOSE`) for statement handles used for catalog queries to release any locks against the catalog tables. Outstanding locks on the catalog tables can prevent `CREATE`, `DROP` or `ALTER` statements from executing.

## Using column names of function generated result sets

The column names of the result sets generated by catalog and information functions can change as the X/Open and ISO standards evolve. The *position* of the columns, however, does not change.

Any application dependency should be based on the column position (*icol* parameter) and not the name.

## Making use of dynamic SQL statement caching

To make use of *dynamic caching* (when the server caches a prepared version of a dynamic SQL statement), the application must use the same statement handle for the same SQL statement.

For example, if an application routinely uses a set of 10 SQL statements, 10 statement handles should be allocated and associated with each of those statements. Do not free the statement handle while the statement can still be executed. (The transaction can still be rolled back or committed without affecting any of the prepared statements). The application continues to prepare and execute the statements in a normal manner. The DB2 server determines if the prepare is actually needed.

To reduce function call overhead, the statement can be prepared once, and executed repeatedly throughout the application.



## Optimizing insertion and retrieval of data

The methods described in “Using arrays to input parameter values” on page 375 and “Retrieving a result set into an array” on page 379 optimize the network flow.

Use these methods as much as possible.

## Optimizing for large object data

Use LOB data types and the supporting functions for long strings whenever possible. Unlike LONG VARCHAR, LONG VARBINARY, and LONG VARGRAPHIC data types, LOB data values can use LOB locators and functions such as SQLGetPosition() and SQLGetSubString() to manipulate large data values at the server.

## Using SQLDriverConnect instead of SQLConnect

Using SQLDriverConnect() gives an application the flexibility to override any or all of the initialization keyword values specified for the target data source.

## Turning off statement scanning

DB2 ODBC by default, scans each SQL statement searching for vendor escape clause sequences.

If the application does not generate SQL statements that contain vendor escape clause sequences ( “Using vendor escape clauses” on page 404), then the SQL\_NO\_SCAN statement option should be set to SQL\_NOSCAN\_ON at the connection level so that DB2 ODBC does not perform a scan for vendor escape clauses.

## Casting distinct types

If a parameter marker is used in the predicate of a query statement, and the parameter is a distinct type, the statement must use a CAST function to cast either the parameter marker or the distinct type.

For example, assume the following distinct type and table are defined:

```
CREATE DISTINCT TYPE CNUM AS INTEGER WITH COMPARISONS
CREATE TABLE CUSTOMER (
    Cust_Num    CNUM NOT NULL,
    First_Name  CHAR(30) NOT NULL,
    Last_Name   CHAR(30) NOT NULL,
    Phone_Num   CHAR(20) WITH DEFAULT,
    PRIMARY KEY (Cust_Num) )
```

The following statement fails because the parameter marker cannot be type CNUM and the comparison fails due to incompatible types:

```
SELECT first_name, last_name, phone_num FROM customer
       where cust_num = ?
```

Alternatively the parameter marker can be cast to INTEGER. The server can then apply the INTEGER to CNUM conversion:

```
SELECT first_name, last_name, phone_num FROM customer
       where cust_num = cast( ? as integer )
```

See *DB2 SQL Reference* for more information about parameter markers (PREPARE statement) and casting (CAST function).

---

## Chapter 7. Problem diagnosis

This section provides guidelines for working with the DB2 ODBC traces and information about general diagnosis, debugging, and abends. You can obtain traces for DB2 ODBC applications and diagnostics and DB2 ODBC stored procedures.

---

### Tracing

DB2 ODBC provides two traces that differ in purpose:

- An application trace intended for debugging user applications, described in “Application trace.”
- A service trace for problem diagnosis, described in “Diagnostic trace” on page 415.

### Application trace

The DB2 ODBC application trace is enabled using the APPLTRACE and APPLTRACEFILENAME keywords in the DB2 ODBC initialization file.

The APPLTRACE keyword is intended for customer application debugging. This trace records data information at the DB2 ODBC API interface; it is specifically designed to trace ODBC API calls. The trace is written to the file specified on the APPLTRACEFILENAME keyword. We strongly recommend that you use this trace

#### Specifying the trace file name

You can use a JCL DD card format or an OS/390 UNIX HFS file name format to specify the APPLTRACEFILENAME keyword setting. The primary use of the JCL DD card format is write to an MVS preallocated sequential data set. You can also specify OS/390 UNIX HFS files on a DD statement. The OS/390 UNIX HFS file name format is used strictly for writing to HFS files.

**JCL DD card format:** The JCL DD card format is APPLTRACEFILENAME="DD:ddname". The ddname value is the name of the DD card specified in your job or TSO logon procedure.

*Examples:* Assume the keyword setting is APPLTRACEFILENAME="DD:APPLDD". You can use the following JCL DD statement examples in your job or TSO logon procedure to specify the MVS trace data set.

*Example 1:* Write to preallocated MVS sequential data set USER01.MYTRACE.

```
//APPLDD DD DISP=SHR,DSN=USER01.MYTRACE
```

*Example 2:* Write to preallocated OS/390 UNIX HFS file MYTRACE in directory /usr/db2.

```
//APPLDD DD PATH='/usr/db2/MYTRACE'
```

*Example 3:* Allocate OS/390 UNIX HFS file MYTRACE in directory /usr/db2 specifying permission for the file owner to read from (SIRUSR) and write to (SIWUSR) the trace file:

```
//APPLDD DD PATH='/usr/db2/MYTRACE',  
PATHOPTS=(ORDWR,OCREAT,OTRUNC),  
PATHMODE=(SIRUSR,SIWUSR)
```

**OS/390 UNIX HFS file name format:** The OS/390 UNIX HFS file name format is APPLTRACEFILENAME=hfs\_filename. The hfs\_filename value specifies the path and file name for the HFS file. The HFS file does not have to be preallocated. If the file name does not exist in the specified directory, the file is dynamically allocated.

*Examples:* The following examples use the APPLTRACEFILENAME keyword to specify an OS/390 UNIX HFS trace file.

*Example 1:* Create and write to HFS file named APPLTRC1 in the fully qualified directory /usr/db2.

```
APPLTRACEFILENAME=/usr/db2/APPLTRC1
```

*Example 2:* Create and write to HFS file named APPLTRC1 in the current working directory of the application.

```
APPLTRACEFILENAME=./APPLTRC1
```

*Example 3:* Create and write to HFS file named APPLTRC1 in the parent directory of the current working directory.

```
APPLTRACEFILENAME=../APPLTRC1
```

### **Application trace output**

The following example of application trace output shows how DB2 ODBC follows the APIs invoked, indicates values used, data pointers, etc. Errors are also indicated.

```

SQLAllocEnv( phEnv=&6b7e77c )
SQLAllocEnv( phEnv=1 )
----> SQL_SUCCESS

SQLAllocConnect( hEnv=1, phDbc=&6b7e778 )
SQLAllocConnect( phDbc=1 )
----> SQL_SUCCESS

SQLConnect( hDbc=1, szDSN=NULL Pointer, cbDSN=0, szUID=NULL Pointer, cbUID=0,
szAuthStr=NULL Pointer, cbAuthStr=0 )
SQLConnect( )
----> SQL_SUCCESS

SQLAllocStmt( hDbc=1, phStmt=&6b7e774 )
SQLAllocStmt( phStmt=1 )
----> SQL_SUCCESS

SQLExecDirect( hStmt=1, pszSqlStr="SELECT NAME FROM SYSIBM.SYSPLAN", cbSqlStr=-3 )
SQLExecDirect( )
----> SQL_SUCCESS

SQLFetch( hStmt=1 )
SQLFetch( )
----> SQL_SUCCESS

SQLTransact( hEnv=1, hDbc=1, fType=SQL_COMMIT )
SQLTransact( )
----> SQL_SUCCESS

SQLFreeStmt( hStmt=1, fOption=SQL_DROP )
SQLFreeStmt( )
----> SQL_SUCCESS

SQLDisconnect( hDbc=1 )
SQLDisconnect( )
----> SQL_SUCCESS

SQLFreeConnect( hDbc=1 )
SQLFreeConnect( )
----> SQL_SUCCESS

SQLFreeEnv( hEnv=1 )
SQLFreeEnv( )
----> SQL_SUCCESS

```

For more information about how to specify the APPLTRACE and APPLTRACEFILENAME keywords, see “DB2 ODBC initialization file” on page 60.

## Diagnostic trace

The DB2 ODBC diagnostic trace captures information to use in DB2 ODBC problem determination. The trace is intended for use under the direction of the IBM Support Center; it is not intended to assist in debugging user written DB2 ODBC applications. You can view this trace to obtain information about the general flow of an application, such as commit information. However, this trace is intended for IBM service information only and is therefore subject to change.

You can activate the diagnostic trace by either the DSNAOTRC command or the DIAGTRACE keyword in the DB2 ODBC initialization file.

If you activate the diagnostic trace using the `DIAGTRACE` keyword in the initialization file, you must also allocate a `DSNAOTRC` DD statement in your job or TSO logon procedure. You can use one of the following methods to allocate a `DSNAOTRC` DD statement:

- Specify a `DSNAOTRC` DD JCL statement in your job or TSO logon procedure
- Use the TSO/E `ALLOCATE` command
- Use dynamic allocation in your ODBC application

### Specifying the diagnostic trace file

The diagnostic trace data can be written to an MVS sequential data set or an OS/390 UNIX HFS file.

An MVS data set must be preallocated with the following data set attributes:

- Sequential data set organization
- Fixed-block 80 record format

When you execute an ODBC application in OS/390 UNIX and activate the diagnostic trace using the `DIAGTRACE` keyword in the initialization file, DB2 writes the diagnostic data to a dynamically allocated file, `DD:DSNAOTRC`. This file is located in the current working directory of the application if the `DSNAOTRC` DD statement is not available to the ODBC application. You can format `DD:DSNAOTRC` using the trace formatting program.

**Examples:** The following JCL examples use a `DSNAOTRC` JCL DD card to specify the diagnostic trace file.

*Example 1:* Write to preallocated MVS sequential data set `USER01.DIAGTRC`.

```
//DSNAOTRC DD      DISP=SHR,DSN=USER01.DIAGTRC
```

*Example 2:* Write to preallocated OS/390 UNIX HFS file `DIAGTRC` in directory `/usr/db2`.

```
//DSNAOTRC      DD      PATH='/usr/db2/DIAGTRC'
```

*Example 3:* Allocate OS/390 UNIX HFS file `DIAGTRC` in directory `/usr/db2` specifying permission for the file owner to read from (`SIRUSR`) and write to (`SIWUSR`) the trace file.

```
//DSNAOTRC      DD      PATH='/usr/db2/DIAGTRC',  
                    PATHOPTS=(ORDWR,OCREAT,OTRUNC),  
                    PATHMODE=(SIRUSR,SIWUSR)
```

For more information about the `DIAGTRACE` keyword, see “DB2 ODBC initialization file” on page 60.

### Using the diagnostic trace command: `DSNAOTRC`

You can use the `DSNAOTRC` command to:

- Manually start or stop the recording of memory resident diagnostic trace records.
- Query the current status of the diagnostic trace.
- Capture the memory resident trace table to an MVS data set or OS/390 UNIX HFS file.

- Format the DB2 ODBC diagnostic trace.

**Special OS/390 UNIX considerations:** You can issue the DSNAOTRC command from the OS/390 UNIX shell command line to activate the diagnostic trace prior to executing an ODBC application. Under the direction of IBM support only, you must store the dsnaotrc program load module in an OS/390 UNIX HFS file.

Use the TSO/E command, OPUTX, to store the dsnaotrc load module in an HFS file. The following example uses the OPUTX command to store load module dsnaotrc from MVS partitioned data set DB2A.DSNLOAD to HFS file DSNAOTRC in directory /usr/db2:

```
OPUTX 'DB2A.DSNLOAD(DSNAOTRC)' /usr/db2/dsnaotrc
```

After storing the dsnaotrc program module in an HFS file, follow these steps at the OS/390 UNIX shell to activate, dump, and format the diagnostic trace:

1. Enable the shared address space environment variable for the OS/390 UNIX shell. Issue the following export statement at the command line or specify it in your \$HOME/.profile:

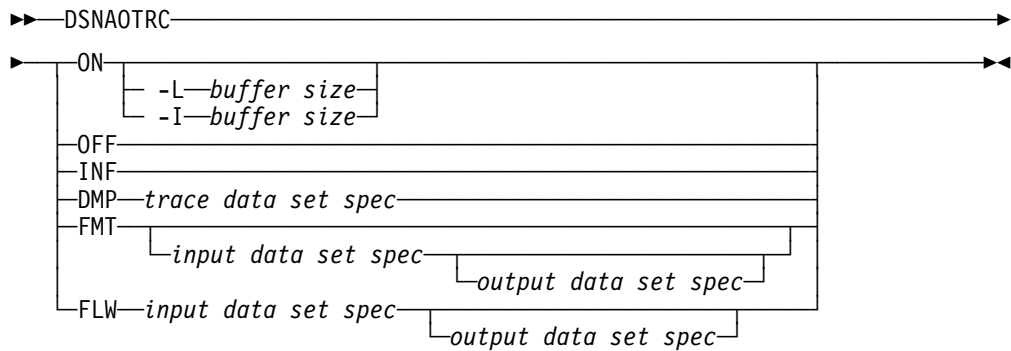
```
export _BPX_SHAREAS=YES
```

Setting this environment variable allows the OMVS command and the OS/390 UNIX shell to run in the same TSO address space.

2. Go to the directory that contains the dsnaotrc module.
3. Verify that execute permission is established for the dsnaotrc load module. If execute permission was not granted, use the chmod command to set execute permission for the dsnaotrc load module.
4. Issue dsnaotrc on. The options for activating the diagnostic trace are optional.
5. Execute the ODBC application.
6. Issue dsnaotrc dmp "raw\_trace\_file". The raw\_trace\_file value is the name of the output file to which DB2 writes the raw diagnostic trace data.
7. Issue dsnaotrc off to deactivate the diagnostic trace.
8. Issue dsnaotrc fmt "raw\_trace\_file" "fmt\_trace\_file" to format the raw trace data records from input file "raw\_trace\_file" to output file "fmt\_trace\_file".

After successfully formatting the diagnostic trace data, delete the dsnaotrc program module from your OS/390 UNIX directory. Do not attempt to maintain a private copy of the dsnaotrc program module in your HFS directory.

## Syntax



## Option descriptions

### ON

Start the DB2 ODBC diagnostic trace.

#### **-L** *buffer size*

L = Last. The trace wraps; it captures the last, most current trace records.

*buffer size* is the number of bytes to allocate for the trace buffer. This value is required. The buffer size is rounded to a multiple of 65536 (64K).

#### **-I** *buffer size*

I = Initial. The trace does not wrap; it captures the initial trace records.

*buffer size* is the number of bytes to allocate for the trace buffer. This value is required. The buffer size is rounded to a multiple of 65536 (64K).

### OFF

Stop the DB2 ODBC diagnostic trace.

### INF

Display information about the currently active DB2 ODBC diagnostic trace.

### DMP

Dump the currently active DB2 ODBC diagnostic trace.

#### *trace data set spec*

Specifies the MVS data set or OS/390 UNIX HFS file to which DB2 writes the raw DB2 ODBC diagnostic trace data. The data set specification can be either an MVS data set name, an OS/390 UNIX HFS file name, or a currently allocated JCL DD card name.

### FMT

Generate a formatted detail report of the DB2 ODBC diagnostic trace contents.

### FLW

Generate a formatted flow report of the DB2 ODBC diagnostic trace contents.

#### *input data set spec*

The data set that contains the raw DB2 ODBC diagnostic trace data to be formatted. This is the data set that was generated as the result of a DSNAOTRC DMP command or due to the DSNAOTRC DD card if the trace was started by using the DIAGTRACE initialization keyword. The data set specification can be either an MVS data set name, an OS/390 UNIX HFS file name, or a currently allocated JCL DD card name. If this



parameter is not specified, then DSNAOTRC attempts to format the memory resident DSNAOTRC that is currently active.

#### *output data set spec*

The data set to which the formatted DB2 ODBC diagnostic trace records are written. The data set specification can be either an MVS data set name, an OS/390 UNIX HFS file name, or a currently allocated JCL DD card name. If you specify an MVS data set or OS/390 UNIX HFS file that does not exist, DB2 allocates it dynamically. If this parameter is not specified, the output is written to standard output ("STDOUT").

**Examples:** The following examples show how to code the data set specifications.

- Trace data set specification:

*Example 1:* Currently allocated JCL DD card name TRACEDD.

```
DSNAOTRC DMP DD:TRACEDD
```

*Example 2:* MVS sequential data set USER01.DIAGTRC.

```
DSNAOTRC DMP "USER01.DIAGTRC"
```

*Example 3:* OS/390 UNIX HFS file named DIAGTRC in directory /usr/db2:

```
DSNAOTRC DMP "/usr/db2/DIAGTRC"
```

- Input data set specification:

*Example 1:* Currently allocated JCL DD card name INPDD.

```
DSNAOTRC FLW DD:INPDD output-dataset-spec
```

*Example 2:* MVS sequential data set USER01.DIAGTRC.

```
DSNAOTRC FLW "USER01.DIAGTRC" output-dataset-spec
```

*Example 3:* OS/390 UNIX HFS file DIAGTRC in directory /usr/db2.

```
DSNAOTRC FLW "/usr/db2/DIAGTRC" output-dataset-spec
```

- Output data set specification:

*Example 1:* Currently allocated JCL DD card name OUTPDD.

```
DSNAOTRC FLW input-dataset-spec DD:OUTPDD
```

*Example 2:* MVS sequential data set USER01.TRCFLOW.

```
DSNAOTRC FLW input-dataset-spec "USER01.TRCFLOW"
```

*Example 3:* OS/390 UNIX HFS file TRCFLOW in directory /usr/db2.

```
DSNAOTRC FLW input-dataset-spec "/usr/db2/TRCFLOW"
```

## Stored procedure trace

This section describes the steps required to obtain an application trace or a diagnostic trace of a DB2 ODBC stored procedure. DB2 ODBC stored procedures run in either a DB2-established stored procedures address space or a WLM-established address space. Both the main application that calls the stored procedure (client application), and the stored procedure itself, can be either a DB2 ODBC application or a standard DB2 precompiled application.

If the client application and the stored procedure are DB2 ODBC application programs, you can trace:

- A client application only

- A stored procedure only
- Both the client application and stored procedure

More than one address spaces can not share write access to a single data set. Therefore, you must use the appropriate JCL DD statements to allocate a unique trace data set for each stored procedures address space that uses the DB2 ODBC application trace or diagnostic trace.

## Tracing a client application

This section explains how to obtain an application trace and a diagnostic trace for a client application.

**Application trace:** Follow these steps to obtain an application trace.

1. Set APPLTRACE=1 and APPLTRACEFILENAME="DD:DDNAME" in the common section of the DB2 ODBC initialization file as follows:

```
[COMMON]
APPLTRACE=1
APPLTRACEFILENAME="DD:APPLTRC"
```

DDNAME is the name of an OS/390 JCL DD statement specified in the JCL for the application job or your TSO logon procedure.

2. Specify an OS/390 JCL DD statement in the JCL for the application job or your TSO logon procedure. The DD statement references a pre-allocated OS/390 sequential data set with DCB attributes RECFM=VBA,LRECL=137, an OS/390 UNIX HFS file to contain the client application trace, as shown in the following examples:

```
//APPLTRC DD DISP=SHR,DSN=CLI.APPLTRC
//APPLTRC DD PATH='/u/cli/appltrc'
```

**Diagnostic trace:** When tracing only the client application, you can activate the diagnostic trace by using the DIAGTRACE keyword in the DB2 ODBC initialization file or the DSNAOTRC command. See “Diagnostic trace” on page 415 for information about obtaining a diagnostic trace of the client application.

## Tracing a stored procedure

This section explains how to obtain an application trace and a diagnostic trace for a stored procedure.

**Application trace:** Follow these steps to obtain an application trace.

1. Set APPLTRACE=1 and APPLTRACEFILENAME="DD:DDNAME" in the common section of the DB2 ODBC initialization file as follows:

```
[COMMON]
APPLTRACE=1
APPLTRACEFILENAME="DD:APPLTRC"
```

DDNAME is the name of an OS/390 JCL DD statement specified in the JCL for the stored procedures address space.

2. Specify an OS/390 JCL DD statement in the JCL for the stored procedures address space. The DD statement references a pre-allocated OS/390 sequential data set with DCB attributes RECFM=VBA,LRECL=137 or OS/390 UNIX HFS file to contain the client application trace, as shown in the following examples:

```
//APPLTRC DD DISP=SHR,DSN=CLI.APPLTRC
```

```
//APPLTRC DD PATH='/u/cli/appltrc'
```

**Diagnostic trace:** Follow these steps to obtain a diagnostic trace.

1. Set DIAGTRACE=1, DIAGTRACE\_BUFFER\_SIZE=nnnnnn, and DIAGTRACE\_NO\_WRAP=0|1 in the common section of the DB2 ODBC initialization file. For example:

```
[COMMON]
DIAGTRACE=1
DIAGTRACE_BUFFER_SIZE=2000000
DIAGTRACE_NO_WRAP=1
```

nnnnnn is the number of bytes to allocate for the diagnostic trace buffer.

2. Specify an OS/390 DSNAOINI JCL DD statement in the JCL for the stored procedures address space. The DD statement references the DB2 ODBC initialization file, as shown in the following examples:

```
//DSNAOINI DD DISP=SHR,DSN=CLI.DSNAOINI
//DSNAOINI DD PATH='/u/cli/dsnaoini'
```

3. Specify an OS/390 DSNAOTRC JCL DD statement in the JCL for the stored procedures space. The DD statement references a pre-allocated OS/390 sequential data set with DCB attributes RECFM=FB,LRECL=80, or an OS/390 UNIX HFS file to contain the unformatted diagnostic data, as shown in the following examples:

```
//DSNAOTRC DD DISP=SHR,DSN=CLI.DIAGTRC
//DSNAOTRC DD PATH='/u/cli/diagtrc'
```

4. Execute the client application that calls the stored procedure.
5. After the DB2 ODBC stored procedure executes, stop the stored procedures address space.
  - For DB2-established address spaces, use the DB2 command, STOP PROCEDURE.
  - For WLM-established address spaces operating in WLM goal mode, use the MVS command, "VARY WLM,APPLENV=name,QUIESCE". name is the WLM application environment name.
  - For WLM-established address spaces operating in WLM compatibility mode, use the MVS command, "CANCEL address-space-name". address-space-name is the name of the WLM-established address space.
6. You can submit either the formatted or unformatted diagnostic trace data to the IBM Support Center. To format the raw trace data at your site, run the DSNAOTRC FMT or DSNAOTRC FLW command against the diagnostic trace data set.

## Tracing both a client application and a stored procedure

This section explains how to obtain an application trace and a diagnostic trace for both a client application and a stored procedure.

**Application trace:** Follow these steps to obtain an application trace.

1. Set APPLTRACE=1 and APPLTRACEFILENAME="DD:DDNAME" in the common section of the DB2 ODBC initialization file as follows:

```
[COMMON]
APPLTRACE=1
APPLTRACEFILENAME="DD:APPLTRC"
```

DDNAME is the name of an OS/390 JCL DD statement specified in both the JCL for the client application job and the stored procedures address space.

2. Specify an OS/390 JCL DD statement in the JCL for the client application. The DD statement references a pre-allocated OS/390 sequential data set with DCB attributes RECFM=VBA,LRECL=137, or an OS/390 UNIX HFS file to contain the client application trace, as shown in the following examples:

```
//APPLTRC DD DISP=SHR,DSN=CLI.APPLTRC.CLIENT
//APPLTRC DD PATH='/u/cli/appltrc.client'
```

You must allocate a separate application trace data set, or an HFS file for the client application. Do not attempt to write to the same application trace data set or HFS file used for the stored procedure.

3. Specify an OS/390 JCL DD statement in the JCL for the stored procedures address space. The DD statement references a pre-allocated OS/390 sequential data set, or an OS/390 UNIX HFS file to contain the stored procedure application trace, as shown in the following examples:

```
//APPLTRC DD DISP=SHR,DSN=CLI.APPLTRC.SPROC
//APPLTRC DD PATH='/u/cli/appltrc.sproc'
```

You must allocate a separate trace data set or HFS file for the stored procedure. Do not attempt to write to the same application trace data set or HFS file used for the client application.

**Diagnostic trace:** Follow these steps to obtain a diagnostic trace.

1. Set DIAGTRACE=1, DIAGTRACE\_BUFFER\_SIZE=nnnnnn, and DIAGTRACE\_NO\_WRAP=0|1 in the common section of the DB2 ODBC initialization file. For example:

```
[COMMON]
DIAGTRACE=1
DIAGTRACE_BUFFER_SIZE=2000000
DIAGTRACE_NO_WRAP=1
```

nnnnnn is the number of bytes to allocate for the diagnostic trace buffer.

2. Specify an OS/390 DSNAOINI JCL DD statement in the JCL for the stored procedures address space. The DD statement references the DB2 ODBC initialization file, as shown in the following examples:
3. Specify an OS/390 DSNAOTRC JCL DD statement in JCL for the client application job. The DD statement references a pre-allocated OS/390 sequential data set with DCB attributes RECFM=FB,LRECL=80, or an OS/390 UNIX HFS file to contain the unformatted diagnostic data, as shown in the following examples:

```
//DSNAOTRC DD DISP=SHR,DSN=CLI.DIAGTRC.CLIENT
//DSNAOTRC DD PATH='/u/cli/diagtrc.client'
```

4. Specify an OS/390 DSNAOTRC JCL DD statement in the JCL for the stored procedures address space. The DD statement references a pre-allocated

OS/390 sequential data set with DCB attributes RECFM=FB,LRECL=80, or an OS/390 UNIX HFS file to contain the stored procedure's unformatted diagnostic data, as shown in the following examples:

```
//DSNAOTRC DD DISP=SHR,DSN=CLI.DIAGTRC.SPROC
```

```
//DSNAOTRC DD PATH='/u/cli/diagtrc.sproc'
```

5. Execute the client application that calls the stored procedure.
6. After the DB2 ODBC stored procedure executes, stop the stored procedures address space.
  - For DB2-established address spaces, use the DB2 command, STOP PROCEDURE.
  - For WLM-established address spaces operating in WLM goal mode, use the MVS command, "VARY WLM,APPLENV=name,QUIESCE". name is the WLM application environment name.
  - For WLM-established address spaces operating in WLM compatibility mode, use the MVS command, "CANCEL address-space-name". address-space-name is the name of the WLM-established address space.
7. You can submit either the formatted or unformatted diagnostic trace data to the IBM Support Center. To format the raw trace data at your site, run the DSNAOTRC FMT or DSNAOTRC FLW command against the client application's diagnostic trace data set and the stored procedure's diagnostic trace data set.

---

## Debugging

You can debug DB2 for OS/390 ODBC applications debug tool shipped with your the C or C++ language compiler. For detailed instructions on debugging DB2 stored procedures, including DB2 ODBC stored procedures, see Section 7 of *DB2 Application Programming and SQL Guide*.

## Abnormal termination

Language Environment reports abends since DB2 ODBC runs under Language Environment. Typically, Language Environment reports the type of abend that occurs and the function that is active in the address space at the time of the abend.

DB2 ODBC has no facility for abend recovery. When an abend occurs, DB2 ODBC terminates. DBMSs follow the normal recovery process for any outstanding DB2 unit of work.

"CEE" is the prefix for all Language Environment messages. If the prefix of the active function is "CLI", then DB2 ODBC had control during the abend which indicates that this can be a DB2 ODBC, a DB2, or a user error.

The following example shows an abend:

```
CEE3250C The system or user abend S04E R=00000000 was issued.  
From entry point CLI_mvscallProcedure(CLI_CONNECTINFO*,...  
+091A2376 at address 091A2376...
```

In this message, you can determine what caused the abend as follows:

- "CEE" indicates that Language Environment is reporting the abend.
- The entry point shows that DB2 ODBC is the active module.
- Abend code "S04E" means that this is a DB2 system abend.

For further information on debugging, see *OS/390 Language Environment for OS/390 & VM Debugging Guide*. For further information on the DB2 recovery process, see Section 4 (Volume 1) of *DB2 Administration Guide*.

## Internal error code

DB2 ODBC provides an internal error code for ODBC diagnosis that is intended for use under the guidance of IBM service. This unique error location, ERRLOC, is a good tool for APAR searches. The following example of a failed `SQLAllocConnect()` shows an error location:

```
DB2 ODBC Sample SQLError Information
DB2 ODBC Sample SQLSTATE           : 58004
DB2 ODBC Sample Native Error Code  : -99999
DB2 ODBC Sample Error message text:
  {DB2 for OS/390}{ODBC Driver} SQLSTATE=58004 ERRLOC=2:170:4;
  RRS "IDENTIFY" failed using DB2 system:V61A,
  RC=08 and REASON=00f30091
```

# Appendix A. DB2 ODBC and ODBC

This appendix explains the differences between DB2 ODBC and ODBC in the following areas:

- “DB2 ODBC and ODBC drivers”
- “ODBC APIs and data types” on page 426
- “Isolation levels” on page 428

For a complete list of functions that DB2 ODBC and ODBC support, see Table 10 on page 74.

## DB2 ODBC and ODBC drivers

This section discusses the support provided by the ODBC driver, and how it differs from DB2 ODBC.

Figure 20 below compares DB2 ODBC and the DB2 ODBC driver.

1. An ODBC driver under the ODBC driver manager
2. DB2 ODBC, callable interface designed for DB2 specific applications.

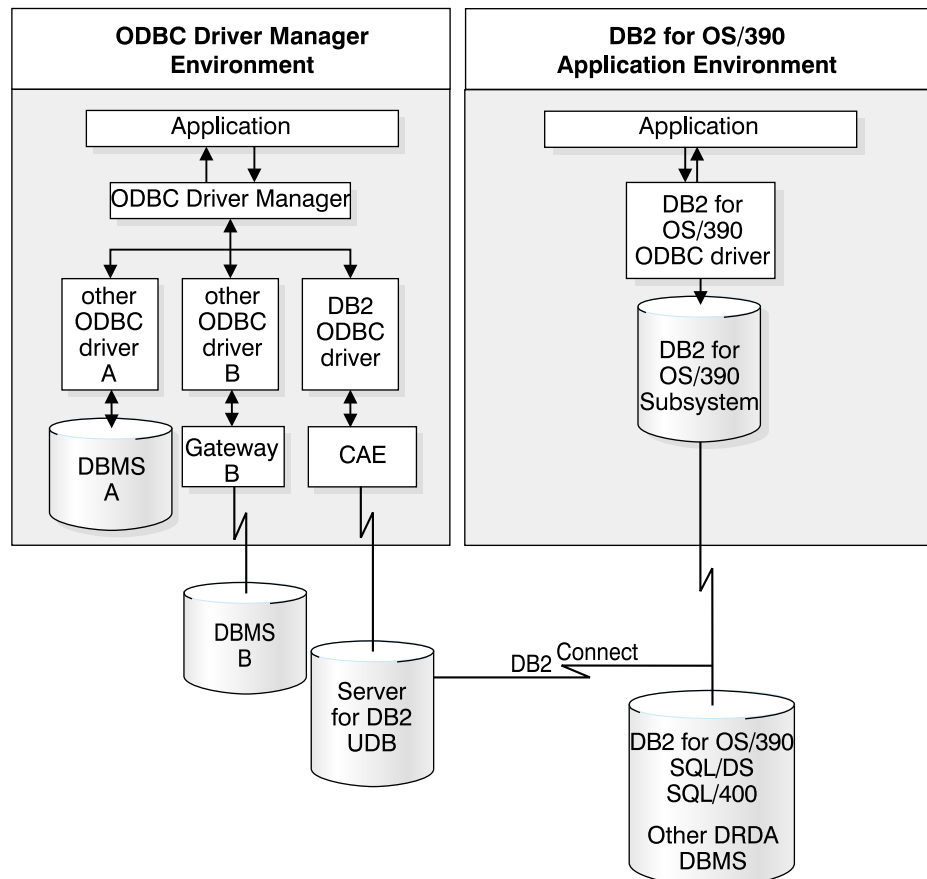


Figure 20. DB2 ODBC and ODBC.

In an ODBC environment, the driver manager provides the interface to the application. It also dynamically loads the necessary *driver* for the database server

## ODBC APIs and data types

to which the application connects. It is the driver that implements the ODBC function set, with the exception of some extended functions implemented by the driver manager.

The DB2 ODBC driver does not execute in this environment. Rather, DB2 ODBC is a self-sufficient driver which supports a subset of the functions provided by the ODBC driver.

DB2 ODBC applications interface directly with the ODBC driver which executes within the application address space. Applications do not interface with a driver manager. The capabilities provided to the application are a subset of the Microsoft ODBC Version 2 specifications.

---

## ODBC APIs and data types

Table 148 summarizes the ODBC Version 2 application programming interfaces, ODBC SQL data types and ODBC C data types and whether those functions and data types are supported by DB2 ODBC. Table 10 on page 74 provides a complete list of functions supported by DB2 ODBC and ODBC 2.0.

*Table 148 (Page 1 of 2). DB2 ODBC support*

<b>ODBC features</b>	<b>DB2 ODBC</b>
Core level functions	All
Level 1 functions	All
Level 2 functions	All, except for: <ul style="list-style-type: none"><li>• SQLBrowseConnect()</li><li>• SQLDrivers()</li><li>• SQLSetPos()</li><li>• SQLSetScrollOptions()</li></ul>
Additional DB2 ODBC functions	<ul style="list-style-type: none"><li>• SQLCancel()</li><li>• SQLSetConnection()</li><li>• SQLGetEnvAttr()</li><li>• SQLSetEnvAttr()</li><li>• SQLSetColAttributes()</li><li>• SQLGetSQLCA()</li></ul>
Minimum SQL data types	<ul style="list-style-type: none"><li>• SQL_CHAR</li><li>• SQL_LONGVARCHAR</li><li>• SQL_VARCHAR</li></ul>
Core SQL data types	<ul style="list-style-type: none"><li>• SQL_DECIMAL</li><li>• SQL_NUMERIC</li><li>• SQL_SMALLINT</li><li>• SQL_INTEGER</li><li>• SQL_REAL</li><li>• SQL_FLOAT</li><li>• SQL_DOUBLE</li></ul>



Table 148 (Page 2 of 2). DB2 ODBC support

ODBC features	DB2 ODBC
Extended SQL data types	<ul style="list-style-type: none"> <li>• SQL_BIT</li> <li>• SQL_TINYINT</li> <li>• SQL_BIGINT (NOT SUPPORTED)</li> <li>• SQL_BINARY</li> <li>• SQL_BLOB</li> <li>• SQL_BLOB_LOCATOR</li> <li>• SQL_CLOB</li> <li>• SQL_CLOB_LOCATOR</li> <li>• SQL_DATE</li> <li>• SQL_DBCLOB</li> <li>• SQL_DBCLOB_LOCATOR</li> <li>• SQL_LONGVARBINARY</li> <li>• SQL_ROWID</li> <li>• SQL_TIME</li> <li>• SQL_TIMESTAMP</li> <li>• SQL_VARBINARY</li> </ul>
ODBC Version 3 SQL data types	<ul style="list-style-type: none"> <li>• SQL_GRAPHIC</li> <li>• SQL_LONGVARGRAPHIC</li> <li>• SQL_VARGRAPHIC</li> </ul>
Core C data types	<ul style="list-style-type: none"> <li>• SQL_C_CHAR</li> <li>• SQL_C_DOUBLE</li> <li>• SQL_C_FLOAT</li> <li>• SQL_C_LONG(SLONG, ULONG)</li> <li>• SQL_C_SHORT (SSHORT, USHORT)</li> </ul>
Extended C data types	<ul style="list-style-type: none"> <li>• SQL_C_BINARY</li> <li>• SQL_C_BIT</li> <li>• SQL_C_BLOB_LOCATOR</li> <li>• SQL_C_CLOB_LOCATOR</li> <li>• SQL_C_DATE</li> <li>• SQL_C_DBCLOB_LOCATOR</li> <li>• SQL_C_TIME</li> <li>• SQL_C_TIMESTAMP</li> <li>• SQL_C_TINYINT</li> </ul>
ODBC Version 3 C data types	<ul style="list-style-type: none"> <li>• SQL_C_DBCHAR (SQL_C_WCHAR)</li> </ul>
Return codes	<ul style="list-style-type: none"> <li>• SQL_SUCCESS</li> <li>• SQL_SUCCESS_WITH_INFO</li> <li>• SQL_NEED_DATA</li> <li>• SQL_NO_DATA_FOUND</li> <li>• SQL_ERROR</li> <li>• SQL_INVALID_HANDLE</li> </ul>
SQLSTATEs	Mapped to X/Open SQLSTATEs with additional IBM SQLSTATEs
Multiple connections per application	Supported but type 1 connections, SQL_CONNECTTYPE = SQL_CONCURRENT_TRANS. Must be on a transaction boundary prior to an SQLConnect or SQLSetConnection.

For more information about ODBC, see *Microsoft ODBC 3.0 Programmer's Reference and SDK Guide*.

---

### Isolation levels

The following table maps IBM RDBMs isolation levels to ODBC transaction isolation levels. The `SQLGetInfo()` function, indicates which isolation levels are available.

*Table 149. Isolation levels under ODBC*

<b>IBM isolation level</b>	<b>ODBC isolation level</b>
Cursor stability	SQL_TXN_READ_COMMITTED
Repeatable read	SQL_TXN_SERIALIZABLE_READ
Read stability;	SQL_TXN_REPEATABLE_READ
Uncommitted read	SQL_TXN_READ_UNCOMMITTED
No commit	(no equivalent in ODBC)

**Note:** `SQLSetConnectOption()` and `SQLSetStmtOption` return `SQL_ERROR` with an `SQLSTATE` of **S1009** if you try to set an unsupported isolation level.

---

## Appendix B. Extended scalar functions

The following functions are defined by ODBC using vendor escape clauses. Each function can be called using the escape clause syntax, or calling the equivalent DB2 function.

These functions are presented in the following categories:

- “String functions” on page 430
- “Date and time functions” on page 431
- “System functions” on page 431

For more information about vendor escape clauses, see “ODBC scalar functions” on page 407.

All errors detected by the following functions, when connected to a DB2 UDB Version 2 server, return SQLSTATE 38552. The text portion of the message is of the form SYSFUN:*nn* where *nn* is one of the following reason codes:

- 01** Numeric value out of range
- 02** Division by zero
- 03** Arithmetic overflow or underflow
- 04** Invalid date format
- 05** Invalid time format
- 06** Invalid timestamp format
- 07** Invalid character representation of a timestamp duration
- 08** Invalid interval type (must be one of 1, 2, 4, 8, 16, 32, 64, 128, 256)
- 09** String too long
- 10** Length or position in string function out of range
- 11** Invalid character representation of a floating point number

---

## String functions

The string functions in this section are supported by DB2 ODBC and defined by ODBC using vendor escape clauses.

**Note:**

- Character string literals used as arguments to scalar functions must be bounded by single quotes.
- Arguments denoted as *string\_exp* can be the name of a column, a string literal, or the result of another scalar function, where the underlying data type can be represented as SQL\_CHAR, SQL\_VARCHAR, or SQL\_LONGVARCHAR.
- Arguments denoted as *start*, *length*, *code* or *count* can be a numeric literal or the result of another scalar function, where the underlying data type is integer based (SQL\_SMALLINT, SQL\_INTEGER).
- The first character in the string is considered to be at position 1.

**ASCII( *string\_exp* )**

Returns the ASCII code value of the leftmost character of *string\_exp* as an integer.

**CONCAT( *string\_exp1*, *string\_exp2* )**

Returns a character string that is the result of concatenating *string\_exp2* to *string\_exp1*.

**INSERT( *string\_exp1*, *start*, *length*, *string\_exp2* )**

Returns a character string where *length* number of characters beginning at *start* is replaced by *string\_exp2* which contains *length* characters.

**LEFT( *string\_exp*,*count* )**

Returns the leftmost *count* of characters of *string\_exp*.

**LENGTH( *string\_exp* )**

Returns the number of characters in *string\_exp*, excluding trailing blanks and the string termination character.

**REPEAT( *string\_exp*, *count* )**

Returns a character string composed of *string\_exp* repeated *count* times.

**RIGHT( *string\_exp*, *count* )**

Returns the rightmost count of characters of *string\_exp*.

**SUBSTRING( *string\_exp*, *start*, *length* )**

Returns a character string that is derived from *string\_exp* beginning at the character position specified by *start* for *length* characters.

---

## Date and time functions

The date and time functions in this section are supported by DB2 ODBC and defined by ODBC using vendor escape clauses.

**Note:**

- Arguments denoted as *timestamp\_exp* can be the name of a column, the result of another scalar function, or a time, date, or timestamp literal.
- Arguments denoted as *date\_exp* can be the name of a column, the result of another scalar function, or a date or timestamp literal, where the underlying data type can be character based, or date or timestamp based.
- Arguments denoted as *time\_exp* can be the name of a column, the result of another scalar function, or a time or timestamp literal, where the underlying data types can be character based, or time or timestamp based.

**CURDATE()**

Returns the current date as a date value.

**CURTIME()**

Returns the current local time as a time value.

**DAYOFMONTH ( *date\_exp* )**

Returns the day of the month in *date\_exp* as an integer value in the range of 1-31.

**HOUR( *time\_exp* )**

Returns the hour in *time\_exp* as an integer value in the range of 0-23.

**MINUTE( *time\_exp* )**

Returns the minute in *time\_exp* as integer value in the range of 0-59.

**MONTH( *date\_exp* )**

Returns the month in *date\_exp* as an integer value in the range of 1-12.

**NOW()**

Returns the current date and time as a timestamp value.

**SECOND( *time\_exp* )**

Returns the second in *time\_exp* as an integer value in the range of 0-59.

---

## System functions

The system functions in this section are supported by DB2 ODBC and defined by ODBC using vendor escape clauses.

- Arguments denoted as *exp* can be the name of a column, the result of another scalar function, or a literal.

## System functions

- Arguments denoted as *value* can be a literal constant.

### **DATABASE()**

Returns the name of the database corresponding to the connection handle (*hdbc*). (The name of the database is also available using `SQLGetInfo()` by specifying the information type `SQL_DATABASE_NAME`.)

### **IFNULL( *exp*, *value* )**

If *exp* is null, *value* is returned. If *exp* is not null, *exp* is returned. The possible data types of *value* must be compatible with the data type of *exp*.

### **USER()**

Returns the user's authorization name. (The user's authorization name is also available using `SQLGetInfo()` by specifying the information type `SQL_USER_NAME`.)

## Appendix C. SQLSTATE cross reference

This table is a cross-reference of all the SQLSTATES listed in the 'Diagnostics' section of each function description in "Chapter 5. Functions" on page 73.

**Note:** DB2 ODBC can also return SQLSTATES generated by the server that are not listed in this table. If the returned SQLSTATE is not listed here, see the documentation for the server for additional SQLSTATE information.

Table 150 (Page 1 of 10). SQLSTATE cross reference

SQLSTATE	Description	Functions
01000	Warning.	<ul style="list-style-type: none"> <li>SQLSetConnectOption()</li> </ul>
01002	Disconnect error.	<ul style="list-style-type: none"> <li>SQLDisconnect()</li> </ul>
01004	Data truncated.	<ul style="list-style-type: none"> <li>SQLColAttributes()</li> <li>SQLDataSources()</li> <li>SQLDescribeCol()</li> <li>SQLDriverConnect()</li> <li>SQLExecDirect()</li> <li>SQLExecute()</li> <li>SQLExtendedFetch()</li> <li>SQLFetch()</li> <li>SQLGetCursorName()</li> <li>SQLGetData()</li> <li>SQLGetInfo()</li> <li>SQLGetSubString()</li> <li>SQLNativeSql()</li> <li>SQLPutData()</li> <li>SQLSetColAttributes()</li> </ul>
01504	The UPDATE or DELETE statement does not include a WHERE clause.	<ul style="list-style-type: none"> <li>SQLExecDirect()</li> <li>SQLExecute()</li> <li>SQLPrepare()</li> </ul>
01508	Statement disqualified for blocking.	<ul style="list-style-type: none"> <li>SQLExecDirect()</li> <li>SQLExecute()</li> <li>SQLPrepare()</li> </ul>
01S00	Invalid connection string attribute.	<ul style="list-style-type: none"> <li>SQLDriverConnect()</li> </ul>
01S01	Error in row.	<ul style="list-style-type: none"> <li>SQLExtendedFetch()</li> </ul>
01S02	Option value changed.	<ul style="list-style-type: none"> <li>SQLDriverConnect()</li> <li>SQLSetConnectOption()</li> <li>SQLSetEnvAttr()</li> <li>SQLSetStmtOption()</li> </ul>
07001	Wrong number of parameters.	<ul style="list-style-type: none"> <li>SQLExecDirect()</li> <li>SQLExecute()</li> <li>SQLParamData()</li> </ul>
07002	Too many columns.	<ul style="list-style-type: none"> <li>SQLExtendedFetch()</li> <li>SQLFetch()</li> </ul>
07005	The statement did not return a result set.	<ul style="list-style-type: none"> <li>SQLColAttributes()</li> <li>SQLDescribeCol()</li> </ul>

Table 150 (Page 2 of 10). SQLSTATE cross reference

SQLSTATE	Description	Functions
07006	Invalid conversion.	<ul style="list-style-type: none"> <li>• SQLBindParameter()</li> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLExtendedFetch()</li> <li>• SQLFetch()</li> <li>• SQLGetData()</li> <li>• SQLSetParam()</li> <li>• SQLGetLength()</li> <li>• SQLGetPosition()</li> <li>• SQLGetSubString()</li> </ul>
08001	Unable to connect to data source.	<ul style="list-style-type: none"> <li>• SQLConnect()</li> <li>• SQLDriverConnect()</li> </ul>
08002	Connection in use.	<ul style="list-style-type: none"> <li>• SQLConnect()</li> <li>• SQLDriverConnect()</li> </ul>
08003	Connection is closed.	<ul style="list-style-type: none"> <li>• SQLAllocStmnt()</li> <li>• SQLDisconnect()</li> <li>• SQLGetConnectOption()</li> <li>• SQLGetInfo()</li> <li>• SQLNativeSql()</li> <li>• SQLSetConnection()</li> <li>• SQLSetConnectOption()</li> <li>• SQLTransact()</li> </ul>
08004	The application server rejected establishment of the connection.	<ul style="list-style-type: none"> <li>• SQLConnect()</li> <li>• SQLDriverConnect()</li> </ul>
08007	Connection failure during transaction.	<ul style="list-style-type: none"> <li>• SQLTransact()</li> </ul>
0F001	The LOB token variable does not currently represent any value.	<ul style="list-style-type: none"> <li>• SQLGetLength()</li> <li>• SQLGetPosition()</li> </ul>
21S01	Insert value list does not match column list.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLPrepare()</li> </ul>
21S02	Degrees of derived table does not match column list.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLPrepare()</li> </ul>
22001	String data right truncation.	<ul style="list-style-type: none"> <li>• SQLPutData()</li> </ul>
22002	Invalid output or indicator buffer specified.	<ul style="list-style-type: none"> <li>• SQLExtendedFetch()</li> <li>• SQLFetch()</li> <li>• SQLGetData()</li> </ul>
22003	Numeric value out of range.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLExtendedFetch()</li> <li>• SQLFetch()</li> <li>• SQLGetData()</li> <li>• SQLPutData()</li> </ul>
22005	Error in assignment.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLExtendedFetch()</li> <li>• SQLFetch()</li> <li>• SQLGetData()</li> <li>• SQLPutData()</li> </ul>



Table 150 (Page 3 of 10). SQLSTATE cross reference

SQLSTATE	Description	Functions
22008	Datetime field overflow.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLExtendedFetch()</li> <li>• SQLFetch()</li> <li>• SQLGetData()</li> </ul>
22011	A substring error occurred.	<ul style="list-style-type: none"> <li>• SQLGetSubString()</li> </ul>
22012	Division by zero is invalid.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLExtendedFetch()</li> <li>• SQLFetch()</li> </ul>
23000	Integrity constraint violation.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> </ul>
24000	Invalid cursor state.	<ul style="list-style-type: none"> <li>• SQLColumns()</li> <li>• SQLColumnPrivileges()</li> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLExtendedFetch()</li> <li>• SQLFetch()</li> <li>• SQLForeignKeys()</li> <li>• SQLGetData()</li> <li>• SQLGetStmtOption()</li> <li>• SQLGetTypeInfo()</li> <li>• SQLPrepare()</li> <li>• SQLPrimaryKeys()</li> <li>• SQLProcedures()</li> <li>• SQLProcedureColumns()</li> <li>• SQLSetColAttributes()</li> <li>• SQLSetStmtOption()</li> <li>• SQLSpecialColumns()</li> <li>• SQLStatistics()</li> <li>• SQLTables()</li> <li>• SQLTablePrivileges()</li> </ul>
24504	The cursor identified in the UPDATE, DELETE, SET, or GET statement is not positioned on a row.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> </ul>
25000	Invalid transaction state.	<ul style="list-style-type: none"> <li>• SQLDisconnect()</li> </ul>
25501	Invalid transaction state.	<ul style="list-style-type: none"> <li>• SQLDisconnect()</li> </ul>
28000	Invalid authorization specification.	<ul style="list-style-type: none"> <li>• SQLConnect()</li> <li>• SQLDriverConnect()</li> </ul>
34000	Invalid cursor name.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLPrepare()</li> <li>• SQLSetCursorName()</li> </ul>
37XXX	Invalid SQL syntax.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLPrepare()</li> </ul>
37000	Invalid SQL syntax.	<ul style="list-style-type: none"> <li>• SQLNativeSql()</li> </ul>
38552	Error in function listed in SYSFUN schema	All scalar functions, and SQLFetch() where the SQL statement references an ODBC scalar function.

## SQLSTATE cross reference

Table 150 (Page 4 of 10). SQLSTATE cross reference

SQLSTATE	Description	Functions
40001	Transaction rollback.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLPrepare()</li> </ul>
40003	Communication link failure.	<ul style="list-style-type: none"> <li>• SQLAllocStmt()</li> <li>• SQLBindCol()</li> <li>• SQLBindParameter()</li> <li>• SQLCancel()</li> <li>• SQLColumns()</li> <li>• SQLColumnPrivileges()</li> <li>• SQLColAttributes()</li> <li>• SQLDescribeCol()</li> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLExtendedFetch()</li> <li>• SQLFetch()</li> <li>• SQLForeignKeys()</li> <li>• SQLFreeStmt()</li> <li>• SQLGetConnectOption()</li> <li>• SQLGetCursorName()</li> <li>• SQLGetData()</li> <li>• SQLGetFunctions()</li> <li>• SQLGetInfo()</li> <li>• SQLGetLength()</li> <li>• SQLGetPosition()</li> <li>• SQLGetStmtOption()</li> <li>• SQLGetSubString()</li> <li>• SQLGetTypeInfo()</li> <li>• SQLMoreResults()</li> <li>• SQLNumParams()</li> <li>• SQLNumResultCols()</li> <li>• SQLParamData()</li> <li>• SQLParamOptions()</li> <li>• SQLPrepare()</li> <li>• SQLPrimaryKeys()</li> <li>• SQLProcedures()</li> <li>• SQLProcedureColumns()</li> <li>• SQLPutData()</li> <li>• SQLRowCount()</li> <li>• SQLSetColAttributes()</li> <li>• SQLSetConnectOption()</li> <li>• SQLSetCursorName()</li> <li>• SQLSetParam()</li> <li>• SQLSetStmtOption()</li> <li>• SQLSpecialColumns()</li> <li>• SQLStatistics()</li> <li>• SQLTables()</li> <li>• SQLTablePrivileges()</li> </ul>
42XXX	Syntax error or access rule violation	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLPrepare()</li> </ul>
425XX	Syntax error or access rule violation	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLPrepare()</li> </ul>
42601	PARMLIST syntax error.	<ul style="list-style-type: none"> <li>• SQLProcedureColumns()</li> </ul>
42818	The operands of an operator or function are not compatible.	<ul style="list-style-type: none"> <li>• SQLGetPosition()</li> </ul>

Table 150 (Page 5 of 10). SQLSTATE cross reference

SQLSTATE	Description	Functions
42895	The value of a host variable in the EXECUTE or OPEN statement cannot be used because of its data type	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> </ul>
44000	Integrity constraint violation.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> </ul>
56084	LOB data is not supported in DRDA.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLExtendedFetch()</li> <li>• SQLFetch()</li> </ul>
58004	Unexpected system failure.	<ul style="list-style-type: none"> <li>• SQLAllocConnect()</li> <li>• SQLAllocStmt()</li> <li>• SQLBindCol()</li> <li>• SQLBindParameter()</li> <li>• SQLColAttributes()</li> <li>• SQLConnect()</li> <li>• SQLDriverConnect()</li> <li>• SQLDataSources()</li> <li>• SQLDescribeCol()</li> <li>• SQLDisconnect()</li> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLExtendedFetch()</li> <li>• SQLFetch()</li> <li>• SQLFreeConnect()</li> <li>• SQLFreeEnv()</li> <li>• SQLFreeStmt()</li> <li>• SQLGetCursorName()</li> <li>• SQLGetData()</li> <li>• SQLGetFunctions()</li> <li>• SQLGetInfo()</li> <li>• SQLGetLength()</li> <li>• SQLGetPosition()</li> <li>• SQLGetSubString()</li> <li>• SQLMoreResults()</li> <li>• SQLNumResultCols()</li> <li>• SQLPrepare()</li> <li>• SQLRowCount()</li> <li>• SQLSetCursorName()</li> <li>• SQLSetParam()</li> <li>• SQLTransact()</li> </ul>
S0001	Database object already exists.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLPrepare()</li> </ul>
S0002	Database object does not exist.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLPrepare()</li> </ul>
S0011	Index already exists.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLPrepare()</li> </ul>
S0012	Index not found.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLPrepare()</li> </ul>

Table 150 (Page 6 of 10). SQLSTATE cross reference

SQLSTATE	Description	Functions
S0021	Column already exists.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLPrepare()</li> </ul>
S0022	Column not found.	<ul style="list-style-type: none"> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLPrepare()</li> </ul>
S1000	General error.	<ul style="list-style-type: none"> <li>• SQLDataSources()</li> <li>• SQLDescribeParam()</li> <li>• SQLDriverConnect()</li> <li>• SQLSetColAttributes()</li> <li>• SQLSetConnection()</li> <li>• SQLSetStmtOption()</li> </ul>
S1001	Memory allocation failure.	All functions.
S1002	Invalid column number.	<ul style="list-style-type: none"> <li>• SQLBindCol()</li> <li>• SQLColAttributes()</li> <li>• SQLDescribeCol()</li> <li>• SQLFetch()</li> <li>• SQLGetData()</li> <li>• SQLSetColAttributes()</li> </ul>
S1003	Program type out of range.	<ul style="list-style-type: none"> <li>• SQLBindCol()</li> <li>• SQLBindParameter()</li> <li>• SQLGetData()</li> <li>• SQLGetLength()</li> <li>• SQLGetSubString()</li> <li>• SQLSetParam()</li> </ul>
S1004	SQL data type out of range.	<ul style="list-style-type: none"> <li>• SQLBindParameter()</li> <li>• SQLGetTypeInfo()</li> <li>• SQLSetColAttributes()</li> <li>• SQLSetParam()</li> </ul>
S1009	Invalid argument value.	<ul style="list-style-type: none"> <li>• SQLAllocConnect()</li> <li>• SQLAllocStmt()</li> <li>• SQLBindParameter()</li> <li>• SQLColumnPrivileges()</li> <li>• SQLConnect()</li> <li>• SQLDriverConnect()</li> <li>• SQLExecDirect()</li> <li>• SQLForeignKeys()</li> <li>• SQLGetConnectOption()</li> <li>• SQLGetData()</li> <li>• SQLGetFunctions()</li> <li>• SQLGetInfo()</li> <li>• SQLGetLength()</li> <li>• SQLGetPosition()</li> <li>• SQLGetStmtOption()</li> <li>• SQLGetSubString()</li> <li>• SQLNativeSql()</li> <li>• SQLNumParams()</li> <li>• SQLNumResultCols()</li> <li>• SQLPrepare()</li> <li>• SQLPutData()</li> <li>• SQLSetConnectOption()</li> <li>• SQLSetCursorName()</li> <li>• SQLSetEnvAttr()</li> <li>• SQLSetParam()</li> <li>• SQLSetStmtOption()</li> </ul>

Table 150 (Page 7 of 10). SQLSTATE cross reference

SQLSTATE	Description	Functions
S1010	Function sequence error.	<ul style="list-style-type: none"> <li>• SQLBindCol()</li> <li>• SQLBindParameter()</li> <li>• SQLColumns()</li> <li>• SQLColAttributes()</li> <li>• SQLDescribeCol()</li> <li>• SQLDescribeParam()</li> <li>• SQLDisconnect()</li> <li>• SQLExecute()</li> <li>• SQLExtendedFetch()</li> <li>• SQLFetch()</li> <li>• SQLForeignKeys()</li> <li>• SQLFreeConnect()</li> <li>• SQLFreeEnv()</li> <li>• SQLFreeStmt()</li> <li>• SQLGetCursorName()</li> <li>• SQLGetData()</li> <li>• SQLGetFunctions()</li> <li>• SQLGetStmtOption()</li> <li>• SQLGetTypeInfo()</li> <li>• SQLMoreResults()</li> <li>• SQLNumParams()</li> <li>• SQLNumResultCols()</li> <li>• SQLParamData()</li> <li>• SQLParamOptions()</li> <li>• SQLPrepare()</li> <li>• SQLPrimaryKeys()</li> <li>• SQLProcedures()</li> <li>• SQLProcedureColumns()</li> <li>• SQLPutData()</li> <li>• SQLRowCount()</li> <li>• SQLSetColAttributes()</li> <li>• SQLSetConnectOption()</li> <li>• SQLSetCursorName()</li> <li>• SQLSetParam()</li> <li>• SQLSetStmtOption()</li> <li>• SQLSpecialColumns()</li> <li>• SQLStatistics()</li> <li>• SQLTables()</li> <li>• SQLTablePrivileges()</li> </ul>
S1011	Operation invalid at this time.	<ul style="list-style-type: none"> <li>• SQLSetConnectOption()</li> <li>• SQLSetEnvAttr()</li> <li>• SQLSetStmtOption()</li> </ul>
S1012	Invalid transaction code.	<ul style="list-style-type: none"> <li>• SQLTransact()</li> </ul>

Table 150 (Page 8 of 10). SQLSTATE cross reference

SQLSTATE	Description	Functions
S1013	Unexpected memory handling error.	<ul style="list-style-type: none"> <li>• SQLAllocConnect()</li> <li>• SQLAllocStmt()</li> <li>• SQLBindCol()</li> <li>• SQLBindParameter()</li> <li>• SQLCancel()</li> <li>• SQLColAttributes()</li> <li>• SQLConnect()</li> <li>• SQLDriverConnect()</li> <li>• SQLDataSources()</li> <li>• SQLDescribeCol()</li> <li>• SQLDisconnect()</li> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLExtendedFetch()</li> <li>• SQLFetch()</li> <li>• SQLFreeConnect()</li> <li>• SQLFreeEnv()</li> <li>• SQLGetCursorName()</li> <li>• SQLGetData()</li> <li>• SQLGetFunctions()</li> <li>• SQLGetLength()</li> <li>• SQLGetPosition()</li> <li>• SQLGetSubString()</li> <li>• SQLMoreResults()</li> <li>• SQLNumParams()</li> <li>• SQLNumResultCols()</li> <li>• SQLPrepare()</li> <li>• SQLRowCount()</li> <li>• SQLSetColAttributes()</li> <li>• SQLSetCursorName()</li> <li>• SQLSetParam()</li> <li>• SQLTransact()</li> </ul>
S1014	No more handles.	<ul style="list-style-type: none"> <li>• SQLAllocConnect()</li> <li>• SQLAllocStmt()</li> <li>• SQLColumns()</li> <li>• SQLColumnPrivileges()</li> <li>• SQLExecDirect()</li> <li>• SQLForeignKeys()</li> <li>• SQLPrepare()</li> <li>• SQLPrimaryKeys()</li> <li>• SQLProcedures()</li> <li>• SQLProcedureColumns()</li> <li>• SQLSpecialColumns()</li> <li>• SQLStatistics()</li> <li>• SQLTables()</li> <li>• SQLTablePrivileges()</li> </ul>

Table 150 (Page 9 of 10). SQLSTATE cross reference

SQLSTATE	Description	Functions
S1090	Invalid string or buffer length.	<ul style="list-style-type: none"> <li>• SQLBindCol()</li> <li>• SQLBindParameter()</li> <li>• SQLColumns()</li> <li>• SQLColumnPrivileges()</li> <li>• SQLColAttributes()</li> <li>• SQLConnect()</li> <li>• SQLDriverConnect()</li> <li>• SQLDataSources()</li> <li>• SQLDescribeCol()</li> <li>• SQLDriverConnect()</li> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLParamData()</li> <li>• SQLForeignKeys()</li> <li>• SQLGetCursorName()</li> <li>• SQLGetData()</li> <li>• SQLGetInfo()</li> <li>• SQLGetPosition()</li> <li>• SQLGetSubString()</li> <li>• SQLNativeSql()</li> <li>• SQLPrepare()</li> <li>• SQLPrimaryKeys()</li> <li>• SQLProcedures()</li> <li>• SQLProcedureColumns()</li> <li>• SQLPutData()</li> <li>• SQLSetColAttributes()</li> <li>• SQLSetCursorName()</li> <li>• SQLSpecialColumns()</li> <li>• SQLStatistics()</li> <li>• SQLTables()</li> <li>• SQLTablePrivileges()</li> </ul>
S1091	Descriptor type out of range.	<ul style="list-style-type: none"> <li>• SQLColAttributes()</li> </ul>
S1092	Option type out of range.	<ul style="list-style-type: none"> <li>• SQLFreeStmt()</li> <li>• SQLGetConnectOption()</li> <li>• SQLGetEnvAttr()</li> <li>• SQLExecDirect()</li> <li>• SQLExecute()</li> <li>• SQLExtendedFetch()</li> <li>• SQLFetch()</li> <li>• SQLGetStmtOption()</li> <li>• SQLParamData()</li> <li>• SQLSetConnectOption()</li> <li>• SQLSetEnvAttr()</li> <li>• SQLSetStmtOption()</li> </ul>
S1093	Invalid parameter number.	<ul style="list-style-type: none"> <li>• SQLBindParameter()</li> <li>• SQLDescribeParam()</li> <li>• SQLSetParam()</li> </ul>
S1094	Invalid scale value.	<ul style="list-style-type: none"> <li>• SQLBindParameter()</li> <li>• SQLSetColAttributes()</li> <li>• SQLSetParam()</li> </ul>
S1096	Information type out of range.	<ul style="list-style-type: none"> <li>• SQLGetInfo()</li> </ul>
S1097	Column type out of range.	<ul style="list-style-type: none"> <li>• SQLSpecialColumns()</li> </ul>
S1098	Scope type out of range.	<ul style="list-style-type: none"> <li>• SQLSpecialColumns()</li> </ul>
S1099	Nullable type out of range.	<ul style="list-style-type: none"> <li>• SQLSetColAttributes()</li> <li>• SQLSpecialColumns()</li> </ul>
S1100	Uniqueness option type out of range.	<ul style="list-style-type: none"> <li>• SQLStatistics()</li> </ul>
S1101	Accuracy option type out of range.	<ul style="list-style-type: none"> <li>• SQLStatistics()</li> </ul>

Table 150 (Page 10 of 10). SQLSTATE cross reference

SQLSTATE	Description	Functions
S1103	Direction option out of range.	<ul style="list-style-type: none"> <li>SQLDataSources()</li> <li>SQLGetInfo()</li> </ul>
S1104	Invalid precision value.	<ul style="list-style-type: none"> <li>SQLBindParameter()</li> <li>SQLSetColAttributes()</li> <li>SQLSetParam()</li> </ul>
S1105	Invalid parameter type.	<ul style="list-style-type: none"> <li>SQLBindParameter()</li> </ul>
S1106	Fetch type out of range.	<ul style="list-style-type: none"> <li>SQLExtendedFetch()</li> </ul>
S1107	Row value out of range.	<ul style="list-style-type: none"> <li>SQLParamOptions()</li> </ul>
S1110	Invalid driver completion.	<ul style="list-style-type: none"> <li>SQLDriverConnect()</li> </ul>
S1501	Invalid data source name.	<ul style="list-style-type: none"> <li>SQLConnect()</li> <li>SQLDriverConnect()</li> </ul>
S1503	Invalid file name length.	<ul style="list-style-type: none"> <li>SQLExecDirect()</li> <li>SQLExecute()</li> <li>SQLParamData()</li> </ul>
S1506	Error closing a file.	<ul style="list-style-type: none"> <li>SQLCancel()</li> <li>SQLFreeStmt()</li> <li>SQLParamData()</li> </ul>
S1509	Error deleting a file.	<ul style="list-style-type: none"> <li>SQLParamData()</li> </ul>
S1C00	Driver not capable.	<ul style="list-style-type: none"> <li>SQLBindCol()</li> <li>SQLBindParameter()</li> <li>SQLColumns()</li> <li>SQLColumnPrivileges()</li> <li>SQLColAttributes()</li> <li>SQLDescribeCol()</li> <li>SQLDescribeParam()</li> <li>SQLExtendedFetch()</li> <li>SQLFetch()</li> <li>SQLForeignKeys()</li> <li>SQLGetConnectOption()</li> <li>SQLGetData()</li> <li>SQLGetInfo()</li> <li>SQLGetLength()</li> <li>SQLGetPosition()</li> <li>SQLGetStmtOption()</li> <li>SQLGetSubString()</li> <li>SQLPrimaryKeys()</li> <li>SQLProcedures()</li> <li>SQLProcedureColumns()</li> <li>SQLSetConnectOption()</li> <li>SQLSetEnvAttr()</li> <li>SQLSetParam()</li> <li>SQLSetStmtOption()</li> <li>SQLSpecialColumns()</li> <li>SQLStatistics()</li> <li>SQLTables()</li> <li>SQLTablePrivileges()</li> </ul>



---

## Appendix D. Data conversion

This section contains tables used for data conversion between C and SQL data types. This includes:

- Precision, scale, length, and display size of each data type
- Conversion from SQL to C data types
- Conversion from C to SQL data types

For a list of SQL and C data types, their symbolic types, and the default conversions, see Table 4 on page 39 and Table 5 on page 40. Supported conversions are shown in Table 8 on page 42.

### Data type attributes

Information is shown for the following data type attributes:

- “Precision”
- “Scale” on page 446
- “Length” on page 447
- “Display size” on page 448

### Precision

The precision of a numeric column or parameter refers to the maximum number of digits used by the data type of the column or parameter. The precision of a non-numeric column or parameter generally refers to the maximum length or the defined length of the column or parameter. The following table defines the precision for each SQL data type.

Table 151. Precision

fSqlType	Precision
SQL_CHAR SQL_VARCHAR SQL_CLOB	The defined length of the column or parameter. For example, the precision of a column defined as CHAR(10) is 10.
SQL_LONGVARCHAR	The maximum length of the column or parameter. <sup>a</sup>
SQL_DECIMAL SQL_NUMERIC	The defined maximum number of digits. For example, the precision of a column defined as NUMERIC(10,3) is 10.
SQL_SMALLINT <sup>b</sup>	5
SQL_INTEGER <sup>b</sup>	10
SQL_FLOAT <sup>b</sup>	15
SQL_REAL <sup>b</sup>	7
SQL_ROWID	40
SQL_DOUBLE <sup>b</sup>	15
SQL_BINARY SQL_VARBINARY SQL_BLOB	The defined length of the column or parameter. For example, the precision of a column defined as CHAR(10) FOR BIT DATA, is 10.
SQL_LONGVARBINARY	The maximum length of the column or parameter.
SQL_DATE <sup>b</sup>	10 (the number of characters in the yyyy-mm-dd format).
SQL_TIME <sup>b</sup>	8 (the number of characters in the hh:mm:ss format).
SQL_TIMESTAMP	The number of characters in the "yyyy-mm-dd hh:mm:ss[.fff[fff]]" or "yyyy-mm-dd.hh.mm.ss.fff[fff]" format used by the TIMESTAMP data type. For example, if a timestamp does not use seconds or fractional seconds, the precision is 16 (the number of characters in the "yyyy-mm-dd hh:mm" format). If a timestamp uses thousandths of a second, the precision is 26 (the number of characters in the "yyyy-mm-dd hh:mm:ss.fffff" format). The maximum for fractional seconds is 6 digits.
SQL_GRAPHIC SQL_VARGRAPHIC SQL_DBCLOB	The defined length of the column or parameter. For example, the precision of a column defined as GRAPHIC(10) is 10.
SQL_LONGVARGRAPHIC	The maximum length of the column or parameter.

**Note:**

- <sup>a</sup> When defining the precision of a parameter of this data type with `SQLBindParameter()` or `SQLSetParam()`, `cbParamDef` should be set to the total length of the data, not the precision as defined in this table.
- <sup>b</sup> The `cbParamDef` argument of `SQLBindParameter()` or `SQLSetParam()` is ignored for this data type.

### Scale

The scale of a numeric column or parameter refers to the maximum number of digits to the right of the decimal point. Note that, for approximate floating point number columns or parameters, the scale is undefined, since the number of digits to the right of the decimal place is not fixed. The following table defines the scale for each SQL data type.

Table 152. Scale

<b>fSqlType</b>	<b>Scale</b>
SQL_CHAR SQL_VARCHAR SQL_LONGVARCHAR SQL_CLOB	Not applicable.
SQL_DECIMAL SQL_NUMERIC	The defined number of digits to the right of the decimal place. For example, the scale of a column defined as NUMERIC(10, 3) is 3.
SQL_SMALLINT SQL_INTEGER	0
SQL_REAL SQL_FLOAT SQL_DOUBLE	Not applicable.
SQL_ROWID	Not applicable.
SQL_BINARY SQL_VARBINARY SQL_LONGVARBINARY SQL_BLOB	Not applicable.
SQL_DATE SQL_TIME	Not applicable.
SQL_TIMESTAMP	The number of digits to the right of the decimal point in the "yyyy-mm-dd hh:mm:ss[fff[fff]]" format. For example, if the TIMESTAMP data type uses the "yyyy-mm-dd hh:mm:ss.fff" format, the scale is 3. The maximum for fractional seconds is 6 digits.
SQL_GRAPHIC SQL_VARGRAPHIC SQL_LONGVARGRAPHIC SQL_DBCLOB	Not applicable.

## Length

The length of a column is the maximum number of *bytes* returned to the application when data is transferred to its default C data type. For character data, the length does not include the null termination byte. Note that the length of a column can be different than the number of bytes required to store the data on the data source. For a list of default C data types, see the "Default C Data Types" section.

The following table defines the length for each SQL data type.

Table 153. Length

fSqlType	Length
SQL_CHAR SQL_VARCHAR SQL_CLOB	The defined length of the column. For example, the length of a column defined as CHAR(10) is 10.
SQL_LONGVARCHAR	The maximum length of the column.
SQL_DECIMAL SQL_NUMERIC	The maximum number of digits plus two. Since these data types are returned as character strings, characters are needed for the digits, a sign, and a decimal point. For example, the length of a column defined as NUMERIC(10,3) is 12.
SQL_SMALLINT	2 (two bytes).
SQL_INTEGER	4 (four bytes).
SQL_REAL	4 (four bytes).
SQL_ROWID	40
SQL_FLOAT	8 (eight bytes).
SQL_DOUBLE	8 (eight bytes).
SQL_BINARY SQL_VARBINARY SQL_BLOB	The defined length of the column. For example, the length of a column defined as CHAR(10) FOR BIT DATA is 10.
SQL_LONGVARBINARY	The maximum length of the column.
SQL_DATE SQL_TIME	6 (the size of the DATE_STRUCT or TIME_STRUCT structure).
SQL_TIMESTAMP	16 (the size of the TIMESTAMP_STRUCT structure).
SQL_GRAPHIC SQL_VARGRAPHIC SQL_DBCLOB	The defined length of the column times 2. For example, the length of a column defined as GRAPHIC(10) is 20.
SQL_LONGVARGRAPHIC	The maximum length of the column times 2.

### Display size

The display size of a column is the maximum number of *bytes* needed to display data in character form. The following table defines the display size for each SQL data type.

Table 154. Display size

<b>fSqlType</b>	<b>Display Size</b>
SQL_CHAR SQL_VARCHAR SQL_CLOB	The defined length of the column. For example, the display size of a column defined as CHAR(10) is 10.
SQL_LONGVARCHAR	The maximum length of the column.
SQL_DECIMAL SQL_NUMERIC	The precision of the column plus two (a sign, precision digits, and a decimal point). For example, the display size of a column defined as NUMERIC(10,3) is 12.
SQL_SMALLINT	6 (a sign and 5 digits).
SQL_INTEGER	11 (a sign and 10 digits).
SQL_REAL	13 (a sign, 7 digits, a decimal point, the letter E, a sign, and 2 digits).
SQL_ROWID	40
SQL_FLOAT SQL_DOUBLE	22 (a sign, 15 digits, a decimal point, the letter E, a sign, and 3 digits).
SQL_BINARY SQL_VARBINARY SQL_BLOB	The defined length of the column times 2 (each binary byte is represented by a 2 digit hexadecimal number). For example, the display size of a column defined as CHAR(10) FOR BIT DATA is 20.
SQL_LONGVARBINARY	The maximum length of the column times 2.
SQL_DATE	10 (a date in the format yyyy-mm-dd).
SQL_TIME	8 (a time in the format hh:mm:ss).
SQL_TIMESTAMP	19 (if the scale of the timestamp is 0) or 20 plus the scale of the timestamp (if the scale is greater than 0). This is the number of characters in the "yyyy-mm-dd hh:mm:ss[fff[fff]]" or "yyyy-mm-dd.hh.mm.ss.fff[fff]" format. For example, the display size of a column storing thousandths of a second is 23 (the number of characters in "yyyy-mm-dd hh:mm:ss.ffff"). The maximum for fractional seconds is 6 digits.
SQL_GRAPHIC SQL_VARGRAPHIC SQL_DBCLOB	The defined length of the column or parameter. For example, the display size of a column defined as GRAPHIC(10) is 20.
SQL_LONGVARGRAPHIC	The maximum length of the column or parameter.

---

## Converting data from SQL to C data types

For a given SQL data type:

- The first column of the table lists the legal input values of the *fCType* argument in `SQLBindCol()` and `SQLGetData()`.
- The second column lists the outcomes of a test, often using the *cbValueMax* argument specified in `SQLBindCol()` or `SQLGetData()`, which the driver performs to determine if it can convert the data.
- The third and fourth columns list the values (for each outcome) of the *rgbValue* and *pcbValue* arguments specified in the `SQLBindCol()` or `SQLGetData()` after the driver has attempted to convert the data.
- The last column lists the `SQLSTATE` returned for each outcome by `SQLFetch()`, `SQLExtendedFetch()`, or `SQLGetData()`.

The tables list the conversions defined by ODBC to be valid for a given SQL data type.

If the *fCType* argument in `SQLBindCol()` or `SQLGetData()` contains a value not shown in the table for a given SQL data type, `SQLFetch()`, or `SQLGetData()` returns the `SQLSTATE 07006` (restricted data type attribute violation).

If the *fCType* argument contains a value shown in the table but which specifies a conversion not supported by the driver, `SQLFetch()`, or `SQLGetData()` returns `SQLSTATE S1C00` (driver not capable).

Though it is not shown in the tables, the *pcbValue* argument contains `SQL_NULL_DATA` when the SQL data value is `NULL`. For an explanation of the use of *pcbValue* when multiple calls are made to retrieve data, see `SQLGetData()`.

When SQL data is converted to character C data, the character count returned in *pcbValue* does not include the null termination byte. If *rgbValue* is a null pointer, `SQLBindCol()` or `SQLGetData()` returns `SQLSTATE S1009` (Invalid argument value).

In the following tables:

### Length of data

The total length of the data after it has been converted to the specified C data type (excluding the null termination byte if the data was converted to a string). This is true even if data is truncated before it is returned to the application.

### Significant digits

The minus sign (if needed) and the digits to the left of the decimal point.

### Display size

The total number of bytes needed to display data in the character format.

## Converting character SQL data to C data

The character SQL data types are:

SQL\_CHAR  
 SQL\_VARCHAR  
 SQL\_LONGVARCHAR  
 SQL\_CLOB

Table 155. Converting character SQL data to C data

fCType	Test	rgbValue	pcbValue	SQLSTATE
SQL_C_CHAR	Length of data < cbValueMax	Data	Length of data	00000
	Length of data >= cbValueMax	Truncated data	Length of data	01004
SQL_C_BINARY	Length of data <= cbValueMax	Data	Length of data	00000
	Length of data > cbValueMax	Truncated data	Length of data	01004
SQL_C_SHORT	Data converted without truncation <sup>a</sup>	Data	Size of the C data type	00000
SQL_C_LONG				
SQL_C_FLOAT				
SQL_C_DOUBLE	Data converted with truncation, but without loss of significant digits <sup>a</sup>	Data	Size of the C data type	01004
SQL_C_TINYINT				
SQL_C_BIT	Conversion of data would result in loss of significant digits <sup>a</sup>	Untouched	Size of the C data type	22003
	Data is not a number <sup>a</sup>	Untouched	Size of the C data type	22005
SQL_C_DATE	Data value is a valid date <sup>a</sup>	Data	6 <sup>b</sup>	00000
	Data value is not a valid date <sup>a</sup>	Untouched	6 <sup>b</sup>	22008
SQL_C_TIME	Data value is a valid time <sup>a</sup>	Data	6 <sup>b</sup>	00000
	Data value is not a valid time <sup>a</sup>	Untouched	6 <sup>b</sup>	22008
SQL_C_TIMESTAMP	Data value is a valid timestamp <sup>a</sup>	Data	16 <sup>b</sup>	00000
	Data value is not a valid timestamp <sup>a</sup>	Untouched	16 <sup>b</sup>	22008

**Note:**

- <sup>a</sup> The value of *cbValueMax* is ignored for this conversion. The driver assumes that the size of *rgbValue* is the size of the C data type.
- <sup>b</sup> This is the size of the corresponding C data type.

SQLSTATE **00000** is not returned by `SQLError()`, rather it is indicated when the function returns `SQL_SUCCESS`.



## Converting graphic SQL data to C data

The graphic SQL data types are:

SQL\_GRAPHIC  
 SQL\_VARGRAPHIC  
 SQL\_LONGVARGRAPHIC  
 SQL\_DBCLOB

Table 156. Converting graphic SQL data to C data

fCType	Test	rgbValue	pcbValue	SQLSTATE
SQL_C_CHAR	Number of double byte characters * 2 <= cbValueMax	Data	Length of data(octets)	00000
	Number of double byte characters * 2 <= cbValueMax	Truncated data, to the nearest even byte that is less than <i>cbValueMax</i> .	Length of data(octets)	01004
SQL_C_DBCHAR	Number of double byte characters * 2 < cbValueMax	Data	Length of data(octets)	00000
	Number of double byte characters * 2 >= cbValueMax	Truncated <i>cbValueMax</i> . data, to the nearest even byte that is less than <i>cbValueMax</i> .	Length of data(octets)	01004

**Note:**

SQLSTATE **00000** is not returned by `SQLError()`, rather it is indicated when the function returns `SQL_SUCCESS`.

## Converting numeric SQL data to C data

The numeric SQL data types are:

SQL\_DECIMAL  
 SQL\_NUMERIC  
 SQL\_SMALLINT  
 SQL\_INTEGER  
 SQL\_REAL  
 SQL\_FLOAT  
 SQL\_DOUBLE

Table 157. Converting numeric SQL data to C data

fCType	Test	rgbValue	pcbValue	SQLSTATE
SQL_C_CHAR	Display size < cbValueMax	Data	Length of data	00000
	Number of significant digits < cbValueMax	Truncated data	Length of data	01004
	Number of significant digits >= cbValueMax	Untouched	Length of data	22003
SQL_C_SHORT SQL_C_LONG	Data converted without truncation <sup>a</sup>	Data	Size of the C data type	00000
SQL_C_FLOAT SQL_C_DOUBLE SQL_C_TINYINT	Data converted with truncation, but without loss of significant digits <sup>a</sup>	Truncated data	Size of the C data type	01004
SQL_C_BIT	Conversion of data would result in loss of significant digits <sup>a</sup>	Untouched	Size of the C data type	22003

**Note:**

<sup>a</sup> The value of *cbValueMax* is ignored for this conversion. The driver assumes that the size of *rgbValue* is the size of the C data type.

SQLSTATE **00000** is not returned by `SQLERROR()`, rather it is indicated when the function returns `SQL_SUCCESS`.

## Converting binary SQL data to C data

The binary SQL data types are:

SQL\_BINARY  
 SQL\_VARBINARY  
 SQL\_LONGVARBINARY  
 SQL\_BLOB

Table 158. Converting binary SQL data to C data

<b>fCType</b>	<b>Test</b>	<b>rgbValue</b>	<b>pcbValue</b>	<b>SQLSTATE</b>
SQL_C_CHAR	(Length of data) < cbValueMax	Data	Length of data	N/A
	(Length of data) >= cbValueMax	Truncated data	Length of data	01004
SQL_C_BINARY	Length of data <= cbValueMax	Data	Length of data	N/A
	Length of data > cbValueMax	Truncated data	Length of data	01004

### Converting date SQL data to C data

The date SQL data type is:

SQL\_DATE

Table 159. Converting date SQL data to C data

fCType	Test	rgbValue	pcbValue	SQLSTATE
SQL_C_CHAR	cbValueMax >= 11	Data	10	00000
	cbValueMax < 11	Untouched	10	22003
SQL_C_DATE	None <sup>a</sup>	Data	6 <sup>b</sup>	00000
SQL_C_TIMESTAMP	None <sup>a</sup>	Data <sup>c</sup>	16 <sup>b</sup>	00000
SQL_C_BINARY	Length of data <= cbValueMax	Data	Length of data	00000
	Length of data > cbValueMax	Untouched	Untouched	22003

**Note:**

- <sup>a</sup> The value of *cbValueMax* is ignored for this conversion. The driver assumes that the size of *rgbValue* is the size of the C data type.
- <sup>b</sup> This is the size of the corresponding C data type.
- <sup>c</sup> The time fields of the `TIMESTAMP_STRUCT` structure are set to zero.

SQLSTATE **00000** is not returned by `SQLERROR()`, rather it is indicated when the function returns `SQL_SUCCESS`.

When the date SQL data type is converted to the character C data type, the resulting string is in the "yyyy-mm-dd" format.

## Converting time SQL data to C data

The time SQL data type is:

SQL\_TIME

Table 160. Converting time SQL data to C data

fCType	Test	rgbValue	pcbValue	SQLSTATE
SQL_C_CHAR	cbValueMax >= 9	Data	8	00000
	cbValueMax < 9	Untouched	8	22003
SQL_C_TIME	None <sup>a</sup>	Data	6 <sup>b</sup>	00000
SQL_C_TIMESTAMP	None <sup>a</sup>	Data <sup>c</sup>	16 <sup>b</sup>	00000

**Note:**

- <sup>a</sup> The value of *cbValueMax* is ignored for this conversion. The driver assumes that the size of *rgbValue* is the size of the C data type.
- <sup>b</sup> This is the size of the corresponding C data type.
- <sup>c</sup> The date fields of the `TIMESTAMP_STRUCT` structure are set to the current system date of the machine that the application is running, and the time fraction is set to zero.

SQLSTATE **00000** is not returned by `SQLError()`, rather it is indicated when the function returns `SQL_SUCCESS`.

When the time SQL data type is converted to the character C data type, the resulting string is in the "hh:mm:ss" format.

## Converting timestamp SQL data to C data

The timestamp SQL data type is:

SQL\_TIMESTAMP

Table 161. Converting timestamp SQL data to C data

fCType	Test	rgbValue	pcbValue	SQLSTATE
SQL_C_CHAR	Display size < cbValueMax	Data	Length of data	00000
	19 <= cbValueMax <= Display size	Truncated data <sup>b</sup>	Length of data	01004
	cbValueMax < 19	Untouched	Length of data	22003
SQL_C_DATE	None <sup>a</sup>	Truncated data <sup>c</sup>	6 <sup>e</sup>	01004
SQL_C_TIME	None <sup>a</sup>	Truncated data <sup>d</sup>	6 <sup>e</sup>	01004
SQL_C_TIMESTAMP	None <sup>a</sup>	Data	16 <sup>e</sup>	00000
	Fractional seconds portion of timestamp is truncated. <sup>a</sup>	Data <sup>b</sup>	16	01004

### Note:

- a The value of *cbValueMax* is ignored for this conversion. The driver assumes that the size of *rgbValue* is the size of the C data type.
- b The fractional seconds of the timestamp are truncated.
- c The time portion of the timestamp is deleted.
- d The date portion of the timestamp is deleted.
- e This is the size of the corresponding C data type.

SQLSTATE **00000** is not returned by `SQLERROR()`, rather it is indicated when the function returns `SQL_SUCCESS`.

When the timestamp SQL data type is converted to the character C data type, the resulting string is in the "yyyy-mm-dd hh:mm:ss[.fff[fff]]" format (regardless of the precision of the timestamp SQL data type).

### Converting row ID SQL data to C data

The row ID SQL data type is:

SQL\_ROWID

Table 162. Converting row ID SQL data to C data

<b>fCType</b>	<b>Test</b>	<b>rgbValue</b>	<b>pcbValue</b>	<b>SQLSTATE</b>
SQL_C_CHAR	Length of data < = cbValueMax	Data	Length of data	00000
	Length of data > cbValueMax	Truncated data	Length of data	01004

## SQL to C data conversion examples

Table 163. SQL to C data conversion examples

SQL Data Type	SQL Data Value	C Data Type	cbValueMax	rgbValue	SQL STATE
SQL_CHAR	abcdef	SQL_C_CHAR	7	abcdef\0 <sup>a</sup>	00000
SQL_CHAR	abcdef	SQL_C_CHAR	6	abcde\0 <sup>a</sup>	01004
SQL_DECIMAL	1234.56	SQL_C_CHAR	8	1234.56\0 <sup>a</sup>	00000
SQL_DECIMAL	1234.56	SQL_C_CHAR	5	1234\0 <sup>a</sup>	01004
SQL_DECIMAL	1234.56	SQL_C_CHAR	4	---	22003
SQL_DECIMAL	1234.56	SQL_C_FLOAT	Ignored	1234.56	00000
SQL_DECIMAL	1234.56	SQL_C_SHORT	Ignored	1234	01004
SQL_DATE	1992-12-31	SQL_C_CHAR	11	1992-12-31\0 <sup>a</sup>	00000
SQL_DATE	1992-12-31	SQL_C_CHAR	10	---	22003
SQL_DATE	1992-12-31	SQL_C_TIMESTAMP	Ignored	1992,12,31, 0,0,0,0 <sup>b</sup>	00000
SQL_TIMESTAMP	1992-12-31 23:45:55.12	SQL_C_CHAR	23	1992-12-31 23:45:55.12\0 <sup>a</sup>	00000
SQL_TIMESTAMP	1992-12-31 23:45:55.12	SQL_C_CHAR	22	1992-12-31 23:45:55.1\0 <sup>a</sup>	01004
SQL_TIMESTAMP	1992-12-31 23:45:55.12	SQL_C_CHAR	18	---	22003

**Note:**

<sup>a</sup> "\0" represents a null termination character.

<sup>b</sup> The numbers in this list are the numbers stored in the fields of the `TIMESTAMP_STRUCT` structure.

SQLSTATE **00000** is not returned by `SQLERROR()`, rather it is indicated when the function returns `SQL_SUCCESS`.



---

## Converting data from C to SQL data types

For a given C data type:

- The first column of the table lists the legal input values of the *fSqlType* argument in `SQLBindParameter()` or `SQLSetParam()`.
- The second column lists the outcomes of a test, often using the length of the parameter data as specified in the *pcbValue* argument in `SQLBindParameter()` or `SQLSetParam()`, which the driver performs to determine if it can convert the data.
- The third column lists the SQLSTATE returned for each outcome by `SQLExecDirect()` or `SQLExecute()`.

**Note:** Data is sent to the data source only if the SQLSTATE is 00000 (success).

The tables list the conversions defined by ODBC to be valid for a given SQL data type.

If the *fSqlType* argument in `SQLBindParameter()` or `SQLSetParam()` contains a value not shown in the table for a given C data type, SQLSTATE 07006 is returned (Restricted data type attribute violation).

If the *fSqlType* argument contains a value shown in the table but which specifies a conversion not supported by the driver, `SQLBindParameter()` or `SQLSetParam()` returns SQLSTATE S1C00 (Driver not capable).

If the *rgbValue* and *pcbValue* arguments specified in `SQLBindParameter()` or `SQLSetParam()` are both null pointers, that function returns SQLSTATE S1009 (Invalid argument value).

### Length of data

The total length of the data after it has been converted to the specified SQL data type (excluding the null termination byte if the data was converted to a string). This is true even if data is truncated before it is sent to the data source.

### Column length

The maximum number of bytes returned to the application when data is transferred to its default C data type. For character data, the length does not include the null termination byte.

### Display size

The maximum number of bytes needed to display data in character form.

### Significant digits

The minus sign (if needed) and the digits to the left of the decimal point.

## Converting character C data to SQL data

The character C data type is:

SQL\_C\_CHAR

Table 164. Converting character C data to SQL data

fSQLType	Test	SQLSTATE
SQL_CHAR	Length of data <= Column length	00000
SQL_VARCHAR	Length of data > Column length	01004
SQL_LONGVARCHAR		
SQL_CLOB		
SQL_DECIMAL	Data converted without truncation	00000
SQL_NUMERIC	Data converted with truncation, but without loss of significant digits	01004
SQL_SMALLINT		
SQL_INTEGER	Conversion of data would result in loss of significant digits	22003
SQL_REAL	Data value is not a numeric value	22005
SQL_FLOAT		
SQL_DOUBLE		
SQL_BINARY	(Length of data) < Column length	N/A
SQL_VARBINARY	(Length of data) >= Column length	01004
SQL_LONGVARBINARY		
SQL_BLOB	Data value is not a hexadecimal value	22005
SQL_ROWID	Length of data <= Column length	00000
	Length of data > Column length	01004
SQL_DATE	Data value is a valid date	00000
	Data value is not a valid date	22008
SQL_TIME	Data value is a valid time	00000
	Data value is not a valid time	22008
	Data value is a valid timestamp; time portion is non-zero	01004
SQL_TIMESTAMP	Data value is a valid timestamp	00000
	Data value is a valid timestamp; fractional seconds portion is non-zero	01004
	Data value is not a valid timestamp	22008
	Data value is a valid timestamp; fractional seconds portion is non-zero	01004
SQL_GRAPHIC	Length of data / 2 <= Column length	00000
SQL_VARGRAPHIC	Length of data / 2 < Column length	01004
SQL_LONGVARGRAPHIC		
SQL_DBCLOB		

**Note:**

SQLSTATE 00000 is not returned by SQL\_Error(), rather it is indicated when the function returns SQL\_SUCCESS.

## Converting numeric C data to SQL data

The numeric C data types are:

```
SQL_C_SHORT
SQL_C_LONG
SQL_C_FLOAT
SQL_C_DOUBLE
SQL_C_TINYINT
SQL_C_BIT
```

Table 165. Converting numeric C data to SQL data

<b>fSQLType</b>	<b>Test</b>	<b>SQLSTATE</b>
SQL_DECIMAL	Data converted without truncation	00000
SQL_NUMERIC	Data converted with truncation, but without loss of significant digits	01004
SQL_SMALLINT	Conversion of data would result in loss of significant digits	22003
SQL_INTEGER		
SQL_REAL		
SQL_FLOAT		
SQL_DOUBLE		
SQL_CHAR	Data converted without truncation.	00000
SQL_VARCHAR	Conversion of data would result in loss of significant digits.	22003

**Note:**

SQLSTATE **00000** is not returned by `SQLError()`, rather it is indicated when the function returns `SQL_SUCCESS`.

### Converting binary C data to SQL data

The binary C data type is:

SQL\_C\_BINARY

Table 166. Converting binary C data to SQL data

<b>fSQLType</b>	<b>Test</b>	<b>SQLSTATE</b>
SQL_CHAR	Length of data <= Column length	N/A
SQL_VARCHAR	Length of data > Column length	01004
SQL_LONGVARCHAR		
SQL_CLOB		
SQL_BINARY	Length of data <= Column length	N/A
SQL_VARBINARY	Length of data > Column length	01004
SQL_LONGVARBINARY		
SQL_BLOB		

## Converting DBCHAR C data to SQL data

The double byte C data type is:

SQL\_C\_DBCHAR

Table 167. Converting DBCHAR C data to SQL data

<b>fSQLType</b>	<b>Test</b>	<b>SQLSTATE</b>
SQL_CHAR	Length of data <= Column length x 2	N/A
SQL_VARCHAR	Length of data > Column length x 2	01004
SQL_LONGVARCHAR		
SQL_CLOB		
SQL_BINARY	Length of data <= Column length x 2	N/A
SQL_VARBINARY	Length of data > Column length x 2	01004
SQL_LONGVARBINARY		
SQL_BLOB		

### Converting date C data to SQL data

The date C data type is:

SQL\_C\_DATE

Table 168. Converting date C data to SQL data

<b>fSQLType</b>	<b>Test</b>	<b>SQLSTATE</b>
SQL_CHAR	Column length >= 10	00000
SQL_VARCHAR	Column length < 10	22003
SQL_DATE	Data value is a valid date	00000
	Data value is not a valid date	22008
SQL_TIMESTAMP <sup>a</sup>	Data value is a valid date	00000
	Data value is not a valid date	22008

**Note:**

<sup>a</sup> The time component of TIMESTAMP is set to zero.

SQLSTATE **00000** is not returned by `SQLERROR()`, rather it is indicated when the function returns `SQL_SUCCESS`.

## Converting time C data to SQL data

The time C data type is:

SQL\_C\_TIME

Table 169. Converting time C data to SQL data

fSQLType	Test	SQLSTATE
SQL_CHAR	Column length >= 8	00000
SQL_VARCHAR	Column length < 8	22003
SQL_TIME	Data value is a valid time	00000
	Data value is not a valid time	22008
SQL_TIMESTAMP <sup>a</sup>	Data value is a valid time	00000
	Data value is not a valid time	22008

**Note:**

- <sup>a</sup> The date component of TIMESTAMP is set to the system date of the machine at which the application is running.

SQLSTATE **00000** is not returned by `SQLERROR()`, rather it is indicated when the function returns `SQL_SUCCESS`.

## Converting timestamp C data to SQL data

The timestamp C data type is:

SQL\_C\_TIMESTAMP

Table 170. Converting timestamp C data to SQL data

fSQLType	Test	SQLSTATE
SQL_CHAR	Column length >= Display size	00000
SQL_VARCHAR	19 <= Column length < Display size <sup>a</sup>	01004
	Column length < 19	22003
SQL_DATE	Data value is a valid date <sup>b</sup>	01004
	Data value is not a valid date	22008
SQL_TIME	Data value is a valid time <sup>c</sup>	01004
	Data value is not a valid time	22008
	Fractional seconds fields are non-zero	01004
SQL_TIMESTAMP	Data value is a valid timestamp	00000
	Data value is not a valid timestamp	22008

**Note:**

- a The fractional seconds of the timestamp are truncated.
- b The time portion of the timestamp is deleted.
- c The date portion of the timestamp is deleted.

SQLSTATE **00000** is not returned by `SQLERROR()`, rather it is indicated when the function returns `SQL_SUCCESS`.



## C to SQL data conversion examples

Table 171. C to SQL data conversion examples

C Data Type	C Data Value	SQL Data Type	Column Length	SQL Data Value	SQLSTATE
SQL_C_CHAR	abcdef\0	SQL_CHAR	6	abcdef	00000
SQL_C_CHAR	abcdef\0	SQL_CHAR	5	abcde	01004
SQL_C_CHAR	1234.56\0	SQL_DECIMAL	6	1234.56	00000
SQL_C_CHAR	1234.56\0	SQL_DECIMAL	5	1234.5	01004
SQL_C_CHAR	1234.56\0	SQL_DECIMAL	3	---	22003
SQL_C_FLOAT	1234.56	SQL_FLOAT	Not applicable	1234.56	00000
SQL_C_FLOAT	1234.56	SQL_INTEGER	Not applicable	1234	01004

**Note:**

SQLSTATE **00000** is not returned by `SQLError()`, rather it is indicated when the function returns `SQL_SUCCESS`.



## Appendix E. Example code

This section provides an extensive example of a DB2 ODBC client application (APD29) calling a DB2 ODBC stored procedure (SPD29). It includes very fundamental processing of query result sets (a query cursor opened in a stored procedure and return to client for fetching). For completeness, the CREATE TABLE, data INSERTs and CREATE PROCEDURE definition is provided.

### STEP 1. Create table

```
printf("\nAPDDL SQLExecDirect stmt=%d",__LINE__);
strcpy((char *)sqlstmt,
"CREATE TABLE TABLE2A (INT4 INTEGER,SMINT SMALLINT,FLOAT8 FLOAT");
strcat((char *)sqlstmt,
",DEC312 DECIMAL(31,2),CHR10 CHARACTER(10),VCHR20 VARCHAR(20)");
strcat((char *)sqlstmt,
",LVCHR LONG VARCHAR,CHRSB CHAR(10),CHRBIT CHAR(10) FOR BIT DATA");
strcat((char *)sqlstmt,
",DDATE DATE,TTIME TIME,TSTMP TIMESTAMP)");
printf("\nAPDDL sqlstmt=%s",sqlstmt);
rc=SQLExecDirect(hstmt,sqlstmt,SQL_NTS);
if( rc != SQL_SUCCESS ) goto dberror;
```

### STEP 2. Insert 101 rows into table

```
/* insert 100 rows into table2a */
for (jx=1;jx<=100 ;jx++ ) {
printf("\nAPDIN SQLExecDirect stmt=%d",__LINE__);
strcpy((char *)sqlstmt,"insert into table2a values(");
sprintf((char *)sqlstmt+strlen((char *)sqlstmt),"%ld",jx);
strcat((char *)sqlstmt,
",4,8.2E+30,1515151515151.51,'CHAR','VCHAR','LVCCHAR','SBCS'");
strcat((char *)sqlstmt,
", 'MIXED','01/01/1991','3:33 PM','1999-09-09-09.09.09.090909')");
printf("\nAPDIN sqlstmt=%s",sqlstmt);
rc=SQLExecDirect(hstmt,sqlstmt,SQL_NTS);
if( rc != SQL_SUCCESS ) goto dberror;
} /* endfor */
```

### STEP 3. Define stored procedure with CREATE PROCEDURE SQL statement

```
CREATE PROCEDURE SPD29
(INOUT INTEGER)
PROGRAM TYPE MAIN
EXTERNAL NAME SPD29
COLLID DSNAOCLI
LANGUAGE C
RESULT SET 2
MODIFIES SQL DATA
PARAMETER STYLE GENERAL
NO WLM ENVIRONMENT;
```

STEP 4. Stored procedure

```

/*START OF SPD29*****
/* PRAGMA TO CALL PLI SUBRTN CSPSUB TO ISSUE CONSOLE MSGS */
#pragma options (rent)
#pragma runopts(plist(os))
/*****
/* Include the 'C' include files */
/*****
#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include "sqlcli1.h"
#include <sqlca.h>
#include <decimal.h>
#include <wchar.h>
/*****
/* Variables for COMPARE routines */
/*****
#ifndef NULL
#define NULL 0
#endif

SQLHENV henv = SQL_NULL_HENV;
SQLHDBC hdbc = SQL_NULL_HDBC;
SQLHSTMT hstmt = SQL_NULL_HSTMT;
SQLHSTMT hstmt2 = SQL_NULL_HSTMT;
SQLRETURN rc = SQL_SUCCESS;
SQLINTEGER id;
SQLCHAR name[51];
SQLINTEGER namelen, intlen, colcount;
SQLSMALLINT scale;
struct sqlca sqlca;
SQLCHAR server[18];
SQLCHAR uid[30];
SQLCHAR pwd[30];
SQLCHAR sqlstmt[500];
SQLCHAR sqlstmt2[500];
SQLSMALLINT pcpair=0;
SQLSMALLINT pccol=0;
SQLCHAR cursor[19];
SQLSMALLINT cursor_len;

SQLINTEGER SPCODE;
struct {
    SQLSMALLINT LEN;
    SQLCHAR DATA_200"; } STMTSQL;

SQLSMALLINT H1SMINT;
SQLINTEGER H1INT4;
SQLDOUBLE H1FLOAT8;
SQLDOUBLE H1DEC312;
SQLCHAR H1CHR10[11];
SQLCHAR H1VCHR20[21];
SQLCHAR H1LVCHR[21];
SQLCHAR H1CHRSB[11];
SQLCHAR H1CHRBIT[11];
SQLCHAR H1DDATE[11];
SQLCHAR H1TTIME[9];
SQLCHAR H1TSTMP[27];

```

SQLSMALLINT	I1SMINT;
SQLSMALLINT	I1INT4;
SQLSMALLINT	I1FLOAT8;
SQLSMALLINT	I1DEC312;
SQLSMALLINT	I1CHR10;
SQLSMALLINT	I1VCHR20;
SQLSMALLINT	I1LVCHR;
SQLSMALLINT	I1CHRSB;
SQLSMALLINT	I1CHRBIT;
SQLSMALLINT	I1DDATE;
SQLSMALLINT	I1TTIME;
SQLSMALLINT	I1TSTMP;
SQLINTEGER	LEN_H1SMINT;
SQLINTEGER	LEN_H1INT4;
SQLINTEGER	LEN_H1FLOAT8;
SQLINTEGER	LEN_H1DEC312;
SQLINTEGER	LEN_H1CHR10;
SQLINTEGER	LEN_H1VCHR20;
SQLINTEGER	LEN_H1LVCHR;
SQLINTEGER	LEN_H1CHRSB;
SQLINTEGER	LEN_H1CHRBIT;
SQLINTEGER	LEN_H1DDATE;
SQLINTEGER	LEN_H1TTIME;
SQLINTEGER	LEN_H1TSTMP;
SQLSMALLINT	H2SMINT;
SQLINTEGER	H2INT4;
SQLDOUBLE	H2FLOAT8;
SQLCHAR	H2CHR10[11];
SQLCHAR	H2VCHR20[21];
SQLCHAR	H2LVCHR[21];
SQLCHAR	H2CHRSB[11];
SQLCHAR	H2CHRBIT[11];
SQLCHAR	H2DDATE[11];
SQLCHAR	H2TTIME[9];
SQLCHAR	H2TSTMP[27];
SQLSMALLINT	I2SMINT;
SQLSMALLINT	I2INT4;
SQLSMALLINT	I2FLOAT8;
SQLSMALLINT	I2CHR10;
SQLSMALLINT	I2VCHR20;
SQLSMALLINT	I2LVCHR;
SQLSMALLINT	I2CHRSB;
SQLSMALLINT	I2CHRBIT;
SQLSMALLINT	I2DDATE;
SQLSMALLINT	I2TTIME;
SQLSMALLINT	I2TSTMP;
SQLINTEGER	LEN_H2SMINT;
SQLINTEGER	LEN_H2INT4;
SQLINTEGER	LEN_H2FLOAT8;
SQLINTEGER	LEN_H2CHR10;
SQLINTEGER	LEN_H2VCHR20;
SQLINTEGER	LEN_H2LVCHR;
SQLINTEGER	LEN_H2CHRSB;
SQLINTEGER	LEN_H2CHRBIT;
SQLINTEGER	LEN_H2DDATE;
SQLINTEGER	LEN_H2TTIME;
SQLINTEGER	LEN_H2TSTMP;

## Example code

```
SQLCHAR locsite[18] = "stlec1";
SQLCHAR remsite[18] = "stlec1b";

SQLCHAR    spname[8];
SQLINTEGER ix,jx,locix;
SQLINTEGER result;
SQLCHAR    state_blank[6] = "    ";

SQLRETURN
check_error(SQLHENV henv,
            SQLHDBC hdbc,
            SQLHSTMT hstmt,
            SQLRETURN frc);

SQLRETURN
prt_sqlca();
/*****
/* Main Program
*****/
SQLINTEGER
main(SQLINTEGER argc, SQLCHAR *argv[] )
{
    printf("\nSPD29 INITIALIZATION");
    scale = 0;
    rc=0;

    rc=0;
    SPCODE=0;

    /* argv0 = sp module name */
    if (argc != 2)
    {
        printf("SPD29 parm number error\n ");
        printf("SPD29 EXPECTED =%d\n",3);
        printf("SPD29 received =%d\n",argc);
        goto dberror;
    }
    strcpy((char *)spname,(char *)argv[0]);
    result = strcmp((char *)spname,"SPD29",5);
    if (result != 0)
    {
        printf("SPD29 argv0 sp name error\n ");
        printf("SPD29 compare rresult =%i\n",result);
        printf("SPD29 expected =%s\n","SPD29");
        printf("SPD29 received spname=%s\n",spname);
        printf("SPD29 received argv0 =%s\n",argv[0]);
        goto dberror;
    }
    /* get input spcode value */
    SPCODE = *(SQLINTEGER *) argv[1];
    printf("\nSPD29 SQLAllocEnv      number=      1\n");
    henv=0;
    rc = SQLAllocEnv(&henv);
    if( rc != SQL_SUCCESS ) goto dberror;
    printf("\nSPD29-henv=%i",henv);
    /*****
    printf("\nSPD29 SQLAllocConnect
    hdbc=0;
    SQLAllocConnect(henv, &hdbc);
    if( rc != SQL_SUCCESS ) goto dberror;
    printf("\nSPD29-hdbc=%i",hdbc);
    *****/
}
```

```

/*****/
/* Make sure no autocommits after cursors are allocated, commits */
/* cause sp failure. No-autocommit could also be specified in the*/
/* INI file. */
/* Also, sp could be defined with COMMIT_ON_RETURN in the */
/* DB2 catalog table SYSIBM.SYSPROCEDURES, but be wary that this */
/* removes control from the client appl to control commit scope. */
/*****/
printf("\nSPD29 SQLSetConnectOption-no autocommits in stored procs");
rc = SQLSetConnectOption(hdbc,SQL_AUTOCOMMIT,SQL_AUTOCOMMIT_OFF);
if( rc != SQL_SUCCESS ) goto dberror;
/*****/
printf("\nSPD29 SQLConnect NULL connect in stored proc ");
strcpy((char *)uid,"cliuser");
strcpy((char *)pwd,"password");
printf("\nSPD29 server=%s",NULL);
printf("\nSPD29 uid=%s",uid);
printf("\nSPD29 pwd=%s",pwd);
rc=SQLConnect(hdbc, NULL, 0, uid, SQL_NTS, pwd, SQL_NTS);
if( rc != SQL_SUCCESS ) goto dberror;
/*****/
/* Start SQL statements *****/
/*****/
switch(SPCODE)
{
/*****/
/* CASE(SPCODE=0) do nothing and return *****/
/*****/
case 0:
break;
case 1:
/*****/
/* CASE(SPCODE=1) *****/
/* -sqlprepare/sqlexecute insert int4=200 *****/
/* -sqlexecdirect insert int4=201 *****/
/* *validated in client appl that inserts occur *****/
/*****/
SPCODE=0;

printf("\nSPD29 SQLAllocStmt \n");
hstmt=0;
rc=SQLAllocStmt(hdbc, &hstmt);
if( rc != SQL_SUCCESS ) goto dberror;
printf("\nSPD29-hstmt=%i\n",hstmt);

printf("\nSPD29 SQLPrepare \n");
strcpy((char *)sqlstmt,
"insert into TABLE2A(int4) values(?)");
printf("\nSPD29 sqlstmt=%s",sqlstmt);
rc=SQLPrepare(hstmt,sqlstmt,SQL_NTS);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nSPD29 SQLNumParams \n");
rc=SQLNumParams(hstmt,&pcpar);
if( rc != SQL_SUCCESS) goto dberror;
if (pcpar!=1) {
printf("\nSPD29 incorrect pcpar=%d",pcpar);
goto dberror;
}
}

```

## Example code

```
printf("\nSPD29 SQLBindParameter  int4          \n");
H1INT4=200;
LEN_H1INT4=sizeof(H1INT4);
rc=SQLBindParameter(hstmt,1,SQL_PARAM_INPUT,SQL_C_LONG,
SQL_INTEGER,0,0,&H1INT4,0,(SQLINTEGER *)&LEN_H1INT4);
if( rc != SQL_SUCCESS) goto dberror;

printf("\nSPD29 SQLExecute          \n");
rc=SQLExecute(hstmt);
if( rc != SQL_SUCCESS) goto dberror;

printf("\nSPD29 SQLFreeStmt          \n");
rc=SQLFreeStmt(hstmt, SQL_DROP);
if( rc != SQL_SUCCESS ) goto dberror;
/*****
printf("\nAPDIN SQLAllocStmt  stmt=%d",__LINE__);
hstmt=0;
rc=SQLAllocStmt(hdbc, &hstmt);
if( rc != SQL_SUCCESS ) goto dberror;
printf("\nAPDIN-hstmt=%i\n",hstmt);

jx=201;
printf("\nAPDIN SQLExecDirect stmt=%d",__LINE__);
strcpy((char *)sqlstmt,"insert into table2a values(");
sprintf((char *)sqlstmt+strlen((char *)sqlstmt),"%ld",jx);
strcat((char *)sqlstmt,
",4,8.2E+30,1515151515151.51,'CHAR','VCHAR','LVCHAR','SBCS'");
strcat((char *)sqlstmt,
", 'MIXED', '01/01/1991', '3:33 PM', '1999-09-09-09.09.09.090909'");
printf("\nAPDIN sqlstmt=%s",sqlstmt);
rc=SQLExecDirect(hstmt,sqlstmt,SQL_NTS);
if( rc != SQL_SUCCESS ) goto dberror;

break;
/*****
case 2:
/*****
/* CASE(SPCODE=2)          *****/
/* -sqlprepare/sqlexecute select int4  from table2a          *****/
/* -sqlprepare/sqlexecute select chr10 from table2a          *****/
/* *qrs cursors should be allocated and left open by CLI     *****/
/*****
SPCODE=0;

/* generate 1st queryresultset */
printf("\nSPD29 SQLAllocStmt          \n");
hstmt=0;
rc=SQLAllocStmt(hdbc, &hstmt);
if( rc != SQL_SUCCESS ) goto dberror;
printf("\nSPD29-hstmt=%i\n",hstmt);

printf("\nSPD29 SQLPrepare          \n");
strcpy((char *)sqlstmt,
"SELECT INT4 FROM TABLE2A");
printf("\nSPD29 sqlstmt=%s",sqlstmt);
rc=SQLPrepare(hstmt,sqlstmt,SQL_NTS);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nSPD29 SQLExecute          \n");
rc=SQLExecute(hstmt);
if( rc != SQL_SUCCESS ) goto dberror;
```



```

/* allocate 2nd stmt handle for 2nd queryresultset */
/* generate 2nd queryresultset */
printf("\nSPD29 SQLAllocStmt          \n");
hstmt=0;
rc=SQLAllocStmt(hdbc, &hstmt2);
if( rc != SQL_SUCCESS ) goto dberror;
printf("\nSPD29-hstmt2=%i\n",hstmt2);

printf("\nSPD29 SQLPrepare          \n");
strcpy((char *)sqlstmt2,
"SELECT CHR10 FROM TABLE2A");
printf("\nSPD29 sqlstmt2=%s",sqlstmt2);
rc=SQLPrepare(hstmt2,sqlstmt2,SQL_NTS);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nSPD29 SQLExecute          \n");
rc=SQLExecute(hstmt2);
if( rc != SQL_SUCCESS ) goto dberror;

/*leave queryresultset cursor open for fetch back at client appl */

break;
/*****/
default:
{
printf("SPD29 INPUT SPCODE INVALID\n");
printf("SPD29...EXPECTED SPCODE=0-2\n");
printf("SPD29...RECIEVED SPCODE=%i\n",SPCODE);
goto dberror;
break;
}
}
/*****/
/* End SQL statements *****/
/*****/
/*Be sure NOT to put a SQLTransact with SQL_COMMIT in a DB2/MVS */
/* stored procedure. Commit is not allowed in a DB2/MVS */
/* stored procedure. Use SQLTransact with SQL_ROLLBACK to */
/* force a must rollback condition for this sp and calling */
/* client application. */
/*****/
printf("\nSPD29 SQLDisconnect      number=    4\n");
rc=SQLDisconnect(hdbc);
if( rc != SQL_SUCCESS ) goto dberror;
/*****/
printf("\nSPD29 SQLFreeConnect      number=    5\n");
rc = SQLFreeConnect(hdbc);
if( rc != SQL_SUCCESS ) goto dberror;
/*****/
printf("\nSPD29 SQLFreeEnv          number=    6\n");
rc = SQLFreeEnv(henv);
if( rc != SQL_SUCCESS ) goto dberror;
/*****/
goto pgmend;

dberror:
printf("\nSPD29 entry dberror label");
printf("\nSPD29 rc=%d",rc);
check_error(henv,hdbc,hstmt,rc);
printf("\nSPD29 SQLFreeEnv          number=    7\n");
rc = SQLFreeEnv(henv);

```

## Example code

```
printf("\nSPD29 rc=%d",rc);
rc=12;
rc=12;
SPCODE=12;
goto pgmend;

pgmend:

printf("\nSPD29  TERMINATION  ");
if (rc!=0)
{
printf("\nSPD29 WAS NOT SUCCESSFUL");
printf("\nSPD29 SPCODE = %i", SPCODE    );
printf("\nSPD29 rc      = %i", rc    );
}
else
{
printf("\nSPD29 WAS SUCCESSFUL");
}
/* assign output spcode value */
*(SQLINTEGER *) argv[1] = SPCODE;
exit;
} /*END MAIN*/
/*****
** check_error - call print_error(), checks severity of return code
*****/
SQLRETURN
check_error(SQLHENV henv,
            SQLHDBC hdbc,
            SQLHSTMT hstmt,
            SQLRETURN frc )
{
SQLCHAR      buffer[SQL_MAX_MESSAGE_LENGTH + 1];
SQLCHAR      cli_sqlstate[SQL_SQLSTATE_SIZE + 1];
SQLINTEGER   cli_sqlcode;
SQLSMALLINT  length;

printf("\nSPD29 entry check_error rtn");

switch (frc) {
case SQL_SUCCESS:
break;
case SQL_INVALID_HANDLE:
printf("\nSPD29 check_error> SQL_INVALID_HANDLE ");
case SQL_ERROR:
printf("\nSPD29 check_error> SQL_ERROR ");
break;
case SQL_SUCCESS_WITH_INFO:
printf("\nSPD29 check_error> SQL_SUCCESS_WITH_INFO");
break;
case SQL_NO_DATA_FOUND:
printf("\nSPD29 check_error> SQL_NO_DATA_FOUND ");
break;
default:
printf("\nSPD29 check_error> Invalid rc from api rc=%i",frc);
break;
} /*end switch*/
```

```

printf("\nSPD29 SQLError ");
while ((rc=SQLError(henv, hdbc, hstmt, cli_sqlstate, &cli_sqlcode,
  buffer,SQL_MAX_MESSAGE_LENGTH + 1, &length)) == SQL_SUCCESS) {
  printf("      SQLSTATE: %s", cli_sqlstate);
  printf("Native Error Code: %ld", cli_sqlcode);
  printf("%s ", buffer);
};
if (rc!=SQL_NO_DATA_FOUND)
  printf("SQLError api call failed rc=%d",rc);

printf("\nSPD29 SQLGetSQLCA ");
rc = SQLGetSQLCA(henv, hdbc, hstmt, &sqlca);
if( rc == SQL_SUCCESS )
  prt_sqlca();
else
  printf("\n SPD29-check_error SQLGetSQLCA failed rc=%i",rc);

return (frc);
}
/*****
/*          P r i n t   S Q L C A          */
*****/
SQLRETURN
prt_sqlca()
{
  SQLINTEGER i;
  printf("\nSPD29 entry prts_sqlca rtn");
  printf("\r\r*** Printing the SQLCA:\r");
  printf("\nSQLCAID .... %s",sqlca.sqlcaid);
  printf("\nSQLCABC .... %d",sqlca.sqlcabc);
  printf("\nSQLCODE .... %d",sqlca.sqlcode);
  printf("\nSQLERRML ... %d",sqlca.sqlerrml);
  printf("\nSQLERRMC ... %s",sqlca.sqlerrmc);
  printf("\nSQLERRP ... %s",sqlca.sqlerrp);
  for (i = 0; i < 6; i++)
    printf("\nSQLERRD%d ... %d",i+1,sqlca.sqlerrd??(i??));
  for (i = 0; i < 10; i++)
    printf("\nSQLWARN%d ... %c",i,sqlca.sqlwarn[i]);
  printf("\nSQLWARNA ... %c",sqlca.sqlwarn[10]);
  printf("\nSQLSTATE ... %s",sqlca.sqlstate);

  return(0);
}
/* End of prtsqlca */
/*****
/*END OF SPD29 *****/

```

## Example code

### STEP 5. Client application

```
/******  
/*START OF SPD29*****  
/* SCEANRIO PSEUDOCODE: */  
/* APD29(CLI CODE CLIENT APPL) */  
/* -CALL SPD29 (CLI CODE STORED PROCEDURE APPL) */  
/* -SPCODE=0 */  
/* -PRINTF MSGS (CHECK SDSF FOR SPAS ADDR TO VERIFY) */  
/* -SPCODE=1 */  
/* -PRINTF MSGS (CHECK SDSF FOR SPAS ADDR TO VERIFY) */  
/* -SQLPREPARE/EXECUTE INSERT INT4=200 */  
/* -SQLEXECDIRECT INSERT INT4=201 */  
/* -SPCODE=2 */  
/* -PRINTF MSGS (CHECK SDSF FOR SPAS ADDR TO VERIFY) */  
/* -SQLPREPARE/EXECUTE SELECT INT4 FROM TABLE2A */  
/* -SQLPREPARE/EXECUTE SELECT CHR10 FROM TABLE2A */  
/* (CLI CURSORS OPENED 'WITH RETURN')... */  
/* -RETURN */  
/* -FETCH QRS FROM SP CURSOR */  
/* -COMMIT */  
/* -VERIFY INSERTS BY SPD29 */  
/******  
/* Include the 'C' include files */  
/******  
#include <stdio.h>  
#include <string.h>  
#include <stdlib.h>  
#include "sqlcli.h"  
#include <sqlca.h>  
/******  
/* Variables for COMPARE routines */  
/******  
#ifndef NULL  
#define NULL 0  
#endif  
  
SQLHENV henv = SQL_NULL_HENV;  
SQLHDBC hdbc = SQL_NULL_HDBC;  
SQLHSTMT hstmt = SQL_NULL_HSTMT;  
SQLRETURN rc = SQL_SUCCESS;  
SQLINTEGER id;  
SQLCHAR name[51];  
SQLINTEGER namelen, intlen, colcount;  
SQLSMALLINT scale;  
struct sqlca sqlca;  
SQLCHAR server[18];  
SQLCHAR uid[30];  
SQLCHAR pwd[30];  
SQLCHAR sqlstmt[250];  
SQLSMALLINT pcpair=0;  
SQLSMALLINT pccol=0;  
  
SQLINTEGER SPCODE;  
struct {  
SQLSMALLINT LEN;  
SQLCHAR DATA[200]; } STMTSQL;
```

SQLSMALLINT	H1SMINT;
SQLINTEGER	H1INT4;
SQLDOUBLE	H1FLOAT8;
SQLDOUBLE	H1DEC312;
SQLCHAR	H1CHR10[11];
SQLCHAR	H1VCHR20[21];
SQLCHAR	H1LVCHR[21];
SQLCHAR	H1CHRSB[11];
SQLCHAR	H1CHRBIT[11];
SQLCHAR	H1DDATE[11];
SQLCHAR	H1TTIME[9];
SQLCHAR	H1TSTMP[27];

SQLSMALLINT	I1SMINT;
SQLSMALLINT	I1INT4;
SQLSMALLINT	I1FLOAT8;
SQLSMALLINT	I1DEC312;
SQLSMALLINT	I1CHR10;
SQLSMALLINT	I1VCHR20;
SQLSMALLINT	I1LVCHR;
SQLSMALLINT	I1CHRSB;
SQLSMALLINT	I1CHRBIT;
SQLSMALLINT	I1DDATE;
SQLSMALLINT	I1TTIME;
SQLSMALLINT	I1TSTMP;

SQLINTEGER	LNH1SMINT;
SQLINTEGER	LNH1INT4;
SQLINTEGER	LNH1FLOAT8;
SQLINTEGER	LNH1DEC312;
SQLINTEGER	LNH1CHR10;
SQLINTEGER	LNH1VCHR20;
SQLINTEGER	LNH1LVCHR;
SQLINTEGER	LNH1CHRSB;
SQLINTEGER	LNH1CHRBIT;
SQLINTEGER	LNH1DDATE;
SQLINTEGER	LNH1TTIME;
SQLINTEGER	LNH1TSTMP;

SQLSMALLINT	H2SMINT;
SQLINTEGER	H2INT4;
SQLDOUBLE	H2FLOAT8;
SQLCHAR	H2CHR10[11];
SQLCHAR	H2VCHR20[21];
SQLCHAR	H2LVCHR[21];
SQLCHAR	H2CHRSB[11];
SQLCHAR	H2CHRBIT[11];
SQLCHAR	H2DDATE[11];
SQLCHAR	H2TTIME[9];
SQLCHAR	H2TSTMP[27];

SQLSMALLINT	I2SMINT;
SQLSMALLINT	I2INT4;
SQLSMALLINT	I2FLOAT8;
SQLSMALLINT	I2CHR10;
SQLSMALLINT	I2VCHR20;
SQLSMALLINT	I2LVCHR;
SQLSMALLINT	I2CHRSB;
SQLSMALLINT	I2CHRBIT;
SQLSMALLINT	I2DDATE;
SQLSMALLINT	I2TTIME;
SQLSMALLINT	I2TSTMP;

## Example code

```
SQLINTEGER          LNH2SMINT;
SQLINTEGER          LNH2INT4;
SQLINTEGER          LNH2FLOAT8;
SQLINTEGER          LNH2CHR10;
SQLINTEGER          LNH2VCHR20;
SQLINTEGER          LNH2LVCHR;
SQLINTEGER          LNH2CHRSB;
SQLINTEGER          LNH2CHRBIT;
SQLINTEGER          LNH2DDATE;
SQLINTEGER          LNH2TTIME;
SQLINTEGER          LNH2TSTMP;

SQLCHAR locsite[18] = "stlec1";
SQLCHAR remsite[18] = "stlec1b";

SQLINTEGER  ix,jx,locix;
SQLINTEGER  result;
SQLCHAR    state_blank[6] = "    ";

SQLRETURN
check_error(SQLHENV henv,
            SQLHDBC hdbc,
            SQLHSTMT hstmt,
            SQLRETURN frc);

SQLRETURN
prt_sqlca();
/*****
/* Main Program */
*****/
SQLINTEGER
main()
{
    printf("\nAPD29 INITIALIZATION");
    scale = 0;
    rc=0;

    printf("\nAPD29 SQLAllocEnv  stmt=%d",__LINE__);
    henv=0;
    rc = SQLAllocEnv(&henv);
    if( rc != SQL_SUCCESS ) goto dberror;
    printf("\nAPD29-henv=%i",henv);

    for (locix=1;locix<=2;locix++)
    {
        /* Start SQL statements *****/
        /*****
        printf("\nAPD29 SQLAllocConnect          ");
        hdbc=0;
        SQLAllocConnect(henv, &hdbc);
        if( rc != SQL_SUCCESS ) goto dberror;
        printf("\nAPD29-hdbc=%i",hdbc);
        /*****
        printf("\nAPD29 SQLConnect          ");
        if (locix == 1)
        {
            strcpy((char *)server,(char *)locsite);
        }
        else
        {
            strcpy((char *)server,(char *)remsite);
        }
    }
}
```

```

strcpy((char *)uid,"cliuser");
strcpy((char *)pwd,"password");
printf("\nAPD29 server=%s",server);
printf("\nAPD29 uid=%s",uid);
printf("\nAPD29 pwd=%s",pwd);
rc=SQLConnect(hdbc, server, SQL_NTS, uid, SQL_NTS, pwd, SQL_NTS);
if( rc != SQL_SUCCESS ) goto dberror;
/*****
/* CASE(SPCODE=0) QRS RETURNED=0 COL=0 ROW=0 */
*****/
printf("\nAPD29 SQLAllocStmt stmt=%d",__LINE__);
hstmt=0;
rc=SQLAllocStmt(hdbc, &hstmt);
if( rc != SQL_SUCCESS ) goto dberror;
printf("\nAPD29-hstmt=%i\n",hstmt);

SPCODE=0;
printf("\nAPD29 call sp SPCODE =%i\n",SPCODE);
printf("\nAPD29 SQLPrepare stmt=%d",__LINE__);
strcpy((char*)sqlstmt,"CALL SPD29(?)");
printf("\nAPD29 sqlstmt=%s",sqlstmt);
rc=SQLPrepare(hstmt,sqlstmt,SQL_NTS);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nAPD29 SQLBindParameter stmt=%d",__LINE__);
rc = SQLBindParameter(hstmt,
                        1,
                        SQL_PARAM_INPUT_OUTPUT,
                        SQL_C_LONG,
                        SQL_INTEGER,
                        0,
                        0,
                        &SPCODE,
                        0,
                        NULL);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nAPD29 SQLExecute stmt=%d",__LINE__);
rc=SQLExecute(hstmt);
if( rc != SQL_SUCCESS ) goto dberror;
if( SPCODE != 0 )
{
    printf("\nAPD29 SPCODE not zero, spcode=%i\n",SPCODE);
    goto dberror;
}

printf("\nAPD29 SQLTransact stmt=%d",__LINE__);
rc=SQLTransact(henv, hdbc, SQL_COMMIT);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nAPD29 SQLFreeStmt stmt=%d",__LINE__);
rc=SQLFreeStmt(hstmt, SQL_DROP);
if( rc != SQL_SUCCESS ) goto dberror;
/*****
/* CASE(SPCODE=1) QRS RETURNED=0 COL=0 ROW=0 */
*****/
printf("\nAPD29 SQLAllocStmt stmt=%d",__LINE__);
hstmt=0;
rc=SQLAllocStmt(hdbc, &hstmt);
if( rc != SQL_SUCCESS ) goto dberror;
printf("\nAPD29-hstmt=%i\n",hstmt);

```

## Example code

```
SPCODE=1;
printf("\nAPD29 call sp SPCODE =%i\n",SPCODE);
printf("\nAPD29 SQLPrepare stmt=%d",__LINE__);
strcpy((char*)sqlstmt,"CALL SPD29(?)");
printf("\nAPD29 sqlstmt=%s",sqlstmt);
rc=SQLPrepare(hstmt,sqlstmt,SQL_NTS);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nAPD29 SQLBindParameter stmt=%d",__LINE__);
rc = SQLBindParameter(hstmt,
                      1,
                      SQL_PARAM_INPUT_OUTPUT,
                      SQL_C_LONG,
                      SQL_INTEGER,
                      0,
                      0,
                      &SPCODE,
                      0,
                      NULL);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nAPD29 SQLExecute stmt=%d",__LINE__);
rc=SQLExecute(hstmt);
if( rc != SQL_SUCCESS ) goto dberror;
if( SPCODE != 0 )
{
    printf("\nAPD29 SPCODE not zero, spcode=%i\n",SPCODE);
    goto dberror;
}

printf("\nAPD29 SQLTransact stmt=%d",__LINE__);
rc=SQLTransact(henv, hdbc, SQL_COMMIT);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nAPD29 SQLFreeStmt stmt=%d",__LINE__);
rc=SQLFreeStmt(hstmt, SQL_DROP);
if( rc != SQL_SUCCESS ) goto dberror;
/*****
/* CASE(SPCODE=2) QRS RETURNED=2 COL=1(int4/chr10) ROW=100+ */
*****/
printf("\nAPD29 SQLAllocStmt          number= 18\n");
hstmt=0;
rc=SQLAllocStmt(hdbc, &hstmt);
if( rc != SQL_SUCCESS ) goto dberror;
printf("\nAPD29-hstmt=%i\n",hstmt);

SPCODE=2;
printf("\nAPD29 call sp SPCODE =%i\n",SPCODE);
printf("\nAPD29 SQLPrepare          number= 19\n");
strcpy((char*)sqlstmt,"CALL SPD29(?)");
printf("\nAPD29 sqlstmt=%s",sqlstmt);
rc=SQLPrepare(hstmt,sqlstmt,SQL_NTS);
if( rc != SQL_SUCCESS ) goto dberror;
```



```

printf("\nAPD29 SQLBindParameter      number= 20\n");
rc = SQLBindParameter(hstmt,
                      1,
                      SQL_PARAM_INPUT_OUTPUT,
                      SQL_C_LONG,
                      SQL_INTEGER,
                      0,
                      0,
                      &SPCODE,
                      0,
                      NULL);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nAPD29 SQLExecute            number= 21\n");
rc=SQLExecute(hstmt);
if( rc != SQL_SUCCESS ) goto dberror;
if( SPCODE != 0 )
{
    printf("\nAPD29 spcode incorrect");
    goto dberror;
}

printf("\nAPD29 SQLNumResultCols      number= 22\n");
rc=SQLNumResultCols(hstmt,&pccol);
if (pccol!=1)
{
    printf("APD29 col count wrong=%i\n",pccol);
    goto dberror;
}

printf("\nAPD29 SQLBindCol            number= 23\n");
rc=SQLBindCol(hstmt,
              1,
              SQL_C_LONG,
              (SQLPOINTER) &H1INT4,
              (SQLINTEGER)sizeof(SQLINTEGER),
              (SQLINTEGER *) &LNH1INT4
              );
if( rc != SQL_SUCCESS ) goto dberror;

jx=0;
printf("\nAPD29 SQLFetch              number= 24\n");
while ((rc = SQLFetch(hstmt)) == SQL_SUCCESS)
{
    jx++;
    printf("\nAPD29 fetch loop jx =%i\n",jx);
    if ( (H1INT4<=0) || (H1INT4>=202)
        || (LNH1INT4!=4 && LNH1INT4!=-1) )
    { /* data error */
        printf("\nAPD29 H1INT4=%i\n",H1INT4);
        printf("\nAPD29 LNH1INT4=%i\n",LNH1INT4);
        goto dberror;
    }
    printf("\nAPD29 SQLFetch            number= 24\n");
} /* end while loop */

if( rc != SQL_NO_DATA_FOUND )
{
    printf("\nAPD29 invalid end of data\n");
    goto dberror;
}

```

## Example code

```
printf("\nAPD29 SQLMoreResults      number= 25\n");
rc=SQLMoreResults(hstmt);
if(rc != SQL_SUCCESS) goto dberror;

printf("\nAPD29 SQLNumResultCols    number= 26\n");
rc=SQLNumResultCols(hstmt,&pccol);
if (pccol!=1) {
    printf("APD29 col count wrong=%i\n",pccol);
    goto dberror;
}

printf("\nAPD29 SQLBindCol          number= 27\n");
rc=SQLBindCol(hstmt,
              1,
              SQL_C_CHAR,
              (SQLPOINTER) H1CHR10,
              (SQLINTEGER)sizeof(H1CHR10),
              (SQLINTEGER *) &LNH1CHR10
              );
if( rc != SQL_SUCCESS ) goto dberror;

jx=0;
while ((rc = SQLFetch(hstmt)) == SQL_SUCCESS)
{
    jx++;
    printf("\nAPD29 fetch loop jx =%i\n",jx);
    result=strcmp((char *)H1CHR10,"CHAR");
    if ( (result!=0)
        || (LNH1INT4!=4 && LNH1INT4!=-1) )
    {
        printf("\nAPD29 H1CHR10=%s\n",H1CHR10);
        printf("\nAPD29 result=%i\n",result);
        printf("\nAPD29 LNH1CHR10=%i\n",LNH1CHR10);
        printf("\nAPD29 strlen(H1CHR10)=%i\n",strlen((char *)H1CHR10));
        goto dberror;
    }
    printf("\nAPD29 SQLFetch          number= 24\n");
} /* end while loop */

if( rc != SQL_NO_DATA_FOUND )
    goto dberror;

printf("\nAPD29 SQLMoreResults      number= 29\n");
rc=SQLMoreResults(hstmt);
if( rc != SQL_NO_DATA_FOUND) goto dberror;

printf("\nAPD29 SQLTransact         number= 30\n");
rc=SQLTransact(henv, hdbc, SQL_COMMIT);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nAPD29 SQLFreeStmnt        number= 31\n");
rc=SQLFreeStmnt(hstmt, SQL_DROP);
if( rc != SQL_SUCCESS ) goto dberror;
/*****
printf("\nAPD29 SQLDisconnect stmt=%d",__LINE__);
rc=SQLDisconnect(hdbc);
if( rc != SQL_SUCCESS ) goto dberror;
*****/
printf("\nSQLFreeConnect stmt=%d",__LINE__);
rc=SQLFreeConnect(hdbc);
if( rc != SQL_SUCCESS ) goto dberror;
/*****
/* End SQL statements *****/
```

```

} /* end for each site perform these stmts */

for (locix=1;locix<=2;locix++)
{
  /******
  printf("\nAPD29 SQLAllocConnect          ");
  hdbc=0;
  SQLAllocConnect(henv, &hdbc);
  if( rc != SQL_SUCCESS ) goto dberror;
  printf("\nAPD29-hdbc=%i",hdbc);
  /******
  printf("\nAPD29 SQLConnect          ");
  if (locix == 1)
  {
    strcpy((char *)server,(char *)locsite);
  }
  else
  {
    strcpy((char *)server,(char *)remsite);
  }

  strcpy((char *)uid,"cliuser");
  strcpy((char *)pwd,"password");
  printf("\nAPD29 server=%s",server);
  printf("\nAPD29 uid=%s",uid);
  printf("\nAPD29 pwd=%s",pwd);
  rc=SQLConnect(hdbc, server, SQL_NTS, uid, SQL_NTS, pwd, SQL_NTS);
  if( rc != SQL_SUCCESS ) goto dberror;
  /******
  /* Start validate SQL statements *****/
  /******
  printf("\nAPD01 SQLAllocStmt          \n");
  hstmt=0;
  rc=SQLAllocStmt(hdbc, &hstmt);
  if( rc != SQL_SUCCESS ) goto dberror;
  printf("\nAPD01-hstmt=%i\n",hstmt);

  printf("\nAPD01 SQLExecDirect          \n");
  strcpy((char *)sqlstmt,
  "SELECT INT4 FROM TABLE2A WHERE INT4=200");
  printf("\nAPD01 sqlstmt=%s",sqlstmt);
  rc=SQLExecDirect(hstmt,sqlstmt,SQL_NTS);
  if( rc != SQL_SUCCESS ) goto dberror;

  printf("\nAPD01 SQLBindCol          \n");
  rc=SQLBindCol(hstmt,
                1,
                SQL_C_LONG,
                (SQLPOINTER) &H1INT4,;
                (SQLINTEGER)sizeof(SQLINTEGER),
                (SQLINTEGER *) &LNH1INT4
                );
  if( rc != SQL_SUCCESS ) goto dberror;

  printf("\nAPD01 SQLFetch          \n");
  rc=SQLFetch(hstmt);
  if( rc != SQL_SUCCESS ) goto dberror;
  if ((H1INT4!=200) || (LNH1INT4!=4))
  {
    printf("\nAPD01 H1INT4=%i\n",H1INT4);
    printf("\nAPD01 LNH1INT4=%i\n",LNH1INT4);
    goto dberror;
  }
}

```

## Example code

```
printf("\nAPD01 SQLTransact\n");
rc=SQLTransact(henv, hdbc, SQL_COMMIT);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nAPD01 SQLFreeStmt\n");
rc=SQLFreeStmt(hstmt, SQL_CLOSE);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nAPD01 SQLExecDirect\n");
strcpy((char *)sqlstmt,
"SELECT INT4 FROM TABLE2A WHERE INT4=201");
printf("\nAPD01 sqlstmt=%s",sqlstmt);
rc=SQLExecDirect(hstmt,sqlstmt,SQL_NTS);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nAPD01 SQLFetch\n");
rc=SQLFetch(hstmt);
if( rc != SQL_SUCCESS ) goto dberror;
if ((H1INT4!=201) || (LNH1INT4!=4))
{
printf("\nAPD01 H1INT4=%i\n",H1INT4);
printf("\nAPD01 LNH1INT4=%i\n",LNH1INT4);
goto dberror;
}

printf("\nAPD01 SQLTransact\n");
rc=SQLTransact(henv, hdbc, SQL_COMMIT);
if( rc != SQL_SUCCESS ) goto dberror;

printf("\nAPD01 SQLFreeStmt\n");
rc=SQLFreeStmt(hstmt, SQL_DROP);
if( rc != SQL_SUCCESS ) goto dberror;
/*****
/* End validate SQL statements *****/
/*****
printf("\nAPD29 SQLDisconnect stmt=%d",__LINE__);
rc=SQLDisconnect(hdbc);
if( rc != SQL_SUCCESS ) goto dberror;
/*****
printf("\nSQLFreeConnect stmt=%d",__LINE__);
rc=SQLFreeConnect(hdbc);
if( rc != SQL_SUCCESS ) goto dberror;

} /* end for each site perform these stmts */
/*****
printf("\nSQLFreeEnv stmt=%d",__LINE__);
rc=SQLFreeEnv(henv);
if( rc != SQL_SUCCESS ) goto dberror;
/*****
goto pgmend;

dberror:
printf("\nAPD29 entry dberror label");
printf("\nAPD29 rc=%d",rc);
check_error(henv,hdbc,hstmt,rc);
printf("\nAPDXX SQLFreeEnv number= 6\n");
rc=SQLFreeEnv(henv);
printf("\nAPDXX FREEENV rc =%d",rc);
rc=12;
printf("\nAPDXX DBERROR set rc =%d",rc);
goto pgmend;
```

```

pgmend:
    printf("\nAPD29 TERMINATION ");
    if (rc!=0)
    {
        printf("\nAPD29 WAS NOT SUCCESSFUL");
        printf("\nAPD29 SPCODE = %i", SPCODE );
        printf("\nAPD29 rc = %i", rc );
    }
    else
        printf("\nAPD29 WAS SUCCESSFUL");

    return(rc);

} /*END MAIN*/
/*****
** check_error - call print_error(), checks severity of return code
*****/
SQLRETURN
check_error(SQLHENV henv,
            SQLHDBC hdbc,
            SQLHSTMT hstmt,
            SQLRETURN frc )
{

    SQLCHAR          buffer_SQL_MAX_MESSAGE_LENGTH + 1";
    SQLCHAR          cli_sqlstate_SQL_SQLSTATE_SIZE + 1";
    SQLINTEGER       cli_sqlcode;
    SQLSMALLINT     length;

    printf("\nAPD29 entry check_error rtn");

    switch (frc) {
    case SQL_SUCCESS:
        break;
    case SQL_INVALID_HANDLE:
        printf("\nAPD29 check_error> SQL_INVALID_HANDLE ");
    case SQL_ERROR:
        printf("\nAPD29 check_error> SQL_ERROR ");
        break;
    case SQL_SUCCESS_WITH_INFO:
        printf("\nAPD29 check_error> SQL_SUCCESS_WITH_INFO");
        break;
    case SQL_NO_DATA_FOUND:
        printf("\nAPD29 check_error> SQL_NO_DATA_FOUND ");
        break;
    default:
        printf("\nAPD29 check_error> Invalid rc from api rc=%i",frc);
        break;
    } /*end switch*/

    printf("\nAPD29 SQLError ");
    while ((rc=SQLError(henv, hdbc, hstmt, cli_sqlstate, &cli_sqlcode,
        buffer,SQL_MAX_MESSAGE_LENGTH + 1, &length)) == SQL_SUCCESS) {
        printf("          SQLSTATE: %s", cli_sqlstate);
        printf("Native Error Code: %ld", cli_sqlcode);
        printf("%s ", buffer);
    };
    if (rc!=SQL_NO_DATA_FOUND)
        printf("SQLError api call failed rc=%d",rc);

```

## Example code

```
printf("\nAPD29 SQLGetSQLCA ");
rc = SQLGetSQLCA(henv, hdbc, hstmt, &sqlca);
if( rc == SQL_SUCCESS )
    prt_sqlca();
else
    printf("\n APD29-check_error SQLGetSQLCA failed rc=%i",rc);

return (frc);
}
/*****
/*          P r i n t   S Q L C A
*****/
SQLRETURN
prt_sqlca()
{
    SQLINTEGER i;
    printf("\n1APD29 entry prts_sqlca rtn");
    printf("\r\r*** Printing the SQLCA:\r");
    printf("\nSQLCAID .... %s",sqlca.sqlcaid);
    printf("\nSQLCABC .... %d",sqlca.sqlcabc);
    printf("\nSQLCODE .... %d",sqlca.sqlcode);
    printf("\nSQLERRML ... %d",sqlca.sqlerrml);
    printf("\nSQLERRMC ... %s",sqlca.sqlerrmc);
    printf("\nSQLERRP ... %s",sqlca.sqlerrp);
    for (i = 0; i < 6; i++)
        printf("\nSQLERRD%d ... %d",i+1,sqlca.sqlerrd??(i??));
    for (i = 0; i < 10; i++)
        printf("\nSQLWARN%d ... %c",i,sqlca.sqlwarn[i]);
    printf("\nSQLWARNA ... %c",sqlca.sqlwarn[10]);
    printf("\nSQLSTATE ... %s",sqlca.sqlstate);

    return(0);
}
/* End of prtsqlca */
/*END OF APD29*****/
```

---

## Appendix F. 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.

Any references in this publication to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is as your own risk.

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.

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

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

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

---

## Programming interface information

This book is intended to help the customer write applications that use DB2 ODBC to access IBM DB2 for OS/390 servers. This book documents General-use Programming Interface and Associated Guidance Information provided by DATABASE 2 for OS/390 (DB2 for OS/390).

General-use programming interfaces allow the customer to write programs that obtain the services of DB2 for OS/390.



---

## Trademarks

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

AIX	IMS
AS/400	IMS/ESA
BookManager	Language Environment
CICS	MVS/ESA
CICS/ESA	Net.Data
CICS/MVS	OpenEdition
DATABASE 2	OS/2
DataPropagator	OS/390
DB2	OS/400
DB2 Universal Database	Parallel Sysplex
DFSMS	QMF
DFSMSHsm	RACF
Distributed Relational Database Architecture	SQL/DS
DRDA	VTAM
IBM	

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 or registered 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, and service names may be trademarks or service marks of others.



## Glossary

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

### A

**address space.** A range of virtual storage pages that is identified by a number (ASID) and a collection of segment and page tables that map the virtual pages to real pages of the computer's memory.

**address space connection.** The result of connecting an allied address space to DB2. Each address space that contains a task that is connected to DB2 has exactly one address space connection, even though more than one task control block (TCB) can be present. See also *allied address space* and *task control block*.

**allied address space.** An area of storage that is external to DB2 and that is connected to DB2. An allied address space is capable of requesting DB2 services.

**American National Standards Institute (ANSI).** An organization consisting of producers, consumers, and general interest groups, that establishes the procedures by which accredited organizations create and maintain voluntary industry standards in the United States.

**ANSI.** American National Standards Institute.

**API.** Application programming interface.

**application.** A program or set of programs that performs a task; for example, a payroll application.

**application plan.** The control structure that is produced during the bind process. DB2 uses the application plan to process SQL statements that it encounters during statement execution.

**application programming interface (API).** A functional interface that is supplied by the operating system or by a separately orderable licensed program that allows an application program that is written in a high-level language to use specific data or functions of the operating system or licensed program.

**application requester (AR).** See *requester*.

**AR.** Application requester. See *requester*.

**ASCII.** An encoding scheme that is used to represent strings in many environments, typically on PCs and workstations. Contrast with *EBCDIC*.

**attachment facility.** An interface between DB2 and TSO, IMS, CICS, or batch address spaces. An attachment facility allows application programs to access DB2.

**authorization ID.** A string that can be verified for connection to DB2 and to which a set of privileges are allowed. It can represent an individual, an organizational group, or a function, but DB2 does not determine this representation.

### 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.

**bind.** The process by which the output from the DB2 precompiler is converted to a usable control structure (which is called a package or an application plan). During the process, access paths to the data are selected and some authorization checking is performed.

**automatic bind.** (More correctly *automatic rebind*). A process by which SQL statements are bound automatically (without a user issuing a BIND command) when an application process begins execution and the bound application plan or package it requires is not valid.

**dynamic bind.** A process by which SQL statements are bound as they are entered.

**incremental bind.** A process by which SQL statements are bound during the execution of an application process, because they could not be bound during the bind process, and VALIDATE(RUN) was specified.

**static bind.** A process by which SQL statements are bound after they have been precompiled. All static SQL statements are prepared for execution at the same time.

**BLOB.** Binary large object.

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

### C

**CAF.** Call attachment facility.

**call attachment facility (CAF).** A DB2 attachment facility for application programs that run in TSO or MVS batch. The CAF is an alternative to the DSN command processor and provides greater control over the execution environment.

**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.

**cast function.** A function that is used to convert instances of a (source) data type into instances of a different (target) data type. In general, a cast function has the name of the target data type. It has one single argument whose type is the source data type; its return type is the target data type.

**catalog.** In DB2, a collection of tables that contains descriptions of objects such as tables, views, and indexes.

**catalog table.** Any table in the DB2 catalog.

**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.

**CLI.** Call level interface.

**client.** See *requester*.

**CLOB.** Character large object.

**column function.** An SQL operation that derives its result from a collection of values across one or more rows. Contrast with *scalar function*.

**commit.** The operation that ends a unit of work by releasing locks so that the database changes that are made by that unit of work can be perceived by other processes.

**connection handle.** The data object containing information that is associated with a connection that is managed by DB2 ODBC. This includes general status information, transaction status, and diagnostic information.

**constant.** A language element that specifies an unchanging value. Constants are classified as string constants or numeric constants. Contrast with *variable*.

**context.** The application's logical connection to the data source and associated internal DB2 ODBC connection information that allows the application to direct its operations to a data source. A DB2 ODBC context represents a DB2 thread.

**cursor.** A named control structure that an application program uses to point to a row of interest within some set of rows, and to retrieve rows from the set, possibly making updates or deletions.

### D

**database management system (DBMS).** A software system that controls the creation, organization, and modification of a database and the access to the data stored within it.

**data source.** A local or remote relational or non-relational data manager that is capable of supporting data access via an ODBC driver that supports the ODBC APIs. In the case of DB2 for OS/390, the data sources are always relational database managers.

**DBCLOB.** Double-byte character large object.

**DBMS.** Database management system.

**DB2 for VSE & VM.** The IBM DB2 relational database management system for the VSE and VM operating systems.

**DB2 thread.** The DB2 structure that describes an application's connection, traces its progress, processes resource functions, and delimits its accessibility to DB2 resources and services.

# **dimension.** A data category such as time, products, or markets. The elements of a dimension are referred to as members. Dimensions offer a very concise, intuitive way of organizing and selecting data for retrieval, exploration, and analysis. See also *dimension table*.

# **dimension table.** The representation of a dimension in a star schema. Each row in a dimension table represents all of the attributes for a particular member of the dimension. See also *dimension*, *star schema*, and *star join*.

**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.

**DRDA.** Distributed Relational Database Architecture.

**dynamic SQL.** SQL statements that are prepared and executed within an application program while the program is executing. In dynamic SQL, the SQL source is contained in host language variables rather than being coded into the application program. The SQL statement can change several times during the application program's execution.

**E**

**EBCDIC.** Extended binary coded decimal interchange code. An encoding scheme that is used to represent character data in the OS/390, MVS, VM, VSE, and OS/400® environments. Contrast with *ASCII*.

**embedded SQL.** SQL statements that are coded within an application program. See *static SQL*.

**enclave.** In Language Environment, an independent collection of routines, one of which is designated as the main routine. An enclave is similar to a program or run unit.

**environment.** A collection of names of logical and physical resources that are used to support the performance of a function.

**environment handle.** In DB2 ODBC, the data object that contains global information regarding the state of the application. An environment handle must be allocated before a connection handle can be allocated. Only one environment handle can be allocated per application.

**equi-join.** A join operation in which the join-condition has the form *expression = expression*.

**external function.** A function for which the body is written in a programming language that takes scalar argument values and produces a scalar result for each invocation. Contrast with *sourced function* and *built-in function*.

**F**

**foreign key.** A key that is specified in the definition of a referential constraint. Because of the foreign key, the table is a dependent table. The key must have the same number of columns, with the same descriptions, as the primary key of the parent table.

**full outer join.** The result of a join operation that includes the matched rows of both tables that are being joined and preserves the unmatched rows of both tables. See also *join*.

**function.** A specific purpose of an entity or its characteristic action such as a column function or scalar function. (See also *column function* and *scalar function*.)

Functions can be user-defined, built-in, or generated by DB2. (See *built-in function*, *cast function*, *external function*, *sourced function*, and *user-defined function*.)

**H**

**handle.** In DB2 ODBC, a variable that refers to a data structure and associated resources. See *statement handle*, *connection handle*, and *environment handle*.

**I**

**initialization file.** For DB2 ODBC applications, a file containing values that can be set to adjust the performance of the database manager.

**inner join.** The result of a join operation that includes only the matched rows of both tables being joined. See also *join*.

**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*.

**left outer join.** The result of a join operation that includes the matched rows of both tables that are being joined, and that preserves the unmatched rows of the first table. See also *join*.

**linkage editor.** A computer program for creating load modules from one or more object modules or load modules by resolving cross references among the modules and, if necessary, adjusting addresses.

**link-edit.** The action of creating a loadable computer program using a linkage editor.

## load module • Pthread

**load module.** A program unit that is suitable for loading into main storage for execution. The output of a linkage editor.

**LOB.** Large object.

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

**local.** A way of referring to any object that the local DB2 subsystem maintains. A *local table*, for example, is a table that is maintained by the local DB2 subsystem. Contrast with *remote*.

## M

**multithreading.** Multiple TCBs that are executing one copy of DB2 ODBC code concurrently (sharing a processor) or in parallel (on separate central processors).

**mutex.** Pthread mutual exclusion; a lock. A Pthread mutex variable is used as a locking mechanism to allow serialization of critical sections of code by temporarily blocking the execution of all but one thread.

**MVS/ESA.** Multiple Virtual Storage/Enterprise Systems Architecture.

## N

**NUL.** In C, a single character that denotes the end of the string.

**null.** A special value that indicates the absence of information.

**NUL-terminated host variable.** A varying-length host variable in which the end of the data is indicated by the presence of a NUL terminator.

**NUL terminator.** In C, the value that indicates the end of a string. For character strings, the NUL terminator is X'00'.

## O

**ODBC.** Open Database Connectivity.

**ODBC driver.** A dynamically-linked library (DLL) that implements ODBC function calls and interacts with a data source.

**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.

**OS/390.** Operating System/390.

**outer join.** The result of a join operation that includes the matched rows of both tables that are being joined and preserves some or all of the unmatched rows of the tables that are being joined. See also *join*.

## P

**plan.** See *application plan*.

**plan name.** The name of an application plan.

**POSIX.** Portable Operating System Interface. The IEEE operating system interface standard, which defines the Pthread standard of threading. See *Pthread*.

**precompilation.** A processing of application programs containing SQL statements that takes place before compilation. SQL statements are replaced with statements that are recognized by the host language compiler. Output from this precompilation includes source code that can be submitted to the compiler and the database request module (DBRM) that is input to the bind process.

**prepare.** The first phase of a two-phase commit process in which all participants are requested to prepare for commit.

**prepared SQL statement.** A named object that is the executable form of an SQL statement that has been processed by the PREPARE statement.

**primary key.** In a relational database, a unique, nonnull key that is part of the definition of a table. A table cannot be defined as a parent unless it has a unique key or primary key.

**Pthread.** The POSIX threading standard model for splitting an application into subtasks. The Pthread standard includes functions for creating threads, terminating threads, synchronizing threads through locking, and other thread control facilities.

## R

**RDBMS.** Relational database management system.

**reentrant.** Executable code that can reside in storage as one shared copy for all threads. Reentrant code is not self-modifying and provides separate storage areas for each thread. Reentrancy is a compiler and operating system concept, and reentrancy alone is not enough to guarantee logically consistent results when multithreading. See *threadsafe*.

**relational database management system (RDBMS).**

A collection of hardware and software that organizes and provides access to a relational database.

**remote.** Any object that is maintained by a remote DB2 subsystem (that is, by a DB2 subsystem other than the local one). A *remote view*, for example, is a view that is maintained by a remote DB2 subsystem. Contrast with *local*.

**requester.** The source of a request to a remote RDBMS, the system that requests the data. A requester is sometimes called an *application requester (AR)*.

**result set.** The set of rows that a stored procedure returns to a client application.

**result set locator.** A 4-byte value that DB2 uses to uniquely identify a query result set that a stored procedure returns.

**result table.** The set of rows that are specified by a SELECT statement.

**right outer join.** The result of a join operation that includes the matched rows of both tables that are being joined and preserves the unmatched rows of the second join operand. See also *join*.

**rollback.** The process of restoring data changed by SQL statements to the state at its last commit point. All locks are freed. Contrast with *commit*.

## S

**scalar function.** An SQL operation that produces a single value from another value and is expressed as a function name, followed by a list of arguments that are enclosed in parentheses. Contrast with *column function*.

**sourced function.** A function that is implemented by another built-in or user-defined function that is already known to the database manager. This function can be a scalar function or a column (aggregating) function; it returns a single value from a set of values (for example,

MAX or AVG). Contrast with *external function* and *built-in function*.

**source type.** An existing type that is used to internally represent a distinct type.

**SQL.** Structured Query Language.

**SQL authorization ID (SQL ID).** The authorization ID that is used for checking dynamic SQL statements in some situations.

**SQL communication area (SQLCA).** A structure that is used to provide an application program with information about the execution of its SQL statements.

**SQL descriptor area (SQLDA).** A structure that describes input variables, output variables, or the columns of a result table.

**SQLCA.** SQL communication area.

**SQLDA.** SQL descriptor area.

**SQL/DS.** Structured Query Language/Data System. This product is now obsolete and has been replaced by DB2 for VSE & VM.

# **star join.** A method of joining a dimension column of a fact table to the key column of the corresponding dimension table. See also *join*, *dimension*, and *star schema*.

# **star schema.** The combination of a fact table (which contains most of the data) and a number of dimension tables. See also *star join*, *dimension*, and *dimension table*.

**statement handle.** In DB2 ODBC, the data object that contains information about an SQL statement that is managed by DB2 ODBC. This includes information such as dynamic arguments, bindings for dynamic arguments and columns, cursor information, result values, and status information. Each statement handle is associated with the connection handle.

**static SQL.** SQL statements, embedded within a program, that are prepared during the program preparation process (before the program is executed). After being prepared, the SQL statement does not change (although values of host variables that are specified by the statement might change).

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

**Structured Query Language (SQL).** A standardized language for defining and manipulating data in a relational database.

## T

**table.** A named data object consisting of a specific number of columns and some number of unordered rows. See also *base table* or *temporary table*.

**task control block (TCB).** A control block that is used to communicate information about tasks within an address space that are connected to DB2. An address space can support many task connections (as many as one per task), but only one address space connection. See also *address space connection*.

**TCB.** Task control block (MVS).

**TCP/IP.** A network communication protocol that computer systems use to exchange information across telecommunication links.

# **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*.

**threadsafe.** A characteristic of code that allows multithreading both by providing private storage areas for each thread, and by properly serializing shared (global) storage areas.

**timestamp.** A seven-part value that consists of a date and time. The timestamp is expressed in years, months, days, hours, minutes, seconds, and microseconds.

**trace.** A DB2 facility that provides the ability to monitor and collect DB2 monitoring, auditing, performance, accounting, statistics, and serviceability (global) data.

## U

**UDF.** User-defined function.

**UDT.** User-defined data type. In DB2 for OS/390, the term *distinct type* is used instead of user-defined function.

**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*.

## V

**variable.** A data element that specifies a value that can be changed. A COBOL elementary data item is an example of a variable. Contrast with *constant*.

## X

**X/Open.** An independent, worldwide open systems organization that is supported by most of the world's largest information systems suppliers, user organizations, and software companies. X/Open's goal is to increase the portability of applications by combining existing and emerging standards.



---

# Bibliography

## DB2 Universal Database Server for OS/390 Version 6 Product Libraries:

### DB2 Universal Database for OS/390

- *DB2 Administration Guide*, SC26-9003
- *DB2 Application Programming and SQL Guide*, SC26-9004
- *DB2 Application Programming Guide and Reference for Java™*, SC26-9018
- *DB2 ODBC Guide and Reference*, SC26-9005
- *DB2 Command Reference*, SC26-9006
- *DB2 Data Sharing: Planning and Administration*, SC26-9007
- *DB2 Data Sharing Quick Reference Card*, SX26-3843
- *DB2 Diagnosis Guide and Reference*, LY36-3736
- *DB2 Diagnostic Quick Reference Card*, LY36-3737
- *DB2 Image, Audio, and Video Extenders Administration and Programming*, SC26-9650
- *DB2 Installation Guide*, GC26-9008
- *DB2 Licensed Program Specifications*, GC26-9009
- *DB2 Messages and Codes*, GC26-9011
- *DB2 Master Index*, SC26-9010
- *DB2 Reference for Remote DRDA Requesters and Servers*, SC26-9012
- *DB2 Reference Summary*, SX26-3844
- *DB2 Release Planning Guide*, SC26-9013
- *DB2 SQL Reference*, SC26-9014
- *DB2 Text Extender Administration and Programming*, SC26-9651
- *DB2 Utility Guide and Reference*, SC26-9015
- *DB2 What's New?* GC26-9017
- *DB2 Program Directory*, GI10-8182

### DB2 Administration Tool

- *DB2 Administration Tool for OS/390 User's Guide*, SC26-9847

### DB2 Buffer Pool Tool

- *DB2 Buffer Pool Tool for OS/390 User's Guide and Reference*, SC26-9306

### DB2 DataPropagator

- *DB2 Replication Guide and Reference*, SC26-9642

## Net.Data for OS/390

The following books are available at  
# <http://www.ibm.com/software/net.data/library.html>:

- *Net.Data Library: Administration and Programming Guide for OS/390*
- *Net.Data Library: Language Environment Interface Reference*
- *Net.Data Library: Messages and Codes*
- *Net.Data Library: Reference*

## DB2 PM for OS/390

- *DB2 PM for OS/390 Batch User's Guide*, SC26-9167
- *DB2 PM for OS/390 Command Reference*, SC26-9166
- *DB2 PM for OS/390 General Information*, GC26-9172
- *DB2 PM for OS/390 Installation and Customization*, SC26-9171
- *DB2 PM for OS/390 Messages*, SC26-9169
- *DB2 PM for OS/390 Online Monitor User's Guide*, SC26-9168
- *DB2 PM for OS/390 Report Reference Volume 1*, SC26-9164
- *DB2 PM for OS/390 Report Reference Volume 2*, SC26-9165
- *DB2 PM for OS/390 Using the Workstation Online Monitor*, SC26-9170
- *DB2 PM for OS/390 Program Directory*, GI10-8183

## Query Management Facility

- *Query Management Facility: Developing QMF Applications*, SC26-9579
- *Query Management Facility: Getting Started with QMF on Windows*, SC26-9582
- *Query Management Facility: High Performance Option User's Guide for OS/390*, SC26-9581
- *Query Management Facility: Installing and Managing QMF on OS/390*, GC26-9575
- *Query Management Facility: Installing and Managing QMF on Windows*, GC26-9583
- *Query Management Facility: Introducing QMF*, GC26-9576
- *Query Management Facility: Messages and Codes*, GC26-9580
- *Query Management Facility: Reference*, SC26-9577
- *Query Management Facility: Using QMF*, SC26-9578

## Ada/370

- *IBM Ada/370 Language Reference, SC09-1297*
- *IBM Ada/370 Programmer's Guide, SC09-1414*
- *IBM Ada/370 SQL Module Processor for DB2 Database Manager User's Guide, SC09-1450*

## APL2

- *APL2 Programming Guide, SH21-1072*
- *APL2 Programming: Language Reference, SH21-1061*
- *APL2 Programming: Using Structured Query Language (SQL), SH21-1057*

## AS/400

- *DB2 for OS/400 SQL Programming, SC41-4611*
- *DB2 for OS/400 SQL Reference, SC41-4612*

## BASIC

- *IBM BASIC/MVS Language Reference, GC26-4026*
- *IBM BASIC/MVS Programming Guide, SC26-4027*

## BookManager READ/MVS

- *BookManager READ/MVS V1R3: Installation Planning & Customization, SC38-2035*

## C/370

- *IBM SAA AD/Cycle C/370 Programming Guide, SC09-1841*
- *IBM SAA AD/Cycle C/370 Programming Guide for Language Environment/370, SC09-1840*
- *IBM SAA AD/Cycle C/370 User's Guide, SC09-1763*
- *SAA CPI C Reference, SC09-1308*

## Character Data Representation Architecture

- *Character Data Representation Architecture Overview, GC09-2207*
- *Character Data Representation Architecture Reference and Registry, SC09-2190*

## CICS/ESA

- *CICS/ESA Application Programming Guide, SC33-1169*
- *CICS for MVS/ESA Application Programming Reference, SC33-1170*
- *CICS for MVS/ESA CICS-RACF Security Guide, SC33-1185*
- *CICS for MVS/ESA CICS-Supplied Transactions, SC33-1168*
- *CICS for MVS/ESA Customization Guide, SC33-1165*
- *CICS for MVS/ESA Data Areas, LY33-6083*
- *CICS for MVS/ESA Installation Guide, SC33-1163*
- *CICS for MVS/ESA Intercommunication Guide, SC33-1181*

- *CICS for MVS/ESA Messages and Codes, GC33-1177*
- *CICS for MVS/ESA Operations and Utilities Guide, SC33-1167*
- *CICS/ESA Performance Guide, SC33-1183*
- *CICS/ESA Problem Determination Guide, SC33-1176*
- *CICS for MVS/ESA Resource Definition Guide, SC33-1166*
- *CICS for MVS/ESA System Definition Guide, SC33-1164*
- *CICS for MVS/ESA System Programming Reference, GC33-1171*

## CICS/MVS

- *CICS/MVS Application Programmer's Reference, SC33-0512*
- *CICS/MVS Facilities and Planning Guide, SC33-0504*
- *CICS/MVS Installation Guide, SC33-0506*
- *CICS/MVS Operations Guide, SC33-0510*
- *CICS/MVS Problem Determination Guide, SC33-0516*
- *CICS/MVS Resource Definition (Macro), SC33-0509*
- *CICS/MVS Resource Definition (Online), SC33-0508*

## IBM C/C++ for MVS/ESA

- *IBM C/C++ for MVS/ESA Library Reference, SC09-1995*
- *IBM C/C++ for MVS/ESA Programming Guide, SC09-1994*

## IBM COBOL

- *IBM COBOL Language Reference, SC26-4769*
- *IBM COBOL for MVS & VM Programming Guide, SC26-4767*

## Conversion Guide

- *IMS-DB and DB2 Migration and Coexistence Guide, GH21-1083*

## Cooperative Development Environment

- *CoOperative Development Environment/370: Debug Tool, SC09-1623*

## Data Extract (DXT)

- *Data Extract Version 2: General Information, GC26-4666*
- *Data Extract Version 2: Planning and Administration Guide, SC26-4631*

## DataPropagator NonRelational

- *DataPropagator NonRelational MVS/ESA Administration Guide, SH19-5036*
- *DataPropagator NonRelational MVS/ESA Reference, SH19-5039*

## Data Facility Data Set Services

- *Data Facility Data Set Services: User's Guide and Reference*, SC26-4388

## Database Design

- *DB2 Design and Development Guide*, Gabrielle Wiorkowski and David Kull, Addison Wesley, ISBN 0-20158-049-8
- *Handbook of Relational Database Design*, C. Fleming and B. Von Halle, Addison Wesley, ISBN 0-20111-434-8

## DataHub

- *IBM DataHub General Information*, GC26-4874

## DB2 Connect

- *DB2 Connect Enterprise Edition for OS/2 and Windows NT: Quick Beginnings*, GC09-2828
- *DB2 Connect Personal Edition Quick Beginnings*, GC09-2830
- *DB2 Connect User's Guide*, SC09-2838

## DB2 Server for VSE & VM

- *DB2 Server for VM: DBS Utility*, SC09-2394
- *DB2 Server for VSE: DBS Utility*, SC09-2395

## DB2 Universal Database (UDB)

- *DB2 UDB Administration Guide Volume 1: Design and Implementation*, SC09-2839
- *DB2 UDB Administration Guide Volume 2: Performance*, SC09-2840
- *DB2 UDB Administrative API Reference*, SC09-2841
- *DB2 UDB Application Building Guide*, SC09-2842
- *DB2 UDB Application Development Guide*, SC09-2845
- *DB2 UDB Call Level Interface Guide and Reference*, SC09-2843
- *DB2 UDB SQL Getting Started*, SC09-2856
- *DB2 UDB SQL Reference Volume 1*, SC09-2847
- *DB2 UDB SQL Reference Volume 2*, SC09-2848

## Device Support Facilities

- *Device Support Facilities User's Guide and Reference*, GC35-0033

## DFSMS/MVS

- *DFSMS/MVS: Access Method Services for the Integrated Catalog*, SC26-4906
- *DFSMS/MVS: Access Method Services for VSAM Catalogs*, SC26-4905
- *DFSMS/MVS: Administration Reference for DFSMSdss*, SC26-4929
- *DFSMS/MVS: DFSMSHsm Managing Your Own Data*, SH21-1077
- *DFSMS/MVS: Diagnosis Reference for DFSMSdftp*, LY27-9606

- *DFSMS/MVS Storage Management Library: Implementing System-Managed Storage*, SC26-3123
- *DFSMS/MVS: Macro Instructions for Data Sets*, SC26-4913
- *DFSMS/MVS: Managing Catalogs*, SC26-4914
- *DFSMS/MVS: Program Management*, SC26-4916
- *DFSMS/MVS: Storage Administration Reference for DFSMSdftp*, SC26-4920
- *DFSMS/MVS: Using Advanced Services*, SC26-4921
- *DFSMS/MVS: Utilities*, SC26-4926
- *MVS/DFP: Using Data Sets*, SC26-4749

## DFSORT

- *DFSORT Application Programming: Guide*, SC33-4035

## Distributed Relational Database

- *Data Stream and OPA Reference*, SC31-6806
- *IBM SQL Reference*, SC26-8416
- *Open Group Technical Standard (the Open Group presently makes the following books available through its Web site at <http://www.opengroup.org>):*
  - *DRDA Volume 1: Distributed Relational Database Architecture (DRDA)*, ISBN 1-85912-295-7
  - # – *DRDA Version 2 Volume 2: Formatted Data Object Content Architecture*, available only on Web
  - #
  - #
  - *DRDA Volume 3: Distributed Database Management (DDM) Architecture*, ISBN 1-85912-206-X

## Domain Name System

- *DNS and BIND, Third Edition*, Paul Albitz and Cricket Liu, O'Reilly, SR23-8771

## Education

- *IBM Dictionary of Computing*, McGraw-Hill, ISBN 0-07031-489-6
- *1999 IBM All-in-One Education and Training Catalog*, GR23-8105

## Enterprise System/9000 and Enterprise System/3090

- *Enterprise System/9000 and Enterprise System/3090 Processor Resource/System Manager Planning Guide*, GA22-7123

## High Level Assembler

- *High Level Assembler for MVS and VM and VSE Language Reference*, SC26-4940
- *High Level Assembler for MVS and VM and VSE Programmer's Guide*, SC26-4941

## Parallel Sysplex Library

- *OS/390 Parallel Sysplex Application Migration*, GC28-1863
- *System/390 MVS Sysplex Hardware and Software Migration*, GC28-1862
- *OS/390 Parallel Sysplex Overview: An Introduction to Data Sharing and Parallelism*, GC28-1860
- *OS/390 Parallel Sysplex Systems Management*, GC28-1861
- *OS/390 Parallel Sysplex Test Report*, GC28-1963
- *System/390 9672/9674 System Overview*, GA22-7148

## ICSF/MVS

- *ICSF/MVS General Information*, GC23-0093

## IMS/ESA

- *IMS Batch Terminal Simulator General Information*, GH20-5522
- *IMS/ESA Administration Guide: System*, SC26-8013
- *IMS/ESA Administration Guide: Transaction Manager*, SC26-8731
- *IMS/ESA Application Programming: Database Manager*, SC26-8727
- *IMS/ESA Application Programming: Design Guide*, SC26-8016
- *IMS/ESA Application Programming: Transaction Manager*, SC26-8729
- *IMS/ESA Customization Guide*, SC26-8020
- *IMS/ESA Installation Volume 1: Installation and Verification*, SC26-8023
- *IMS/ESA Installation Volume 2: System Definition and Tailoring*, SC26-8024
- *IMS/ESA Messages and Codes*, SC26-8028
- *IMS/ESA Operator's Reference*, SC26-8030
- *IMS/ESA Utilities Reference: System*, SC26-8035

## ISPF

- *ISPF V4 Dialog Developer's Guide and Reference*, SC34-4486
- *ISPF V4 Messages and Codes*, SC34-4450
- *ISPF V4 Planning and Customizing*, SC34-4443
- *ISPF V4 User's Guide*, SC34-4484

## Language Environment

- *Debug Tool User's Guide and Reference*, SC09-2137

## National Language Support

- *National Language Support Reference Volume 2*, SE09-8002

## NetView

- *NetView Installation and Administration Guide*, SC31-8043
- *NetView User's Guide*, SC31-8056

## ODBC

- *Microsoft ODBC 3.0 Programmer's Reference and SDK Guide*, Microsoft Press, ISBN 1-55615-658-8

## OS/390

- *OS/390 C/C++ Programming Guide*, SC09-2362
- *OS/390 C/C++ Run-Time Library Reference*, SC28-1663
- *OS/390 C/C++ User's Guide*, SC09-2361
- *OS/390 eNetwork Communications Server: IP Configuration*, SC31-8513
- *OS/390 Hardware Configuration Definition Planning*, GC28-1750
- *OS/390 Information Roadmap*, GC28-1727
- *OS/390 Introduction and Release Guide*, GC28-1725
- *OS/390 JES2 Initialization and Tuning Guide*, SC28-1791
- *OS/390 JES3 Initialization and Tuning Guide*, SC28-1802
- *OS/390 Language Environment for OS/390 & VM Concepts Guide*, GC28-1945
- *OS/390 Language Environment for OS/390 & VM Customization*, SC28-1941
- *OS/390 Language Environment for OS/390 & VM Debugging Guide*, SC28-1942
- *OS/390 Language Environment for OS/390 & VM Programming Guide*, SC28-1939
- *OS/390 Language Environment for OS/390 & VM Programming Reference*, SC28-1940
- *OS/390 MVS Diagnosis: Procedures*, LY28-1082
- *OS/390 MVS Diagnosis: Reference*, SY28-1084
- *OS/390 MVS Diagnosis: Tools and Service Aids*, LY28-1085
- *OS/390 MVS Initialization and Tuning Guide*, SC28-1751
- *OS/390 MVS Initialization and Tuning Reference*, SC28-1752
- *OS/390 MVS Installation Exits*, SC28-1753
- *OS/390 MVS JCL Reference*, GC28-1757
- *OS/390 MVS JCL User's Guide*, GC28-1758
- *OS/390 MVS Planning: Global Resource Serialization*, GC28-1759
- *OS/390 MVS Planning: Operations*, GC28-1760
- *OS/390 MVS Planning: Workload Management*, GC28-1761
- *OS/390 MVS Programming: Assembler Services Guide*, GC28-1762
- *OS/390 MVS Programming: Assembler Services Reference*, GC28-1910
- *OS/390 MVS Programming: Authorized Assembler Services Guide*, GC28-1763
- *OS/390 MVS Programming: Authorized Assembler Services Reference, Volumes 1-4*, GC28-1764, GC28-1765, GC28-1766, GC28-1767
- *OS/390 MVS Programming: Callable Services for High-Level Languages*, GC28-1768
- *OS/390 MVS Programming: Extended Addressability Guide*, GC28-1769

- *OS/390 MVS Programming: Sysplex Services Guide, GC28-1771*
- *OS/390 MVS Programming: Sysplex Services Reference, GC28-1772*
- *OS/390 MVS Programming: Workload Management Services, GC28-1773*
- *OS/390 MVS Routing and Descriptor Codes, GC28-1778*
- *OS/390 MVS Setting Up a Sysplex, GC28-1779*
- *OS/390 MVS System Codes, GC28-1780*
- *OS/390 MVS System Commands, GC28-1781*
- *OS/390 MVS System Messages Volume 1, GC28-1784*
- *OS/390 MVS System Messages Volume 2, GC28-1785*
- *OS/390 MVS System Messages Volume 3, GC28-1786*
- *OS/390 MVS System Messages Volume 4, GC28-1787*
- *OS/390 MVS System Messages Volume 5, GC28-1788*
- *OS/390 MVS Using the Subsystem Interface, SC28-1789*
- *OS/390 Security Server (RACF) Auditor's Guide, SC28-1916*
- *OS/390 Security Server (RACF) Command Language Reference, SC28-1919*
- *OS/390 Security Server (RACF) General User's Guide, SC28-1917*
- *OS/390 Security Server (RACF) Introduction, GC28-1912*
- *OS/390 Security Server (RACF) Macros and Interfaces, SK2T-6700 (OS/390 Collection Kit ), SK27-2180 (OS/390 Security Server Information Package )*
- *OS/390 Security Server (RACF) Security Administrator's Guide, SC28-1915*
- *OS/390 Security Server (RACF) System Programmer's Guide, SC28-1913*
- *OS/390 SMP/E Reference, SC28-1806*
- *OS/390 SMP/E User's Guide, SC28-1740*
- *OS/390 RMF User's Guide, SC28-1949*
- *OS/390 TSO/E CLISTS, SC28-1973*
- *OS/390 TSO/E Command Reference, SC28-1969*
- *OS/390 TSO/E Customization, SC28-1965*
- *OS/390 TSO/E Messages, GC28-1978*
- *OS/390 TSO/E Programming Guide, SC28-1970*
- *OS/390 TSO/E Programming Services, SC28-1971*
- *OS/390 TSO/E User's Guide, SC28-1968*
- *OS/390 DCE Administration Guide, SC28-1584*
- *OS/390 DCE Introduction, GC28-1581*
- *OS/390 DCE Messages and Codes, SC28-1591*
- *OS/390 UNIX System Services Command Reference, SC28-1892*
- *OS/390 UNIX System Services Planning, SC28-1890*
- *OS/390 UNIX System Services User's Guide, SC28-1891*
- *OS/390 UNIX System Services Programming: Assembler Callable Services Reference, SC28-1899*

j

### **PL/I for MVS & VM**

- *IBM PL/I MVS & VM Language Reference, SC26-3114*
- *IBM PL/I MVS & VM Programming Guide, SC26-3113*

### **OS PL/I**

- *OS PL/I Programming Language Reference, SC26-4308*
- *OS PL/I Programming Guide, SC26-4307*

### **Prolog**

- *IBM SAA AD/Cycle Prolog/MVS & VM Programmer's Guide, SH19-6892*

### **Remote Recovery Data Facility**

- *Remote Recovery Data Facility Program Description and Operations, LY37-3710*

### **Storage Management**

- *DFSMS/MVS Storage Management Library: Implementing System-Managed Storage, SC26-3123*
- *MVS/ESA Storage Management Library: Leading a Storage Administration Group, SC26-3126*
- *MVS/ESA Storage Management Library: Managing Data, SC26-3124*
- *MVS/ESA Storage Management Library: Managing Storage Groups, SC26-3125*
- *MVS Storage Management Library: Storage Management Subsystem Migration Planning Guide, SC26-4659*

### **System/370 and System/390**

- *ESA/370 Principles of Operation, SA22-7200*
- *ESA/390 Principles of Operation, SA22-7201*
- *System/390 MVS Sysplex Hardware and Software Migration, GC28-1210*

### **System Network Architecture (SNA)**

- *SNA Formats, GA27-3136*
- *SNA LU 6.2 Peer Protocols Reference, SC31-6808*
- *SNA Transaction Programmer's Reference Manual for LU Type 6.2, GC30-3084*
- *SNA/Management Services Alert Implementation Guide, GC31-6809*

### **TCP/IP**

- *IBM TCP/IP for MVS: Customization & Administration Guide, SC31-7134*
- *IBM TCP/IP for MVS: Diagnosis Guide, LY43-0105*
- *IBM TCP/IP for MVS: Messages and Codes, SC31-7132*

- *IBM TCP/IP for MVS: Planning and Migration Guide, SC31-7189*

#### **VS COBOL II**

- *VS COBOL II Application Programming Guide for MVS and CMS, SC26-4045*
- *VS COBOL II Application Programming: Language Reference, GC26-4047*
- *VS COBOL II Installation and Customization for MVS, SC26-4048*

#### **VS FORTRAN**

- *VS FORTRAN Version 2: Language and Library Reference, SC26-4221*

- *VS FORTRAN Version 2: Programming Guide for CMS and MVS, SC26-4222*

#### **VTAM**

- *Planning for NetView, NCP, and VTAM, SC31-8063*
- *VTAM for MVS/ESA Diagnosis, LY43-0069*
- *VTAM for MVS/ESA Messages and Codes, SC31-6546*
- *VTAM for MVS/ESA Network Implementation Guide, SC31-6548*
- *VTAM for MVS/ESA Operation, SC31-6549*
- *VTAM for MVS/ESA Programming, SC31-6550*
- *VTAM for MVS/ESA Programming for LU 6.2, SC31-6551*
- *VTAM for MVS/ESA Resource Definition Reference, SC31-6552*

---

# Index

## Special Characters

— 370  
% 370

## A

abends 423  
allocate functions  
    AllocConnect 79  
    AllocEnv 83  
    AllocStmt 85  
application  
    compile 57  
    execute 59  
    execution steps 57  
    link-edit 58  
    multithreaded 394  
    pre-link 58  
    preparation 56  
    requirements 57  
    sample, DSNTJ8 57  
    tasks 21  
    trace 413  
application variables  
    binding 29  
APPLTRACE keyword  
    description 63  
    use of 415  
APPLTRACEFILENAME keyword 63  
array input 375  
array output 379  
ASCII scalar function 430  
attributes  
    connection 361  
    environment 361  
    querying and setting 361  
    statement 361  
AUTOCOMMIT keyword 63

## B

BINARY  
    conversion to C 453  
BindCol, function 87  
binding  
    application variables  
        columns 31  
        parameter markers 29  
    packages 53  
    plan 54  
    sample, DSNTIJCL 52, 54  
    stored procedures 54

BindParameter, function 93  
BITDATA keyword 64

## C

caching  
    dynamic SQL statement 410  
Cancel, function 104  
case sensitivity 44  
catalog  
    functions  
        description 369  
        limiting use of 410  
        querying 369  
CD-ROM, books on 4  
CHAR  
    conversion to C 450  
    display size 448  
    length 447  
    precision 444  
    scale 446  
character strings 43, 44  
CLISCHEMA keyword 64  
    use of 371  
ColAttributes, function 106  
COLLECTIONID keyword 64  
column-wise binding 379, 380  
ColumnPrivileges, function 112  
Columns, function 117  
commit 32  
common server 14  
compile, application 57  
CONCAT scalar function 430  
configuring DB2 ODBC 49  
CONNECT 142  
    *See also* SQLDriverConnect  
    type 1, type 2 24  
Connect, function 123  
connection attributes (options)  
    description 361  
    setting 408  
connection handle 15  
    allocating 23  
    AllocConnect, function 79  
    Free, function 184  
    freeing 23  
connection string 362  
connectivity  
    ODBC model 23  
    requirements 51  
CONNECTTYPE keyword 64

- contexts
  - multiple 396
- coordinated distributed transactions 363
- coordinated transactions, establishing 366
- core level functions 13
- CURDATE scalar function 431
- CURRENTSERVER, plan bind option 55
- CURRENTSQLID keyword 65
- cursor
  - definition 32
  - use in DB2 ODBC 15
- CURSORHOLD keyword 66
- CURTIME scalar function 431

## D

- data conversion
  - C data types 38
  - C to SQL data types 459
  - data types 36
  - default data types 38
  - description 41
  - display size of SQL data types 448
  - length of SQL data types 447
  - precision of SQL data types 444
  - scale of SQL data types 446
  - SQL data types 38
  - SQL to C data types 449
- data source information, querying 45
- data types
  - C 38, 41
  - generic 41
  - ODBC 41
  - SQL 38
- data-at-execute 373
- DATABASE scalar function 432
- DataSources, function 129
- DATE
  - conversion to C 454
  - display size 448
  - length 447
  - precision 444
  - scale 446
- DAYOFMONTH scalar function 431
- DB2 books online 4
- DB2 DataPropagator, managing shadow catalog 372
- DB2 ODBC
  - advantages of using 18
  - application requirements 56
  - components 56
  - configuring 49
  - diagnostic trace 415
  - initialization file 60
  - installing 49
  - migration considerations 72
  - shadow catalog 371

- DB2 ODBC (*continued*)
  - stored procedures 389
- DBNAME keyword 66
- DBRMs
  - binding to packages 52
- debugging 423
- DECIMAL
  - conversion to C 452
  - display size 448
  - length 447
  - precision 444
  - scale 446
- deferred arguments 29
- DescribeCol, function 132
- diagnosis
  - description 34
  - trace 415
- DIAGTRACE keyword
  - description 66
  - use of 415
- DIAGTRACE\_BUFFER\_SIZE keyword 67
- DIAGTRACE\_NO\_WRAP keyword 67
- Disconnect, function 140
- DISCONNECT, plan bind option 55
- display size of SQL data types 448
- distributed transactions 363
- DOUBLE
  - conversion to C 452
  - display size 448
  - length 447
  - precision 444
  - scale 446
- driver
  - DB2 ODBC 425
  - ODBC 425
- driver manager 425
- DriverConnect, function 142
- DSN8OIVP, JCL sample 57
- DSNAOINI
  - DD card 60
- DSNTEJ8, application sample 57
- DSNTIJCL, bind sample 52, 54
- DYNAMICRULES, package bind option 53

## E

- embedded SQL
  - comparison to DB2 ODBC 15
  - mixing with DB2 ODBC 401
- environment
  - attributes (options) 361
  - information, querying 45
  - OS/390 UNIX, setup 55
  - runtime 49
- environment handle
  - allocating 22



- environment handle (*continued*)
  - AllocEnv, function 83
  - description 15
  - Free, function 186
  - freeing 22
- environmental variables
  - OS/390 UNIX 60
- error code, internal 424
- Error, function 148
- escape clauses, vendor 404
- example
  - stored procedure 469
- examples
  - array INSERT 378
  - catalog functions 370
- ExecDirect, function 154
- execute direct 28
- execute statement 28
- execute, application 57, 59
- Execute, function 159
- export statements
  - OS/390 UNIX 60
- ExtendedFetch, function 162

## F

- FAR pointers 73
- Fetch, function 169
- fetching data in pieces 374
- FLOAT
  - conversion to C 452
  - display size 448
  - length 447
  - precision 444
  - scale 446
- ForeignKeys, function 175
- free functions
  - FreeConnect 184
  - FreeEnv 186
  - FreeStmt 188
- function list, ODBC 74
- functions
  - ODBC list 74

## G

- GetConnectOption, function 191
- GetCursorName, function 193
- GetData, function 199
- GetFunctions, function 214
- GetInfo, function 219
- GetNumResultCols, function 271
- GetSQLCA, function 242
- GetStmtOption, function 249
- GetTypeInfo, function 254

- GRAPHIC
  - conversion to C 451
  - keyword 67

## H

- handles
  - connection handle 15, 22
  - environment handle 15, 22
  - statement handle 15
- HOURLY scalar function 431

## I

- IFNULL scalar function 432
- initialization
  - file
    - common errors 62, 407
    - defaults, changing 362
    - description 60
    - specifying 60
  - task 21
- INSERT scalar function 430
- installation 49
- INTEGER
  - conversion to C 452
  - display size 448
  - length 447
  - precision 444
  - scale 446
- internal error code 424
- introduction to DB2 ODBC 13
- INVALID\_HANDLE 35
- isolation levels, ODBC 428
- ISOLATION, package bind option 53

## K

- keywords, initialization 61, 63

## L

- LE threads
  - multiple 394
- LEFT scalar function 430
- length of SQL data types 447
- LENGTH scalar function 430
- library
  - online 4
- link-edit, application 58
- long data
  - retrieving in pieces 373
  - sending in pieces 373
- LONGVARBINARY
  - conversion to C 453
- LONGVARCHAR
  - conversion to C 450

## LONGVARCHAR (continued)

- display size 448
- length 447
- precision 444
- scale 446

## LONGVARGRAPHIC

- conversion to C 451

## M

- MAXCONN keyword 68
- metadata characters 370
- migration considerations 72
- MINUTE scalar function 431
- MONTH scalar function 231
- MoreResults, function 262
- MULTICONTEXT keyword 68
- multiple contexts 396
- multithreaded applications 394
- MVSATTACHTYPE keyword 69
- MVSDEFAULTSSID keyword 69

## N

- native error code 36
- NativeSQL, function 266
- notices, legal 489
- NOW scalar function 431
- null connect 391
- null-terminated strings 43
- NUMERIC
  - conversion to C 452
  - display size 448
  - length 447
  - precision 444
  - scale 446
- NumParams, function 269
- NumResultCols, function 271

## O

- ODBC
  - and DB2 ODBC 13, 425
  - connectivity 23
  - core level functions 13
  - function list 74
  - isolation levels 428
  - vendor escape clauses 404
- online books 4
- OPTIMIZEFORNROWS keyword 69
- options
  - connection 361
  - environment 361
  - querying and setting 361
  - statement 361

## OS/390 UNIX

- compile application 58
- environment setup 55
- environmental variable 60
- execute application 59
- export statements 60
- pre-link, link-edit application 59
- special considerations 57

## P

- packages, binding 52
- ParamData, function 273
- parameter markers
  - array input 375
  - binding 29
  - use of 15
- ParamOptions, function 275
- pattern-values 370
- plan, binding 54
- PLANNAME keyword 69
- pointers, FAR 73
- portability 18
- pre-link, application 58
- precision of SQL data types 444
- Prepare, function 277
- prepare, statement 28
- PrimaryKey, function 285
- ProcedureCols, function 290
- Procedures, function 300
- PutData, function 304

## Q

- query statements, processing 30
- querying
  - data source information 45
  - environment information 45
  - shadow catalog 371
  - system catalog information 369

## R

- REAL
  - conversion to C 452
  - display size 448
  - length 447
  - precision 444
  - scale 446
- registering stored procedures 390
- remote site
  - binding packages 54
- REPEAT scalar function 430
- result sets
  - function generated 410
  - retrieving into array 379

- result sets (*continued*)
  - returning from stored procedures 392
- retrieving multiple rows 379
- return codes 35
- RIGHT scalar function 430
- rollback 32
- row-wise binding 379, 380
- RowCount, function 307
- ROWID
  - conversion to C 457
  - display size 448
  - length 447
  - precision 444
  - scale 446
- rowset 162
- runtime environment
  - setting up 52
  - support 49

## S

- samples
  - DSN8OIVP 57
  - DSNTEJ8 57
  - DSNTIJCL 52, 54
- scalar functions 407
- scale of SQL data types 446
- SCHEMALIST keyword 69
- search arguments 370
- SECOND scalar function 431
- SELECT 30
- SetColAttributes, function 309
- SetConnectOption, function 316
- SetCursorName, function 322
- SetParam, function 328
- SetStmtOption, function 334
- shadow catalog
  - managing 372
  - querying 371
- SMALLINT
  - conversion to C 452
  - display size 448
  - length 447
  - precision 444
  - scale 446
- SMP/E jobs 49
- softcopy publications 4
- SpecialColumns, function 339
- SQL
  - dynamically prepared 16
  - parameter markers 29
  - preparing and executing statements 28
  - query statements 30
  - SELECT 30
  - statements
    - DELETE 32
    - UPDATE 32

- SQL (*continued*)
  - VALUES 30
- SQL Access Group 13
- SQL\_ACCESS\_MODE 316
- SQL\_ATTR\_OUTPUT\_NTS 325
- SQL\_AUTOCOMMIT 316, 408
- SQL\_BIND\_TYPE 335
- SQL\_C\_BINARY
  - conversion from SQL 462
- SQL\_C\_BIT
  - conversion from SQL 461
- SQL\_C\_CHAR
  - conversion from SQL 460
- SQL\_C\_DATE
  - conversion from SQL 464
- SQL\_C\_DBCHAR
  - conversion from SQL 463
- SQL\_C\_DOUBLE
  - conversion from SQL 461
- SQL\_C\_FLOAT
  - conversion from SQL 461
- SQL\_C\_LONG
  - conversion from SQL 461
- SQL\_C\_SHORT
  - conversion from SQL 461
- SQL\_C\_TIME
  - conversion from SQL 465
- SQL\_C\_TIMESTAMP
  - conversion from SQL 466
- SQL\_C\_TINYINT
  - conversion from SQL 461
- SQL\_CLOSE\_BEHAVIOR 335
- SQL\_COLUMN\_AUTO\_INCREMENT 107
- SQL\_COLUMN\_CASE\_SENSITIVE 107
- SQL\_COLUMN\_CATALOG\_NAME 107
- SQL\_COLUMN\_COUNT 107
- SQL\_COLUMN\_DISPLAY\_SIZE 107
- SQL\_COLUMN\_DISTINCT\_TYPE 107
- SQL\_COLUMN\_LABEL 107
- SQL\_COLUMN\_LENGTH 107
- SQL\_COLUMN\_MONEY 107
- SQL\_COLUMN\_NAME 107
- SQL\_COLUMN\_NULLABLE 107
- SQL\_COLUMN\_OWNER\_NAME 107
- SQL\_COLUMN\_PRECISION 107
- SQL\_COLUMN\_QUALIFIER\_NAME 107
- SQL\_COLUMN\_SCALE 107
- SQL\_COLUMN\_SCHEMA\_NAME 107
- SQL\_COLUMN\_SEARCHABLE 108
- SQL\_COLUMN\_TABLE\_NAME 108
- SQL\_COLUMN\_TYPE 108
- SQL\_COLUMN\_TYPE\_NAME 108
- SQL\_COLUMN\_UNSIGNED 108
- SQL\_COLUMN\_UPDATABLE 108
- SQL\_CONCURRENCY 335

SQL\_CONCURRENT\_TRANS 364  
 SQL\_CONNECTTYPE 316, 325, 364  
 SQL\_COORDINATED\_TRANS 364  
 SQL\_CURRENT\_SCHEMA 316  
 SQL\_CURSOR\_HOLD 335, 408  
 SQL\_CURSOR\_TYPE 335  
 SQL\_DATA\_AT\_EXEC 373  
 SQL\_ERROR 35  
 SQL\_MAX\_LENGTH 335  
 SQL\_MAX\_ROWS 335, 408  
 SQL\_MAXCONN 316, 325  
 SQL\_NEED\_DATA 35  
 SQL\_NO\_DATA\_FOUND 35  
 SQL\_NODESCRIBE 335  
 SQL\_NOSCAN 335, 411  
 SQL\_NTS 43  
 SQL\_PARAMOPT\_ATOMIC 316  
 SQL\_RETRIEVE\_DATA 335  
 SQL\_ROWSET\_SIZE 335, 380  
 SQL\_STMTTXN\_ISOLATION 335, 409  
 SQL\_SUCCESS 35  
 SQL\_SUCCESS\_WITH\_INFO 35  
 SQL\_TXN\_ISOLATION 316, 335, 408  
 SQLAllocConnect, function  
     description 79  
 SQLAllocEnv, function  
     description 83  
 SQLAllocStmt, function  
     description 85  
     overview 27  
 SQLBindCol, function  
     description 87  
     overview 27, 31  
 SQLBindParameter, function  
     description 93  
     overview 31  
 SQLCancel, function  
     description 104  
     use in data-at-execute 374  
 SQLColAttributes, function  
     description 106  
     overview 27, 31  
 SQLColumnPrivileges, function  
     description 112  
 SQLColumns, function  
     description 117  
 SQLConnect, function  
     description 123  
 SQLDataSources, function  
     description 129  
     overview 27  
 SQLDescribeCol, function  
     description 132  
     overview 27, 31  
 SQLDescribeParam, function  
     description 138  
 SQLDisconnect, function  
     description 140  
 SQLDriverConnect, function  
     description 142  
 SQLDriverConnect() 362  
 SQLError, function  
     description 148  
 SQLError, package bind option 53  
 SQLExecDirect, function  
     description 154  
     overview 27, 28  
 SQLExecute, function  
     description 159  
     overview 27, 28  
 SQLExtendedFetch, function  
     description 162  
 SQLFetch, function  
     description 169  
     overview 27, 31  
 SQLForeignKeys, function  
     description 175  
 SQLFreeConnect, function  
     description 184  
 SQLFreeEnv, function  
     description 186  
 SQLFreeStmt, function  
     description 188  
     overview 27  
 SQLGetConnectOption, function  
     description 191  
 SQLGetCursorName, function  
     description 193  
 SQLGetData, function  
     description 199  
     overview 27, 31  
 SQLGetEnvAttr, function  
     description 212  
 SQLGetFunctions, function  
     description 214  
 SQLGetInfo, function  
     description 219  
 SQLGetLength, function  
     description 235  
 SQLGetPosition, function  
     description 238  
 SQLGetSQLCA, function  
     description 242  
 SQLGetStmtOption, function  
     description 249  
 SQLGetSubString, function  
     description 251  
 SQLGetTypeInfo, function  
     description 254  
 SQLMoreResults, function  
     description 262  
     use of 376

- SQLNativeSql, function
  - description 266
- SQLNumParams, function
  - description 269
- SQLNumResultCols, function
  - description 271
  - overview 27, 31
- SQLParamData, function
  - description 273
  - use in data-at-execute 373
- SQLParamOptions, function
  - description 275
- SQLPrepare, function
  - description 277
  - overview 27, 28
- SQLPrimaryKeys, function
  - description 285
- SQLProcedureColumns, function
  - description 290
- SQLProcedures, function
  - description 300
- SQLPutData, function
  - description 304
  - use in data-at-execute 373
- SQLRowCount, function
  - description 307
  - overview 27
- SQLSetColAttributes, function
  - description 309
- SQLSetConnection, function
  - description 314
- SQLSetConnectOption, function
  - description 316
- SQLSetCursorName, function
  - description 322
- SQLSetEnvAttr, function
  - description 325
- SQLSetParam, function
  - description 328
  - overview 27, 28, 31
- SQLSetStmtOption, function
  - description 334
- SQLSpecialColumns, function
  - description 339
- SQLSTATE
  - description 35
  - format of 35
  - function cross reference 433
  - in DB2 ODBC 15
- SQLStatistics, function
  - description 344
- SQLTablePrivileges, function
  - description 349
- SQLTables, function
  - description 353
- SQLTransact, function
  - description 357
  - overview 27, 31, 32
- statement attributes (options)
  - description 361
  - setting 408
- statement handle 15
  - allocating 28
  - Free, function 188
  - freeing 34
  - maximum number of 28
- statement handle functions
  - AllocStmt 85
- Statistics, function 344
- stored procedures
  - binding 54
  - catalog table
    - registering 390
  - example 469
  - ODBC escape clause 406
  - returning result sets 392
  - tracing 419
  - using with DB2 ODBC 389
- string
  - arguments 43, 44
  - null-termination 43
  - truncation 44
- SUBSTRING scalar function 430
- subsystem, defining 59
- SYSSHEMA keyword 70
- system catalogs, querying 369

**T**

- Table, function 353
- TablePrivileges, function 349
- TABLETYPE keyword 71
- termination task 21
- threads
  - multiple 394
- THREADSAFE keyword
  - description 71
- TIME
  - conversion to C 455
  - display size 448
  - length 447
  - precision 444
  - scale 446
- TIMESTAMP
  - conversion to C 456
  - display size 448
  - length 447
  - precision 444
  - scale 446
- trace
  - application 413

trace (*continued*)  
  DB2 ODBC diagnosis 415  
  keywords, using 415  
  stored procedure 419  
Transact, function 357  
transaction  
  isolation levels, ODBC 428  
  management 32  
  processing 21  
truncation 44  
TXNISOLATION keyword 71

## U

UNDERSCORE keyword 72  
USER scalar function 432

## V

VALUES 30  
VARBINARY  
  conversion to C 453  
VARCHAR  
  conversion to C 450  
  display size 448  
  length 447  
  precision 444  
  scale 446  
VARGRAPHIC  
  conversion to C 451  
vendor escape clauses 404

## W

writing DB2 ODBC applications 21

## X

X/Open CAE 36  
X/Open Company 13  
X/Open SQL CLI 13



---

## How to send your comments

DB2 Universal Database for OS/390  
ODBC Guide and Reference  
Version 6

Publication No. SC26-9005-01

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

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

- Send your comments from the Web. Visit the DB2 for OS/390 Web site at:

**<http://www.ibm.com/software/db2os390>**

The Web site has a feedback page that you can use to send comments.

- Complete the readers' comment form at the back of the book and return it by mail, by fax (800-426-7773 for the United States and Canada), or by giving it to an IBM representative.



---

# Readers' Comments

**DB2 Universal Database for OS/390  
ODBC Guide and Reference  
Version 6**

**Publication No. SC26-9005-01**

How satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Technically accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Grammatically correct and consistent	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Graphically well designed	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

May we contact you to discuss your comments?  Yes  No

\_\_\_\_\_  
Name

\_\_\_\_\_  
Address

\_\_\_\_\_  
Company or Organization


\_\_\_\_\_  
Phone No.




Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE  
NECESSARY  
IF MAILED IN THE  
UNITED STATES



# BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation  
Department BWE/H3  
PO Box 49023  
San Jose, CA 95161-9945



Fold and Tape

Please do not staple

Fold and Tape





Program Number: 5645-DB2



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

SC26-9005-01

