

ADABAS D

ODBC DRIVER



Manual Order Number: ESD612-041WOU

This document is applicable to ADABAS D Version 6.1.1 PE and to all subsequent releases, unless otherwise indicated in new editions or technical newsletters.

Specifications contained herein are subject to change and these changes will be reported in subsequent revisions or editions.

Readers' comments are welcomed. Comments may be addressed to the Documentation Department at the address on the back cover.

Reprinted with permission of Microsoft Corporation.

© 14.04.2025, Microsoft Corporation

All rights reserved

Printed in the Federal Republic of Germany

The SOFTWARE AG documentation often refers to numerous hardware and software products by their trade names. In most, if not all cases, these designations are claimed as trademarks or registered trademarks by their respective companies.

About This Manual

The Microsoft® Open Database Connectivity™ (ODBC) interface is a C programming language interface for database connectivity. ADABAS D provides an ODBC-compliant CALL interface (API) on all ADABAS D platforms, that is Unix, Windows NT, Windows 3.1 and OS/2. This manual addresses the following questions:

- What features does ODBC offer?
- How do applications use the interface?

Organization of this Manual

This manual is a partial reprint of :

Microsoft ODBC 2.0 Programmer's Reference and SDK Guide,
Microsoft Press
ISBN 1-55615-658-8

Parts which are not relevant for the ADABAS CALL interface are omitted in this reprint, which is organized in the following way :

- Part 1 *Introduction to ODBC, (omitted)*
- Part 2 *Developing Applications*, containing information for developing applications using the ODBC interface;
- Part 3 *Developing Drivers, (omitted)*
- Part 4 *Installing and Configuring ODBC Software, (omitted)*
- Part 5 *API Reference*, containing syntax and semantic information for all ODBC functions.

Document Conventions

This manual uses the following typographic conventions.

Format	Used for
WIN.INI	Uppercase letters indicate filenames, SQL statements, macro names, and terms used at the operating-system command level.
RETCODE SQLFetch(hdbc)	This font is used for sample command lines and program code.
<i>argument</i>	Italicized words indicate information that the user or the application must provide, or word emphasis.
SQLTransact	Bold type indicates that syntax must be typed exactly as shown, including function names.
[]	Brackets indicate optional items; if in bold text, brackets must be included in the syntax.
	A vertical bar separates two mutually exclusive choices in a syntax line.
{ }	Braces delimit a set of mutually exclusive choices in a syntax line; if in bold text, braces must be included in the syntax.
...	An ellipsis indicates that arguments can be repeated several times.
. . . .	A column of three dots indicates continuation of previous lines of code.

P A R T 1

Introduction to ODBC

C H A P T E R S 1 - 2

- omitted -

PART 2

Developing Applications

Guidelines for Calling ODBC Functions

This chapter describes the general characteristics of ODBC functions, determining driver conformance levels, the role of the Driver Manager, ODBC function arguments, and the values ODBC functions return.

General Information

Each ODBC function name starts with the prefix "SQL." Each function accepts one or more arguments. Arguments are defined as input (to the driver) or output (from the driver).

C programs that call ODBC functions must include the `sql.h`, `sql.h`, and `WINDOWS.H` header files. These files define Environment and ODBC constants and types and provide function prototypes for all ODBC functions.

Note: For portability reasons, the file `WINDOWS.H` is distributed on all non-MS-WINDOWS platforms (equivalent to `MS windows.h`). It only contains a platform-specific subset of the required prototypes and constants.

Determining Driver Conformance Levels

ODBC defines conformance levels for drivers in two areas: the ODBC API and the ODBC SQL grammar (which includes the ODBC SQL data types). These levels establish standard sets of functionality. By inquiring the conformance levels supported by a driver, an application can easily determine if the driver provides the necessary functionality.

Note: The following sections refer to **SQLGetInfo** and **SQLGetTypeInfo**, which are part of the Level 1 API conformance level. Although it is strongly recommended that drivers support this conformance level, drivers are not required to do so. If these functions are not supported, an application developer must consult the driver documentation to determine its conformance levels.

Determining API Conformance Levels

ODBC functions are divided into core functions, which are defined in the X/Open and SQL Access Group Call Level Interface specification, and two levels of extension functions, with which ODBC extends this specification. To determine the function conformance level of a driver, an application calls **SQLGetInfo** with the `SQL_ODBC_SAG_CLI_CONFORMANCE` and `SQL_ODBC_API_CONFORMANCE` flags. Note that a driver can support one or more extension functions but not conform to ODBC extension Level 1 or 2. To determine if a driver supports a particular function, an application calls **SQLGetFunctions**. Note that **SQLGetFunctions** is implemented by the Driver Manager and can be called for any driver, regardless of its level.

Determining SQL Conformance Levels

The ODBC SQL grammar, which includes SQL data types, is divided into a minimum grammar, a core grammar, which corresponds to the X/Open and SQL Access Group SQL CAE specification (1992), and an extended grammar, which provides common extensions to SQL. To determine the SQL conformance level of a driver, an application calls **SQLGetInfo** with the `SQL_ODBC_SQL_CONFORMANCE` flag. To determine whether a driver supports a specific SQL extension, an application calls **SQLGetInfo** with a flag for that extension. For more information, see Appendix C, "SQL Grammar." To determine whether a driver supports a specific SQL data type, an application calls **SQLGetTypeInfo**.

Note: See the ADABAS README file on the specific platform which API and SQL conformance levels are supported by ADABAS ODBC driver.

Using the Driver Manager (Unix, OS/2)

On non-WINDOWS platforms, a Driver Manager for ODBC is not needed.

Chapter 3 Guidelines for Calling ODBC Functions

The ODBC driver is provided as a static library directly linked with the application.

Using the Driver Manager (Windows)

The Driver Manager is a DLL that provides access to ODBC drivers. An application typically links with the Driver Manager import library (ODBC.LIB) to gain access to the Driver Manager.

Whenever an application calls an ODBC function, the Driver Manager performs one of the following actions:

- For **SQLDataSources** and **SQLDrivers**, the Driver Manager processes the call. It does not pass the call to the driver.
- For **SQLGetFunctions**, the Driver Manager passes the call to the driver associated with the connection. If the driver does not support **SQLGetFunctions**, the Driver Manager processes the call.
- For **SQLAllocEnv**, **SQLAllocConnect**, **SQLSetConnectOption**, **SQLFreeConnect**, and **SQLFreeEnv**, the Driver Manager processes the call. The Driver Manager calls **SQLAllocEnv**, **SQLAllocConnect**, and **SQLSetConnectOption** in the driver when the application calls a function to connect to the data source (**SQLConnect**, **SQLDriverConnect**, or **SQLBrowseConnect**). The Driver Manager calls **SQLFreeConnect** and **SQLFreeEnv** in the driver when the application calls **SQLDisconnect**.
- For **SQLConnect**, **SQLDriverConnect**, **SQLBrowseConnect**, and **SQLError**, the Driver Manager performs initial processing then passes the call to the driver associated with the connection.
- For any other ODBC function, the Driver Manager passes the call to the driver associated with the connection.

If requested, the Driver Manager records each called function in a trace file. The name of each function is recorded, along with the values of the input arguments and the names of the output arguments (as listed in the function definitions).

Calling ODBC Functions

The following paragraphs describe general characteristics of ODBC functions.

Buffers

An application passes data to a driver in an input buffer. The driver returns data to an application in an output buffer. The application must allocate memory for both input and output buffers. (If the application will use the buffer to retrieve string data, the buffer must contain space for the null termination byte.)

Note that some functions accept pointers to buffers that are later used by other functions. The application must ensure that these pointers remain valid until all applicable functions have used them. For example, the argument *rgbValue* in **SQLBindCol** points to an output buffer in which **SQLFetch** returns the data for a column.

Input Buffers

An application passes the address and length of an input buffer to a driver. The length of the buffer must be one of the following values:

- A length greater than or equal to zero. This is the actual length of the data in the input buffer. For character data, a length of zero indicates that the data is an empty (zero length) string. Note that this is different from a null pointer. If the application specifies the length of character data, the character data does not need to be null-terminated.
- **SQL_NTS**. This specifies that a character data value is null-terminated.
- **SQL_NULL_DATA**. This tells the driver to ignore the value in the input buffer and use a NULL data value instead. It is only valid when the input buffer is used to provide the value of a parameter in an SQL statement.

The operation of ODBC functions on character data containing embedded null characters is undefined, and is not recommended for maximum interoperability.

Unless it is specifically prohibited in a function description, the address of an input buffer may be a null pointer. When the address of an input buffer is a null pointer, the value of the corresponding buffer length argument is ignored.

For more information about input buffers, see "Converting Data from C to SQL Data Types" in Appendix D, "Data Types."

Output Buffers

An application passes the following arguments to a driver, so that it can return data in an output buffer:

- The address of the buffer in which the driver returns the data (the output buffer). Unless it is specifically prohibited in a function description, the address of an output buffer can be a null pointer. In this case, the driver does not return anything in the buffer and, in the absence of other errors, returns **SQL_SUCCESS**.
If necessary, the driver converts data before returning it. The driver always null-terminates character data before returning it.
- The length of the buffer. This is ignored by the driver if the returned data has a fixed length in C, such as an integer, real number, or date structure.

- The address of a variable in which the driver returns the length of the data (the length buffer). The returned length of the data is `SQL_NULL_DATA` if the data is a NULL value in a result set. Otherwise, it is the number of bytes of data available to return. If the driver converts the data, it is the number of bytes after the conversion. For character data, it does not include the null termination byte added by the driver.

If the output buffer is too small, the driver attempts to truncate the data. If the truncation does not cause a loss of significant data, the driver returns the truncated data in the output buffer, returns the length of the available data (as opposed to the length of the truncated data) in the length buffer, and returns `SQL_SUCCESS_WITH_INFO`. If the truncation causes a loss of significant data, the driver leaves the output and length buffers untouched and returns `SQL_ERROR`. The application calls **SQLError** to retrieve information about the truncation or the error.

For more information about output buffers, see "Converting Data from SQL to C Data Types" in Appendix D, "Data Types."

Environment, Connection, and Statement Handles

When so requested by an application, the ODBC Driver allocates storage for information about the ODBC environment, each connection, and each SQL statement. The handles to these storage areas are returned to the application. The application then uses one or more of them in each call to an ODBC function.

The ODBC interface defines three types of handles:

- The **environment handle** identifies memory storage for global information, including the valid connection handles and the current active connection handle. ODBC defines the environment handle as a variable of type `HENV`. An application uses a single environment handle; it must request this handle prior to connecting to a data source.
- **Connection handles** identify memory storage for information about a particular connection. ODBC defines connection handles as variables of type `HDBC`. An application must request a connection handle prior to connecting to a data source. Each connection handle is associated with the environment handle. The environment handle can, however, have multiple connection handles associated with it.
- **Statement handles** identify memory storage for information about an SQL statement. ODBC defines statement handles as variables of type `HSTMT`. An application must request a statement handle prior to submitting SQL requests. Each statement handle is associated with exactly one connection

handle. Each connection handle can, however, have multiple statement handles associated with it.

For more information about requesting a connection handle, see Chapter 5, "Connecting to the SERVERDB." For more information about requesting a statement handle, see Chapter 6, "Executing SQL Statements."

Using Data Types

Data stored on a data source has an SQL data type, which may be specific to that data source. A driver maps data source-specific SQL data types to ODBC SQL data types, which are defined in the ODBC SQL grammar, and driver-specific SQL data types. (A driver returns these mappings through **SQLGetTypeInfo**. It also uses the ODBC SQL data types to describe the data types of columns and parameters in **SQLColAttributes**, **SQLDescribeCol**, and **SQLDescribeParam**.)

Each SQL data type corresponds to an ODBC C data type. By default, the driver assumes that the C data type of a storage location corresponds to the SQL data type of the column or parameter to which the location is bound. If the C data type of a storage location is not the *default* C data type, the application can specify the correct C data type with the *fCType* argument in **SQLBindCol**, **SQLGetData**, or **SQLBindParameter**. Before returning data from the data source, the driver converts it to the specified C data type. Before sending data to the data source, the driver converts it from the specified C data type.

For more information about data types, see Appendix D, "Data Types". The C data types are defined in SQL.H and SQLEXT.H.

ODBC Function Return Codes

When an application calls an ODBC function, the driver executes the function and returns a predefined code. These return codes indicate success, warning, or failure status. The return codes are:

SQL_SUCCESS	SQL_INVALID_HANDLE
SQL_SUCCESS_WITH_INFO	SQL_STILL_EXECUTING
SQL_NO_DATA_FOUND	SQL_NEED_DATA
SQL_ERROR	

Basic Application Steps

To interact with a data source, a simple application:

1. Connects to the data source. It specifies the data source name and any additional information needed to complete the connection.
2. Processes one or more SQL statements:
 - The application places the SQL text string in a buffer. If the statement includes parameter markers, it sets the parameter values.
 - If the statement returns a result set, the application assigns a cursor name for the statement or allows the driver to do so.
 - The application submits the statement for prepared or immediate execution.
 - If the statement creates a result set, the application can inquire about the attributes of the result set, such as the number of columns and the name and type of a specific column. It assigns storage for each column in the result set and fetches the results.
 - If the statement causes an error, the application retrieves error information from the driver and takes appropriate action.
3. Ends each transaction by committing it or rolling it back.
4. Terminates the connection when it has finished interacting with the data source.

The following diagram lists the ODBC function calls that an application makes to connect to a data source, process SQL statements, and disconnect from the data source. Depending on its needs, an application may call other ODBC functions.

Connecting to the SERVERDB

About Data Sources (Windows)

A data source consists of the data a user wants to access, its associated DBMS, the platform on which the DBMS resides, and the network (if any) used to access that platform. Each data source requires that a driver provide certain information in order to connect to it. At the core level, this is defined to be the name of the data source, a user ID, and a password. ODBC extensions allow drivers to specify additional information, such as a network address or additional passwords.

The connection information for each data source is stored in the ODBC.INI file or registry, which is created during installation and maintained with an administration program. A section in this file lists the available data sources. Additional sections describe each data source in detail, specifying the driver name, a description, and any additional information the driver needs to connect to the data source.

For example :

```
[ODBC Data Sources]
sqlberlin:DB611=ADABAS D
[sqlberlin:DB611]
Driver=%DBROOT%\BIN\sqlodbc.dll
```

Initializing the ODBC Environment

Before an application can use any other ODBC function, it must initialize the ODBC interface and associate an environment handle with the environment. To initialize the interface and allocate an environment handle, an application:

1. Declares a variable of type HENV. For example, the application could use the declaration:
`HENV henv1;`
2. Calls **SQLAllocEnv** and passes it the address of the variable. The driver initializes the ODBC environment, allocates memory to store information about the environment, and returns the environment handle in the variable.

These steps should be performed only once by an application; **SQLAllocEnv** supports one or more connections to data sources.

Allocating a Connection Handle

Before an application can connect to a driver, it must allocate a connection handle for the connection. To allocate a connection handle, an application:

1. Declares a variable of type HDBC. For example, the application could use the declaration:
`HDBC hdbc1;`
2. Calls **SQLAllocConnect** and passes it the address of the variable. The driver allocates memory to store information about the connection and returns the connection handle in the variable.

Connecting to the Database (Unix, OS/2)

This chapter briefly describes how to connect to the SERVERDB. The ODBC driver can administer up to eight concurrent database sessions. To open a database session, various parameters are required which are detailed in the following.

1. SERVERNODE: The name of the machine where the database is installed. Usually, this name must be available in the file '/etc/hosts'. (See: User Manual for UNIX, ...).
2. SERVERDB-Name: The name of the database installed on the SERVERNODE.
3. User-Name: Name of an existing database user.
4. Password: Password of the user.

Connecting to a database session is done by calling the function `SQLConnect`. The parameter `szDSN` must be set to the SERVERNODE and SERVERDB parameters. SERVERNODE and SERVERDB must be separated from each other by a colon.

```
SQLConnect(hdbc, "sqlberlin:DB611", SQL_NTS, ...);
```

Note: Any parameters specified here are case sensitive.

Before the `SQLConnect` function can be called, the ODBC Environment must be initialized. How this is done, is described in the next section.

Connecting to the Database (Windows)

Next, the application specifies a specific driver and data source. It passes the following information to the driver in a call to **SQLConnect**:

- **Data source name** The name of the data source being requested by the application.
- **User ID** The login ID or account name for access to the data source, if appropriate (optional).
- **Authentication string (password)** A character string associated with the user ID that allows access to the data source (optional).

When an application calls **SQLConnect**, the Driver Manager uses the data source name to read the name of the driver DLL from the appropriate section of the ODBC.INI file or registry. It then loads the driver DLL and passes the **SQLConnect** arguments to it. If the driver needs additional information to connect to the data source, it reads this information from the same section of the ODBC.INI file.

If the application specifies a data source name that is not in the ODBC.INI file or registry, or if the application does not specify a data source name, the Driver Manager searches for the default data source specification. If it finds the default data source, it loads the default driver DLL and passes the application-specified data source name to it. If there is no default data source, the Driver Manager returns an error.

Additional Functions

ODBC also provides the following functions related to connections, drivers, and data sources. For more information about these functions, see Chapter 22, "ODBC Function Reference."

Function	Description
SQLGetFunctions	Retrieves functions supported by a driver. This function allows an application to determine at run time whether a particular function is supported by a driver.
SQLGetInfo	Retrieves general information about a driver and data source, including filenames, versions, conformance levels, and capabilities.
SQLGetTypeInfo	Retrieves the SQL data types supported by a driver and data source.
SQLSetConnectOptio	These functions set or retrieve connection options, such

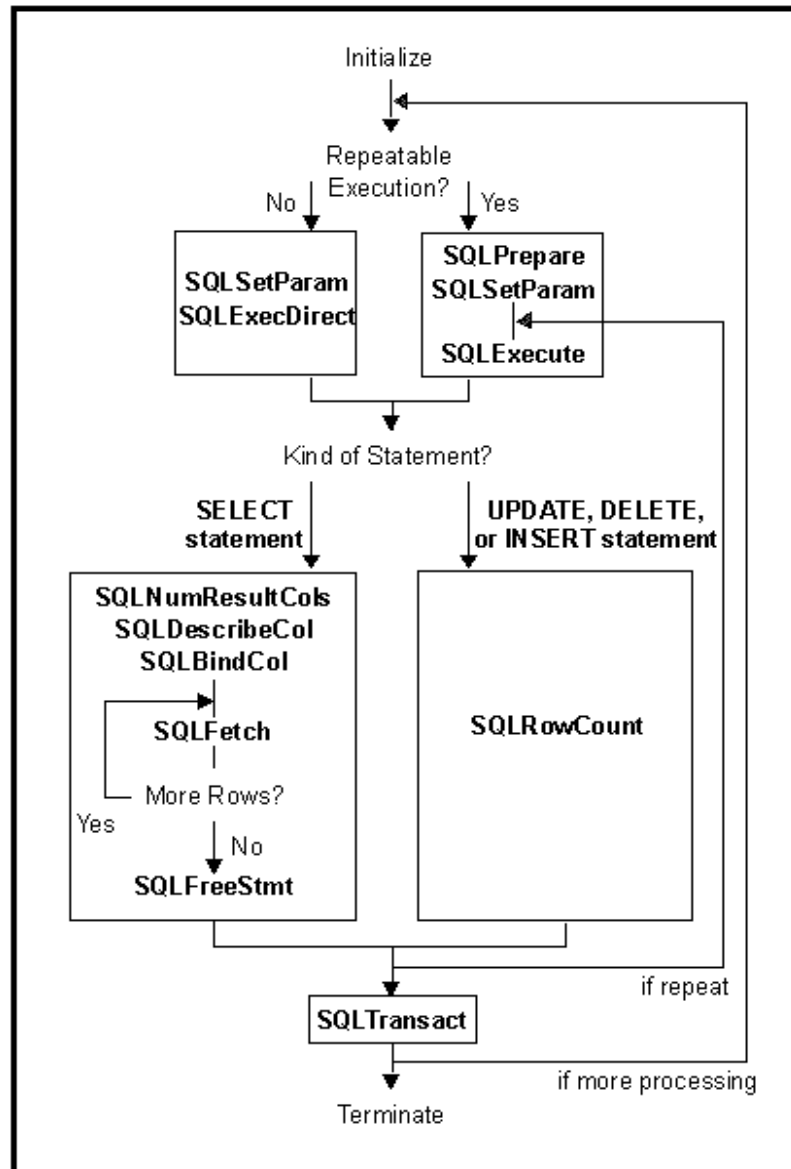
n as the data source access mode, automatic transaction
SQLGetConnectOptio commitment, timeout values, function tracing, data
n translation options, and transaction isolation.

Executing SQL Statements

An application can submit any SQL statement supported by a data source. ODBC defines a standard syntax for SQL statements (listed in Appendix C, "SQL Grammar"). For maximum interoperability, an application should only submit SQL statements that use this syntax; the driver will translate these statements to the syntax used by the data source. If an application submits an SQL statement that does not use the ODBC syntax, the driver passes it directly to the data source.

Note: For **CREATE TABLE** and **ALTER TABLE** statements, applications should use the data type name returned by **SQLGetTypeInfo** in the **TYPE_NAME** column, rather than the data type name defined in the SQL grammar.

The following diagram shows a simple sequence of ODBC function calls to execute SQL statements. Note that statements can be executed a single time with **SQLExecDirect** or prepared with **SQLPrepare** and executed multiple times with **SQLExecute**. Note also that an application calls **SQLTransact** to commit or roll back a transaction.



Allocating a Statement Handle

Before an application can submit an SQL statement, it must allocate a statement handle for the statement. To allocate a statement handle, an application:

1. Declares a variable of type HSTMT. For example, the application could use the declaration:

```
HSTMT hstmt1;
```

2. Calls **SQLAllocStmt** and passes it the address of the variable and the connected *hdbc* with which to associate the statement. The driver allocates memory to store information about the statement, associates the statement handle with the *hdbc*, and returns the statement handle in the variable.

Executing an SQL Statement

An application can submit an SQL statement for execution in two ways:

- **Prepared** Call **SQLPrepare** and then call **SQLExecute**.
- **Direct** Call **SQLExecDirect**.

These options are similar, though not identical to, the prepared and immediate options in embedded SQL. For a comparison of the ODBC functions and embedded SQL, see Appendix E, "Comparison Between Embedded SQL and ODBC."

Prepared Execution

An application should prepare a statement before executing it if either of the following is true:

- The application will execute the statement more than once, possibly with intermediate changes to parameter values.
- The application needs information about the result set prior to execution.

A prepared statement executes faster than an unprepared statement because the data source compiles the statement, produces an access plan, and returns an access plan identifier to the driver. The data source minimizes processing time as it does not have to produce an access plan each time it executes the statement. Network traffic is minimized because the driver sends the access plan identifier to the data source instead of the entire statement.

Important: Committing or rolling back a transaction, either by calling **SQLTransact** or by using the `SQL_AUTOCOMMIT` connection option, can cause the data source to delete the access plans for all *hstmts* on an *hdbc*. For more information, see the `SQL_CURSOR_COMMIT_BEHAVIOR` and `SQL_CURSOR_ROLLBACK_BEHAVIOR` information types in **SQLGetInfo**.

Part 2 Developing Applications

To prepare and execute an SQL statement, an application:

1. Calls **SQLPrepare** to prepare the statement.
2. Sets the values of any statement parameters. For more information, see "Setting Parameter Values" later in this chapter.
3. Retrieves information about the result set, if necessary. For more information, see "Determining the Characteristics of a Result Set" in Chapter 7, "Retrieving Results."
4. Calls **SQLExecute** to execute the statement.
5. Repeats steps 2 through 4 as necessary.

Direct Execution

An application should execute a statement directly if both of the following are true:

- The application will execute the statement only once.
- The application does not need information about the result set prior to execution.

To execute an SQL statement directly, an application:

1. Sets the values of any statement parameters. For more information, see "Setting Parameter Values" later in this chapter.
2. Calls **SQLExecDirect** to execute the statement.

Setting Parameter Values

An SQL statement can contain parameter markers that indicate values that the driver retrieves from the application at execution time. For example, an application might use the following statement to insert a row of data into the EMPLOYEE table:

```
INSERT INTO EMPLOYEE (NAME, AGE, HIREDATE) VALUES (?, ?, ?)
```

An application uses parameter markers instead of literal values if:

- It needs to execute the same prepared statement several times with different parameter values.
- The parameter values are not known when the statement is prepared.
- The parameter values need to be converted from one data type to another.

To set a parameter value, an application performs the following steps in any order:

- Calls **SQLBindParameter** to bind a storage location to a parameter marker and specify the data types of the storage location and the column associated with the parameter, as well as the precision and scale of the parameter.
- Places the parameter's value in the storage location.

These steps can be performed before or after a statement is prepared, but must be performed before a statement is executed.

Parameter values must be placed in storage locations in the C data types specified in **SQLBindParameter**. For example:

Parameter Value	SQL Data Type	C Data Type	Stored Value
ABC	SQL_CHAR	SQL_C_CHAR	ABC\0 ^a
10	SQL_INTEGER	SQL_C_SLONG	10
10	SQL_INTEGER	SQL_C_CHAR	10\0 ^a
1 P.M.	SQL_TIME	SQL_C_TIME	13,0,0 ^b
1 P.M.	SQL_TIME	SQL_C_CHAR	{t '13:00:00'}\0 ^{a,c}

^a The numbers in this list are the numbers stored in the fields of the TIME_STRUCT structure.

^b The string uses the ODBC date escape clause. For more information, see "Date, Time, and Timestamp Data" later in this chapter.

Storage locations remain bound to parameter markers until the application calls **SQLFreeStmt** with the SQL_RESET_PARAMS option or the SQL_DROP option. An application can bind a different storage area to a parameter marker at any time by calling **SQLBindParameter**. An application can also change the value in a storage location at any time. When a statement is executed, the driver uses the current values in the most recently defined storage locations.

Performing Transactions

In *auto-commit* mode, every SQL statement is a complete transaction, which is automatically committed. In *manual-commit* mode, a transaction consists of one or more statements. In manual-commit mode, when an application submits an SQL statement and no transaction is open, the driver implicitly begins a transaction. The transaction remains open until the application commits or rolls back the transaction with **SQLTransact**.

If a driver supports the `SQL_AUTOCOMMIT` connection option, the default transaction mode is auto-commit; otherwise, it is manual-commit. An application calls **SQLSetConnectOption** to switch between manual-commit and auto-commit mode. Note that if an application switches from manual-commit to auto-commit mode, the driver commits any open transactions on the connection.

Applications should call **SQLTransact**, rather than submitting a **COMMIT** or **ROLLBACK** statement, to commit or roll back a transaction. The result of a **COMMIT** or **ROLLBACK** statement depends on the driver and its associated data source.

Important: Committing or rolling back a transaction, either by calling **SQLTransact** or by using the `SQL_AUTOCOMMIT` connection option, can cause the data source to close the cursors and delete the access plans for all *hstmts* on an *hdbc*. For more information, see the `SQL_CURSOR_COMMIT_BEHAVIOR` and `SQL_CURSOR_ROLLBACK_BEHAVIOR` information types in **SQLGetInfo**.

ODBC Extensions for SQL Statements

ODBC extends the X/Open and SQL Access Group Call Level Interface to provide additional functions related to SQL statements. ODBC also extends the X/Open and SQL Access Group SQL CAE specification (1992) to provide common extensions to SQL. The remainder of this chapter describes these functions and SQL extensions.

To determine if a driver supports a specific function, an application calls **SQLGetFunctions**. To determine if a driver supports a specific ODBC extension to SQL, such as outer joins or procedure invocation, an application calls **SQLGetInfo**.

Retrieving Information About the Data Source's Catalog

The following functions, known as catalog functions, return information about a data source's catalog:

- **SQLTables** returns the names of tables stored in a data source.
- **SQLTablePrivileges** returns the privileges associated with one or more tables.
- **SQLColumns** returns the names of columns in one or more tables.

- **SQLColumnPrivileges** returns the privileges associated with each column in a single table.
- **SQLPrimaryKeys** returns the names of columns that comprise the primary key of a single table.
- **SQLForeignKeys** returns the names of columns in a single table that are foreign keys. It also returns the names of columns in other tables that refer to the primary key of the specified table.
- **SQLSpecialColumns** returns information about the optimal set of columns that uniquely identify a row in a single table or the columns in that table that are automatically updated when any value in the row is updated by a transaction.
- **SQLStatistics** returns statistics about a single table and the indexes associated with that table.
- **SQLProcedures** returns the names of procedures stored in a data source.
- **SQLProcedureColumns** returns a list of the input and output parameters, as well as the names of columns in the result set, for one or more procedures.

Each function returns the information as a result set. An application retrieves these results by calling **SQLBindCol** and **SQLFetch**.

Sending Parameter Data at Execution Time

To send parameter data at statement execution time, such as for parameters of the SQL_LONGVARCHAR or SQL_LONGVARBINARY types, an application uses the following three functions:

- **SQLBindParameter**
- **SQLParamData**
- **SQLPutData**

To indicate that it plans to send parameter data at statement execution time, an application calls **SQLBindParameter** and sets the *pcbValue* buffer for the parameter to the result of the SQL_LEN_DATA_AT_EXEC(*length*) macro. If the *fSqlType* argument is SQL_LONGVARBINARY or SQL_LONGVARCHAR and the driver returns "Y" for the SQL_NEED_LONG_DATA_LEN information type in **SQLGetInfo**, *length* is the total number of bytes of data to be sent for the parameter; otherwise, it is ignored.

The application sets the *rgbValue* argument to a value that, at run time, can be used to retrieve the data. For example, *rgbValue* might point to a storage

location that will contain the data at statement execution time or to a file that contains the data. The driver returns the value to the application at statement execution time.

When the driver processes a call to **SQLExecute** or **SQLExecDirect** and the statement being executed includes a data-at-execution parameter, the driver returns **SQL_NEED_DATA**. To send the parameter data, the application:

1. Calls **SQLParamData**, which returns *rgbValue* (as set with **SQLBindParameter**) for the first data-at-execution parameter.
2. Calls **SQLPutData** one or more times to send data for the parameter. (More than one call will be needed if the data value is larger than the buffer; multiple calls are allowed only if the C data type is character or binary and the SQL data type is character, binary, or data source–specific.)
3. Calls **SQLParamData** again to indicate that all data has been sent for the parameter. If there is another data-at-execution parameter, the driver returns *rgbValue* for that parameter and **SQL_NEED_DATA** for the function return code. Otherwise, it returns **SQL_SUCCESS** for the function return code.
4. Repeats steps 2 and 3 for the remaining data-at-execution parameters.

For additional information, see the description of **SQLBindParameter** in Chapter 22, "ODBC Function Reference."

Specifying Arrays of Parameter Values

To specify multiple sets of parameter values for a single SQL statement, an application calls **SQLParamOptions**. For example, if there are ten sets of column values to insert into a table—and the same SQL statement can be used for all ten operations—the application can set up an array of values, then submit a single **INSERT** statement.

If an application uses **SQLParamOptions**, it must allocate enough memory to handle the arrays of values.

Executing Functions Asynchronously

By default, a driver executes ODBC functions synchronously; the driver does not return control to an application until a function call completes. If a driver supports asynchronous execution, however, an application can request asynchronous execution for the functions listed below. (All of these functions either submit requests to a data source or retrieve data. These operations may require extensive processing.)

SQLColAttributes	SQLForeignKeys	SQLProcedureColumns
SQLColumnPrivileges	SQLGetData	SQLProcedures
SQLColumns	SQLGetTypeInfo	SQLPutData
SQLDescribeCol	SQLMoreResults	SQLSetPos
SQLDescribeParam	SQLNumParams	SQLSpecialColumns
SQLExecDirect	SQLNumResultCols	SQLStatistics
SQLExecute	SQLParamData	SQLTablePrivileges
SQLExtendedFetch	SQLPrepare	SQLTables
SQLFetch	SQLPrimaryKeys	

Asynchronous execution is performed on a statement-by-statement basis. To execute a statement asynchronously, an application:

1. Calls **SQLSetStmtOption** with the SQL_ASYNC_ENABLE option to enable asynchronous execution for an *hstmt*. (To enable asynchronous execution for all *hstmts* associated with an *hdbc*, an application calls **SQLSetConnectOption** with the SQL_ASYNC_ENABLE option.)
2. Calls one of the functions listed earlier in this section and passes it the *hstmt*. The driver begins asynchronous execution of the function and returns SQL_STILL_EXECUTING.

Note: If the application calls a function that cannot be executed asynchronously, the driver executes the function synchronously.

3. Performs other operations while the function is executing asynchronously. The application can call any function with a different *hstmt* or an *hdbc* not associated with the original *hstmt*. With the original *hstmt* and the *hdbc* associated with that *hstmt*, the application can only call the original function, **SQLAllocStmt**, **SQLCancel**, or **SQLGetFunctions**.
4. Calls the asynchronously executing function to check if it has finished. While the arguments must be valid, the driver ignores all of them except the *hstmt* argument. For example, suppose an application called **SQLExecDirect** to execute a **SELECT** statement asynchronously. When the application calls **SQLExecDirect** again, the return value indicates the status of the **SELECT** statement, even if the *szSqlStr* argument contains an **INSERT** statement.

If the function is still executing, the driver returns SQL_STILL_EXECUTING and the application must repeat steps 3 and 4. If the function has finished, the driver returns a different code, such as SQL_SUCCESS or SQL_ERROR. For information about canceling a function executing

asynchronously, see "Terminating Statement Processing" in Chapter 9, "Terminating Transactions and Connections."

5. Repeats steps 2 through 4 as needed.

To disable asynchronous execution for an *hstmt*, an application calls **SQLSetStmtOption** with the SQL_ASYNC_ENABLE option. To disable asynchronous execution for all *hstmts* associated with an *hdbc*, an application calls **SQLSetConnectOption** with the SQL_ASYNC_ENABLE option.

Using ODBC Extensions to SQL

ODBC defines the following extensions to SQL, which are common to most DBMS's:

- Date, time, and timestamp data
- Scalar functions such as numeric, string, and data type conversion functions
- **LIKE** predicate escape characters
- Outer joins
- Procedures

The syntax defined by ODBC for these extensions uses the escape clause provided by the X/Open and SQL Access Group SQL CAE specification (1992) to cover vendor-specific extensions to SQL. Its format is:

--(vendor*(*vendor-name*), *product*(*product-name*) extension *)--**

For the ODBC extensions to SQL, *product-name* is always "ODBC", since the product defining them is ODBC. *Vendor-name* is always "Microsoft", since ODBC is a Microsoft product. ODBC also defines a shorthand syntax for these extensions:

{*extension*}

Most DBMS's provide the same extensions to SQL as does ODBC. Because of this, an application may be able to submit an SQL statement using one of these extensions in either of two ways:

- Use the syntax defined by ODBC. An application that uses the ODBC syntax will be interoperable among DBMS's.
- Use the syntax defined by the DBMS. An application that uses DBMS-specific syntax will not be interoperable among DBMS's.

Due to the difficulty in implementing some ODBC extensions to SQL, such as outer joins, a driver might only implement those ODBC extensions that are supported by its associated DBMS. To determine whether the driver and data source support all the ODBC extensions to SQL, an application calls **SQLGetInfo** with the SQL_ODBC_SQL_CONFORMANCE flag. For information about how an application determines whether a specific extension is supported, see the section that describes the extension.

Note: Many DBMS's provide extensions to SQL other than those defined by ODBC. To use one of these extensions, an application uses the DBMS-specific syntax. The application will not be interoperable among DBMS's.

Date, Time, and Timestamp Data

The escape clauses ODBC uses 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' *)--
```

where **d** indicates *value* is a date in the "yyyy-mm-dd" format, **t** indicates *value* is a time in the "hh:mm:ss" format, and **ts** indicates *value* is a timestamp in the "yyyy-mm-dd hh:mm:ss[.f...]" format. The shorthand syntax for date, time, and timestamp data is:

```
{d 'value'}  
{t 'value'}  
{ts 'value'}
```

For example, each of the following statements updates the birthday of John Smith in the EMPLOYEE table. The first statement uses the escape clause syntax. The second statement uses the shorthand syntax. The third statement uses the native syntax for a DATE column in DEC's Rdb and is not interoperable among DBMS's.

```
UPDATE EMPLOYEE  
SET BIRTHDAY=--(*vendor(Microsoft),product(ODBC) d '1967-01-15' *)--  
WHERE NAME='Smith, John'
```

```
UPDATE EMPLOYEE  
SET BIRTHDAY={d '1967-01-15'}  
WHERE NAME='Smith, John'
```

```
UPDATE EMPLOYEE  
SET BIRTHDAY='15-Jan-1967'  
WHERE NAME='Smith, John'
```

The ODBC escape clauses for date, time, and timestamp literals can be used in parameters with a C data type of SQL_C_CHAR. For example, the following statement uses a parameter to update the birthday of John Smith in the EMPLOYEE table:

```
UPDATE EMPLOYEE SET BIRTHDAY=? WHERE NAME='Smith, John'
```

A storage location of type SQL_C_CHAR bound to the parameter might contain any of the following values. The first value uses the escape clause syntax. The second value uses the shorthand syntax. The third value uses the native syntax for a DATE column in DEC's Rdb and is not interoperable among DBMS's.

```
--(*vendor(Microsoft),product(ODBC) d '1967-01-15' *)--  
"{d '1967-01-15'"
```

```
"15-Jan-1967"
```

An application can also send date, time, or timestamp values as parameters using the C structures defined by the C data types SQL_C_DATE, SQL_C_TIME, and SQL_C_TIMESTAMP.

To determine if a data source supports date, time, or timestamp data, an application calls **SQLGetTypeInfo**. If a driver supports date, time, or timestamp data, it must also support the escape clauses for date, time, or timestamp literals.

Scalar Functions

Scalar functions—such as string length, absolute value, or current date—can be used on columns of a result set and on columns that restrict rows of a result set. The escape clause ODBC uses for scalar functions is:

```
--(*vendor(Microsoft),product(ODBC) fn scalar-function *)--
```

where *scalar-function* is one of the functions listed in Appendix F, "Scalar Functions." The shorthand syntax for scalar functions is:

```
{fn scalar-function}
```

For example, each of the following statements creates the same result set of uppercase employee names. The first statement uses the escape clause syntax. The second statement uses the shorthand syntax. The third statement uses the native syntax for Ingres™ for OS/2 and is not interoperable among DBMS's.

```
SELECT --(*vendor(Microsoft),product(ODBC) fn UCASE(NAME) *)--  
FROM EMPLOYEE
```

```
SELECT {fn UCASE(NAME)} FROM EMPLOYEE
```

```
SELECT uppercase(NAME) FROM EMPLOYEE
```

An application can mix scalar functions that use native syntax and scalar functions that use ODBC syntax. For example, the following statement creates a result set of last names of employees in the EMPLOYEE table. (Names in the EMPLOYEE table are stored as a last name, a comma, and a first name.) The statement uses the ODBC scalar function **SUBSTRING** and the SQL Server scalar function **CHARINDEX** and will only execute correctly on SQL Server.

```
SELECT {fn SUBSTRING(NAME, 1, CHARINDEX(',', NAME) - 1)} FROM EMPLOYEE
```

To determine which scalar functions are supported by a data source, an application calls **SQLGetInfo** with the SQL_NUMERIC_FUNCTIONS, SQL_STRING_FUNCTIONS, SQL_SYSTEM_FUNCTIONS, and SQL_TIMEDATE_FUNCTIONS flags.

Data Type Conversion Function

ODBC defines a special scalar function, **CONVERT**, that requests that the data source convert data from one SQL data type to another SQL data type. The escape clause ODBC uses for the **CONVERT** function is:

```
--(*vendor(Microsoft),product(ODBC)  
  fn CONVERT(value_exp, data_type) *)--
```

where *value_exp* is a column name, the result of another scalar function, or a literal value, and *data_type* is a keyword that matches the **#define** name used by an ODBC SQL data type (as defined in Appendix D, "Data Types"). The shorthand syntax for the **CONVERT** function is:

```
{fn CONVERT(value_exp, data_type)}
```

For example, the following statement creates a result set of the names and ages of all employees in their twenties. It uses the **CONVERT** function to convert each employee's age from type SQL_SMALLINT to type SQL_CHAR. Each resulting character string is compared to the pattern "2%" to determine if the employee's age is in the twenties.

```
SELECT NAME, AGE FROM EMPLOYEE  
WHERE {fn CONVERT(AGE,SQL_CHAR)} LIKE '2%'
```

To determine if the **CONVERT** function is supported by a data source, an application calls **SQLGetInfo** with the SQL_CONVERT_FUNCTIONS flag. For more information about the **CONVERT** function, see Appendix F, "Scalar Functions."

LIKE Predicate Escape Characters

In a **LIKE** predicate, the percent character (%) matches zero or more of any character and the underscore character (_) matches any one character. The percent and underscore characters can be used as literals in a **LIKE** predicate 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 data source. The shorthand syntax for the **LIKE** predicate escape character is:

{escape 'escape-character'}

For example, each of the following statements creates the same result set of department names that start with the characters "%AAA". The first statement uses the escape clause syntax. The second statement uses the shorthand syntax. The third statement uses the native syntax for Ingres and is not interoperable among DBMS's. Note that the second percent character in each **LIKE** predicate is a wild-card character that matches zero or more of any character.

```
SELECT NAME FROM DEPT WHERE NAME  
LIKE '\%AAA%' --(*vendor(Microsoft),product(ODBC) escape '\*')--
```

```
SELECT NAME FROM DEPT WHERE NAME LIKE '\%AAA%' {escape '\'}
```

```
SELECT NAME FROM DEPT WHERE NAME LIKE '\%AAA%' ESCAPE '\'
```

To determine whether **LIKE** predicate escape characters are supported by a data source, an application calls **SQLGetInfo** with the **SQL_LIKE_ESCAPE_CLAUSE** information type.

Outer Joins

ODBC supports the ANSI SQL-92 left outer join syntax. The escape clause ODBC uses for outer joins is:

--(*vendor(Microsoft),product(ODBC) oj outer-join *)--

where *outer-join* is:

table-reference **LEFT OUTER JOIN** {*table-reference* | *outer-join*}
ON *search-condition*

table-reference specifies a table name, and *search-condition* specifies the join condition between the *table-references*. The shorthand syntax for outer joins is:

{oj outer-join}

Chapter 9 Terminating Transactions and Connections

An outer join request must appear after the **FROM** keyword and before the **WHERE** clause (if one exists). For complete syntax information, see Appendix C, "SQL Grammar."

For example, each of the following statements creates the same result set of the names and departments of employees working on project 544. The first statement uses the escape clause syntax. The second statement uses the shorthand syntax. The third statement uses the native syntax for Oracle and is not interoperable among DBMS's.

```
SELECT EMPLOYEE.NAME, DEPT.DEPTNAME  
FROM --(*vendor(Microsoft),product(ODBC) oj  
      EMPLOYEE LEFT OUTER JOIN DEPT ON EMPLOYEE.DEPTID=DEPT.DEPTID*)--  
WHERE EMPLOYEE.PROJID=544
```

Part 2 Developing Applications

```
SELECT EMPLOYEE.NAME, DEPT.DEPTNAME
FROM {oj EMPLOYEE LEFT OUTER JOIN DEPT
      ON EMPLOYEE.DEPTID=DEPT.DEPTID}
WHERE EMPLOYEE.PROJID=544
```

```
SELECT EMPLOYEE.NAME, DEPT.DEPTNAME
FROM EMPLOYEE, DEPT
WHERE (EMPLOYEE.PROJID=544) AND (EMPLOYEE.DEPTID = DEPT.DEPTID (+))
```

To determine the level of outer joins a data source supports, an application calls **SQLGetInfo** with the `SQL_OUTER_JOINS` flag. Data sources can support two-table outer joins, partially support multi-table outer joins, fully support multi-table outer joins, or not support outer joins.

Procedures

An application can call a procedure in place of an SQL statement. The escape clause ODBC uses for calling a procedure is:

```
--(*vendor(Microsoft),product(ODBC)
  [?=] call procedure-name[([parameter][,parameter]...)] *)--
```

where *procedure-name* specifies the name of a procedure stored on the data source and *parameter* specifies a procedure parameter. A procedure can have zero or more parameters and can return a value. The shorthand syntax for procedure invocation is:

```
{[?]=}call procedure-name[([parameter][,parameter]...)]}
```

For output parameters, *parameter* must be a parameter marker. For input and input/output parameters, *parameter* can be a literal, a parameter marker, or not specified. If *parameter* is a literal or is not specified for an input/output parameter, the driver discards the output value. If *parameter* is not specified for an input or input/output parameter, the procedure uses the default value of the parameter as the input value; the procedure also uses the default value if *parameter* is a parameter marker and the *pcbValue* argument in **SQLBindParameter** is `SQL_DEFAULT_PARAM`. If a procedure call includes parameter markers (including the "=?", parameter marker for the return value), the application must bind each marker by calling **SQLBindParameter** prior to calling the procedure.

Note: For some data sources, *parameter* cannot be a literal value. For all data sources, it can be a parameter marker. For maximum interoperability, applications should always use a parameter marker for *parameter*.

If an application specifies a return value parameter for a procedure that does not return a value, the driver sets the *pcbValue* buffer specified in **SQLBindParameter** for the parameter to SQL_NULL_DATA. If the application omits the return value parameter for a procedure returns a value, the driver ignores the value returned by the procedure.

If a procedure returns a result set, the application retrieves the data in the result set in the same manner as it retrieves data from any other result set.

For example, each of the following statements uses the procedure EMPS_IN_PROJ to create the same result set of names of employees working on a project. The first statement uses the escape clause syntax. The second statement uses the shorthand syntax. For an example of code that calls a procedure, see **SQLProcedures** in Chapter 22, "ODBC Function Reference."

```
--(*vendor(Microsoft),product(ODBC) call EMPS_IN_PROJ(?)*)--  
  
{call EMPS_IN_PROJ(?)}
```

To determine if a data source supports procedures, an application calls **SQLGetInfo** with the SQL_PROCEDURES information type. To retrieve a list of the procedures stored in a data source, an application calls **SQLProcedures**. To retrieve a list of the input, input/output, and output parameters, as well as the return value and the columns that make up the result set (if any) returned by a procedure, an application calls **SQLProcedureColumns**.

Additional Extension Functions

ODBC also provides the following functions related to SQL statements. For more information about these functions, see Chapter 22, "ODBC Function Reference."

Function	Description
SQLDescribeParam	Retrieves information about prepared parameters.
SQLNativeSql	Retrieves the SQL statement as processed by the data source, with escape sequences translated to SQL code used by the data source.
SQLNumParams	Retrieves the number of parameters in an SQL statement.
SQLSetStmtOption SQLSetConnectOption n SQLGetStmtOption	These functions set or retrieve statement options, such as asynchronous processing, orientation for binding rowsets, maximum amount of variable length data to return, maximum number of result set rows to return, and query timeout value. Note that SQLSetConnectOption sets options for all

statements in a connection.

Retrieving Results

A **SELECT** statement is used to retrieve data that meets a given set of specifications. For example, **SELECT * FROM EMPLOYEE WHERE EMPNAME = "Jones"** is used to retrieve all columns of all rows in EMPLOYEE where the employee's name is Jones. ODBC extension functions also can retrieve data. For example, **SQLColumns** retrieves data about columns in the data source. These sets of data, called result sets, can contain zero or more rows.

Note that other SQL statements, such as **GRANT** or **REVOKE**, do not return result sets. For these statements, the return code from **SQLExecute** or **SQLExecDirect** is usually the only source of information as to whether the statement was successful. (For **INSERT**, **UPDATE**, and **DELETE** statements, an application can call **SQLRowCount** to return the number of affected rows.)

The steps an application takes to process a result set depends on what is known about it.

- **Known result set** The application knows the exact form of the SQL statement, and therefore the result set, at compile time. For example, the query **SELECT EMPNO, EMPNAME FROM EMPLOYEE** returns two specific columns.
- **Unknown result set** The application does not know the exact form of the SQL statement, and therefore the result set, at compile time. For example, the ad hoc query **SELECT * FROM EMPLOYEE** returns all currently defined columns in the EMPLOYEE table. The application may not be able to predict the format of these results prior to execution.

Assigning Storage for Results (Binding)

An application can assign storage for results before or after it executes an SQL statement. If an application prepares or executes the SQL statement first, it can

inquire about the result set before it assigns storage for results. For example, if the result set is unknown, the application must retrieve the number of columns before it can assign storage for them.

To associate storage for a column of data, an application calls **SQLBindCol** and passes it the following information:

- The data type to which the data is to be converted. For more information, see "Converting Data from SQL to C Data Types" in Appendix D, "Data Types."
- The address of an output buffer for the data. The application must allocate this buffer and it must be large enough to hold the data in the form to which it is converted.
- The length of the output buffer. This value is ignored if the returned data has a fixed width in C, such as an integer, real number, or date structure.
- The address of a storage buffer in which to return the number of bytes of available data.

Determining the Characteristics of a Result Set

To determine the characteristics of a result set, an application can:

- Call **SQLNumResultCols** to determine how many columns a request returned.
- Call **SQLColAttributes** or **SQLDescribeCol** to describe a column in the result set.

If the result set is unknown, an application can use the information returned by these functions to bind the columns in the result set. An application can call these functions at any time after a statement is prepared or executed. Note that, although **SQLRowCount** can sometimes return the number of rows in a result set, it is not guaranteed to do so. Few data sources support this functionality and interoperable applications should not rely on it.

Note: For optimal performance, an application should call **SQLColAttributes**, **SQLDescribeCol**, and **SQLNumResultCols** after a statement is executed. In data sources that emulate statement preparation, these functions sometimes execute more slowly before a statement is executed because the information returned by them is not readily available until after the statement is executed.

Fetching Result Data

To retrieve a row of data from the result set, an application:

1. Calls **SQLBindCol** to bind the columns of the result set to storage locations if it has not already done so.
2. Calls **SQLFetch** to move to the next row in the result set and retrieve data for all bound columns.

The following diagram shows the operations an application uses to retrieve data from the result set:

Using Cursors

To keep track of its position in the result set, a driver maintains a cursor. The cursor is so named because it indicates the current position in the result set, just as the cursor on a CRT screen indicates current position.

Each time an application calls **SQLFetch**, the driver moves the cursor to the next row and returns that row. The cursor supported by the core ODBC functions only scrolls forward, one row at a time. (To reretrieve a row of data that it has already retrieved from the result set, the application must close the cursor by calling **SQLFreeStmt** with the **SQL_CLOSE** option, reexecute the **SELECT** statement, and fetch rows with **SQLFetch** until the target row is retrieved.)

Important: Committing or rolling back a transaction, either by calling **SQLTransact** or by using the SQL_AUTOCOMMIT connection option, can cause the data source to close the cursors for all *hstmts* on an *hdbc*. For more information, see the SQL_CURSOR_COMMIT_BEHAVIOR and SQL_CURSOR_ROLLBACK_BEHAVIOR information types in **SQLGetInfo**.

ODBC Extensions for Results

ODBC extends the X/Open and SQL Access Group Call Level Interface to provide additional functions related to retrieving results. The remainder of this chapter describes these functions. To determine if a driver supports a specific function, an application calls **SQLGetFunctions**.

Retrieving Data from Unbound Columns

To retrieve data from unbound columns—that is, columns for which storage has not been assigned with **SQLBindCol**—an application uses **SQLGetData**. The application first calls **SQLFetch** or **SQLExtendedFetch** to position the cursor on the next row. It then calls **SQLGetData** to retrieve data from specific unbound columns.

An application may retrieve data from both bound and unbound columns in the same row. It calls **SQLBindCol** to bind as many columns as desired. It calls **SQLFetch** or **SQLExtendedFetch** to position the cursor on the next row of the result set and retrieve all bound columns. It then calls **SQLGetData** to retrieve data from unbound columns.

If the data type of a column is character, binary, or data source-specific and the column contains more data than can be retrieved in a single call, an application may call **SQLGetData** more than once for that column, as long as the data is being transferred to a buffer of type SQL_C_CHAR or SQL_C_BINARY. For example, data of the SQL_LONGVARIABLE and SQL_LONGVARCHAR types may need to be retrieved in several parts.

For maximum interoperability, an application should only call **SQLGetData** for columns to the right of the rightmost bound column and then only in left-to-right order. To determine if a driver can return data with **SQLGetData** for any column (including unbound columns before the last bound column and any bound columns) or in any order, an application calls **SQLGetInfo** with the SQL_GETDATA_EXTENSIONS option.

Assigning Storage for Rowsets (Binding)

In addition to binding individual rows of data, an application can call **SQLBindCol** to assign storage for a *rowset* (one or more rows of data). By default, rowsets are bound in column-wise fashion. They can also be bound in row-wise fashion.

To specify how many rows of data are in a rowset, an application calls **SQLSetStmtOption** with the `SQL_ROWSET_SIZE` option.

Column-Wise Binding

To assign storage for column-wise bound results, an application performs the following steps for each column to be bound:

1. Allocates an array of data storage buffers. The array has as many elements as there are rows in the rowset.
2. Allocates an array of storage buffers to hold the number of bytes available to return for each data value. The array has as many elements as there are rows in the rowset.
3. Calls **SQLBindCol** and specifies the address of the data array, the size of one element of the data array, the address of the number-of-bytes array, and the type to which the data will be converted. When data is retrieved, the driver will use the array element size to determine where to store successive rows of data in the array.

Row-Wise Binding

To assign storage for row-wise bound results, an application performs the following steps:

1. Declares a structure that can hold a single row of retrieved data and the associated data lengths. (For each column to be bound, the structure contains one field to contain data and one field to contain the number of bytes of data available to return.)
2. Allocates an array of these structures. This array has as many elements as there are rows in the rowset.
3. Calls **SQLBindCol** for each column to be bound. In each call, the application specifies the address of the column's data field in the first array element, the size of the data field, the address of the column's number-of-bytes field in the first array element, and the type to which the data will be converted.
4. Calls **SQLSetStmtOption** with the `SQL_BIND_TYPE` option and specifies the size of the structure. When the data is retrieved, the driver will use the

structure size to determine where to store successive rows of data in the array.

Retrieving Rowset Data

Before it retrieves rowset data, an application calls **SQLSetStmtOption** with the `SQL_ROWSET_SIZE` option to specify the number of rows in the rowset. It then binds columns in the rowset with **SQLBindCol**. The rowset may be bound in column-wise or row-wise fashion. For more information, see "Assigning Storage for Rowsets (Binding)" previous in this chapter.

To retrieve rowset data, an application calls **SQLExtendedFetch**.

For maximum interoperability, an application should not use **SQLGetData** to retrieve data from unbound columns in a block (more than one row) of data that has been retrieved with **SQLExtendedFetch**. To determine if a driver can return data with **SQLGetData** from a block of data, an application calls **SQLGetInfo** with the `SQL_GETDATA_EXTENSIONS` option.

Using Block and Scrollable Cursors

As originally designed, cursors in SQL only scroll forward through a result set, returning one row at a time. However, interactive applications often require forward and backward scrolling, absolute or relative positioning within the result set, and the ability to retrieve and update blocks of data, or *rowsets*.

To retrieve and update rowset data, ODBC provides a *block* cursor attribute. To allow an application to scroll forwards or backwards through the result set, or move to an absolute or relative position in the result set, ODBC provides a *scrollable* cursor attribute. Cursors may have one or both attributes.

Block Cursors

An application calls **SQLSetStmtOption** with the `SQL_ROWSET_SIZE` option to specify the rowset size. The application can call **SQLSetStmtOption** to change the rowset size at any time. Each time the application calls **SQLExtendedFetch**, the driver returns the next *rowset size* rows of data. After the data is returned, the cursor points to the first row in the rowset. By default, the rowset size is one.

Scrollable Cursors

Applications have different needs in their ability to sense changes in the tables underlying a result set. For example, when balancing financial data, an accountant needs data that appears static; it is impossible to balance books when the data is continually changing. When selling concert tickets, a clerk needs up-to-the minute, or dynamic, data on which tickets are still available.

Various cursor models are designed to meet these needs, each of which requires different sensitivities to changes in the tables underlying the result set.

Static Cursors

At one extreme are *static* cursors, to which the data in the underlying tables appears to be static. The membership, order, and values in the result set used by a static cursor are generally fixed when the cursor is opened. Rows updated, deleted, or inserted by other users (including other cursors in the same application) are not detected by the cursor until it is closed and then reopened; the `SQL_STATIC_SENSITIVITY` information type returns whether the cursor can detect rows it has updated, deleted, or inserted.

Static cursors are commonly implemented by taking a snapshot of the data or locking the result set. Note that in the former case, the cursor diverges from the underlying tables as other users make changes; in the latter case, other users are prohibited from changing the data.

Dynamic Cursors

At the other extreme are *dynamic* cursors, to which the data appears to be dynamic. The membership, order, and values in the result set used by a dynamic cursor are ever-changing. Rows updated, deleted, or inserted by all users (the cursor, other cursors in the same application, and other applications) are detected by the cursor when data is next fetched. Although ideal for many situations, dynamic cursors are difficult to implement.

Keyset-Driven Cursors

Between static and dynamic cursors are *keyset-driven* cursors, which have some of the attributes of each. Like static cursors, the membership and ordering of the result set of a keyset-driven cursor is generally fixed when the cursor is opened. Like dynamic cursors, most changes to the values in the underlying result set are visible to the cursor when data is next fetched.

When a keyset-driven cursor is opened, the driver saves the keys for the entire result set, thus fixing the membership and order of the result set. As the cursor scrolls through the result set, the driver uses the keys in this *keyset* to retrieve the current data values for each row in the rowset. Because data values are retrieved only when the cursor scrolls to a given row, updates to that row by other users (including other cursors in the same application) after the cursor was opened are visible to the cursor.

If the cursor scrolls to a row of data that has been deleted by other users (including other cursors in the same application), the row appears as a *hole* in the result set, since the key is still in the keyset but the row is no longer in the result set. Updating the key values in a row is considered to be deleting the existing row and inserting a new row; therefore, rows of data for which the key

values have been changed also appear as holes. When the driver encounters a hole in the result set, it returns a status code of `SQL_ROW_DELETED` for the row.

Rows of data inserted into the result set by other users (including other cursors in the same application) after the cursor was opened are not visible to the cursor, since the keys for those rows are not in the keyset.

The `SQL_STATIC_SENSITIVITY` information type returns whether the cursor can detect rows it has deleted or inserted. Because updating key values in a keyset-driven cursor is considered to be deleting the existing row and inserting a new row, keyset-driven cursors can always detect rows they have updated.

Mixed (Keyset/Dynamic) Cursors

If a result set is large, it may be impractical for the driver to save the keys for the entire result set. Instead, the application can use a *mixed* cursor. In a mixed cursor, the keyset is smaller than the result set, but larger than the rowset.

Within the boundaries of the keyset, a mixed cursor is keyset-driven, that is, the driver uses keys to retrieve the current data values for each row in the rowset. When a mixed cursor scrolls beyond the boundaries of the keyset, it becomes dynamic, that is, the driver simply retrieves the next *rowset size* rows of data. The driver then constructs a new keyset, which contains the new rowset.

For example, assume a result set has 1000 rows and uses a mixed cursor with a keyset size of 100 and a rowset size of 10. When the cursor is opened, the driver (depending on the implementation) saves keys for the first 100 rows and retrieves data for the first 10 rows. If another user deletes row 11 and the cursor then scrolls to row 11, the cursor will detect a hole in the result set; the key for row 11 is in the keyset but the data is no longer in the result set. This is the same behavior as a keyset-driven cursor. However, if another user deletes row 101 and the cursor then scrolls to row 101, the cursor will not detect a hole; the key for the row 101 is not in the keyset. Instead, the cursor will retrieve the data for the row that was originally row 102. This is the same behavior as a dynamic cursor.

Specifying the Cursor Type

To specify the cursor type, an application calls **SQLSetStmtOption** with the `SQL_CURSOR_TYPE` option. The application can specify a cursor that only scrolls forward, a static cursor, a dynamic cursor, a keyset-driven cursor, or a mixed cursor. If the application specifies a mixed cursor, it also specifies the size of the keyset used by the cursor.

Note: To use the ODBC cursor library, an application calls **SQLSetConnectOption** with the `SQL_ODBC_CURSORS` option before it

connects to the data source. The cursor library supports block scrollable cursors. It also supports positioned update and delete statements. For more information, see Appendix G, "ODBC Cursor Library."

Unless the cursor is a forward-only cursor, an application calls **SQLExtendedFetch** to scroll the cursor backwards, forwards, or to an absolute or relative position in the result set. The application calls **SQLSetPos** to refresh the row currently pointed to by the cursor.

Specifying Cursor Concurrency

Concurrency is the ability of more than one user to use the same data at the same time. A transaction is *serializable* if it is performed in a manner in which it appears as if no other transactions operate on the same data at the same time. For example, assume one transaction doubles data values and another adds 1 to data values. If the transactions are serializable and both attempt to operate on the values 0 and 10 at the same time, the final values will be 1 and 21 or 2 and 22, depending on which transaction is performed first. If the transactions are not serializable, the final values will be 1 and 21, 2 and 22, 1 and 22, or 2 and 21; the sets of values 1 and 22, and 2 and 21, are the result of the transactions acting on each value in a different order.

Serializability is considered necessary to maintain database integrity. For cursors, it is most easily implemented at the expense of concurrency by locking the result set. A compromise between serializability and concurrency is *optimistic concurrency control*. In a cursor using optimistic concurrency control, the driver does not lock rows when it retrieves them. When the application requests an update or delete operation, the driver or data source checks if the row has changed. If the row has not changed, the driver or data source prevents other transactions from changing the row until the operation is complete. If the row has changed, the transaction containing the update or delete operation fails.

To specify the concurrency used by a cursor, an application calls **SQLSetStmtOption** with the SQL_CONCURRENCY option. The application can specify that the cursor is read-only, locks the result set, uses optimistic concurrency control and compares row versions to determine if a row has changed, or uses optimistic concurrency control and compares data values to determine if a row has changed. The application calls **SQLSetPos** to lock the row currently pointed to by the cursor, regardless of the specified cursor concurrency.

Using Bookmarks

A bookmark is a 32-bit value that an application uses to return to a row. The application does not request that the driver places a bookmark on a row;

instead, the application requests a bookmark that it can use to return to a row. For example, if a bookmark is a row number, an application requests the row number of a row and stores it. Later, the application passes this row number back to the driver and requests that the driver return to the row.

Before opening the cursor, an application must call **SQLSetStmtOption** with the `SQL_USE_BOOKMARKS` option to inform the driver it will use bookmarks. After opening the cursor, the application retrieves bookmarks either from column 0 of the result set or by calling **SQLGetStmtOption** with the `SQL_GET_BOOKMARK` option. To retrieve a bookmark from the result set, the application either binds column 0 and calls **SQLExtendedFetch** or calls **SQLGetData**; in either case, the *fcType* argument must be set to `SQL_C_BOOKMARK`. To return to the row specified by a bookmark, the application calls **SQLExtendedFetch** with a fetch type of `SQL_FETCH_BOOKMARK`.

If a bookmark requires more than 32 bits, such as when it is a key value, the driver maps the bookmarks requested by the application to 32-bit binary values. The 32-bit binary values are then returned to the application. Because this mapping may require considerable memory, applications should only bind column 0 of the result set if they will actually use bookmarks for most rows. Otherwise, they should call **SQLGetStmtOption** with the SQL_BOOKMARK statement option or call **SQLGetData** for column 0.

Before an application opens a cursor with which it will use bookmarks, it:

- Calls **SQLSetStmtOption** with the SQL_USE_BOOKMARKS option and a value of SQL_UB_ON.

To retrieve a bookmark for the current row, an application:

- Retrieves the value from column 0 of the rowset. The application can either call **SQLBindCol** to bind column 0 before it calls **SQLExtendedFetch** or call **SQLGetData** to retrieve the data after it calls **SQLExtendedFetch**. In either case, the *fCType* argument must be SQL_C_BOOKMARK.

Note: To determine whether it can call **SQLGetData** for a block (more than one row) of data and whether it can call **SQLGetData** for a column before the last bound column, an application calls **SQLGetInfo** with the SQL_GETDATA_EXTENSIONS information type.

– Or –

Calls **SQLSetPos** with the SQL_POSITION option to position the cursor on the row and calls **SQLGetStmtOption** with the SQL_BOOKMARK option to retrieve the bookmark.

To return to the row specified by a bookmark (or a row a certain number of rows from the bookmark), an application:

- Calls **SQLExtendedFetch** with the *irow* argument set to the bookmark and the *fFetchType* argument set to SQL_FETCH_BOOKMARK. The driver returns the rowset starting with the row identified by the bookmark.

Modifying Result Set Data

ODBC provides two ways to modify data in the result set. Positioned update and delete statements are similar to such statements in embedded SQL. Calls to **SQLSetPos** allow an application to update, delete, or add new data without executing SQL statements.

Executing Positioned Update and Delete Statements

An application can update or delete the row in the result set currently pointed to by the cursor. This is known as a positioned update or delete statement. After executing a **SELECT** statement to create a result set, an application calls **SQLFetch** one or more times to position the cursor on the row to be updated or deleted. Alternatively, it fetches the rowset with **SQLExtendedFetch** and positions the cursor on the desired row by calling **SQLSetPos** with the **SQL_POSITION** option. To update or delete the row, the application then executes an SQL statement with the following syntax on a different *hstmt*:

```
UPDATE table-name
SET column-identifier = {expression | NULL}
    [, column-identifier = {expression | NULL}]...
WHERE CURRENT OF cursor-name
DELETE FROM table-name WHERE CURRENT OF cursor-name
```

Positioned update and delete statements require cursor names. An application can name a cursor with **SQLSetCursorName**. If the application has not named the cursor by the time the driver executes a **SELECT** statement, the driver generates a cursor name. To retrieve the cursor name for an *hstmt*, an application calls **SQLGetCursorName**.

To execute a positioned update or delete statement, an application must follow these guidelines:

- The **SELECT** statement that creates the result set must use a **FOR UPDATE** clause.
- The cursor name used in the **UPDATE** or **DELETE** statement must be the same as the cursor name associated with the **SELECT** statement.
- The application must use different *hstmts* for the **SELECT** statement and the **UPDATE** or **DELETE** statement.
- The *hstmts* for the **SELECT** statement and the **UPDATE** or **DELETE** statement must be on the same connection.

To determine if a data source supports positioned update and delete statements, an application calls **SQLGetInfo** with the **SQL_POSITIONED_STATEMENTS** option. For an example of code that performs a positioned update in a rowset, see **SQLSetPos** in Chapter 22, "ODBC Function Reference."

Note: In ODBC 1.0, positioned update, positioned delete, and **SELECT FOR UPDATE** statements were part of the core SQL grammar; in ODBC 2.0, they are part of the extended grammar. Applications that use the SQL conformance level to determine whether these statements are supported also need to check the version number of the driver to correctly interpret the information. In

particular, applications that use these features with ODBC 1.0 drivers need to explicitly check for these capabilities in ODBC 2.0 drivers.

Modifying Data with SQLSetPos

To add, update, and delete rows of data, an application calls **SQLSetPos** and specifies the operation, the row number, and how to lock the row. Where new rows of data are added to the result set, and whether they are visible to the cursor is data source-defined.

The row number determines both the number of the row in the rowset to update or delete and the index of the row in the rowset buffers from which to retrieve data to add or update. If the row number is 0, the operation affects all of the rows in the rowset.

SQLSetPos retrieves the data to update or add from the rowset buffers. It only updates those columns in a row that have been bound with **SQLBindCol** and do not have a length of SQL_IGNORE. However, it cannot add a new row of data unless all of the columns in the row are bound, are nullable, or have a default value.

Note: The rowset buffers are used both to send and retrieve data. To avoid overwriting existing data when it adds a new row of data, an application can allocate an extra row at the end of the rowset buffers to use as an add buffer.

To add a new row of data to the result set, an application:

1. Places the data for each column in the *rgbValue* buffers specified with **SQLBindCol**. To avoid overwriting an existing row of data, the application should allocate an extra row of the rowset buffers to use as an add buffer.
2. Places the length of each column in the *pcbValue* buffer specified with **SQLBindCol**; this only needs to be done for columns with an *fCType* of SQL_C_CHAR or SQL_C_BINARY. To use the default value for a column, the application specifies a length of SQL_IGNORE.

Note: To add a new row of data to a result set, one of the following two conditions must be met:

- All columns in the underlying tables must be bound with **SQLBindCol**.
 - All unbound columns and all bound columns for which the specified length is **SQL_IGNORE** must accept NULL values or have default values.
-

To determine if a row in a result set accepts NULL values, an application calls **SQLColAttributes**. To determine if a data source supports non-nullable columns, an application calls **SQLGetInfo** with the **SQL_NON_NULLABLE** flag.

3. Calls **SQLSetPos** with the *fOption* argument set to **SQL_ADD**. The *irow* argument determines the row in the rowset buffers from which the data is retrieved. For information about how an application sends data for data-at-execution columns, see **SQLSetPos** in Chapter 22, "ODBC Function Reference."

After the row is added, the row the cursor points to is unchanged.

Note: Columns for long data types, such as **SQL_LONGVARCHAR** and **SQL_LONGVARBINARY**, are generally not bound. However, if an application uses **SQLSetPos** to send data for these columns, it must bind them with **SQLBindCol**. Unless the driver returns the **SQL_GD_BOUND** bit for the **SQL_GETDATA_EXTENSIONS** information type, the application must unbind them before calling **SQLGetData** to retrieve data from them.

To update a row of data, an application:

1. Modifies the data of each column to be updated in the *rgbValue* buffer specified with **SQLBindCol**.
2. Places the length of each column to be updated in the *pcbValue* buffer specified with **SQLBindCol**. This only needs to be done for columns with an *fCType* of **SQL_C_CHAR** or **SQL_C_BINARY**.
3. Sets the value of the *pcbValue* buffer for each bound column that is not to be updated to **SQL_IGNORE**.
4. Calls **SQLSetPos** with the *fOption* argument set to **SQL_UPDATE**. The *irow* argument specifies the number of the row in the rowset to modify and the index of row in the rowset buffer from which to retrieve the data. The cursor points to this row after it is updated.

For information about how an application sends data for data-at-execution columns, see **SQLSetPos** in Chapter 22, "ODBC Function Reference."

To delete a row of data, an application:

- Calls **SQLSetPos** with the *fOption* argument set to **SQL_DELETE**. The *irrow* argument specifies the number of the row in the rowset to delete. The cursor points to this row after it is deleted.

Note: The application cannot perform any positioned operations, such as executing a positioned update or delete statement or calling **SQLGetData**, on a deleted row.

To determine what operations a data source supports for **SQLSetPos**, an application calls **SQLGetInfo** with the **SQL_POS_OPERATIONS** flag.

Processing Multiple Results

SELECT statements return result sets. **UPDATE**, **INSERT**, and **DELETE** statements return a count of affected rows. If any of these statements are batched, submitted with arrays of parameters, or in procedures, they can return multiple result sets or counts.

To process a batch of statements, statement with arrays of parameters, or procedure returning multiple result sets or row counts, an application:

1. Calls **SQLExecute** or **SQLExecDirect** to execute the statement or procedure.
2. Calls **SQLRowCount** to determine the number of rows affected by an **UPDATE**, **INSERT**, or **DELETE** statement. For statements or procedures that return result sets, the application calls functions to determine the characteristics of the result set and retrieve data from the result set.
3. Calls **SQLMoreResults** to determine if another result set or row count is available.
4. Repeats steps 2 and 3 until **SQLMoreResults** returns **SQL_NO_DATA_FOUND**.

Retrieving Status and Error Information

This chapter defines the ODBC return codes and error handling protocol. The return codes indicate whether a function succeeded, succeeded but returned a warning, or failed. The error handling protocol defines how the components in an ODBC connection construct and return error messages through **SQLError**.

The protocol describes:

- Use of the error text to identify the source of an error.
- Rules to ensure consistent and useful error information.
- Responsibility for setting the ODBC SQLSTATE based on the native error.

Function Return Codes

When an application calls an ODBC function, the driver executes the function and returns a predefined code. These return codes indicate success, warning, or failure status. The following table defines the return codes.

Return Code	Description
SQL_SUCCESS	Function completed successfully; no additional information is available.
SQL_SUCCESS_WITH_INFO	Function completed successfully, possibly with a nonfatal error. The application can call SQLError to retrieve additional information.
SQL_NO_DATA_FOUND	All rows from the result set have been fetched.
SQL_ERROR	Function failed. The application can call SQLError to retrieve error information.
SQL_INVALID_HANDLE	Function failed due to an invalid environment handle, connection handle, or statement handle. This indicates a programming error. No additional information is available from SQLError .
Return Code	Description

SQL_STILL_EXECUTING	A function that was started asynchronously is still executing.
SQL_NEED_DATA	While processing a statement, the driver determined that the application needs to send parameter data values.

The application is responsible for taking the appropriate action based on the return code.

Retrieving Error Messages

If an ODBC function other than **SQLError** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an application can call **SQLError** to obtain additional information. The application may need to call **SQLError** more than once to retrieve all the error messages from a function, since a function may return more than one error message. When the application calls a different function, the error messages from the previous function are deleted.

Additional error or status information can come from one of two sources:

- Error or status information from an ODBC function, indicating that a programming error was detected.
- Error or status information from the data source, indicating that an error occurred during SQL statement processing.

The information returned by **SQLError** is in the same format as that provided by SQLSTATE in the X/Open and SQL Access Group SQL CAE specification (1992). Note that **SQLError** never returns error information about itself.

ODBC Error Messages

ODBC defines a layered architecture to connect an application to a data source. At its simplest, an ODBC connection requires two components: the Driver Manager and a driver.

A more complex connection might include more components: the Driver Manager, a number of drivers, and a (possibly different) number of DBMS's. The connection might cross computing platforms and operating systems and use a variety of networking protocols.

As the complexity of an ODBC connection increases, so does the importance of providing consistent and complete error messages to the application, its users, and support personnel. Error messages must not only explain the error, but also provide the identity of the component in which it occurred. The identity of the component is particularly important to support personnel when an

application uses ODBC components from more than one vendor. Because **SQLException** does not return the identity of the component in which the error occurred, this information must be embedded in the error text.

Error Text Format

Error messages returned by **SQLException** come from two sources: data sources and components in an ODBC connection. Typically, data sources do not directly support ODBC. Consequently, if a component in an ODBC connection receives an error message from a data source, it must identify the data source as the source of the error. It must also identify itself as the component that received the error.

If the source of an error is the component itself, the error message must explain this. Therefore, the error text returned by **SQLException** has two different formats: one for errors that occur in a data source and one for errors that occur in other components in an ODBC connection.

For errors that do not occur in a data source, the error text must use the format:

[vendor-identifier][ODBC-component-identifier]component-supplied-text

For errors that occur in a data source, the error text must use the format:

**[vendor-identifier][ODBC-component-identifier][data-source-identifier]
data-source-supplied-text**

The following table shows the meaning of each element.

Element	Meaning
<i>vendor-identifier</i>	Identifies the vendor of the component in which the error occurred or that received the error directly from the data source.
<i>ODBC-component-identifier</i>	Identifies the component in which the error occurred or that received the error directly from the data source.
<i>data-source-identifier</i>	Identifies the data source. For single-tier drivers, this is typically a file format, such as Xbase ¹ . For multiple-tier drivers, this is the DBMS product.
<i>component-supplied-text</i>	Generated by the ODBC component.
<i>data-source-supplied-text</i>	Generated by the data source.

¹ In this case, the driver is acting as both the driver and the data source.

Note that the brackets ([]) are included in the error text; they do not indicate optional items.

Sample Error Messages

The following are examples of how various components in an ODBC connection might generate the text of error messages and how various drivers might return them to the application with **SQLError**. Note that these examples do not represent actual implementations of the error handling protocol. For more information on how an individual driver has implemented the protocol, see the documentation for that driver.

Single-Tier Driver

A single-tier driver acts both as an ODBC driver and as a data source. It can therefore generate errors both as a component in an ODBC connection and as a data source. Because it also is the component that interfaces with the Driver Manager, it formats and returns arguments for **SQLError**.

For example, if a Microsoft driver for dBASE® could not allocate sufficient memory, it might return the following arguments for **SQLError**:

```
szSQLState      = "S1001"
pfNativeError   = NULL
szErrorMsg      = "[Microsoft][ODBC dBASE Driver]Unable to
                  ➔ allocate sufficient memory."
pcbErrorMsg     = 67
```

Because this error was not related to the data source, the driver only added prefixes to the error text for the vendor ([Microsoft]) and the driver ([ODBC dBASE Driver]).

If the driver could not find the file EMPLOYEE.DBF, it might return the following arguments for **SQLError**:

```
szSQLState      = "S0002"
pfNativeError   = NULL
szErrorMsg      = "[Microsoft][ODBC dBASE Driver][dBASE]Invalid file
                  ➔ name;file EMPLOYEE.DBF not found."
pcbErrorMsg     = 83
```

Because this error was related to the data source, the driver added the file format of the data source ([dBASE]) as a prefix to the error text. Because the driver was also the component that interfaced with the data source, it added prefixes for the vendor ([Microsoft]) and the driver ([ODBC dBASE Driver]).

Multiple-Tier Driver

A multiple-tier driver sends requests to a DBMS and returns information to the application through the Driver Manager. Because it is the component that

interfaces with the Driver Manager, it formats and returns arguments for **SQLException**.

For example, if a Microsoft driver for DEC's Rdb using SQL/Services encountered a duplicate cursor name, it might return the following arguments for **SQLError**:

```
szSQLState      = "3C000"
pfNativeError   = NULL
szErrorMsg      = "[Microsoft][ODBC Rdb Driver]Duplicate cursor name:
                  ➔ EMPLOYEE_CURSOR."
pcbErrorMsg     = 67
```

Because the error occurred in the driver, it added prefixes to the error text for the vendor ([Microsoft]) and the driver ([ODBC Rdb Driver]).

If the DBMS could not find the table EMPLOYEE, the driver might format and return the following arguments for **SQLError**:

```
szSQLState      = "S0002"
pfNativeError   = -1
szErrorMsg      = "[Microsoft][ODBC Rdb Driver][Rdb]
                  ➔ %SQL-F-RELNOTDEF, Table EMPLOYEE is not defined
                  ➔ in schema."
pcbErrorMsg     = 92
```

Because the error occurred in the data source, the driver added a prefix for the data source identifier ([Rdb]) to the error text. Because the driver was the component that interfaced with the data source, it added prefixes for its vendor ([Microsoft]) and identifier ([ODBC Rdb Driver]) to the error text.

Gateways

In a gateway architecture, a driver sends requests to a gateway that supports ODBC. The gateway sends the requests to a DBMS. Because it is the component that interfaces with the Driver Manager, the driver formats and returns arguments for **SQLError**.

For example, if DEC based a gateway to Rdb on Microsoft Open Data Services, and Rdb could not find the table EMPLOYEE, the gateway might generate the following error text:

```
"[S0002][-1][DEC][ODS Gateway][Rdb]%SQL-F-RELNOTDEF, Table EMPLOYEE
➔ is not defined in schema."
```

Because the error occurred in the data source, the gateway added a prefix for the data source identifier ([Rdb]) to the error text. Because the gateway was the component that interfaced with the data source, it added prefixes for its vendor ([DEC]) and identifier ([ODS Gateway]) to the error text. Note that it also added the SQLSTATE value and the Rdb error code to the beginning of the error text. This permitted it to preserve the semantics of its own message structure and still supply the ODBC error information to the driver.

Because the gateway driver is the component that interfaces with the Driver Manager, it would use the preceding error text to format and return the following arguments for **SQLError**:

```
szSQLState      = "S0002"
pfNativeError   = -1
szErrorMsg      = "[DEC][ODS Gateway][Rdb]%SQL-F-RELNOTDEF, Table
                  ➔ EMPLOYEE is not defined in schema."
pcbErrorMsg     = 81
```

Driver Manager

The Driver Manager can also generate error messages. For example, if an application passed an invalid argument value to **SQLDataSources**, the Driver Manager might format and return the following arguments for **SQLError**:

```
szSQLState      = "S1009"
pfNativeError   = NULL
szErrorMsg      = "[Microsoft][ODBC DLL]Invalid argument value:
                  ➔ SQLDataSources."
pcbErrorMsg     = 60
```

Because the error occurred in the Driver Manager, it added prefixes to the error text for its vendor ([Microsoft]) and its identifier ([ODBC DLL]).

Processing Error Messages

Applications should provide users with all the error information available through **SQLError**: the ODBC SQLSTATE, the native error code, the error text, and the source of the error. The application may parse the error text to separate the text from the information identifying the source of the error. It is the application's responsibility to take appropriate action based on the error or provide the user with a choice of actions.

Note: All messages are generated directly from the ODBC Driver under non Windows platforms.

Terminating Transactions and Connections

The ODBC interface provides functions that terminate statements, transactions, and connections, and free statement (*hstmt*), connection (*hdbc*), and environment (*henv*) handles.

Terminating Statement Processing

To free resources associated with a statement handle, an application calls **SQLFreeStmt**. The **SQLFreeStmt** function has four options:

- **SQL CLOSE** Closes the cursor, if one exists, and discards pending results. The application can use the statement handle again later.
- **SQL DROP** Closes the cursor if one exists, discards pending results, and frees all resources associated with the statement handle.
- **SQL UNBIND** Frees all return buffers bound by **SQLBindCol** for the statement handle.
- **SQL RESET PARAMS** Frees all parameter buffers requested by **SQLBindParameter** for the statement handle.

To cancel a statement that is executing asynchronously, an application:

- Calls **SQLCancel**. When and if the statement is actually canceled is driver- and data source-dependent.
- Calls the function that was executing the statement asynchronously. If the statement is still executing, the function returns **SQL_STILL_EXECUTING**; if it was successfully canceled, the function returns **SQL_ERROR** and **SQLSTATE S1008** (Operation canceled); if it completed normal execution, the function returns any valid return code, such as **SQL_SUCCESS** or **SQL_ERROR**.

- Calls **SQLError** if the function returned SQL_ERROR. If the driver successfully canceled the function, the SQLSTATE will be S1008 (Operation canceled).

Terminating Transactions

An application calls **SQLTransact** to commit or roll back the current transaction.

Terminating Connections

To terminate a connection to a driver and data source, an application performs the following steps:

1. Calls **SQLDisconnect** to close the connection. The application can then use the handle to reconnect to the same data source or to a different data source.
2. Calls **SQLFreeConnect** to free the connection handle and free all resources associated with the handle.
3. Calls **SQLFreeEnv** to free the environment handle and free all resources associated with the handle.

Constructing an ODBC Application

This chapter provides two examples of C-language source code for ODBC-enabled applications.

Sample Application Code

The following sections contain two ODBC examples that are written in the C programming language:

- An example that uses static SQL functions to create a table, add data to it, and select the inserted data.
- An example of interactive, ad-hoc query processing.

Static SQL Example

The following example constructs SQL statements within the application. The example comments include equivalent embedded SQL calls for illustrative purposes.

```
#include "SQL.H"
#include <string.h>

#ifdef NULL
#define NULL 0
#endif

#define MAX_NAME_LEN 50
#define MAX_STMT_LEN 100
int print_err(HDBC hdbc, HSTMT hstmt);
```


SQLAllocStmt

```
int example1(server, uid, pwd)
  UCHAR * server;
  UCHAR * uid;
  UCHAR * pwd;
{
  HENV henv;
  HDBC hdbc;
  HSTMT hstmt;

  SDWORD id;
  UCHAR name[MAX_NAME_LEN + 1];
  UCHAR create[MAX_STMT_LEN]
  UCHAR insert[MAX_STMT_LEN]
  UCHAR select[MAX_STMT_LEN]
  SDWORD namelen;
  RETCODE rc;

  /* EXEC SQL CONNECT TO :server USER :uid USING :pwd; */
  /* Allocate an environment handle.          */
  /* Allocate a connection handle.            */
  /* Connect to a data source.                 */
  /* Allocate a statement handle.              */

  SQLAllocEnv(&henv);
  SQLAllocConnect(henv, &hdbc);
  rc = SQLConnect(hdbc, server, SQL_NTS, uid, SQL_NTS, pwd, SQL_NTS);
  if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(print_err(hdbc, SQL_NULL_HSTMT));
  SQLAllocStmt(hdbc, &hstmt);

  /* EXEC SQL CREATE TABLE NAMEID (ID integer, NAME varchar(50)); */
  /* Execute the SQL statement.          */

  lstrcpy(create, "CREATE TABLE NAMEID (ID INTEGER, NAME
    VARCHAR(50))");
  rc = SQLExecDirect(hstmt, create, SQL_NTS);
  if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(print_err(hdbc, hstmt));

  /* EXEC SQL COMMIT WORK;          */
  /* Commit the table creation.      */

  /* Note that the default transaction mode for drivers that support */
  /* SQLSetConnectOption is auto-commit and SQLTransact has no effect. */

  SQLTransact(hdbc, SQL_COMMIT);
```

```

/* EXEC SQL INSERT INTO NAMEID VALUES (:id, :name); */
/* Show the use of the SQLPrepare/SQLExecute method: */
/* Prepare the insertion and bind parameters.      */
/* Assign parameter values.                        */
/* Execute the insertion.                          */

lstrcpy(insert, "INSERT INTO NAMEID VALUES (?, ?)");
if (SQLPrepare(hstmt, insert, SQL_NTS) != SQL_SUCCESS)
    return(print_err(hdbc, hstmt));
SQLBindParameter(hstmt, 1, SQL_PARAM_INPUT, SQL_C_SLONG, SQL_INTEGER,
    0, 0, &id, 0, NULL);
SQLBindParameter(hstmt, 2, SQL_PARAM_INPUT, SQL_C_CHAR, SQL_VARCHAR,
    MAX_NAME_LEN, 0, name, 0, NULL);
id=500;
lstrcpy(name, "Babbage");
if (SQLExecute(hstmt) != SQL_SUCCESS)
    return(print_err(hdbc, hstmt));

/* EXEC SQL COMMIT WORK; */
/* Commit the insertion. */

SQLTransact(hdbc, SQL_COMMIT);

/* EXEC SQL DECLARE c1 CURSOR FOR SELECT ID, NAME FROM NAMEID; */
/* EXEC SQL OPEN c1; */
/* Show the use of the SQLExecDirect method. */
/* Execute the selection. */
/* Note that the application does not declare a cursor. */

lstrcpy(select, "SELECT ID, NAME FROM NAMEID");
if (SQLExecDirect(hstmt, select, SQL_NTS) != SQL_SUCCESS)
    return(print_err(hdbc, hstmt));

/* EXEC SQL FETCH c1 INTO :id, :name; */
/* Bind the columns of the result set with SQLBindCol. */
/* Fetch the first row. */

SQLBindCol(hstmt, 1, SQL_C_SLONG, &id, 0, NULL);
SQLBindCol(hstmt, 2, SQL_C_CHAR, name, (SDWORD)sizeof(name), &namelen);
SQLFetch(hstmt);

/* EXEC SQL COMMIT WORK; */
/* Commit the transaction. */

SQLTransact(hdbc, SQL_COMMIT);

/* EXEC SQL CLOSE c1; */
/* Free the statement handle. */

```

SQLAllocStmt

```
SQLFreeStmt(hstmt, SQL_DROP);

/* EXEC SQL DISCONNECT;      */
/* Disconnect from the data source. */
/* Free the connection handle.    */
/* Free the environment handle.   */

SQLDisconnect(hdbc);
SQLFreeConnect(hdbc);
SQLFreeEnv(henv);

return(0);
}
```

Interactive Ad Hoc Query Example

The following example illustrates how an application can determine the nature of the result set prior to retrieving results.

```
#include "SQL.H"
#include <string.h>
#include <stdlib.h>

#define MAXCOLS 100
#define max(a,b) (a>b?a:b)

int  print_err(HDBC hdbc, HSTMT hstmt);
UDWORD display_size(SWORD coltype, UDWORD collen, UCHAR *colname);

example2(server, uid, pwd, sqlstr)
UCHAR * server;
UCHAR * uid;
UCHAR * pwd;
UCHAR * sqlstr;
{
    int  i;
    HENV henv;
    HDBC hdbc;
    HSTMT hstmt;
    UCHAR errmsg[256];
    UCHAR colname[32];
    SWORD coltype;
    SWORD colnamelen;
    SWORD nullable;
    UDWORD collen[MAXCOLS];
    SWORD scale;
    SDWORD outlen[MAXCOLS];
```

```
    UCHAR * data[MAXCOLS];
    SWORD  nresultcols;
    SDWORD rowcount;
    RETCODE rc;

    /* Allocate environment and connection handles. */
    /* Connect to the data source.                  */
    /* Allocate a statement handle.                  */
    SQLAllocEnv(&henv);
    SQLAllocConnect(henv, &hdbc);
    rc = SQLConnect(hdbc, server, SQL_NTS, uid, SQL_NTS, pwd, SQL_NTS);
    if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
        return(print_err(hdbc, SQL_NULL_HSTMT));
    SQLAllocStmt(hdbc, &hstmt);
```

SQLAllocStmt

```
/* Execute the SQL statement. */
if (SQLExecDirect(hstmt, sqlstr, SQL_NTS) != SQL_SUCCESS)
    return(print_err(hdbc, hstmt));

/* See what kind of statement it was. If there are no result
/* columns, the statement is not a SELECT statement. If the
/* number of affected rows is greater than 0, the statement was
/* probably an UPDATE, INSERT, or DELETE statement, so print the
/* number of affected rows. If the number of affected rows is 0,
/* the statement is probably a DDL statement, so print that the
/* operation was successful and commit it. */

SQLNumResultCols(hstmt, &nresultcols);
if (nresultcols == 0) {
    SQLRowCount(hstmt, &rowcount);
    if (rowcount > 0) {
        printf("%d rows affected.\n", rowcount);
    } else {
        printf("Operation successful.\n");
    }
    SQLTransact(hdbc, SQL_COMMIT);

/* Otherwise, display the column names of the result set and use the
/* display_size() function to compute the length needed by each data
/* type. Next, bind the columns and specify all data will be
/* converted to char. Finally, fetch and print each row, printing
/* truncation messages as necessary. */

} else {
    for (i = 0; i < nresultcols; i++) {
        SQLDescribeCol(hstmt, i + 1, colname, (SWORD)sizeof(colname),
            &colnamelen, &coltype, &collen[i], &scale,
            &nullable);
        collen[i] = display_size(coltype, collen[i], colname);
        printf("%s", collen[i], collen[i], colname);
        data[i] = (UCHAR *) malloc(collen[i] + 1);
        SQLBindCol(hstmt, i + 1, SQL_C_CHAR, data[i], collen[i],
            &outlen[i]);
    }
    while (TRUE) {
        rc = SQLFetch(hstmt);
        if (rc == SQL_SUCCESS || rc == SQL_SUCCESS_WITH_INFO) {
            errmsg[0] = '\0';
            for (i = 0; i < nresultcols; i++)
                if (outlen[i] == SQL_NULL_DATA) {
                    strcpy(data[i], "NULL");
                } else if (outlen[i] >= collen[i]) {
```

```

        sprintf(&errmsg[strlen(errmsg)],
            "%d chars truncated, col %d\n",
            *outlen[i] - collen[i] + 1,
            colnum);
    }
    printf("%*.*s ", collen[i], collen[i], data[i]);
}
printf("\n%s", errmsg);
} else {
    break;
}
}
}

/* Free the data buffers. */
for (i = 0; i < nresultcols; i++) {
    free(data[i]);
}

SQLFreeStmt(hstmt, SQL_DROP); /* Free the statement handle. */
SQLDisconnect(hdbc);          /* Disconnect from the data source. */
SQLFreeConnect(hdbc);         /* Free the connection handle. */
SQLFreeEnv(henv);             /* Free the environment handle. */

return(0);
}

/*****
/* The following function is included for completeness, but */
/* is not relevant for understanding the function of ODBC. */
*****/

#define MAX_NUM_PRECISION 15

/* Define max length of char string representation of number as: */
/* = max(precision) + leading sign + E + exp sign + max exp length */
/* = 15      + 1      + 1 + 1      + 2      */
/* = 15 + 5                                         */

#define MAX_NUM_STRING_SIZE (MAX_NUM_PRECISION + 5)

UDWORD display_size(coltype, collen, colname)
SWORD coltype;
UDWORD collen;
UCHAR * colname;
{
    switch (coltype) {

```


SQLAllocStmt

```
case SQL_CHAR:
case SQL_VARCHAR:
    return(max(collen, strlen(colname)));

case SQL_SMALLINT:
    return(max(6, strlen(colname)));

case SQL_INTEGER:
    return(max(11, strlen(colname)));

case SQL_DECIMAL:
case SQL_NUMERIC:
case SQL_REAL:
case SQL_FLOAT:
case SQL_DOUBLE:
    return(max(MAX_NUM_STRING_SIZE, strlen(colname)));

/* Note that this function only supports the core data types. */
default:
    printf("Unknown datatype, %d\n", coltype);
    return(0);
}
}
```

P A R T 3

Developing Drivers

SQLBrowseConnect

C H A P T E R S 1 1 - 1 8

- omitted -

Installing and Configuring ODBC Software

- omitted -

API Reference

Function Summary

This chapter summarizes the functions used by ODBC-enabled applications and related software:

- ODBC functions
- Setup DLL functions
- Installer DLL functions
- Translation DLL functions

ODBC Function Summary

The following table lists ODBC functions, grouped by type of task, and includes the conformance designation and a brief description of the purpose of each function. For more information about the syntax and semantics for each function, see Chapter 22, “ODBC Function Reference.”

An application can call the **SQLGetInfo** function to obtain conformance information about a driver. To obtain information about support for a specific function in a driver, an application can call **SQLGetFunctions**.

Task	Function Name	Conformance	Purpose
Connecting to a Data Source	SQLAllocEnv	Core	Obtains an environment handle. One environment handle is used for one or more connections.
	SQLAllocConnect	Core	Obtains a connection handle.
	SQLConnect	Core	Connects to a specific driver by data source name, user ID, and password.
SQLDisconnect			

SQLDriverConnect (Windows)

	SQLDriverConnect	Level 1	Connects to a specific driver by connection string or requests that the Driver Manager and driver display connection dialog boxes for the user.
	SQLBrowseConnect	Level 2	Returns successive levels of connection attributes and valid attribute values. When a value has been specified for each connection attribute, connects to the data source.
Obtaining Information about a Driver and Data Source	SQLDataSources (Windows)	Level 2	Returns the list of available data sources.
	SQLDrivers (Windows)	Level 2	Returns the list of installed drivers and their attributes.
	SQLGetInfo	Level 1	Returns information about a specific driver and data source.
	SQLGetFunctions	Level 1	Returns supported driver functions.
	SQLGetTypeInfo	Level 1	Returns information about supported data types.
Setting and Retrieving Driver Options	SQLSetConnectOption	Level 1	Sets a connection option.
	SQLGetConnectOption	Level 1	Returns the value of a connection option.
	SQLSetStmtOption	Level 1	Sets a statement option.
	SQLGetStmtOption	Level 1	Returns the value of a statement option.
Task	Function Name	Conformance	Purpose
Preparing SQL Requests	SQLAllocStmt	Core	Allocates a statement handle.
	SQLPrepare	Core	Prepares an SQL statement for later execution.
	SQLBindParameter	Level 1	Assigns storage for a parameter in an SQL statement.
	SQLParamOptions	Level 2	Specifies the use of multiple values for parameters.
	SQLGetCursorName	Core	Returns the cursor name associated with a statement handle.
	SQLSetCursorName	Core	Specifies a cursor name.
Submitting Requests	SQLSetScrollOptions	Level 2	Sets options that control cursor behavior.
	SQLExecute	Core	Executes a prepared statement.
	SQLExecDirect	Core	Executes a statement.

	SQLNativeSql	Level 2	Returns the text of an SQL statement as translated by the driver.
	SQLDescribeParam	Level 2	Returns the description for a specific parameter in a statement.
	SQLNumParams	Level 2	Returns the number of parameters in a statement.
	SQLParamData	Level 1	Used in conjunction with SQLPutData to supply parameter data at execution time. (Useful for long data values.)
	SQLPutData	Level 1	Send part or all of a data value for a parameter. (Useful for long data values.)

Task	Function Name	Conformance	Purpose
Retrieving Results and Information about Results	SQLRowCount	Core	Returns the number of rows affected by an insert, update, or delete request.
	SQLNumResultCols	Core	Returns the number of columns in the result set.
	SQLDescribeCol	Core	Describes a column in the result set.
	SQLColAttributes	Core	Describes attributes of a column in the result set.
	SQLBindCol	Core	Assigns storage for a result column and specifies the data type.
	SQLFetch	Core	Returns a result row.
	SQLExtendedFetch	Level 2	Returns multiple result rows.
	SQLGetData	Level 1	Returns part or all of one column of one row of a result set. (Useful for long data values.)
	SQLSetPos	Level 2	Positions a cursor within a fetched block of data.
	SQLMoreResults	Level 2	Determines whether there are more result sets available and, if so, initializes processing for the next result set.
	SQLError	Core	Returns additional error or status information.

Task	Function Name	Conformance	Purpose
Obtaining information about the data source's system tables	SQLColumnPrivileges	Level 2	Returns a list of columns and associated privileges for one or more tables.
SQLDriverConnect (Windows)			

SQLDisconnect

(catalog functions)	SQLColumns	Level 1	Returns the list of column names in specified tables.
	SQLForeignKeys	Level 2	Returns a list of column names that comprise foreign keys, if they exist for a specified table.
	SQLPrimaryKeys	Level 2	Returns the list of column name(s) that comprise the primary key for a table.
	SQLProcedureColumns	Level 2	Returns the list of input and output parameters, as well as the columns that make up the result set for the specified procedures.
	SQLProcedures	Level 2	Returns the list of procedure names stored in a specific data source.
	SQLSpecialColumns	Level 1	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	Level 1	Returns statistics about a single table and the list of indexes associated with the table.
	SQLTablePrivileges	Level 2	Returns a list of tables and the privileges associated with each table.
	SQLTables	Level 1	Returns the list of table names stored in a specific data source.
Terminating a Statement	SQLFreeStmt	Core	Ends statement processing and closes the associated cursor, discards pending results, and, optionally, frees all resources associated with the statement handle.
	SQLCancel	Core	Cancels an SQL statement.
	SQLTransact	Core	Commits or rolls back a transaction.
Terminating a Connection	SQLDisconnect	Core	Closes the connection.
	SQLFreeConnect	Core	Releases the connection handle.
	SQLFreeEnv	Core	Releases the environment handle.

ODBC Function Reference

The following pages describe each ODBC function in alphabetic order. Each function is defined as a C programming language function. Descriptions include the following:

- Purpose
- ODBC version
- Conformance level
- Syntax
- Arguments
- Return values
- Diagnostics
- Comments about usage and implementation
- Code example
- References to related functions

Error handling is described in the **SQLError** function description. The text associated with SQLSTATE values is included to provide a description of the condition, but is not intended to prescribe specific text.

Arguments

All function arguments use a naming convention of the following form:

`[[prefix]...]tag[qualifier][suffix]`

Optional elements are enclosed in square brackets ([]). The following prefixes are used:

Prefix	Description
--------	-------------

c	Count of
h	Handle of
i	Index of
p	Pointer to
rg	Range (array) of

The following tags are used:

Tag	Description
-----	-------------

b	Byte
col	Column (of a result set)
dbc	Database connection
env	Environment
f	Flag (enumerated type)
par	Parameter (of an SQL statement)
row	Row (of a result set)
stmt	Statement
sz	Character string (array of characters, terminated by zero)
v	Value of unspecified type

Prefixes and tags combine to correspond roughly to the ODBC C types listed below. Flags (f) and byte counts (cb) do not distinguish between SWORD, UWORD, SDWORD, and UDWORD.

Combined	Prefix	Tag	ODBC C Type(s)	Description
cb	c	b	SWORD, SDWORD, UDWORD	Count of bytes
crow	c	row	SDWORD, UDWORD, UWORD	Count of rows
f	—	f	SWORD, UWORD	Flag
hdbc	h	dbc	HDBC	Connection handle
henv	h	env	HENV	Environment handle
hstmt	h	stmt	HSTMT	Statement handle
hwnd	h	wnd	HWND	Window handle
Combined	Prefix	Tag	ODBC C Type(s)	Description

ib	i	b	SWORD	Byte index
icol	i	col	UWORD	Column index
ipar	i	par	UWORD	Parameter index
irow	i	row	SDWORD, UWORD	Row index
pcb	pc	b	SWORD FAR *, SDWORD FAR *, UDWORD FAR *	Pointer to byte count
pccol	pc	col	SWORD FAR *	Pointer to column count
pcpar	pc	par	SWORD FAR *	Pointer to parameter count
pcrow	pc	row	SDWORD FAR *, UDWORD FAR *	Pointer to row count
pf	p	f	SWORD, SDWORD, UWORD	Pointer to flag
phdbc	ph	dbc	HDBC FAR *	Pointer to connection handle
phenv	ph	env	HENV FAR *	Pointer to environment handle
phstmt	ph	stmt	HSTMT FAR *	Pointer to statement handle
pib	pi	b	SWORD FAR *	Pointer to byte index
pirow	pi	row	UDWORD FAR *	Pointer to row index
prgb	prg	b	PTR FAR *	Pointer to range (array) of bytes
pv	p	v	PTR	Pointer to value of unspecified type
rgb	rg	b	PTR	Range (array) of bytes
rgf	rg	f	UWORD FAR *	Range (array) of flags
sz	—	sz	UCHAR FAR *	String, zero terminated
v	—	v	UDWORD	Value of unspecified type

Qualifiers are used to distinguish specific variables of the same type. Qualifiers consist of the concatenation of one or more capitalized English words or abbreviations.

ODBC defines one value for the suffix *Max*, which denotes that the variable represents the largest value of its type for a given situation.

For example, the argument *cbErrorMsgMax* contains the largest possible byte count for an error message; in this case, the argument corresponds to the size in bytes of the argument *szErrorMsg*, a character string buffer. The argument *pcbErrorMsg* is a pointer to the count of bytes available to return in the argument *szErrorMsg*, not including the null termination character.

ODBC Include Files

The files SQL.H and SQLEXT.H contain function prototypes for all of the ODBC functions. They also contain all type definitions and **#define** names used by ODBC.

Diagnostics

The diagnostics provided with each function list the SQLSTATEs that may be returned for the function by the Driver Manager or a driver. Drivers can, however, return additional SQLSTATEs arising out of implementation-specific situations.

The character string value returned for an SQLSTATE consists of a two-character class value followed by a three-character subclass value. A class value of "01" indicates a warning and is accompanied by a return code of SQL_SUCCESS_WITH_INFO. Class values other than "01", except for the class "IM", indicate an error and are accompanied by a return code of SQL_ERROR. The class "IM" is specific to warnings and errors that derive from the implementation of ODBC itself. The subclass value "000" in any class is for implementation-defined conditions within the given class. The assignment of class and subclass values is defined by ANSI SQL-92.

Tables and Views

In ODBC functions, tables and views are interchangeable. The term *table* is used for both tables and views, except where view is used explicitly.

Catalog Functions

ODBC supports a set of functions that return information about the data source's system tables or catalog. These are sometimes referred to collectively as the *catalog functions*. For more information about catalog functions, see "Retrieving Information About the Data Source's Catalog" in Chapter 6, "Executing SQL Statements," and "Returning Information About the Data Source's Catalog" in Chapter 14, "Processing an SQL Statement." The catalog functions are:

SQLColumnPrivileges

SQLColumns

SQLForeignKeys

SQLProcedures

SQLSpecialColumns

SQLStatistics

SQLExecDirect

SQLPrimaryKeys
SQLProcedureColumns

SQLTablePrivileges
SQLTables

Search Pattern Arguments

Each catalog function returns information in the form of a result set. The information returned by a function may be constrained by a search pattern passed as an argument to that function. These search patterns can contain the metacharacters underscore (`_`) and percent (`%`) and a driver-defined escape character as follows:

- The underscore character represents any single character.
- The percent character represents any sequence of zero or more characters.
- The escape character permits the underscore and percent metacharacters to be used as literal characters in search patterns. To use a metacharacter as a literal character in the search pattern, precede it with the escape character. To use the escape character as a literal character in the search pattern, include it twice. To obtain the escape character for a driver, an application must call **SQLGetInfo** with the **SQL_SEARCH_PATTERN_ESCAPE** option.
- All other characters represent themselves.

For example, if the search pattern for a table name is `"%A%"`, the function will return all tables with names that contain the character "A". If the search pattern for a table name is `"B__"` ("B" followed by two underscores), the function will return all tables with names that are three characters long and start with the character "B". If the search pattern for a table name is `"%"`, the function will return all tables.

Suppose the search pattern escape character for a driver is a backslash (`\`). If the search pattern for a table name is `"ABC\%"`, the function will return the table named "ABC%." If the search pattern for a table name is `"\\%"`, the function will return all tables with names that start with a backslash. Failing to precede a metacharacter used as a literal with an escape character may return more results than expected. For example, if a table identifier, "MY_TABLE" was returned as the result of a call to **SQLTables** and an application wanted to retrieve a list of columns for "MY_TABLE" using **SQLColumns**, **SQLColumns** would return all of the tables that matched MY_TABLE, such as MY_TABLE, MY1TABLE, MY2TABLE, and so on, unless the escape character precedes the underscore.

Note: A zero-length search pattern matches the empty string. A search pattern argument that is a null pointer means the search will not be constrained for that argument. (A null pointer and a search string of “%” should return the same values.)

SQLAllocConnect



Core **SQLAllocConnect** allocates memory for a connection handle within the environment identified by *henv*.

Syntax RETCODE SQLAllocConnect(*henv*, *phdbc*)

The SQLAllocConnect function accepts the following arguments.

Type	Argument	Use	Description
HENV	<i>henv</i>	Input	Environment handle.
HDBC FAR *	<i>phdbc</i>	Output	Pointer to storage for the connection handle.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

If **SQLAllocConnect** returns SQL_ERROR, it will set the *hdbc* referenced by *phdbc* to SQL_NULL_HDBC. To obtain additional information, the application can call **SQLError** with the specified *henv* and with *hdbc* and *hstmt* set to SQL_NULL_HDBC and SQL_NULL_HSTMT, respectively.

Diagnostics When **SQLAllocConnect** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLAllocConnect** and explains each one in the context of this function; the notation “(DM)” precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
S1000	General error	An error occurred for which there was no

SQLExecDirect

		specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
	S1001	Memory allocation failure (DM) The Driver Manager was unable to allocate memory for the connection handle. The driver was unable to allocate memory for the connection handle.
	SQLSTATE	Error
	S1009	Invalid argument value (DM) The argument <i>phdbc</i> was a null pointer.
Comments	<p>A connection handle references information such as the valid statement handles on the connection and whether a transaction is currently open. To request a connection handle, an application passes the address of an <i>hdbc</i> to SQLAllocConnect. The driver allocates memory for the connection information and stores the value of the associated handle in the <i>hdbc</i>. On operating systems that support multiple threads, applications can use the same <i>hdbc</i> on different threads and drivers must therefore support safe, multithreaded access to this information. The application passes the <i>hdbc</i> value in all subsequent calls that require an <i>hdbc</i>.</p> <p>The Driver Manager processes the SQLAllocConnect function and calls the driver's SQLAllocConnect function when the application calls SQLConnect, SQLBrowseConnect, or SQLDriverConnect. (For more information, see the description of the SQLConnect function.)</p> <p>If the application calls SQLAllocConnect with a pointer to a valid <i>hdbc</i>, the driver overwrites the <i>hdbc</i> without regard to its previous contents.</p>	
Code Example	See SQLBrowseConnect and SQLConnect .	
Related Functions	For information about	See
	Connecting to a data source	SQLConnect
	Freeing a connection handle	SQLFreeConnect

SQLAllocEnv

ODBC

Core **SQLAllocEnv** allocates memory for an environment handle and initializes the ODBC call level interface for use by an application. An application must call **SQLAllocEnv** prior to calling any other ODBC function.

Syntax RETCODE **SQLAllocEnv**(*phenv*)

The **SQLAllocEnv** function accepts the following argument.

Type	Argument	Use	Description
HENV FAR *	<i>phenv</i>	Output	Pointer to storage for the environment handle.

Returns SQL_SUCCESS or SQL_ERROR.

If **SQLAllocEnv** returns SQL_ERROR, it will set the *henv* referenced by *phenv* to SQL_NULL_HENV. In this case, the application can assume that the error was a memory allocation error.

Diagnostics A driver cannot return SQLSTATE values directly after the call to **SQLAllocEnv**, since no valid handle will exist with which to call **SQLError**.

There are two levels of **SQLAllocEnv** functions, one within the Driver Manager and one within each driver. The Driver Manager does not call the driver-level function until the application calls **SQLConnect**, **SQLBrowseConnect**, or **SQLDriverConnect**. If an error occurs in the driver-level **SQLAllocEnv** function, then the Driver Manager–level **SQLConnect**, **SQLBrowseConnect**, or **SQLDriverConnect** function returns SQL_ERROR. A subsequent call to **SQLError** with *henv*, SQL_NULL_HDBC, and SQL_NULL_HSTMT returns SQLSTATE IM004 (Driver's **SQLAllocEnv** failed), followed by one of the following errors from the driver:

- SQLSTATE S1000 (General error).
- A driver-specific SQLSTATE value, ranging from S1000 to S19ZZ. For example, SQLSTATE S1001 (Memory allocation failure) indicates that the Driver Manager's call to the driver-level **SQLAllocEnv** returned SQL_ERROR, and the Driver Manager's *henv* was set to SQL_NULL_HENV.

For additional information about the flow of function calls between the Driver Manager and a driver, see the **SQLConnect** function description.

Comments

An environment handle references global information such as valid connection handles and active connection handles. To request an environment handle, an application passes the address of an *henv* to **SQLAllocEnv**.

The driver allocates memory for the environment information and stores the value of the associated handle in the *henv*. On operating systems that support multiple threads, applications can use the same *henv* on different threads and drivers must therefore support safe, multithreaded access to this information. The application passes the *henv* value in all subsequent calls that require an *henv*.

There should never be more than one *henv* allocated at one time and the application should not call **SQLAllocEnv** when there is a current valid *henv*. If the application calls **SQLAllocEnv** with a pointer to a valid *henv*, the driver overwrites the *henv* without regard to its previous contents.

When the Driver Manager processes the **SQLAllocEnv** function, it checks the **Trace** keyword in the [ODBC] section of the ODBC.INI file or the ODBC subkey in the registry. If it is set to 1, the Driver Manager enables tracing for all applications on Windows 3.1 or for the current application on Windows NT.

Code Example

See **SQLBrowseConnect** and **SQLConnect**.

Related Functions**For information about****See**

Allocating a connection handle

SQLAllocConnect

Connecting to a data source

SQLConnect

Freeing an environment handle

SQLFreeEnv

SQLAllocStmt



Core **SQLAllocStmt** allocates memory for a statement handle and associates the statement handle with the connection specified by *hdbc*.

An application must call **SQLAllocStmt** prior to submitting SQL statements.

Syntax RETCODE **SQLAllocStmt**(*hdbc*, *phstmt*)

The **SQLAllocStmt** function accepts the following arguments.

Type	Argument	Use	Description
HDBC	<i>hdbc</i>	Input	Connection handle.
HSTMT FAR *	<i>phstmt</i>	Output	Pointer to storage for the statement handle.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_INVALID_HANDLE, or SQL_ERROR.

If **SQLAllocStmt** returns SQL_ERROR, it will set the *hstmt* referenced by *phstmt* to SQL_NULL_HSTMT. The application can then obtain additional information by calling **SQLError** with the *hdbc* and SQL_NULL_HSTMT.

Diagnostics When **SQLAllocStmt** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLAllocStmt** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
08003	Connection not open	(DM) The connection specified by the <i>hdbc</i> argument was not open. The connection process must be completed successfully (and the connection must be open) for the driver to allocate an <i>hstmt</i> .
IM001	Driver does not	(DM) The driver associated with the <i>hdbc</i>

support this function		does not support the function.
SQLSTATE	Error	Description
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	(DM) The Driver Manager was unable to allocate memory for the statement handle. The driver was unable to allocate memory for the statement handle.
S1009	Invalid argument value	(DM) The argument <i>phstmt</i> was a null pointer.

Comments

A statement handle references statement information, such as network information, SQLSTATE values and error messages, cursor name, number of result set columns, and status information for SQL statement processing.

To request a statement handle, an application connects to a data source and then passes the address of an *hstmt* to **SQLAllocStmt**. The driver allocates memory for the statement information and stores the value of the associated handle in the *hstmt*. On operating systems that support multiple threads, applications can use the same *hstmt* on different threads and drivers must therefore support safe, multithreaded access to this information. The application passes the *hstmt* value in all subsequent calls that require an *hstmt*.

If the application calls **SQLAllocStmt** with a pointer to a valid *hstmt*, the driver overwrites the *hstmt* without regard to its previous contents.

Code Example

See **SQLBrowseConnect**, **SQLConnect**, and **SQLSetCursorName**.

Related Functions For information about**See**

Executing an SQL statement	SQLExecDirect
Executing a prepared SQL statement	SQLExecute
Freeing a statement handle	SQLFreeStmt
Preparing a statement for execution	SQLPrepare

SQLExtendedFetch

SQLBindCol



Core

SQLBindCol assigns the storage and data type for a column in a result set, including:

- A storage buffer that will receive the contents of a column of data
- The length of the storage buffer
- A storage location that will receive the actual length of the column of data returned by the fetch operation
- Data type conversion

Syntax

RETCODE **SQLBindCol**(*hstmt*, *icol*, *fCType*, *rgbValue*, *cbValueMax*, *pcbValue*)

The **SQLBindCol** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UWORD	<i>icol</i>	Input	Column number of result data, ordered sequentially left to right, starting at 1. A column number of 0 is used to retrieve a bookmark for the row; bookmarks are not supported by ODBC 1.0 drivers or by SQLFetch .
Type	Argument	Use	Description
SWORD	<i>fCType</i>	Input	The C data type of the result data. This must be one of the following values: SQL_C_BINARY SQL_C_BIT SQL_C_BOOKMARK SQL_C_CHAR SQL_C_DATE SQL_C_DEFAULT SQL_C_DOUBLE SQL_C_FLOAT SQL_C_SLONG SQL_C_SSHORT SQL_C_STINYINT SQL_C_TIME SQL_C_TIMESTAMP SQL_C_ULONG

SQL_C_USHORT
SQL_C_UTINYINT

SQL_C_DEFAULT specifies that data be transferred to its default C data type.

Note: Drivers must also support the following values of *fCType* from ODBC 1.0. Applications must use these values, rather than the ODBC 2.0 values, when calling an ODBC 1.0 driver:

SQL_C_LONG
SQL_C_SHORT
SQL_C_TINYINT

For more information, see "ODBC 1.0 C Data Types" in Appendix D, "Data Types."

For information about how data is converted, see "Converting Data from SQL to C Data Types" in Appendix D, "Data Types."

Type	Argument	Use	Description
PTR	<i>rgbValue</i>	Input	<p>Pointer to storage for the data. If <i>rgbValue</i> is a null pointer, the driver unbinds the column. (To unbind all columns, an application calls SQLFreeStmt with the SQL_UNBIND option.)</p> <hr/> <p>Note: If a null pointer was passed for <i>rgbValue</i> in ODBC 1.0, the driver returned SQLSTATE S1009 (Invalid argument value); individual columns could not be unbound.</p> <hr/>
SDWORD	<i>cbValueMax</i>	Input	<p>Maximum length of the <i>rgbValue</i> buffer. For character data, <i>rgbValue</i> must also include space for the null-termination byte. For more information about length, see "Precision, Scale, Length, and Display Size" in Appendix D, "Data Types."</p>

SQLFetch

SQLFetch

Type	Argument	Use	Description
SDWORD FAR *	<i>pcbValue</i>	Input	<p>SQL_NULL_DATA or the number of bytes (excluding the null termination byte for character data) available to return in <i>rgbValue</i> prior to calling SQLExtendedFetch or SQLFetch, or SQL_NO_TOTAL if the number of available bytes cannot be determined.</p> <p>For character data, if the number of bytes available to return is SQL_NO_TOTAL or is greater than or equal to <i>cbValueMax</i>, the data in <i>rgbValue</i> is truncated to <i>cbValueMax</i> – 1 bytes and is null-terminated by the driver.</p> <p>For binary data, if the number of bytes available to return is SQL_NO_TOTAL or is greater than <i>cbValueMax</i>, the data in <i>rgbValue</i> is truncated to <i>cbValueMax</i> bytes.</p> <p>For all other data types, the value of <i>cbValueMax</i> is ignored and the driver assumes the size of <i>rgbValue</i> is the size of the C data type specified with <i>fCType</i>.</p> <p>For more information about the value returned in <i>pcbValue</i> for each <i>fCType</i>, see "Converting Data from SQL to C Data Types" in Appendix D, "Data Types."</p>

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, SQL_INVALID_HANDLE.

Diagnostics When **SQLBindCol** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLBindCol** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable 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 0 and the driver was an ODBC 1.0 driver. 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	(DM) The argument <i>fCType</i> was not a valid data type or SQL_C_DEFAULT. The argument <i>icol</i> was 0 and the argument <i>fCType</i> was not SQL_C_BOOKMARK.
S1009	Invalid argument value	The driver supported ODBC 1.0 and the argument <i>rgbValue</i> was a null pointer.
S1010	Function sequence error	(DM) An asynchronously executing function was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1090	Invalid string or buffer length	(DM) The value specified for the argument <i>cbValueMax</i> was less than 0.

SQLSTATE Error		Description
S1C00	Driver not capable	<p>The driver does not support the data type specified in the argument <i>fCType</i>.</p> <p>The argument <i>icol</i> was 0 and the driver does not support bookmarks.</p> <p>The driver only supports ODBC 1.0 and the argument <i>fCType</i> was one of the following:</p> <p>SQL_C_STINYINT SQL_C_UTINYINT SQL_C_SSHORT SQL_C_USHORT SQL_C_SLONG SQL_C_ULONG</p>
S1C00	Driver not capable	<p>The driver does not support the data type specified in the argument <i>fCType</i>.</p> <p>The argument <i>icol</i> was 0 and the driver does not support bookmarks.</p> <p>The driver only supports ODBC 1.0 and the argument <i>fCType</i> was one of the following:</p> <p>SQL_C_STINYINT SQL_C_UTINYINT SQL_C_SSHORT SQL_C_USHORT SQL_C_SLONG SQL_C_ULONG</p>

Comments

The ODBC interface provides two ways to retrieve a column of data:

- **SQLBindCol** assigns the storage location for a column of data before the data is retrieved. When **SQLFetch** or **SQLExtendedFetch** is called, the driver places the data for all bound columns in the assigned locations.
- **SQLGetData** (an extended function) assigns a storage location for a column of data after **SQLFetch** or **SQLExtendedFetch** has been called. It also places the data for the requested column in the assigned location. Because it can retrieve data from a column in parts, **SQLGetData** can be used to retrieve long data values.

An application may choose to bind every column with **SQLBindCol**, to do no binding and retrieve data only with **SQLGetData**, or to use a combination of the two. However, unless the driver provides extended

functionality, **SQLGetData** can only be used to retrieve data from columns that occur after the last bound column.

An application calls **SQLBindCol** to pass the pointer to the storage buffer for a column of data to the driver and to specify how or if the data will be converted. It is the application's responsibility to allocate enough storage for the data. If the buffer will contain variable length data, the application must allocate as much storage as the maximum length of the bound column or the data may be truncated. For a list of valid data conversion types, see "Converting Data from SQL to C Data Types" in Appendix D, "Data Types."

At fetch time, the driver processes the data for each bound column according to the arguments specified in **SQLBindCol**. First, it converts the data according to the argument *fCType*. Next, it fills the buffer pointed to by *rgbValue*. Finally, it stores the available number of bytes in *pcbValue*; this is the number of bytes available prior to calling **SQLFetch** or **SQLExtendedFetch**.

- If **SQL_MAX_LENGTH** has been specified with **SQLSetStmtOption** and the available number of bytes is greater than **SQL_MAX_LENGTH**, the driver stores **SQL_MAX_LENGTH** in *pcbValue*.
- If the data is truncated because of **SQL_MAX_LENGTH**, but the user's buffer was large enough for **SQL_MAX_LENGTH** bytes of data, **SQL_SUCCESS** is returned.

Note: The **SQL_MAX_LENGTH** statement option is intended to reduce network traffic and may not be supported by all drivers. To guarantee that data is truncated, an application should allocate a buffer of the desired size and specify this size in the *cbValueMax* argument.

- If the user's buffer causes the truncation, the driver returns **SQL_SUCCESS_WITH_INFO** and **SQLSTATE 01004** (Data truncated) for the fetch function.
- If the data value for a column is **NULL**, the driver sets *pcbValue* to **SQL_NULL_DATA**.
- If the number of bytes available to return cannot be determined in advance, the driver sets *pcbValue* to **SQL_NO_TOTAL**.

When an application uses **SQLExtendedFetch** to retrieve more than one row of data, it only needs to call **SQLBindCol** once for each column of the result set (just as when it binds a column in order to retrieve a single **SQLFetch**

row of data with **SQLFetch**). The **SQLExtendedFetch** function coordinates the placement of each row of data into subsequent locations in the rowset buffers. For additional information about binding rowset buffers, see the "Comments" topic for **SQLExtendedFetch**.

An application can call **SQLBindCol** to bind a column to a new storage location, regardless of whether data has already been fetched. The new binding replaces the old binding. Note that the new binding does not apply to data already fetched; the next time data is fetched, the data will be placed in the new storage location.

To unbind a single bound column, an application calls **SQLBindCol** and specifies a null pointer for *rgbValue*; if *rgbValue* is a null pointer and the column is not bound, **SQLBindCol** returns SQL_SUCCESS. To unbind all bound columns, an application calls **SQLFreeStmt** with the SQL_UNBIND option.

Code Example

In the following example, an application executes a **SELECT** statement to return a result set of the employee names, ages, and birthdays, which is sorted by birthday. It then calls **SQLBindCol** to bind the columns of data to local storage locations. Finally, the application fetches each row of data with **SQLFetch** and prints each employee's name, age, and birthday.

For more code examples, see **SQLColumns**, **SQLExtendedFetch**, and **SQLSetPos**.

```
#define NAME_LEN 30
#define BDAY_LEN 11

UCHAR szName[NAME_LEN], szBirthday[BDAY_LEN];
WORD sAge;
SDWORD cbName, cbAge, cbBirthday;

retcode = SQLExecDirect(hstmt,
    "SELECT NAME, AGE, BIRTHDAY FROM EMPLOYEE ORDER BY 3, 2, 1",
    SQL_NTS);
if (retcode == SQL_SUCCESS) {

    /* Bind columns 1, 2, and 3 */

    SQLBindCol(hstmt, 1, SQL_C_CHAR, szName, NAME_LEN, &cbName);
    SQLBindCol(hstmt, 2, SQL_C_SHORT, &sAge, 0, &cbAge);
    SQLBindCol(hstmt, 3, SQL_C_CHAR, szBirthday, BDAY_LEN, &cbBirthday);

    /* Fetch and print each row of data. On */
```

```
/* an error, display a message and exit. */

while (TRUE) {
    retcode = SQLFetch(hstmt);
    if (retcode == SQL_ERROR || retcode == SQL_SUCCESS_WITH_INFO) {
        show_error();
    }
    if (retcode == SQL_SUCCESS || retcode == SQL_SUCCESS_WITH_INFO){
        fprintf(out, "%-*s %-2d %-*s", NAME_LEN-1, szName,
            sAge, BDAY_LEN-1, szBirthday);
    } else {
        break;
    }
}
}
```

Related Functions For information about**See**

Returning information about a column in a result set	SQLDescribeCol
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
Fetching a row of data	SQLFetch
Freeing a statement handle	SQLFreeStmt
Fetching part or all of a column of data	SQLGetData (extension)
Returning the number of result set columns	SQLNumResultCols

SQLBindParameter



Level 1

SQLBindParameter binds a buffer to a parameter marker in an SQL statement.

Note: This function replaces the ODBC 1.0 function **SQLSetParam**. For more information, see "Comments."

Syntax

RETCODE **SQLBindParameter**(*hstmt*, *ipar*, *fParamType*, *fCType*, *fSqlType*, *cbColDef*, *ibScale*, *rgbValue*, *cbValueMax*, *pcbValue*)

The **SQLBindParameter** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UWORD	<i>ipar</i>	Input	Parameter number, ordered sequentially left to right, starting at 1.
SWORD	<i>fParamType</i>	Input	The type of the parameter. For more information, see "fParamType Argument" in "Comments."
SWORD	<i>fCType</i>	Input	The C data type of the parameter. For more information, see "fCType Argument" in "Comments."
SWORD	<i>fSqlType</i>	Input	The SQL data type of the parameter. For more information, see "fSqlType Argument" in "Comments."
UDWORD	<i>cbColDef</i>	Input	The precision of the column or expression of the corresponding parameter marker. For more information, see "cbColDef Argument" in "Comments."
SWORD	<i>ibScale</i>	Input	The scale of the column or expression of the corresponding parameter marker. For further information concerning scale, see "Precision, Scale, Length, and Display Size," in Appendix D, "Data Types."
PTR	<i>rgbValue</i>	Input/ Output	A pointer to a buffer for the parameter's data.

Type	Argument	Use	Description
SDWORD	<i>cbValueMax</i>	Input	Maximum length of the <i>rgbValue</i> buffer. For more information, see "cbValueMax Argument" in "Comments."
SDWORD FAR *	<i>pcbValue</i>	Input/ Output	A pointer to a buffer for the parameter's length. For more information, see "pcbValue Argument" in "Comments."

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLBindParameter** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLBindParameter** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATEs returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
07006	Restricted data type attribute violation	The data value identified by the <i>fCType</i> argument cannot be converted to the data type identified by the <i>fSqlType</i> argument.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.

SQLForeignKeys

SQLForeignKeys

S1003	Program type out of range	(DM) The value specified by the argument <i>fCType</i> was not a valid data type or SQL_C_DEFAULT.
SQLSTATE	Error	Description
S1004	SQL data type out of range	(DM) The value specified for the argument <i>fSqlType</i> was in the block of numbers reserved for ODBC SQL data type indicators but was not a valid ODBC SQL data type indicator.
S1009	Invalid argument value	(DM) The argument <i>rgbValue</i> was a null pointer, the argument <i>pcbValue</i> was a null pointer, and the argument <i>fParamType</i> was not SQL_PARAM_OUTPUT.
S1010	Function sequence error	(DM) An asynchronously executing function was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1090	Invalid string or buffer length	(DM) The value specified for the argument <i>cbValueMax</i> was less than 0.
S1093	Invalid parameter number	(DM) The value specified for the argument <i>ipar</i> was less than 1. The value specified for the argument <i>ipar</i> was greater than the maximum number of parameters supported by the data source.
S1094	Invalid scale value	The value specified for the argument <i>ibScale</i> was outside the range of values supported by the data source for a column of the SQL data type specified by the <i>fSqlType</i> argument.
S1104	Invalid precision value	The value specified for the argument <i>cbColDef</i> was outside the range of values supported by the data source for a column of the SQL data type specified by the <i>fSqlType</i> argument.
S1105	Invalid parameter type	(DM) The value specified for the argument <i>fParamType</i> was invalid (see

"Comments").

The value specified for the argument *fParamType* was SQL_PARAM_OUTPUT and the parameter did not mark a return value from a procedure or a procedure parameter.

The value specified for the argument *fParamType* was SQL_PARAM_INPUT and the parameter marked the return value from a procedure.

SQLSTATE Error		Description
S1C00	Driver not capable	<p>The driver or data source does not support the conversion specified by the combination of the value specified for the argument <i>fCType</i> and the driver-specific value specified for the argument <i>fSqlType</i>. The value specified for the argument <i>fSqlType</i> was a valid ODBC SQL data type indicator for the version of ODBC supported by the driver, but was not supported by the driver or data source.</p> <p>The value specified for the argument <i>fSqlType</i> was in the range of numbers reserved for driver-specific SQL data type indicators, but was not supported by the driver or data source. The driver only supports ODBC 1.0 and the argument <i>fCType</i> was one of the following:</p> <p>SQL_C_STINYINT SQL_C_UTINYINT SQL_C_SSHORT SQL_C_USHORT SQL_C_SLONG SQL_C_ULONG</p>

Comments

An application calls **SQLBindParameter** to bind each parameter marker in an SQL statement. Bindings remain in effect until the application calls **SQLBindParameter** again or until the application calls **SQLFreeStmt** with the SQL_DROP or SQL_RESET_PARAMS option. For more information concerning parameter data types and parameter markers, see "Parameter Data Types" and "Parameter Markers" in Appendix C, "SQL Grammar."

fParamType Argument

The *fParamType* argument specifies the type of the parameter. All parameters in SQL statements that do not call procedures, such as **INSERT** statements, are input parameters. Parameters in procedure calls can be input, input/output, or output parameters. (An application calls **SQLProcedureColumns** to determine the type of a parameter in a procedure call; parameters in procedure calls whose type cannot be determined are assumed to be input parameters.)

The *fParamType* argument is one of the following values:

- **SQL_PARAM_INPUT**. The parameter marks a parameter in an SQL statement that does not call a procedure, such as an **INSERT** statement, or it marks an input parameter in a procedure; these are collectively known as *input parameters*. For example, the parameters in **INSERT INTO Employee VALUES (?, ?, ?)** and **{call AddEmp(?, ?, ?)}** are input parameters.

When the statement is executed, the driver sends data for the parameter to the data source; the *rgbValue* buffer must contain a valid input value or the *pcbValue* buffer must contain **SQL_NULL_DATA**, **SQL_DATA_AT_EXEC**, or the result of the **SQL_LEN_DATA_AT_EXEC** macro.

If an application cannot determine the type of a parameter in a procedure call, it sets *fParamType* to **SQL_PARAM_INPUT**; if the data source returns a value for the parameter, the driver discards it.

- **SQL_PARAM_INPUT_OUTPUT**. The parameter marks an input/output parameter in a procedure. For example, the parameter in **{call GetEmpDept(?)}** is an input/output parameter that accepts an employee's name and returns the name of the employee's department.

When the statement is executed, the driver sends data for the parameter to the data source; the *rgbValue* buffer must contain a valid input value or the *pcbValue* buffer must contain **SQL_NULL_DATA**, **SQL_DATA_AT_EXEC**, or the result of the **SQL_LEN_DATA_AT_EXEC** macro. After the statement is executed, the driver returns data for the parameter to the application; if the data source does not return a value for an input/output parameter, the driver sets the *pcbValue* buffer to **SQL_NULL_DATA**.

Note: When an ODBC 1.0 application calls **SQLSetParam** in an ODBC 2.0 driver, the Driver Manager converts this to a call to **SQLBindParameter** in which the *fParamType* argument is set to **SQL_PARAM_INPUT_OUTPUT**.

- **SQL_PARAM_OUTPUT**. The parameter marks the return value of a procedure or an output parameter in a procedure; these are collectively known as *output parameters*. For example, the parameter in **{?=call GetNextEmpID}** is an output parameter that returns the next employee ID.

After the statement is executed, the driver returns data for the parameter to the application, unless the *rgbValue* and *pcbValue* arguments are both null pointers, in which case the driver discards the output value. If the data source does not return a value for an output parameter, the driver sets the *pcbValue* buffer to **SQL_NULL_DATA**.

fCType Argument

The C data type of the parameter. This must be one of the following values:

SQL_C_BINARY	SQL_C_SSHORT
SQL_C_BIT	SQL_C_STINYINT
SQL_C_CHAR	SQL_C_TIME
SQL_C_DATE	SQL_C_TIMESTAMP
SQL_C_DEFAULT	SQL_C_ULONG
SQL_C_DOUBLE	SQL_C_USHORT
SQL_C_FLOAT	SQL_C_UTINYINT
SQL_C_SLONG	

SQL_C_DEFAULT specifies that the parameter value be transferred from the default C data type for the SQL data type specified with *fSqlType*. For more information, see "Default C Data Types" and "Converting Data from C to SQL Data Types" and "Converting Data from SQL to C Data Types" in Appendix D, "Data Types."

Note: Drivers must also support the following values of *fCType* from ODBC 1.0. Applications must use these values, instead of the ODBC 2.0 values, when calling an ODBC 1.0 driver:

SQL_C_LONG
SQL_C_SHORT
SQL_C_TINYINT

For more information, see "ODBC 1.0 C Data Types" in Appendix D, "Data Types."

fSqlType Argument

This must be one of the following values:

SQL_BIGINT	SQL_DECIMAL
SQL_BINARY	SQL_DOUBLE
SQL_BIT	SQL_FLOAT
SQL_CHAR	SQL_INTEGER
SQL_DATE	SQL_LONGVARBINARY
SQL_LONGVARCHAR	SQL_TIMESTAMP
SQL_NUMERIC	SQL_TINYINT

SQL_REAL	SQL_VARBINARY
SQL_SMALLINT	SQL_VARCHAR
SQL_TIME	

or a driver-specific value. Values greater than SQL_TYPE_DRIVER_START are reserved by ODBC; values less than or equal to SQL_TYPE_DRIVER_START are driver-specific.

For information about how data is converted, see "Converting Data from C to SQL Data Types" and "Converting Data from SQL to C Data Types" in Appendix D, "Data Types."

cbColDef Argument

The *cbColDef* argument specifies the precision of the column or expression corresponding to the parameter marker, unless all of the following are true:

- An ODBC 2.0 application calls **SQLBindParameter** in an ODBC 1.0 driver or an ODBC 1.0 application calls **SQLSetParam** in an ODBC 2.0 driver. (Note that the Driver Manager converts these calls.)
- The *fSqlType* argument is SQL_LONGVARBINARY or SQL_LONGVARCHAR.
- The data for the parameter will be sent with **SQLPutData**.

In this case, the *cbColDef* argument contains the total number of bytes that will be sent for the parameter. For more information, see "Passing Parameter Values" and SQL_DATA_AT_EXEC in "pcbValue Argument."

rgbValue Argument

The *rgbValue* argument points to a buffer that, when **SQLExecute** or **SQLExecDirect** is called, contains the actual data for the parameter. The data must be in the form specified by the *fCType* argument.

If *rgbValue* points to a character string that contains a literal quote character ('), the driver ensures that each literal quote is translated into the form required by the data source. For example, if the data source required that embedded literal quotes be doubled, the driver would replace each quote character (') with two quote characters (' '). If *pcbValue* is the result of the SQL_LEN_DATA_AT_EXEC(*length*) macro or SQL_DATA_AT_EXEC, then *rgbValue* is an application-defined 32-bit value that is associated with the parameter. It is returned to the application through **SQLParamData**. For example, *rgbValue* might be a token such as a parameter number, a pointer to data, or a pointer to a

SQLForeignKeys

structure that the application used to bind input parameters. Note, however, that if the parameter is an input/output parameter, *rgbValue* must be a pointer to a buffer where the output value will be stored. If **SQLParamOptions** was called to specify multiple values for the parameter, the application can use the value of the *pirow* argument in **SQLParamOptions** in conjunction with the *rgbValue*. For example, *rgbValue* might point to an array of values and the application might use *pirow* to retrieve the correct value from the array. For more information, see "Passing Parameter Values."

If the *fParamType* argument is SQL_PARAM_INPUT_OUTPUT or SQL_PARAM_OUTPUT, *rgbValue* points to a buffer in which the driver returns the output value. If the procedure returns one or more result sets, the *rgbValue* buffer is not guaranteed to be set until all results have been fetched. (If *fParamType* is SQL_PARAM_OUTPUT and *rgbValue* and *pcbValue* are both null pointers, the driver discards the output value.)

If the application calls **SQLParamOptions** to specify multiple values for each parameter, *rgbValue* points to an array. A single SQL statement processes the entire array of input values for an input or input/output parameter and returns an array of output values for an input/output or output parameter.

cbValueMax Argument

For character and binary C data, the *cbValueMax* argument specifies the length of the *rgbValue* buffer (if it is a single element) or the length of an element in the *rgbValue* array (if the application calls **SQLParamOptions** to specify multiple values for each parameter). If the application specifies multiple values, *cbValueMax* is used to determine the location of values in the *rgbValue* array, both on input and on output. For input/output and output parameters, it is used to determine whether to truncate character and binary C data on output:

- For character C data, if the number of bytes available to return is greater than or equal to *cbValueMax*, the data in *rgbValue* is truncated to *cbValueMax* – 1 bytes and is null-terminated by the driver.
- For binary C data, if the number of bytes available to return is greater than *cbValueMax*, the data in *rgbValue* is truncated to *cbValueMax* bytes.

For all other types of C data, the *cbValueMax* argument is ignored. The length of the *rgbValue* buffer (if it is a single element) or the length of an element in the *rgbValue* array (if the application calls **SQLParamOptions**

to specify multiple values for each parameter) is assumed to be the length of the C data type.

Note: When an ODBC 1.0 application calls **SQLSetParam** in an ODBC 2.0 driver, the Driver Manager converts this to a call to **SQLBindParameter** in which the *cbValueMax* argument is always `SQL_SETPARAM_VALUE_MAX`. Because the Driver Manager returns an error if an ODBC 2.0 application sets *cbValueMax* to `SQL_SETPARAM_VALUE_MAX`, an ODBC 2.0 driver can use this to determine when it is called by an ODBC 1.0 application.

When an ODBC 2.0 application calls **SQLBindParameter** in an ODBC 1.0 driver, the Driver Manager converts this to a call to **SQLSetParam** and discards the *cbValueMax* argument.

In **SQLSetParam**, the way in which an application specifies the length of the *rgbValue* buffer so that the driver can return character or binary data and the way in which an application sends an array of character or binary parameter values to the driver are driver-defined. If an ODBC 2.0 application uses this functionality in an ODBC 1.0 driver, it must use the semantics defined by that driver. If an ODBC 2.0 driver supported this functionality as an ODBC 1.0 driver, it must continue to support this functionality for ODBC 1.0 applications.

pcbValue Argument

The *pcbValue* argument points to a buffer that, when **SQLExecute** or **SQLExecDirect** is called, contains one of the following:

- The length of the parameter value stored in *rgbValue*. This is ignored except for character or binary C data.
- `SQL_NTS`. The parameter value is a null-terminated string.
- `SQL_NULL_DATA`. The parameter value is NULL.
- `SQL_DEFAULT_PARAM`. A procedure is to use the default value of a parameter, rather than a value retrieved from the application. This value is valid only in a procedure call, and then only if the *fParamType* argument is `SQL_PARAM_INPUT` or `SQL_PARAM_INPUT_OUTPUT`. When *pcbValue* is `SQL_DEFAULT_PARAM`, the *fCType*, *fSqlType*, *cbColDef*, *ibScale*, *cbValueMax* and *rgbValue* arguments are ignored for input parameters and are used only to define the output parameter value for input/output parameters.

Note: This value was introduced in ODBC 2.0.

SQLForeignKeys

SQLForeignKeys

- The result of the `SQL_LEN_DATA_AT_EXEC(length)` macro. The data for the parameter will be sent with **SQLPutData**. If the *fSqlType* argument is `SQL_LONGVARBINARY`, `SQL_LONGVARCHAR`, or a long, data source-specific data type and the driver returns "Y" for the `SQL_NEED_LONG_DATA_LEN` information type in **SQLGetInfo**, *length* is the number of bytes of data to be sent for the parameter; otherwise, *length* must be a nonnegative value and is ignored. For more information, see "Passing Parameter Values."

For example, to specify that 10,000 bytes of data will be sent with **SQLPutData** for an SQL_LONGVARCHAR parameter, an application sets *pcbValue* to SQL_LEN_DATA_AT_EXEC(10000).

Note: This macro was introduced in ODBC 2.0.

- SQL_DATA_AT_EXEC. The data for the parameter will be sent with **SQLPutData**. This value is used by ODBC 2.0 applications when calling ODBC 1.0 drivers and by ODBC 1.0 applications when calling ODBC 2.0 drivers. For more information, see "Passing Parameter Values."

If *pcbValue* is a null pointer, the driver assumes that all input parameter values are non-NULL and that character and binary data are null-terminated. If *fParamType* is SQL_PARAM_OUTPUT and *rgbValue* and *pcbValue* are both null pointers, the driver discards the output value.

Note: Application developers are strongly discouraged from specifying a null pointer for *pcbValue* when the data type of the parameter is SQL_C_BINARY. For SQL_C_BINARY data, a driver sends only the data preceding an occurrence of the null-termination character, 0x00. To ensure that a driver does not unexpectedly truncate SQL_C_BINARY data, *pcbValue* should contain a pointer to a valid length value.

If the *fParamType* argument is SQL_PARAM_INPUT_OUTPUT or SQL_PARAM_OUTPUT, *pcbValue* points to a buffer in which the driver returns SQL_NULL_DATA, the number of bytes available to return in *rgbValue* (excluding the null termination byte of character data), or SQL_NO_TOTAL if the number of bytes available to return cannot be determined. If the procedure returns one or more result sets, the *pcbValue* buffer is not guaranteed to be set until all results have been fetched.

If the application calls **SQLParamOptions** to specify multiple values for each parameter, *pcbValue* points to an array of SDWORD values. These can be any of the values listed earlier in this section and are processed with a single SQL statement.

Passing Parameter Values

An application can pass the value for a parameter either in the *rgbValue* buffer or with one or more calls to **SQLPutData**. Parameters whose data is passed with **SQLPutData** are known as *data-at-execution* parameters. These are commonly used to send data for SQL_LONGVARBINARY and

SQLForeignKeys

SQL_LONGVARCHAR parameters and can be mixed with other parameters.

To pass parameter values, an application:

1. Calls **SQLBindParameter** for each parameter to bind buffers for the parameter's value (*rgbValue* argument) and length (*pcbValue* argument). For data-at-execution parameters, *rgbValue* is an application-defined 32-bit value such as a parameter number or a pointer to data. The value will be returned later and can be used to identify the parameter.
2. Places values for input and input/output parameters in the *rgbValue* and *pcbValue* buffers:
 - For normal parameters, the application places the parameter value in the *rgbValue* buffer and the length of that value in the *pcbValue* buffer.
 - For data-at-execution parameters, the application places the result of the `SQL_LEN_DATA_AT_EXEC(length)` macro (when calling an ODBC 2.0 driver) or `SQL_DATA_AT_EXEC` (when calling an ODBC 1.0 driver) in the *pcbValue* buffer.
3. Calls `SQLExecute` or `SQLExecDirect` to execute the SQL statement.
 - If there are no data-at-execution parameters, the process is complete.
 - If there are any data-at-execution parameters, the function returns `SQL_NEED_DATA`.
4. Calls **SQLParamData** to retrieve the application-defined value specified in the *rgbValue* argument for the first data-at-execution parameter to be processed.

Note: Although data-at-execution parameters are similar to data-at-execution columns, the value returned by `SQLParamData` is different for each.

Data-at-execution parameters are parameters in an SQL statement for which data will be sent with `SQLPutData` when the statement is executed with `SQLExecDirect` or `SQLExecute`. They are bound with `SQLBindParameter`. The value returned by `SQLParamData` is a 32-bit value passed to `SQLBindParameter` in the *rgbValue* argument.

Data-at-execution columns are columns in a rowset for which data will be sent with `SQLPutData` when a row is updated or added with `SQLSetPos`. They are bound with `SQLBindCol`. The value returned by `SQLParamData` is the address of the row in the *rgbValue* buffer that is being processed.

5. Calls **SQLPutData** one or more times to send data for the parameter. More than one call is needed if the data value is larger than the *rgbValue* buffer specified in **SQLPutData**; note that multiple calls to **SQLPutData** for the same parameter are allowed only when sending character C data to a column with a character, binary, or data source-specific data type or when sending binary C data to a column with a character, binary, or data source-specific data type.
6. Calls **SQLParamData** again to signal that all data has been sent for the parameter.
 - If there are more data-at-execution parameters, **SQLParamData** returns **SQL_NEED_DATA** and the application-defined value for the next data-at-execution parameter to be processed. The application repeats steps 5 and 6.
 - If there are no more data-at-execution parameters, the process is complete. If the statement was successfully executed, **SQLParamData** returns **SQL_SUCCESS** or **SQL_SUCCESS_WITH_INFO**; if the execution failed, it returns **SQL_ERROR**. At this point, **SQLParamData** can return any **SQLSTATE** that can be returned by the function used to execute the statement (**SQLExecDirect** or **SQLExecute**).

Output values for any input/output or output parameters will be available in the *rgbValue* and *pcbValue* buffers after the application retrieves any result sets generated by the statement.

After **SQLExecute** or **SQLExecDirect** returns **SQL_NEED_DATA**, and before data is sent for all data-at-execution parameters, the statement is canceled, or an error occurs in **SQLParamData** or **SQLPutData**, the application can only call **SQLCancel**, **SQLGetFunctions**, **SQLParamData**, or **SQLPutData** with the *hstmt* or the *hdbc* associated with the *hstmt*. If it calls any other function with the *hstmt* or the *hdbc* associated with the *hstmt*, the function returns **SQL_ERROR** and **SQLSTATE** S1010 (Function sequence error).

If the application calls **SQLCancel** while the driver still needs data for data-at-execution parameters, the driver cancels statement execution; the application can then call **SQLExecute** or **SQLExecDirect** again. If the application calls **SQLParamData** or **SQLPutData** after canceling the statement, the function returns **SQL_ERROR** and **SQLSTATE** S1008 (Operation canceled).

Conversion of Calls to and from **SQLSetParam**

When an ODBC 1.0 application calls **SQLSetParam** in an ODBC 2.0 driver, the ODBC 2.0 Driver Manager maps the call as follows:

Call by ODBC 1.0 Application

Call to ODBC 2.0 Driver

SQLForeignKeys

SQLForeignKeys

```
SQLSetParam(  
    hstmt, ipar,  
    fCType, fSqlType, cbColDef, ibScale,  
    rgbValue,  
    pcbValue);
```

```
SQLBindParameter(  
    hstmt, ipar, SQL_PARAM_INPUT_OUTPUT,  
    fCType, fSqlType, cbColDef, ibScale,  
    rgbValue, SQL_SETPARAM_VALUE_MAX,  
    pcbValue);
```

When an ODBC 2.0 application calls **SQLBindParameter** in an ODBC 1.0 driver, the ODBC 2.0 Driver Manager maps the calls as follows:

Call by ODBC 2.0 Application	Call to ODBC 1.0 Driver
SQLBindParameter(hstmt, ipar, fParamType, fCType, fSqlType, cbColDef, ibScale, rgbValue, cbValueMax, pcbValue);	SQLSetParam(hstmt, ipar, fCType, fSqlType, cbColDef, ibScale, rgbValue, pcbValue);

Code Example

In the following example, an application prepares an SQL statement to insert data into the EMPLOYEE table. The SQL statement contains parameters for the NAME, AGE, and BIRTHDAY columns. For each parameter in the statement, the application calls **SQLBindParameter** to specify the ODBC C data type and the SQL data type of the parameter and to bind a buffer to each parameter. For each row of data, the application assigns data values to each parameter and calls **SQLExecute** to execute the statement.

For more code examples, see **SQLParamOptions**, **SQLProcedures**, **SQLPutData**, and **SQLSetPos**.

```
#define NAME_LEN 30  
  
UCHAR    szName[NAME_LEN];  
SWORD    sAge;  
SDWORD    cbName = SQL_NTS, cbAge = 0, cbBirthday = 0;  
DATE_STRUCT dsBirthday;  
  
retcode = SQLPrepare(hstmt,  
    "INSERT INTO EMPLOYEE (NAME, AGE, BIRTHDAY) VALUES (?, ?, ?)",  
    SQL_NTS);  
  
if (retcode == SQL_SUCCESS) {  
  
    /* Specify data types and buffers. */  
    /* for Name, Age, Birthday parameter data. */  
  
    SQLBindParameter(hstmt, 1, SQL_PARAM_INPUT, SQL_C_CHAR,  
        SQL_CHAR, NAME_LEN, 0, szName, 0, &cbName);  
    SQLBindParameter(hstmt, 2, SQL_PARAM_INPUT, SQL_C_SSHORT,  
        SQL_SMALLINT, 0, 0, &sAge, 0, &cbAge);
```

```
SQLBindParameter(hstmt, 3, SQL_PARAM_INPUT, SQL_C_DATE,
                 SQL_DATE, 0, 0, &dsBirthday, 0, &cbBirthday);

strcpy(szName, "Smith, John D."); /* Specify first row of */
sAge = 40; /* parameter data */
dsBirthday.year = 1952;
dsBirthday.month = 2;
dsBirthday.day = 29;
retcode = SQLExecute(hstmt); /* Execute statement with */
/* first row */
strcpy(szName, "Jones, Bob K."); /* Specify second row of */
sAge = 52; /* parameter data */
dsBirthday.year = 1940;
dsBirthday.month = 3;
dsBirthday.day = 31;
SQLExecute(hstmt); /* Execute statement with */
/* second row */
}
```

**Related
Functions****For information about****See**

Returning information about a parameter in a statement	SQLDescribeParam (extension)
Executing an SQL statement	SQLExecDirect
Executing a prepared SQL statement	SQLExecute
Returning the number of statement parameters	SQLNumParams (extension)
Returning the next parameter to send data for	SQLParamData (extension)
Specifying multiple parameter values	SQLParamOptions (extension)
Sending parameter data at execution time	SQLPutData (extension)

SQLBrowseConnect



Extension Level 2 **SQLBrowseConnect** supports an iterative method of discovering and enumerating the attributes and attribute values required to connect to a data source. Each call to **SQLBrowseConnect** returns successive levels of attributes and attribute values. When all levels have been enumerated, a connection to the data source is completed and a complete connection string is returned by **SQLBrowseConnect**. A return code of **SQL_SUCCESS** or **SQL_SUCCESS_WITH_INFO** indicates that all connection information has been specified and the application is now connected to the data source.

Syntax

RETCODE **SQLBrowseConnect**(*hdbc*, *szConnStrIn*, *cbConnStrIn*, *szConnStrOut*, *cbConnStrOutMax*, *pcbConnStrOut*)

The **SQLBrowseConnect** function accepts the following arguments:

Type	Argument	Use	Description
HDBC	<i>hdbc</i>	Input	Connection handle.
UCHAR FAR *	<i>szConnStrIn</i>	Input	Browse request connection string (see "szConnStrIn Argument" in "Comments").
SWORD	<i>cbConnStrIn</i>	Input	Length of <i>szConnStrIn</i> .
UCHAR FAR *	<i>szConnStrOut</i>	Output	Pointer to storage for the browse result connection string (see "szConnStrOut Argument" in "Comments").
SWORD	<i>cbConnStrOutMax</i>	Input	Maximum length of the <i>szConnStrOut</i> buffer.
SWORD FAR *	<i>pcbConnStrOut</i>	Output	The total number of bytes (excluding the null termination byte) available to return in <i>szConnStrOut</i> . If the number of bytes available to return is greater than or equal to <i>cbConnStrOutMax</i> , the connection string in <i>szConnStrOut</i> is truncated to <i>cbConnStrOutMax</i> – 1 bytes.

Returns

SQL_SUCCESS, **SQL_SUCCESS_WITH_INFO**, **SQL_NEED_DATA**, **SQL_ERROR**, or **SQL_INVALID_HANDLE**.

Diagnostics

When **SQLBrowseConnect** returns **SQL_ERROR**, **SQL_SUCCESS_WITH_INFO**, or **SQL_NEED_DATA**, an associated **SQLSTATE** value may be obtained by calling **SQLError**. The following

table lists the SQLSTATE values commonly returned by
SQLBrowseConnect and

SQLFreeConnect

explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATEs returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	The buffer <i>szConnStrOut</i> was not large enough to return entire browse result connection string, so the string was truncated. The argument <i>pcbConnStrOut</i> contains the length of the untruncated browse result connection string. (Function returns SQL_SUCCESS_WITH_INFO.)
01S00	Invalid connection string attribute	An invalid attribute keyword was specified in the browse request connection string (<i>szConnStrIn</i>). (Function returns SQL_NEED_DATA.) An attribute keyword was specified in the browse request connection string (<i>szConnStrIn</i>) that does not apply to the current connection level. (Function returns SQL_NEED_DATA.)
08001	Unable to connect to data source	The driver was unable to establish a connection with the data source.
08002	Connection in use	(DM) The specified <i>hdbc</i> had already been used to establish a connection with a data source and the connection was open.
08004	Data source rejected establishment of connection	The data source rejected the establishment of the connection for implementation defined reasons.
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was attempting to connect failed before the function completed processing.
28000	Invalid authorization specification	Either the user identifier or the authorization string or both as specified in the browse request connection string (<i>szConnStrIn</i>) violated restrictions defined by the data source.
SQLSTATE	Error	Description

IM001	Driver does not support this function	(DM) The driver corresponding to the specified data source name does not support the function.
IM002	Data source not found and no default driver specified	(DM) The data source name specified in the browse request connection string (<i>szConnStrIn</i>) was not found in the ODBC.INI file or registry nor was there a default driver specification. (DM) The ODBC.INI file could not be found.
IM003	Specified driver could not be loaded	(DM) The driver listed in the data source specification in the ODBC.INI file or registry, or specified by the DRIVER keyword was not found or could not be loaded for some other reason.
IM004	Driver's SQLAllocEnv failed	(DM) During SQLBrowseConnect , the Driver Manager called the driver's SQLAllocEnv function and the driver returned an error.
IM005	Driver's SQLAllocConnect failed	(DM) During SQLBrowseConnect , the Driver Manager called the driver's SQLAllocConnect function and the driver returned an error.
IM006	Driver's SQLSetConnectOption failed	(DM) During SQLBrowseConnect , the Driver Manager called the driver's SQLSetConnectOption function and the driver returned an error.
IM009	Unable to load translation DLL	The driver was unable to load the translation DLL that was specified for the data source or for the connection.
IM010	Data source name too long	(DM) The attribute value for the DSN keyword was longer than SQL_MAX_DSN_LENGTH characters.
IM011	Driver name too long	(DM) The attribute value for the DRIVER keyword was longer than 255 characters.
IM012	DRIVER keyword syntax error	(DM) The keyword-value pair for the DRIVER keyword contained a syntax error.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.

SQLSTATE Error

Description

SQLFreeConnect

S1001	Memory allocation failure	<p>(DM) The Driver Manager was unable to allocate memory required to support execution or completion of the function.</p> <p>The driver was unable to allocate memory required to support execution or completion of the function.</p>
S1090	Invalid string or buffer length	<p>(DM) The value specified for argument <i>cbConnStrIn</i> was less than 0 and was not equal to SQL_NTS.</p> <p>(DM) The value specified for argument <i>cbConnStrOutMax</i> was less than 0.</p>
S1T00	Timeout expired	<p>The timeout period expired before the connection to the data source completed. The timeout period is set through SQLSetConnectOption, SQL_LOGIN_TIMEOUT.</p>

szConnStrIn Argument

```

connection-string ::= attribute[:] | attribute; connection-string
attribute ::= attribute-keyword=attribute-value | DRIVER={attribute-value}
(The braces are literal; the application must specify them.)
attribute-keyword ::= DSN | UID | PWD
                        | driver-defined-attribute-keyword
attribute-value ::= character-string
driver-defined-attribute-keyword ::= identifier

```

Note: The **DRIVER** keyword was introduced in ODBC 2.0 and is not supported by ODBC 1.0 drivers.

ADABAS D

browse request connection string, the Driver Manager and driver use whichever keyword appears first.

szConnStrOut Argument

The browse result connection string is a list of connection attributes. A connection attribute consists of an attribute keyword and a corresponding attribute value. The browse result connection string has the following Syntax

```
connection-string ::= attribute[;] | attribute; connection-string
attribute ::= [*]attribute-keyword=attribute-value
attribute-keyword ::= ODBC-attribute-keyword
                    | driver-defined-attribute-keyword
ODBC-attribute-keyword = {UID | PWD}[;localized-identifier]
driver-defined-attribute-keyword ::= identifier[;localized-identifier]
attribute-value ::= {attribute-value-list} | ?
(The braces are literal; they are returned by the driver.)
attribute-value-list ::= character-string | character-string, attribute-value-list
```

where *character-string* has zero or more characters; *identifier* and *localized-identifier* have one or more characters; *attribute-keyword* is case insensitive; and *attribute-value* may be case sensitive. Because of connection string and initialization file grammar, keywords, localized identifiers, and attribute values that contain the characters `[]{};,?*=!@` should be avoided. Because of the registry grammar, keywords and data source names cannot contain the backslash (\) character.

The browse result connection string syntax is used according to the following semantic rules:

- If an asterisk (*) precedes an *attribute-keyword*, the *attribute* is optional, and may be omitted in the next call to **SQLBrowseConnect**.
- The attribute keywords **UID** and **PWD** have the same meaning as defined in **SQLDriverConnect**.
- A *driver-defined-attribute-keyword* names the kind of attribute for which an attribute value may be supplied. For example, it might be **SERVER**, **DATABASE**, **HOST**, or **DBMS**.
- *ODBC-attribute-keywords* and *driver-defined-attribute-keywords* include a localized or user-friendly version of the keyword. This might be used by applications as a label in a dialog box. However, **UID**, **PWD**, or the *identifier* alone must be used when passing a browse request string to the driver.

- The *{attribute-value-list}* is an enumeration of actual values valid for the corresponding *attribute-keyword*. Note that the braces ({}) do not indicate a list of choices; they are returned by the driver. For example, it might be a list of server names or a list of database names.
- If the *attribute-value* is a single question mark (?), a single value corresponds to the *attribute-keyword*. For example, UID=JohnS; PWD=Sesame.
- Each call to **SQLBrowseConnect** returns only the information required to satisfy the next level of the connection process. The driver associates state information with the connection handle so that the context can always be determined on each call.

Using SQLBrowseConnect

SQLBrowseConnect requires an allocated *hdbc*. The Driver Manager loads the driver that was specified in or that corresponds to the data source name specified in the initial browse request connection string; for information on when this occurs, see the "Comments" section in **SQLConnect**. It may establish a connection with the data source during the browsing process. If **SQLBrowseConnect** returns **SQL_ERROR**, outstanding connections are terminated and the *hdbc* is returned to an unconnected state.

When **SQLBrowseConnect** is called for the first time on an *hdbc*, the browse request connection string must contain the **DSN** keyword or the **DRIVER** keyword. If the browse request connection string contains the **DSN** keyword, the Driver Manager locates a corresponding data source specification in the ODBC.INI file or registry:

- If the Driver Manager finds the corresponding data source specification, it loads the associated driver DLL; the driver can retrieve information about the data source from the ODBC.INI file or registry.
- If the Driver Manager cannot find the corresponding data source specification, it locates the default data source specification and loads the associated driver DLL; the driver can retrieve information about the default data source from the ODBC.INI file or registry.
- If the Driver Manager cannot find the corresponding data source specification and there is no default data source specification, it returns **SQL_ERROR** with **SQLSTATE** IM002 (Data source not found and no default driver specified).

If the browse request connection string contains the **DRIVER** keyword, the Driver Manager loads the specified driver; it does not attempt to

locate a data source in the ODBC.INI file or registry. Because the **DRIVER** keyword does not use information from the ODBC.INI file or registry, the driver must define enough keywords so that a driver can connect to a data source using only the information in the browse request connection strings.

On each call to **SQLBrowseConnect**, the application specifies the connection attribute values in the browse request connection string. The driver returns successive levels of attributes and attribute values in the browse result connection string; it returns **SQL_NEED_DATA** as long as there are connection attributes that have not yet been enumerated in the browse request connection string. The application uses the contents of the browse result connection string to build the browse request connection string for the next call to **SQLBrowseConnect**. Note that the application cannot use the contents of previous browse result connection strings when building the current browse request connection string; that is, it cannot specify different values for attributes set in previous levels.

When all levels of connection and their associated attributes have been enumerated, the driver returns **SQL_SUCCESS**, the connection to the data source is complete, and a complete connection string is returned to the application. The connection string is suitable to use in conjunction with **SQLDriverConnect** with the **SQL_DRIVER_NOPROMPT** option to establish another connection.

SQLBrowseConnect also returns **SQL_NEED_DATA** if there are recoverable, nonfatal errors during the browse process, for example, an invalid password supplied by the application or an invalid attribute keyword supplied by the application. When **SQL_NEED_DATA** is returned and the browse result connection string is unchanged, an error has occurred and the application must call **SQLError** to return the **SQLSTATE** for browse-time errors. This permits the application to correct the attribute and continue the browse.

An application may terminate the browse process at any time by calling **SQLDisconnect**. The driver will terminate any outstanding connections and return the *hdbc* to an unconnected state.

For more information, see "Connection Browsing With **SQLBrowseConnect**" in Chapter 5, "Establishing Connections."

If a driver supports **SQLBrowseConnect**, the driver keyword section of the ODBC.INF file for the driver must contain the **ConnectFunctions** keyword with the third character set to "Y". For more information, see "Driver Keyword Sections" in Chapter 19, "Installing ODBC Software."

SQLFreeConnect

Code Example In the following example, an application calls **SQLBrowseConnect** repeatedly. Each time **SQLBrowseConnect** returns `SQL_NEED_DATA`, it passes back information about the data it needs in *szConnStrOut*. The application passes *szConnStrOut* to its routine **GetUserInput** (not shown). **GetUserInput** parses the information, builds and displays a dialog box, and returns the information entered by the user in *szConnStrIn*. The application passes the user's information to the driver in the next call to **SQLBrowseConnect**. After the application has provided all necessary information for the driver to connect to the data source, **SQLBrowseConnect** returns `SQL_SUCCESS` and the application proceeds.

For example, to connect to the data source My Source, the following actions might occur. First, the application passes the following string to **SQLBrowseConnect**:

```
"DSN=My Source"
```

The Driver Manager loads the driver associated with the data source My Source. It then calls the driver's **SQLBrowseConnect** function with the same arguments it received from the application. The driver returns the following string in *szConnStrOut*.

```
"HOST:Server={red,blue,green};UID:ID=?;PWD:Password=?"
```

The application passes this string to its **GetUserInput** routine, which builds a dialog box that asks the user to select the red, blue, or green server, and to enter a user ID and password. The routine passes the following user-specified information back in *szConnStrIn*, which the application passes to **SQLBrowseConnect**:

```
"HOST=red;UID=Smith;PWD=Sesame"
```

SQLBrowseConnect uses this information to connect to the red server as Smith with the password Sesame, then returns the following string in *szConnStrOut*:

```
"*DATABASE:Database={master,model,empdata}"
```

The application passes this string to its **GetUserInput** routine, which builds a dialog box that asks the user to select a database. The user selects empdata and the application calls **SQLBrowseConnect** a final time with the string:

```
"DATABASE=empdata"
```

This is the final piece of information the driver needs to connect to the data source; **SQLBrowseConnect** returns `SQL_SUCCESS` and *szConnStrOut* contains the completed connection string:

```
"DSN=My Source;HOST=red;UID=Smith;PWD=Sesame;DATABASE=empdata"
```


SQLFreeConnect

```
#define BRWS_LEN 100
HENV  henv;
HDBC  hdbc;
HSTMT hstmt;
RETCODE retcode;
UCHAR szConnStrIn[BRWS_LEN], szConnStrOut[BRWS_LEN];
SWORD cbConnStrOut;

retcode = SQLAllocEnv(&henv);          /* Environment handle */
if (retcode == SQL_SUCCESS) {
    retcode = SQLAllocConnect(henv, &hdbc); /* Connection handle */
    if (retcode == SQL_SUCCESS) {
        /* Call SQLBrowseConnect until it returns a value other than */
        /* SQL_NEED_DATA (pass the data source name the first time). */
        /* If SQL_NEED_DATA is returned, call GetUserInput (not */
        /* shown) to build a dialog from the values in szConnStrOut. */
        /* The user-supplied values are returned in szConnStrIn, */
        /* which is passed in the next call to SQLBrowseConnect. */

        lstrcpy(szConnStrIn, "DSN=MyServer");
        do {
            retcode = SQLBrowseConnect(hstmt, szConnStrIn, SQL_NTS,
                                       szConnStrOut, BRWS_LEN, &cbConnStrOut)
            if (retcode == SQL_NEED_DATA)
                GetUserInput(szConnStrOut, szConnStrIn);
        } while (retcode == SQL_NEED_DATA);

        if (retcode == SQL_SUCCESS || retcode == SQL_SUCCESS_WITH_INFO){

            /* Process data after successful connection */

            retcode = SQLAllocStmt(hdbc, &hstmt);
            if (retcode == SQL_SUCCESS) {
                ...;
                ...;
                ...;
                SQLFreeStmt(hstmt, SQL_DROP);
            }
            SQLDisconnect(hdbc);
        }
    }
    SQLFreeConnect(hdbc);
}
SQLFreeEnv(henv);
```

**Related
Functions****For information about****See**

Allocating a connection handle

SQLAllocConnect

Connecting to a data source

SQLConnect

Disconnecting from a data source

SQLDisconnectConnecting to a data source using a
connection string or dialog box**SQLDriverConnect** (extension)

Returning driver descriptions and attributes

SQLDrivers (extension)

Freeing a connection handle

SQLFreeConnect

SQLCancel

ADBC

Core **SQLCancel** cancels the processing on an *hstmt*.

Syntax RETCODE **SQLCancel**(*hstmt*)

The **SQLCancel** function accepts the following argument.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLCancel** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLCancel** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
70100	Operation aborted	The data source was unable to process the cancel request.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.

Comments **SQLCancel** can cancel the following types of processing on an *hstmt*:

- A function running asynchronously on the *hstmt*.

- A function on an *hstmt* that needs data.
- A function running on the *hstmt* on another thread.

If an application calls **SQLCancel** when no processing is being done on the *hstmt*, **SQLCancel** has the same effect as **SQLFreeStmt** with the `SQL_CLOSE` option; this behavior is defined only for completeness and applications should call **SQLFreeStmt** to close cursors.

Canceling Asynchronous Processing

After an application calls a function asynchronously, it calls the function repeatedly to determine whether it has finished processing. If the function is still processing, it returns `SQL_STILL_EXECUTING`. If the function has finished processing, it returns a different code.

After any call to the function that returns `SQL_STILL_EXECUTING`, an application can call **SQLCancel** to cancel the function. If the cancel request is successful, the driver returns `SQL_SUCCESS`. This message does not indicate that the function was actually canceled; it indicates that the cancel request was processed. When or if the function is actually canceled is driver- and data source-dependent. The application must continue to call the original function until the return code is not `SQL_STILL_EXECUTING`. If the function was successfully canceled, the return code is `SQL_ERROR` and `SQLSTATE S1008` (Operation canceled). If the function completed its normal processing, the return code is `SQL_SUCCESS` or `SQL_SUCCESS_WITH_INFO` if the function succeeded or `SQL_ERROR` and a `SQLSTATE` other than `S1008` (Operation canceled) if the function failed.

Canceling Functions that Need Data

After **SQLExecute** or **SQLExecDirect** returns `SQL_NEED_DATA` and before data has been sent for all data-at-execution parameters, an application can call **SQLCancel** to cancel the statement execution. After the statement has been canceled, the application can call **SQLExecute** or **SQLExecDirect** again. For more information, see **SQLBindParameter**.

After **SQLSetPos** returns `SQL_NEED_DATA` and before data has been sent for all data-at-execution columns, an application can call **SQLCancel** to cancel the operation. After the operation has been canceled, the application can call **SQLSetPos** again; canceling does not affect the cursor state or the current cursor position. For more information, see **SQLSetPos**.

Canceling Functions in Multithreaded Applications

In a multithreaded application, the application can cancel a function that is running synchronously on an *hstmt*. To cancel the function, the application calls **SQLCancel** with the same *hstmt* as that used by the target function, but on a different thread. As in canceling a function running asynchronously, the return code of the **SQLCancel** only indicates whether the driver processed the request successfully. The return code of the original function indicates whether it completed normally or was canceled.

Related Functions

For information about

See

Assigning storage for a parameter	SQLBindParameter
Executing an SQL statement	SQLExecDirect
Executing a prepared SQL statement	SQLExecute
Freeing a statement handle	SQLFreeStmt
Positioning the cursor in a rowset	SQLSetPos (extension)
Returning the next parameter to send data for	SQLParamData (extension)
Sending parameter data at execution time	SQLPutData (extension)

SQLColAttributes



Core **SQLColAttributes** returns descriptor information for a column in a result set; it cannot be used to return information about the bookmark column (column 0). Descriptor information is returned as a character string, a 32-bit descriptor-dependent value, or an integer value.

Syntax RETCODE **SQLColAttributes**(*hstmt*, *icol*, *fDescType*, *rgbDesc*, *cbDescMax*, *pcbDesc*, *pfDesc*)

The **SQLColAttributes** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UWORD	<i>icol</i>	Input	Column number of result data, ordered sequentially from left to right, starting at 1. Columns may be described in any order.
UWORD	<i>fDescType</i>	Input	A valid descriptor type (see "Comments").
PTR	<i>rgbDesc</i>	Output	Pointer to storage for the descriptor information. The format of the descriptor information returned depends on the <i>fDescType</i> .
SWORD	<i>cbDescMax</i>	Input	Maximum length of the <i>rgbDesc</i> buffer.
SWORD FAR *	<i>pcbDesc</i>	Output	Total number of bytes (excluding the null termination byte for character data) available to return in <i>rgbDesc</i> . For character data, if the number of bytes available to return is greater than or equal to <i>cbDescMax</i> , the descriptor information in <i>rgbDesc</i> is truncated to <i>cbDescMax</i> – 1 bytes and is null-terminated by the driver. For all other types of data, the value of <i>cbValueMax</i> is ignored and the driver assumes the size of <i>rgbValue</i> is 32 bits.
SDWORD FAR *	<i>pfDesc</i>	Output	Pointer to an integer value to contain descriptor information for numeric descriptor types, such as

SQL_COLUMN_LENGTH.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLColAttributes** returns either SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLColAttributes** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	The buffer <i>rgbDesc</i> was not large enough to return the entire string value, so the string value was truncated. The argument <i>pcbDesc</i> contains the length of the untruncated string value. (Function returns SQL_SUCCESS_WITH_INFO.)
24000	Invalid cursor state	The statement associated with the <i>hstmt</i> did not return a result set. There were no columns to describe.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1002	Invalid column number	(DM) The value specified for the argument <i>icol</i> was 0 and the argument <i>fDescType</i> was not SQL_COLUMN_COUNT. The value specified for the argument <i>icol</i> was greater than the number of columns in the result set and the argument <i>fDescType</i>

SQLFreeStmt

SQLFreeStmt

SQLSTATE Error		Description
S1008	Operation canceled	was not SQL_COLUMN_COUNT. Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1010	Function sequence error	(DM) The function was called prior to calling SQLPrepare or SQLExecDirect for the <i>hstmt</i> . (DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1090	Invalid string or buffer length	(DM) The value specified for the argument <i>cbDescMax</i> was less than 0.
S1091	Descriptor type out of range	(DM) The value specified for the argument <i>fDescType</i> was in the block of numbers reserved for ODBC descriptor types but was not valid for the version of ODBC supported by the driver (see "Comments").
S1C00	Driver not capable	The value specified for the argument <i>fDescType</i> was in the range of numbers reserved for driver-specific descriptor types but was not supported by the driver.
S1T00	Timeout expired	The timeout period expired before the data source returned the requested information. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

SQLColAttributes can return any SQLSTATE that can be returned by **SQLPrepare** or **SQLExecute** when called after **SQLPrepare** and before **SQLExecute** depending on when the data source evaluates the SQL statement associated with the *hstmt*.

Comments

SQLColAttributes returns information either in *pfDesc* or in *rgbDesc*. Integer information is returned in *pfDesc* as a 32-bit, signed value; all other formats of information are returned in *rgbDesc*. When information is returned in *pfDesc*, the driver ignores *rgbDesc*, *cbDescMax*, and *pcbDesc*. When information is returned in *rgbDesc*, the driver ignores *pfDesc*.

The currently defined descriptor types, the version of ODBC in which they were introduced, and the arguments in which information is returned for them are shown below; it is expected that more descriptor types will be defined to take advantage of different data sources. Descriptor types from 0 to 999 are reserved by ODBC; driver developers must reserve values greater than or equal to `SQL_COLUMN_DRIVER_START` for driver-specific use.

A driver must return a value for each of the descriptor types defined in the following table. If a descriptor type does not apply to a driver or data source, then, unless otherwise stated, the driver returns 0 in *pcbDesc* or an empty string in *rgbDesc*.

<i>fDescType</i>	Information returned in	Description
SQL_COLUMN_AUTO_INCREMENT (ODBC 1.0)	<i>pfDesc</i>	TRUE if the column is autoincrement. FALSE if the column is not autoincrement or is not numeric. Auto increment is valid for numeric data type columns only. An application can insert values into an autoincrement column, but cannot update values in the column.
SQL_COLUMN_CASE_SENSITIVE (ODBC 1.0)	<i>pfDesc</i>	TRUE if the column is treated as case sensitive for collations and comparisons. FALSE if the column is not treated as case sensitive for collations and comparisons or is noncharacter.
SQL_COLUMN_COUNT (ODBC 1.0)	<i>pfDesc</i>	Number of columns available in the result set. The <i>icol</i> argument is ignored.
SQL_COLUMN_DISPLAY_SIZE (ODBC 1.0)	<i>pfDesc</i>	Maximum number of characters required to display data from the column. For more information on display size, see "Precision, Scale, Length, and Display Size" in Appendix D, "Data Types."
SQL_COLUMN_LABEL (ODBC 2.0)	<i>rgbDesc</i>	The column label or title. For example, a column named EmpName might be labeled Employee Name.

SQLFreeStmt

SQLFreeStmt

<p>If a column does not have a label, the column name is returned. If the column is unlabeled and unnamed, an empty string is returned.</p>		
<i>fDescType</i>	Information returned in	Description
SQL_COLUMN_LENGTH (ODBC 1.0)	<i>pfDesc</i>	The length in bytes of data transferred on an SQLGetData or SQLFetch operation if SQL_C_DEFAULT is specified. For numeric data, this size may be different than the size of the data stored on the data source. For more length information, see "Precision, Scale, Length, and Display Size" in Appendix D, "Data Types."
SQL_COLUMN_MONEY (ODBC 1.0)	<i>pfDesc</i>	TRUE if the column is money data type. FALSE if the column is not money data type.
SQL_COLUMN_NAME (ODBC 1.0)	<i>rgbDesc</i>	The column name. If the column is unnamed, an empty string is returned.
SQL_COLUMN_NULLABLE (ODBC 1.0)	<i>pfDesc</i>	SQL_NO_NULLS if the column does not accept NULL values. SQL_NULLABLE if the column accepts NULL values. SQL_NULLABLE_UNKNOWN if it is not known if the column accepts NULL values.
SQL_COLUMN_OWNER_NAME (ODBC 2.0)	<i>rgbDesc</i>	The owner of the table that contains the column. The returned value is implementation-defined if the column is an expression or if the column is part of a view. If the data source does not support owners or the owner name cannot be determined, an empty string is returned.
SQL_COLUMN_PRECISION (ODBC 1.0)	<i>pfDesc</i>	The precision of the column on the data source. For more information on precision, see "Precision, Scale, Length, and Display Size" in Appendix D, "Data Types."
SQL_COLUMN_QUALIFIER_NAME (ODBC 2.0)	<i>rgbDesc</i>	The qualifier of the table that contains the column. The returned value is implementation-defined if the column is an expression or if the column is part of a view. If the data source does not support qualifiers or the qualifier name cannot be determined, an empty string is returned.
SQL_COLUMN_QUALIFIER_NAME	<i>rgbDesc</i>	The qualifier of the table that contains the

column. The returned value is implementation-defined if the column is an expression or if the column is part of a view. If the data source does not support qualifiers or the qualifier name cannot be determined, an empty string is returned.

<i>fDescType</i>	Information returned in	Description
SQL_COLUMN_SCALE (ODBC 1.0)	<i>pfDesc</i>	The scale of the column on the data source. For more information on scale, see "Precision, Scale, Length, and Display Size" in Appendix D, "Data Types."
SQL_COLUMN_SEARCHABLE (ODBC 1.0)	<i>pfDesc</i>	<p>SQL_UNSEARCHABLE if the column cannot be used in a WHERE clause.</p> <p>SQL_LIKE_ONLY if the column can be used in a WHERE clause only with the LIKE predicate.</p> <p>SQL_ALL_EXCEPT_LIKE if the column can be used in a WHERE clause with all comparison operators except LIKE.</p> <p>SQL_SEARCHABLE if the column can be used in a WHERE clause with any comparison operator.</p> <p>Columns of type SQL_LONGVARCHAR and SQL_LONGVARIABLE usually return SQL_LIKE_ONLY.</p>
SQL_COLUMN_TABLE_NAME (ODBC 2.0)	<i>rgbDesc</i>	<p>The name of the table that contains the column. The returned value is implementation-defined if the column is an expression or if the column is part of a view.</p> <p>If the table name cannot be determined, an empty string is returned.</p>
SQL_COLUMN_TYPE (ODBC 1.0)	<i>pfDesc</i>	SQL data type. This can be an ODBC SQL data type or a driver-specific SQL data type. For a list of valid ODBC SQL data types, see "SQL Data Types" in Appendix D, "Data Types." For information about driver-specific SQL data types, see the driver's documentation.
SQL_COLUMN_TYPE_NAME (ODBC 1.0)	<i>rgbDesc</i>	<p>Data source-dependent data type name; for example, "CHAR", "VARCHAR", "MONEY", "LONG VARBINARY", or "CHAR () FOR BIT DATA".</p> <p>If the type is unknown, an empty string is returned.</p>

SQLFreeStmt

SQL_COLUMN_UNSIGNED (ODBC 1.0)	<i>pfDesc</i>	TRUE if the column is unsigned (or not numeric). FALSE if the column is signed.
<i>fDescType</i>	Information returned in	Description
SQL_COLUMN_UPDATABLE (ODBC 1.0)	<i>pfDesc</i>	Column is described by the values for the defined constants: SQL_ATTR_READONLY SQL_ATTR_WRITE SQL_ATTR_READWRITE_UNKNOWN SQL_COLUMN_UPDATABLE describes the updatability of the column in the result set. Whether a column is updatable can be based on the data type, user privileges, and the definition of the result set itself. If it is unclear whether a column is updatable, SQL_ATTR_READWRITE_UNKNOWN should be returned.

This function is an extensible alternative to **SQLDescribeCol**. **SQLDescribeCol** returns a fixed set of descriptor information based on ANSI-89 SQL. **SQLColAttributes** allows access to the more extensive set of descriptor information available in ANSI SQL-92 and DBMS vendor extensions.

Related Functions

For information about	See
Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Returning information about a column in a result set	SQLDescribeCol
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
Fetching a row of data	SQLFetch

SQLColumnPrivileges



Extension Level 2 SQLColumnPrivileges returns a list of columns and associated privileges for the specified table. The driver returns the information as a result set on the specified *hstmt*.

Syntax

RETCODE **SQLColumnPrivileges**(*hstmt*, *szTableQualifier*, *cbTableQualifier*, *szTableOwner*, *cbTableOwner*, *szTableName*, *cbTableName*, *szColumnName*, *cbColumnName*)

The **SQLColumnPrivileges** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UCHAR FAR *	<i>szTableQualifier</i>	Input	Table qualifier. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have qualifiers.
SWORD	<i>cbTableQualifier</i>	Input	Length of <i>szTableQualifier</i> .
UCHAR FAR *	<i>szTableOwner</i>	Input	Owner name. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have owners.
SWORD	<i>cbTableOwner</i>	Input	Length of <i>szTableOwner</i> .
UCHAR FAR *	<i>szTableName</i>	Input	Table name.
SWORD	<i>cbTableName</i>	Input	Length of <i>szTableName</i> .
UCHAR FAR *	<i>szColumnName</i>	Input	String search pattern for column names.
SWORD	<i>cbColumnName</i>	Input	Length of <i>szColumnName</i> .

Returns

SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics

When **SQLColumnPrivileges** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLColumnPrivileges** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return

SQLGetConnectOption

SQLGetConnectOption

code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called. A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
SQLSTATE	Error	Description
S1010	Function sequence error	(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called.

		(DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1090	Invalid string or buffer length	(DM) The value of one of the name length arguments was less than 0, but not equal to SQL_NTS. The value of one of the name length arguments exceeded the maximum length value for the corresponding qualifier or name (see "Comments").
S1C00	Driver not capable	A table qualifier was specified and the driver or data source does not support qualifiers. A table owner was specified and the driver or data source does not support owners. A string search pattern was specified for the column name and the data source does not support search patterns for that argument. The combination of the current settings of the SQL_CONCURRENCY and SQL_CURSOR_TYPE statement options was not supported by the driver or data source.
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

Comments

SQLColumnPrivileges returns the results as a standard result set, ordered by TABLE_QUALIFIER, TABLE_OWNER, TABLE_NAME, COLUMN_NAME, and PRIVILEGE. The following table lists the columns in the result set.

Note: **SQLColumnPrivileges** might not return privileges for all columns. For example, a driver might not return information about privileges for pseudo-columns, such as Oracle ROWID. Applications can use any valid column, regardless of whether it is returned by **SQLColumnPrivileges**.

The lengths of VARCHAR columns shown in the table are maximums; the actual lengths depend on the data source. To determine the actual lengths of the TABLE_QUALIFIER, TABLE_OWNER, TABLE_NAME,

SQLGetConnectOption

SQLGetConnectOption

and COLUMN_NAME columns, an application can call **SQLGetInfo** with the SQL_MAX_QUALIFIER_NAME_LEN, SQL_MAX_OWNER_NAME_LEN, SQL_MAX_TABLE_NAME_LEN, and SQL_MAX_COLUMN_NAME_LEN options.

Column Name	Data Type	Comments
TABLE_QUALIFIER	Varchar(128)	Table qualifier identifier; NULL if not applicable to the data source. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those tables that do not have qualifiers.
TABLE_OWNER	Varchar(128)	Table owner identifier; NULL if not applicable to the data source. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those tables that do not have owners.
TABLE_NAME	Varchar(128) not NULL	Table identifier.
COLUMN_NAME	Varchar(128) not NULL	Column identifier.
GRANTOR	Varchar(128)	Identifier of the user who granted the privilege; NULL if not applicable to the data source.
GRANTEE	Varchar(128) not NULL	Identifier of the user to whom the privilege was granted.
Column Name	Data Type	Comments
PRIVILEGE	Varchar(128) not NULL	Identifies the column privilege. May be one of the following or others supported by the data source when implementation-defined: SELECT: The grantee is permitted to retrieve data for the column. INSERT: The grantee is permitted to provide data for the column in new rows that are inserted into the associated table. UPDATE: The grantee is permitted to update data in the column. REFERENCES: The grantee is permitted to refer to the column within a constraint

		(for example, a unique, referential, or table check constraint).
IS_GRANTABLE	Varchar(3)	Indicates whether the grantee is permitted to grant the privilege to other users; "YES", "NO", or NULL if unknown or not applicable to the data source.

The *szColumnName* argument accepts a search pattern. For more information about valid search patterns, see "Search Pattern Arguments" earlier in this chapter.

Code Example For a code example of a similar function, see **SQLColumns**.

Related Functions

For information about	See
Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Returning the columns in a table or tables	SQLColumns (extension)
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
Fetching a row of data	SQLFetch
Returning privileges for a table or tables	SQLTablePrivileges (extension)
Returning a list of tables in a data source	SQLTables (extension)

SQLColumns



Extension Level 1 **SQLColumns** returns the list of column names in specified tables. The driver returns this information as a result set on the specified *hstmt*.

Syntax RETCODE **SQLColumns**(*hstmt*, *szTableQualifier*, *cbTableQualifier*, *szTableOwner*, *cbTableOwner*, *szTableName*, *cbTableName*, *szColumnName*, *cbColumnName*)

The **SQLColumns** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UCHAR FAR *	<i>szTableQualifier</i>	Input	Qualifier name. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, SQLGetConnectOption

SQLGetConnectOption

			an empty string ("") denotes those tables that do not have qualifiers.
SWORD	<i>cbTableQualifier</i>	Input	Length of <i>szTableQualifier</i> .
UCHAR FAR *	<i>szTableOwner</i>	Input	String search pattern for owner names. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have owners.
SWORD	<i>cbTableOwner</i>	Input	Length of <i>szTableOwner</i> .
UCHAR FAR *	<i>szTableName</i>	Input	String search pattern for table names.
SWORD	<i>cbTableName</i>	Input	Length of <i>szTableName</i> .
UCHAR FAR *	<i>szColumnName</i>	Input	String search pattern for column names.
SWORD	<i>cbColumnName</i>	Input	Length of <i>szColumnName</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLColumns** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLColumns** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called. A cursor was open on the <i>hstmt</i> but

		SQLFetch or SQLExtendedFetch had not been called.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1010	Function sequence error	(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA . This function was called before data was sent for all data-at-execution parameters or columns.
SQLSTATE Error		Description
S1090	Invalid string or buffer length	(DM) The value of one of the name length arguments was less than 0, but not equal to SQL_NTS . The value of one of the name length arguments exceeded the maximum length value for the corresponding qualifier or name. The maximum length of each qualifier or name may be obtained by calling SQLGetInfo with the <i>flInfoType</i> values (see "Comments").
S1C00	Driver not capable	A table qualifier was specified and the SQLGetConnectOption

SQLGetConnectOption

		driver or data source does not support qualifiers.
		A table owner was specified and the driver or data source does not support owners.
		A string search pattern was specified for the table owner, table name, or column name and the data source does not support search patterns for one or more of those arguments.
		The combination of the current settings of the SQL_CONCURRENCY and SQL_CURSOR_TYPE statement options was not supported by the driver or data source.
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

Comments

This function is typically used before statement execution to retrieve information about columns for a table or tables from the data source's catalog. Note by contrast, that the functions **SQLColAttributes** and **SQLDescribeCol** describe the columns in a result set and that the function **SQLNumResultCols** returns the number of columns in a result set.

Note: **SQLColumns** might not return all columns. For example, a driver might not return information about pseudo-columns, such as Oracle ROWID. Applications can use any valid column, regardless of whether it is returned by **SQLColumns**.

SQLColumns returns the results as a standard result set, ordered by TABLE_QUALIFIER, TABLE_OWNER, and TABLE_NAME. The following table lists the columns in the result set. Additional columns beyond column 12 (REMARKS) can be defined by the driver.

The lengths of VARCHAR columns shown in the table are maximums; the actual lengths depend on the data source. To determine the actual lengths of the TABLE_QUALIFIER, TABLE_OWNER, TABLE_NAME, and COLUMN_NAME columns, an application can call **SQLGetInfo** with the SQL_MAX_QUALIFIER_NAME_LEN, SQL_MAX_OWNER_NAME_LEN, SQL_MAX_TABLE_NAME_LEN, and SQL_MAX_COLUMN_NAME_LEN options.

Column Name	Data Type	Comments
-------------	-----------	----------

TABLE_QUALIFIER	Varchar(128)	Table qualifier identifier; NULL if not applicable to the data source. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those tables that do not have qualifiers.
TABLE_OWNER	Varchar(128)	Table owner identifier; NULL if not applicable to the data source. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those tables that do not have owners.
TABLE_NAME	Varchar(128) not NULL	Table identifier.
COLUMN_NAME	Varchar(128) not NULL	Column identifier.
DATA_TYPE	Smallint not NULL	SQL data type. This can be an ODBC SQL data type or a driver-specific SQL data type. For a list of valid ODBC SQL data types, see "SQL Data Types" in Appendix D, "Data Types." For information about driver-specific SQL data types, see the driver's documentation.
TYPE_NAME	Varchar(128) not NULL	Data source-dependent data type name; for example, "CHAR", "VARCHAR", "MONEY", "LONG VARBINARY", or "CHAR () FOR BIT DATA".
PRECISION	Integer	The precision of the column on the data source. For precision information, see "Precision, Scale, Length, and Display Size" in Appendix D, "Data Types."
Column Name	Data Type	Comments
LENGTH	Integer	The length in bytes of data transferred on an SQLGetData or SQLFetch operation if SQL_C_DEFAULT is specified. For numeric data, this size may be different than the size of the data stored on the data source. This value is the same as the PRECISION column for character or binary data. For more information about length, see "Precision, Scale, Length, and Display Size" in Appendix D, "Data Types."

SQLGetConnectOption

SQLGetConnectOption

SCALE	Smallint	The scale of the column on the data source. For more scale information, see "Precision, Scale, Length, and Display Size" in Appendix D, "Data Types." NULL is returned for data types where scale is not applicable.
RADIX	Smallint	<p>For numeric data types, either 10 or 2. If it is 10, the values in PRECISION and SCALE give the number of decimal digits allowed for the column. For example, a DECIMAL(12,5) column would return a RADIX of 10, a PRECISION of 12, and a SCALE of 5; A FLOAT column could return a RADIX of 10, a PRECISION of 15 and a SCALE of NULL.</p> <p>If it is 2, the values in PRECISION and SCALE give the number of bits allowed in the column. For example, a FLOAT column could return a RADIX of 2, a PRECISION of 53, and a SCALE of NULL.</p> <p>NULL is returned for data types where radix is not applicable.</p>
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> <p>SQL_NULLABLE_UNKNOWN if it is not known if the column accepts NULL values.</p>
REMARKS	Varchar(254)	A description of the column.

The *szTableOwner*, *szTableName*, and *szColumnName* arguments accept search patterns. For more information about valid search patterns, see "Search Pattern Arguments" earlier in this chapter.

Code Example In the following example, an application declares storage locations for the result set returned by **SQLColumns**. It calls **SQLColumns** to return a result set that describes each column in the EMPLOYEE table. It then calls **SQLBindCol** to bind the columns in the result set to the storage locations. Finally, the application fetches each row of data with **SQLFetch** and processes it.

```
#define STR_LEN 128+1
#define REM_LEN 254+1
```

```

/* Declare storage locations for result set data */

UCHAR szQualifier[STR_LEN], szOwner[STR_LEN];
UCHAR szTableName[STR_LEN], szColName[STR_LEN];
UCHAR szTypeName[STR_LEN], szRemarks[REM_LEN];
SDWORD Precision, Length;
SWORD DataType, Scale, Radix, Nullable;

/* Declare storage locations for bytes available to return */

SDWORD cbQualifier, cbOwner, cbTableName, cbColName;
SDWORD cbTypeName, cbRemarks, cbDataType, cbPrecision;
SDWORD cbLength, cbScale, cbRadix, cbNullable;

retcode = SQLColumns(hstmt,
    NULL, 0, /* All qualifiers */
    NULL, 0, /* All owners */
    "EMPLOYEE", SQL_NTS, /* EMPLOYEE table */
    NULL, 0); /* All columns */
if (retcode == SQL_SUCCESS) {

    /* Bind columns in result set to storage locations */

    SQLBindCol(hstmt, 1, SQL_C_CHAR, szQualifier, STR_LEN, &cbQualifier);
    SQLBindCol(hstmt, 2, SQL_C_CHAR, szOwner, STR_LEN, &cbOwner);
    SQLBindCol(hstmt, 3, SQL_C_CHAR, szTableName, STR_LEN, &cbTableName);
    SQLBindCol(hstmt, 4, SQL_C_CHAR, szColName, STR_LEN, &cbColName);
    SQLBindCol(hstmt, 5, SQL_C_SHORT, &DataType, 0, &cbDataType);
    SQLBindCol(hstmt, 6, SQL_C_CHAR, szTypeName, STR_LEN, &cbTypeName);
    SQLBindCol(hstmt, 7, SQL_C_LONG, &Precision, 0, &cbPrecision);
    SQLBindCol(hstmt, 8, SQL_C_LONG, &Length, 0, &cbLength);
    SQLBindCol(hstmt, 9, SQL_C_SHORT, &Scale, 0, &cbScale);
    SQLBindCol(hstmt, 10, SQL_C_SHORT, &Radix, 0, &cbRadix);
    SQLBindCol(hstmt, 11, SQL_C_SHORT, &Nullable, 0, &cbNullable);
    SQLBindCol(hstmt, 12, SQL_C_CHAR, szRemarks, REM_LEN, &cbRemarks);

    while(TRUE) {
        retcode = SQLFetch(hstmt);
        if (retcode == SQL_ERROR || retcode == SQL_SUCCESS_WITH_INFO) {
            show_error();
        }
        if (retcode == SQL_SUCCESS || retcode == SQL_SUCCESS_WITH_INFO){
            ...; /* Process fetched data */
        } else {
            break;
        }
    }
}

```

SQLGetConnectOption

Related Functions	For information about	See
	Assigning storage for a column in a result set	SQLBindCol
	Canceling statement processing	SQLCancel
	Returning privileges for a column or columns	SQLColumnPrivileges (extension)
	Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
	Fetching a row of data	SQLFetch
	Returning table statistics and indexes	SQLStatistics (extension)
	Returning a list of tables in a data source	SQLTables (extension)
	Returning privileges for a table or tables	SQLTablePrivileges (extension)

SQLConnect



Core **SQLConnect** loads a driver and establishes a connection to a data source. The connection handle references storage of all information about the connection, including status, transaction state, and error information.

Syntax RETCODE **SQLConnect**(*hdbc*, *szDSN*, *cbDSN*, *szUID*, *cbUID*, *szAuthStr*, *cbAuthStr*)

The **SQLConnect** function accepts the following arguments.

Type	Argument	Use	Description
HDBC	<i>hdbc</i>	Input	Connection handle.
UCHAR FAR *	<i>szDSN</i>	Input	Data source name.
SWORD	<i>cbDSN</i>	Input	Length of <i>szDSN</i> .
UCHAR FAR *	<i>szUID</i>	Input	User identifier.
SWORD	<i>cbUID</i>	Input	Length of <i>szUID</i> .
UCHAR FAR *	<i>szAuthStr</i>	Input	Authentication string (typically the password).
SWORD	<i>cbAuthStr</i>	Input	Length of <i>szAuthStr</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics

When **SQLConnect** returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, an associated **SQLSTATE** value may be obtained by calling **SQLError**. The following table lists the **SQLSTATE** values commonly returned by **SQLConnect** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of **SQLSTATE**s returned by the Driver Manager. The return code associated with each **SQLSTATE** value is **SQL_ERROR**, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO .)
08001	Unable to connect to data source	The driver was unable to establish a connection with the data source.
08002	Connection in use	(DM) The specified <i>hdbc</i> had already been used to establish a connection with a data source and the connection was still open.
SQLSTATE	Error	Description
08004	Data source rejected establishment of connection	The data source rejected the establishment of the connection for implementation-defined reasons.
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was attempting to connect failed before the function completed processing.
28000	Invalid authorization specification	The value specified for the argument <i>szUID</i> or the value specified for the argument <i>szAuthStr</i> violated restrictions defined by the data source.
IM001	Driver does not support this function	(DM) The driver specified by the data source name does not support the function.
IM002	Data source not found and no default driver specified	(DM) The data source name specified in the argument <i>szDSN</i> was not found in the ODBC.INI file or registry, nor was there a default driver specification. (DM) The ODBC.INI file could not be found.
IM003	Specified driver could not be loaded	(DM) The driver listed in the data source specification in the ODBC.INI file or registry was not found or could not be loaded for some other reason.

SQLGetConnectOption

SQLGetConnectOption

IM004	Driver's SQLAllocEnv failed	(DM) During SQLConnect , the Driver Manager called the driver's SQLAllocEnv function and the driver returned an error.
IM005	Driver's SQLAllocConnect failed	(DM) During SQLConnect , the Driver Manager called the driver's SQLAllocConnect function and the driver returned an error.
IM006	Driver's SQLSetConnectOption failed	(DM) During SQLConnect , the Driver Manager called the driver's SQLSetConnectOption function and the driver returned an error. (Function returns SQL_SUCCESS_WITH_INFO).
IM009	Unable to load translation DLL	The driver was unable to load the translation DLL that was specified for the data source.
SQLSTATE	Error	Description
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	(DM) The Driver Manager was unable to allocate memory required to support execution or completion of the function. The driver was unable to allocate memory required to support execution or completion of the function.
S1090	Invalid string or buffer length	(DM) The value specified for argument <i>cbDSN</i> was less than 0, but not equal to SQL_NTS. (DM) The value specified for argument <i>cbDSN</i> exceeded the maximum length for a data source name. (DM) The value specified for argument <i>cbUID</i> was less than 0, but not equal to SQL_NTS. (DM) The value specified for argument <i>cbAuthStr</i> was less than 0, but not equal to SQL_NTS.
S1T00	Timeout expired	The timeout period expired before the connection to the data source completed. The timeout period is set through

**SQLSetConnectOption,
SQL_LOGIN_TIMEOUT.****Comments**

The Driver Manager does not load a driver until the application calls a function (**SQLConnect**, **SQLDriverConnect**, or **SQLBrowseConnect**) to connect to the driver. Until that point, the Driver Manager works with its own handles and manages connection information. When the application calls a connection function, the Driver Manager checks if a driver is currently loaded for the specified *hdbc*:

- If a driver is not loaded, the Driver Manager loads the driver and calls **SQLAllocEnv**, **SQLAllocConnect**, **SQLSetConnectOption** (if the application specified any connection options), and the connection function in the driver. The Driver Manager returns SQLSTATE IM006 (Driver's SQLSetConnectOption failed) and SQL_SUCCESS_WITH_INFO for the connection function if the driver returned an error for **SQLSetConnectOption**.
- If the specified driver is already loaded on the *hdbc*, the Driver Manager only calls the connection function in the driver. In this case, the driver must ensure that all connection options for the *hdbc* maintain their current settings.
- If a different driver is loaded, the Driver Manager calls **SQLFreeConnect** and **SQLFreeEnv** in the loaded driver and then unloads that driver. It then performs the same operations as when a driver is not loaded.

The driver then allocates handles and initializes itself.

Note: To resolve the addresses of the ODBC functions exported by the driver, the Driver Manager checks if the driver exports a dummy function with the ordinal 199. If it does not, the Driver Manager resolves the addresses by name. If it does, the Driver Manager resolves the addresses of the ODBC functions by ordinal, which is faster. The ordinal values of the ODBC functions must match the values of the *fFunction* argument in **SQLGetFunctions**; all other exported functions (such as **WEP**) must have ordinal values outside the range 1–199.

When the application calls **SQLDisconnect**, the Driver Manager calls **SQLDisconnect** in the driver. However, it does not unload the driver. This keeps the driver in memory for applications that repeatedly connect to and disconnect from a data source. When the application calls **SQLFreeConnect**, the Driver Manager calls **SQLFreeConnect** and **SQLFreeEnv** in the driver and then unloads the driver.

SQLGetConnectOption

SQLGetConnectOption

An ODBC application can establish more than one connection.

Code Example In the following example, an application allocates environment and connection handles. It then connects to the EmpData data source with the user ID JohnS and the password Sesame and processes data. When it has finished processing data, it disconnects from the data source and frees the handles.

```
HENV henv;
HDBC hdbc;
HSTMT hstmt;
RETCODE retcode;

retcode = SQLAllocEnv(&henv);          /* Environment handle */
if (retcode == SQL_SUCCESS) {
    retcode = SQLAllocConnect(henv, &hdbc); /* Connection handle */
    if (retcode == SQL_SUCCESS) {

        /* Set login timeout to 5 seconds. */

        SQLSetConnectOption(hdbc, SQL_LOGIN_TIMEOUT, 5);

        /* Connect to data source */

        retcode = SQLConnect(hdbc, "EmpData", SQL_NTS,
                               "JohnS", SQL_NTS,
                               "Sesame", SQL_NTS);

        if (retcode == SQL_SUCCESS || retcode == SQL_SUCCESS_WITH_INFO){

            /* Process data after successful connection */

            retcode = SQLAllocStmt(hdbc, &hstmt); /* Statement handle */

            if (retcode == SQL_SUCCESS) {
                ...;
                ...;
                ...;
                SQLFreeStmt(hstmt, SQL_DROP);
            }
            SQLDisconnect(hdbc);
        }
        SQLFreeConnect(hdbc);
    }
    SQLFreeEnv(henv);
}
```

**Related
Functions****For information about****See**

Allocating a connection handle

SQLAllocConnect

Allocating a statement handle

SQLAllocStmtDiscovering and enumerating values required
to connect to a data source**SQLBrowseConnect**
(extension)

Disconnecting from a data source

SQLDisconnectConnecting to a data source using a
connection string or dialog box**SQLDriverConnect** (extension)

Returning the setting of a connection option

SQLGetConnectOption
(extension)

Setting a connection option

SQLSetConnectOption
(extension)

SQLDescribeCol



Core

SQLDescribeCol returns the result descriptor — column name, type, precision, scale, and nullability — for one column in the result set; it cannot be used to return information about the bookmark column (column 0).

Syntax

RETCODE **SQLDescribeCol**(*hstmt*, *icol*, *szColName*, *cbColNameMax*, *pcbColName*, *pfSqlType*, *pcbColDef*, *pibScale*, *pfNullable*)

The **SQLDescribeCol** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UWORD	<i>icol</i>	Input	Column number of result data, ordered sequentially left to right, starting at 1.
UCHAR FAR *	<i>szColName</i>	Output	Pointer to storage for the column name. If the column is unnamed or the column name cannot be determined, the driver returns an empty string.
SWORD	<i>cbColNameMax</i>	Input	Maximum length of the <i>szColName</i> buffer.
SWORD FAR *	<i>pcbColName</i>	Output	Total number of bytes (excluding the null termination byte) available to return in <i>szColName</i> . If the number of bytes available to return is greater than or equal to <i>cbColNameMax</i> , the column name in <i>szColName</i> is truncated to <i>cbColNameMax</i> – 1 bytes.
Type	Argument	Use	Description
SWORD FAR *	<i>pfSqlType</i>	Output	The SQL data type of the column. This must be one of the following values: SQL_BIGINT SQL_BINARY SQL_BIT SQL_CHAR SQL_DATE SQL_DECIMAL SQL_DOUBLE

SQL_FLOAT
 SQL_INTEGER
 SQL_LONGVARBINARY
 SQL_LONGVARCHAR
 SQL_NUMERIC
 SQL_REAL
 SQL_SMALLINT
 SQL_TIME
 SQL_TIMESTAMP
 SQL_TINYINT
 SQL_VARBINARY
 SQL_VARCHAR

or a driver-specific SQL data type. If the data type cannot be determined, the driver returns 0.

For more information, see "SQL Data Types" in Appendix D, "Data Types." For information about driver-specific SQL data types, see the driver's documentation.

UDWORD FAR * <i>pcbColDef</i>	Output	The precision of the column on the data source. If the precision cannot be determined, the driver returns 0. For more information on precision, see “Precision, Scale, Length, and Display Size” in Appendix D, “Data Types.”	
SWORD FAR * <i>pibScale</i>	Output	The scale of the column on the data source. If the scale cannot be determined or is not applicable, the driver returns 0. For more information on scale, see “Precision, Scale, Length, and Display Size” in Appendix D, “Data Types.”	
Type	Argument	Use	Description
SWORD FAR * <i>pfNullable</i>	Output	Indicates whether the column allows NULL values. One of the following values: SQL_NO_NULLS: The column does not allow NULL values. SQL_NULLABLE: The column allows NULL values. SQL_NULLABLE_UNKNOWN: The driver cannot determine if the	

SQLGetCursorName

SQLGetCursorName

column allows NULL values.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLDescribeCol** returns either SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLDescribeCol** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	The buffer <i>szColName</i> was not large enough to return the entire column name, so the column name was truncated. The argument <i>pcbColName</i> contains the length of the untruncated column name. (Function returns SQL_SUCCESS_WITH_INFO.)
24000	Invalid cursor state	The statement associated with the <i>hstmt</i> did not return a result set. There were no columns to describe.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
SQLSTATE	Error	Description
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1002	Invalid column number	(DM) The value specified for the argument <i>icol</i> was 0. The value specified for the argument <i>icol</i> was greater than the number of columns in

		the result set.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1010	Function sequence error	(DM) The function was called prior to calling SQLPrepare or SQLExecDirect for the <i>hstmt</i> . (DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1090	Invalid string or buffer length	(DM) The value specified for argument <i>cbColNameMax</i> was less than 0.
S1T00	Timeout expired	<div>BEGIN BREAK</div> The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

SQLDescribeCol can return any SQLSTATE that can be returned by **SQLPrepare** or **SQLExecute** when called after **SQLPrepare** and before **SQLExecute** depending on when the data source evaluates the SQL statement associated with the *hstmt*.

Comments

An application typically calls **SQLDescribeCol** after a call to **SQLPrepare** and before or after the associated call to **SQLExecute**. An application can also call **SQLDescribeCol** after a call to **SQLExecDirect**. **SQLDescribeCol** retrieves the column name, type, and length generated by a **SELECT** statement. If the column is an expression, *szColName* is either an empty string or a driver-defined name.

SQLGetCursorName

Note: ODBC supports SQL_NULLABLE_UNKNOWN as an extension, even though the X/Open and SQL Access Group Call Level Interface specification does not specify the option for **SQLDescribeCol**.

Related Functions

For information about

See

Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Returning information about a column in a result set	SQLColAttributes
Fetching a row of data	SQLFetch
Returning the number of result set columns	SQLNumResultCols
Preparing a statement for execution	SQLPrepare

SQLDescribeParam



Extension Level 2 **SQLDescribeParam** returns the description of a parameter marker associated with a prepared SQL statement.

Syntax RETCODE **SQLDescribeParam**(*hstmt*, *ipar*, *pfSqlType*, *pcbColDef*, *pibScale*, *pfNullable*)

The **SQLDescribeParam** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UWORD	<i>ipar</i>	Input	Parameter marker number ordered sequentially left to right, starting at 1.
SWORD FAR *	<i>pfSqlType</i>	Output	The SQL data type of the parameter. This must be one of the following values: SQL_BIGINT SQL_BINARY SQL_BIT SQL_CHAR SQL_DATE SQL_DECIMAL SQL_DOUBLE SQL_FLOAT SQL_INTEGER SQL_LONGVARBINARY SQL_LONGVARCHAR SQL_NUMERIC SQL_REAL SQL_SMALLINT SQL_TIME SQL_TIMESTAMP SQL_TINYINT SQL_VARBINARY SQL_VARCHAR or a driver-specific SQL data type. For more information, see “SQL Data Types” in Appendix D, “Data Types.” For information about driver-specific SQL data types, see the driver’s documentation.
Type	Argument	Use	Description
UDWORD FAR	<i>pcbColDef</i>	Output	The precision of the column or

SQLGetData

SQLGetData

*		expression of the corresponding parameter marker as defined by the data source. For further information concerning precision, see "Precision, Scale, Length, and Display Size," in Appendix D, "Data Types."
SWORD FAR * <i>piScale</i>	Output	The scale of the column or expression of the corresponding parameter as defined by the data source. For more information on scale, see "Precision, Scale, Length, and Display Size," in Appendix D, "Data Types."
SWORD FAR * <i>pfNullable</i>	Output	Indicates whether the parameter allows NULL values. One of the following: SQL_NO_NULLS: The parameter does not allow NULL values (this is the default value). SQL_NULLABLE: The parameter allows NULL values. SQL_NULLABLE_UNKNOWN: The driver cannot determine if the parameter allows NULL values.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLDescribeParam** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLDescribeParam** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.

SQLSTATE	Error	Description
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation error	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1010	Function sequence error	(DM) The function was called prior to calling SQLPrepare or SQLExecDirect for the <i>hstmt</i> . (DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1093	Invalid parameter number	(DM) The value specified for the argument <i>ipar</i> was 0. The value specified for the argument <i>ipar</i> was greater than the number of parameters in the associated SQL statement.
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

SQLGetData

Comments

Parameter markers are numbered from left to right in the order they appear in the SQL statement.

SQLDescribeParam does not return the type (input, input/output, or output) of a parameter in an SQL statement. Except in calls to procedures, all parameters in SQL statements are input parameters. To determine the type of each parameter in a call to a procedure, an application calls **SQLProcedureColumns**.

Related Functions

For information about

See

Canceling statement processing

SQLCancel

Executing a prepared SQL statement

SQLExecute

Preparing a statement for execution

SQLPrepare

Assigning storage for a parameter

SQLBindParameter

SQLDisconnect



Core **SQLDisconnect** closes the connection associated with a specific connection handle.

Syntax RETCODE **SQLDisconnect**(*hdbc*)

The **SQLDisconnect** function accepts the following argument.

Type	Argument	Use	Description
HDBC	<i>hdbc</i>	Input	Connection handle.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLDisconnect** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLDisconnect** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01002	Disconnect error	An error occurred during the disconnect. However, the disconnect succeeded. (Function returns SQL_SUCCESS_WITH_INFO.)
08003	Connection not open	(DM) The connection specified in the argument <i>hdbc</i> was not open.
25000	Invalid transaction state	There was a transaction in process on the connection specified by the argument <i>hdbc</i> . The transaction remains active.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hdbc</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i>

SQLGetFunctions

SQLGetFunctions

SQLSTATE Error		Description
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1010	Function sequence error	(DM) An asynchronously executing function was called for an <i>hstmt</i> associated with the <i>hdbc</i> and was still executing when SQLDisconnect was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for an <i>hstmt</i> associated with the <i>hdbc</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.

Comments If an application calls **SQLDisconnect** after **SQLBrowseConnect** returns SQL_NEED_DATA and before it returns a different return code, the driver cancels the connection browsing process and returns the *hdbc* to an unconnected state.

If an application calls **SQLDisconnect** while there is an incomplete transaction associated with the connection handle, the driver returns SQLSTATE 25000 (Invalid transaction state), indicating that the transaction is unchanged and the connection is open. An incomplete transaction is one that has not been committed or rolled back with **SQLTransact**.

If an application calls **SQLDisconnect** before it has freed all *hstmts* associated with the connection, the driver frees those *hstmts* after it successfully disconnects from the data source. However, if one or more of the *hstmts* associated with the connection are still executing asynchronously, **SQLDisconnect** will return SQL_ERROR with a SQLSTATE value of S1010 (Function sequence error).

Code Example See **SQLBrowseConnect** and **SQLConnect**.

Related Functions	For information about	See
	Allocating a connection handle	SQLAllocConnect
	Connecting to a data source	SQLConnect
	Connecting to a data source using a	SQLDriverConnect (extension)

connection string or dialog box

Freeing a connection handle

SQLFreeConnect

Executing a commit or rollback operation

SQLTransact

SQLDriverConnect (Windows)



Extension Level 1 **SQLDriverConnect** is an alternative to **SQLConnect**. It supports data sources that require more connection information than the three arguments in **SQLConnect**; dialog boxes to prompt the user for all connection information; and data sources that are not defined in the ODBC.INI file or registry.

SQLDriverConnect provides the following connection options:

- Establish a connection using a connection string that contains the data source name, one or more user IDs, one or more passwords, and other information required by the data source.
- Establish a connection using a partial connection string or no additional information; in this case, the Driver Manager and the driver can each prompt the user for connection information.
- Establish a connection to a data source that is not defined in the ODBC.INI file or registry. If the application supplies a partial connection string, the driver can prompt the user for connection information.

Once a connection is established, **SQLDriverConnect** returns the completed connection string. The application can use this string for subsequent connection requests.

Syntax

RETCODE **SQLDriverConnect**(*hdbc*, *hwnd*, *szConnStrIn*, *cbConnStrIn*, *szConnStrOut*, *cbConnStrOutMax*, *pcbConnStrOut*, *fDriverCompletion*)

The **SQLDriverConnect** function accepts the following arguments:

Type	Argument	Use	Description
HDBC	<i>hdbc</i>	Input	Connection handle.
HWND	<i>hwnd</i>	Input	Window handle. The application can pass the handle of the parent window, if applicable, or a null pointer if either the window handle is not applicable or if SQLDriverConnect will not present any dialog boxes.
UCHAR FAR *	<i>szConnStrIn</i>	Input	A full connection string (see the syntax in "Comments"), a partial connection string, or an empty string.

Type	Argument	Use	Description
SWORD	<i>cbConnStrIn</i>	Input	Length of <i>szConnStrIn</i> .
UCHAR FAR *	<i>szConnStrOut</i>	Output	Pointer to storage for the completed connection string. Upon successful connection to the target data source, this buffer contains the completed connection string. Applications should allocate at least 255 bytes for this buffer.
SWORD	<i>cbConnStrOutMax</i>	Input	Maximum length of the <i>szConnStrOut</i> buffer.
SWORD FAR *	<i>pcbConnStrOut</i>	Output	Pointer to the total number of bytes (excluding the null termination byte) available to return in <i>szConnStrOut</i> . If the number of bytes available to return 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.
UWORD	<i>fDriverCompletion</i>	Input	Flag which indicates whether Driver Manager or driver must prompt for more connection information: SQL_DRIVER_PROMPT, SQL_DRIVER_COMPLETE, SQL_DRIVER_COMPLETE_REQUIRED, or SQL_DRIVER_NOPROMPT. (See "Comments," for additional information.)

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_NO_DATA_FOUND, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLDriverConnect** returns either SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLDriverConnect** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE Error	Description
----------------	-------------

SQLGetInfo

SQLGetInfo

01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	The buffer <i>szConnStrOut</i> was not large enough to return the entire connection string, so the connection string was truncated. The argument <i>pcbConnStrOut</i> contains the length of the untruncated connection string. (Function returns SQL_SUCCESS_WITH_INFO.)
01S00	Invalid connection string attribute	An invalid attribute keyword was specified in the connection string (<i>szConnStrIn</i>) but the driver was able to connect to the data source anyway. (Function returns SQL_SUCCESS_WITH_INFO.)
08001	Unable to connect to data source	The driver was unable to establish a connection with the data source.
08002	Connection in use	(DM) The specified <i>hdbc</i> had already been used to establish a connection with a data source and the connection was still open.
08004	Data source rejected establishment of connection	The data source rejected the establishment of the connection for implementation-defined reasons.
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was attempting to connect failed before the function completed processing.
28000	Invalid authorization specification	Either the user identifier or the authorization string or both as specified in the connection string (<i>szConnStrIn</i>) violated restrictions defined by the data source.
IM001	Driver does not support this function	(DM) The driver corresponding to the specified data source name does not support the function.
IM002	Data source not found and no default driver specified	(DM) The data source name specified in the connection string (<i>szConnStrIn</i>) was not found in the ODBC.INI file or registry and there was no default driver specification. (DM) The ODBC.INI file could not be found.

SQLSTATE Error

Description

IM003	Specified driver could not be loaded	(DM) The driver listed in the data source specification in the ODBC.INI file or registry, or specified by the DRIVER keyword, was not found or could not be loaded for some other reason.
IM004	Driver's SQLAllocEnv failed	(DM) During SQLDriverConnect , the Driver Manager called the driver's SQLAllocEnv function and the driver returned an error.
IM005	Driver's SQLAllocConnect failed	(DM) During SQLDriverConnect , the Driver Manager called the driver's SQLAllocConnect function and the driver returned an error.
IM006	Driver's SQLSetConnectOption failed	(DM) During SQLDriverConnect , the Driver Manager called the driver's SQLSetConnectOption function and the driver returned an error.
IM007	No data source or driver specified; dialog prohibited	No data source name or driver was specified in the connection string and <i>fDriverCompletion</i> was SQL_DRIVER_NOPROMPT.
IM008	Dialog failed	(DM) The Driver Manager attempted to display the SQL Data Sources dialog box and failed. The driver attempted to display its login dialog box and failed.
IM009	Unable to load translation DLL	The driver was unable to load the translation DLL that was specified for the data source or for the connection.
IM010	Data source name too long	(DM) The attribute value for the DSN keyword was longer than SQL_MAX_DSN_LENGTH characters.
IM011	Driver name too long	(DM) The attribute value for the DRIVER keyword was longer than 255 characters.
IM012	DRIVER keyword syntax error	(DM) The keyword-value pair for the DRIVER keyword contained a syntax error.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.

SQLGetInfo

SQLSTATE	Error	Description
S1001	Memory allocation failure	The Driver Manager was unable to allocate memory required to support execution or completion of the function. The driver was unable to allocate memory required to support execution or completion of the function.
S1090	Invalid string or buffer length	(DM) The value specified for argument <i>cbConnStrIn</i> was less than 0 and was not equal to SQL_NTS. (DM) The value specified for argument <i>cbConnStrOutMax</i> was less than 0.
S1110	Invalid driver completion	(DM) The value specified for the argument <i>fDriverCompletion</i> was not equal to SQL_DRIVER_PROMPT, SQL_DRIVER_COMPLETE, SQL_DRIVER_COMPLETE_REQUIRED or SQL_DRIVER_NOPROMPT.
S1T00	Timeout expired	The timeout period expired before the connection to the data source completed. The timeout period is set through SQLSetConnectOption , SQL_LOGIN_TIMEOUT.

Comments

Connection Strings

A connection string has the following Syntax

```

connection-string ::= empty-string[:] | attribute[:] | attribute; connection-
string
empty-string ::=
attribute ::= attribute-keyword=attribute-value | DRIVER={attribute-value}
(The braces ({} ) are literal; the application must specify them.)
attribute-keyword ::= DSN | UID | PWD
                    | driver-defined-attribute-keyword
attribute-value ::= character-string
driver-defined-attribute-keyword ::= identifier

```

where *character-string* has zero or more characters; *identifier* has one or more characters; *attribute-keyword* is case insensitive; *attribute-value* may be case sensitive; and the value of the **DSN** keyword does not consist solely of blanks. Because of connection string and initialization file grammar, keywords and attribute values that contain the characters **{ } 0 , ; ? * = ! @** should be avoided. Because of the registry grammar,

keywords and data source names cannot contain the backslash (\) character.

Note: The **DRIVER** keyword was introduced in ODBC 2.0 and is not supported by ODBC 1.0 drivers.

The connection string may include any number of driver-defined keywords. Because the **DRIVER** keyword does not use information from the ODBC.INI file or registry, the driver must define enough keywords so that a driver can connect to a data source using only the information in the connection string. (For more information, see “Driver Guidelines,” later in this section.) The driver defines which keywords are required in order to connect to the data source.

If any keywords are repeated in the connection string, the driver uses the value associated with the first occurrence of the keyword. If the **DSN** and **DRIVER** keywords are included in the same connection string, the Driver Manager and the driver use whichever keyword appears first. The following table describes the attribute values of the **DSN**, **DRIVER**, **UID**, and **PWD** keywords.

Keyword	Attribute value description
DSN	Name of a data source as returned by SQLDataSources or the data sources dialog box of SQLDriverConnect .
DRIVER	Description of the driver as returned by the SQLDrivers function. For example, Rdb or SQL Server.
UID	A user ID.
PWD	The password corresponding to the user ID, or an empty string if there is no password for the user ID (PWD=;).

Driver Manager Guidelines

The Driver Manager constructs a connection string to pass to the driver in the *szConnStrIn* argument of the driver's **SQLDriverConnect** function. Note that the Driver Manager does not modify the *szConnStrIn* argument passed to it by the application.

If the connection string specified by the application contains the **DSN** keyword or does not contain either the **DSN** or **DRIVER** keywords, the action of the Driver Manager is based on the value of the *fDriverCompletion* argument:

- **SQL_DRIVER_PROMPT**: The Driver Manager displays the Data

SQLGetInfo

Sources dialog box. It constructs a connection string from the data source name returned by the dialog box and any other keywords passed to it by the application. If the data source name returned by the dialog box is empty, the Driver Manager specifies the keyword-value pair DSN=Default.

- **SQL_DRIVER_COMPLETE** or **SQL_DRIVER_COMPLETE_REQUIRED**: If the connection string specified by the application includes the **DSN** keyword, the Driver Manager copies the connection string specified by the application. Otherwise, it takes the same actions as it does when *fDriverCompletion* is **SQL_DRIVER_PROMPT**.
- **SQL_DRIVER_NOPROMPT**: The Driver Manager copies the connection string specified by the application.

If the connection string specified by the application contains the **DRIVER** keyword, the Driver Manager copies the connection string specified by the application.

Using the connection string it has constructed, the Driver Manager determines which driver to use, loads that driver, and passes the connection string it has constructed to the driver; for more information about the interaction of the Driver Manager and the driver, see the “Comments” section in **SQLConnect**. If the connection string contains the **DSN** keyword or does not contain either the **DSN** or the **DRIVER** keyword, the Driver Manager determines which driver to use as follows:

1. If the connection string contains the **DSN** keyword, the Driver Manager retrieves the driver associated with the data source from the ODBC.INI file or registry.
2. If the connection string does not contain the **DSN** keyword or the data source is not found, the Driver Manager retrieves the driver associated with the Default data source from the ODBC.INI file or registry. However, the Driver Manager does not change the value of the **DSN** keyword in the connection string.
3. If the data source is not found and the Default data source is not found, the Driver Manager returns **SQL_ERROR** with **SQLSTATE** IM002 (Data source not found and no default driver specified).

Driver Guidelines

The driver checks if the connection string passed to it by the Driver Manager contains the **DSN** or **DRIVER** keyword. If the connection string contains the **DRIVER** keyword, the driver cannot retrieve information

about the data source from the ODBC.INI file or registry. If the connection string contains the **DSN** keyword or does not contain either the **DSN** or the **DRIVER** keyword, the driver can retrieve information about the data source from the ODBC.INI file or registry as follows:

1. If the connection string contains the **DSN** keyword, the driver retrieves the information for the specified data source.
2. If the connection string does not contain the **DSN** keyword or the specified data source is not found, the driver retrieves the information for the Default data source.

The driver uses any information it retrieves from the ODBC.INI file or registry to augment the information passed to it in the connection string. If the information in the ODBC.INI file or registry duplicates information in the connection string, the driver uses the information in the connection string.

Based on the value of *fDriverCompletion*, the driver prompts the user for connection information, such as the user ID and password, and connects to the data source:

- **SQL_DRIVER_PROMPT**: The driver displays a dialog box, using the values from the connection string and ODBC.INI file or registry (if any) as initial values. When the user exits the dialog box, the driver connects to the data source. It also constructs a connection string from the value of the **DSN** or **DRIVER** keyword in *szConnStrIn* and the information returned from the dialog box. It places this connection string in the buffer referenced by *szConnStrOut*.
- **SQL_DRIVER_COMPLETE** or **SQL_DRIVER_COMPLETE_REQUIRED**: If the connection string contains enough information, and that information is correct, the driver connects to the data source and copies *szConnStrIn* to *szConnStrOut*. If any information is missing or incorrect, the driver takes the same actions as it does when *fDriverCompletion* is **SQL_DRIVER_PROMPT**, except that if *fDriverCompletion* is **SQL_DRIVER_COMPLETE_REQUIRED**, the driver disables the controls for any information not required to connect to the data source.
- **SQL_DRIVER_NOPROMPT**: If the connection string contains enough information, the driver connects to the data source and copies *szConnStrIn* to *szConnStrOut*. Otherwise, the driver returns **SQL_ERROR** for **SQLDriverConnect**.

On successful connection to the data source, the driver also sets *pcbConnStrOut* to the length of *szConnStrOut*.

SQLGetInfo

SQLGetInfo

If the user cancels a dialog box presented by the Driver Manager or the driver, **SQLDriverConnect** returns `SQL_NO_DATA_FOUND`.

For information about how the Driver Manager and the driver interact during the connection process, see **SQLConnect**.

If a driver supports **SQLDriverConnect**, the driver keyword section of the ODBC.INF file for the driver must contain the **ConnectFunctions** keyword with the second character set to "Y".

Connection Options

The `SQL_LOGIN_TIMEOUT` connection option, set using **SQLSetConnectOption**, defines the number of seconds to wait for a login request to complete before returning to the application. If the user is prompted to complete the connection string, a waiting period for each login request begins after the user has dismissed each dialog box.

The driver opens the connection in `SQL_MODE_READ_WRITE` access mode by default. To set the access mode to `SQL_MODE_READ_ONLY`, the application must call **SQLSetConnectOption** with the `SQL_ACCESS_MODE` option prior to calling **SQLDriverConnect**.

If a default translation DLL is specified in the ODBC.INI file or registry for the data source, the driver loads it. A different translation DLL can be loaded by calling **SQLSetConnectOption** with the `SQL_TRANSLATE_DLL` option. A translation option can be specified by calling **SQLSetConnectOption** with the `SQL_TRANSLATE_OPTION` option.

Related Functions

For information about	See
Allocating a connection handle	SQLAllocConnect
Discovering and enumerating values required to connect to a data source	SQLBrowseConnect (extension)
Connecting to a data source	SQLConnect
Disconnecting from a data source	SQLDisconnect
Returning driver descriptions and attributes	SQLDrivers (extension)
Freeing a connection handle	SQLFreeConnect
Setting a connection option	SQLSetConnectOption (extension)

SQLDrivers (Windows)



Extension Level 2 **SQLDrivers** lists driver descriptions and driver attribute keywords. This function is implemented solely by the Driver Manager.

Syntax

RETCODE **SQLDrivers**(*henv*, *fDirection*, *szDriverDesc*, *cbDriverDescMax*, *pcbDriverDesc*, *szDriverAttributes*, *cbDrvrAttrMax*, *pcbDrvrAttr*)

The **SQLDrivers** function accepts the following arguments:

Type	Argument	Use	Description
HENV	<i>henv</i>	Input	Environment handle.
UWORD	<i>fDirection</i>	Input	Determines whether the Driver Manager fetches the next driver description in the list (SQL_FETCH_NEXT) or whether the search starts from the beginning of the list (SQL_FETCH_FIRST).
UCHAR FAR *	<i>szDriverDesc</i>	Output	Pointer to storage for the driver description.
SWORD	<i>cbDriverDescMax</i>	Input	Maximum length of the <i>szDriverDesc</i> buffer.
SWORD FAR *	<i>pcbDriverDesc</i>	Output	Total number of bytes (excluding the null termination byte) available to return in <i>szDriverDesc</i> . If the number of bytes available to return is greater than or equal to <i>cbDriverDescMax</i> , the driver description in <i>szDriverDesc</i> is truncated to <i>cbDriverDescMax</i> – 1 bytes.
UCHAR FAR *	<i>szDriverAttributes</i>	Output	Pointer to storage for the list of driver attribute value pairs (see "Comments").
SWORD	<i>cbDrvrAttrMax</i>	Input	Maximum length of the <i>szDriverAttributes</i> buffer.
SWORD FAR *	<i>pcbDrvrAttr</i>	Output	Total number of bytes (excluding the null termination byte) available to return in <i>szDriverAttributes</i> . If the number of bytes available to return is greater than or equal to <i>cbDrvrAttrMax</i> , the list of attribute value pairs in <i>szDriverAttributes</i> is

SQLGetStmtOption

SQLGetStmtOption

truncated to *cbDrvrAttrMax* – 1 bytes.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_NO_DATA_FOUND, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLDrivers** returns either SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLDrivers** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATEs returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	(DM) Driver Manager-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	(DM) The buffer <i>szDriverDesc</i> was not large enough to return the entire driver description, so the description was truncated. The argument <i>pcbDriverDesc</i> contains the length of the entire driver description. (Function returns SQL_SUCCESS_WITH_INFO.) (DM) The buffer <i>szDriverAttributes</i> was not large enough to return the entire list of attribute value pairs, so the list was truncated. The argument <i>pcbDrvrAttr</i> contains the length of the untruncated list of attribute value pairs. (Function returns SQL_SUCCESS_WITH_INFO.)
S1000	General error	(DM) An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	(DM) The Driver Manager was unable to allocate memory required to support execution or completion of the function.
S1090	Invalid string or buffer length	(DM) The value specified for argument <i>cbDriverDescMax</i> was less than 0. (DM) The value specified for argument

		<i>cbDrvAttrMax</i> was less than 0 or equal to 1.
SQLSTATE	Error	Description
S1103	Direction option out of range	(DM) The value specified for the argument <i>fDirection</i> was not equal to SQL_FETCH_FIRST or SQL_FETCH_NEXT.

Comments

SQLDrivers returns the driver description in the *szDriverDesc* argument. It returns additional information about the driver in the *szDriverAttributes* argument as a list of keyword-value pairs. Each pair is terminated with a null byte, and the entire list is terminated with a null byte (that is, two null bytes mark the end of the list). For example, a dBASE driver might return the following list of attributes ("\\0" represents a null byte):

FileUsage=1\\0FileExtns=*.dbf\\0\\0

If *szDriverAttributes* is not large enough to hold the entire list, the list is truncated, **SQLDrivers** returns SQLSTATE 01004 (Data truncated), and the length of the list (excluding the final null termination byte) is returned in *pcbDrvAttr*.

Driver attribute keywords are added from the ODBC.INF file when the driver is installed.

An application can call **SQLDrivers** multiple times to retrieve all driver descriptions. The Driver Manager retrieves this information from the ODBCINST.INI file or the registry. When there are no more driver descriptions, **SQLDrivers** returns SQL_NO_DATA_FOUND. If **SQLDrivers** is called with SQL_FETCH_NEXT immediately after it returns SQL_NO_DATA_FOUND, it returns the first driver description.

If SQL_FETCH_NEXT is passed to **SQLDrivers** the very first time it is called, **SQLDrivers** returns the first data source name.

Because **SQLDrivers** is implemented in the Driver Manager, it is supported for all drivers regardless of a particular driver's conformance level.

Related Functions

For information about	BEGIN BREAK	See
	Discovering and listing values required to connect to a data source	SQLBrowseConnect (extension) SQLGetStmtOption

Connecting to a data source
 Returning data source names
 Connecting to a data source using a
 connection string or dialog box

SQLConnect
SQLDataSources (extension)
SQLDriverConnect (extension)

SQLError



Core **SQLError** returns error or status information.

Syntax RETCODE **SQLError**(*henv*, *hdbc*, *hstmt*, *szSqlState*, *pfNativeError*, *szErrorMsg*, *cbErrorMsgMax*, *pcbErrorMsg*)

The **SQLError** function accepts the following arguments.

Type	Argument	Use	Description
HENV	<i>henv</i>	Input	Environment handle or SQL_NULL_HENV.
HDBC	<i>hdbc</i>	Input	Connection handle or SQL_NULL_HDBC.
HSTMT	<i>hstmt</i>	Input	Statement handle or SQL_NULL_HSTMT.
UCHAR FAR *	<i>szSqlState</i>	Output	SQLSTATE as null-terminated string. For a list of SQLSTATES, see Appendix A, "ODBC Error Codes."
SDWORD FAR *	<i>pfNativeError</i>	Output	Native error code (specific to the data source).
UCHAR FAR *	<i>szErrorMsg</i>	Output	Pointer to storage for the error message text.
SWORD	<i>cbErrorMsgMax</i>	Input	Maximum length of the <i>szErrorMsg</i> buffer. This must be less than or equal to SQL_MAX_MESSAGE_LENGTH – 1.
SWORD FAR *	<i>pcbErrorMsg</i>	Output	Pointer to the total number of bytes (excluding the null termination byte) available to return in <i>szErrorMsg</i> . If the number of bytes available to return is greater than or equal to <i>cbErrorMsgMax</i> , the error message text in <i>szErrorMsg</i> is truncated to <i>cbErrorMsgMax</i> – 1 bytes.

Returns

SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_NO_DATA_FOUND, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics

SQLError does not post error values for itself. **SQLError** returns SQL_NO_DATA_FOUND when it is unable to retrieve any error information, (in which case *szSqlState* equals 00000). If **SQLError** cannot access error values for any reason that would normally return SQL_ERROR, **SQLError** returns SQL_ERROR but does not post any error values. If the buffer for the error message is too short, **SQLError** returns SQL_SUCCESS_WITH_INFO but, again, does not return a SQLSTATE value for **SQLError**.

To determine that a truncation occurred in the error message, an application can compare *cbErrorMsgMax* to the actual length of the message text written to *pcbErrorMsg*.

Comments

An application typically calls **SQLError** when a previous call to an ODBC function returns SQL_ERROR or SQL_SUCCESS_WITH_INFO. However, any ODBC function can post zero or more errors each time it is called, so an application can call **SQLError** after any ODBC function call.

SQLError retrieves an error from the data structure associated with the rightmost non-null handle argument. An application requests error information as follows:

- To retrieve errors associated with an environment, the application passes the corresponding *henv* and includes SQL_NULL_HDBC and SQL_NULL_HSTMT in *hdbc* and *hstmt*, respectively. The driver returns the error status of the ODBC function most recently called with the same *henv*.
- To retrieve errors associated with a connection, the application passes the corresponding *hdbc* plus an *hstmt* equal to SQL_NULL_HSTMT. In such a case, the driver ignores the *henv* argument. The driver returns the error status of the ODBC function most recently called with the *hdbc*.
- To retrieve errors associated with a statement, an application passes the corresponding *hstmt*. If the call to **SQLError** contains a valid *hstmt*, the driver ignores the *hdbc* and *henv* arguments. The driver returns the error status of the ODBC function most recently called with the *hstmt*.
- To retrieve multiple errors for a function call, an application calls **SQLError** multiple times. For each error, the driver returns SQL_SUCCESS and removes that error from the list of available errors.

SQLGetStmtOption

When there is no additional information for the rightmost non-null handle, **SQLError** returns SQL_NO_DATA_FOUND. In this case, *szSqlState* equals 00000 (Success), *pfNativeError* is undefined, *pcbErrorMsg* equals 0, and *szErrorMsg* contains a single null termination byte (unless *cbErrorMsgMax* equals 0).

The Driver Manager stores error information in its *henv*, *hdbc*, and *hstmt* structures. Similarly, the driver stores error information in its *henv*, *hdbc*, and *hstmt* structures. When the application calls **SQLError**, the Driver Manager checks if there are any errors in its structure for the specified handle. If there are errors for the specified handle, it returns the first error; if there are no errors, it calls **SQLError** in the driver.

The Driver Manager can store up to 64 errors with an *henv* and its associated *hdbcs* and *hstmts*. When this limit is reached, the Driver Manager discards any subsequent errors posted on the Driver Manager's *henv*, *hdbcs*, or *hstmts*. The number of errors that a driver can store is driver-dependent.

An error is removed from the structure associated with a handle when **SQLError** is called for that handle and returns that error. All errors stored for a given handle are removed when that handle is used in a subsequent function call. For example, errors on an *hstmt* that were returned by **SQLExecDirect** are removed when **SQLExecDirect** or **SQLTables** is called with that *hstmt*. The errors stored on a given handle are not removed as the result of a call to a function using an associated handle of a different type. For example, errors on an *hdbc* that were returned by **SQLNativeSql** are not removed when **SQLError** or **SQLExecDirect** is called with an *hstmt* associated with that *hdbc*.

For more information about error codes, see Appendix A, "ODBC Error Codes."

Related Functions None.

Note: The Driver Manager does not exist on non-Windows systems. The Driver returns the errors directly.

SQLExecDirect



Core **SQLExecDirect** executes a preparable statement, using the current values of the parameter marker variables if any parameters exist in the statement. **SQLExecDirect** is the fastest way to submit an SQL statement for one-time execution.

Syntax RETCODE **SQLExecDirect**(*hstmt*, *szSqlStr*, *cbSqlStr*)

The **SQLExecDirect** function uses the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UCHAR FAR *	<i>szSqlStr</i>	Input	SQL statement to be executed.
SDWORD	<i>cbSqlStr</i>	Input	Length of <i>szSqlStr</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_NEED_DATA, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLExecDirect** returns either SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLExecDirect** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	The argument <i>szSqlStr</i> contained an SQL statement that contained a character or binary parameter or literal and the value exceeded the maximum length of the associated table column. The argument <i>szSqlStr</i> contained an SQL statement that contained a numeric parameter or literal and the fractional part of the value was truncated. The argument <i>szSqlStr</i> contained an SQL statement that contained a date or time parameter or literal and a timestamp value was truncated.

SQLGetTypeInfo

SQLGetTypeInfo

SQLSTATE	Error	Description
01006	Privilege not revoked	The argument <i>szSqlStr</i> contained a REVOKE statement and the user did not have the specified privilege. (Function returns SQL_SUCCESS_WITH_INFO.)
01S03	No rows updated or deleted	The argument <i>szSqlStr</i> contained a positioned update or delete statement and no rows were updated or deleted. (Function returns SQL_SUCCESS_WITH_INFO.)
01S04	More than one row updated or deleted	The argument <i>szSqlStr</i> contained a positioned update or delete statement and more than one row was updated or deleted. (Function returns SQL_SUCCESS_WITH_INFO.)
07001	Wrong number of parameters	The number of parameters specified in SQLBindParameter was less than the number of parameters in the SQL statement contained in the argument <i>szSqlStr</i> .
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
21S01	Insert value list does not match column list	The argument <i>szSqlStr</i> contained an INSERT statement and the number of values to be inserted did not match the degree of the derived table.
21S02	Degree of derived table does not match column list	The argument <i>szSqlStr</i> contained a CREATE VIEW statement and the number of names specified is not the same degree as the derived table defined by the query specification.
22003	Numeric value out of range	The argument <i>szSqlStr</i> contained an SQL statement which contained a numeric parameter or literal and the value caused the whole (as opposed to fractional) part of the number to be truncated when assigned to the associated table column.
22005	Error in assignment	The argument <i>szSqlStr</i> contained an SQL statement that contained a parameter or literal and the value was incompatible with the data type of the associated table column.
SQLSTATE	Error	Description
22008	Datetime field overflow	The argument <i>szSqlStr</i> contained an SQL statement that contained a date, time, or timestamp parameter or literal and the value

		was, respectively, an invalid date, time, or timestamp.
22012	Division by zero	The argument <i>szSqlStr</i> contained an SQL statement which contained an arithmetic expression which caused division by zero.
23000	Integrity constraint violation	The argument <i>szSqlStr</i> contained an SQL statement which contained a parameter or literal. The parameter value was NULL for a column defined as NOT NULL in the associated table column, a duplicate value was supplied for a column constrained to contain only unique values, or some other integrity constraint was violated.
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called. A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called. The argument <i>szSqlStr</i> contained a positioned update or delete statement and the cursor was positioned before the start of the result set or after the end of the result set.
34000	Invalid cursor name	The argument <i>szSqlStr</i> contained a positioned update or delete statement and the cursor referenced by the statement being executed was not open.
37000	Syntax error or access violation	The argument <i>szSqlStr</i> contained an SQL statement that was not preparable or contained a syntax error.
40001	Serialization failure	The transaction to which the SQL statement contained in the argument <i>szSqlStr</i> belonged was terminated to prevent deadlock.
42000	Syntax error or access violation	The user did not have permission to execute the SQL statement contained in the argument <i>szSqlStr</i> .
SQLSTATE	Error	Description
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S0001	Base table or view already exists	The argument <i>szSqlStr</i> contained a CREATE TABLE or CREATE VIEW statement and the table name or view name specified already exists.

SQLGetTypeInfo

S0002	Table or view not found	<p>The argument <i>szSqlStr</i> contained a DROP TABLE or a DROP VIEW statement and the specified table name or view name did not exist.</p> <p>The argument <i>szSqlStr</i> contained an ALTER TABLE statement and the specified table name did not exist.</p> <p>The argument <i>szSqlStr</i> contained a CREATE VIEW statement and a table name or view name defined by the query specification did not exist.</p> <p>The argument <i>szSqlStr</i> contained a CREATE INDEX statement and the specified table name did not exist.</p> <p>The argument <i>szSqlStr</i> contained a GRANT or REVOKE statement and the specified table name or view name did not exist.</p> <p>The argument <i>szSqlStr</i> contained a SELECT statement and a specified table name or view name did not exist.</p> <p>The argument <i>szSqlStr</i> contained a DELETE, INSERT, or UPDATE statement and the specified table name did not exist.</p> <p>The argument <i>szSqlStr</i> contained a CREATE TABLE statement and a table specified in a constraint (referencing a table other than the one being created) did not exist.</p>
S0011	Index already exists	The argument <i>szSqlStr</i> contained a CREATE INDEX statement and the specified index name already existed.
S0012	Index not found	The argument <i>szSqlStr</i> contained a DROP INDEX statement and the specified index name did not exist.
SQLSTATE	Error	Description
S0021	Column already exists	The argument <i>szSqlStr</i> contained 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	The argument <i>szSqlStr</i> contained a CREATE INDEX statement and one or more of the column names specified in the column list did not exist.

		<p>The argument <i>szSqlStr</i> contained a GRANT or REVOKE statement and a specified column name did not exist.</p> <p>The argument <i>szSqlStr</i> contained a SELECT, DELETE, INSERT, or UPDATE statement and a specified column name did not exist.</p> <p>The argument <i>szSqlStr</i> contained a CREATE TABLE statement and a column specified in a constraint (referencing a table other than the one being created) did not exist.</p>
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLERROR in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	<p>Asynchronous processing was enabled for the <i>hstmt</i>. The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i>. Then the function was called again on the <i>hstmt</i>.</p> <p>The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.</p>
S1009	Invalid argument value	(DM) The argument <i>szSqlStr</i> was a null pointer.
SQLSTATE	Error	Description
S1010	Function sequence error	<p>(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called.</p> <p>(DM) SQLExecute, SQLExecDirect, or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.</p>
S1090	Invalid string or buffer length	<p>(DM) The argument <i>cbSqlStr</i> was less than or equal to 0, but not equal to SQL_NTS.</p> <p>A parameter value, set with SQLBindParameter, was a null pointer and</p> <p style="text-align: right;">SQLGetTypeInfo</p>

SQLGetTypeInfo

		the parameter length value was not 0, SQL_NULL_DATA, SQL_DATA_AT_EXEC, or less than or equal to SQL_LEN_DATA_AT_EXEC_OFFSET. A parameter value, set with SQLBindParameter , was not a null pointer and the parameter length value was less than 0, but was not SQL_NTS, SQL_NULL_DATA, SQL_DATA_AT_EXEC, or less than or equal to SQL_LEN_DATA_AT_EXEC_OFFSET.
S1109	Invalid cursor position	The argument <i>szSqlStr</i> contained a positioned update or delete statement and the cursor was positioned (by SQLSetPos or SQLExtendedFetch) on a row for which the value in the <i>rgfRowStatus</i> array in SQLExtendedFetch was SQL_ROW_DELETED or SQL_ROW_ERROR.
S1C00	Driver not capable	The combination of the current settings of the SQL_CONCURRENCY and SQL_CURSOR_TYPE statement options was not supported by the driver or data source.
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

Comments

The application calls **SQLExecDirect** to send an SQL statement to the data source. The driver modifies the statement to use the form of SQL used by the data source, then submits it to the data source. In particular, the driver modifies the escape clauses used to define ODBC-specific SQL. For a description of SQL statement grammar, see Appendix C, “SQL Grammar.”

The application can include one or more parameter markers in the SQL statement. To include a parameter marker, the application embeds a question mark (?) into the SQL statement at the appropriate position.

If the SQL statement is a **SELECT** statement, and if the application called **SQLSetCursorName** to associate a cursor with an *hstmt*, then the driver uses the specified cursor. Otherwise, the driver generates a cursor name.

If the data source is in manual-commit mode (requiring explicit transaction initiation), and a transaction has not already been initiated, the driver initiates a transaction before it sends the SQL statement.

If an application uses **SQLExecDirect** to submit a **COMMIT** or **ROLLBACK** statement, it will not be interoperable between DBMS products. To commit or roll back a transaction, call **SQLTransact**.

If **SQLExecDirect** encounters a data-at-execution parameter, it returns **SQL_NEED_DATA**. The application sends the data using **SQLParamData** and **SQLPutData**. See **SQLBindParameter**, **SQLParamOptions**, **SQLParamData**, and **SQLPutData** for more information.

Code Example

See **SQLBindCol**, **SQLExtendedFetch**, **SQLGetData**, and **SQLProcedures**.

Related Functions**For information about****See**

Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Executing a prepared SQL statement	SQLExecute
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
Fetching a row of data	SQLFetch
Returning a cursor name	SQLGetCursorName
Fetching part or all of a column of data	SQLGetData (extension)
Returning the next parameter to send data for	SQLParamData (extension)
For information about	See
Preparing a statement for execution	SQLPrepare
Sending parameter data at execution time	SQLPutData (extension)
Setting a cursor name	SQLSetCursorName
Setting a statement option	SQLSetStmtOption (extension)
Executing a commit or rollback operation	SQLTransact

SQLExecute

ADBC

Core **SQLExecute** executes a prepared statement, using the current values of the parameter marker variables if any parameter markers exist in the statement.

Syntax RETCODE **SQLExecute**(*hstmt*)

The **SQLExecute** statement accepts the following argument.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_NEED_DATA, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLExecute** returns either SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLExecute** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	The prepared statement associated with the <i>hstmt</i> contained a character or binary parameter or literal and the value exceeded the maximum length of the associated table column. The prepared statement associated with the <i>hstmt</i> contained a numeric parameter or literal and the fractional part of the value was truncated. The prepared statement associated with the <i>hstmt</i> contained a date or time parameter or literal and a timestamp value was truncated.
01006	Privilege not revoked	The prepared statement associated with the <i>hstmt</i> was REVOKE and the user did not have the specified privilege. (Function returns SQL_SUCCESS_WITH_INFO.)
SQLSTATE	Error	Description

01S03	No rows updated or deleted	The prepared statement associated with the <i>hstmt</i> was a positioned update or delete statement and no rows were updated or deleted. (Function returns SQL_SUCCESS_WITH_INFO.)
01S04	More than one row updated or deleted	The prepared statement associated with the <i>hstmt</i> was a positioned update or delete statement and more than one row was updated or deleted. (Function returns SQL_SUCCESS_WITH_INFO.)
07001	Wrong number of parameters	The number of parameters specified in SQLBindParameter was less than the number of parameters in the prepared statement associated with the <i>hstmt</i> .
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
22003	Numeric value out of range	The prepared statement associated with the <i>hstmt</i> contained a numeric parameter and the parameter value caused the whole (as opposed to fractional) part of the number to be truncated when assigned to the associated table column.
22005	Error in assignment	The prepared statement associated with the <i>hstmt</i> contained a parameter and the value was incompatible with the data type of the associated table column.
22008	Datetime field overflow	The prepared statement associated with the <i>hstmt</i> contained a date, time, or timestamp parameter or literal and the value was, respectively, an invalid date, time, or timestamp.
22012	Division by zero	The prepared statement associated with the <i>hstmt</i> contained an arithmetic expression which caused division by zero.
23000	Integrity constraint violation	The prepared statement associated with the <i>hstmt</i> contained a parameter. The parameter value was NULL for a column defined as NOT NULL in the associated table column, a duplicate value was supplied for a column constrained to contain only unique values, or some other integrity constraint was violated.
SQLSTATE	Error	Description
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLMoreResults

		<p>SQLFetch or SQLExtendedFetch had been called.</p> <p>A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called.</p> <p>The prepared statement associated with the <i>hstmt</i> contained a positioned update or delete statement and the cursor was positioned before the start of the result set or after the end of the result set.</p>
40001	Serialization failure	The transaction to which the prepared statement associated with the <i>hstmt</i> belonged was terminated to prevent deadlock.
42000	Syntax error or access violation	The user did not have permission to execute the prepared statement associated with the <i>hstmt</i> .
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	<p>Asynchronous processing was enabled for the <i>hstmt</i>. The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i>. Then the function was called again on the <i>hstmt</i>.</p> <p>The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.</p>
SQLSTATE	Error	Description
S1010	Function sequence error	<p>(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called.</p> <p>(DM) SQLExecute, SQLExecDirect, or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all</p>

		data-at-execution parameters or columns. (DM) The <i>hstmt</i> was not prepared. Either the <i>hstmt</i> was not in an executed state, or a cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called. The <i>hstmt</i> was not prepared. It was in an executed state and either no result set was associated with the <i>hstmt</i> or SQLFetch or SQLExtendedFetch had not been called.
S1090	Invalid string or buffer length	A parameter value, set with SQLBindParameter , was a null pointer and the parameter length value was not 0, SQL_NULL_DATA, SQL_DATA_AT_EXEC, or less than or equal to SQL_LEN_DATA_AT_EXEC_OFFSET. A parameter value, set with SQLBindParameter , was not a null pointer and the parameter length value was less than 0, but was not SQL_NTS, SQL_NULL_DATA, or SQL_DATA_AT_EXEC, or less than or equal to SQL_LEN_DATA_AT_EXEC_OFFSET.
S1109	Invalid cursor position	The prepared statement was a positioned update or delete statement and the cursor was positioned (by SQLSetPos or SQLExtendedFetch) on a row for which the value in the <i>rgfRowStatus</i> array in SQLExtendedFetch was SQL_ROW_DELETED or SQL_ROW_ERROR.
S1C00	Driver not capable	The combination of the current settings of the SQL_CONCURRENCY and SQL_CURSOR_TYPE statement options was not supported by the driver or data source.
SQLSTATE Error		Description
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

SQLExecute can return any SQLSTATE that can be returned by **SQLPrepare** based on when the data source evaluates the SQL statement associated with the *hstmt*.

SQLMoreResults

Comments

SQLExecute executes a statement prepared by **SQLPrepare**. Once the application processes or discards the results from a call to **SQLExecute**, the application can call **SQLExecute** again with new parameter values.

To execute a **SELECT** statement more than once, the application must call **SQLFreeStmt** with the **SQL_CLOSE** parameter before reissuing the **SELECT** statement.

If the data source is in manual-commit mode (requiring explicit transaction initiation), and a transaction has not already been initiated, the driver initiates a transaction before it sends the SQL statement.

If an application uses **SQLPrepare** to prepare and **SQLExecute** to submit a **COMMIT** or **ROLLBACK** statement, it will not be interoperable between DBMS products. To commit or roll back a transaction, call **SQLTransact**.

If **SQLExecute** encounters a data-at-execution parameter, it returns **SQL_NEED_DATA**. The application sends the data using **SQLParamData** and **SQLPutData**. See **SQLBindParameter**, **SQLParamOptions**, **SQLParamData**, and **SQLPutData** for more information.

Code Example

See **SQLBindParameter**, **SQLParamOptions**, **SQLPutData**, and **SQLSetPos**.

Related Functions

For information about	See
-----------------------	-----

Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Executing an SQL statement	SQLExecDirect
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
Fetching a row of data	SQLFetch
Freeing a statement handle	SQLFreeStmt
Returning a cursor name	SQLGetCursorName
Fetching part or all of a column of data	SQLGetData (extension)

For information about	See
Returning the next parameter to send data for	SQLParamData (extension)
Preparing a statement for execution	SQLPrepare
Sending parameter data at execution time	SQLPutData (extension)
Setting a cursor name	SQLSetCursorName
Setting a statement option	SQLSetStmtOption (extension)

SQLExtendedFetch

ADBC

Extension Level 2 **SQLExtendedFetch** extends the functionality of **SQLFetch** in the following ways:

- It returns rowset data (one or more rows), in the form of an array, for each bound column.
- It scrolls through the result set according to the setting of a scroll-type argument.

SQLExtendedFetch works in conjunction with **SQLSetStmtOption**.

To fetch one row of data at a time in a forward direction, an application should call **SQLFetch**.

For more information about scrolling through result sets, see “Using Block and Scrollable Cursors” in Chapter 7, “Retrieving Results.”

Syntax

RETCODE **SQLExtendedFetch**(*hstmt*, *fFetchType*, *irow*, *pcrow*, *rgfRowStatus*)
The **SQLExtendedFetch** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UWORD	<i>fFetchType</i>	Input	Type of fetch. For more information, see the “Comments” section.
SDWORD	<i>irow</i>	Input	Number of the row to fetch. For more information, see the “Comments” section.
UDWORD FAR *	<i>pcrow</i>	Output	Number of rows actually fetched.
UWORD FAR *	<i>rgfRowStatus</i>	Output	An array of status values. For more information, see the “Comments” section.

Returns

SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_NO_DATA_FOUND, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics

When **SQLExtendedFetch** returns either SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLExtendedFetch** and explains each one in the context of this function; the notation “(DM)” precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
----------	-------	-------------

01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	The data returned for one or more columns was truncated. String values are right truncated. For numeric values, the fractional part of number was 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.)
07006	Restricted data type attribute violation	A data value could not be converted to the C data type specified by <i>fCType</i> in SQLBindCol .
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
22003	Numeric value out of range	Returning the numeric value (as numeric or string) for one or more columns would have caused the whole (as opposed to fractional) part of the number to be truncated. Returning the binary value for one or more columns would have caused a loss of binary significance. For more information, see Appendix D, "Data Types."
22012	Division by zero	A value from an arithmetic expression was returned which resulted in division by zero.
24000	Invalid cursor state	The <i>hstmt</i> was in an executed state but no result set was associated with the <i>hstmt</i> .
40001	Serialization failure	The transaction in which the fetch was executed was terminated to prevent deadlock.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hdbc</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
SQLSTATE Error		Description

SQLNativeSql

S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1002	Invalid column number	A column number specified in the binding for one or more columns was greater than the number of columns in the result set. Column 0 was bound with SQLBindCol and the SQL_USE_BOOKMARKS statement option was set to SQL_UB_OFF.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1010	Function sequence error	(DM) The specified <i>hstmt</i> was not in an executed state. The function was called without first calling SQLExecDirect , SQLExecute , or a catalog function.. (DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns. (DM) SQLExtendedFetch was called for an <i>hstmt</i> after SQLFetch was called and before SQLFreeStmt was called with the SQL_CLOSE option.
S1106	Fetch type out of range	(DM) The value specified for the argument <i>fFetchType</i> was invalid (see "Comments"). The value of the SQL_CURSOR_TYPE statement option was SQL_CURSOR_FORWARD_ONLY and the value of argument <i>fFetchType</i> was not SQL_FETCH_NEXT.

SQLSTATE	Error	Description
S1107	Row value out of range	The value specified with the <code>SQL_CURSOR_TYPE</code> statement option was <code>SQL_CURSOR_KEYSET_DRIVEN</code> , but the value specified with the <code>SQL_KEYSET_SIZE</code> statement option was greater than 0 and less than the value specified with the <code>SQL_ROWSET_SIZE</code> statement option.
S1111	Invalid bookmark value	The argument <i>fFetchType</i> was <code>SQL_FETCH_BOOKMARK</code> and the bookmark specified in the <i>irow</i> argument was not valid.
S1C00	Driver not capable	Driver or data source does not support the specified fetch type. The driver or 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. This error only applies when the SQL data type of the column was mapped to a driver-specific SQL data type. The argument <i>fFetchType</i> was <code>SQL_FETCH_RESUME</code> and the driver supports ODBC 2.0.
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , <code>SQL_QUERY_TIMEOUT</code> .

Comments **SQLExtendedFetch** returns one rowset of data to the application. An application cannot mix calls to **SQLExtendedFetch** and **SQLFetch** for the same cursor.

An application specifies the number of rows in the rowset by calling **SQLSetStmtOption** with the `SQL_ROWSET_SIZE` statement option.

Binding

If any columns in the result set have been bound with **SQLBindCol**, the driver converts the data for the bound columns as necessary and stores it in the locations bound to those columns. The result set can be bound in a column-wise (the default) or row-wise fashion.

Column-Wise Binding

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.) For each column to be bound, the application:

1. Allocates an array of data storage buffers. The array has as many elements as there are rows in the rowset, plus an additional element if the application will search for key values or append new rows of data. Each buffer's size is the maximum size of the C data that can be returned for the column. For example, when the C data type is `SQL_C_DEFAULT`, each buffer's size is the column length. When the C data type is `SQL_C_CHAR`, each buffer's size is the display size of the data. For more information, see "Converting Data from SQL to C Data Types" and "Precision, Scale, Length, and Display Size" in Appendix D, "Data Types."
2. Allocates an array of `SDWORD`s to hold the number of bytes available to return for each row in the column. The array has as many elements as there are rows in the rowset.
3. Calls **SQLBindCol**:
 - The *rgbValue* argument specifies the address of the data storage array.
 - The *cbValueMax* argument specifies the size of each buffer in the data storage array.
 - The *pcbValue* argument specifies the address of the number-of-bytes array.

When the application calls **SQLExtendedFetch**, the driver retrieves the data and the number of bytes available to return and stores them in the buffers allocated by the application:

- For each bound column, the driver stores the data in the *rgbValue* buffer bound to the column. It stores the first row of data at the start of the buffer and each subsequent row of data at an offset of *cbValueMax* bytes from the data for the previous row.
- For each bound column, the driver stores the number of bytes available to return in the *pcbValue* buffer bound to the column. This is the number of bytes available prior to calling **SQLExtendedFetch**. (If the number of bytes available to return cannot be determined in advance, the driver sets *pcbValue* to `SQL_NO_TOTAL`. If the data for the column is `NULL`, the driver sets *pcbValue* to `SQL_NULL_DATA`.) It stores the number of bytes available to return for the first row at the start of the buffer and the number of bytes available to return for each subsequent row at an offset of **sizeof(SDWORD)** from the value for the previous row.

Row-Wise Binding

To bind a result set in row-wise fashion, an application:

1. Declares a structure that can hold a single row of retrieved data and the associated data lengths. For each bound column, the structure contains one field for the data and one `SDWORD` field for the number of bytes available to return. The data field's size is the maximum size of the C data that can be returned for the column.
2. Calls **SQLSetStmtOption** with *fOption* set to `SQL_BIND_TYPE` and *vParam* set to the size of the structure.
3. Allocates an array of these structures. The array has as many elements as there are rows in the rowset, plus an additional element if the application will search for key values or append new rows of data.
4. Calls **SQLBindCol** for each column to be bound:
 - The *rgbValue* argument specifies the address of the column's data field in the first array element.
 - The *cbValueMax* argument specifies the size of the column's data field.
 - The *pcbValue* argument specifies the address of the column's number-of-bytes field in the first array element.

When the application calls **SQLExtendedFetch**, the driver retrieves the data and the number of bytes available to return and stores them in the buffers allocated by the application:

- For each bound column, the driver stores the first row of data at the address specified by *rgbValue* for the column and each subsequent row of data at an offset of *vParam* bytes from the data for the previous row.
- For each bound column, the driver stores the number of bytes available to return for the first row at the address specified by *pcbValue* and the number of bytes available to return for each subsequent row at an offset of *vParam* bytes from the value for the previous row. This is the number of bytes available prior to calling **SQLExtendedFetch**. (If the number of bytes available to return cannot be determined in advance, the driver sets *pcbValue* to `SQL_NO_TOTAL`. If the data for the column is `NULL`, the driver sets *pcbValue* to `SQL_NULL_DATA`.)

Positioning the Cursor

The following operations require a cursor position:

- Positioned update and delete statements.

SQLNativeSql

- Calls to **SQLGetData**.
- Calls to **SQLSetPos** with the SQL_DELETE, SQL_REFRESH, and SQL_UPDATE options.

An application can specify a cursor position when it calls **SQLSetPos**. Before it executes a positioned update or delete statement or calls **SQLGetData**, the application must position the cursor by calling **SQLExtendedFetch** to retrieve a rowset; the cursor points to the first row in the rowset. To position the cursor to a different row in the rowset, the application calls **SQLSetPos**.

The following table shows the rowset and return code returned when the application requests different rowsets.

Requested Rowset	Return Code	Cursor Position	Returned Rowset
Before start of result set	SQL_NO_DATA_FOUND	Before start of result set	None. The contents of the rowset buffers are undefined.
Overlaps start of result set	SQL_SUCCESS	Row 1 of rowset	First rowset in result set.
Within result set	SQL_SUCCESS	Row 1 of rowset	Requested rowset.
Overlaps end of result set	SQL_SUCCESS	Row 1 of rowset	For rows in the rowset that overlap the result set, data is returned. For rows in the rowset outside the result set, the contents of the <i>rgbValue</i> and <i>pcbValue</i> buffers are undefined and the <i>rgfRowStatus</i> array contains SQL_ROW_NOROW.
After end of result set	SQL_NO_DATA_FOUND	After end of result set	None. The contents of the rowset buffers are undefined.

For example, suppose a result set has 100 rows and the rowset size is 5. The following table shows the rowset and return code returned by **SQLExtendedFetch** for different values of *irow* when the fetch type is SQL_FETCH_RELATIVE:

Current Rowset	<i>irow</i>	Return Code	New Rowset
1 to 5	-5	SQL_NO_DATA_FOUND	None.
1 to 5	-3	SQL_SUCCESS	1 to 5
96 to 100	5	SQL_NO_DATA_FOUND	None.
96 to 100	3	SQL_SUCCESS	99 and 100. For rows 3, 4, SQLNativeSql

and 5 in the rowset, the *rgfRowStatusArray* is set to SQL_ROW_NOROW.

Before **SQLExtendedFetch** is called the first time, the cursor is positioned before the start of the result set.

For the purpose of moving the cursor, deleted rows (that is, rows with an entry in the *rgfRowStatus* array of SQL_ROW_DELETED) are treated no differently than other rows. For example, calling **SQLExtendedFetch** with *fFetchType* set to SQL_FETCH_ABSOLUTE and *irow* set to 15 returns the rowset starting at row 15, even if the *rgfRowStatus* array for row 15 is SQL_ROW_DELETED.

Processing Errors

If an error occurs that pertains to the entire rowset, such as SQLSTATE S1T00 (Timeout expired), the driver returns SQL_ERROR and the appropriate SQLSTATE. The contents of the rowset buffers are undefined and the cursor position is unchanged.

If an error occurs that pertains to a single row, the driver:

- Sets the element in the *rgfRowStatus* array for the row to SQL_ROW_ERROR.
- Posts SQLSTATE 01S01 (Error in row) in the error queue.
- Posts zero or more additional SQLSTATES for the error after SQLSTATE 01S01 (Error in row) in the error queue.

After it has processed the error or warning, the driver continues the operation for the remaining rows in the rowset and returns SQL_SUCCESS_WITH_INFO. Thus, for each error that pertains to a single row, the error queue contains SQLSTATE 01S01 (Error in row) followed by zero or more additional SQLSTATES.

After it has processed the error, the driver fetches the remaining rows in the rowset and returns SQL_SUCCESS_WITH_INFO. Thus, for each row that returned an error, the error queue contains SQLSTATE 01S01 (Error in row) followed by zero or more additional SQLSTATES.

If the rowset contains rows that have already been fetched, the driver is not required to return SQLSTATES for errors that occurred when the rows were first fetched. It is, however, required to return SQLSTATE 01S01 (Error in row) for each row in which an error originally occurred and to return SQL_SUCCESS_WITH_INFO. For example, a static cursor that maintains a cache might cache row status information (so it can

determine which rows contain errors) but might not cache the SQLSTATE associated with those errors.

Error rows do not affect relative cursor movements. For example, suppose the result set size is 100 and the rowset size is 10. If the current rowset is rows 11 through 20 and the element in the *rgfRowStatus* array for row 11 is SQL_ROW_ERROR, calling **SQLExtendedFetch** with the SQL_FETCH_NEXT fetch type still returns rows 21 through 30.

If the driver returns any warnings, such as SQLSTATE 01004 (Data truncated), it returns warnings that apply to the entire rowset or to unknown rows in the rowset before it returns error information applying to specific rows. It returns warnings for specific rows along with any other error information about those rows.

***fFetchType* Argument**

The *fFetchType* argument specifies how to move through the result set. It is one of the following values:

SQL_FETCH_NEXT	SQL_FETCH_ABSOLUTE
SQL_FETCH_FIRST	SQL_FETCH_RELATIVE
SQL_FETCH_LAST	SQL_FETCH_BOOKMARK
SQL_FETCH_PRIOR	

If the value of the SQL_CURSOR_TYPE statement option is SQL_CURSOR_FORWARD_ONLY, the *fFetchType* argument must be SQL_FETCH_NEXT.

Note: In ODBC 1.0, **SQLExtendedFetch** supported the SQL_FETCH_RESUME fetch type. In ODBC 2.0, SQL_FETCH_RESUME is obsolete and the Driver Manager returns SQLSTATE S1C00 (Driver not capable) if an application specifies it for an ODBC 2.0 driver.

The SQL_FETCH_BOOKMARK fetch type was introduced in ODBC 2.0; the Driver Manager returns SQLSTATE S1106 (Fetch type out of range) if it is specified for an ODBC 1.0 driver.

Moving by Row Position

SQLExtendedFetch supports the following values of the *fFetchType* argument to move relative to the current rowset:

<i>fFetchType</i> Argument	Action
SQL_FETCH_NEXT	The driver returns the next rowset. If the cursor is positioned before the start of the result set, this is equivalent to SQL_FETCH_FIRST.
SQL_FETCH_PRIOR	The driver returns the prior rowset. If the cursor is positioned after the end of the result set, this is equivalent to SQL_FETCH_LAST.
SQL_FETCH_RELATIVE	The driver returns the rowset <i>irow</i> rows from the start of the current rowset. If <i>irow</i> equals 0, the driver refreshes the current rowset. If the cursor is positioned before the start of the result set and <i>irow</i> is greater than 0 or if the cursor is positioned after the end of the result set and <i>irow</i> is less than 0, this is equivalent to SQL_FETCH_ABSOLUTE.

It supports the following values of the *fFetchType* argument to move to an absolute position in the result set:

<i>fFetchType</i> Argument	Action
SQL_FETCH_FIRST	The driver returns the first rowset in the result set.
SQL_FETCH_LAST	The driver returns the last complete rowset in the result set.
SQL_FETCH_ABSOLUTE	<p>If <i>irow</i> is greater than 0, the driver returns the rowset starting at row <i>irow</i>.</p> <p>If <i>irow</i> equals 0, the driver returns SQL_NO_DATA_FOUND and the cursor is positioned before the start of the result set.</p> <p>If <i>irow</i> is less than 0, the driver returns the rowset starting at row $n+irow+1$, where <i>n</i> is the number of rows in the result set. For example, if <i>irow</i> is -1, the driver returns the rowset starting at the last row in the result set. If the result set size is 10 and <i>irow</i> is -10, the driver returns the rowset starting at the first row in the result set.</p>

Positioning to a Bookmark

When an application calls **SQLExtendedFetch** with the SQL_FETCH_BOOKMARK fetch type, the driver retrieves the rowset starting with the row specified by the bookmark in the *irow* argument.

To inform the driver that it will use bookmarks, the application calls **SQLSetStmtOption** with the `SQL_USE_BOOKMARKS` option before opening the cursor. To retrieve the bookmark for a row, the application either positions the cursor on the row and calls **SQLGetStmtOption** with the `SQL_GET_BOOKMARK` option, or retrieves the bookmark from column 0 of the result set. If the application retrieves a bookmark from column 0 of the result set, it must set *fCType* in **SQLBindCol** or **SQLGetData** to `SQL_C_BOOKMARK`. The application stores the bookmarks for those rows in each rowset to which it will return later.

Bookmarks are 32-bit binary values; if a bookmark requires more than 32 bits, such as when it is a key value, the driver maps the bookmarks requested by the application to 32-bit binary values. The 32-bit binary values are then returned to the application. Because this mapping may require considerable memory, applications should only bind column 0 of the result set if they will actually use bookmarks for most rows. Otherwise, applications should call **SQLGetStmtOption** with the `SQL_GET_BOOKMARK` statement option or call **SQLGetData** for column 0.

***row* Argument**

For the `SQL_FETCH_ABSOLUTE` fetch type, **SQLExtendedFetch** returns the rowset starting at the row number specified by the *row* argument.

For the `SQL_FETCH_RELATIVE` fetch type, **SQLExtendedFetch** returns the rowset starting *row* rows from the first row in the current rowset.

For the `SQL_FETCH_BOOKMARK` fetch type, the *row* argument specifies the bookmark that marks the first row in the requested rowset.

The *row* argument is ignored for the `SQL_FETCH_NEXT`, `SQL_FETCH_PRIOR`, `SQL_FETCH_FIRST`, and `SQL_FETCH_LAST`, fetch types.

***rgfRowStatus* Argument**

In the *rgfRowStatus* array, **SQLExtendedFetch** returns any changes in status to each row since it was last retrieved from the data source. Rows may be unchanged (`SQL_ROW_SUCCESS`), updated (`SQL_ROW_UPDATED`), deleted (`SQL_ROW_DELETED`), added (`SQL_ROW_ADDED`), or were unretrievable due to an error (`SQL_ROW_ERROR`). For static cursors, this information is available for all rows. For keyset, mixed, and dynamic cursors, this information is only available for rows in the keyset; the driver does not save data outside the keyset and therefore cannot compare the newly retrieved data to anything.

Note: Some drivers cannot detect changes to data. To determine whether a driver can detect changes to refetched rows, an application calls **SQLGetInfo** with the SQL_ROW_UPDATES option.

The number of elements 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 driver sets remaining status elements to `SQL_ROW_NOROW`.

When an application calls **SQLSetPos** with *fOption* set to `SQL_DELETE` or `SQL_UPDATE`, **SQLSetPos** changes the *rgfRowStatus* array for the changed row to `SQL_ROW_DELETED` or `SQL_ROW_UPDATED`.

Note: For keyset, mixed, and dynamic cursors, if a key value is updated, the row of data is considered to have been deleted and a new row added.

Code Example The following two examples show how an application could use column-wise or row-wise binding to bind storage locations to the same result set.

For more code examples, see **SQLSetPos**.

Column-Wise Binding

In the following example, an application declares storage locations for column-wise bound data and the returned numbers of bytes. Because column-wise binding is the default, there is no need, as in the row-wise binding example, to request column-wise binding with **SQLSetStmtOption**. However, the application does call **SQLSetStmtOption** to specify the number of rows in the rowset.

The application then executes a **SELECT** statement to return a result set of the employee names and birthdays, which is sorted by birthday. It calls **SQLBindCol** to bind the columns of data, passing the addresses of storage locations for both the data and the returned numbers of bytes. Finally, the application fetches the rowset data with **SQLExtendedFetch** and prints each employee's name and birthday.

```
#define ROWS 100
#define NAME_LEN 30
#define BDAY_LEN 11

UCHAR   szName[ROWS][NAME_LEN], szBirthday[ROWS][BDAY_LEN];
SWORD   sAge[ROWS];
SDWORD  cbName[ROWS], cbAge[ROWS], cbBirthday[ROWS];

UDWORD  crow, irow;
UWORD   rgfRowStatus[ROWS];

SQLSetStmtOption(hstmt, SQL_CONCURRENCY, SQL_CONCUR_READ_ONLY);
SQLSetStmtOption(hstmt, SQL_CURSOR_TYPE, SQL_CURSOR_KEYSET_DRIVEN);
SQLSetStmtOption(hstmt, SQL_ROWSET_SIZE, ROWS);
```

SQLNativeSql

```

retcode = SQLExecDirect(hstmt,
    "SELECT NAME, AGE, BIRTHDAY FROM EMPLOYEE ORDER BY 3, 2, 1",
    SQL_NTS);
if (retcode == SQL_SUCCESS) {
    SQLBindCol(hstmt, 1, SQL_C_CHAR, szName, NAME_LEN, cbName);
    SQLBindCol(hstmt, 2, SQL_C_SSHORT, sAge, 0, cbAge);
    SQLBindCol(hstmt, 3, SQL_C_CHAR, szBirthday, BDAY_LEN,
        cbBirthday);

    /* Fetch the rowset data and print each row. */
    /* On an error, display a message and exit. */

    while (TRUE) {
        retcode = SQLExtendedFetch(hstmt, SQL_FETCH_NEXT, 1, &crow,
            rgfRowStatus);
        if (retcode == SQL_ERROR || retcode == SQL_SUCCESS_WITH_INFO) {
            show_error();
        }
        if (retcode == SQL_SUCCESS || retcode == SQL_SUCCESS_WITH_INFO){
            for (irow = 0; irow < crow; irow++) {
                if (rgfRowStatus[irow] != SQL_ROW_DELETED &&
                    rgfRowStatus[irow] != SQL_ROW_ERROR)
                    fprintf(out, "%-*s %-2d %*s",
                        NAME_LEN-1, szName[irow], sAge[irow],
                        BDAY_LEN-1, szBirthday[irow]);
            }
        } else {
            break;
        }
    }
}

```

Row-Wise Binding

In the following example, an application declares an array of structures to hold row-wise bound data and the returned numbers of bytes. Using **SQLSetStmtOption**, it requests row-wise binding and passes the size of the structure to the driver. The driver will use this size to find successive storage locations in the array of structures. Using **SQLSetStmtOption**, it specifies the size of the rowset.

The application then executes a **SELECT** statement to return a result set of the employee names and birthdays, which is sorted by birthday. It calls **SQLBindCol** to bind the columns of data, passing the addresses of storage locations for both the data and the returned numbers of bytes. Finally, the application fetches the rowset data with **SQLExtendedFetch** and prints each employee's name and birthday.

```
#define ROWS 100
#define NAME_LEN 30
#define BDAY_LEN 11

typedef struct {
    UCHAR    szName[NAME_LEN];
    SDWORD   cbName;
    SWORD    sAge;
    SDWORD   cbAge;
    UCHAR    szBirthday[BDAY_LEN];
    SDWORD   cbBirthday;
}    EmpTable;

EmpTable rget[ROWS];
UDWORD   crow, irow;
UWORD    rgfRowStatus[ROWS];

SQLSetStmtOption(hstmt, SQL_BIND_TYPE, sizeof(EmpTable));
SQLSetStmtOption(hstmt, SQL_CONCURRENCY, SQL_CONCUR_READ_ONLY);
SQLSetStmtOption(hstmt, SQL_CURSOR_TYPE, SQL_CURSOR_KEYSET_DRIVEN);
SQLSetStmtOption(hstmt, SQL_ROWSET_SIZE, ROWS);
retcode = SQLExecDirect(hstmt,
    "SELECT NAME, AGE, BIRTHDAY FROM EMPLOYEE ORDER BY 3, 2, 1",
    SQL_NTS);

if (retcode == SQL_SUCCESS) {
    SQLBindCol(hstmt, 1, SQL_C_CHAR, rget[0].szName, NAME_LEN,
        &rget[0].cbName);
    SQLBindCol(hstmt, 2, SQL_C_SSHORT, &rget[0].sAge, 0,
        &rget[0].cbAge);
    SQLBindCol(hstmt, 3, SQL_C_CHAR, rget[0].szBirthday, BDAY_LEN,
        &rget[0].cbBirthday);

    /* Fetch the rowset data and print each row. */
    /* On an error, display a message and exit. */

    while (TRUE) {
        retcode = SQLExtendedFetch(hstmt, SQL_FETCH_NEXT, 1, &crow,
            rgfRowStatus);
        if (retcode == SQL_ERROR || retcode == SQL_SUCCESS_WITH_INFO) {
            show_error();
        }
    }
}
```

SQLNativeSql

```
        if (retcode == SQL_SUCCESS || retcode == SQL_SUCCESS_WITH_INFO){
            for (irow = 0; irow < crow; irow++) {
                if (rgfRowStatus[irow] != SQL_ROW_DELETED &&
                    rgfRowStatus[irow] != SQL_ROW_ERROR)
                    fprintf(out, "%-*s %-2d %*s",
                        NAME_LEN-1, rget[irow].szName, rget[irow].sAge,
                        BDAY_LEN-1, rget[irow].szBirthday);
            }
        } else {
            break;
        }
    }
}
```

Related Functions

For information about	See
Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Returning information about a column in a result set	SQLDescribeCol
Executing an SQL statement	SQLExecDirect
Executing a prepared SQL statement	SQLExecute
Returning the number of result set columns	SQLNumResultCols
Positioning the cursor in a rowset	SQLSetPos (extension)
Setting a statement option	SQLSetStmtOption (extension)

SQLFetch



Core **SQLFetch** fetches a row of data from a result set. The driver returns data for all columns that were bound to storage locations with **SQLBindCol**.

Syntax RETCODE **SQLFetch**(*hstmt*)

The **SQLFetch** function accepts the following argument.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_NO_DATA_FOUND, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLFetch** returns either SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLFetch** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	The data returned for one or more columns was truncated. String values are right truncated. For numeric values, the fractional part of number was truncated. (Function returns SQL_SUCCESS_WITH_INFO.)
07006	Restricted data type attribute violation	The data value could not be converted to the data type specified by <i>fCType</i> in SQLBindCol .
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
SQLSTATE	Error	Description
22003	Numeric value out of range	Returning the numeric value (as numeric or string) for one or more columns would have caused the whole (as opposed to fractional) part of the number to be truncated.

Returning the binary value for one or more
SQLNumParams

SQLNumParams

		columns would have caused a loss of binary significance. For more information, see “Converting Data from SQL to C Data Types” in Appendix D, “Data Types.”
22012	Division by zero	A value from an arithmetic expression was returned which resulted in division by zero.
24000	Invalid cursor state	The <i>hstmt</i> was in an executed state but no result set was associated with the <i>hstmt</i> .
40001	Serialization failure	The transaction in which the fetch was executed was terminated to prevent deadlock.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLERROR in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1002	Invalid column number	A column number specified in the binding for one or more columns was greater than the number of columns in the result set. A column number specified in the binding for a column was 0; SQLFetch cannot be used to retrieve bookmarks.
SQLSTATE	Error	Description
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1010	Function sequence error	(DM) The specified <i>hstmt</i> was not in an executed state. The function was called without first calling SQLExecDirect , SQLExecute ,

		or a catalog function..
		(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called.
		(DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
		(DM) SQLExtendedFetch was called for an <i>hstmt</i> after SQLFetch was called and before SQLFreeStmt was called with the SQL_CLOSE option.
S1C00	Driver not capable	The driver or 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. This error only applies when the SQL data type of the column was mapped to a driver-specific SQL data type.
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

Comments

SQLFetch positions the cursor on the next row of the result set. Before **SQLFetch** is called the first time, the cursor is positioned before the start of the result set. When the cursor is positioned on the last row of the result set, **SQLFetch** returns SQL_NO_DATA_FOUND and the cursor is positioned after the end of the result set. An application cannot mix calls to **SQLExtendedFetch** and **SQLFetch** for the same cursor.

If the application called **SQLBindCol** to bind columns, **SQLFetch** stores data into the locations specified by the calls to **SQLBindCol**. If the application does not call **SQLBindCol** to bind any columns, **SQLFetch** doesn't return any data; it just moves the cursor to the next row. An application can call **SQLGetData** to retrieve data that is not bound to a storage location.

The driver manages cursors during the fetch operation and places each value of a bound column into the associated storage. The driver follows these guidelines when performing a fetch operation:

- **SQLFetch** accesses column data in left-to-right order.
- After each fetch, *pcbValue* (specified in **SQLBindCol**) contains the number of bytes available to return for the column. This is the number of bytes available prior to calling **SQLFetch**. If the number of bytes

SQLNumParams

SQLNumParams

available to return cannot be determined in advance, the driver sets *pcbValue* to SQL_NO_TOTAL. (If SQL_MAX_LENGTH has been specified with **SQLSetStmtOption** and the number of bytes available to return is greater than SQL_MAX_LENGTH, *pcbValue* contains SQL_MAX_LENGTH.)

Note: The SQL_MAX_LENGTH statement option is intended to reduce network traffic and may not be supported by all drivers. To guarantee that data is truncated, an application should allocate a buffer of the desired size and specify this size in the *cbValueMax* argument.

- If *rgbValue* is not large enough to hold the entire result, the driver stores part of the value and returns SQL_SUCCESS_WITH_INFO. A subsequent call to **SQLError** indicates that a truncation occurred. The application can compare *pcbValue* to *cbValueMax* (specified in **SQLBindCol**) to determine which column or columns were truncated. If *pcbValue* is greater than or equal to *cbValueMax*, then truncation occurred.
- If the data value for the column is NULL, the driver stores SQL_NULL_DATA in *pcbValue*.

SQLFetch is valid only after a call that returns a result set.

For information about conversions allowed by **SQLBindCol** and **SQLGetData**, see “Converting Data from SQL to C Data Types” in Appendix D, “Data Types.”

Code Example See **SQLBindCol**, **SQLColumns**, **SQLGetData**, and **SQLProcedures**.

Related Functions

For information about	See
-----------------------	-----

Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Returning information about a column in a result set	SQLDescribeCol
Executing an SQL statement	SQLExecDirect
Executing a prepared SQL statement	SQLExecute
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
Freeing a statement handle	SQLFreeStmt

Fetching part or all of a column of data
Returning the number of result set columns
Preparing a statement for execution

SQLGetData (extension)

SQLNumResultCols

SQLPrepare

SQLForeignKeys

ADBC

Extension Level 2 SQLForeignKeys can return:

- A list of foreign keys in the specified table (columns in the specified table that refer to primary keys in other tables).
- A list of foreign keys in other tables that refer to the primary key in the specified table.

The driver returns each list as a result set on the specified *hstmt*.

Syntax

RETCODE SQLForeignKeys(*hstmt*, *szPkTableQualifier*, *cbPkTableQualifier*, *szPkTableOwner*, *cbPkTableOwner*, *szPkTableName*, *cbPkTableName*, *szFkTableQualifier*, *cbFkTableQualifier*, *szFkTableOwner*, *cbFkTableOwner*, *szFkTableName*, *cbFkTableName*)

The SQLForeignKeys function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UCHAR FAR *	<i>szPkTableQualifier</i>	Input	Primary key table qualifier. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have qualifiers.
SWORD	<i>cbPkTableQualifier</i>	Input	Length of <i>szPkTableQualifier</i> .
UCHAR FAR *	<i>szPkTableOwner</i>	Input	Primary key owner name. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have owners.
SWORD	<i>cbPkTableOwner</i>	Input	Length of <i>szPkTableOwner</i> .
UCHAR FAR *	<i>szPkTableName</i>	Input	Primary key table name.
SWORD	<i>cbPkTableName</i>	Input	Length of <i>szPkTableName</i> .
Type	Argument	Use	Description
UCHAR FAR *	<i>szFkTableQualifier</i>	Input	Foreign key table qualifier. If a driver supports qualifiers for some tables but not for others, such as when

the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have qualifiers.

SWORD	<i>cbFkTableQualifier</i>	Input	Length of <i>szFkTableQualifier</i> .
UCHAR FAR *	<i>szFkTableOwner</i>	Input	Foreign key owner name. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have owners.
SWORD	<i>cbFkTableOwner</i>	Input	Length of <i>szFkTableOwner</i> .
UCHAR FAR *	<i>szFkTableName</i>	Input	Foreign key table name.
SWORD	<i>cbFkTableName</i>	Input	Length of <i>szFkTableName</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLForeignKeys** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLForeignKeys** and explains each one in the context of this function; the notation “(DM)” precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called. A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called.
SQLSTATE	Error	Description
IM001	Driver does not support	(DM) The driver associated with the <i>hstmt</i> does

SQLNumResultCols

SQLNumResultCols

	this function	not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLERROR in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1009	Invalid argument value	(DM) The arguments <i>szPkTableName</i> and <i>szFkTableName</i> were both null pointers.
S1010	Function sequence error	(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1090	Invalid string or buffer length	(DM) The value of one of the name length arguments was less than 0, but not equal to SQL_NTS. The value of one of the name length arguments exceeded the maximum length value for the corresponding qualifier or name (see "Comments").
SQLSTATE Error		Description
S1C00	Driver not capable	A table qualifier was specified and the driver or data source does not support qualifiers. A table owner was specified and the driver or data source does not support owners.

		The combination of the current settings of the SQL_CONCURRENCY and SQL_CURSOR_TYPE statement options was not supported by the driver or data source.
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

Comments

If *szPkTableName* contains a table name, **SQLForeignKeys** returns a result set containing the primary key of the specified table and all of the foreign keys that refer to it.

If *szFkTableName* contains a table name, **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 most.

SQLForeignKeys returns results as a standard result set. If the foreign keys associated with a primary key are requested, the result set is ordered by FKTABLE_QUALIFIER, FKTABLE_OWNER, FKTABLE_NAME, and KEY_SEQ. If the primary keys associated with a foreign key are requested, the result set is ordered by PKTABLE_QUALIFIER, PKTABLE_OWNER, PKTABLE_NAME, and KEY_SEQ. The following table lists the columns in the result set.

The lengths of VARCHAR columns shown in the table are maximums; the actual lengths depend on the data source. To determine the actual lengths of the TABLE_QUALIFIER, TABLE_OWNER, TABLE_NAME, and COLUMN_NAME columns, an application can call **SQLGetInfo** with the SQL_MAX_QUALIFIER_NAME_LEN, SQL_MAX_OWNER_NAME_LEN, SQL_MAX_TABLE_NAME_LEN, and SQL_MAX_COLUMN_NAME_LEN options.

Column Name	Data Type	Comments
PKTABLE_QUALIFIER	Varchar(128)	Primary key table qualifier identifier; NULL if not applicable to the data source. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different

SQLNumResultCols

SQLNumResultCols

		DBMSs, it returns an empty string ("" for those tables that do not have qualifiers.
PKTABLE_OWNER	Varchar(128)	Primary key table owner identifier; NULL if not applicable to the data source. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("" for those tables that do not have owners.
PKTABLE_NAME	Varchar(128) not NULL	Primary key table identifier.
PKCOLUMN_NAME	Varchar(128) not NULL	Primary key column identifier.
FKTABLE_QUALIFIER	Varchar(128)	Foreign key table qualifier identifier; NULL if not applicable to the data source. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("" for those tables that do not have qualifiers.
FKTABLE_OWNER	Varchar(128)	Foreign key table owner identifier; NULL if not applicable to the data source. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("" for those tables that do not have owners.
FKTABLE_NAME	Varchar(128) not NULL	Foreign key table identifier.
FKCOLUMN_NAME	Varchar(128) not NULL	Foreign key column identifier.
KEY_SEQ	Smallint not NULL	Column sequence number in key (starting with 1).

Column Name	<small>BEGIN BREAK</small> Data Type	Comments
UPDATE_RULE	Smallint	Action to be applied to the foreign key

		when the SQL operation is UPDATE : SQL_CASCADE SQL_RESTRICT SQL_SET_NULL NULL if not applicable to the data source.
DELETE_RULE	Smallint	Action to be applied to the foreign key when the SQL operation is DELETE : SQL_CASCADE SQL_RESTRICT SQL_SET_NULL NULL if not applicable to the data source.
FK_NAME	Varchar(128)	Foreign key identifier. NULL if not applicable to the data source.
PK_NAME	Varchar(128)	Primary key identifier. NULL if not applicable to the data source.

Note: The FK_NAME and PK_NAME columns were added in ODBC 2.0. ODBC 1.0 drivers may return different, driver-specific columns with the same column numbers.

Code Example This example uses four tables:

SALES_ORDER	SALES_LINE	CUSTOMER	EMPLOYEE
SALES_ID	SALES_ID	CUSTOMER_ID	EMPLOYEE_ID
CUSTOMER_ID	LINE_NUMBER	CUST_NAME	NAME
EMPLOYEE_ID	PART_ID	ADDRESS	AGE
TOTAL_PRICE	QUANTITY	PHONE	BIRTHDAY
	PRICE		

In the SALES_ORDER table, CUSTOMER_ID identifies the customer to whom the sale has been made. It is a foreign key that refers to CUSTOMER_ID in the CUSTOMER table. EMPLOYEE_ID identifies the employee who made the sale. It is a foreign key that refers to EMPLOYEE_ID in the EMPLOYEE table.

In the SALES_LINE table, SALES_ID identifies the sales order with which the line item is associated. It is a foreign key that refers to SALES_ID in the SALES_ORDER table.

This example calls **SQLPrimaryKeys** to get the primary key of the SALES_ORDER table. The result set will have one row and the significant columns are:

SQLNumResultCols

SQLNumResultCols

TABLE_NAME	COLUMN_NAME	KEY_SEQ
SALES_ORDER	SALES_ID	1

Next, the example calls **SQLForeignKeys** to get the foreign keys in other tables that reference the primary key of the SALES_ORDER table. The result set will have one row and the significant columns are:

PKTABLE_NAME	PKCOLUMN_NAME	FKTABLE_NAME	FKCOLUMN_NAME	KEY_SEQ
--------------	---------------	--------------	---------------	---------

SALES_ORDER	SALES_ID	SALES_LINE	SALES_ID	1
-------------	----------	------------	----------	---

Finally, the example calls **SQLForeignKeys** to get the foreign keys in the SALES_ORDER table that refer to the primary keys of other tables. The result set will have two rows and the significant columns are:

PKTABLE_NAME	PKCOLUMN_NAME	FKTABLE_NAME	FKCOLUMN_NAME	KEY_SEQ
CUSTOMER	CUSTOMER_ID	SALES_ORDER	CUSTOMER_ID	1
EMPLOYEE	EMPLOYEE_ID	SALES_ORDER	EMPLOYEE_ID	1

```
#define TAB_LEN SQL_MAX_TABLE_NAME_LEN + 1
#define COL_LEN SQL_MAX_COLUMN_NAME_LEN + 1

LPSTR  szTable;          /* Table to display */

UCHAR  szPkTable[TAB_LEN]; /* Primary key table name */
UCHAR  szFkTable[TAB_LEN]; /* Foreign key table name */
UCHAR  szPkCol[COL_LEN];   /* Primary key column */
UCHAR  szFkCol[COL_LEN];   /* Foreign key column */

HSTMT   hstmt;
SDWORD  cbPkTable, cbPkCol, cbFkTable, cbFkCol, cbKeySeq;
SWORD   iKeySeq;
RETCODE retcode;

/* Bind the columns that describe the primary and foreign keys. */
/* Ignore the table owner, name, and qualifier for this example. */

SQLBindCol(hstmt, 3, SQL_C_CHAR, szPkTable, TAB_LEN, &cbPkTable);
SQLBindCol(hstmt, 4, SQL_C_CHAR, szPkCol, COL_LEN, &cbPkCol);
SQLBindCol(hstmt, 5, SQL_C_SHORT, &iKeySeq, TAB_LEN, &cbKeySeq);
SQLBindCol(hstmt, 7, SQL_C_CHAR, szFkTable, TAB_LEN, &cbFkTable);
SQLBindCol(hstmt, 8, SQL_C_CHAR, szFkCol, COL_LEN, &cbFkCol);
```

```

strcpy(szTable, "SALES_ORDER");

/* Get the names of the columns in the primary key.      */

retcode = SQLPrimaryKeys(hstmt,
                        NULL, 0,      /* Table qualifier */
                        NULL, 0,      /* Table owner    */
                        szTable, SQL_NTS); /* Table name     */

while ((retcode == SQL_SUCCESS) || (retcode == SQL_SUCCESS_WITH_INFO)) {

    /* Fetch and display the result set. This will be a list of the */
    /* columns in the primary key of the SALES_ORDER table.          */

    retcode = SQLFetch(hstmt);
    if (retcode == SQL_SUCCESS || retcode != SQL_SUCCESS_WITH_INFO)
        fprintf(out, "Column: %s   Key Seq: %hd\n", szPkCol, iKeySeq);
}

/* Close the cursor (the hstmt is still allocated).      */

SQLFreeStmt(hstmt, SQL_CLOSE);

/* Get all the foreign keys that refer to SALES_ORDER primary key. */

retcode = SQLForeignKeys(hstmt,
                        NULL, 0,      /* Primary qualifier */
                        NULL, 0,      /* Primary owner     */
                        szTable, SQL_NTS, /* Primary table      */
                        NULL, 0,      /* Foreign qualifier  */
                        NULL, 0,      /* Foreign owner      */
                        NULL, 0);     /* Foreign table      */

while ((retcode == SQL_SUCCESS) || (retcode == SQL_SUCCESS_WITH_INFO)) {

    /* Fetch and display the result set. This will be all of the */
    /* foreign keys in other tables that refer to the SALES_ORDER */
    /* primary key.                                                */

    retcode = SQLFetch(hstmt);
    if (retcode == SQL_SUCCESS || retcode == SQL_SUCCESS_WITH_INFO)
        fprintf(out, "%s ( %s ) <-- %s ( %s )\n", szPkTable,
                szPkCol, szFkTable, szFkCol);
}

/* Close the cursor (the hstmt is still allocated).      */

SQLFreeStmt(hstmt, SQL_CLOSE);

```

SQLNumResultCols

```
/* Get all the foreign keys in the SALES_ORDER table. */

retcode = SQLForeignKeys(hstmt,
    NULL, 0, /* Primary qualifier */
    NULL, 0, /* Primary owner */
    NULL, 0, /* Primary table */
    NULL, 0, /* Foreign qualifier */
    NULL, 0, /* Foreign owner */
    szTable, SQL_NTS); /* Foreign table */

while ((retcode == SQL_SUCCESS) || (retcode == SQL_SUCCESS_WITH_INFO)) {

    /* Fetch and display the result set. This will be all of the
    /* primary keys in other tables that are referred to by foreign
    /* keys in the SALES_ORDER table. */

    retcode = SQLFetch(hstmt);
    if (retcode == SQL_SUCCESS || retcode == SQL_SUCCESS_WITH_INFO)
        fprintf(out, "%-s ( %-s )--> %-s ( %-s )\n", szFkTable, szFkCol,
            szPkTable, szPkCol);
}
/* Free the hstmt. */
SQLFreeStmt(hstmt, SQL_DROP);
```

Related Functions

For information about

See

Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
Fetching a row of data	SQLFetch
Returning the columns of a primary key	SQLPrimaryKeys (extension)
Returning table statistics and indexes	SQLStatistics (extension)

SQLFreeConnect



Core **SQLFreeConnect** releases a connection handle and frees all memory associated with the handle.

Syntax RETCODE **SQLFreeConnect**(*hdbc*)

The **SQLFreeConnect** function accepts the following argument.

Type	Argument	Use	Description
HDBC	<i>hdbc</i>	Input	Connection handle.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLFreeConnect** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLFreeConnect** and explains each one in the context of this function; the notation “(DM)” precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1010	Function sequence error	(DM) The function was called prior to calling SQLDisconnect for the <i>hdbc</i> .

Comments Prior to calling **SQLFreeConnect**, an application must call **SQLDisconnect** for the *hdbc*. Otherwise, **SQLFreeConnect** returns SQL_ERROR and the *hdbc* remains valid. Note that **SQLDisconnect** automatically drops any *hstmts* open on the *hdbc*.

SQLParamData

Code Example See **SQLBrowseConnect** and **SQLConnect**.

Related Functions

For information about

See

Allocating a statement handle

SQLAllocConnect

Connecting to a data source

SQLConnect

Disconnecting from a data source

SQLDisconnect

Connecting to a data source using a connection string or dialog box

SQLDriverConnect (extension)

Freeing an environment handle

SQLFreeEnv

Freeing a statement handle

SQLFreeStmt

SQLFreeEnv



Core **SQLFreeEnv** frees the environment handle and releases all memory associated with the environment handle.

Syntax RETCODE **SQLFreeEnv**(*henv*)

The **SQLFreeEnv** function accepts the following argument.

Type	Argument	Use	Description
HENV	<i>henv</i>	Input	Environment handle.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLFreeEnv** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLFreeEnv** and explains each one in the context of this function; the notation “(DM)” precedes the descriptions of SQLSTATEs returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1010	Function sequence error	(DM) There was at least one <i>hdbc</i> in an allocated or connected state. Call SQLDisconnect and SQLFreeConnect for each <i>hdbc</i> before calling SQLFreeEnv .

Comments Prior to calling **SQLFreeEnv**, an application must call **SQLFreeConnect** for any *hdbc* allocated under the *henv*. Otherwise, **SQLFreeEnv** returns SQL_ERROR and the *henv* and any active *hdbc* remains valid. When the Driver Manager processes the **SQLFreeEnv** function, it checks the **TraceAutoStop** keyword in the [ODBC] section of the ODBC.INI file or the ODBC subkey of the registry. If it is set to 1, the Driver Manager disables

SQLParamOptions

tracing for all applications and sets the **Trace** keyword in the [ODBC] section of the ODBC.INI file or the ODBC subkey of the registry to 0.

Code Example See **SQLBrowseConnect** and **SQLConnect**.

Related Functions	For information about	See
--------------------------	------------------------------	------------

Allocating an environment handle	SQLAllocEnv
----------------------------------	--------------------

Freeing a connection handle	SQLFreeConnect
-----------------------------	-----------------------

Note: The Driver Manager does not exist on non-Windows systems.

SQLFreeStmt



Core **SQLFreeStmt** stops processing associated with a specific *hstmt*, closes any open cursors associated with the *hstmt*, discards pending results, and, optionally, frees all resources associated with the statement handle.

Syntax RETCODE **SQLFreeStmt**(*hstmt*, *fOption*)

The **SQLFreeStmt** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle
UWORD	<i>fOption</i>	Input	One of the following options: SQL_CLOSE: Close the cursor associated with <i>hstmt</i> (if one was defined) and discard all pending results. The application can reopen this cursor later by executing a SELECT statement again with the same or different parameter values. If no cursor is open, this option has no effect for the application. SQL_DROP: Release the <i>hstmt</i> , free all resources associated with it, close the cursor (if one is open), and discard all pending rows. This option terminates all access to the <i>hstmt</i> . The <i>hstmt</i> must be reallocated to be reused. SQL_UNBIND: Release all column buffers bound by SQLBindCol for the given <i>hstmt</i> . SQL_RESET_PARAMS: Release all parameter buffers set by SQLBindParameter for the given <i>hstmt</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLFreeStmt** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLFreeStmt** and explains each one in the context of this function; the notation “(DM)” precedes the descriptions of SQLSTATES returned by the Driver

SQLPrepare

Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1010	Function sequence error	(DM) An asynchronously executing function was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1092	Option type out of range	(DM) The value specified for the argument <i>fOption</i> was not: SQL_CLOSE SQL_DROP SQL_UNBIND SQL_RESET_PARAMS

Comments An application can call **SQLFreeStmt** to terminate processing of a **SELECT** statement with or without canceling the statement handle.

The SQL_DROP option frees all resources that were allocated by the **SQLAllocStmt** function.

Code Example See **SQLBrowseConnect** and **SQLConnect**.

Related Functions	For information about	See
--------------------------	------------------------------	------------

Allocating a statement handle	SQLAllocStmt
Canceling statement processing	SQLCancel
Setting a cursor name	SQLSetCursorName

SQLGetConnectOption



Extension Level 1 **SQLGetConnectOption** returns the current setting of a connection option.

Syntax RETCODE **SQLGetConnectOption**(*hdbc*, *fOption*, *pvParam*)

The **SQLGetConnectOption** function accepts the following arguments:

Type	Argument	Use	Description
HDBC	<i>hdbc</i>	Input	Connection handle.
UWORD	<i>fOption</i>	Input	Option to retrieve.
PTR	<i>pvParam</i>	Output	Value associated with <i>fOption</i> . Depending on the value of <i>fOption</i> , a 32-bit integer value or a pointer to a null-terminated character string will be returned in <i>pvParam</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_NO_DATA_FOUND, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLGetConnectOption** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLGetConnectOption** and explains each one in the context of this function; the notation “(DM)” precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
08003	Connection not open	(DM) An <i>fOption</i> value was specified that required an open connection.
IM001	Driver does not support this function	(DM) The driver corresponding to the <i>hdbc</i> does not support the function.
S1000	General error	An error occurred for which there was no

SQLPrepare

SQLPrepare

specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by **SQLError** in the argument *szErrorMsg* describes the error and its cause.

SQLSTATE	Error	Description
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1010	Function sequence error	(DM) SQLBrowseConnect was called for the <i>hdbc</i> and returned SQL_NEED_DATA. This function was called before SQLBrowseConnect returned SQL_SUCCESS_WITH_INFO or SQL_SUCCESS.
S1092	Option type out of range	(DM) The value specified for the argument <i>fOption</i> was in the block of numbers reserved for ODBC connection and statement options, but was not valid for the version of ODBC supported by the driver.
S1C00	Driver not capable	The value specified for the argument <i>fOption</i> was a valid ODBC connection option for the version of ODBC supported by the driver, but was not supported by the driver. The value specified for the argument <i>fOption</i> was in the block of numbers reserved for driver-specific connection and statement options, but was not supported by the driver.

Comments

For a list of options, see **SQLSetConnectOption**. Note that if *fOption* specifies an option that returns a string, *pvParam* must be a pointer to storage for the string. The maximum length of the string will be SQL_MAX_OPTION_STRING_LENGTH bytes (excluding the null termination byte).

Depending on the option, an application does not need to establish a connection prior to calling **SQLGetConnectOption**. However, if **SQLGetConnectOption** is called and the specified option does not have a default and has not been set by a prior call to **SQLSetConnectOption**, **SQLGetConnectOption** will return SQL_NO_DATA_FOUND.

While an application can set statement options using **SQLSetConnectOption**, an application cannot use **SQLGetConnectOption** to retrieve statement option

values; it must call `SQLGetStmtOption` to retrieve the setting of statement options.

Related Functions**For information about****See**

Returning the setting of a statement option

SQLGetStmtOption (extension)

Setting a connection option

SQLSetConnectOption (extension)

Setting a statement option

SQLSetStmtOption (extension)

SQLGetCursorName

ADBC

Core **SQLGetCursorName** returns the cursor name associated with a specified *hstmt*.

Syntax RETCODE **SQLGetCursorName**(*hstmt*, *szCursor*, *cbCursorMax*, *pcbCursor*)

The **SQLGetCursorName** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UCHAR FAR *	<i>szCursor</i>	Output	Pointer to storage for the cursor name.
SWORD	<i>cbCursorMax</i>	Input	Length of <i>szCursor</i> .
SWORD FAR *	<i>pcbCursor</i>	Output	Total number of bytes (excluding the null termination byte) available to return in <i>szCursor</i> . If the number of bytes available to return is greater than or equal to <i>cbCursorMax</i> , the cursor name in <i>szCursor</i> is truncated to <i>cbCursorMax</i> – 1 bytes.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLGetCursorName** returns either SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLGetCursorName** and explains each one in the context of this function; the notation “(DM)” precedes the descriptions of SQLSTATEs returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	The buffer <i>szCursor</i> was not large enough to return the entire cursor name, so the cursor name was truncated. The argument <i>pcbCursor</i> contains the length of the untruncated cursor name. (Function returns SQL_SUCCESS_WITH_INFO.)
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
SQLSTATE	Error	Description

S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1010	Function sequence error	(DM) An asynchronously executing function was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1015	No cursor name available	(DM) There was no open cursor on the <i>hstmt</i> and no cursor name had been set with SQLSetCursorName .
S1090	Invalid string or buffer length	(DM) The value specified in the argument <i>cbCursorMax</i> was less than 0.

Comments

The only ODBC SQL statements that use a cursor name are positioned update and delete (for example, **UPDATE** *table-name* ...**WHERE CURRENT OF** *cursor-name*). If the application does not call **SQLSetCursorName** to define a cursor name, on execution of a **SELECT** statement the driver generates a name that begins with the letters SQL_CUR and does not exceed 18 characters in length.

SQLGetCursorName returns the name of a cursor regardless of whether the name was created explicitly or implicitly.

A cursor name that is set either explicitly or implicitly remains set until the *hstmt* with which it is associated is dropped, using **SQLFreeStmt** with the SQL_DROP option.

SQLPrimaryKeys

Related Functions

For information about

See

Executing an SQL statement

SQLExecDirect

Executing a prepared SQL statement

SQLExecute

Preparing a statement for execution

SQLPrepare

Setting a cursor name

SQLSetCursorName

Setting cursor scrolling options

SQLSetScrollOptions (extension)

SQLGetData



Extension Level 1 **SQLGetData** returns result data for a single unbound column in the current row. The application must call **SQLFetch**, or **SQLExtendedFetch** and (optionally) **SQLSetPos** to position the cursor on a row of data before it calls **SQLGetData**. It is possible to use **SQLBindCol** for some columns and use **SQLGetData** for others within the same row. This function can be used to retrieve character or binary data values in parts from a column with a character, binary, or data source–specific data type (for example, data from SQL_LONGVARIABLE or SQL_LONGVARCHAR columns).

Syntax

RETCODE **SQLGetData**(*hstmt*, *icol*, *fCType*, *rgbValue*, *cbValueMax*, *pcbValue*)

The **SQLGetData** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UWORD	<i>icol</i>	Input	Column number of result data, ordered sequentially left to right, starting at 1. A column number of 0 is used to retrieve a bookmark for the row; bookmarks are not supported by ODBC 1.0 drivers or SQLFetch .

Type	Argument	Use	Description
SWORD	<i>fCType</i>	Input	The C data type of the result data. This must be one of the following values: SQL_C_BINARY SQL_C_BIT SQL_C_BOOKMARK SQL_C_CHAR SQL_C_DATE SQL_C_DEFAULT SQL_C_DOUBLE SQL_C_FLOAT SQL_C_SLONG SQL_C_SSHORT SQL_C_STINYINT SQL_C_TIME SQL_C_TIMESTAMP SQL_C_ULONG

SQLProcedureColumns

SQLProcedureColumns

SQL_C_USHORT
SQL_C_UTINYINT

SQL_C_DEFAULT specifies that data be converted to its default C data type.

Note: Drivers must also support the following values of *fCType* from ODBC 1.0. Applications must use these values, rather than the ODBC 2.0 values, when calling an ODBC 1.0 driver:

SQL_C_LONG
SQL_C_SHORT
SQL_C_TINYINT

For information about how data is converted, see "Converting Data from SQL to C Data Types" in Appendix D, "Data Types."

PTR	<i>rgbValue</i>	Output	Pointer to storage for the data.
Type	Argument	Use	Description
SDWORD	<i>cbValueMax</i>	Input	Maximum length of the <i>rgbValue</i> buffer. For character data, <i>rgbValue</i> must also include space for the null-termination byte. For character and binary C data, <i>cbValueMax</i> determines the amount of data that can be received in a single call to SQLGetData . For all other types of C data, <i>cbValueMax</i> is ignored; the driver assumes that the size of <i>rgbValue</i> is the size of the C data type specified with <i>fCType</i> and returns the entire data value. For more information about length, see "Precision, Scale, Length, and Display Size" in Appendix D, "Data Types."
SDWORD FAR * <i>pcbValue</i>		Output	SQL_NULL_DATA, the total number of bytes (excluding the null termination byte for character data) available to return in <i>rgbValue</i> prior to the current call to SQLGetData , or SQL_NO_TOTAL if the number of

available bytes cannot be determined.

For character data, if *pcbValue* is SQL_NO_TOTAL or is greater than or equal to *cbValueMax*, the data in *rgbValue* is truncated to *cbValueMax* – 1 bytes and is null-terminated by the driver.

For binary data, if *pcbValue* is SQL_NO_TOTAL or is greater than *cbValueMax*, the data in *rgbValue* is truncated to *cbValueMax* bytes.

For all other data types, the value of *cbValueMax* is ignored and the driver assumes the size of *rgbValue* is the size of the C data type specified with *fCType*.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_NO_DATA_FOUND, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLGetData** returns either SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLGetData** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	All of the data for the specified column, <i>icol</i> , could not be retrieved in a single call to the function. The argument <i>pcbValue</i> contains the length of the data remaining in the specified column prior to the current call to SQLGetData . (Function returns SQL_SUCCESS_WITH_INFO.) For more information on using multiple calls to SQLGetData for a single column, see "Comments."
07006	Restricted data type attribute violation	The data value cannot be converted to the C data type specified by the argument <i>fCType</i> .
08S01	Communication link	The communication link between the driver and

SQLProcedureColumns

SQLProcedureColumns

	failure	the data source to which the driver was connected failed before the function completed processing.
22003	Numeric value out of range	Returning the numeric value (as numeric or string) for the column would have caused the whole (as opposed to fractional) part of the number to be truncated. Returning the binary value for the column would have caused a loss of binary significance. For more information, see Appendix D, "Data Types."
22005	Error in assignment	The data for the column was incompatible with the data type into which it was to be converted. For more information, see Appendix D, "Data Types."
22008	Datetime field overflow	The data for the column was not a valid date, time, or timestamp value. For more information, see Appendix D, "Data Types."
SQLSTATE	Error	Description
24000	Invalid cursor state	(DM) The <i>hstmt</i> was in an executed state but no result set was associated with the <i>hstmt</i> . (DM) A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called. A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called, but the cursor was positioned before the start of the result set or after the end of the result set.
IM001	Driver does not support this function	(DM) The driver corresponding to the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.

S1002	Invalid column number	<p>The value specified for the argument <i>icol</i> was 0 and the driver was an ODBC 1.0 driver.</p> <p>The value specified for the argument <i>icol</i> was 0 and SQLFetch was used to fetch the data.</p> <p>The value specified for the argument <i>icol</i> was 0 and the SQL_USE_BOOKMARKS statement option was set to SQL_UB_OFF.</p> <p>The specified column was greater than the number of result columns.</p> <p>The specified column was bound through a call to SQLBindCol. This description does not apply to drivers that return the SQL_GD_BOUND bitmask for the SQL_GETDATA_EXTENSIONS option in SQLGetInfo.</p>
SQLSTATE Error		Description
		<p>The specified column was at or before the last bound column specified through SQLBindCol. This description does not apply to drivers that return the SQL_GD_ANY_COLUMN bitmask for the SQL_GETDATA_EXTENSIONS option in SQLGetInfo.</p> <p>The application has already called SQLGetData for the current row. The column specified in the current call was before the column specified in the preceding call. This description does not apply to drivers that return the SQL_GD_ANY_ORDER bitmask for the SQL_GETDATA_EXTENSIONS option in SQLGetInfo.</p>
S1003	Program type out of range	<p>(DM) The argument <i>fCType</i> was not a valid data type or SQL_C_DEFAULT.</p> <p>The argument <i>icol</i> was 0 and the argument <i>fCType</i> was not SQL_C_BOOKMARK.</p>
S1008	Operation canceled	<p>Asynchronous processing was enabled for the <i>hstmt</i>. The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i>. Then the function was called again on the <i>hstmt</i>.</p> <p>The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread</p>

SQLProcedureColumns

		in a multithreaded application.
S1009	Invalid argument value	(DM) The argument <i>rgbValue</i> was a null pointer.
S1010	Function sequence error	<p>(DM) The specified <i>hstmt</i> was not in an executed state. The function was called without first calling SQLExecDirect, SQLExecute, or a catalog function.</p> <p>(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called.</p> <p>(DM) SQLExecute, SQLExecDirect, or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.</p>
SQLSTATE	Error	Description
S1090	Invalid string or buffer length	(DM) The value specified for argument <i>cbValueMax</i> was less than 0.
S1109	Invalid cursor position	The cursor was positioned (by SQLSetPos or SQLExtendedFetch) on a row for which the value in the <i>rgfRowStatus</i> array in SQLExtendedFetch was SQL_ROW_DELETED or SQL_ROW_ERROR.
S1C00	Driver not capable	<p>The driver or data source does not support use of SQLGetData with multiple rows in SQLExtendedFetch. This description does not apply to drivers that return the SQL_GD_BLOCK bitmask for the SQL_GETDATA_EXTENSIONS option in SQLGetInfo.</p> <p>The driver or data source does not support the conversion specified by the combination of the <i>fCType</i> argument and the SQL data type of the corresponding column. This error only applies when the SQL data type of the column was mapped to a driver-specific SQL data type. The argument <i>icol</i> was 0 and the driver does not support bookmarks.</p> <p>The driver only supports ODBC 1.0 and the argument <i>fCType</i> was one of the following:</p> <p>SQL_C_STINYINT SQL_C_UTINYINT</p>

SQL_C_SSHORT
 SQL_C_USHORT
 SQL_C_SLONG
 SQL_C_ULONG

S1T00 Timeout expired The timeout period expired before the data source returned the result set. The timeout period is set through **SQLSetStmtOption**, **SQL_QUERY_TIMEOUT**.

Comments With each call, the driver sets *pcbValue* to the number of bytes that were available in the result column prior to the current call to **SQLGetData**. (If **SQL_MAX_LENGTH** has been set with **SQLSetStmtOption**, and the total number of bytes available on the first call is greater than **SQL_MAX_LENGTH**, the available number of bytes is set to **SQL_MAX_LENGTH**. Note that the **SQL_MAX_LENGTH** statement option is intended to reduce network traffic and may not be supported by all drivers. To guarantee that data is truncated, an application should allocate a buffer of the desired size and specify this size in the *cbValueMax* argument.) If the total number of bytes in the result column cannot be determined in advance, the driver sets *pcbValue* to **SQL_NO_TOTAL**. If the data value for the column is NULL, the driver stores **SQL_NULL_DATA** in *pcbValue*.

SQLGetData can convert data to a different data type. The result and success of the conversion is determined by the rules for assignment specified in "Converting Data from SQL to C Data Types" in Appendix D, "Data Types."

If more than one call to **SQLGetData** is required to retrieve data from a single column with a character, binary, or data source-specific data type, the driver returns **SQL_SUCCESS_WITH_INFO**. A subsequent call to **SQLGetInfo** returns **SQLSTATE 01004** (Data truncated). The application can then use the same column number to retrieve subsequent parts of the data until **SQLGetData** returns **SQL_SUCCESS**, indicating that all data for the column has been retrieved. **SQLGetData** will return **SQL_NO_DATA_FOUND** when it is called for a column after all of the data has been retrieved and before data is retrieved for a subsequent column. The application can ignore excess data by proceeding to the next result column.

Note: An application can use **SQLGetData** to retrieve data from a column in parts only when retrieving character C data from a column with a character, binary, or data source-specific data type or when retrieving binary C data from a column with a character, binary, or data source-specific data type. If **SQLGetData** is called more than one time in a row for a column under any other conditions, it returns **SQL_NO_DATA_FOUND** for all calls after the first.

For maximum interoperability, applications should call **SQLGetData** only for unbound columns with numbers greater than the number of the last bound column. Within a single row of data, the column number in each call to **SQLGetData** should be greater than or equal to the column number in the previous call (that is, data should be retrieved in increasing order of column number). As extended functionality, drivers can return data through **SQLGetData** from bound columns, from columns before the last bound column, or from columns in any order. To determine whether a driver supports these extensions, an application calls **SQLGetInfo** with the `SQL_GETDATA_EXTENSIONS` option.

Furthermore, applications that use **SQLExtendedFetch** to retrieve data should call **SQLGetData** only when the rowset size is 1. As extended functionality, drivers can return data through **SQLGetData** when the rowset size is greater than 1. The application calls **SQLSetPos** to position the cursor on a row and calls **SQLGetData** to retrieve data from an unbound column. To determine whether a driver supports this extension, an application calls **SQLGetInfo** with the `SQL_GETDATA_EXTENSIONS` option.

Code Example

In the following example, an application executes a **SELECT** statement to return a result set of the employee names, ages, and birthdays sorted by birthday, age, and name. For each row of data, it calls **SQLFetch** to position the cursor to the next row. It calls **SQLGetData** to retrieve the fetched data; the storage locations for the data and the returned number of bytes are specified in the call to **SQLGetData**. Finally, it prints each employee's name, age, and birthday.

```
#define NAME_LEN 30
#define BDAY_LEN 11

UCHAR    szName[NAME_LEN], szBirthday[BDAY_LEN];
WORD     sAge;
SDWORD   cbName, cbAge, cbBirthday;

retcode = SQLExecDirect(hstmt,
    "SELECT NAME, AGE, BIRTHDAY FROM EMPLOYEE ORDER BY 3, 2, 1",
    SQL_NTS);

if (retcode == SQL_SUCCESS) {
    while (TRUE) {
        retcode = SQLFetch(hstmt);
        if (retcode == SQL_ERROR || retcode == SQL_SUCCESS_WITH_INFO) {
            show_error();
        }
        if (retcode == SQL_SUCCESS || retcode == SQL_SUCCESS_WITH_INFO){

            /* Get data for columns 1, 2, and 3 */
```

```
/* Print the row of data */

SQLGetData(hstmt, 1, SQL_C_CHAR, szName, NAME_LEN, &cbName);
SQLGetData(hstmt, 2, SQL_C_SSHORT, &sAge, 0, &cbAge);
SQLGetData(hstmt, 3, SQL_C_CHAR, szBirthday, BDAY_LEN,
           &cbBirthday);
fprintf(out, "%-*s %-2d %*s", NAME_LEN-1, szName, sAge,
        BDAY_LEN-1, szBirthday);
} else {
    break;
}
}
}
```


SQLProcedureColumns

Related Functions	For information about	See
	Assigning storage for a column in a result set	SQLBindCol
	Canceling statement processing	SQLCancel
	Executing an SQL statement	SQLExecDirect
	Executing a prepared SQL statement	SQLExecute
	Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
	Fetching a row of data	SQLFetch
	Sending parameter data at execution time	SQLPutData (extension)

SQLGetFunctions



Extension Level 1 **SQLGetFunctions** returns information about whether a driver supports a specific ODBC function. This function is implemented in the Driver Manager; it can also be implemented in drivers. If a driver implements **SQLGetFunctions**, the Driver Manager calls the function in the driver. Otherwise, it executes the function itself.

Syntax

RETCODE **SQLGetFunctions**(*hdbc*, *fFunction*, *pfExists*)

The **SQLGetFunctions** function accepts the following arguments:

Type	Argument	Use	Description
HDBC	<i>hdbc</i>	Input	Connection handle.
UWORD	<i>fFunction</i>	Input	SQL_API_ALL_FUNCTIONS or a #define value that identifies the ODBC function of interest. For a list of #define values that identify ODBC functions, see the tables in "Comments."
UWORD FAR *	<i>pfExists</i>	Output	If <i>fFunction</i> is SQL_API_ALL_FUNCTIONS, <i>pfExists</i> points to a UWORD array with 100 elements. The array is indexed by #define values used by <i>fFunction</i> to identify each ODBC function; some elements of the array are unused and reserved for future use. An element is TRUE if it identifies an ODBC function supported by the driver. It is FALSE if it identifies an ODBC function not supported by the driver or does not identify an ODBC function.

Note: The *fFunction* value SQL_API_ALL_FUNCTIONS was added in ODBC 2.0.

If *fFunction* identifies a single ODBC function, *pfExists* points to single UWORD. *pfExists* is TRUE if the specified function is supported by the driver; otherwise, it is FALSE.

Returns

SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics

When **SQLGetFunctions** returns `SQL_ERROR` or `SQL_SUCCESS_WITH_INFO`, an associated `SQLSTATE` value may be obtained by calling **SQLError**. The following table lists the `SQLSTATE` values commonly returned by **SQLGetFunctions** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of `SQLSTATES` returned by the Driver Manager. The return code associated with each `SQLSTATE` value is `SQL_ERROR`, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns <code>SQL_SUCCESS_WITH_INFO</code> .)
S1000	General error	An error occurred for which there was no specific <code>SQLSTATE</code> and for which no implementation-specific <code>SQLSTATE</code> was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1010	Function sequence error	(DM) SQLGetFunctions was called before SQLConnect , SQLBrowseConnect , or SQLDriverConnect . (DM) SQLBrowseConnect was called for the <i>hdbc</i> and returned <code>SQL_NEED_DATA</code> . This function was called before SQLBrowseConnect returned <code>SQL_SUCCESS_WITH_INFO</code> or <code>SQL_SUCCESS</code> .
S1095	Function type out of range	(DM) An invalid <i>fFunction</i> value was specified.

Comments

SQLGetFunctions always returns that **SQLGetFunctions**, **SQLDataSources**, and **SQLDrivers** are supported. It does this because these functions are implemented in the Driver Manager.

The following table lists valid values for *fFunction* for ODBC core functions.

SQL_API_SQLALLOCCONNECT	SQL_API_SQLFETCH
SQL_API_SQLALLOCENV	SQL_API_SQLFREECONNECT
SQL_API_SQLALLOCSTMT	SQL_API_SQLFREEENV
SQL_API_SQLBINDCOL	SQL_API_SQLFREESTMT
SQL_API_SQLCANCEL	SQL_API_SQLGETCURSORNAME
SQL_API_SQLCOLATTRIBUTES	SQL_API_SQLNUMRESULTCOLS
SQL_API_SQLCONNECT	SQL_API_SQLPREPARE
SQL_API_SQLDESCRIBECOL	SQL_API_SQLROWCOUNT
SQL_API_SQLDISCONNECT	SQL_API_SQLSETCURSORNAME
SQL_API_SQLERROR	SQL_API_SQLSETPARAM
SQL_API_SQLEXECDIRECT	SQL_API_SQLTRANSACT
SQL_API_SQLEXECUTE	

Note: For ODBC 1.0 drivers, **SQLGetFunctions** returns TRUE in *pfExists* if *fFunction* is SQL_API_SQLBINDPARAMETER or SQL_API_SQLSETPARAM and the driver supports **SQLSetParam**. For ODBC 2.0 drivers, **SQLGetFunctions** returns TRUE in *pfExists* if *fFunction* is SQL_API_SQLSETPARAM or SQL_API_SQLBINDPARAMETER and the driver supports **SQLBindParameter**.

The following table lists valid values for *fFunction* for ODBC extension level 1 functions.

SQL_API_SQLBINDPARAMETER	SQL_API_SQLGETTYPEINFO
SQL_API_SQLCOLUMNS	SQL_API_SQLPARAMDATA
SQL_API_SQLDRIVERCONNECT	SQL_API_SQLPUTDATA
SQL_API_SQLGETCONNECTOPTION	SQL_API_SQLSETCONNECTOPTION
SQL_API_SQLGETDATA	SQL_API_SQLSETSTMTOPTION
SQL_API_SQLGETFUNCTIONS	SQL_API_SQLSPECIALCOLUMNS
SQL_API_SQLGETINFO	SQL_API_SQLSTATISTICS
SQL_API_SQLGETSTMTOPTION	SQL_API_SQLTABLES

The following table lists valid values for *fFunction* for ODBC extension level 2 functions.

SQL_API_SQLBROWSECONNECT	SQL_API_SQLNUMPARAMS
--------------------------	----------------------

SQLProcedures

SQL_API_SQLCOLUMNPRIVILEGES	SQL_API_SQLPARAMOPTIONS
SQL_API_SQLDATASOURCES	SQL_API_SQLPRIMARYKEYS
SQL_API_SQLDESCRIBEPARAM	SQL_API_SQLPROCEDURECOLUMNS
SQL_API_SQLDRIVERS	SQL_API_SQLPROCEDURES
SQL_API_SQLEXTENDEDFETCH	SQL_API_SQLSETPOS
SQL_API_SQLFOREIGNKEYS	SQL_API_SQLSETSCROLLOPTIONS
SQL_API_SQLMORERESULTS	SQL_API_SQLTABLEPRIVILEGES
SQL_API_SQLNATIVESQL	

Code Example The following two examples show how an application uses **SQLGetFunctions** to determine if a driver supports **SQLTables**, **SQLColumns**, and **SQLStatistics**. If the driver does not support these functions, the application disconnects from the driver. The first example calls **SQLGetFunctions** once for each function.

```
UWORD TablesExists, ColumnsExists, StatisticsExists;
```

```
SQLGetFunctions(hdbc, SQL_API_SQLTABLES, &TablesExists);
SQLGetFunctions(hdbc, SQL_API_SQLCOLUMNS, &ColumnsExists);
SQLGetFunctions(hdbc, SQL_API_SQLSTATISTICS, &StatisticsExists);
```

```
if (TablesExists && ColumnsExists && StatisticsExists) {
```

```
    /* Continue with application */
```

```
}
SQLDisconnect(hdbc);
```

The second example calls **SQLGetFunctions** a single time and passes it an array in which **SQLGetFunctions** returns information about all ODBC functions.

```
UWORD fExists[100];
```

```
SQLGetFunctions(hdbc, SQL_API_ALL_FUNCTIONS, fExists);
```

```
if (fExists[SQL_API_SQLTABLES] &&
    fExists[SQL_API_SQLCOLUMNS] &&
    fExists[SQL_API_SQLSTATISTICS]) {
```

```
    /* Continue with application */
```

```
}
SQLDisconnect(hdbc);
```

**Related
Functions****For information about****See**

Returning the setting of a connection option

SQLGetConnectOption
(extension)

Returning information about a driver or data source

SQLGetInfo (extension)

Returning the setting of a statement option

SQLGetStmtOption (extension)

SQLGetInfo



Extension Level 1 **SQLGetInfo** returns general information about the driver and data source associated with an *hdbc*.

Syntax RETCODE **SQLGetInfo**(*hdbc*, *fInfoType*, *rgbInfoValue*, *cbInfoValueMax*, *pcbInfoValue*)

The **SQLGetInfo** function accepts the following arguments.

Type	Argument	Use	Description
HDBC	<i>hdbc</i>	Input	Connection handle.
UWORD	<i>fInfoType</i>	Input	Type of information. <i>fInfoType</i> must be a value representing the type of interest (see "Comments").
PTR	<i>rgbInfoValue</i>	Output	Pointer to storage for the information. Depending on the <i>fInfoType</i> requested, the information returned will be one of the following: a null-terminated character string, a 16-bit integer value, a 32-bit flag, or a 32-bit binary value.
SWORD	<i>cbInfoValueMax</i>	Input	Maximum length of the <i>rgbInfoValue</i> buffer.
SWORD FAR *	<i>pcbInfoValue</i>	Output	The total number of bytes (excluding the null termination byte for character data) available to return in <i>rgbInfoValue</i> . For character data, if the number of bytes available to return is greater than or equal to <i>cbInfoValueMax</i> , the information in <i>rgbInfoValue</i> is truncated to <i>cbInfoValueMax</i> – 1 bytes and is null-terminated by the driver. For all other types of data, the value of <i>cbValueMax</i> is ignored and the driver assumes the size of <i>rgbValue</i> is 32 bits.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLGetInfo** returns either SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values

commonly returned by **SQLGetInfo** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	The buffer <i>rgbInfoValue</i> was not large enough to return all of the requested information, so the information was truncated. The argument <i>pcbInfoValue</i> contains the length of the requested information in its untruncated form. (Function returns SQL_SUCCESS_WITH_INFO.)
08003	Connection not open	(DM) The type of information requested in <i>fInfoType</i> requires an open connection. Of the information types reserved by ODBC, only SQL_ODBC_VER can be returned without an open connection.
22003	Numeric value out of range	Returning the requested information would have caused a loss of numeric or binary significance.
IM001	Driver does not support this function	(DM) The driver corresponding to the <i>hdbc</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value	(DM) The <i>fInfoType</i> was SQL_DRIVER_HSTMT, and the value pointed to by <i>rgbInfoValue</i> was not a valid statement handle.
S1090	Invalid string or buffer length	(DM) The value specified for argument <i>cbInfoValueMax</i> was less than 0.
S1096	Information type out of range	(DM) The value specified for the argument <i>fOption</i> was in the block of numbers reserved for ODBC information types, but was not valid

SQLPutData

SQLPutData

		for the version of ODBC supported by the driver.
SQLSTATE	Error	Description
S1C00	Driver not capable	The value specified for the argument <i>fOption</i> was in the range of numbers reserved for driver-specific information types, but was not supported by the driver.
S1T00	Timeout expired	The timeout period expired before the data source returned the requested information. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT .

Comments

The currently defined information types are shown below; it is expected that more will be defined to take advantage of different data sources. Information types from 0 to 999 are reserved by ODBC; driver developers must reserve values greater than or equal to **SQL_INFO_DRIVER_START** for driver-specific use.

The format of the information returned in *rgbInfoValue* depends on the *fInfoType* requested. **SQLGetInfo** will return information in one of five different formats:

- A null-terminated character string,
- A 16-bit integer value,
- A 32-bit bitmask,
- A 32-bit integer value,
- Or a 32-bit binary value.

The format of each of the following information types is noted in the type's description. The application must cast the value returned in *rgbInfoValue* accordingly. For an example of how an application could retrieve data from a 32-bit bitmask, see "Code Example."

A driver must return a value for each of the information types defined in the following tables. If an information type does not apply to the driver or data source, then the driver returns one of the following values:

Format of <i>rgbInfoValue</i>	Returned value
Character string ("Y" or "N")	"N"
Character string (not "Y" or "N")	Empty string
16-bit integer	0
32-bit bitmask or 32-bit binary value	0L

For example, if a data source does not support procedures, **SQLGetInfo** returns the following values for the values of *fInfoType* that are related to procedures:

<i>fInfoType</i>	Returned value
SQL_PROCEDURES	"N"
SQL_ACCESSIBLE_PROCEDURES	"N"
SQL_MAX_PROCEDURE_NAME_LEN	0
SQL_PROCEDURE_TERM	Empty string

SQLGetInfo returns SQLSTATE S1096 (Invalid argument value) for values of *fInfoType* that are in the range of information types reserved for use by ODBC but are not defined by the version of ODBC supported by the driver. To determine what version of ODBC a driver conforms to, an application calls **SQLGetInfo** with the SQL_DRIVER_ODBC_VER information type. **SQLGetInfo** returns SQLSTATE S1C00 (Driver not capable) for values of *fInfoType* that are in the range of information types reserved for driver-specific use but are not supported by the driver.

Note: Application developers should be aware that ODBC 1.0 drivers might return SQL_ERROR and SQLSTATE S1C00 (Driver not capable) for values of *fInfoType* that were defined in ODBC 1.0 but do not apply to the driver or the data source.

Information Types

This section lists the information types supported by **SQLGetInfo**. Information types are grouped categorically and listed alphabetically.

Driver Information

The following values of *fInfoType* return information about the ODBC driver, such as the number of active statements, the data source name, and the API conformance levels.

SQL_ACTIVE_CONNECTIONS	SQL_DRIVER_VER
SQL_ACTIVE_STATEMENTS	SQL_FETCH_DIRECTION
SQL_DATA_SOURCE_NAME	SQL_FILE_USAGE
SQL_DRIVER_HDBC	SQL_GETDATA_EXTENSIONS
SQL_DRIVER_HENV	SQL_LOCK_TYPES
SQL_DRIVER_HLIB	SQL_ODBC_API_CONFORMANCE
	SQL_ODBC_SAG_CLI_CONFORMANC

SQLPutData

SQLPutData

SQL_DRIVER_HSTMT	E
SQL_DRIVER_NAME	SQL_ODBC_VER
SQL_DRIVER_ODBC_VER	SQL_POS_OPERATIONS
SQL_ROW_UPDATES	SQL_SERVER_NAME
SQL_SEARCH_PATTERN_ESCAPE	

DBMS Product Information

The following values of *fnInfoType* return information about the DBMS product, such as the DBMS name and version.

SQL_DATABASE_NAME
SQL_DBMS_NAME
SQL_DBMS_VER

Data Source Information

The following values of *fnInfoType* return information about the data source, such as cursor characteristics and transaction capabilities.

SQL_ACCESSIBLE_PROCEDURES	SQL_NULL_COLLATION
SQL_ACCESSIBLE_TABLES	SQL_OWNER_TERM
SQL_BOOKMARK_PERSISTENCE	SQL_PROCEDURE_TERM
SQL_CONCAT_NULL_BEHAVIOR	SQL_QUALIFIER_TERM
SQL_CURSOR_COMMIT_BEHAVIOR	SQL_SCROLL_CONCURRENCY
SQL_CURSOR_ROLLBACK_BEHAVIOR	SQL_SCROLL_OPTIONS
SQL_DATA_SOURCE_READ_ONLY	SQL_STATIC_SENSITIVITY
SQL_DEFAULT_TXN_ISOLATION	SQL_TABLE_TERM
SQL_MULT_RESULT_SETS	SQL_TXN_CAPABLE
SQL_MULTIPLE_ACTIVE_TXN	SQL_TXN_ISOLATION_OPTIONS
SQL_NEED_LONG_DATA_LEN	SQL_USER_NAME

Supported SQL

The following values of *fnInfoType* return information about the SQL statements supported by the data source. These information types do not exhaustively

describe the entire ODBC SQL grammar. Instead, they describe those parts of the grammar for which data sources commonly offer different levels of support.

Applications should determine the general level of supported grammar from the SQL_ODBC_SQL_CONFORMANCE information type and use the other information types to determine variations from the stated conformance level.

SQL_ALTER_TABLE	SQL_OUTER_JOINS
SQL_COLUMN_ALIAS	SQL_OWNER_USAGE
SQL_CORRELATION_NAME	SQL_POSITIONED_STATEMENTS
SQL_EXPRESSIONS_IN_ORDERBY	SQL_PROCEDURES
SQL_GROUP_BY	SQL_QUALIFIER_LOCATION
SQL_IDENTIFIER_CASE	SQL_QUALIFIER_NAME_SEPARATOR
SQL_IDENTIFIER_QUOTE_CHAR	SQL_QUALIFIER_USAGE
SQL_KEYWORDS	SQL_QUOTED_IDENTIFIER_CASE
SQL_LIKE_ESCAPE_CLAUSE	SQL_SPECIAL_CHARACTERS
SQL_NON_NULLABLE_COLUMNS	SQL_SUBQUERIES
SQL_ODBC_SQL_CONFORMANCE	SQL_UNION
SQL_ODBC_SQL_OPT_IEF	
SQL_ORDER_BY_COLUMNS_IN_SELECT	

SQL Limits

The following values of *InfoType* return information about the limits applied to identifiers and clauses in SQL statements, such as the maximum lengths of identifiers and the maximum number of columns in a select list. Limitations may be imposed by either the driver or the data source.

SQL_MAX_BINARY_LITERAL_LEN	SQL_MAX_OWNER_NAME_LEN
SQL_MAX_CHAR_LITERAL_LEN	SQL_MAX_PROCEDURE_NAME_LEN
SQL_MAX_COLUMN_NAME_LEN	SQL_MAX_QUALIFIER_NAME_LEN
SQL_MAX_COLUMNS_IN_GROUP_BY	SQL_MAX_ROW_SIZE
SQL_MAX_COLUMNS_IN_ORDER_BY	SQL_MAX_ROW_SIZE_INCLUDES_LONG

SQLPutData

SQL_MAX_COLUMNS_IN_INDEX	SQL_MAX_STATEMENT_LEN
SQL_MAX_COLUMNS_IN_SELECT	SQL_MAX_TABLE_NAME_LEN
SQL_MAX_COLUMNS_IN_TABLE	SQL_MAX_TABLES_IN_SELECT
SQL_MAX_CURSOR_NAME_LEN	SQL_MAX_USER_NAME_LEN
SQL_MAX_INDEX_SIZE	

Scalar Function Information

The following values of *fnInfoType* return information about the scalar functions supported by the data source and the driver. For more information about scalar functions, see Appendix F, "Scalar Functions."

SQL_CONVERT_FUNCTIONS	SQL_TIMEDATE_ADD_INTERVALS
SQL_NUMERIC_FUNCTIONS	SQL_TIMEDATE_DIFF_INTERVALS
SQL_STRING_FUNCTIONS	SQL_TIMEDATE_FUNCTIONS
SQL_SYSTEM_FUNCTIONS	

Conversion Information

The following values of *fnInfoType* return a list of the SQL data types to which the data source can convert the specified SQL data type with the **CONVERT** scalar function.

SQL_CONVERT_BIGINT	SQL_CONVERT_LONGVARCHA R
SQL_CONVERT_BINARY	
SQL_CONVERT_BIT	SQL_CONVERT_NUMERIC
SQL_CONVERT_CHAR	SQL_CONVERT_REAL
SQL_CONVERT_DATE	SQL_CONVERT_SMALLINT
SQL_CONVERT_DECIMAL	SQL_CONVERT_TIME
SQL_CONVERT_DOUBLE	SQL_CONVERT_TIMESTAMP
SQL_CONVERT_FLOAT	SQL_CONVERT_TINYINT
SQL_CONVERT_INTEGER	SQL_CONVERT_VARBINARY
SQL_CONVERT_LONGVARBINAR Y	SQL_CONVERT_VARCHAR

Information Type Descriptions

The following table alphabetically lists each information type, the version of ODBC in which it was introduced, and its description.

InfoType	Returns
SQL_ACCESSIBLE_PROCEDURES (ODBC 1.0)	A character string: "Y" if the user can execute all procedures returned by SQLProcedures , "N" if there may be procedures returned that the user cannot execute.
SQL_ACCESSIBLE_TABLES (ODBC 1.0)	A character string: "Y" if the user is guaranteed SELECT privileges to all tables returned by SQLTables , "N" if there may be tables returned that the user cannot access.
SQL_ACTIVE_CONNECTIONS	A 16-bit integer value specifying the maximum number of active SQLPutData

SQLPutData

(ODBC 1.0)

hdbcs that the driver can support. This value can reflect a limitation imposed by either the driver or the data source. If there is no specified limit or the limit is unknown, this value is set to zero.

SQL_ACTIVE_STATEMENTS
(ODBC 1.0)

A 16-bit integer value specifying the maximum number of active *hstmts* that the driver can support for an *hdbc*. This value can reflect a limitation imposed by either the driver or the data source. If there is no specified limit or the limit is unknown, this value is set to zero.

SQL_ALTER_TABLE
(ODBC 2.0)

A 32-bit bitmask enumerating the clauses in the **ALTER TABLE** statement supported by the data source.

The following bitmask is used to determine which clauses are supported:

SQL_AT_ADD_COLUMN
SQL_AT_DROP_COLUMN

SQL_BOOKMARK_PERSISTENCE
(ODBC 2.0)

A 32-bit bitmask enumerating the operations through which bookmarks persist.

The following bitmasks are used in conjunction with the flag to determine through which options bookmarks persist:

SQL_BP_CLOSE = Bookmarks are valid after an application calls **SQLFreeStmt** with the SQL_CLOSE option to close the cursor associated with an *hstmt*.

SQL_BP_DELETE = The bookmark for a row is valid after that row has been deleted.

SQL_BP_DROP = Bookmarks are valid after an *hstmt* an application calls **SQLFreeStmt** with the SQL_DROP option to drop an *hstmt*.

SQL_BP_SCROLL = Bookmarks are valid after any scrolling operation (call to **SQLExtendedFetch**). Because all bookmarks must remain valid after **SQLExtendedFetch** is called, this value can be used by applications to determine whether bookmarks are supported.

InfoType

Returns

SQL_BP_TRANSACTION = Bookmarks are valid after an application commits or rolls back a transaction.

SQL_BP_UPDATE = The bookmark for a row is valid after any column in that row has been updated, including key columns.

SQL_BP_OTHER_HSTMT = A bookmark associated with one *hstmt* can be used with another *hstmt*.

SQL_COLUMN_ALIAS
(ODBC 2.0)

A character string: "Y" if the data source supports column aliases; otherwise, "N".

SQL_CONCAT_NULL_BEHAVIOR
(ODBC 1.0)

A 16-bit integer value indicating how the data source handles the concatenation of NULL valued character data type columns with non-NULL valued character data type columns:

SQL_CB_NULL = Result is NULL valued.

SQL_CB_NON_NULL = Result is concatenation of non-NULL valued column or columns.

InfoType

Returns

SQL_CONVERT_BIGINT
SQL_CONVERT_BINARY
SQL_CONVERT_BIT
SQL_CONVERT_CHAR
SQL_CONVERT_DATE
SQL_CONVERT_DECIMAL
SQL_CONVERT_DOUBLE
SQL_CONVERT_FLOAT
SQL_CONVERT_INTEGER
SQL_CONVERT_LONGVARBINARY
SQL_CONVERT_LONGVARCHAR
SQL_CONVERT_NUMERIC
SQL_CONVERT_REAL
SQL_CONVERT_SMALLINT
SQL_CONVERT_TIME
SQL_CONVERT_TIMESTAMP
SQL_CONVERT_TINYINT
SQL_CONVERT_VARBINARY
SQL_CONVERT_VARCHAR
(ODBC 1.0)

A 32-bit bitmask. The bitmask indicates the conversions supported by the data source with the CONVERT scalar function for data of the type named in the *fInfoType*. If the bitmask equals zero, the data source does not support any conversions for data of the named type, including conversion to the same data type.

For example, to find out if a data source supports the conversion of SQL_INTEGER data to the SQL_BIGINT data type, an application calls **SQLGetInfo** with the *fInfoType* of SQL_CONVERT_INTEGER. The application ANDs the returned bitmask with SQL_CVT_BIGINT. If the resulting value is nonzero, the conversion is supported.

The following bitmasks are used to determine which conversions are supported:

SQL_CVT_BIGINT
SQL_CVT_BINARY
SQL_CVT_BIT
SQL_CVT_CHAR
SQL_CVT_DATE
SQL_CVT_DECIMAL
SQL_CVT_DOUBLE
SQL_CVT_FLOAT
SQL_CVT_INTEGER
SQL_CVT_LONGVARBINARY
SQL_CVT_LONGVARCHAR
SQL_CVT_NUMERIC
SQL_CVT_REAL
SQL_CVT_SMALLINT
SQL_CVT_TIME
SQL_CVT_TIMESTAMP
SQL_CVT_TINYINT
SQL_CVT_VARBINARY
SQL_CVT_VARCHAR

SQL_CONVERT_FUNCTIONS
(ODBC 1.0)

A 32-bit bitmask enumerating the scalar conversion functions supported by the driver and associated data source.

The following bitmask is used to determine which conversion functions are supported:

SQL_FN_CVT_CONVERT

SQL_CORRELATION_NAME

A 16-bit integer indicating if table correlation names are

SQLPutData

SQLPutData

(ODBC 1.0)

supported:

SQL_CN_NONE = Correlation names are not supported.

SQL_CN_DIFFERENT = Correlation names are supported, but must differ from the names of the tables they represent.

SQL_CN_ANY = Correlation names are supported and can be any valid user-defined name.

InfoType

Returns

SQL_CURSOR_COMMIT_BEHAVIOR
(ODBC 1.0)

A 16-bit integer value indicating how a **COMMIT** operation affects cursors and prepared statements in the data source:

SQL_CB_DELETE = Close cursors and delete prepared statements. To use the cursor again, the application must reprepare and reexecute the *hstmt*.

SQL_CB_CLOSE = Close cursors. For prepared statements, the application can call **SQLExecute** on the *hstmt* without calling **SQLPrepare** again.

SQL_CB_PRESERVE = Preserve cursors in the same position as before the **COMMIT** operation. The application can continue to fetch data or it can close the cursor and reexecute the *hstmt* without reparing it.

SQL_CURSOR_ROLLBACK_BEHAVIOR
(ODBC 1.0)

A 16-bit integer value indicating how a **ROLLBACK** operation affects cursors and prepared statements in the data source:

SQL_CB_DELETE = Close cursors and delete prepared statements. To use the cursor again, the application must reprepare and reexecute the *hstmt*.

SQL_CB_CLOSE = Close cursors. For prepared statements, the application can call **SQLExecute** on the *hstmt* without calling **SQLPrepare** again.

SQL_CB_PRESERVE = Preserve cursors in the same position as before the **ROLLBACK** operation. The application can continue to fetch data or it can close the cursor and reexecute the *hstmt* without reparing it.

SQL_DATA_SOURCE_NAME
(ODBC 1.0)

A character string with the data source name used during connection. If the application called **SQLConnect**, this is the value of the *szDSN* argument. If the application called **SQLDriverConnect** or **SQLBrowseConnect**, this is the value of the DSN keyword in the connection string passed to the driver. If the connection string did not contain the DSN keyword (such as when it contains the DRIVER keyword), this is an empty string.

SQL_DATA_SOURCE_READ_ONLY
(ODBC 1.0)

A character string. "Y" if the data source is set to READ ONLY mode, "N" if it is otherwise.

This characteristic pertains only to the data source itself, it is not a characteristic of the driver that enables access to the data source.

InfoType

Returns

ADABAS D SQL_DATABASE_NAME (ODBC 1.0)	<div>280</div> <p>A character string with the name of the current database in use, if the data source defines a named object called "database."</p> <hr/> <p>Note: In ODBC 2.0, this value of <i>InfoType</i> has been replaced by the SQL_CURRENT_QUALIFIER connection option. ODBC 2.0 drivers should continue to support the SQL_DATABASE_NAME information type, and ODBC 2.0 applications should only use it with ODBC 1.0 drivers.</p> <hr/>
SQL_DBMS_NAME (ODBC 1.0) SQL_DBMS_VER (ODBC 1.0)	<p>A character string with the name of the DBMS product accessed by the driver.</p> <p>A character string indicating the version of the DBMS product accessed by the driver. The version is of the form <i>###.###.####</i>, where the first two digits are the major version, the next two digits are the minor version, and the last four digits are the release version. The driver must render the DBMS product version in this form, but can also append the DBMS product-specific version as well. For example, "04.01.0000 Rdb 4.1".</p>
SQL_DEFAULT_TXN_ISOLATION (ODBC 1.0)	<p>A 32-bit integer that indicates the default transaction isolation level supported by the driver or data source, or zero if the data source does not support transactions. The following terms are used to define transaction isolation levels:</p> <p>Dirty Read Transaction 1 changes a row. Transaction 2 reads the changed row before transaction 1 commits the change. If transaction 1 rolls back the change, transaction 2 will have read a row that is considered to have never existed.</p> <p>Nonrepeatable Read Transaction 1 reads a row. Transaction 2 updates or deletes that row and commits this change. If transaction 1 attempts to reread the row, it will receive different row values or discover that the row has been deleted.</p> <p>Phantom Transaction 1 reads a set of rows that satisfy some search criteria. Transaction 2 inserts a row that matches the search criteria. If transaction 1 reexecutes the statement that read the rows, it receives a different set of rows.</p> <p>If the data source supports transactions, the driver returns one of the following bitmasks:</p> <p>SQL_TXN_READ_UNCOMMITTED = Dirty reads, nonrepeatable reads, and phantoms are possible.</p> <p>SQL_TXN_READ_COMMITTED = Dirty reads are not possible. Nonrepeatable reads and phantoms are possible.</p> <p>SQL_TXN_REPEATABLE_READ = Dirty reads and nonrepeatable reads are not possible. Phantoms are possible.</p>
InfoType	Returns
	SQL_TXN_SERIALIZABLE = Transactions are serializable.

SQLPutData

	<p>Dirty reads, nonrepeatable reads, and phantoms are not possible.</p> <p>SQL_TXN_VERSIONING = Transactions are serializable, but higher concurrency is possible than with SQL_TXN_SERIALIZABLE. Dirty reads are not possible. Typically, SQL_TXN_SERIALIZABLE is implemented by using locking protocols that reduce concurrency and SQL_TXN_VERSIONING is implemented by using a non-locking protocol such as record versioning. Oracle's Read Consistency isolation level is an example of SQL_TXN_VERSIONING.</p>
SQL_DRIVER_HDBC SQL_DRIVER_HENV (ODBC 1.0)	<p>A 32-bit value, the driver's environment handle or connection handle, determined by the argument <i>hdbc</i>.</p> <p>These information types are implemented by the Driver Manager alone.</p>
SQL_DRIVER_HLIB (ODBC 2.0)	<p>A 32-bit value, the library handle returned to the Driver Manager when it loaded the driver DLL. The handle is only valid for the <i>hdbc</i> specified in the call to SQLGetInfo.</p> <p>This information type is implemented by the Driver Manager alone.</p>
SQL_DRIVER_HSTMT (ODBC 1.0)	<p>A 32-bit value, the driver's statement handle determined by the Driver Manager statement handle, which must be passed on input in <i>rgbInfoValue</i> from the application. Note that in this case, <i>rgbInfoValue</i> is both an input and an output argument. The input <i>hstmt</i> passed in <i>rgbInfoValue</i> must have been an <i>hstmt</i> allocated on the argument <i>hdbc</i>.</p> <p>This information type is implemented by the Driver Manager alone.</p>
SQL_DRIVER_NAME (ODBC 1.0)	<p>A character string with the filename of the driver used to access the data source.</p>
SQL_DRIVER_ODBC_VER (ODBC 2.0)	<p>A character string with the version of ODBC that the driver supports. The version is of the form <i>##.##</i>, where the first two digits are the major version and the next two digits are the minor version. SQL_SPEC_MAJOR and SQL_SPEC_MINOR define the major and minor version numbers. For the version of ODBC described in this manual, these are 2 and 0, and the driver should return "02.00".</p> <p>If a driver supports SQLGetInfo but does not support this value of the <i>fInfoType</i> argument, the Driver Manager returns "01.00".</p>
InfoType	Returns
SQL_DRIVER_VER (ODBC 1.0)	<p>A character string with the version of the driver and, optionally a description of the driver. At a minimum, the version is of the form <i>##.##.####</i>, where the first two digits are the major version, the</p>

	next two digits are the minor version, and the last four digits are the release version.
SQL_EXPRESSIONS_IN_ORDERBY (ODBC 1.0)	A character string: "Y" if the data source supports expressions in the ORDER BY list; "N" if it does not.
SQL_FETCH_DIRECTION (ODBC 1.0)	A 32-bit bitmask enumerating the supported fetch direction options.
The information type was introduced in ODBC 1.0; each bitmask is labeled with the version in which it was introduced.	The following bitmasks are used in conjunction with the flag to determine which options are supported:
	SQL_FD_FETCH_NEXT (ODBC 1.0)
	SQL_FD_FETCH_FIRST (ODBC 1.0)
	SQL_FD_FETCH_LAST (ODBC 1.0)
	SQL_FD_FETCH_PRIOR (ODBC 1.0)
	SQL_FD_FETCH_ABSOLUTE (ODBC 1.0)
	SQL_FD_FETCH_RELATIVE (ODBC 1.0)
	SQL_FD_FETCH_RESUME (ODBC 1.0)
	SQL_FD_FETCH_BOOKMARK (ODBC 2.0)
SQL_FILE_USAGE (ODBC 2.0)	A 16-bit integer value indicating how a single-tier driver directly treats files in a data source:
	SQL_FILE_NOT_SUPPORTED = The driver is not a single-tier driver. For example, an ORACLE driver is a two-tier driver.
	SQL_FILE_TABLE = A single-tier driver treats files in a data source as tables. For example, an Xbase driver treats each Xbase file as a table.
	SQL_FILE_QUALIFIER = A single-tier driver treats files in a data source as a qualifier. For example, a Microsoft Access driver treats each Microsoft Access file as a complete database.
	An application might use this to determine how users will select data. For example, Xbase users often think of data as stored in files, while ORACLE and Microsoft Access users generally think of data as stored in tables.
	When a user selects an Xbase data source, the application could display the Windows File Open common dialog box; when the user selects a Microsoft Access or ORACLE data source, the application could display a custom Select Table dialog box.

InfoType**Returns**

SQL_GETDATA_EXTENSIONS
(ODBC 2.0)

A 32-bit bitmask enumerating extensions to **SQLGetData**. The following bitmasks are used in conjunction with the flag to determine what common extensions the driver supports for **SQLGetData**:

SQL_GD_ANY_COLUMN = **SQLGetData** can be called for any unbound column, including those before the last bound column. Note that the columns must be called in order of ascending column number unless SQL_GD_ANY_ORDER is also returned.

SQLPutData

SQLPutData

SQL_GROUP_BY
(ODBC 2.0)

SQL_GD_ANY_ORDER = **SQLGetData** can be called for unbound columns in any order. Note that **SQLGetData** can only be called for columns after the last bound column unless SQL_GD_ANY_COLUMN is also returned.

SQL_GD_BLOCK = **SQLGetData** can be called for an unbound column in any row in a block (more than one row) of data after positioning to that row with **SQLSetPos**.

SQL_GD_BOUND = **SQLGetData** can be called for bound columns as well as unbound columns. A driver cannot return this value unless it also returns SQL_GD_ANY_COLUMN.

SQLGetData is only required to return data from unbound columns that occur after the last bound column, are called in order of increasing column number, and are not in a row in a block of rows.

A 16-bit integer value specifying the relationship between the columns in the **GROUP BY** clause and the non-aggregated columns in the select list:

SQL_GB_NOT_SUPPORTED = **GROUP BY** clauses are not supported.

SQL_GB_GROUP_BY_EQUALS_SELECT = The **GROUP BY** clause must contain all non-aggregated columns in the select list. It cannot contain any other columns. For example, **SELECT DEPT, MAX(SALARY) FROM EMPLOYEE GROUP BY DEPT**.

SQL_GB_GROUP_BY_CONTAINS_SELECT = The **GROUP BY** clause must contain all non-aggregated columns in the select list. It can contain columns that are not in the select list. For example, **SELECT DEPT, MAX(SALARY) FROM EMPLOYEE GROUP BY DEPT, AGE**.

SQL_GB_NO_RELATION = The columns in the **GROUP BY** clause and the select list are not related. The meaning of non-grouped, non-aggregated columns in the select list is data source-dependent. For example, **SELECT DEPT, SALARY FROM EMPLOYEE GROUP BY DEPT, AGE**.

InfoType

SQL_IDENTIFIER_CASE
(ODBC 1.0)

Returns

A 16-bit integer value as follows:

SQL_IC_UPPER = Identifiers in SQL are case insensitive and are stored in upper case in system catalog.

SQL_IC_LOWER = Identifiers in SQL are case insensitive and are stored in lower case in system catalog.

SQL_IC_SENSITIVE = Identifiers in SQL are case sensitive and are stored in mixed case in system catalog.

SQL_IC_MIXED = Identifiers in SQL are case insensitive and

	are stored in mixed case in system catalog.
SQL_IDENTIFIER_QUOTE_CHAR (ODBC 1.0)	The character string used as the starting and ending delimiter of a quoted (delimited) identifiers in SQL statements. (Identifiers passed as arguments to ODBC functions do not need to be quoted.) If the data source does not support quoted identifiers, a blank is returned.
SQL_KEYWORDS (ODBC 2.0)	A character string containing a comma-separated list of all data source-specific keywords. This list does not contain keywords specific to ODBC or keywords used by both the data source and ODBC. For a list of ODBC keywords, see "List of Reserved Keywords" in Appendix C, "SQL Grammar." The #define value SQL_ODBC_KEYWORDS contains a comma-separated list of ODBC keywords.
SQL_LIKE_ESCAPE_CLAUSE (ODBC 2.0)	A character string: "Y" if the data source supports an escape character for the percent character (%) and underscore character (_) in a LIKE predicate and the driver supports the ODBC syntax for defining a LIKE predicate escape character; "N" otherwise.
SQL_LOCK_TYPES (ODBC 2.0)	A 32-bit bitmask enumerating the supported lock types for the <i>fLock</i> argument in SQLSetPos . The following bitmasks are used in conjunction with the flag to determine which lock types are supported: SQL_LCK_NO_CHANGE SQL_LCK_EXCLUSIVE SQL_LCK_UNLOCK
SQL_MAX_BINARY_LITERAL_LEN (ODBC 2.0)	A 32-bit integer value specifying the maximum length (number of hexadecimal characters, excluding the literal prefix and suffix returned by SQLGetTypeInfo) of a binary literal in an SQL statement. For example, the binary literal 0xFFAA has a length of 4. If there is no maximum length or the length is unknown, this value is set to zero.
InfoType	Returns
SQL_MAX_CHAR_LITERAL_LEN (ODBC 2.0)	A 32-bit integer value specifying the maximum length (number of characters, excluding the literal prefix and suffix returned by SQLGetTypeInfo) of a character literal in an SQL statement. If there is no maximum length or the length is unknown, this value is set to zero.
SQL_MAX_COLUMN_NAME_LEN (ODBC 1.0)	A 16-bit integer value specifying the maximum length of a column name in the data source. If there is no maximum length or the length is unknown, this value is set to zero.
SQL_MAX_COLUMNS_IN_GROUP_BY (ODBC 2.0)	A 16-bit integer value specifying the maximum number of columns allowed in a GROUP BY clause. If there is no specified limit or the limit is unknown, this value is set to zero.

SQLPutData

SQL_MAX_COLUMNS_IN_INDEX (ODBC 2.0)	A 16-bit integer value specifying the maximum number of columns allowed in an index. If there is no specified limit or the limit is unknown, this value is set to zero.
SQL_MAX_COLUMNS_IN_ORDER_BY (ODBC 2.0)	A 16-bit integer value specifying the maximum number of columns allowed in an ORDER BY clause. If there is no specified limit or the limit is unknown, this value is set to zero.
SQL_MAX_COLUMNS_IN_SELECT (ODBC 2.0)	A 16-bit integer value specifying the maximum number of columns allowed in a select list. If there is no specified limit or the limit is unknown, this value is set to zero.
SQL_MAX_COLUMNS_IN_TABLE (ODBC 2.0)	A 16-bit integer value specifying the maximum number of columns allowed in a table. If there is no specified limit or the limit is unknown, this value is set to zero.
SQL_MAX_CURSOR_NAME_LEN (ODBC 1.0)	A 16-bit integer value specifying the maximum length of a cursor name in the data source. If there is no maximum length or the length is unknown, this value is set to zero.
SQL_MAX_INDEX_SIZE (ODBC 2.0)	A 32-bit integer value specifying the maximum number of bytes allowed in the combined fields of an index. If there is no specified limit or the limit is unknown, this value is set to zero.
SQL_MAX_OWNER_NAME_LEN (ODBC 1.0)	A 16-bit integer value specifying the maximum length of an owner name in the data source. If there is no maximum length or the length is unknown, this value is set to zero.
SQL_MAX_PROCEDURE_NAME_LEN (ODBC 1.0)	A 16-bit integer value specifying the maximum length of a procedure name in the data source. If there is no maximum length or the length is unknown, this value is set to zero.
SQL_MAX_QUALIFIER_NAME_LEN (ODBC 1.0)	A 16-bit integer value specifying the maximum length of a qualifier name in the data source. If there is no maximum length or the length is unknown, this value is set to zero.
SQL_MAX_ROW_SIZE (ODBC 2.0)	A 32-bit integer value specifying the maximum length of a single row in a table. If there is no specified limit or the limit is unknown, this value is set to zero.
InfoType	Returns
SQL_MAX_ROW_SIZE_INCLUDES_LONG (ODBC 2.0)	A character string: "Y" if the maximum row size returned for the SQL_MAX_ROW_SIZE information type includes the length of all SQL_LONGVARCHAR and SQL_LONGVARBINARY columns in the row; "N" otherwise.
SQL_MAX_STATEMENT_LEN (ODBC 2.0)	A 32-bit integer value specifying the maximum length (number of characters, including white space) of an SQL statement. If there is no maximum length or the length is unknown, this value is set to zero.
SQL_MAX_TABLE_NAME_LEN (ODBC 1.0)	A 16-bit integer value specifying the maximum length of a table name in the data source. If there is no maximum length or the length is unknown, this value is set to zero.

ADABAS D		286
SQL_MAX_TABLES_IN_SELECT (ODBC 2.0)	A 16-bit integer value specifying the maximum number of tables allowed in the FROM clause of a SELECT statement. If there is no specified limit or the limit is unknown, this value is set to zero.	
SQL_MAX_USER_NAME_LEN (ODBC 2.0)	A 16-bit integer value specifying the maximum length of a user name in the data source. If there is no maximum length or the length is unknown, this value is set to zero.	
SQL_MULT_RESULT_SETS (ODBC 1.0)	A character string: "Y" if the data source supports multiple result sets, "N" if it does not.	
SQL_MULTIPLE_ACTIVE_TXN (ODBC 1.0)	A character string: "Y" if active transactions on multiple connections are allowed, "N" if only one connection at a time can have an active transaction.	
SQL_NEED_LONG_DATA_LEN (ODBC 2.0)	A character string: "Y" if the data source needs the length of a long data value (the data type is SQL_LONGVARCHAR, SQL_LONGVARIABLE, or a long, data source-specific data type) before that value is sent to the data source, "N" if it does not. For more information, see SQLBindParameter and SQLSetPos .	
SQL_NON_NULLABLE_COLUMNS (ODBC 1.0)	<p>A 16-bit integer specifying whether the data source supports non-nullable columns:</p> <p>SQL_NNC_NULL = All columns must be nullable.</p> <p>SQL_NNC_NON_NULL = Columns may be non-nullable (the data source supports the NOT NULL column constraint in CREATE TABLE statements).</p>	
SQL_NULL_COLLATION (ODBC 2.0)	<p>A 16-bit integer value specifying where NULLs are sorted in a list:</p> <p>SQL_NC_END = NULLs are sorted at the end of the list, regardless of the sort order.</p> <p>SQL_NC_HIGH = NULLs are sorted at the high end of the list.</p> <p>SQL_NC_LOW = NULLs are sorted at the low end of the list.</p> <p>SQL_NC_START = NULLs are sorted at the start of the list, regardless of the sort order.</p>	
InfoType	Returns	
SQL_NUMERIC_FUNCTIONS (ODBC 1.0)	A 32-bit bitmask enumerating the scalar numeric functions supported by the driver and associated data source.	
The information type was introduced in ODBC 1.0; each bitmask is labeled with the version in which it was introduced.	<p>The following bitmasks are used to determine which numeric functions are supported:</p> <p>SQL_FN_NUM_ABS (ODBC 1.0)</p> <p>SQL_FN_NUM_ACOS (ODBC 1.0)</p> <p>SQL_FN_NUM_ASIN (ODBC 1.0)</p> <p>SQL_FN_NUM_ATAN (ODBC 1.0)</p> <p>SQL_FN_NUM_ATAN2 (ODBC 1.0)</p> <p>SQL_FN_NUM_CEILING (ODBC 1.0)</p>	

SQLPutData

	SQL_FN_NUM_COS (ODBC 1.0)
	SQL_FN_NUM_COT (ODBC 1.0)
	SQL_FN_NUM_DEGREES (ODBC 2.0)
	SQL_FN_NUM_EXP (ODBC 1.0)
	SQL_FN_NUM_FLOOR (ODBC 1.0)
	SQL_FN_NUM_LOG (ODBC 1.0)
	SQL_FN_NUM_LOG10 (ODBC 2.0)
	SQL_FN_NUM_MOD (ODBC 1.0)
	SQL_FN_NUM_PI (ODBC 1.0)
	SQL_FN_NUM_POWER (ODBC 2.0)
	SQL_FN_NUM_RADIANS (ODBC 2.0)
	SQL_FN_NUM_RAND (ODBC 1.0)
	SQL_FN_NUM_ROUND (ODBC 2.0)
	SQL_FN_NUM_SIGN (ODBC 1.0)
	SQL_FN_NUM_SIN (ODBC 1.0)
	SQL_FN_NUM_SQRT (ODBC 1.0)
	SQL_FN_NUM_TAN (ODBC 1.0)
	SQL_FN_NUM_TRUNCATE (ODBC 2.0)
SQL_ODBC_API_CONFORMANCE (ODBC 1.0)	<p>A 16-bit integer value indicating the level of ODBC conformance:</p> <p>SQL_OAC_NONE = None</p> <p>SQL_OAC_LEVEL1 = Level 1 supported</p> <p>SQL_OAC_LEVEL2 = Level 2 supported</p> <p>(For a list of functions and conformance levels, see Chapter 21, "Function Summary.")</p>
SQL_ODBC_SAG_CLI_CONFORMANCE (ODBC 1.0)	<p>A 16-bit integer value indicating compliance to the functions of the SAG specification:</p> <p>SQL_OSCC_NOT_COMPLIANT = Not SAG-compliant; one or more core functions are not supported</p> <p>SQL_OSCC_COMPLIANT = SAG-compliant</p>
InfoType	Returns
SQL_ODBC_SQL_CONFORMANCE (ODBC 1.0)	<p>A 16-bit integer value indicating SQL grammar supported by the driver:</p> <p>SQL_OSC_MINIMUM = Minimum grammar supported</p> <p>SQL_OSC_CORE = Core grammar supported</p> <p>SQL_OSC_EXTENDED = Extended grammar supported</p>
SQL_ODBC_SQL_OPT_IEF (ODBC 1.0)	A character string: "Y" if the data source supports the optional Integrity Enhancement Facility; "N" if it does not.
SQL_ODBC_VER (ODBC 1.0)	A character string with the version of ODBC to which the Driver Manager conforms. The version is of the form ##.##, where the first two digits are the major version and the next two digits are the minor version. This is implemented solely in the Driver Manager.

SQL_ORDER_BY_COLUMNS_IN_SELECT
(ODBC 2.0)

A character string: "Y" if the columns in the **ORDER BY** clause must be in the select list; otherwise, "N".

SQL_OUTER_JOINS
(ODBC 1.0)

A character string:

The information type was introduced in ODBC 1.0; each return value is labeled with the version in which it was introduced.

"N" = No. The data source does not support outer joins. (ODBC 1.0)

"Y" = Yes. The data source supports two-table outer joins, and the driver supports the ODBC outer join syntax except for nested outer joins. However, columns on the left side of the comparison operator in the ON clause must come from the left-hand table in the outer join, and columns on the right side of the comparison operator must come from the right-hand table. (ODBC 1.0)

"P" = Partial. The data source partially supports nested outer joins, and the driver supports the ODBC outer join syntax. However, columns on the left side of the comparison operator in the ON clause must come from the left-hand table in the outer join and columns on the right side of the comparison operator must come from the right-hand table. Also, the right-hand table of an outer join cannot be included in an inner join. (ODBC 2.0)

"F" = Full. The data source fully supports nested outer joins, and the driver supports the ODBC outer join syntax. (ODBC 2.0)

SQL_OWNER_TERM
(ODBC 1.0)

A character string with the data source vendor's name for an owner; for example, "owner", "Authorization ID", or "Schema".

InfoType

Returns

SQL_OWNER_USAGE
(ODBC 2.0)

A 32-bit bitmask enumerating the statements in which owners can be used:

SQL_OU_DML_STATEMENTS = Owners are supported in all Data Manipulation Language statements: **SELECT**, **INSERT**, **UPDATE**, **DELETE**, and, if supported, **SELECT FOR UPDATE** and positioned update and delete statements.

SQL_OU_PROCEDURE_INVOCATION = Owners are supported in the ODBC procedure invocation statement.

SQL_OU_TABLE_DEFINITION = Owners are supported in all table definition statements: **CREATE TABLE**, **CREATE VIEW**, **ALTER TABLE**, **DROP TABLE**, and **DROP VIEW**.

SQL_OU_INDEX_DEFINITION = Owners are supported in all index definition statements: **CREATE INDEX** and **DROP INDEX**.

SQL_OU_PRIVILEGE_DEFINITION = Owners are supported in all privilege definition statements: **GRANT** and **REVOKE**.

SQL_POS_OPERATIONS
(ODBC 2.0)

A 32-bit bitmask enumerating the supported operations in **SQLSetPos**.

SQLPutData

	<p>The following bitmasks are used to in conjunction with the flag to determine which options are supported:</p> <p>SQL_POS_POSITION SQL_POS_REFRESH SQL_POS_UPDATE SQL_POS_DELETE SQL_POS_ADD</p>
SQL_POSITIONED_STATEMENTS (ODBC 2.0)	<p>A 32-bit bitmask enumerating the supported positioned SQL statements.</p> <p>The following bitmasks are used to determine which statements are supported:</p> <p>SQL_PS_POSITIONED_DELETE SQL_PS_POSITIONED_UPDATE SQL_PS_SELECT_FOR_UPDATE</p>
SQL_PROCEDURE_TERM (ODBC 1.0)	<p>A character string with the data source vendor's name for a procedure; for example, "database procedure", "stored procedure", or "procedure".</p>
SQL_PROCEDURES (ODBC 1.0)	<p>A character string: "Y" if the data source supports procedures and the driver supports the ODBC procedure invocation syntax; "N" otherwise.</p>
InfoType	Returns
SQL_QUALIFIER_LOCATION (ODBC 2.0)	<p>A 16-bit integer value indicating the position of the qualifier in a qualified table name:</p> <p>SQL_QL_START SQL_QL_END</p> <p>For example, an Xbase driver returns SQL_QL_START because the directory (qualifier) name is at the start of the table name, as in \EMPDATA\EMP.DBF. An ORACLE Server driver returns SQL_QL_END, because the qualifier is at the end of the table name, as in ADMIN.EMP@EMPDATA.</p>
SQL_QUALIFIER_NAME_SEPARATOR (ODBC 1.0)	<p>A character string: the character or characters that the data source defines as the separator between a qualifier name and the qualified name element that follows it.</p>
SQL_QUALIFIER_TERM (ODBC 1.0)	<p>A character string with the data source vendor's name for a qualifier; for example, "database" or "directory".</p>
SQL_QUALIFIER_USAGE (ODBC 2.0)	<p>A 32-bit bitmask enumerating the statements in which qualifiers can be used.</p> <p>The following bitmasks are used to determine where qualifiers can be used:</p> <p>SQL_QU_DML_STATEMENTS = Qualifiers are supported in all Data Manipulation Language statements: SELECT, INSERT, UPDATE, DELETE, and, if supported, SELECT FOR UPDATE and positioned update and delete statements.</p>

SQL_QUOTED_IDENTIFIER_CASE (ODBC 2.0)	<p>SQL_QU_PROCEDURE_INVOCATION = Qualifiers are supported in the ODBC procedure invocation statement.</p> <p>SQL_QU_TABLE_DEFINITION = Qualifiers are supported in all table definition statements: CREATE TABLE, CREATE VIEW, ALTER TABLE, DROP TABLE, and DROP VIEW.</p> <p>SQL_QU_INDEX_DEFINITION = Qualifiers are supported in all index definition statements: CREATE INDEX and DROP INDEX.</p> <p>SQL_QU_PRIVILEGE_DEFINITION = Qualifiers are supported in all privilege definition statements: GRANT and REVOKE.</p> <p>A 16-bit integer value as follows:</p> <p>SQL_IC_UPPER = Quoted identifiers in SQL are case insensitive and are stored in upper case in system catalog.</p> <p>SQL_IC_LOWER = Quoted identifiers in SQL are case insensitive and are stored in lower case in system catalog.</p> <p>SQL_IC_SENSITIVE = Quoted identifiers in SQL are case sensitive and are stored in mixed case in system catalog.</p> <p>SQL_IC_MIXED = Quoted identifiers in SQL are case insensitive and are stored in mixed case in system catalog.</p>
InfoType	Returns
SQL_ROW_UPDATES (ODBC 1.0)	A character string: "Y" if a keyset-driven or mixed cursor maintains row versions or values for all fetched rows and therefore can detect any changes made to a row by any user since the row was last fetched; otherwise, "N".
SQL_SCROLL_CONCURRENCY (ODBC 1.0)	<p>A 32-bit bitmask enumerating the concurrency control options supported for scrollable cursors.</p> <p>The following bitmasks are used to determine which options are supported:</p> <p>SQL_SCCO_READ_ONLY = Cursor is read only. No updates are allowed.</p> <p>SQL_SCCO_LOCK = Cursor uses the lowest level of locking sufficient to ensure that the row can be updated.</p> <p>SQL_SCCO_OPT_ROWVER = Cursor uses optimistic concurrency control, comparing row versions, such as SQLBase® ROWID or Sybase TIMESTAMP.</p> <p>SQL_SCCO_OPT_VALUES = Cursor uses optimistic concurrency control, comparing values.</p> <p>For information about cursor concurrency, see "Specifying Cursor Concurrency" in Chapter 7, "Retrieving Results."</p>
SQL_SCROLL_OPTIONS (ODBC 1.0)	A 32-bit bitmask enumerating the scroll options supported for scrollable cursors.

SQLPutData

The information type was introduced in ODBC 1.0; each bitmask is labeled with the version in which it was introduced.

The following bitmasks are used to determine which options are supported:

SQL_SO_FORWARD_ONLY = The cursor only scrolls forward. (ODBC 1.0)

SQL_SO_STATIC = The data in the result set is static. (ODBC 2.0)

SQL_SO_KEYSET_DRIVEN = The driver saves and uses the keys for every row in the result set. (ODBC 1.0)

SQL_SO_DYNAMIC = The driver keeps the keys for every row in the rowset (the keyset size is the same as the rowset size). (ODBC 1.0)

SQL_SO_MIXED = The driver keeps the keys for every row in the keyset, and the keyset size is greater than the rowset size. The cursor is keyset-driven inside the keyset and dynamic outside the keyset. (ODBC 1.0)

For information about scrollable cursors, see "Scrollable Cursors" in Chapter 7, "Retrieving Results."

InfoType

Returns

SQL_SEARCH_PATTERN_ESCAPE
(ODBC 1.0)

A character string specifying what the driver supports as an escape character that permits the use of the pattern match metacharacters underscore (`_`) and percent (`%`) as valid characters in search patterns. This escape character applies only for those catalog function arguments that support search strings. If this string is empty, the driver does not support a search-pattern escape character.

This *InfoType* is limited to catalog functions. For a description of the use of the escape character in search pattern strings, see "Search Pattern Arguments" earlier in this chapter.

SQL_SERVER_NAME
(ODBC 1.0)

A character string with the actual data source-specific server name; useful when a data source name is used during **SQLConnect**, **SQLDriverConnect**, and **SQLBrowseConnect**.

SQL_SPECIAL_CHARACTERS
(ODBC 2.0)

A character string containing all special characters (that is, all characters except a through z, A through Z, 0 through 9, and underscore) that can be used in an object name, such as a table, column, or index name, on the data source. For example, `"#$^"`.

SQL_STATIC_SENSITIVITY
(ODBC 2.0)

A 32-bit bitmask enumerating whether changes made by an application to a static or keyset-driven cursor through **SQLSetPos** or positioned update or delete statements can be detected by that application:

SQL_SS_ADDITIONS = Added rows are visible to the cursor; the cursor can scroll to these rows. Where these rows are added to the cursor is driver-dependent.

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.

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. Because updating key values in a keyset-driven cursor is considered to be deleting the existing row and adding a new row, this value is always returned for keyset-driven cursors.

Whether an application can detect changes made to the result set by other users, including other cursors in the same application, depends on the cursor type. For more information, see "Scrollable Cursors" in Chapter 7, "Retrieving Results."

InfoType

Returns

SQL_STRING_FUNCTIONS
(ODBC 1.0)

The information type was introduced in ODBC 1.0; each bitmask is labeled with the version in which it was introduced.

A 32-bit bitmask enumerating the scalar string functions supported by the driver and associated data source.

The following bitmasks are used to determine which string functions are supported:

SQL_FN_STR_ASCII	(ODBC 1.0)
SQL_FN_STR_CHAR	(ODBC 1.0)
SQL_FN_STR_CONCAT	(ODBC 1.0)
SQL_FN_STR_DIFFERENCE	(ODBC 2.0)
SQL_FN_STR_INSERT	(ODBC 1.0)
SQL_FN_STR_LCASE	(ODBC 1.0)
SQL_FN_STR_LEFT	(ODBC 1.0)
SQL_FN_STR_LENGTH	(ODBC 1.0)
SQL_FN_STR_LOCATE	(ODBC 1.0)
SQL_FN_STR_LOCATE_2	(ODBC 2.0)
SQL_FN_STR_LTRIM	(ODBC 1.0)
SQL_FN_STR_REPEAT	(ODBC 1.0)
SQL_FN_STR_REPLACE	(ODBC 1.0)
SQL_FN_STR_RIGHT	(ODBC 1.0)
SQL_FN_STR_RTRIM	(ODBC 1.0)
SQL_FN_STR_SOUNDEX	(ODBC 2.0)
SQL_FN_STR_SPACE	(ODBC 2.0)
SQL_FN_STR_SUBSTRING	(ODBC 1.0)
SQL_FN_STR_UCASE	(ODBC 1.0)

If an application can call the LOCATE scalar function with the *string_exp1*, *string_exp2*, and *start* arguments, the driver returns the SQL_FN_STR_LOCATE bitmask. If an application can call the LOCATE scalar function with only the *string_exp1* and *string_exp2* arguments, the driver returns the SQL_FN_STR_LOCATE_2 bitmask. Drivers that fully support the LOCATE scalar function

SQLPutData

SQLPutData

SQL_SUBQUERIES (ODBC 2.0)	<p>return both bitmasks.</p> <p>A 32-bit bitmask enumerating the predicates that support subqueries:</p> <p>SQL_SQ_CORRELATED_SUBQUERIES SQL_SQ_COMPARISON SQL_SQ_EXISTS SQL_SQ_IN SQL_SQ_QUANTIFIED</p> <p>The SQL_SQ_CORRELATED_SUBQUERIES bitmask indicates that all predicates that support subqueries support correlated subqueries.</p>
InfoType	Returns
SQL_SYSTEM_FUNCTIONS (ODBC 1.0)	<p>A 32-bit bitmask enumerating the scalar system functions supported by the driver and associated data source.</p> <p>The following bitmasks are used to determine which system functions are supported:</p> <p>SQL_FN_SYS_DBNAME SQL_FN_SYS_IFNULL SQL_FN_SYS_USERNAME</p>
SQL_TABLE_TERM (ODBC 1.0)	<p>A character string with the data source vendor's name for a table; for example, "table" or "file".</p>
SQL_TIMESTAMP_ADD_INTERVALS (ODBC 2.0)	<p>A 32-bit bitmask enumerating the timestamp intervals supported by the driver and associated data source for the TIMESTAMPADD scalar function.</p> <p>The following bitmasks are used to determine which intervals are supported:</p> <p>SQL_FN_TSI_FRAC_SECOND SQL_FN_TSI_SECOND SQL_FN_TSI_MINUTE SQL_FN_TSI_HOUR SQL_FN_TSI_DAY SQL_FN_TSI_WEEK SQL_FN_TSI_MONTH SQL_FN_TSI_QUARTER SQL_FN_TSI_YEAR</p>
SQL_TIMESTAMP_DIFF_INTERVALS (ODBC 2.0)	<p>A 32-bit bitmask enumerating the timestamp intervals supported by the driver and associated data source for the TIMESTAMPDIF scalar function.</p> <p>The following bitmasks are used to determine which intervals are supported:</p> <p>SQL_FN_TSI_FRAC_SECOND SQL_FN_TSI_SECOND</p>

SQL_FN_TSI_MINUTE
 SQL_FN_TSI_HOUR
 SQL_FN_TSI_DAY
 SQL_FN_TSI_WEEK
 SQL_FN_TSI_MONTH
 SQL_FN_TSI_QUARTER
 SQL_FN_TSI_YEAR

InfoType**Returns**

SQL_TIMEDATE_FUNCTIONS
 (ODBC 1.0)

The information type was introduced in ODBC 1.0; each bitmask is labeled with the version in which it was introduced.

A 32-bit bitmask enumerating the scalar date and time functions supported by the driver and associated data source.

The following bitmasks are used to determine which date and time functions are supported:

SQL_FN_TD_CURDATE	(ODBC 1.0)
SQL_FN_TD_CURTIME	(ODBC 1.0)
SQL_FN_TD_DAYNAME	(ODBC 2.0)
SQL_FN_TD_DAYOFMONTH	(ODBC 1.0)
SQL_FN_TD_DAYOFWEEK	(ODBC 1.0)
SQL_FN_TD_DAYOFYEAR	(ODBC 1.0)
SQL_FN_TD_HOUR	(ODBC 1.0)
SQL_FN_TD_MINUTE	(ODBC 1.0)
SQL_FN_TD_MONTH	(ODBC 1.0)
SQL_FN_TD_MONTHNAME	(ODBC 2.0)
SQL_FN_TD_NOW	(ODBC 1.0)
SQL_FN_TD_QUARTER	(ODBC 1.0)
SQL_FN_TD_SECOND	(ODBC 1.0)
SQL_FN_TD_TIMESTAMPADD	(ODBC 2.0)
SQL_FN_TD_TIMESTAMPDIFF	(ODBC 2.0)
SQL_FN_TD_WEEK	(ODBC 1.0)
SQL_FN_TD_YEAR	(ODBC 1.0)

SQL_TXN_CAPABLE
 (ODBC 1.0)

The information type was introduced in ODBC 1.0; each return value is labeled with the version in which it was introduced

A 16-bit integer value describing the transaction support in the driver or data source:

SQL_TC_NONE = Transactions not supported. (ODBC 1.0)

SQL_TC_DML = Transactions can only contain Data Manipulation Language (DML) statements (**SELECT**, **INSERT**, **UPDATE**, **DELETE**). Data Definition Language (DDL) statements encountered in a transaction cause an error. (ODBC 1.0)

SQL_TC_DDL_COMMIT = Transactions can only contain DML statements. DDL statements (**CREATE TABLE**, **DROP INDEX**, an so on) encountered in a transaction cause the transaction to be committed. (ODBC 2.0)

SQL_TC_DDL_IGNORE = Transactions can only contain DML statements. DDL statements encountered in a transaction are ignored. (ODBC 2.0)

SQLPutData

SQL_TC_ALL = Transactions can contain DDL statements and DML statements in any order. (ODBC 1.0)

InfoType	Returns
SQL_TXN_ISOLATION_OPTION (ODBC 1.0)	A 32-bit bitmask enumerating the transaction isolation levels available from the driver or data source. The following bitmasks are used in conjunction with the flag to determine which options are supported: SQL_TXN_READ_UNCOMMITTED SQL_TXN_READ_COMMITTED SQL_TXN_REPEATABLE_READ SQL_TXN_SERIALIZABLE SQL_TXN_VERSIONING For descriptions of these isolation levels, see the description of SQL_DEFAULT_TXN_ISOLATION.
SQL_UNION (ODBC 2.0)	A 32-bit bitmask enumerating the support for the UNION clause: SQL_U_UNION = The data source supports the UNION clause. SQL_U_UNION_ALL = The data source supports the ALL keyword in the UNION clause. (SQLGetInfo returns both SQL_U_UNION and SQL_U_UNION_ALL in this case.)
SQL_USER_NAME (ODBC 1.0)	A character string with the name used in a particular database, which can be different than login name.

Code Example **SQLGetInfo** returns lists of supported options as a 32-bit bitmask in *rgbInfoValue*. The bitmask for each option is used in conjunction with the flag to determine whether the option is supported.

For example, an application could use the following code to determine whether the SUBSTRING scalar function is supported by the driver associated with the *hdbc*:

```
UDWORD    fFuncs;

SQLGetInfo(hdbc,
           SQL_STRING_FUNCTIONS,
           (PTR)&fFuncs,
           sizeof(fFuncs),
           NULL);

if (fFuncs & SQL_FN_STR_SUBSTRING) /* SUBSTRING supported */
    ...;
else                               /* SUBSTRING not supported */
    ...;
```

Related	For information about	See
----------------	------------------------------	------------

Returning the setting of a connection option	SQLGetConnectOption (extension)
Determining if a driver supports a function	SQLGetFunctions (extension)
Returning the setting of a statement option	SQLGetStmtOption (extension)
Returning information about a data source's data types	SQLGetTypeInfo (extension)

SQLGetStmtOption



Extension Level 1 **SQLGetStmtOption** returns the current setting of a statement option.

Syntax RETCODE **SQLGetStmtOption**(*hstmt*, *fOption*, *pvParam*)

The **SQLGetStmtOption** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UWORD	<i>fOption</i>	Input	Option to retrieve.
PTR	<i>pvParam</i>	Output	Value associated with <i>fOption</i> . Depending on the value of <i>fOption</i> , a 32-bit integer value or a pointer to a null-terminated character string will be returned in <i>pvParam</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLGetStmtOption** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLGetStmtOption** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
24000	Invalid cursor state	The argument <i>fOption</i> was SQL_ROW_NUMBER or SQL_GET_BOOKMARK and the cursor was not open, or the cursor was positioned before the start of the result set or after the end of the result set.
IM001	Driver does not support this function	(DM) The driver corresponding to the <i>hstmt</i> does not support the function.
SQLSTATE	Error	Description
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no

		implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1010	Function sequence error	(DM) An asynchronously executing function was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1011	Operation invalid at this time	The <i>fOption</i> argument was SQL_GET_BOOKMARK and the value of the SQL_USE_BOOKMARKS statement option was SQL_UB_OFF.
S1092	Option type out of range	(DM) The value specified for the argument <i>fOption</i> was in the block of numbers reserved for ODBC connection and statement options, but was not valid for the version of ODBC supported by the driver.
S1109	Invalid cursor position	The <i>fOption</i> argument was SQL_GET_BOOKMARK or SQL_ROW_NUMBER and the value in the <i>rgfRowStatus</i> array in SQLExtendedFetch for the current row was SQL_ROW_DELETED or SQL_ROW_ERROR.
S1C00	Driver not capable	The value specified for the argument <i>fOption</i> was a valid ODBC statement option for the version of ODBC supported by the driver, but was not supported by the driver. The value specified for the argument <i>fOption</i> was in the block of numbers reserved for driver-specific connection and statement options, but was not supported by the driver.

Comments

The following table lists statement options for which corresponding values can be returned, but not set. The table also lists the version of ODBC in which they were introduced. For a list of options that can be set and retrieved, see **SQLSetStmtOption**. If *fOption* specifies an option that returns a string, *pvParam* must be a pointer to storage for the string. The maximum length

SQLRowCount

of the string will be SQL_MAX_OPTION_STRING_LENGTH bytes (excluding the null termination byte).

<i>fOption</i>	<i>pvParam</i> contents
SQL_GET_BOOKMARK (ODBC 2.0)	<p>A 32-bit integer value that is the bookmark for the current row. Before using this option, an application must set the SQL_USE_BOOKMARKS statement option to SQL_UB_ON, create a result set, and call SQLExtendedFetch.</p> <p>To return to the rowset starting with the row marked by this bookmark, an application calls SQLExtendedFetch with the SQL_FETCH_BOOKMARK fetch type and <i>irow</i> set to this value.</p> <p>Bookmarks are also returned as column 0 of the result set.</p>
SQL_ROW_NUMBER (ODBC 2.0)	<p>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, the driver returns 0.</p>

Related Functions

For information about	See
Returning the setting of a connection option	SQLGetConnectOption (extension)
Setting a connection option	SQLSetConnectOption (extension)
Setting a statement option	SQLSetStmtOption (extension)

SQLGetTypeInfo



Extension Level 1 **SQLGetTypeInfo** returns information about data types supported by the data source. The driver returns the information in the form of an SQL result set.

Important: Applications must use the type names returned in the TYPE_NAME column in **ALTER TABLE** and **CREATE TABLE** statements; they must not use the sample type names listed in Appendix C, "SQL Grammar." **SQLGetTypeInfo** may return more than one row with the same value in the DATA_TYPE column.

Syntax

RETCODE **SQLGetTypeInfo**(*hstmt*, *fSqlType*)

The **SQLGetTypeInfo** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle for the result set.
SWORD	<i>fSqlType</i>	Input	The SQL data type. This must be one of the following values: SQL_BIGINT SQL_BINARY SQL_BIT SQL_CHAR SQL_DATE SQL_DECIMAL SQL_DOUBLE SQL_FLOAT SQL_INTEGER SQL_LONGVARBINARY SQL_LONGVARCHAR SQL_NUMERIC SQL_REAL SQL_SMALLINT SQL_TIME SQL_TIMESTAMP SQL_TINYINT SQL_VARBINARY SQL_VARCHAR or a driver-specific SQL data type. SQL_ALL_TYPES specifies that information about all data types should be returned.

SQLSetConnectOption

Type	Argument	Use	Description
			For information about ODBC SQL data types, see “SQL Data Types” in Appendix D, “Data Types.” For information about driver-specific SQL data types, see the driver’s documentation.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLGetTypeInfo** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLGetTypeInfo** and explains each one in the context of this function; the notation “(DM)” precedes the descriptions of SQLSTATEs returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had not been called. A result set was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called.
IM001	Driver does not support this function	(DM) The driver corresponding to the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
SQLSTATE	Error	Description

S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1004	SQL data type out of range	(DM) The value specified for the argument <i>fSqlType</i> was in the block of numbers reserved for ODBC SQL data type indicators but was not a valid ODBC SQL data type indicator.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> , then the function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1010	Function sequence error	(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1C00	Driver not capable	The value specified for the argument <i>fSqlType</i> was in the range of numbers reserved for driver-specific SQL data type indicators, but was not supported by the driver or data source. The combination of the current settings of the SQL_CONCURRENCY and SQL_CURSOR_TYPE statement options was not supported by the driver or data source.
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

Comments

SQLGetTypeInfo returns the results as a standard result set, ordered by DATA_TYPE and TYPE_NAME. The following table lists the columns in the result set.

Note: **SQLGetTypeInfo** might not return all data types. For example, a driver might not return user-defined data types. Applications can use any valid data type, regardless of whether it is returned by **SQLGetTypeInfo**.

The lengths of VARCHAR columns shown in the table are maximums; the actual lengths depend on the data source.

Column Name	Data Type	Comments
TYPE_NAME	Varchar(128) not NULL	Data source-dependent data type name; for example, "CHAR", "VARCHAR", "MONEY", "LONG VARBINARY", or "CHAR () FOR BIT DATA". Applications must use this name in CREATE TABLE and ALTER TABLE statements.
DATA_TYPE	Smallint not NULL	SQL data type. This can be an ODBC SQL data type or a driver-specific SQL data type. For a list of valid ODBC SQL data types, see "SQL Data Types" in Appendix D, "Data Types." For information about driver-specific SQL data types, see the driver's documentation.
PRECISION	Integer	The maximum precision of the data type on the data source. NULL is returned for data types where precision is not applicable. For more information on precision, see "Precision, Scale, Length, and Display Size" in Appendix D, "Data Types."
LITERAL_PREFIX	Varchar(128)	Character or characters used to prefix a literal; for example, a single quote (') for character data types or 0x for binary data types; NULL is returned for data types where a literal prefix is not applicable.
LITERAL_SUFFIX	Varchar(128)	Character or characters used to terminate a literal; for example, a single quote (') for character data types; NULL is returned for data types where a literal suffix is not applicable.
Column Name	Data Type	Comments
CREATE_PARAMS	Varchar(128)	Parameters for a data type definition. For example, CREATE_PARAMS for DECIMAL would be "precision,scale"; CREATE_PARAMS for VARCHAR

		would equal "max length"; NULL is returned if there are no parameters for the data type definition, for example INTEGER.
		The driver supplies the CREATE_PARAMS text in the language of the country where it is used.
NULLABLE	Smallint not NULL	Whether the data type accepts a NULL value: SQL_NO_NULLS if the data type does not accept NULL values. SQL_NULLABLE if the data type accepts NULL values. SQL_NULLABLE_UNKNOWN if it is not known if the column accepts NULL values.
CASE_SENSITIVE	Smallint not NULL	Whether a character data type is case sensitive in collations and comparisons: TRUE if the data type is a character data type and is case sensitive. FALSE if the data type is not a character data type or is not case sensitive.
SEARCHABLE	Smallint not NULL	How the data type is used in a WHERE clause: SQL_UNSEARCHABLE if the data type cannot be used in a WHERE clause. SQL_LIKE_ONLY if the data type can be used in a WHERE clause only with the LIKE predicate. SQL_ALL_EXCEPT_LIKE if the data type can be used in a WHERE clause with all comparison operators except LIKE . SQL_SEARCHABLE if the data type can be used in a WHERE clause with any comparison operator.
Column Name	Data Type	Comments
UNSIGNED_ATTRIBUTE	Smallint	Whether the data type is unsigned: TRUE if the data type is unsigned. FALSE if the data type is signed. NULL is returned if the attribute is not applicable to the data type or the data type is not numeric.

SQLSetConnectOption

SQLSetConnectOption

MONEY	Smallint not NULL	Whether the data type is a money data type: TRUE if it is a money data type. FALSE if it is not.
AUTO_INCREMENT	Smallint	Whether the data type is autoincrementing: TRUE if the data type is autoincrementing. FALSE if the data type is not autoincrementing. NULL is returned if the attribute is not applicable to the data type or the data type is not numeric. An application can insert values into a column having this attribute, but cannot update the values in the column.
LOCAL_TYPE_NAME	Varchar(128)	Localized version of the data source–dependent name of the data type. NULL is returned if a localized name is not supported by the data source. This name is intended for display only, such as in dialog boxes.
MINIMUM_SCALE	Smallint	The minimum scale of the data type on the data source. If a data type has a fixed scale, the MINIMUM_SCALE and MAXIMUM_SCALE columns both contain this value. For example, an SQL_TIMESTAMP column might have a fixed scale for fractional seconds. NULL is returned where scale is not applicable. For more information, see “Precision, Scale, Length, and Display Size” in Appendix D, “Data Types.”

Column Name	Data Type	Comments
MAXIMUM_SCALE	Smallint	The maximum scale of the data type on the data source. NULL is returned where scale is not applicable. If the maximum scale is not defined separately on the data source, but is instead defined to be the same as the maximum precision, this column contains the same value as the PRECISION column. For more information, see “Precision, Scale, Length, and Display Size” in Appendix D, “Data Types.”

Note: The MINIMUM_SCALE and MAXIMUM_SCALE columns were added in ODBC 2.0. ODBC 1.0 drivers may return different, driver-specific columns with the same column numbers.

Attribute information can apply to data types or to specific columns in a result set. **SQLGetTypeInfo** returns information about attributes associated with data types; **SQLColAttributes** returns information about attributes associated with columns in a result set.

Related Functions

For information about	See
Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Returning information about a column in a result set	SQLColAttributes
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
Fetching a row of data	SQLFetch
Returning information about a driver or data source	SQLGetInfo (extension)

SQLMoreResults



Extension Level 2 **SQLMoreResults** determines whether there are more results available on an *hstmt* containing **SELECT**, **UPDATE**, **INSERT**, or **DELETE** statements and, if so, initializes processing for those results.

Syntax RETCODE **SQLMoreResults**(*hstmt*)

The **SQLMoreResults** function accepts the following argument:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_NO_DATA_FOUND, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLMoreResults** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLMoreResults** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
SQLSTATE	Error	Description
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called

again on the *hstmt*.

The function was called and, before it completed execution, **SQLCancel** was called on the *hstmt* from a different thread in a multithreaded application.

S1010	Function sequence error	(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

Comments

SELECT statements return result sets. **UPDATE**, **INSERT**, and **DELETE** statements return a count of affected rows. If any of these statements are batched, submitted with arrays of parameters, or in procedures, they can return multiple result sets or counts.

If another result set or count is available, **SQLMoreResults** returns SQL_SUCCESS and initializes the result set or count for additional processing. After calling **SQLMoreResults** for **SELECT** statements, an application can call functions to determine the characteristics of the result set and to retrieve data from the result set. After calling **SQLMoreResults** for **UPDATE**, **INSERT**, or **DELETE** statements, an application can call **SQLRowCount**.

If all results have been processed, **SQLMoreResults** returns SQL_NO_DATA_FOUND.

Note that if there is a current result set with unfetched rows, **SQLMoreResults** discards that result set and makes the next result set or count available.

If a batch of statements or a procedure mixes other SQL statements with **SELECT**, **UPDATE**, **INSERT**, and **DELETE** statements, these other statements do not affect **SQLMoreResults**.

For additional information about the valid sequencing of result-processing functions, see Appendix B, "ODBC State Transition Tables."

Related Functions

For information about

See

SQLSetCursorName

SQLSetCursorName

Canceling statement processing

SQLCancel

Fetching a block of data or scrolling through a result set

SQLExtendedFetch (extension)

Fetching a row of data

SQLFetch

Fetching part or all of a column of data

SQLGetData (extension)

SQLNativeSql



Extension Level 2 **SQLNativeSql** returns the SQL string as translated by the driver.

Syntax RETCODE **SQLNativeSql**(*hdbc*, *szSqlStrIn*, *cbSqlStrIn*, *szSqlStr*, *cbSqlStrMax*, *pcbSqlStr*)

The **SQLNativeSql** function accepts the following arguments:

Type	Argument	Use	Description
HDBC	<i>hdbc</i>	Input	Connection handle.
UCHAR FAR *	<i>szSqlStrIn</i>	Input	SQL text string to be translated.
SDWORD	<i>cbSqlStrIn</i>	Input	Length of <i>szSqlStrIn</i> text string.
UCHAR FAR *	<i>szSqlStr</i>	Output	Pointer to storage for the translated SQL string.
SDWORD	<i>cbSqlStrMax</i>	Input	Maximum length of the <i>szSqlStr</i> buffer.
SDWORD FAR *	<i>pcbSqlStr</i>	Output	The total number of bytes (excluding the null termination byte) 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 translated SQL string in <i>szSqlStr</i> is truncated to <i>cbSqlStrMax</i> – 1 bytes.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLNativeSql** returns either SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLNativeSql** and explains each one in the context of this function; the notation “(DM)” precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	The buffer <i>szSqlStr</i> was not large enough to return the entire SQL string, so the SQL string was truncated. The argument <i>pcbSqlStr</i> contains the length of the untruncated SQL string. (Function returns

SQLSetParam

SQLSetParam

		SQL_SUCCESS_WITH_INFO.)
08003	Connection not open	The <i>hdbc</i> was not in a connected state.
37000	Syntax error or access violation	The argument <i>szSqlStrIn</i> contained an SQL statement that was not preparable or contained a syntax error.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hdbc</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1009	Invalid argument value	(DM) The argument <i>szSqlStrIn</i> was a null pointer.
S1090	Invalid string or buffer length	(DM) The argument <i>cbSqlStrIn</i> was less than 0, but not equal to SQL_NTS. (DM) The argument <i>cbSqlStrMax</i> was less than 0 and the argument <i>szSqlStr</i> was not a null pointer..

Related Functions None.

SQLNumParams



Extension Level 2 **SQLNumParams** returns the number of parameters in an SQL statement.

Syntax RETCODE **SQLNumParams**(*hstmt*, *pcpar*)

The **SQLNumParams** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
SWORD FAR *	<i>pcpar</i>	Output	Number of parameters in the statement.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLNumParams** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLNumParams** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
SQLSTATE	Error	Description
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> .

SQLSetPos

SQLSetPos

		The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1010	Function sequence error	(DM) The function was called prior to calling SQLPrepare or SQLExecDirect for the <i>hstmt</i> . (DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

SQLNumParams can only be called after **SQLPrepare** has been called.

Comments

If the statement associated with *hstmt* does not contain parameters, **SQLNumParams** sets *pcpar* to 0.

Related Functions

For information about

See

Returning information about a parameter in a statement

SQLDescribeParam (extension)

Assigning storage for a parameter

SQLBindParameter

SQLNumResultCols



Core **SQLNumResultCols** returns the number of columns in a result set.

Syntax RETCODE **SQLNumResultCols**(*hstmt*, *pccol*)

The **SQLNumResultCols** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
SWORD FAR *	<i>pccol</i>	Output	Number of columns in the result set.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLNumResultCols** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLNumResultCols** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
SQLSTATE	Error	Description
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function SQLSetScrollOptions

SQLSetScrollOptions

		was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1010	Function sequence error	(DM) The function was called prior to calling SQLPrepare or SQLExecDirect for the <i>hstmt</i> . (DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

SQLNumResultCols can return any SQLSTATE that can be returned by **SQLPrepare** or **SQLExecute** when called after **SQLPrepare** and before **SQLExecute** depending on when the data source evaluates the SQL statement associated with the *hstmt*.

Comments

SQLNumResultCols can be called successfully only when the *hstmt* is in the prepared, executed, or positioned state.

If the statement associated with *hstmt* does not return columns, **SQLNumResultCols** sets *pccol* to 0.

Related Functions

For information about	See
Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Returning information about a column in a result set	SQLColAttributes
Returning information about a column in a result set	SQLDescribeCol
Fetching a block of data or scrolling through a	SQLExtendedFetch (extension)

result set

Fetching a row of data

Fetching part or all of a column of data

Setting cursor scrolling options

SQLFetch

SQLGetData (extension)

SQLSetScrollOptions
(extension)

SQLParamData



Extension Level 1 **SQLParamData** is used in conjunction with **SQLPutData** to supply parameter data at statement execution time.

Syntax RETCODE **SQLParamData**(*hstmt*, *prgbValue*)

The **SQLParamData** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
PTR FAR *	<i>prgbValue</i>	Output	Pointer to storage for the value specified for the <i>rgbValue</i> argument in SQLBindParameter (for parameter data) or the address of the <i>rgbValue</i> buffer specified in SQLBindCol (for column data).

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_NEED_DATA, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLParamData** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLParamData** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
SQLSTATE	Error	Description
22026	String data, length mismatch	The SQL_NEED_LONG_DATA_LEN information type in SQLGetInfo was "Y" and less data was sent for a long parameter (the data type was SQL_LONGVARCHAR, SQL_LONGVARIABLE, or a long, data

		<p>source-specific data type) than was specified with the <i>pcbValue</i> argument in SQLBindParameter.</p> <p>The SQL_NEED_LONG_DATA_LEN information type in SQLGetInfo was “Y” and less data was sent for a long column (the data type was SQL_LONGVARCHAR, SQL_LONGVARBINARY, or a long, data source-specific data type) than was specified in the length buffer corresponding to a column in a row of data that was added or updated with SQLSetPos.</p>
IM001	Driver does not support this function	(DM) The driver that corresponds the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	<p>Asynchronous processing was enabled for the <i>hstmt</i>. The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i>. Then the function was called again on the <i>hstmt</i>.</p> <p>The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.</p> <p>SQLExecute, SQLExecDirect, or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. SQLCancel was called before data was sent for all data-at-execution parameters or columns.</p>
SQLSTATE Error		Description
S1010	Function sequence error	<p>(DM) The previous function call was not a call to SQLExecDirect, SQLExecute, or SQLSetPos where the return code was SQL_NEED_DATA or a call to SQLPutData.</p> <p>The previous function call was a call to SQLParamData.</p>

SQLSetStmtOption

SQLSetStmtOption

		(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called.
S1T00	Timeout expired	The timeout period expired before the data source completed processing the parameter value. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT .

If **SQLParamData** is called while sending data for a parameter in an SQL statement, it can return any SQLSTATE that can be returned by the function called to execute the statement (**SQLExecute** or **SQLExecDirect**). If it is called while sending data for a column being updated or added with **SQLSetPos**, it can return any SQLSTATE that can be returned by **SQLSetPos**.

Comments For an explanation of how data-at-execution parameter data is passed at statement execution time, see “Passing Parameter Values” in **SQLBindParameter**. For an explanation of how data-at-execution column data is updated or added, see “Using SQLSetPos” in **SQLSetPos**.

Code Example See **SQLPutData**.

Related Functions	For information about	See
--------------------------	------------------------------	------------

Canceling statement processing	SQLCancel
Returning information about a parameter in a statement	SQLDescribeParam (extension)
Executing an SQL statement	SQLExecDirect
Executing a prepared SQL statement	SQLExecute
Sending parameter data at execution time	SQLPutData (extension)
Assigning storage for a parameter	SQLBindParameter

SQLParamOptions



Extension Level 2 SQLParamOptions allows an application to specify multiple values for the set of parameters assigned by **SQLBindParameter**. The ability to specify multiple values for a set of parameters is useful for bulk inserts and other work that requires the data source to process the same SQL statement multiple times with various parameter values. An application can, for example, specify three sets of values for the set of parameters associated with an **INSERT** statement, and then execute the **INSERT** statement once to perform the three insert operations.

Syntax

RETCODE **SQLParamOptions**(*hstmt*, *crow*, *pirow*)

The **SQLParamOptions** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UDWORD	<i>crow</i>	Input	Number of values for each parameter. If <i>crow</i> is greater than 1, the <i>rgbValue</i> argument in SQLBindParameter points to an array of parameter values and <i>pcbValue</i> points to an array of lengths.
UDWORD FAR *	<i>pirow</i>	Input	Pointer to storage for the current row number. As each row of parameter values is processed, <i>pirow</i> is set to the number of that row. No row number will be returned if <i>pirow</i> is set to a null pointer.

Returns

SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics

When **SQLParamOptions** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLParamOptions** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message.

SQLSpecialColumns

SQLSpecialColumns

		(Function returns SQL_SUCCESS_WITH_INFO.)
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1010	Function sequence error	(DM) An asynchronously executing function was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1107	Row value out of range	(DM) The value specified for the argument <i>crow</i> was equal to 0.

Comments

As a statement executes, the driver sets *pirow* to the number of the current row of parameter values; the first row is row number 1. The contents of *pirow* can be used as follows:

- When **SQLParamData** returns SQL_NEED_DATA for data-at-execution parameters, the application can access the value in *pirow* to determine which row of parameters is being executed.
- When **SQLExecute** or **SQLExecDirect** returns an error, the application can access the value in *pirow* to find out which row of parameters failed.
- When **SQLExecute**, **SQLExecDirect**, **SQLParamData**, or **SQLPutData** succeed, the value in *pirow* is set to *crow*—the total number of rows of parameters processed.

Code Example

In the following example, an application specifies an array of parameter values with **SQLBindParameter** and **SQLParamOptions**. It then inserts those values into a table with a single **INSERT** statement and checks for any errors. If the first row fails, the application rolls back all changes. If any other row fails, the application commits the transaction, skips the

failed row, rebinds the remaining parameters, and continues processing. (Note that **irow** is 1-based and **szData[]** is 0-based, so the **irow** entry of **szData[]** is skipped by rebinding at **szData[irow]**.)

```
#define CITY_LEN 256
SDWORD cbValue[ ] = {SQL_NTS, SQL_NTS, SQL_NTS, SQL_NTS, SQL_NTS};
UCHAR szData[ ][CITY_LEN] = {"Boston","New York","Keokuk","Seattle",
                              "Eugene"};

UDWORD irow;
SQLSetConnectOption(hdbc, SQL_AUTOCOMMIT, 0);
SQLBindParameter(hstmt, 1, SQL_PARAM_INPUT, SQL_C_DEFAULT, SQL_CHAR,
                  CITY_LEN, 0, szData, 0, cbValue);
SQLPrepare(hstmt, "INSERT INTO CITIES VALUES (?)", SQL_NTS);
SQLParamOptions(hstmt, 5, &irow);

while (TRUE) {

    retcode = SQLExecute(hstmt);

    /* Done if execution was successful */

    if (retcode != SQL_ERROR) {
        break;
    }

    /* On an error, print the error. If the error is in row 1, roll */
    /* back the transaction and quit. If the error is in another */
    /* row, commit the transaction and, unless the error is in the */
    /* last row, rebind to the next row and continue processing. */

    show_error();
    if (irow == 1) {
        SQLTransact(henv, hstmt, SQL_ROLLBACK);
        break;
    } else {
        SQLTransact(henv, hstmt, SQL_COMMIT);
        if (irow == 5) {
            break;
        } else {
            SQLBindParameter(hstmt, 1, SQL_PARAM_INPUT,
                              SQL_C_DEFAULT, SQL_CHAR, CITY_LEN, 0,
                              szData[irow], 0, cbValue[irow]);
            SQLParamOptions(hstmt, 5-irow, &irow);
        }
    }
}
```

SQLSpecialColumns

Related Functions	For information about	See
	Returning information about a parameter in a statement	SQLDescribeParam (extension)
	Assigning storage for a parameter	SQLBindParameter

SQLPrepare



Core **SQLPrepare** prepares an SQL string for execution.

Syntax RETCODE **SQLPrepare**(*hstmt*, *szSqlStr*, *cbSqlStr*)

The **SQLPrepare** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UCHAR FAR *	<i>szSqlStr</i>	Input	SQL text string.
SDWORD	<i>cbSqlStr</i>	Input	Length of <i>szSqlStr</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLPrepare** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLPrepare** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATEs returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
21S01	Insert value list does not match column list	The argument <i>szSqlStr</i> contained an INSERT statement and the number of values to be inserted did not match the degree of the derived table.
21S02	Degree of derived table does not match column list	The argument <i>szSqlStr</i> contained a CREATE VIEW statement and the number of names specified is not the same degree as the derived table defined by the query specification.
SQLSTATE	Error	Description

SQLStatistics

22005	Error in assignment	The argument <i>szSqlStr</i> contained an SQL statement that contained a literal or parameter and the value was incompatible with the data type of the associated table column.
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called. A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called.
34000	Invalid cursor name	The argument <i>szSqlStr</i> contained a positioned DELETE or a positioned UPDATE and the cursor referenced by the statement being prepared was not open.
37000	Syntax error or access violation	The argument <i>szSqlStr</i> contained an SQL statement that was not preparable or contained a syntax error.
42000	Syntax error or access violation	The argument <i>szSqlStr</i> contained a statement for which the user did not have the required privileges.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S0001	Base table or view already exists	The argument <i>szSqlStr</i> contained a CREATE TABLE or CREATE VIEW statement and the table name or view name specified already exists.
S0002	Base table not found	The argument <i>szSqlStr</i> contained a DROP TABLE or a DROP VIEW statement and the specified table name or view name did not exist. The argument <i>szSqlStr</i> contained an ALTER TABLE statement and the specified table name did not exist. The argument <i>szSqlStr</i> contained a CREATE VIEW statement and a table name or view name defined by the query specification did not exist. The argument <i>szSqlStr</i> contained a CREATE INDEX statement and the specified table name did not exist.

SQLSTATE Error

Description

		<p>The argument <i>szSqlStr</i> contained a GRANT or REVOKE statement and the specified table name or view name did not exist.</p> <p>The argument <i>szSqlStr</i> contained a SELECT statement and a specified table name or view name did not exist.</p> <p>The argument <i>szSqlStr</i> contained a DELETE, INSERT, or UPDATE statement and the specified table name did not exist.</p> <p>The argument <i>szSqlStr</i> contained a CREATE TABLE statement and a table specified in a constraint (referencing a table other than the one being created) did not exist.</p>
S0011	Index already exists	The argument <i>szSqlStr</i> contained a CREATE INDEX statement and the specified index name already existed.
S0012	Index not found	The argument <i>szSqlStr</i> contained a DROP INDEX statement and the specified index name did not exist.
S0021	Column already exists	The argument <i>szSqlStr</i> contained 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	<p>The argument <i>szSqlStr</i> contained a CREATE INDEX statement and one or more of the column names specified in the column list did not exist.</p> <p>The argument <i>szSqlStr</i> contained a GRANT or REVOKE statement and a specified column name did not exist.</p> <p>The argument <i>szSqlStr</i> contained a SELECT, DELETE, INSERT, or UPDATE statement and a specified column name did not exist.</p> <p>The argument <i>szSqlStr</i> contained a CREATE TABLE statement and a column specified in a constraint (referencing a table other than the one being created) did not exist.</p>

SQLSTATE Error**Description**

S1000	General error	An error occurred for which there was no
-------	---------------	--

SQLStatistics

SQLStatistics

		specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1009	Invalid argument value	(DM) The argument <i>szSqlStr</i> was a null pointer.
S1010	Function sequence error	(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1090	Invalid string or buffer length	(DM) The argument <i>cbSqlStr</i> was less than or equal to 0, but not equal to SQL_NTS.
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

Comments

The application calls **SQLPrepare** to send an SQL statement to the data source for preparation. The application can include one or more parameter markers in the SQL statement. To include a parameter marker, the application embeds a question mark (?) into the SQL string at the appropriate position.

Note: If an application uses **SQLPrepare** to prepare and **SQLExecute** to submit a **COMMIT** or **ROLLBACK** statement, it will not be interoperable between DBMS products. To commit or roll back a transaction, call **SQLTransact**.

The driver modifies the statement to use the form of SQL used by the data source, then submits it to the data source for preparation. In particular, the driver modifies the escape clauses used to define ODBC-specific SQL. (For a description of SQL statement grammar, see Appendix C, “SQL Grammar.”) For the driver, an *hstmt* is similar to a statement identifier in embedded SQL code. If the data source supports statement identifiers, the driver can send a statement identifier and parameter values to the data source.

Once a statement is prepared, the application uses *hstmt* to refer to the statement in later function calls. The prepared statement associated with the *hstmt* may be reexecuted by calling **SQLExecute** until the application frees the *hstmt* with a call to **SQLFreeStmt** with the **SQL_DROP** option or until the *hstmt* is used in a call to **SQLPrepare**, **SQLExecDirect**, or one of the catalog functions (**SQLColumns**, **SQLTables**, and so on). Once the application prepares a statement, it can request information about the format of the result set.

Some drivers cannot return syntax errors or access violations when the application calls **SQLPrepare**. A driver may handle syntax errors and access violations, only syntax errors, or neither syntax errors nor access violations. Therefore, an application must be able to handle these conditions when calling subsequent related functions such as **SQLNumResultCols**, **SQLDescribeCol**, **SQLColAttributes**, and **SQLExecute**.

Depending on the capabilities of the driver and data source and on whether the application has called **SQLBindParameter**, parameter information (such as data types) might be checked when the statement is prepared or when it is executed. For maximum interoperability, an application should unbind all parameters that applied to an old SQL statement before preparing a new SQL statement on the same *hstmt*. This prevents errors that are due to old parameter information being applied to the new statement.

Important: Committing or rolling back a transaction, either by calling **SQLTransact** or by using the SQL_AUTOCOMMIT connection option, can cause the data source to delete the access plans for all *hstmts* on an *hdbc*. For more information, see the SQL_CURSOR_COMMIT_BEHAVIOR and SQL_CURSOR_ROLLBACK_BEHAVIOR information types in **SQLGetInfo**.

Code Example See **SQLBindParameter**, **SQLParamOptions**, **SQLPutData**, and **SQLSetPos**.

Related Functions	For information about	See
	Allocating a statement handle	SQLAllocStmt
	Assigning storage for a column in a result set	SQLBindCol
	Canceling statement processing	SQLCancel
	Executing an SQL statement	SQLExecDirect
	Executing a prepared SQL statement	SQLExecute
	Returning the number of rows affected by a statement	SQLRowCount
	Setting a cursor name	SQLSetCursorName
	Assigning storage for a parameter	SQLBindParameter

SQLPrimaryKeys

ADBC

Extension Level 2 **SQLPrimaryKeys** returns the column names that comprise the primary key for a table. The driver returns the information as a result set. This function does not support returning primary keys from multiple tables in a single call.

Syntax RETCODE **SQLPrimaryKeys**(*hstmt*, *szTableQualifier*, *cbTableQualifier*, *szTableOwner*, *cbTableOwner*, *szTableName*, *cbTableName*)

The **SQLPrimaryKeys** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UCHAR FAR *	<i>szTableQualifier</i>	Input	Qualifier name. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have qualifiers.
SWORD	<i>cbTableQualifier</i>	Input	Length of <i>szTableQualifier</i> .
UCHAR FAR *	<i>szTableOwner</i>	Input	Table owner. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have owners.
SWORD	<i>cbTableOwner</i>	Input	Length of <i>szTableOwner</i> .
UCHAR FAR *	<i>szTableName</i>	Input	Table name.
SWORD	<i>cbTableName</i>	Input	Length of <i>szTableName</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLPrimaryKeys** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLPrimaryKeys** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called. A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1010	Function sequence error	(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
SQLSTATE	Error	Description

SQLTablePrivileges

S1090	Invalid string or buffer length	(DM) The value of one of the name length arguments was less than 0, but not equal to SQL_NTS. The value of one of the name length arguments exceeded the maximum length value for the corresponding qualifier or name.
S1C00	Driver not capable	A table qualifier was specified and the driver or data source does not support qualifiers. A table owner was specified and the driver or data source does not support owners. The combination of the current settings of the SQL_CONCURRENCY and SQL_CURSOR_TYPE statement options was not supported by the driver or data source.
S1T00	Timeout expired	The timeout period expired before the data source returned the requested result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

Comments

SQLPrimaryKeys returns the results as a standard result set, ordered by TABLE_QUALIFIER, TABLE_OWNER, TABLE_NAME, and KEY_SEQ. The following table lists the columns in the result set.

Note: **SQLPrimaryKeys** might not return all primary keys. For example, a Paradox driver might only return primary keys for files (tables) in the current directory.

The lengths of VARCHAR columns shown in the table are maximums; the actual lengths depend on the data source. To determine the actual lengths of the TABLE_QUALIFIER, TABLE_OWNER, TABLE_NAME, and COLUMN_NAME columns, call **SQLGetInfo** with the SQL_MAX_QUALIFIER_NAME_LEN, SQL_MAX_OWNER_NAME_LEN, SQL_MAX_TABLE_NAME_LEN, and SQL_MAX_COLUMN_NAME_LEN options.

Column Name	Data Type	Comments
TABLE_QUALIFIER	Varchar(128)	Primary key table qualifier identifier; NULL if not applicable to the data source. If a driver supports qualifiers for some tables but not for others, such as

		when the driver retrieves data from different DBMSs, it returns an empty string ("") for those tables that do not have qualifiers.
TABLE_OWNER	Varchar(128)	Primary key table owner identifier; NULL if not applicable to the data source. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those tables that do not have owners.
TABLE_NAME	Varchar(128) not NULL	Primary key table identifier.
COLUMN_NAME	Varchar(128) not NULL	Primary key column identifier.
KEY_SEQ	Smallint not NULL	Column sequence number in key (starting with 1).
PK_NAME	Varchar(128)	Primary key identifier. NULL if not applicable to the data source.

Note: The PK_NAME column was added in ODBC 2.0. ODBC 1.0 drivers may return a different, driver-specific column with the same column number.

Code Example See **SQLForeignKeys**.

Related Functions

For information about

See

Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
Fetching a row of data	SQLFetch
Returning the columns of foreign keys	SQLForeignKeys (extension)
Returning table statistics and indexes	SQLStatistics (extension)

SQLProcedureColumns

ADBC

Extension Level 2 **SQLProcedureColumns** returns the list of input and output parameters, as well as the columns that make up the result set for the specified procedures. The driver returns the information as a result set on the specified *hstmt*.

Syntax **RETCODE SQLProcedureColumns**(*hstmt*, *szProcQualifier*, *cbProcQualifier*, *szProcOwner*, *cbProcOwner*, *szProcName*, *cbProcName*, *szColumnName*, *cbColumnName*)

The **SQLProcedureColumns** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UCHAR FAR *	<i>szProcQualifier</i>	Input	Procedure qualifier name. If a driver supports qualifiers for some procedures but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those procedures that do not have qualifiers.
SWORD	<i>cbProcQualifier</i>	Input	Length of <i>szProcQualifier</i> .
UCHAR FAR *	<i>szProcOwner</i>	Input	String search pattern for procedure owner names. If a driver supports owners for some procedures but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those procedures that do not have owners.
SWORD	<i>cbProcOwner</i>	Input	Length of <i>szProcOwner</i> .
UCHAR FAR *	<i>szProcName</i>	Input	String search pattern for procedure names.
SWORD	<i>cbProcName</i>	Input	Length of <i>szProcName</i> .
UCHAR FAR *	<i>szColumnName</i>	Input	String search pattern for column names.
SWORD	<i>cbColumnName</i>	Input	Length of <i>szColumnName</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLProcedureColumns** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be

obtained by calling **SQL_Error**. The following table lists the SQLSTATE values commonly returned by **SQLProcedureColumns** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATEs returned by the Driver Manager. The return code associated with each SQLSTATE value is `SQL_ERROR`, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns <code>SQL_SUCCESS_WITH_INFO</code> .)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called. A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQL_Error in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
SQLSTATE	Error	Description
S1010	Function sequence	(DM) An asynchronously executing

SQLTables

	error	<p>function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called.</p> <p>(DM) SQLExecute, SQLExecDirect, or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.</p>
S1090	Invalid string or buffer length	<p>(DM) The value of one of the name length arguments was less than 0, but not equal to SQL_NTS.</p> <p>The value of one of the name length arguments exceeded the maximum length value for the corresponding qualifier or name.</p>
S1C00	Driver not capable	<p>A procedure qualifier was specified and the driver or data source does not support qualifiers.</p> <p>A procedure owner was specified and the driver or data source does not support owners.</p> <p>A string search pattern was specified for the procedure owner, procedure name, or column name and the data source does not support search patterns for one or more of those arguments.</p> <p>The combination of the current settings of the SQL_CONCURRENCY and SQL_CURSOR_TYPE statement options was not supported by the driver or data source.</p>
S1T00	Timeout expired	<p>The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption, SQL_QUERY_TIMEOUT.</p>

Comments

This function is typically used before statement execution to retrieve information about procedure parameters and columns from the data source's catalog. For more information about stored procedures, see "Using ODBC Extensions to SQL" in Chapter 6, "Executing SQL Statements."

Note: **SQLProcedureColumns** might not return all columns used by a procedure. For example, a driver might only return information about the

parameters used by a procedure and not the columns in a result set it generates.

The *szProcOwner*, *szProcName*, and *szColumnName* arguments accept search patterns. For more information about valid search patterns, see “Search Pattern Arguments” earlier in this chapter.

SQLProcedureColumns returns the results as a standard result set, ordered by `PROCEDURE_QUALIFIER`, `PROCEDURE_OWNER`, `PROCEDURE_NAME`, and `COLUMN_TYPE`. The following table lists the columns in the result set. Additional columns beyond column 13 (REMARKS) can be defined by the driver.

The lengths of VARCHAR columns shown in the table are maximums; the actual lengths depend on the data source. To determine the actual lengths of the `PROCEDURE_QUALIFIER`, `PROCEDURE_OWNER`, `PROCEDURE_NAME`, and `COLUMN_NAME` columns, an application can call **SQLGetInfo** with the `SQL_MAX_QUALIFIER_NAME_LEN`, `SQL_MAX_OWNER_NAME_LEN`, `SQL_MAX_PROCEDURE_NAME_LEN`, and `SQL_MAX_COLUMN_NAME_LEN` options.

Column Name	Data Type	Comments
PROCEDURE_QUALIFIER	Varchar(128)	Procedure qualifier identifier; NULL if not applicable to the data source. If a driver supports qualifiers for some procedures but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those procedures that do not have qualifiers.
PROCEDURE_OWNER	Varchar(128)	Procedure owner identifier; NULL if not applicable to the data source. If a driver supports owners for some procedures but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those procedures that do not have owners.
PROCEDURE_NAME	Varchar(128) not NULL	Procedure identifier.
COLUMN_NAME	Varchar(128) not NULL	Procedure column identifier.
Column Name	Data Type	Comments

SQLTables

COLUMN_TYPE	Smallint not NULL	<p>Defines the procedure column as parameter or a result set column:</p> <p>SQL_PARAM_TYPE_UNKNOWN: The procedure column is a parameter whose type is unknown. (ODBC 1.0)</p> <p>SQL_PARAM_INPUT: The procedure column is an input parameter. (ODBC 1.0)</p> <p>SQL_PARAM_INPUT_OUTPUT: the procedure column is an input/output parameter. (ODBC 1.0)</p> <p>SQL_PARAM_OUTPUT: The procedure column is an output parameter. (ODBC 1.0)</p> <p>SQL_RETURN_VALUE: The procedure column is the return value of the procedure. (ODBC 2.0)</p> <p>SQL_RESULT_COL: The procedure column is a result set column. (ODBC 1.0)</p>
DATA_TYPE	Smallint not NULL	SQL data type. This can be an ODBC SQL data type or a driver-specific SQL data type. For a list of valid ODBC SQL data types, see "SQL Data Types" in Appendix D, "Data Types." For information about driver-specific SQL data types, see the driver's documentation.
TYPE_NAME	Varchar(128) not NULL	Data source-dependent data type name; for example, "CHAR", "VARCHAR" or "LONG".
PRECISION	Integer	The precision of the procedure column on the data source. NULL is returned for data types where precision is not applicable. For more information concerning precision, see "Precision, Scale, Length, and Display Size," in Appendix D, "Data Types."
Column Name	Data Type	Comments
LENGTH	Integer	The length in bytes of data transferred on an SQLGetData or SQLFetch operation if SQL_C_DEFAULT is specified. For numeric data, this size

			may be different than the size of the data stored on the data source. For more information, see "Precision, Scale, Length, and Display Size," in Appendix D, "Data Types."
SCALE	Smallint		The scale of the procedure column on the data source. NULL is returned for data types where scale is not applicable. For more information concerning scale, see "Precision, Scale, Length, and Display Size," in Appendix D, "Data Types."
RADIX	Smallint		For numeric data types, either 10 or 2. If it is 10, the values in PRECISION and SCALE give the number of decimal digits allowed for the column. For example, a DECIMAL(12,5) column would return a RADIX of 10, a PRECISION of 12, and a SCALE of 5; a FLOAT column could return a RADIX of 10, a PRECISION of 15 and a SCALE of NULL. If it is 2, the values in PRECISION and SCALE give the number of bits allowed in the column. For example, a FLOAT column could return a RADIX of 2, a PRECISION of 53, and a SCALE of NULL. NULL is returned for data types where radix is not applicable.
NULLABLE	Smallint not NULL		Whether the procedure column accepts a NULL value: SQL_NO_NULLS: The procedure column does not accept NULL values. SQL_NULLABLE: The procedure column accepts NULL values. SQL_NULLABLE_UNKNOWN: It is not known if the procedure column accepts NULL values.
REMARKS	Varchar(254)		A description of the procedure column.

Code Example See **SQLProcedures**.

Related Functions For information about

See

SQLTables

Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
Fetching a row of data	SQLFetch
Returning a list of procedures in a data source	SQLProcedures (extension)

SQLProcedures



Extension Level 2 SQLProcedures returns the list of procedure names stored in a specific data source. *Procedure* is a generic term used to describe an *executable object*, or a named entity that can be invoked using input and output parameters, and which can return result sets similar to the results returned by SQL **SELECT** expressions.

Syntax RETCODE **SQLProcedures**(*hstmt*, *szProcQualifier*, *cbProcQualifier*, *szProcOwner*, *cbProcOwner*, *szProcName*, *cbProcName*)

The **SQLProcedures** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UCHAR FAR *	<i>szProcQualifier</i>	Input	Procedure qualifier. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have qualifiers.
SWORD	<i>cbProcQualifier</i>	Input	Length of <i>szProcQualifier</i> .
UCHAR FAR *	<i>szProcOwner</i>	Input	String search pattern for procedure owner names. If a driver supports owners for some procedures but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those procedures that do not have owners.
SWORD	<i>cbProcOwner</i>	Input	Length of <i>szProcOwner</i> .
UCHAR FAR *	<i>szProcName</i>	Input	String search pattern for procedure names.
SWORD	<i>cbProcName</i>	Input	Length of <i>szProcName</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLProcedures** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLProcedures** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code

SQLTransact

associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called. A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support this function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
SQLSTATE	Error	Description
S1010	Function sequence error	(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called.

		(DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1090	Invalid string or buffer length	(DM) The value of one of the name length arguments was less than 0, but not equal to SQL_NTS. The value of one of the name length arguments exceeded the maximum length value for the corresponding qualifier or name.
S1C00	Driver not capable	A procedure qualifier was specified and the driver or data source does not support qualifiers. A procedure owner was specified and the driver or data source does not support owners. A string search pattern was specified for the procedure owner or procedure name and the data source does not support search patterns for one or more of those arguments. The combination of the current settings of the SQL_CONCURRENCY and SQL_CURSOR_TYPE statement options was not supported by the driver or data source.
S1T00	Timeout expired	The timeout period expired before the data source returned the requested result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

Comments **SQLProcedures** lists all procedures in the requested range. A user may or may not have permission to execute any of these procedures. To check accessibility, an application can call **SQLGetInfo** and check the SQL_ACCESSIBLE_PROCEDURES information value. Otherwise, the application must be able to handle a situation where the user selects a procedure which it cannot execute.

Note: SQLProcedures might not return all procedures. Applications can use any valid procedure, regardless of whether it is returned by **SQLProcedures**.

SQLProcedures returns the results as a standard result set, ordered by PROCEDURE_QUALIFIER, PROCEDURE_OWNER, and PROCEDURE_NAME. The following table lists the columns in the result set.

The lengths of VARCHAR columns shown in the table are maximums; the actual lengths depend on the data source. To determine the actual lengths of the PROCEDURE_QUALIFIER, PROCEDURE_OWNER, and PROCEDURE_NAME columns, an application can call **SQLGetInfo** with the SQL_MAX_QUALIFIER_NAME_LEN, SQL_MAX_OWNER_NAME_LEN, and SQL_MAX_PROCEDURE_NAME_LEN options.

Column Name	Data Type	Comments
PROCEDURE_QUALIFIER	Varchar(128)	Procedure qualifier identifier; NULL if not applicable to the data source. If a driver supports qualifiers for some procedures but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those procedures that do not have qualifiers.
PROCEDURE_OWNER	Varchar(128)	Procedure owner identifier; NULL if not applicable to the data source. If a driver supports owners for some procedures but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those procedures that do not have owners.
PROCEDURE_NAME	Varchar(128) not NULL	Procedure identifier.
NUM_INPUT_PARAMS	N/A	Reserved for future use. Applications should not rely on the data returned in these result columns.
NUM_OUTPUT_PARAMS	N/A	Reserved for future use. Applications should not rely on the data returned in these result columns.

NUM_RESULT_SETS	N/A	Reserved for future use. Applications should not rely on the data returned in these result columns.
REMARKS	Varchar(254)	A description of the procedure.
Column Name	Data Type	Comments
PROCEDURE_TYPE	Smallint	Defines the procedure type: SQL_PT_UNKNOWN: It cannot be determined whether the procedure returns a value. SQL_PT_PROCEDURE: The returned object is a procedure; that is, it does not have a return value. SQL_PT_FUNCTION: The returned object is a function; that is, it has a return value.

Note: The PROCEDURE_TYPE column was added in ODBC 2.0. ODBC 1.0 drivers might return a different, driver-specific column with the same column number.

The *szProcOwner* and *szProcName* arguments accept search patterns. For more information about valid search patterns, see “Search Pattern Arguments” earlier in this chapter.

Code Example

In this example, an application uses the procedure **AddEmployee** to insert data into the EMPLOYEE table. The procedure contains input parameters for NAME, AGE, and BIRTHDAY columns. It also contains one output parameter that returns a remark about the new employee. The example also shows the use of a return value from a stored procedure. For the return value and each parameter in the procedure, the application calls **SQLBindParameter** to specify the ODBC C data type and the SQL data type of the parameter and to specify the storage location and length of the parameter. The application assigns data values to the storage locations for each parameter and calls **SQLExecDirect** to execute the procedure. If **SQLExecDirect** returns SQL_SUCCESS or SQL_SUCCESS_WITH_INFO, the return value and the value of each output or input/output parameter is automatically put into the storage location defined for the parameter in **SQLBindParameter**.

```
#define NAME_LEN 30
#define REM_LEN 128

UCHAR    szName[NAME_LEN], szRemark[REM_LEN];
WORD     sAge, sEmpld;
```

SQLTransact

```
SDWORD    cbEmpld, cbName, cbAge = 0, cbBirthday = 0, cbRemark;  
DATE_STRUCT dsBirthday;
```

```

/* Define parameter for return value (Employee ID) from procedure. */

SQLBindParameter(hstmt, 1, SQL_PARAM_OUTPUT, SQL_C_SLONG, SQL_INTEGER,
                 0, 0, &sEmpId, 0, &cbEmpId);

/* Define data types and storage locations for Name, Age, Birthday */
/* input parameter data. */

SQLBindParameter(hstmt, 2, SQL_PARAM_INPUT, SQL_C_CHAR, SQL_CHAR,
                 NAME_LEN, 0, szName, 0, &cbName);
SQLBindParameter(hstmt, 3, SQL_PARAM_INPUT, SQL_C_SSHORT, SQL_SMALLINT,
                 0, 0, &sAge, 0, &cbAge);
SQLBindParameter(hstmt, 4, SQL_PARAM_INPUT, SQL_C_DATE, SQL_DATE,
                 0, 0, &dsBirthday, 0, &cbBirthday);

/* Define data types and storage location for Remark output parameter */

SQLBindParameter(hstmt, 5, SQL_PARAM_OUTPUT, SQL_C_CHAR, SQL_CHAR,
                 REM_LEN, 0, szRemark, REM_LEN, &cbRemark);

strcpy(szName, "Smith, John D."); /* Specify first row of */
sAge = 40; /* parameter data. */
dsBirthday.year = 1952;
dsBirthday.month = 2;
dsBirthday.day = 29;
cbName = SQL_NTS;

/* Execute procedure with first row of data. After the procedure */
/* is executed, sEmpId and szRemark will have the values */
/* returned by AddEmployee. */

retcode = SQLExecDirect(hstmt, "{?=call AddEmployee(?,?,?)", SQL_NTS);

strcpy(szName, "Jones, Bob K."); /* Specify second row of */
sAge = 52; /* parameter data */
dsBirthday.year = 1940;
dsBirthday.month = 3;
dsBirthday.day = 31;

/* Execute procedure with second row of data. After the procedure */
/* is executed, sEmpId and szRemark will have the new values */
/* returned by AddEmployee. */

retcode = SQLExecDirect(hstmt,
                        "{?=call AddEmployee(?,?,?)", SQL_NTS);

```

**Related
Functions**

For information about

See

Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
Fetching a row of data	SQLFetch
Returning information about a driver or data source	SQLGetInfo (extension)
Returning the parameters and result set columns of a procedure	SQLProcedureColumns (extension)
Syntax for invoking stored procedures	Chapter 6 , "Executing SQL Statements"

SQLPutData



Extension Level 1 **SQLPutData** allows an application to send data for a parameter or column to the driver at statement execution time. This function can be used to send character or binary data values in parts to a column with a character, binary, or data source–specific data type (for example, parameters of the SQL_LONGVARIABLE or SQL_LONGVARIABLE types).

Syntax

RETCODE **SQLPutData**(*hstmt*, *rgbValue*, *cbValue*)

The **SQLPutData** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
PTR	<i>rgbValue</i>	Input	Pointer to storage for the actual data for the parameter or column. The data must use the C data type specified in the <i>fCType</i> argument of SQLBindParameter (for parameter data) or SQLBindCol (for column data).
SDWORD	<i>cbValue</i>	Input	Length of <i>rgbValue</i> . Specifies the amount of data sent in a call to SQLPutData .

BEGIN BREAK

The amount of data can vary with each call for a given parameter or column. *cbValue* is ignored unless it is SQL_NTS, SQL_NULL_DATA, or SQL_DEFAULT_PARAM; the C data type specified in **SQLBindParameter** or **SQLBindCol** is SQL_C_CHAR or SQL_C_BINARY; or the C data type is SQL_C_DEFAULT and the default C data type for the specified SQL data type is SQL_C_CHAR or SQL_C_BINARY. For all other types of C data, if *cbValue* is not SQL_NULL_DATA or SQL_DEFAULT_PARAM, the driver assumes that the size of *rgbValue* is the size of the C data type specified with *fCType* and sends the entire data value.

For more information, see
“Converting Data from C to SQL
Data Types” in Appendix D, “Data
Types.”

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO,
SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLPutData** returns SQL_ERROR or
SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be
obtained by calling **SQLError**. The following table lists the SQLSTATE
values commonly returned by **SQLPutData** and explains each one in the
context of this function; the notation “(DM)” precedes the descriptions of
SQLSTATEs returned by the Driver Manager. The return code
associated with each SQLSTATE value is SQL_ERROR, unless noted
otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	The data sent for a character or binary parameter or column in one or more calls to SQLPutData exceeded the maximum length of the associated character or binary column. The fractional part of the data sent for a numeric or bit parameter or column was truncated. Timestamp data sent for a date or time parameter or column was truncated.
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
SQLSTATE	Error	Description
22001	String data right truncation	The SQL_NEED_LONG_DATA_LEN information type in SQLGetInfo was “Y” and more data was sent for a long parameter (the data type was SQL_LONGVARCHAR, SQL_LONGVARBINARY, or a long, data source-specific data type) than was specified with the <i>pcbValue</i> argument in

SQLBindParameter.

The `SQL_NEED_LONG_DATA_LEN` information type in **SQLGetInfo** was “Y” and more data was sent for a long column (the data type was `SQL_LONGVARCHAR`, `SQL_LONGVARBINARY`, or a long, data source-specific data type) than was specified in the length buffer corresponding to a column in a row of data that was added or updated with **SQLSetPos**.

22003 Numeric value out of range

SQLPutData was called more than once for a parameter or column and it was not being used to send character C data to a column with a character, binary, or data source-specific data type or to send binary C data to a column with a character, binary, or data source-specific data type.

The data sent for a numeric parameter or column caused the whole (as opposed to fractional) part of the number to be truncated when assigned to the associated table column.

SQLSTATE Error**Description**

22005	Error in assignment	The data sent for a parameter or column was incompatible with the data type of the associated table column.
22008	Datetime field overflow	The data sent for a date, time, or timestamp parameter or column was, respectively, an invalid date, time, or timestamp.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution

S1008	Operation canceled	<p>or completion of the function.</p> <p>Asynchronous processing was enabled for the <i>hstmt</i>. The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i>. Then the function was called again on the <i>hstmt</i>.</p> <p>The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.</p> <p>SQLExecute, SQLExecDirect, or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. SQLCancel was called before data was sent for all data-at-execution parameters or columns.</p>
S1009	Invalid argument value	(DM) The argument <i>rgbValue</i> was a null pointer and the argument <i>cbValue</i> was not 0, SQL_DEFAULT_PARAM, or SQL_NULL_DATA.
SQLSTATE	Error	Description
S1010	Function sequence error	<p>(DM) The previous function call was not a call to SQLPutData or SQLParamData.</p> <p>The previous function call was a call to SQLExecDirect, SQLExecute, or SQLSetPos where the return code was SQL_NEED_DATA.</p> <p>(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called.</p>
S1090	Invalid string or buffer length	The argument <i>rgbValue</i> was not a null pointer and the argument <i>cbValue</i> was less than 0, but not equal to SQL_NTS or SQL_NULL_DATA.
S1T00	Timeout expired	The timeout period expired before the data source completed processing the parameter value. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

Comments

For an explanation of how data-at-execution parameter data is passed at statement execution time, see “Passing Parameter Values” in **SQLBindParameter**. For an explanation of how data-at-execution column data is updated or added, see “Using SQLSetPos” in **SQLSetPos**.

Note: An application can use **SQLPutData** to send data in parts only when sending character C data to a column with a character, binary, or data source–specific data type or when sending binary C data to a column with a character, binary, or data source–specific data type. If **SQLPutData** is called more than once under any other conditions, it returns SQL_ERROR and SQLSTATE 22003 (Numeric value out of range).

Code Example

In the following example, an application prepares an SQL statement to insert data into the EMPLOYEE table. The statement contains parameters for the NAME, ID, and PHOTO columns. For each parameter, the application calls **SQLBindParameter** to specify the C and SQL data types of the parameter. It also specifies that the data for the first and third parameters will be passed at execution time, and passes the values 1 and 3 for later retrieval by **SQLParamData**. These values will identify which parameter is being processed.

The application calls **GetNextID** to get the next available employee ID number. It then calls **SQLExecute** to execute the statement. **SQLExecute** returns **SQL_NEED_DATA** when it needs data for the first and third parameters. The application calls **SQLParamData** to retrieve the value it stored with **SQLBindParameter**; it uses this value to determine which parameter to send data for. For each parameter, the application calls **InitUserData** to initialize the data routine. It repeatedly calls **GetUserData** and **SQLPutData** to get and send the parameter data. Finally, it calls **SQLParamData** to indicate it has sent all the data for the parameter and to retrieve the value for the next parameter. After data has been sent for both parameters, **SQLParamData** returns **SQL_SUCCESS**.

For the first parameter, **InitUserData** does not do anything and **GetUserData** calls a routine to prompt the user for the employee name. For the third parameter, **InitUserData** calls a routine to prompt the user for the name of a file containing a bitmap photo of the employee and opens the file. **GetUserData** retrieves the next **MAX_DATA_LEN** bytes of photo data from the file. After it has retrieved all the photo data, it closes the photo file.

Note that some application routines are omitted for clarity.

```
#define NAME_LEN 30
#define MAX_DATA_LEN 1024
SDWORD cbNameParam, cbID = 0; cbPhotoParam, cbData;
SWORD sID;
PTR pToken, InitValue;
UCHAR Data[MAX_DATA_LEN];

retcode = SQLPrepare(hstmt,
    "INSERT INTO EMPLOYEE (NAME, ID, PHOTO) VALUES (?, ?, ?)",
    SQL_NTS);
if (retcode == SQL_SUCCESS) {

    /* Bind the parameters. For parameters 1 and 3, pass the */
    /* parameter number in rgbValue instead of a buffer address. */

    SQLBindParameter(hstmt, 1, SQL_PARAM_INPUT, SQL_C_CHAR, SQL_CHAR,
        NAME_LEN, 0, 1, 0, &cbNameParam);
    SQLBindParameter(hstmt, 2, SQL_PARAM_INPUT, SQL_C_SSHORT,
        SQL_SMALLINT, 0, 0, &sID, 0, &cbID);
    SQLBindParameter(hstmt, 3, SQL_PARAM_INPUT,
        SQL_C_BINARY, SQL_LONGVARBINARY,
        0, 0, 3, 0, &cbPhotoParam);
```

```

/* Set values so data for parameters 1 and 3 will be passed */
/* at execution. Note that the length parameter in the macro */
/* SQL_LEN_DATA_AT_EXEC is 0. This assumes that the driver */
/* returns "N" for the SQL_NEED_LONG_DATA_LEN information */
/* type in SQLGetInfo. */

cbNameParam = cbPhotoParam = SQL_LEN_DATA_AT_EXEC(0);

SID = GetNextID(); /* Get next available employee ID number. */

retcode = SQLExecute(hstmt);

/* For data-at-execution parameters, call SQLParamData to get the */
/* parameter number set by SQLBindParameter. Call InitUserData. */
/* Call GetUserData and SQLPutData repeatedly to get and put all */
/* data for the parameter. Call SQLParamData to finish processing */
/* this parameter and start processing the next parameter. */

while (retcode == SQL_NEED_DATA) {
    retcode = SQLParamData(hstmt, &pToken);
    if (retcode == SQL_NEED_DATA) {
        InitUserData((SWORD)pToken, InitValue);
        while (GetUserData(InitValue, (SWORD)pToken, Data, &cbData))
            SQLPutData(hstmt, Data, cbData);
    }
}

}

VOID InitUserData(sParam, InitValue)
SWORD sParam;
PTR InitValue;
{
    UCHAR szPhotoFile[MAX_FILE_NAME_LEN];
    switch sParam {
        case 3:

            /* Prompt user for bitmap file containing employee photo. */
            /* OpenPhotoFile opens the file and returns the file handle. */

            PromptPhotoFileName(szPhotoFile);
            OpenPhotoFile(szPhotoFile, (FILE *)InitValue);
            break;
    }
}

```

Part 2 Developing Applications

```
BOOL GetUserData(InitValue, sParam, Data, cbData)
PTR  InitValue;
SWORD sParam;
UCHAR *Data;
SDWORD *cbData;

{

switch sParam {
  case 1:
    /* Prompt user for employee name. */

    PromptEmployeeName(Data);
    *cbData = SQL_NTS;
    return (TRUE);

  case 3:
    /* GetNextPhotoData returns the next piece of photo data and */
    /* the number of bytes of data returned (up to MAX_DATA_LEN). */

    Done = GetNextPhotoData((FILE *)InitValue, Data,
                           MAX_DATA_LEN, &cbData);
    if (Done) {
      ClosePhotoFile((FILE *)InitValue);
      return (TRUE);
    }
    return (FALSE);
}
return (FALSE);
}
```

Related Functions

For information about

See

Canceling statement processing	SQLCancel
Executing an SQL statement	SQLExecDirect
Executing a prepared SQL statement	SQLExecute
Returning the next parameter to send data for	SQLParamData (extension)
Assigning storage for a parameter	SQLBindParameter

!Unexpected End of Expression

SQLRowCount

ADBC

Core **SQLRowCount** returns the number of rows affected by an **UPDATE**, **INSERT**, or **DELETE** statement or by a SQL_UPDATE, SQL_ADD, or SQL_DELETE operation in **SQLSetPos**.

Syntax RETCODE **SQLRowCount**(*hstmt*, *pcrow*)
The **SQLRowCount** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
SDWORD FAR	* <i>pcrow</i>	Output	For UPDATE , INSERT , and DELETE statements and for the SQL_UPDATE, SQL_ADD, and SQL_DELETE operations in SQLSetPos , <i>pcrow</i> is the number of rows affected by the request or –1 if the number of affected rows is not available. For other statements and functions, the driver may define the value of <i>pcrow</i> . For example, some data sources may be able to return the number of rows returned by a SELECT statement or a catalog function before fetching the rows. Note: Many data sources cannot return the number of rows in a result set before fetching them; for maximum interoperability, applications should not rely on this behavior.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLRowCount** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLRowCount** and explains each one in the context of this function; the notation “(DM)” precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code

associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1010	Function sequence error	(DM) The function was called prior to calling SQLExecute , SQLExecDirect , SQLSetPos for the <i>hstmt</i> . (DM) An asynchronously executing function was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.

Comments

If the last executed statement associated with *hstmt* was not an **UPDATE**, **INSERT**, or **DELETE** statement, or if the *fOption* argument in the previous call to **SQLSetPos** was not SQL_UPDATE, SQL_ADD, or SQL_DELETE, the value of *pcrow* is driver-defined.

Related Functions

For information about	See
Executing an SQL statement	SQLExecDirect
Executing a prepared SQL statement	SQLExecute

SQLSetConnectOption



Extension Level 1 **SQLSetConnectOption** sets options that govern aspects of connections.

Syntax RETCODE **SQLSetConnectOption**(*hdbc*, *fOption*, *vParam*)

The **SQLSetConnectOption** function accepts the following arguments:

Type	Argument	Use	Description
HDBC	<i>hdbc</i>	Input	Connection handle.
UWORD	<i>fOption</i>	Input	Option to set, listed in “Comments.”
UDWORD	<i>vParam</i>	Input	Value associated with <i>fOption</i> . Depending on the value of <i>fOption</i> , <i>vParam</i> will be a 32-bit integer value or point to a null-terminated character string.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLSetConnectOption** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLSetConnectOption** and explains each one in the context of this function; the notation “(DM)” precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

The driver can return SQL_SUCCESS_WITH_INFO to provide information about the result of setting an option. For example, setting SQL_ACCESS_MODE to read-only during a transaction might cause the transaction to be committed. The driver could use SQL_SUCCESS_WITH_INFO—and information returned with **SQLError**—to inform the application of the commit action.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01S02	Option value changed	The driver did not support the specified value of the <i>vParam</i> argument and substituted a similar value. (Function

SQLSTATE Error		Description
08002	Connection in use	The argument <i>fOption</i> was SQL_ODBC_CURSORS and the driver was already connected to the data source.
08003	Connection not open	An <i>fOption</i> value was specified that required an open connection, but the <i>hdbc</i> was not in a connected state.
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hdbc</i> does not support the function.
IM009	Unable to load translation DLL	The driver was unable to load the translation DLL that was specified for the connection. This error can only be returned when <i>fOption</i> is SQL_TRANSLATE_DLL.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable 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 was specified for the argument <i>vParam</i> . (The Driver Manager returns this SQLSTATE only for connection and statement options that accept a discrete set of values, such as SQL_ACCESS_MODE or SQL_ASYNC_ENABLE. For all other connection and statement options, the driver must verify the value of the argument <i>vParam</i> .)
SQLSTATE Error		Description
S1010	Function sequence error	(DM) An asynchronously executing function was called for an <i>hstmt</i> associated with the <i>hdbc</i> and was still executing when SQLSetConnectOption was called.

ConfigDSN (Setup)

		(DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for an <i>hstmt</i> associated with the <i>hdbc</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns. (DM) SQLBrowseConnect was called for the <i>hdbc</i> and returned SQL_NEED_DATA. This function was called before SQLBrowseConnect returned SQL_SUCCESS_WITH_INFO or SQL_SUCCESS.
S1011	Operation invalid at this time	The argument <i>fOption</i> was SQL_TXN_ISOLATION and a transaction was open.
S1092	Option type out of range	(DM) The value specified for the argument <i>fOption</i> was in the block of numbers reserved for ODBC connection and statement options, but was not valid for the version of ODBC supported by the driver.
S1C00	Driver not capable	The value specified for the argument <i>fOption</i> was a valid ODBC connection or statement option for the version of ODBC supported by the driver, but was not supported by the driver. The value specified for the argument <i>fOption</i> was in the block of numbers reserved for driver-specific connection and statement options, but was not supported by the driver.
Comments	When <i>fOption</i> is a statement option, SQLSetConnectOption can return any SQLSTATES returned by SQLSetStmtOption .	
	The currently defined options and the version of ODBC in which they were introduced are shown below; it is expected that more will be defined to take advantage of different data sources. Options from 0 to 999 are reserved by ODBC; driver developers must reserve values greater than or equal to SQL_CONNECT_OPT_DVR_START for driver-specific use.	
	An application can call SQLSetConnectOption and include a statement option. The driver sets the statement option for any <i>hstmts</i> associated with the specified <i>hdbc</i> and establishes the statement option as a default for any <i>hstmts</i> later allocated for that <i>hdbc</i> . For a list of statement options, see SQLSetStmtOption .	

All connection and statement options successfully set by the application for the *hdbc* persist until **SQLFreeConnect** is called on the *hdbc*. For example, if an application calls **SQLSetConnectOption** before connecting to a data source, the option persists even if **SQLSetConnectOption** fails in the driver when the application connects to the data source; if an application sets a driver-specific option, the option persists even if the application connects to a different driver on the *hdbc*.

Some connection and statement options support substitution of a similar value if the data source does not support the specified value of *vParam*. In such cases, the driver returns SQL_SUCCESS_WITH_INFO and SQLSTATE 01S02 (Option value changed). For example, if *fOption* is SQL_PACKET_SIZE and *vParam* exceeds the maximum packet size, the driver substitutes the maximum size. To determine the substituted value, an application calls **SQLGetConnectOption** (for connection options) or **SQLGetStmtOption** (for statement options).

The format of information set through *vParam* depends on the specified *fOption*. **SQLSetConnectOption** will accept option information in one of two different formats: a null-terminated character string or a 32-bit integer value. The format of each is noted in the option's description. Character strings pointed to by the *vParam* argument of **SQLSetConnectOption** have a maximum length of SQL_MAX_OPTION_STRING_LENGTH bytes (excluding the null termination byte).

<i>fOption</i>	<i>vParam</i> Contents
SQL_ACCESS_MODE (ODBC 1.0)	A 32-bit integer value. SQL_MODE_READ_ONLY is used by the driver or data source as an indicator that the connection is not required to support SQL statements that cause updates to occur. This mode can be used to optimize locking strategies, transaction management, or other areas as appropriate to the driver or data source. The driver is not required to prevent such statements from being submitted to the data source. The behavior of the driver and data source when asked to process SQL statements that are not read-only during a read-only connection is implementation defined. SQL_MODE_READ_WRITE is the default.
SQL_AUTOCOMMIT (ODBC 1.0)	A 32-bit integer value that specifies whether to use auto-commit or manual-commit mode: SQL_AUTOCOMMIT_OFF = The driver uses manual-commit mode, and the application must explicitly commit or roll back transactions with SQLTransact . SQL_AUTOCOMMIT_ON = The driver uses auto-commit mode. Each statement is committed immediately after it is executed. This is the default. Note that changing from manual-commit mode to

ConfigTranslator (Setup)

ConfigDSN (Setup)

auto-commit mode commits any open transactions on the connection.

Important: Some data sources delete the access plans and close the cursors for all *hstmts* on an *hdbc* each time a statement is committed; autocommit mode can cause this to happen after each statement is executed. For more information, see the `SQL_CURSOR_COMMIT_BEHAVIOR` and `SQL_CURSOR_ROLLBACK_BEHAVIOR` information types in **SQLGetInfo**.

SQL_CURRENT_QUALIFIER (ODBC 2.0)	A null-terminated character string containing the name of the qualifier to be used by the data source. For example, in SQL Server, the qualifier is a database, so the driver sends a USE database statement to the data source, where <i>database</i> is the database specified in <i>vParam</i> . For a single-tier driver, the qualifier might be a directory, so the driver changes its current directory to the directory specified in <i>vParam</i> .
SQL_LOGIN_TIMEOUT (ODBC 1.0)	<p>A 32-bit integer value corresponding to the number of seconds to wait for a login request to complete before returning to the application. The default is driver-dependent and must be nonzero. If <i>vParam</i> is 0, the timeout is disabled and a connection attempt will wait indefinitely.</p> <p>If the specified timeout exceeds the maximum login timeout in the data source, the driver substitutes that value and returns SQLSTATE 01S02 (Option value changed).</p>
fOption	vParam Contents
SQL_ODBC_CURSORS (ODBC 2.0)	<p>A 32-bit option specifying how the Driver Manager uses the ODBC cursor library:</p> <p><code>SQL_CUR_USE_IF_NEEDED</code> = The Driver Manager uses the ODBC cursor library only if it is needed. If the driver supports the <code>SQL_FETCH_PRIOR</code> option in SQLExtendedFetch, the Driver Manager uses the scrolling capabilities of the driver. Otherwise, it uses the ODBC cursor library.</p> <p><code>SQL_CUR_USE_ODBC</code> = The Driver Manager uses the ODBC cursor library.</p> <p><code>SQL_CUR_USE_DRIVER</code> = The Driver Manager uses the scrolling capabilities of the driver. This is the default setting.</p> <p>For more information about the ODBC cursor library, see Appendix G, "ODBC Cursor Library."</p>
SQL_OPT_TRACE (ODBC 1.0)	A 32-bit integer value telling the Driver Manager whether to perform tracing:

SQL_OPT_TRACE_OFF = Tracing off (the default)

SQL_OPT_TRACE_ON = Tracing on

When tracing is on, the Driver Manager writes each ODBC function call to the trace file. On Windows and WOW, the Driver Manager writes to the trace file each time any application calls a function. On Windows NT, the Driver Manager writes to the trace file only for the application that turned tracing on.

Note: When tracing is on, the Driver Manager can return SQLSTATE IM013 (Trace file error) from any function.

An application specifies a trace file with the SQL_OPT_TRACEFILE option. If the file already exists, the Driver Manager appends to the file. Otherwise, it creates the file. If tracing is on and no trace file has been specified, the Driver Manager writes to the file \SQL.LOG. On Windows NT, tracing should only be used for a single application or each application should specify a different trace file. Otherwise, two or more applications will attempt to open the same trace file at the same time, causing an error.

If the **Trace** keyword in the [ODBC] section of the ODBC.INI file (or registry) is set to 1 when an application calls **SQLAllocEnv**, tracing is enabled. On Windows and WOW, it is enabled for all applications; on Windows NT it is enabled only for the application that called **SQLAllocEnv**.

fOption

vParam Contents

SQL_OPT_TRACEFILE
(ODBC 1.0)

A null-terminated character string containing the name of the trace file.

The default value of the SQL_OPT_TRACEFILE option is specified with the TraceFile keyname in the [ODBC] section of the ODBC.INI file (or registry).

SQL_PACKET_SIZE
(ODBC 2.0)

A 32-bit integer value specifying the network packet size in bytes.

Note: Many data sources either do not support this option or can only return the network packet size.

SQL_QUIET_MODE
(ODBC 2.0)

If the specified size exceeds the maximum packet size or is smaller than the minimum packet size, the driver substitutes that value and returns SQLSTATE 01S02 (Option value changed).

A 32-bit window handle (*hwnd*).

If the window handle is a null pointer, the driver does not display any dialog boxes.

If the window handle is not a null pointer, it should be the parent window handle of the application. The driver uses this handle to

ConfigTranslator (Setup)

ConfigDSN (Setup)

display dialog boxes. This is the default.

If the application has not specified a parent window handle for this option, the driver uses a null parent window handle to display dialog boxes or return in **SQLGetConnectOption**.

Note: The SQL_QUIET_MODE connection option does not apply to dialog boxes displayed by **SQLDriverConnect**.

SQL_TRANSLATE_DLL
(ODBC 1.0)

A null-terminated character string containing the name of a DLL containing the functions **SQLDriverToDataSource** and **SQLDataSourceToDriver** that the driver loads and uses to perform tasks such as character set translation. This option may only be specified if the driver has connected to the data source. For more information about translating data, see Chapter 25, "Translation DLL Function Reference."

SQL_TRANSLATE_OPTION
(ODBC 1.0)

A 32-bit flag value that is passed to the translation DLL. This option may only be specified if the driver has connected to the data source.

fOption

vParam Contents

SQL_TXN_ISOLATION
(ODBC 1.0)

A 32-bit bitmask that sets the transaction isolation level for the current *hdbc*. An application must call **SQLTransact** to commit or roll back all open transactions on an *hdbc*, before calling **SQLSetConnectOption** with this option.

The valid values for *vParam* can be determined by calling **SQLGetInfo** with *fInfoType* equal to SQL_TXN_ISOLATION_OPTIONS. The following terms are used to define transaction isolation levels:

Dirty Read Transaction 1 changes a row. Transaction 2 reads the changed row before transaction 1 commits the change. If transaction 1 rolls back the change, transaction 2 will have read a row that is considered to have never existed.

Nonrepeatable Read Transaction 1 reads a row. Transaction 2 updates or deletes that row and commits this change. If transaction 1 attempts to reread the row, it will receive different row values or discover that the row has been deleted.

Phantom Transaction 1 reads a set of rows that satisfy some search criteria. Transaction 2 inserts a row that matches the search criteria. If transaction 1 reexecutes the statement that read the rows, it receives a different set of rows.

vParam must be one of the following values:

SQL_TXN_READ_UNCOMMITTED = Dirty reads, nonrepeatable reads, and phantoms are possible.

SQL_TXN_READ_COMMITTED = Dirty reads are not possible.

Nonrepeatable reads and phantoms are possible.

SQL_TXN_REPEATABLE_READ = Dirty reads and nonrepeatable reads are not possible. Phantoms are possible.

SQL_TXN_SERIALIZABLE = Transactions are serializable. Dirty reads, nonrepeatable reads, and phantoms are not possible.

SQL_TXN_VERSIONING = Transactions are serializable, but higher concurrency is possible than with SQL_TXN_SERIALIZABLE. Dirty reads are not possible. Typically, SQL_TXN_SERIALIZABLE is implemented by using locking protocols that reduce concurrency and SQL_TXN_VERSIONING is implemented by using a non-locking protocol such as record versioning. Oracle's Read Consistency isolation level is an example of SQL_TXN_VERSIONING.

Data Translation

Data translation will be performed for all data flowing between the driver and the data source.

The translation option (set with the SQL_TRANSLATE_OPTION option) can be any 32-bit value. Its meaning depends on the translation DLL being used. A new option can be set at any time. The new option will be applied to the next exchange of data following the call to **SQLSetConnectOption**. A default translation DLL may be specified for the data source in its data source specification in the ODBC.INI file or registry. The default translation DLL is loaded by the driver at connection time. A translation option (SQL_TRANSLATE_OPTION) may be specified in the data source specification as well.

To change the translation DLL for a connection, an application calls **SQLSetConnectOption** with the SQL_TRANSLATE_DLL option after it has connected to the data source. The driver will attempt to load the specified DLL and, if the attempt fails, return SQL_ERROR with the SQLSTATE IM009 (Unable to load translation DLL).

If no translation DLL has been specified in the ODBC initialization file or by calling **SQLSetConnectOption**, the driver will not attempt to translate data. Any value set for the translation option will be ignored.

For more information about translating data, see Chapter 25, "Translation Function DLL Reference."

Code Example See **SQLConnect** and **SQLParamOptions**.

Related Functions For information about See

ConfigTranslator (Setup)

ConfigDSN (Setup)

Returning the setting of a connection option	SQLGetConnectOption (extension)
Returning the setting of a statement option	SQLGetStmtOption (extension)
Setting a statement option	SQLSetStmtOption (extension)

SQLSetCursorName



Core **SQLSetCursorName** associates a cursor name with an active *hstmt*. If an application does not call **SQLSetCursorName**, the driver generates cursor names as needed for SQL statement processing.

Syntax RETCODE **SQLSetCursorName**(*hstmt*, *szCursor*, *cbCursor*)

The **SQLSetCursorName** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UCHAR FAR *	<i>szCursor</i>	Input	Cursor name.
SWORD	<i>cbCursor</i>	Input	Length of <i>szCursor</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLSetCursorName** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLSetCursorName** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
24000	Invalid cursor state	The statement corresponding to <i>hstmt</i> was already in an executed or cursor-positioned state.
34000	Invalid cursor name	The cursor name specified by the argument <i>szCursor</i> was invalid. For example, the cursor name exceeded the maximum length as defined by the driver.
3C000	Duplicate cursor name	The cursor name specified by the argument <i>szCursor</i> already exists.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
SQLSTATE	Error	Description

	S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
	S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
	S1009	Invalid argument value	(DM) The argument <i>szCursor</i> was a null pointer.
	S1010	Function sequence error	(DM) An asynchronously executing function was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
	S1090	Invalid string or buffer length	(DM) The argument <i>cbCursor</i> was less than 0, but not equal to SQL_NTS.
Comments	<p>The only ODBC SQL statements that use a cursor name are a positioned update and delete (for example, UPDATE <i>table-name</i> ...WHERE CURRENT OF <i>cursor-name</i>). If the application does not call SQLSetCursorName to define a cursor name, on execution of a SELECT statement the driver generates a name that begins with the letters SQL_CUR and does not exceed 18 characters in length.</p> <p>All cursor names within the <i>hdbc</i> must be unique. The maximum length of a cursor name is defined by the driver. For maximum interoperability, it is recommended that applications limit cursor names to no more than 18 characters.</p> <p>A cursor name that is set either explicitly or implicitly remains set until the <i>hstmt</i> with which it is associated is dropped, using SQLFreeStmt with the SQL_DROP option.</p>		
Code Example	<p>In the following example, an application uses SQLSetCursorName to set a cursor name for an <i>hstmt</i>. It then uses that <i>hstmt</i> to retrieve results from the EMPLOYEE table. Finally, it performs a positioned update to change the name of 25-year-old John Smith to John D. Smith. Note that the application uses different <i>hstmts</i> for the SELECT and UPDATE statements.</p>		

For more code examples, see **SQLSetPos**.

```
#define NAME_LEN 30

HSTMT    hstmtSelect,
HSTMT    hstmtUpdate;
UCHAR    szName[NAME_LEN];
SWORD    sAge;
SDWORD   cbName;
SDWORD   cbAge;

/* Allocate the statements and set the cursor name */

SQLAllocStmt(hdbc, &hstmtSelect);
SQLAllocStmt(hdbc, &hstmtUpdate);
SQLSetCursorName(hstmtSelect, "C1", SQL_NTS);

/* SELECT the result set and bind its columns to local storage */

SQLExecDirect(hstmtSelect,
              "SELECT NAME, AGE FROM EMPLOYEE FOR UPDATE",
              SQL_NTS);
SQLBindCol(hstmtSelect, 1, SQL_C_CHAR, szName, NAME_LEN, &cbName);
SQLBindCol(hstmtSelect, 2, SQL_C_SSHORT, &sAge, 0, &cbAge);

/* Read through the result set until the cursor is      */
/* positioned on the row for the 25-year-old John Smith */

do
    retcode = SQLFetch(hstmtSelect);
while ((retcode == SQL_SUCCESS || retcode == SQL_SUCCESS_WITH_INFO) &&
      (strcmp(szName, "Smith, John") != 0 || sAge != 25));

/* Perform a positioned update of John Smith's name */

if (retcode == SQL_SUCCESS || retcode == SQL_SUCCESS_WITH_INFO) {
    SQLExecDirect(hstmtUpdate,
                  "UPDATE EMPLOYEE SET NAME=\\"Smith, John D.\\" WHERE CURRENT OF C1",
                  SQL_NTS);
}
```

**Related
Functions****For information about**

See

Executing an SQL statement

SQLExecDirect

Executing a prepared SQL statement

SQLExecute

Returning a cursor name

SQLGetCursorName

Setting cursor scrolling options

SQLSetScrollOptions
(extension)

SQLSetParam

ODBC

Deprecated: In ODBC 2.0, the ODBC 1.0 function **SQLSetParam** has been replaced by **SQLBindParameter**. For more information, see **SQLBindParameter**.

SQLSetPos



Extension Level 2 **SQLSetPos** sets the cursor position in a rowset and allows an application to refresh, update, delete, or add data to the rowset.

Syntax RETCODE **SQLSetPos**(*hstmt*, *irrow*, *fOption*, *fLock*)

The **SQLSetPos** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UWORD	<i>irrow</i>	Input	Position of the row in the rowset on which to perform the operation specified with the <i>fOption</i> argument. If <i>irrow</i> is 0, the operation applies to every row in the rowset. For additional information, see "Comments."
UWORD	<i>fOption</i>	Input	Operation to perform: SQL_POSITION SQL_REFRESH SQL_UPDATE SQL_DELETE SQL_ADD For more information, see "Comments."
UWORD	<i>fLock</i>	Input	Specifies how to lock the row after performing the operation specified in the <i>fOption</i> argument. SQL_LOCK_NO_CHANGE SQL_LOCK_EXCLUSIVE SQL_LOCK_UNLOCK For more information, see "Comments."

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_NEED_DATA, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLSetPos** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLSetPos** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The

return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01004	Data truncated	<p>The argument <i>fOption</i> was SQL_ADD or SQL_UPDATE and the value specified for a character or binary column exceeded the maximum length of the associated table column. (Function returns SQL_SUCCESS_WITH_INFO.)</p> <p>The argument <i>fOption</i> was SQL_ADD or SQL_UPDATE and the fractional part of the value specified for a numeric column was truncated. (Function returns SQL_SUCCESS_WITH_INFO.)</p> <p>The argument <i>fOption</i> was SQL_ADD or SQL_UPDATE and a timestamp value specified for a date or time column was truncated. (Function returns SQL_SUCCESS_WITH_INFO.)</p>
01S01	Error in row	The <i>irow</i> argument was 0 and an error occurred in one or more rows while performing the operation specified with the <i>fOption</i> argument. (Function returns SQL_SUCCESS_WITH_INFO.)
01S03	No rows updated or deleted	The argument <i>fOption</i> was SQL_UPDATE or SQL_DELETE and no rows were updated or deleted. (Function returns SQL_SUCCESS_WITH_INFO.)
01S04	More than one row updated or deleted	The argument <i>fOption</i> was SQL_UPDATE or SQL_DELETE and more than one row was updated or deleted. (Function returns SQL_SUCCESS_WITH_INFO.)
21S02	Degree of derived table does not match column list	The argument <i>fOption</i> was SQL_ADD or SQL_UPDATE and no columns were bound with SQLBindCol .
22003	Numeric value out of range	The argument <i>fOption</i> was SQL_ADD or SQL_UPDATE and the whole part of a numeric value was truncated.
22005	Error in assignment	The argument <i>fOption</i> was SQL_ADD or SQL_UPDATE and a value was incompatible

with the data type of the associated column.

SQLSTATE	Error	Description
22008	Datetime field overflow	The argument <i>fOption</i> was SQL_ADD or SQL_UPDATE and a date, time, or timestamp value was, respectively, an invalid date, time, or timestamp.
23000	Integrity constraint violation	<p>The argument <i>fOption</i> was SQL_ADD or SQL_UPDATE and a value was NULL for a column defined as NOT NULL in the associated column or some other integrity constraint was violated.</p> <p>The argument <i>fOption</i> was SQL_ADD and a column that was not bound with SQLBindCol is defined as NOT NULL or has no default.</p>
24000	Invalid cursor state	<p>(DM) The <i>hstmt</i> was in an executed state but no result set was associated with the <i>hstmt</i>.</p> <p>(DM) A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called.</p> <p>A cursor was open on the <i>hstmt</i> and SQLExtendedFetch had been called, but the cursor was positioned before the start of the result set or after the end of the result set.</p> <p>The argument <i>fOption</i> was SQL_DELETE, SQL_REFRESH, or SQL_UPDATE and the cursor was positioned before the start of the result set or after the end of the result set.</p>
42000	Syntax error or access violation	<p>The driver was unable to lock the row as needed to perform the operation requested in the argument <i>fOption</i>.</p> <p>The driver was unable to lock the row as requested in the argument <i>fLock</i>.</p>
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S0023	No default for column	<p>The <i>fOption</i> argument was SQL_ADD and a column that was not bound did not have a default value and could not be set to NULL.</p> <p>The <i>fOption</i> argument was SQL_ADD, the length specified in the <i>pcbValue</i> buffer</p>

		bound by SQLBindCol was SQL_IGNORE, and the column did not have a default value.
SQLSTATE	Error	Description
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1009	Invalid argument value	(DM) The value specified for the argument <i>fOption</i> was invalid. (DM) The value specified for the argument <i>fLock</i> was invalid. The argument <i>irow</i> was greater than the number of rows in the rowset and the <i>fOption</i> argument was not SQL_ADD. The value specified for the argument <i>fOption</i> was SQL_ADD, SQL_UPDATE, or SQL_DELETE, the value specified for the argument <i>fLock</i> was SQL_LOCK_NO_CHANGE, and the SQL_CONCURRENCY statement option was SQL_CONCUR_READ_ONLY.
SQLSTATE	Error	Description
S1010	Function sequence error	(DM) The specified <i>hstmt</i> was not in an executed state. The function was called without first calling SQLExecDirect , SQLExecute , or a catalog function. (DM) An asynchronously executing function (not this one) was called for the

		<p><i>hstmt</i> and was still executing when this function was called.</p> <p>(DM) SQLExecute, SQLExecDirect, or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.</p>
S1090	Invalid string or buffer length	<p>The <i>fOption</i> argument was SQL_ADD or SQL_UPDATE, a data value was a null pointer, and the column length value was not 0, SQL_DATA_AT_EXEC, SQL_IGNORE, SQL_NULL_DATA, or less than or equal to SQL_LEN_DATA_AT_EXEC_OFFSET.</p> <p>The <i>fOption</i> argument was SQL_ADD or SQL_UPDATE, a data value was not a null pointer, and the column length value was less than 0, but not equal to SQL_DATA_AT_EXEC, SQL_IGNORE, SQL_NTS, or SQL_NULL_DATA, or less than or equal to SQL_LEN_DATA_AT_EXEC_OFFSET.</p>
S1107	Row value out of range	The value specified for the argument <i>irow</i> was greater than the number of rows in the rowset and the <i>fOption</i> argument was not SQL_ADD .
S1109	Invalid cursor position	<p>The cursor associated with the <i>hstmt</i> was defined as forward only, so the cursor could not be positioned within the rowset. See the description for the SQL_CURSOR_TYPE option in SQLSetStmtOption.</p> <p>The <i>fOption</i> argument was SQL_REFRESH, SQL_UPDATE, or SQL_DELETE and the value in the <i>rgfRowStatus</i> array for the row specified by the <i>irow</i> argument was SQL_ROW_DELETED or SQL_ROW_ERROR.</p>

SQLSTATE	Error	Description
S1C00	Driver not capable	The driver or data source does not support the operation requested in the <i>fOption</i> argument or the <i>fLock</i> argument.
S1T00	Timeout expired	The timeout period expired before the

data source returned the result set. The timeout period is set through **SQLSetStmtOption**, **SQL_QUERY_TIMEOUT**.

Comments

***irow* Argument**

The *irow* argument specifies the number of the row in the rowset on which to perform the operation specified by the *fOption* argument. If *irow* is 0, the operation applies to every row in the rowset. Except for the **SQL_ADD** operation, *irow* must be a value from 0 to the number of rows in the rowset. For the **SQL_ADD** operation, *irow* can be any value; generally it is either 0 (to add as many rows as there are in the rowset) or the number of rows in the rowset plus 1 (to add the data from an extra row of buffers allocated for this purpose).

Note: In the C language, arrays are 0-based, while the *irow* argument is 1-based. For example, to update the fifth row of the rowset, an application modifies the rowset buffers at array index 4, but specifies an *irow* of 5.

All operations except for **SQL_ADD** position the cursor on the row specified by *irow*; the **SQL_ADD** operation does not change the cursor position. The following operations require a cursor position:

- Positioned update and delete statements.
- Calls to **SQLGetData**.
- Calls to **SQLSetPos** with the **SQL_DELETE**, **SQL_REFRESH**, and **SQL_UPDATE** options.

For example, if the cursor is positioned on the second row of the rowset, a positioned delete statement deletes that row; if it is positioned on the entire rowset (*irow* is 0), a positioned delete statement deletes every row in the rowset.

An application can specify a cursor position when it calls **SQLSetPos**. Generally, it calls **SQLSetPos** with the **SQL_POSITION** or **SQL_REFRESH** operation to position the cursor before executing a positioned update or delete statement or calling **SQLGetData**.

***fOption* Argument**

The *fOption* argument supports the following operations. To determine which options are supported by a data source, an application calls **SQLGetInfo** with the SQL_POS_OPERATIONS information type.

<i>fOption</i> Argument	Operation
SQL_POSITION	The driver positions the cursor on the row specified by <i>irow</i> . This is the same as the FALSE value of this argument in ODBC 1.0.
SQL_REFRESH	The driver positions the cursor on the row specified by <i>irow</i> and refreshes data in the rowset buffers for that row. For more information about how the driver returns data in the rowset buffers, see the descriptions of row-wise and column-wise binding in SQLExtendedFetch . This is the same as the TRUE value of this argument in ODBC 1.0.
SQL_UPDATE	The driver positions the cursor on the row specified by <i>irow</i> and updates the underlying row of data with the values in the rowset buffers (the <i>rgbValue</i> argument in SQLBindCol). It retrieves the lengths of the data from the number-of-bytes buffers (the <i>pcbValue</i> argument in SQLBindCol). If the length of any column is SQL_IGNORE, the column is not updated. After updating the row, the driver changes the <i>rgfRowStatus</i> array specified in SQLExtendedFetch to SQL_ROW_UPDATED.
SQL_DELETE	The driver positions the cursor on the row specified by <i>irow</i> and deletes the underlying row of data. It changes the <i>rgfRowStatus</i> array specified in SQLExtendedFetch to SQL_ROW_DELETED. After the row has been deleted, positioned update and delete statements, calls to SQLGetData and calls to SQLSetPos with <i>fOption</i> set to anything except SQL_POSITION are not valid for the row. Whether the row remains visible depends on the cursor type. For example, deleted rows are visible to static and keyset-driven cursors but invisible to dynamic cursors.
<i>fOption</i> Argument	Operation
SQL_ADD	The driver adds a new row of data to the data source. Where the row is added to the data source and whether it is visible in the result set is driver-defined.

The driver retrieves the data from the rowset buffers (the *rgbValue* argument in **SQLBindCol**) according to the value of the *row* argument. It retrieves the lengths of the data from the number-of-bytes buffers (the *pcbValue* argument in **SQLBindCol**). Generally, the application allocates an extra row of buffers for this purpose.

For columns not bound to the rowset buffers, the driver uses default values (if they are available) or NULL values (if default values are not available). For columns with a length of SQL_IGNORE, the driver uses default values.

If *row* is less than or equal to the rowset size, the driver changes the *rgfRowStatus* array specified in **SQLExtendedFetch** to SQL_ROW_ADDED after adding the row. At this point, the rowset buffers do not match the cursors for the row. To restore the rowset buffers to match the data in the cursor, an application calls **SQLSetPos** with the SQL_REFRESH option.

This operation does not affect the cursor position.

***fLock* Argument**

The *fLock* argument provides a way for applications to control concurrency and simulate transactions on data sources that do not support them. Generally, data sources that support concurrency levels and transactions will only support the SQL_LOCK_NO_CHANGE value of the *fLock* argument.

The *fLock* argument specifies the lock state of the row after **SQLSetPos** has been executed. To simulate a transaction, an application uses the SQL_LOCK_RECORD macro to lock each of the rows in the transaction. It then uses the SQL_UPDATE_RECORD or SQL_DELETE_RECORD macro to update or delete each row; the driver may temporarily change the lock state of the row while performing the operation specified by the *fOption* argument. Finally, it uses the SQL_LOCK_RECORD macro to unlock each row. For an example of how an application might do this, see the second code example. Note that if the driver is unable to lock the row either to perform the requested operation or to satisfy the *fLock* argument, it returns SQL_ERROR and SQLSTATE 42000 (Syntax error or access violation).

Although the *fLock* argument is specified for an *hstmt*, the lock accords the same privileges to all *hstmts* on the connection. In particular, a lock that is acquired by one *hstmt* on a connection can be unlocked by a different *hstmt* on the same connection.

A row locked through **SQLSetPos** remains locked until the application calls **SQLSetPos** for the row with *fLock* set to SQL_LOCK_UNLOCK or the application calls **SQLFreeStmt** with the SQL_CLOSE or SQL_DROP option.

The *fLock* argument supports the following types of locks. To determine which locks are supported by a data source, an application calls **SQLGetInfo** with the SQL_LOCK_TYPES information type.

<i>fLock</i> Argument	Lock Type
SQL_LOCK_NO_CHANGE	The driver or data source ensures that the row is in the same locked or unlocked state as it was before SQLSetPos was called. This value of <i>fLock</i> allows data sources that do not support explicit row-level locking to use whatever locking is required by the current concurrency and transaction isolation levels. This is the same as the FALSE value of the <i>fLock</i> argument in ODBC 1.0.
SQL_LOCK_EXCLUSIVE	The driver or data source locks the row exclusively. An <i>hstmt</i> on a different <i>hdbc</i> or in a different application cannot be used to acquire any locks on the row. This is the same as the TRUE value of the <i>fLock</i> argument in ODBC 1.0.
SQL_LOCK_UNLOCK	The driver or data source unlocks the row.

For the add, update, and delete operations in **SQLSetPos**, the application uses the *fLock* argument as follows:

- To guarantee that a row does not change after it is retrieved, an application calls **SQLSetPos** with *fOption* set to SQL_REFRESH and *fLock* set to SQL_LOCK_EXCLUSIVE.
- If the application sets *fLock* to SQL_LOCK_NO_CHANGE, the driver guarantees an update, or delete operation will succeed only if the application specified SQL_CONCUR_LOCK for the SQL_CONCURRENCY statement option.
- If the application specifies SQL_CONCUR_ROWVER or SQL_CONCUR_VALUES for the SQL_CONCURRENCY statement

option, the driver compares row versions or values and rejects the operation if the row has changed since the application fetched the row.

- If the application specifies `SQL_CONCUR_READ_ONLY` for the `SQL_CONCURRENCY` statement option, the driver rejects any update or delete operation.

For more information about the `SQL_CONCURRENCY` statement option, see **SQLSetStmtOption**.

Using SQLSetPos

Before an application calls **SQLSetPos**, it must:

1. If the application will call **SQLSetPos** with *fOption* set to `SQL_ADD` or `SQL_UPDATE`, call **SQLBindCol** for each column to specify its data type and associate storage for the column's data and length.
2. Call **SQLExecDirect**, **SQLExecute**, or a catalog function to create a result set.
3. Call **SQLExtendedFetch** to retrieve the data.

To delete data with **SQLSetPos**, an application:

- Calls **SQLSetPos** with *irow* set to the number of the row to delete.

An application can pass the value for a column either in the *rgbValue* buffer or with one or more calls to **SQLPutData**. Columns whose data is passed with **SQLPutData** are known as *data-at-execution* columns. These are commonly used to send data for `SQL_LONGVARIABLE` and `SQL_LONGVARCHAR` columns and can be mixed with other columns.

To update or add data with **SQLSetPos**, an application:

1. Places values in the *rgbValue* and *pcbValue* buffers bound with **SQLBindCol**:
 - For normal columns, the application places the new column value in the *rgbValue* buffer and the length of that value in the *pcbValue* buffer. If the row is being updated and the column is not to be changed, the application places `SQL_IGNORE` in the *pcbValue* buffer.
 - For data-at-execution columns, the application places an application-defined value, such as the column number, in the *rgbValue* buffer. The value can be used later to identify the column.

It places the result of the `SQL_LEN_DATA_AT_EXEC(length)` macro in the *pcbValue* buffer. If the SQL data type of the column is

SQL_LONGVARBINARY, SQL_LONGVARCHAR, or a long, data source-specific data type and the driver returns "Y" for the SQL_NEED_LONG_DATA_LEN information type in **SQLGetInfo**, *length* is the number of bytes of data to be sent for the parameter; otherwise, it must be a nonnegative value and is ignored.

2. Calls **SQLSetPos** or uses an **SQLSetPos** macro to update or add the row of data.
 - If there are no data-at-execution columns, the process is complete.
 - If there are any data-at-execution columns, the function returns SQL_NEED_DATA.
3. Calls **SQLParamData** to retrieve the address of the *rgbValue* buffer for the first data-at-execution column to be processed. The application retrieves the application-defined value from the *rgbValue* buffer.

Note: Although data-at-execution parameters are similar to data-at-execution columns, the value returned by **SQLParamData** is different for each.

Data-at-execution parameters are parameters in an SQL statement for which data will be sent with **SQLPutData** when the statement is executed with **SQLExecDirect** or **SQLExecute**. They are bound with **SQLBindParameter**. The value returned by **SQLParamData** is a 32-bit value passed to **SQLBindParameter** in the *rgbValue* argument.

Data-at-execution columns are columns in a rowset for which data will be sent with **SQLPutData** when a row is updated or added with **SQLSetPos**. They are bound with **SQLBindCol**. The value returned by **SQLParamData** is the address of the row in the *rgbValue* buffer that is being processed.

4. Calls **SQLPutData** one or more times to send data for the column. More than one call is needed if the data value is larger than the *rgbValue* buffer specified in **SQLPutData**; note that multiple calls to **SQLPutData** for the same column are allowed only when sending character C data to a column with a character, binary, or data source-specific data type or when sending binary C data to a column with a character, binary, or data source-specific data type.
5. Calls **SQLParamData** again to signal that all data has been sent for the column.
 - If there are more data-at-execution columns, **SQLParamData** returns SQL_NEED_DATA and the address of the *rgbValue* buffer for the next data-at-execution column to be processed. The application repeats steps 4 and 5.
 - If there are no more data-at-execution columns, the process is complete. If

the statement was executed successfully, **SQLParamData** returns SQL_SUCCESS or SQL_SUCCESS_WITH_INFO; if the execution failed, it returns SQL_ERROR. At this point, **SQLParamData** can return any SQLSTATE that can be returned by **SQLSetPos**.

After **SQLSetPos** returns SQL_NEED_DATA, and before data is sent for all data-at-execution columns, the operation is canceled, or an error occurs in **SQLParamData** or **SQLPutData**, the application can only call **SQLCancel**, **SQLGetFunctions**, **SQLParamData**, or **SQLPutData** with the *hstmt* or the *hdbc* associated with the *hstmt*. If it calls any other function with the *hstmt* or the *hdbc* associated with the *hstmt*, the function returns SQL_ERROR and SQLSTATE S1010 (Function sequence error).

If the application calls **SQLCancel** while the driver still needs data for data-at-execution columns, the driver cancels the operation; the application can then call **SQLSetPos** again; canceling does not affect the cursor state or the current cursor position. If the application calls **SQLParamData** or **SQLPutData** after canceling the operation, the function returns SQL_ERROR and SQLSTATE S1008 (Operation canceled).

Performing Bulk Operations

If the *irow* argument is 0, the driver performs the operation specified in the *fOption* argument for every row in the rowset. If an error occurs that pertains to the entire rowset, such as SQLSTATE S1T00 (Timeout expired), the driver returns SQL_ERROR and the appropriate SQLSTATE. The contents of the rowset buffers are undefined and the cursor position is unchanged.

If an error occurs that pertains to a single row, the driver:

- Sets the element in the *rgfRowStatus* array for the row to SQL_ROW_ERROR.
- Posts SQLSTATE 01S01 (Error in row) in the error queue.
- Posts one or more additional SQLSTATES for the error after SQLSTATE 01S01 (Error in row) in the error queue.

After it has processed the error or warning, the driver continues the operation for the remaining rows in the rowset and returns SQL_SUCCESS_WITH_INFO. Thus, for each row that returned an error, the error queue contains SQLSTATE 01S01 (Error in row) followed by zero or more additional SQLSTATES.

If the driver returns any warnings, such as SQLSTATE 01004 (Data truncated), it returns warnings that apply to the entire rowset or to unknown rows in the rowset before it returns the error information that applies to specific rows. It returns warnings for specific rows along with any other error information about those rows.

SQLSetPos Macros

As an aid to programming, the following macros for calling **SQLSetPos** are defined in the SQLEXT.H file.

Macro name	Function call
SQL_POSITION_TO(<i>hstmt</i> , <i>irow</i>)	SQLSetPos (<i>hstmt</i> , <i>irow</i> , SQL_POSITION, SQL_LOCK_NO_CHANGE)
SQL_LOCK_RECORD(<i>hstmt</i> , <i>irow</i> ,	SQLSetPos (<i>hstmt</i> , <i>irow</i> , SQL_POSITION,

<i>fLock)</i>	<i>fLock)</i>
SQL_REFRESH_RECORD(<i>hstmt</i> , <i>irow</i> , <i>fLock</i>)	SQLSetPos (<i>hstmt</i> , <i>irow</i> , SQL_REFRESH, <i>fLock</i>)
SQL_UPDATE_RECORD(<i>hstmt</i> , <i>irow</i>)	SQLSetPos (<i>hstmt</i> , <i>irow</i> , SQL_UPDATE, SQL_LOCK_NO_CHANGE)
SQL_DELETE_RECORD(<i>hstmt</i> , <i>irow</i>)	SQLSetPos (<i>hstmt</i> , <i>irow</i> , SQL_DELETE, SQL_LOCK_NO_CHANGE)
SQL_ADD_RECORD(<i>hstmt</i> , <i>irow</i>)	SQLSetPos (<i>hstmt</i> , <i>irow</i> , SQL_ADD, SQL_LOCK_NO_CHANGE)

Code Example In the following example, an application allows a user to browse the EMPLOYEE table and update employee birthdays. The cursor is keyset-driven with a rowset size of 20 and uses optimistic concurrency control comparing row versions. After each rowset is fetched, the application prints them and allows the user to select and update an employee's birthday. The application uses **SQLSetPos** to position the cursor on the selected row and performs a positioned update of the row. (Error handling is omitted for clarity.)

```
#define ROWS 20
#define NAME_LEN 30
#define BDAY_LEN 11

UCHAR  szName[ROWS][NAME_LEN], szBirthday[ROWS][BDAY_LEN], szReply[3];
SDWORD cbName[ROWS], cbBirthday[ROWS];
UWORD  rgfRowStatus[ROWS];
UDWORD crow, irow;
HSTMT  hstmtS, hstmtU;

SQLSetStmtOption(hstmtS, SQL_CONCURRENCY, SQL_CONCUR_ROWVER);
SQLSetStmtOption(hstmtS, SQL_CURSOR_TYPE, SQL_CURSOR_KEYSET_DRIVEN);
SQLSetStmtOption(hstmtS, SQL_ROWSET_SIZE, ROWS);
SQLSetCursorName(hstmtS, "C1", SQL_NTS);
SQLExecDirect(hstmtS,
    "SELECT NAME, BIRTHDAY FROM EMPLOYEE FOR UPDATE OF BIRTHDAY",
    SQL_NTS);
SQLBindCol(hstmtS, 1, SQL_C_CHAR, szName, NAME_LEN, cbName);
SQLBindCol(hstmtS, 2, SQL_C_CHAR, szBirthday, BDAY_LEN,
    cbBirthday);

while (SQLExtendedFetch(hstmtS, FETCH_NEXT, 0, &crow, rgfRowStatus) !=
    SQL_ERROR) {
    for (irow = 0; irow < crow; irow++) {
        if (rgfRowStatus[irow] != SQL_ROW_DELETED)
```

```

        printf("%d %-*s %*s\n", irow, NAME_LEN-1, szName[irow],
            BDAY_LEN-1, szBirthday[irow]);
    }
    while (TRUE) {
        printf("\nRow number to update?");
        gets(szReply);
        irow = atoi(szReply);
        if (irow > 0 && irow <= crow) {
            printf("\nNew birthday?");
            gets(szBirthday[irow-1]);
            SQLSetPos(hstmtS, irow, SQL_POSITION, SQL_LOCK_NO_CHANGE);
            SQLPrepare(hstmtU,
                "UPDATE EMPLOYEE SET BIRTHDAY=? WHERE CURRENT OF C1",
                SQL_NTS);
            SQLBindParameter(hstmtU, 1, SQL_PARAM_INPUT,
                SQL_C_CHAR, SQL_DATE,
                BDAY_LEN, 0, szBirthday, 0, NULL);
            SQLExecute(hstmtU);
        } else if (irow == 0) {
            break;
        }
    }
}

```

In the following code fragment, an application simulates a transaction for rows 1 and 2. It locks the rows, updates them, then unlocks them. The code uses the **SQLSetPos** macros.

```

/* Lock rows 1 and 2 */

SQL_LOCK_RECORD(hstmt, 1, SQL_LOCK_EXCLUSIVE);
SQL_LOCK_RECORD(hstmt, 2, SQL_LOCK_EXCLUSIVE);

/* Modify the rowset buffers for rows 1 and 2 (not shown). */
/* Update rows 1 and 2. */
SQL_UPDATE_RECORD(hstmt, 1);
SQL_UPDATE_RECORD(hstmt, 2);

/* Unlock rows 1 and 2 */
SQL_LOCK_RECORD(hstmt, 1, SQL_LOCK_UNLOCK);
SQL_LOCK_RECORD(hstmt, 2, SQL_LOCK_UNLOCK);

```

**Related
Functions**

For information about

See

Part 2 Developing Applications

Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
Setting a statement option	SQLSetStmtOption (extension)

!Unexpected End of Expression

SQLSetScrollOptions

ODBC

Extension Level 2 **SQLSetScrollOptions** sets options that control the behavior of cursors associated with an *hstmt*. **SQLSetScrollOptions** allows the application to specify the type of cursor behavior desired in three areas: concurrency control, sensitivity to changes made by other transactions, and rowset size.

Note: In ODBC 2.0, **SQLSetScrollOptions** has been superseded by the SQL_CURSOR_TYPE, SQL_CONCURRENCY, SQL_KEYSET_SIZE, and SQL_ROWSET_SIZE statement options. ODBC 2.0 drivers must support this function for backwards compatibility; ODBC 2.0 applications should only call this function in ODBC 1.0 drivers.

If an application calls **SQLSetScrollOptions**, a driver must be able to return the values of the aforementioned statement options with **SQLGetStmtOption**. For more information, see **SQLGetStmtOption**.

Syntax

RETCODE **SQLSetScrollOptions**(*hstmt*, *fConcurrency*, *crowKeyset*, *crowRowset*)

The **SQLSetScrollOptions** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UWORD	<i>fConcurrency</i>	Input	Specifies concurrency control for the cursor and must be one of the following values: SQL_CONCUR_READ_ONLY: Cursor is read-only. No updates are allowed. SQL_CONCUR_LOCK: Cursor uses the lowest level of locking sufficient to ensure that the row can be updated. SQL_CONCUR_ROWVER: Cursor uses optimistic concurrency control, comparing row versions, such as SQLBase ROWID or Sybase TIMESTAMP. SQL_CONCUR_VALUES: Cursor uses optimistic concurrency control, comparing values.

Type	Argument	Use	Description
SDWORD	<i>crowKeyset</i>	Input	<p>Number of rows for which to buffer keys. This value must be greater than or equal to <i>crowRowset</i> or one of the following values:</p> <p>SQL_SCROLL_FORWARD_ONLY: The cursor only scrolls forward.</p> <p>SQL_SCROLL_STATIC: The data in the result set is static.</p> <p>SQL_SCROLL_KEYSET_DRIVEN: The driver saves and uses the keys for every row in the result set.</p> <p>SQL_SCROLL_DYNAMIC: The driver sets <i>crowKeyset</i> to the value of <i>crowRowset</i>.</p> <p>If <i>crowKeyset</i> is a value greater than <i>crowRowset</i>, the value defines the number of rows in the keyset that are to be buffered by the driver. This reflects a mixed scrollable cursor; the cursor is keyset driven within the keyset and dynamic outside of the keyset.</p>
UWORD	<i>crowRowset</i>	Input	<p>Number of rows in a rowset. <i>crowRowset</i> defines the number of rows fetched by each call to SQLExtendedFetch; the number of rows that the application buffers.</p>

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLSetScrollOptions** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLSetScrollOptions** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATEs returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns

SQLDataSourceToDriver (Translation)

!Unexpected End of Expression

		SQL_SUCCESS_WITH_INFO.)
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1010	Function sequence error	(DM) The specified <i>hstmt</i> was in a prepared or executed state. The function must be called before calling SQLPrepare or SQLExecDirect . (DM) An asynchronously executing function was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1107	Row value out of range	(DM) The value specified for the argument <i>crowKeyset</i> was less than 1, but was not equal to SQL_SCROLL_FORWARD_ONLY, SQL_SCROLL_STATIC, SQL_SCROLL_KEYSET_DRIVEN, or SQL_SCROLL_DYNAMIC. (DM) The value specified for the argument <i>crowKeyset</i> is greater than 0, but less than <i>crowRowset</i> . (DM) The value specified for the argument <i>crowRowset</i> was 0.
SQLSTATE Error		Description
S1108	Concurrency option out of range	(DM) The value specified for the argument <i>fConcurrency</i> was not equal to SQL_CONCUR_READ_ONLY, SQL_CONCUR_LOCK, SQL_CONCUR_ROWVER, or

		SQL_CONCUR_VALUES.
S1C00	Driver not capable	<p>The driver or data source does not support the concurrency control option specified in the argument <i>fConcurrency</i>.</p> <p>The driver does not support the cursor model specified in the argument <i>crowKeyset</i>.</p>
Comments	<p>If an application calls SQLSetScrollOptions for an <i>hstmt</i>, it must do so before it calls SQLPrepare or SQLExecDirect or creating a result set with a catalog function.</p> <p>The application must specify a buffer in a call to SQLBindCol that is large enough to hold the number of rows specified in <i>crowRowset</i>.</p> <p>If the application does not call SQLSetScrollOptions, <i>crowRowset</i> has a default value of 1, <i>crowKeyset</i> has a default value of SQL_SCROLL_FORWARD_ONLY, and <i>fConcurrency</i> equals SQL_CONCUR_READ_ONLY.</p> <p>For more information concerning scrollable cursors, see “Using Block and Scrollable Cursors” in Chapter 7, “Retrieving Results.”</p>	
Related Functions	For information about	See
<hr/>		
	Assigning storage for a column in a result set	SQLBindCol
	Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
	Positioning the cursor in a rowset	SQLSetPos (extension)
	Setting a statement option	SQLSetStmtOption

SQLSetStmtOption



Extension Level 1 **SQLSetStmtOption** sets options related to an *hstmt*. To set an option for all statements associated with a specific *hdbc*, an application can call **SQLSetConnectOption**.

Syntax RETCODE **SQLSetStmtOption**(*hstmt*, *fOption*, *vParam*)

The **SQLSetStmtOption** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UWORD	<i>fOption</i>	Input	Option to set, listed in "Comments."
UDWORD	<i>vParam</i>	Input	Value associated with <i>fOption</i> . Depending on the value of <i>fOption</i> , <i>vParam</i> will be a 32-bit integer value or point to a null-terminated character string.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLSetStmtOption** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLSetStmtOption** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
01S02	Option value changed	The driver did not support the specified value of the <i>vParam</i> argument and substituted a similar value. (Function returns SQL_SUCCESS_WITH_INFO.)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.

SQLSTATE	Error	Description
24000	Invalid cursor state	The <i>fOption</i> was SQL_CONCURRENCY, SQL_CURSOR_TYPE, SQL_SIMULATE_CURSOR, or SQL_USE_BOOKMARKS and the cursor was open.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable 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 was specified for the argument <i>vParam</i> . (The Driver Manager returns this SQLSTATE only for statement options that accept a discrete set of values, such as SQL_ASYNC_ENABLE. For all other statement options, the driver must verify the value of the argument <i>vParam</i> .)
S1010	Function sequence error	(DM) An asynchronously executing function was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1011	Operation invalid at this time	The <i>fOption</i> was SQL_CONCURRENCY, SQL_CURSOR_TYPE, SQL_SIMULATE_CURSOR, or SQL_USE_BOOKMARKS and the statement was prepared.
S1092	Option type out of range	(DM) The value specified for the argument <i>fOption</i> was in the block of numbers reserved for ODBC connection and statement options, but was not valid for the version of ODBC supported by the driver.

SQLDataSourceToDriver (Translation)

SQLSTATE	Error	Description
S1C00	Driver not capable	<p>The value specified for the argument <i>fOption</i> was a valid ODBC statement option for the version of ODBC supported by the driver, but was not supported by the driver.</p> <p>The value specified for the argument <i>fOption</i> was in the block of numbers reserved for driver-specific connection and statement options, but was not supported by the driver.</p>
Comments	<p>Statement options for an <i>hstmt</i> remain in effect until they are changed by another call to SQLSetStmtOption or the <i>hstmt</i> 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.</p> <p>Some statement options support substitution of a similar value if the data source does not support the specified value of <i>vParam</i>. In such cases, the driver returns SQL_SUCCESS_WITH_INFO and SQLSTATE 01S02 (Option value changed). For example, if <i>fOption</i> is SQL_CONCURRENCY, <i>vParam</i> is SQL_CONCUR_ROWVER, and the data source does not support this, the driver substitutes SQL_CONCUR_VALUES. To determine the substituted value, an application calls SQLGetStmtOption.</p> <p>The currently defined options and the version of ODBC in which they were introduced are shown below; it is expected that more will be defined to take advantage of different data sources. Options from 0 to 999 are reserved by ODBC; driver developers must reserve values greater than or equal to SQL_CONNECT_OPT_DVR_START for driver-specific use.</p> <p>The format of information set with <i>vParam</i> depends on the specified <i>fOption</i>. SQLSetStmtOption accepts option information in one of two different formats: a null-terminated character string or a 32-bit integer value. The format of each is noted in the option's description. This format applies to the information returned for each option in SQLGetStmtOption. Character strings pointed to by the <i>vParam</i> argument of SQLSetStmtOption have a maximum length of SQL_MAX_OPTION_STRING_LENGTH bytes (excluding the null termination byte).</p>	
<i>fOption</i>	<i>vParam</i> Contents	
SQL_ASYNC_ENABLE	A 32-bit integer value that specifies whether a function called with the	

(ODBC 1.0)

specified *hstmt* is executed asynchronously:

SQL_ASYNC_ENABLE_OFF = Off (the default)

SQL_ASYNC_ENABLE_ON = On

Once a function has been called asynchronously, no other functions can be called on the *hstmt* or the *hdbc* associated with the *hstmt* except for the original function, **SQLAllocStmt**, **SQLCancel**, or **SQLGetFunctions**, until the original function returns a code other than SQL_STILL_EXECUTING. Any other function called on the *hstmt* returns SQL_ERROR with an SQLSTATE of S1010 (Function sequence error). Functions can be called on other *hstmts*. For more information, see “Executing Functions Asynchronously” in Chapter 6.

The following functions can be executed asynchronously:

SQLColAttributes	SQLNumParams
SQLColumnPrivileges	SQLNumResultCols
SQLColumns	SQLParamData
SQLDescribeCol	SQLPrepare
SQLDescribeParam	SQLPrimaryKeys
SQLExecDirect	SQLProcedureColumns
SQLExecute	SQLProcedures
SQLExtendedFetch	SQLPutData
SQLFetch	SQLSetPos
SQLForeignKeys	SQLSpecialColumns
SQLGetData	SQLStatistics
SQLGetTypeInfo	SQLTablePrivileges
SQLMoreResults	SQLTables

SQL_BIND_TYPE
(ODBC 1.0)

A 32-bit integer value that sets the binding orientation to be used when **SQLExtendedFetch** is called on the associated *hstmt*. Column-wise binding is selected by supplying the defined constant SQL_BIND_BY_COLUMN for the argument *vParam*. Row-wise binding is selected by supplying a value for *vParam* specifying the length of a structure or an instance of a buffer into which result columns will be bound.

The length specified in *vParam* 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 will point to the beginning of the same column in the next row. When using the **sizeof** operator with structures or unions in ANSI C, this behavior is guaranteed.

Column-wise binding is the default binding orientation for **SQLExtendedFetch**.

*fOption**vParam* ContentsSQL_CONCURRENCY
(ODBC 2.0)

A 32-bit integer value that specifies the cursor concurrency:

SQL_CONCUR_READ_ONLY = Cursor is read-only. No updates are

SQLDriverToDataSource (Translation)

	<p>allowed.</p> <p>SQL_CONCUR_LOCK = Cursor uses the lowest level of locking sufficient to ensure that the row can be updated.</p> <p>SQL_CONCUR_ROWVER = Cursor uses optimistic concurrency control, comparing row versions, such as SQLBase ROWID or Sybase TIMESTAMP.</p> <p>SQL_CONCUR_VALUES = Cursor uses optimistic concurrency control, comparing values.</p> <p>The default value is SQL_CONCUR_READ_ONLY. This option cannot be specified for an open cursor and can also be set through the <i>fConcurrency</i> argument in SQLSetScrollOptions.</p> <p>If the specified concurrency is not supported by the data source, the driver substitutes a different concurrency and returns SQLSTATE 01S02 (Option value changed). For SQL_CONCUR_VALUES, the driver substitutes SQL_CONCUR_ROWVER, and vice versa. For SQL_CONCUR_LOCK, the driver substitutes, in order, SQL_CONCUR_ROWVER or SQL_CONCUR_VALUES.</p>
SQL_CURSOR_TYPE (ODBC 2.0)	<p>A 32-bit integer value that specifies the cursor type:</p> <p>SQL_CURSOR_FORWARD_ONLY = The cursor only scrolls forward.</p> <p>SQL_CURSOR_STATIC = The data in the result set is static.</p> <p>SQL_CURSOR_KEYSET_DRIVEN = The driver saves and uses the keys for the number of rows specified in the SQL_KEYSET_SIZE statement option.</p> <p>SQL_CURSOR_DYNAMIC = The driver only saves and uses the keys for the rows in the rowset.</p> <p>The default value is SQL_CURSOR_FORWARD_ONLY. This option cannot be specified for an open cursor and can also be set through the <i>crowKeyset</i> argument in SQLSetScrollOptions.</p> <p>If the specified cursor type is not supported by the data source, the driver substitutes a different cursor type and returns SQLSTATE 01S02 (Option value changed). For a mixed or dynamic cursor, the driver substitutes, in order, a keyset-driven or static cursor. For a keyset-driven cursor, the driver substitutes a static cursor.</p>
<i>fOption</i>	<i>vParam</i> Contents
SQL_KEYSET_SIZE (ODBC 2.0)	<p>A 32-bit integer value that specifies the number of rows in the keyset for a keyset-driven cursor. If the keyset size is 0 (the default), the cursor is fully keyset-driven. If the keyset size is greater than 0, the cursor is mixed (keyset-driven within the keyset and dynamic outside of the keyset). The default keyset size is 0.</p> <p>If the specified size exceeds the maximum keyset size, the driver substitutes that size and returns SQLSTATE 01S02 (Option value</p>

	changed).
	SQLExtendedFetch returns an error if the keyset size is greater than 0 and less than the rowset size.
SQL_MAX_LENGTH (ODBC 1.0)	<p>A 32-bit integer value that specifies the maximum amount of data that the driver returns from a character or binary column. If <i>vParam</i> is less than the length of the available data, SQLFetch or SQLGetData truncates the data and returns SQL_SUCCESS. If <i>vParam</i> is 0 (the default), the driver attempts to return all available data.</p> <p>If the specified length is less than the minimum amount of data that the data source can return (the minimum is 254 bytes on many data sources), or greater than the maximum amount of data that the data source can return, the driver substitutes that value and returns SQLSTATE 01S02 (Option value changed).</p> <p>This option is intended to reduce network traffic and should only be supported when the data source (as opposed to the driver) in a multiple-tier driver can implement it. To truncate data, an application should specify the maximum buffer length in the <i>cbValueMax</i> argument in SQLBindCol or SQLGetData.</p> <hr/> <p>Note: In ODBC 1.0, this statement option only applied to SQL_LONGVARCHAR and SQL_LONGVARBINARY columns.</p> <hr/>
SQL_MAX_ROWS (ODBC 1.0)	<p>A 32-bit integer value corresponding to the maximum number of rows to return to the application for a SELECT statement. If <i>vParam</i> equals 0 (the default), then the driver returns all rows.</p> <p>This option is intended to reduce network traffic. Conceptually, it is applied when the result set is created and limits the result set to the first <i>vParam</i> rows.</p> <p>If the specified number of rows exceeds the number of rows that can be returned by the data source, the driver substitutes that value and returns SQLSTATE 01S02 (Option value changed).</p>
fOption	vParam Contents
SQL_NOSCAN (ODBC 1.0)	<p>A 32-bit integer value that specifies whether the driver does not scan SQL strings for escape clauses:</p> <p>SQL_NOSCAN_OFF = The driver scans SQL strings for escape clauses (the default).</p> <p>SQL_NOSCAN_ON = The driver does not scan SQL strings for escape clauses. Instead, the driver sends the statement directly to the data source.</p>
SQL_QUERY_TIMEOUT (ODBC 1.0)	<p>A 32-bit integer value corresponding to the number of seconds to wait for an SQL statement to execute before returning to the application. If <i>vParam</i> equals 0 (the default), then there is no time out.</p> <p>If the specified timeout exceeds the maximum timeout in the data</p>
SQLDriverToDataSource (Translation)	

	<p>source or is smaller than the minimum timeout, the driver substitutes that value and returns SQLSTATE 01S02 (Option value changed).</p> <p>Note that the application need not call SQLFreeStmt with the SQL_CLOSE option to reuse the <i>hstmt</i> if a SELECT statement timed out.</p>
SQL_RETRIEVE_DATA (ODBC 2.0)	<p>A 32-bit integer value:</p> <p>SQL_RD_ON = SQLExtendedFetch retrieves data after it positions the cursor to the specified location. This is the default.</p> <p>SQL_RD_OFF = SQLExtendedFetch does not retrieve data after it positions the cursor.</p> <p>By setting SQL_RETRIEVE_DATA to SQL_RD_OFF, an application can verify if a row exists or retrieve a bookmark for the row without incurring the overhead of retrieving rows.</p>
SQL_ROWSET_SIZE (ODBC 2.0)	<p>A 32-bit integer value that specifies the number of rows in the rowset. This is the number of rows returned by each call to SQLExtendedFetch. The default value is 1.</p> <p>If the specified rowset size exceeds the maximum rowset size supported by the data source, the driver substitutes that value and returns SQLSTATE 01S02 (Option value changed).</p> <p>This option can be specified for an open cursor and can also be set through the <i>rowRowset</i> argument in SQLSetScrollOptions.</p>
<i>fOption</i>	<i>vParam Contents</i>
SQL_SIMULATE_CURSOR (ODBC 2.0)	<p>A 32-bit integer value that specifies whether drivers that simulate positioned update and delete statements guarantee that such statements affect only one single row.</p> <p>To simulate positioned update and delete statements, most drivers construct a searched UPDATE or DELETE statement containing a WHERE clause that specifies the value of each column in the current row. Unless these columns comprise a unique key, such a statement may affect more than one row.</p> <p>To guarantee that such statements affect only one row, the driver determines the columns in a unique key and adds these columns to the result set. If an application guarantees that the columns in the result set comprise a unique key, the driver is not required to do so. This may reduce execution time.</p> <p>SQL_SC_NON_UNIQUE = The driver does not guarantee that simulated positioned update or delete statements will affect only one row; it is the application's responsibility to do so. If a statement affects more than one row, SQLExecute or SQLExecDirect returns SQLSTATE 01000 (General warning).</p> <p>SQL_SC_TRY_UNIQUE = The driver attempts to guarantee that simulated positioned update or delete statements affect only one row.</p>

The driver always executes such statements, even if they might affect more than one row, such as when there is no unique key. If a statement affects more than one row, **SQLExecute** or **SQLExecDirect** returns SQLSTATE 01000 (General warning).

SQL_SC_UNIQUE = The driver guarantees that simulated positioned update or delete statements affect only one row. If the driver cannot guarantee this for a given statement, **SQLExecDirect** or **SQLPrepare** returns an error.

If the specified cursor simulation type is not supported by the data source, the driver substitutes a different simulation type and returns SQLSTATE 01S02 (Option value changed). For **SQL_SC_UNIQUE**, the driver substitutes, in order, **SQL_SC_TRY_UNIQUE** or **SQL_SC_NON_UNIQUE**. For **SQL_SC_TRY_UNIQUE**, the driver substitutes **SQL_SC_NON_UNIQUE**.

If a driver does not simulate positioned update and delete statements, it returns SQLSTATE S1C00 (Driver not capable).

SQL_USE_BOOKMARKS
(ODBC 2.0)

A 32-bit integer value that specifies whether an application will use bookmarks with a cursor:

SQL_UB_OFF = Off (the default)

SQL_UB_ON = On

To use bookmarks with a cursor, the application must specify this option with the **SQL_UB_ON** value before opening the cursor.

Code Example See **SQLExtendedFetch**.

**Related
Functions**

For information about

See

Canceling statement processing

SQLCancel

Returning the setting of a connection option

SQLGetConnectOption
(extension)

Returning the setting of a statement option

SQLGetStmtOption (extension)

Setting a connection option

SQLSetConnectOption
(extension)

SQLSpecialColumns



Extension Level 1 **SQLSpecialColumns** retrieves the following information about columns within a specified table:

- The optimal set of columns that uniquely identifies a row in the table.
- Columns that are automatically updated when any value in the row is updated by a transaction.

Syntax

RETCODE **SQLSpecialColumns**(*hstmt*, *fColType*, *szTableQualifier*, *cbTableQualifier*, *szTableOwner*, *cbTableOwner*, *szTableName*, *cbTableName*, *fScope*, *fNullable*)

The **SQLSpecialColumns** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UWORD	<i>fColType</i>	Input	Type of column to return. Must be one of the following values: SQL_BEST_ROWID: Returns the optimal column or set of columns that, by retrieving values from the column or columns, allows any row in the specified table to be uniquely identified. A column can be either a pseudocolumn specifically designed for this purpose (as in Oracle ROWID or Ingres TID) or the column or columns of any unique index for the table. SQL_ROWVER: Returns the column or columns in the specified table, if any, that are automatically updated by the data source when any value in the row is updated by any transaction (as in SQLBase ROWID or Sybase TIMESTAMP).
UCHAR FAR *	<i>szTableQualifier</i>	Input	Qualifier name for the table. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have qualifiers.

Type	Argument	Use	Description
SWORD	<i>cbTableQualifier</i>	Input	Length of <i>szTableQualifier</i> .
UCHAR FAR *	<i>szTableOwner</i>	Input	Owner name for the table. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have owners.
SWORD	<i>cbTableOwner</i>	Input	Length of <i>szTableOwner</i> .
UCHAR FAR *	<i>szTableName</i>	Input	Table name.
SWORD	<i>cbTableName</i>	Input	Length of <i>szTableName</i> .
UWORD	<i>fScope</i>	Input	Minimum required scope of the rowid. The returned rowid may be of greater scope. Must be one of the following: SQL_SCOPE_CURROW: The rowid is guaranteed to be valid only while positioned on that row. A later reselect using rowid may not return a row if the row was updated or deleted by another transaction. SQL_SCOPE_TRANSACTION: The rowid is guaranteed to be valid for the duration of the current transaction. SQL_SCOPE_SESSION: The rowid is guaranteed to be valid for the duration of the session (across transaction boundaries).
UWORD	<i>fNullable</i>	Input	Determines whether to return special columns that can have a NULL value. Must be one of the following: SQL_NO_NULLS: Exclude special columns that can have NULL values. SQL_NULLABLE: Return special columns even if they can have NULL values.

Returns

SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics

When **SQLSpecialColumns** returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, an associated **SQLSTATE** value may be obtained by calling **SQLError**. The following table lists the **SQLSTATE** values commonly returned by **SQLSpecialColumns** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of **SQLSTATE**s returned by the Driver Manager. The return code associated with each **SQLSTATE** value is **SQL_ERROR**, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO .)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called. A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.

SQLSTATE	Error	Description
S1010	Function sequence error	<p>(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called.</p> <p>(DM) SQLExecute, SQLExecDirect, or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.</p>
S1090	Invalid string or buffer length	<p>(DM) The value of one of the length arguments was less than 0, but not equal to SQL_NTS.</p> <p>The value of one of the length arguments exceeded the maximum length value for the corresponding qualifier or name. The maximum length of each qualifier or name may be obtained by calling SQLGetInfo with the <i>fInfoType</i> values: SQL_MAX_QUALIFIER_NAME_LEN, SQL_MAX_OWNER_NAME_LEN, or SQL_MAX_TABLE_NAME_LEN.</p>
S1097	Column type out of range	(DM) An invalid <i>fColType</i> value was specified.
S1098	Scope type out of range	(DM) An invalid <i>fScope</i> value was specified.
S1099	Nullable type out of range	(DM) An invalid <i>fNullable</i> value was specified.
S1C00	Driver not capable	<p>A table qualifier was specified and the driver or data source does not support qualifiers.</p> <p>A table owner was specified and the driver or data source does not support owners.</p> <p>The combination of the current settings of the SQL_CONCURRENCY and SQL_CURSOR_TYPE statement options was not supported by the driver or data source.</p>
S1T00	Timeout expired	<p>The timeout period expired before the data source returned the requested result set. The timeout period is set through SQLSetStmtOption, SQL_QUERY_TIMEOUT.</p>

Comments

SQLSpecialColumns is provided so that applications can provide their own custom scrollable-cursor functionality, similar to that provided by **SQLExtendedFetch** and **SQLSetStmtOption**.

When the *fColType* argument is **SQL_BEST_ROWID**, **SQLSpecialColumns** returns the column or columns that uniquely identify each row in the table. These columns can always be used in a *select-list* or **WHERE** clause. However, **SQLColumns** does not necessarily return these columns. For example, **SQLColumns** might not return the Oracle ROWID pseudo-column ROWID. If there are no columns that uniquely identify each row in the table, **SQLSpecialColumns** returns a rowset with no rows; a subsequent call to **SQLFetch** or **SQLExtendedFetch** on the *hstmt* returns **SQL_NO_DATA_FOUND**.

If the *fColType*, *fScope*, or *fNullable* arguments specify characteristics that are not supported by the data source, **SQLSpecialColumns** returns a result set with no rows (as opposed to the function returning **SQL_ERROR** with **SQLSTATE S1C00** (Driver not capable)). A subsequent call to **SQLFetch** or **SQLExtendedFetch** on the *hstmt* will return **SQL_NO_DATA_FOUND**.

SQLSpecialColumns returns the results as a standard result set, ordered by **SCOPE**. The following table lists the columns in the result set.

The lengths of **VARCHAR** columns shown in the table are maximums; the actual lengths depend on the data source. To determine the actual length of the **COLUMN_NAME** column, an application can call **SQLGetInfo** with the **SQL_MAX_COLUMN_NAME_LEN** option.

Column Name	Data Type	Comments
SCOPE	Smallint	Actual scope of the rowid. Contains one of the following values: SQL_SCOPE_CURROW SQL_SCOPE_TRANSACTION SQL_SCOPE_SESSION NULL is returned when <i>fColType</i> is SQL_ROWVER . For a description of each value, see the description of <i>fScope</i> in the "Syntax" section above.
COLUMN_NAME	Varchar(128) not NULL	Column identifier.

Column Name	Data Type	Comments
DATA_TYPE	Smallint not NULL	SQL data type. This can be an ODBC SQL data type or a driver-specific SQL data type. For a list of valid ODBC SQL data types, see "SQL Data Types" in Appendix D, "Data Types." For information about driver-specific SQL data types, see the driver's documentation.
TYPE_NAME	Varchar(128) not NULL	Data source-dependent data type name; for example, "CHAR", "VARCHAR", "LONG" or "CHAR () BYTE".
PRECISION	Integer	The precision of the column on the data source. NULL is returned for data types where precision is not applicable. For more information concerning precision, see "Precision, Scale, Length, and Display Size," in Appendix D, "Data Types."
LENGTH	Integer	The length in bytes of data transferred on an SQLGetData or SQLFetch operation if SQL_C_DEFAULT is specified. For numeric data, this size may be different than the size of the data stored on the data source. This value is the same as the PRECISION column for character or binary data. For more information, see "Precision, Scale, Length, and Display Size," in Appendix D, "Data Types."
SCALE	Smallint	The scale of the column on the data source. NULL is returned for data types where scale is not applicable. For more information concerning scale, see "Precision, Scale, Length, and Display Size," in Appendix D, "Data Types."
PSEUDO_COLUMN	Smallint	Indicates whether the column is a pseudo-column, such as Oracle ROWID: SQL_PC_UNKNOWN SQL_PC_PSEUDO SQL_PC_NOT_PSEUDO

Note: For maximum interoperability, pseudo-columns should not be quoted

SQLDriverToDataSource (Translation)

with the identifier quote character
returned by **SQLGetInfo**.

Note: The PSEUDO_COLUMN column was added in ODBC 2.0. ODBC 1.0 drivers might return a different, driver-specific column with the same column number.

Once the application retrieves values for SQL_BEST_ROWID, the application can use these values to reselect that row within the defined scope. The **SELECT** statement is guaranteed to return either no rows or one row.

If an application reselects a row based on the rowid column or columns and the row is not found, then the application can assume that the row was deleted or the rowid columns were modified. The opposite is not true: even if the rowid has not changed, the other columns in the row may have changed.

Columns returned for column type SQL_BEST_ROWID are useful for applications that need to scroll forwards and backwards within a result set to retrieve the most recent data from a set of rows. The column or columns of the rowid are guaranteed not to change while positioned on that row.

The column or columns of the rowid may remain valid even when the cursor is not positioned on the row; the application can determine this by checking the SCOPE column in the result set.

Columns returned for column type SQL_ROWVER are useful for applications that need the ability to check if any columns in a given row have been updated while the row was reselected using the rowid. For example, after reselecting a row using rowid, the application can compare the previous values in the SQL_ROWVER columns to the ones just fetched. If the value in a SQL_ROWVER column differs from the previous value, the application can alert the user that data on the display has changed.

Code Example For a code example of a similar function, see **SQLColumns**.

Related Functions

For information about

See

Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Returning the columns in a table or tables	SQLColumns (extension)
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)

SQLDriverToDataSource (Translation)

Fetching a row of data

SQLFetch

Returning the columns of a primary key

SQLPrimaryKeys (extension)

SQLStatistics



Extension Level 1 SQLStatistics retrieves a list of statistics about a single table and the indexes associated with the table. The driver returns the information as a result set.

Syntax RETCODE **SQLStatistics**(*hstmt*, *szTableQualifier*, *cbTableQualifier*, *szTableOwner*, *cbTableOwner*, *szTableName*, *cbTableName*, *fUnique*, *fAccuracy*)

The **SQLStatistics** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UCHAR FAR *	<i>szTableQualifier</i>	Input	Qualifier name. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have qualifiers.
SWORD	<i>cbTableQualifier</i>	Input	Length of <i>szTableQualifier</i> .
UCHAR FAR *	<i>szTableOwner</i>	Input	Owner name. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have owners.
SWORD	<i>cbTableOwner</i>	Input	Length of <i>szTableOwner</i> .
UCHAR FAR *	<i>szTableName</i>	Input	Table name.
SWORD	<i>cbTableName</i>	Input	Length of <i>szTableName</i> .
UWORD	<i>fUnique</i>	Input	Type of index: SQL_INDEX_UNIQUE or SQL_INDEX_ALL.
UWORD	<i>fAccuracy</i>	Input	The importance of the CARDINALITY and PAGES columns in the result set: SQL_ENSURE requests that the driver unconditionally retrieve the statistics. SQL_QUICK requests that the driver retrieve results only if they are readily available from the server. In this case, the driver does not ensure that the values are current.

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLStatistics** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLStatistics** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATEs returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called. A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
SQLSTATE	Error	Description
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function

		<p>was called again on the <i>hstmt</i>.</p> <p>The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.</p>
S1010	Function sequence error	<p>(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called.</p> <p>(DM) SQLExecute, SQLExecDirect, or SQLSetPos was called for the <i>hstmt</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.</p>
S1090	Invalid string or buffer length	<p>(DM) The value of one of the name length arguments was less than 0, but not equal to SQL_NTS.</p> <p>The value of one of the name length arguments exceeded the maximum length value for the corresponding qualifier or name.</p>
S1100	Uniqueness option type out of range	(DM) An invalid <i>fUnique</i> value was specified.
S1101	Accuracy option type out of range	(DM) An invalid <i>fAccuracy</i> value was specified.
S1C00	Driver not capable	<p>A table qualifier was specified and the driver or data source does not support qualifiers.</p> <p>A table owner was specified and the driver or data source does not support owners.</p> <p>The combination of the current settings of the SQL_CONCURRENCY and SQL_CURSOR_TYPE statement options was not supported by the driver or data source.</p>
S1T00	Timeout expired	<p>The timeout period expired before the data source returned the requested result set. The timeout period is set through SQLSetStmtOption, SQL_QUERY_TIMEOUT.</p>

Comments

SQLStatistics returns information about a single table as a standard result set, ordered by NON_UNIQUE, TYPE, INDEX_QUALIFIER, INDEX_NAME, and SEQ_IN_INDEX. The result set combines statistics information for the table with information about each index. The following table lists the columns in the result set.

Note: **SQLStatistics** might not return all indexes. For example, an Xbase driver might only return indexes in files in the current directory. Applications can use any valid index, regardless of whether it is returned by **SQLStatistics**.

The lengths of VARCHAR columns shown in the table are maximums; the actual lengths depend on the data source. To determine the actual lengths of the TABLE_QUALIFIER, TABLE_OWNER, TABLE_NAME, and COLUMN_NAME columns, an application can call **SQLGetInfo** with the SQL_MAX_QUALIFIER_NAME_LEN, SQL_MAX_OWNER_NAME_LEN, SQL_MAX_TABLE_NAME_LEN, and SQL_MAX_COLUMN_NAME_LEN options.

Column Name	Data Type	Comments
TABLE_QUALIFIER	Varchar(128)	Table qualifier identifier of the table to which the statistic or index applies; NULL if not applicable to the data source. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those tables that do not have qualifiers.
TABLE_OWNER	Varchar(128)	Table owner identifier of the table to which the statistic or index applies; NULL if not applicable to the data source. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those tables that do not have owners.
TABLE_NAME	Varchar(128) not NULL	Table identifier of the table to which the statistic or index applies.
Column Name	Data Type	Comments
NON_UNIQUE	Smallint	Indicates whether the index prohibits duplicate values: TRUE if the index values can be nonunique. FALSE if the index values must be unique.

		NULL is returned if TYPE is SQL_TABLE_STAT.
INDEX_QUALIFIER	Varchar(128)	The identifier that is used to qualify the index name doing a DROP INDEX ; NULL is returned if an index qualifier is not supported by the data source or if TYPE is SQL_TABLE_STAT. If a non-null value is returned in this column, it must be used to qualify the index name on a DROP INDEX statement; otherwise the TABLE_OWNER name should be used to qualify the index name.
INDEX_NAME	Varchar(128)	Index identifier; NULL is returned if TYPE is SQL_TABLE_STAT.
TYPE	Smallint not NULL	Type of information being returned: SQL_TABLE_STAT indicates a statistic for the table. SQL_INDEX_CLUSTERED indicates a clustered index. SQL_INDEX_HASHED indicates a hashed index. SQL_INDEX_OTHER indicates another type of index.
SEQ_IN_INDEX	Smallint	Column sequence number in index (starting with 1); NULL is returned if TYPE is SQL_TABLE_STAT.
COLUMN_NAME	Varchar(128)	Column identifier. If the column is based on an expression, such as SALARY + BENEFITS, the expression is returned; if the expression cannot be determined, an empty string is returned. If the index is a filtered index, each column in the filter condition is returned; this may require more than one row. NULL is returned if TYPE is SQL_TABLE_STAT.
Column Name	Data Type	Comments
COLLATION	Char(1)	Sort sequence for the column; "A" for ascending; "D" for descending; NULL is returned if column sort sequence is not supported by the data source or if TYPE is SQL_TABLE_STAT.
CARDINALITY	Integer	Cardinality of table or index; number of rows in table if TYPE is SQL_TABLE_STAT; number of unique

		values in the index if TYPE is not SQL_TABLE_STAT; NULL is returned if the value is not available from the data source.
PAGES	Integer	Number of pages used to store the index or table; number of pages for the table if TYPE is SQL_TABLE_STAT; number of pages for the index if TYPE is not SQL_TABLE_STAT; NULL is returned if the value is not available from the data source, or if not applicable to the data source.
FILTER_CONDITION	Varchar(128)	If the index is a filtered index, this is the filter condition, such as SALARY > 30000; if the filter condition cannot be determined, this is an empty string. NULL if the index is not a filtered index, it cannot be determined whether the index is a filtered index, or TYPE is SQL_TABLE_STAT.

Note: The FILTER_CONDITION column was added in ODBC 2.0. ODBC 1.0 drivers might return a different, driver-specific column with the same column number.

If the row in the result set corresponds to a table, the driver sets TYPE to SQL_TABLE_STAT and sets NON_UNIQUE, INDEX_QUALIFIER, INDEX_NAME, SEQ_IN_INDEX, COLUMN_NAME, and COLLATION to NULL. If CARDINALITY or PAGES are not available from the data source, the driver sets them to NULL.

Code Example For a code example of a similar function, see **SQLColumns**.

Related Functions	For information about	See
	Assigning storage for a column in a result set	SQLBindCol
	Canceling statement processing	SQLCancel
	Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
	Fetching a row of data	SQLFetch
	Returning the columns of foreign keys	SQLForeignKeys (extension)
	Returning the columns of a primary key	SQLPrimaryKeys (extension)

SQLTablePrivileges



Extension Level 2 SQLTablePrivileges returns a list of tables and the privileges associated with each table. The driver returns the information as a result set on the specified *hstmt*.

Syntax RETCODE **SQLTablePrivileges**(*hstmt*, *szTableQualifier*, *cbTableQualifier*, *szTableOwner*, *cbTableOwner*, *szTableName*, *cbTableName*)

The **SQLTablePrivileges** function accepts the following arguments.

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle.
UCHAR FAR *	<i>szTableQualifier</i>	Input	Table qualifier. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have qualifiers.
SWORD	<i>cbTableQualifier</i>	Input	Length of <i>szTableQualifier</i> .
UCHAR FAR *	<i>szTableOwner</i>	Input	String search pattern for owner names. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have owners.
SWORD	<i>cbTableOwner</i>	Input	Length of <i>szTableOwner</i> .
UCHAR FAR *	<i>szTableName</i>	Input	String search pattern for table names.
SWORD	<i>cbTableName</i>	Input	Length of <i>szTableName</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLTablePrivileges** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLTablePrivileges** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

Appendixes

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns <code>SQL_SUCCESS_WITH_INFO</code> .)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called. A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1010	Function sequence error	(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and returned <code>SQL_NEED_DATA</code> . This function was called before data was sent for all data-at-execution parameters or columns.

SQLSTATE Error		Description
S1090	Invalid string or buffer length	(DM) The value of one of the name length arguments was less than 0, but not equal to SQL_NTS. The value of one of the name length arguments exceeded the maximum length value for the corresponding qualifier or name.
S1C00	Driver not capable	A table qualifier was specified and the driver or data source does not support qualifiers. A table owner was specified and the driver or data source does not support owners. A string search pattern was specified for the table owner, table name, or column name and the data source does not support search patterns for one or more of those arguments. The combination of the current settings of the SQL_CONCURRENCY and SQL_CURSOR_TYPE statement options was not supported by the driver or data source.
S1T00	Timeout expired	The timeout period expired before the data source returned the result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

Comments The *szTableOwner* and *szTableName* arguments accept search patterns. For more information about valid search patterns, see “Search Pattern Arguments” earlier in this chapter.

SQLTablePrivileges returns the results as a standard result set, ordered by TABLE_QUALIFIER, TABLE_OWNER, TABLE_NAME, and PRIVILEGE. The following table lists the columns in the result set.

Note: **SQLTablePrivileges** might not return privileges for all tables. For example, an Xbase driver might only return privileges for files (tables) in the current directory. Applications can use any valid table, regardless of whether it is returned by **SQLTablePrivileges**.

Appendixes

The lengths of VARCHAR columns shown in the table are maximums; the actual lengths depend on the data source. To determine the actual lengths of the TABLE_QUALIFIER, TABLE_OWNER, and TABLE_NAME columns, an application can call **SQLGetInfo** with the SQL_MAX_QUALIFIER_NAME_LEN, SQL_MAX_OWNER_NAME_LEN, and SQL_MAX_TABLE_NAME_LEN options.

Column Name	Data Type	Comments
TABLE_QUALIFIER	Varchar(128)	Table qualifier identifier; NULL if not applicable to the data source. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those tables that do not have qualifiers.
TABLE_OWNER	Varchar(128)	Table owner identifier; NULL if not applicable to the data source. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those tables that do not have owners.
TABLE_NAME	Varchar(128) not NULL	Table identifier.
GRANTOR	Varchar(128)	Identifier of the user who granted the privilege; NULL if not applicable to the data source.
GRANTEE	Varchar(128) not NULL	Identifier of the user to whom the privilege was granted.
PRIVILEGE	Varchar(128) not NULL	Identifies the table privilege. May be one of the following or a data source-specific privilege. SELECT: The grantee is permitted to retrieve data for one or more columns of the table. INSERT: The grantee is permitted to insert new rows containing data for one or more columns into to the table. UPDATE: The grantee is permitted to update the data in one or more columns of the table. DELETE: The grantee is permitted to delete rows of data from the table.

REFERENCES: The grantee is permitted to refer to one or more columns of the table within a constraint (for example, a unique, referential, or table check constraint).

Column Name	Data Type	Comments
		The scope of action permitted the grantee by a given table privilege is data source-dependent. For example, the UPDATE privilege might permit the grantee to update all columns in a table on one data source and only those columns for which the grantor has the UPDATE privilege on another data source.
IS_GRANTABLE	Varchar(3)	Indicates whether the grantee is permitted to grant the privilege to other users; "YES", "NO", or NULL if unknown or not applicable to the data source.

Code Example For a code example of a similar function, see **SQLColumns**.

Related Function	For information about	See
	Assigning storage for a column in a result set	SQLBindCol
	Canceling statement processing	SQLCancel
	Returning privileges for a column or columns	SQLColumnPrivileges (extension)
	Returning the columns in a table or tables	SQLColumns (extension)
	Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
	Fetching a row of data	SQLFetch
	Returning table statistics and indexes	SQLStatistics (extension)
	Returning a list of tables in a data source	SQLTables (extension)

SQLTables



Extension Level 1 SQLTables returns the list of table names stored in a specific data source. The driver returns the information as a result set.

Syntax RETCODE **SQLTables**(*hstmt*, *szTableQualifier*, *cbTableQualifier*, *szTableOwner*, *cbTableOwner*, *szTableName*, *cbTableName*, *szTableType*, *cbTableType*)

The **SQLTables** function accepts the following arguments:

Type	Argument	Use	Description
HSTMT	<i>hstmt</i>	Input	Statement handle for retrieved results.
UCHAR FAR *	<i>szTableQualifier</i>	Input	Qualifier name. If a driver supports qualifiers for some tables but not for others, such as when a driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have qualifiers.
SWORD	<i>cbTableQualifier</i>	Input	Length of <i>szTableQualifier</i> .
UCHAR FAR *	<i>szTableOwner</i>	Input	String search pattern for owner names.
SWORD	<i>cbTableOwner</i>	Input	Length of <i>szTableOwner</i> .
UCHAR FAR *	<i>szTableName</i>	Input	String search pattern for table names. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, an empty string ("") denotes those tables that do not have owners.
SWORD	<i>cbTableName</i>	Input	Length of <i>szTableName</i> .
UCHAR FAR *	<i>szTableType</i>	Input	List of table types to match.
SWORD	<i>cbTableType</i>	Input	Length of <i>szTableType</i> .

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_STILL_EXECUTING, SQL_ERROR or SQL_INVALID_HANDLE.

Diagnostics When **SQLTables** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLTables** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of

Appendixes

SQLSTATEs returned by the Driver Manager. The return code associated with each SQLSTATE value is `SQL_ERROR`, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns <code>SQL_SUCCESS_WITH_INFO</code> .)
08S01	Communication link failure	The communication link between the driver and the data source to which the driver was connected failed before the function completed processing.
24000	Invalid cursor state	(DM) A cursor was open on the <i>hstmt</i> and SQLFetch or SQLExtendedFetch had been called. A cursor was open on the <i>hstmt</i> but SQLFetch or SQLExtendedFetch had not been called.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hstmt</i> does not support the function.
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1008	Operation canceled	Asynchronous processing was enabled for the <i>hstmt</i> . The function was called and before it completed execution, SQLCancel was called on the <i>hstmt</i> . Then the function was called again on the <i>hstmt</i> . The function was called and, before it completed execution, SQLCancel was called on the <i>hstmt</i> from a different thread in a multithreaded application.
S1010	Function sequence error	(DM) An asynchronously executing function (not this one) was called for the <i>hstmt</i> and was still executing when this function was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for the <i>hstmt</i> and

returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.

SQLSTATE Error		Description
S1090	Invalid string or buffer length	(DM) The value of one of the name length arguments was less than 0, but not equal to SQL_NTS. The value of one of the name length arguments exceeded the maximum length value for the corresponding qualifier or name.
S1C00	Driver not capable	A table qualifier was specified and the driver or data source does not support qualifiers. A table owner was specified and the driver or data source does not support owners. A string search pattern was specified for the table owner or table name and the data source does not support search patterns for one or more of those arguments. The combination of the current settings of the SQL_CONCURRENCY and SQL_CURSOR_TYPE statement options was not supported by the driver or data source.
S1T00	Timeout expired	The timeout period expired before the data source returned the requested result set. The timeout period is set through SQLSetStmtOption , SQL_QUERY_TIMEOUT.

Comments

SQLTables lists all tables in the requested range. A user may or may not have SELECT privileges to any of these tables. To check accessibility, an application can:

- Call **SQLGetInfo** and check the SQL_ACCESSIBLE_TABLES info value.
- Call **SQLTablePrivileges** to check the privileges for each table.

Otherwise, the application must be able to handle a situation where the user selects a table for which SELECT privileges are not granted.

The *szTableOwner* and *szTableName* arguments accept search patterns. For more information about valid search patterns, see “Search Pattern Arguments” earlier in this chapter.

To support enumeration of qualifiers, owners, and table types, **SQLTables** defines the following special semantics for the *szTableQualifier*, *szTableOwner*, *szTableName*, and *szTableType* arguments:

- If *szTableQualifier* is a single percent character (%) and *szTableOwner* and *szTableName* are empty strings, then the result set contains a list of valid qualifiers for the data source. (All columns except the TABLE_QUALIFIER column contain NULLs.)
- If *szTableOwner* is a single percent character (%) and *szTableQualifier* and *szTableName* are empty strings, then the result set contains a list of valid owners for the data source. (All columns except the TABLE_OWNER column contain NULLs.)
- If *szTableType* is a single percent character (%) and *szTableQualifier*, *szTableOwner*, 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 comma-separated, values for the types of interest; each value may be enclosed in single quotes (') or unquoted. For example, “TABLE',VIEW” or “TABLE, VIEW”. If the data source does not support a specified table type, **SQLTables** does not return any results for that type.

SQLTables returns the results as a standard result set, ordered by TABLE_TYPE, TABLE_QUALIFIER, TABLE_OWNER, and TABLE_NAME. The following table lists the columns in the result set.

Note: **SQLTables** might not return all qualifiers, owners, or tables. For example, an Xbase driver, for which a qualifier is a directory, might only return the current directory instead of all directories on the system. It might also only return files (tables) in the current directory. Applications can use any valid qualifier, owner, or table, regardless of whether it is returned by **SQLTables**.

The lengths of VARCHAR columns shown in the table are maximums; the actual lengths depend on the data source. To determine the actual

lengths of the TABLE_QUALIFIER, TABLE_OWNER, and TABLE_NAME columns, an application can call **SQLGetInfo** with the SQL_MAX_QUALIFIER_NAME_LEN, SQL_MAX_OWNER_NAME_LEN, and SQL_MAX_TABLE_NAME_LEN options.

Column Name	Data Type	Comments
TABLE_QUALIFIER	Varchar(128)	Table qualifier identifier; NULL if not applicable to the data source. If a driver supports qualifiers for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those tables that do not have qualifiers.
TABLE_OWNER	Varchar(128)	Table owner identifier; NULL if not applicable to the data source. If a driver supports owners for some tables but not for others, such as when the driver retrieves data from different DBMSs, it returns an empty string ("") for those tables that do not have owners.
TABLE_NAME	Varchar(128)	Table identifier.
TABLE_TYPE	Varchar(128)	Table type identifier; one of the following: "TABLE", "VIEW", "SYSTEM", "ALIAS", "SYNONYM" or a data source – specific type identifier.
REMARKS	Varchar(254)	A description of the table.

Code Example For a code example of a similar function, see **SQLColumns**.

Related Functions

For information about	See
Assigning storage for a column in a result set	SQLBindCol
Canceling statement processing	SQLCancel
Returning privileges for a column or columns	SQLColumnPrivileges (extension)
Returning the columns in a table or tables	SQLColumns (extension)
Fetching a block of data or scrolling through a result set	SQLExtendedFetch (extension)
Fetching a row of data	SQLFetch
Returning table statistics and indexes	SQLStatistics (extension)
Returning privileges for a table or tables	SQLTablePrivileges (extension)

Appendixes

SQLTransact



Core **SQLTransact** requests a commit or rollback operation for all active operations on all *hstmts* associated with a connection. **SQLTransact** can also request that a commit or rollback operation be performed for all connections associated with the *henv*.

Syntax RETCODE **SQLTransact**(*henv*, *hdbc*, *fType*)

The **SQLTransact** function accepts the following arguments.

Type	Argument	Use	Description
HENV	<i>henv</i>	Input	Environment handle.
HDBC	<i>hdbc</i>	Input	Connection handle.
UWORD	<i>fType</i>	Input	One of the following two values: SQL_COMMIT SQL_ROLLBACK

Returns SQL_SUCCESS, SQL_SUCCESS_WITH_INFO, SQL_ERROR, or SQL_INVALID_HANDLE.

Diagnostics When **SQLTransact** returns SQL_ERROR or SQL_SUCCESS_WITH_INFO, an associated SQLSTATE value may be obtained by calling **SQLError**. The following table lists the SQLSTATE values commonly returned by **SQLTransact** and explains each one in the context of this function; the notation "(DM)" precedes the descriptions of SQLSTATES returned by the Driver Manager. The return code associated with each SQLSTATE value is SQL_ERROR, unless noted otherwise.

SQLSTATE	Error	Description
01000	General warning	Driver-specific informational message. (Function returns SQL_SUCCESS_WITH_INFO.)
08003	Connection not open	(DM) The <i>hdbc</i> was 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.
IM001	Driver does not support this function	(DM) The driver associated with the <i>hdbc</i> does not support the function.

SQLSTATE	Error	Description
S1000	General error	An error occurred for which there was no specific SQLSTATE and for which no implementation-specific SQLSTATE was defined. The error message returned by SQLError in the argument <i>szErrorMsg</i> describes the error and its cause.
S1001	Memory allocation failure	The driver was unable to allocate memory required to support execution or completion of the function.
S1010	Function sequence error	(DM) An asynchronously executing function was called for an <i>hstmt</i> associated with the <i>hdbc</i> and was still executing when SQLTransact was called. (DM) SQLExecute , SQLExecDirect , or SQLSetPos was called for an <i>hstmt</i> associated with the <i>hdbc</i> and returned SQL_NEED_DATA. This function was called before data was sent for all data-at-execution parameters or columns.
S1012	Invalid transaction operation code	(DM) The value specified for the argument <i>fType</i> was neither SQL_COMMIT nor SQL_ROLLBACK.
S1C00	Driver not capable	The driver or data source does not support the ROLLBACK operation.

Comments

If *hdbc* is SQL_NULL_HDBC and *henv* is a valid environment handle, then the Driver Manager will attempt to commit or roll back transactions on all *hdbcs* that are in a connected state. The Driver Manager calls **SQLTransact** in the driver associated with each *hdbc*. The Driver Manager will return SQL_SUCCESS only if it receives SQL_SUCCESS for each *hdbc*. If the Driver Manager receives SQL_ERROR on one or more *hdbcs*, it will return SQL_ERROR to the application. To determine which connection(s) failed during the commit or rollback operation, the application can call **SQLError** for each *hdbc*.

Note: The Driver Manager does not simulate a global transaction across all *hdbcs* and therefore does not use two-phase commit protocols.

If *hdbc* is a valid connection handle, *henv* is ignored and the Driver Manager calls **SQLTransact** in the driver for the *hdbc*.

If *hdbc* is SQL_NULL_HDBC and *henv* is SQL_NULL_HENV, **SQLTransact** returns SQL_INVALID_HANDLE.

If *fType* is SQL_COMMIT, **SQLTransact** issues a commit request for all active operations on any *hstmt* associated with an affected *hdbc*. If *fType* is SQL_ROLLBACK, **SQLTransact** issues a rollback request for all active operations on any *hstmt* associated with an affected *hdbc*. If no transactions are active, **SQLTransact** returns SQL_SUCCESS with no effect on any data sources.

If the driver is in manual-commit mode (by calling **SQLSetConnectOption** with the SQL_AUTOCOMMIT option set to zero), a new transaction is implicitly started when an SQL statement that can be contained within a transaction is executed against the current data source.

To determine how transaction operations affect cursors, an application calls **SQLGetInfo** with the SQL_CURSOR_ROLLBACK_BEHAVIOR and SQL_CURSOR_COMMIT_BEHAVIOR options.

If the SQL_CURSOR_ROLLBACK_BEHAVIOR or SQL_CURSOR_COMMIT_BEHAVIOR value equals SQL_CB_DELETE, **SQLTransact** closes and deletes all open cursors on all *hstmts* associated with the *hdbc* and discards all pending results. **SQLTransact** leaves any *hstmt* present in an allocated (unprepared) state; the application can reuse them for subsequent SQL requests or can call **SQLFreeStmt** to deallocate them.

If the SQL_CURSOR_ROLLBACK_BEHAVIOR or SQL_CURSOR_COMMIT_BEHAVIOR value equals SQL_CB_CLOSE, **SQLTransact** closes all open cursors on all *hstmts* associated with the *hdbc*. **SQLTransact** leaves any *hstmt* present in a prepared state; the application can call **SQLExecute** for an *hstmt* associated with the *hdbc* without first calling **SQLPrepare**.

If the SQL_CURSOR_ROLLBACK_BEHAVIOR or SQL_CURSOR_COMMIT_BEHAVIOR value equals SQL_CB_PRESERVE, **SQLTransact** does not affect open cursors associated with the *hdbc*. Cursors remain at the row they pointed to prior to the call to **SQLTransact**.

For drivers and data sources that support transactions, calling **SQLTransact** with either SQL_COMMIT or SQL_ROLLBACK when no transaction is active will return SQL_SUCCESS (indicating that there is no work to be committed or rolled back) and have no effect on the data source.

Drivers or data sources that do not support transactions (**SQLGetInfo** *fOption* SQL_TXN_CAPABLE is 0) are effectively always in autocommit

mode. Therefore, calling **SQLTransact** with SQL_COMMIT will return SQL_SUCCESS. However, calling **SQLTransact** with SQL_ROLLBACK will result in SQLSTATE S1C00 (Driver not capable), indicating that a rollback can never be performed.

Code Example See **SQLParamOptions**.

Related Function	For information about	See
	Returning information about a driver or data source	SQLGetInfo (extension)
	Freeing a statement handle	SQLFreeStmt

CHAPTER 23

This chapter describes the syntax of the driver setup DLL API, which consists of a single function (**ConfigDSN**). **ConfigDSN** may be either in the driver DLL or in a separate setup DLL.

It also describes the syntax of the translator setup DLL API, which consists of a single function (**ConfigTranslator**). **ConfigTranslator** may be either in the translator DLL or in a separate setup DLL.

Each function is labeled with the version of ODBC in which it was introduced.

For information on argument naming conventions, see Chapter 22, “ODBC Function Reference”.

ConfigDSN



Purpose **ConfigDSN** adds, modifies, or deletes data sources from the ODBC.INI file (or registry). It may prompt the user for connection information. It can be in the driver DLL or a separate setup DLL.

Syntax **BOOL ConfigDSN**(*hwndParent*, *fRequest*, *lpszDriver*, *lpszAttributes*)

The **ConfigDSN** function accepts the following arguments.

Type	Argument	Use	Description
HWND	<i>hwndParent</i>	Input	Parent window handle. The function will not display any dialog boxes if the handle is null.
UINT	<i>fRequest</i>	Input	Type of request. <i>fRequest</i> must contain one of the following values: ODBC_ADD_DSN: add a new data source. ODBC_CONFIG_DSN: configure (modify) an existing data source. ODBC_REMOVE_DSN: remove an existing data source.
LPCSTR	<i>lpszDriver</i>	Input	Driver description (usually the name of the associated DBMS) presented to users instead of the physical driver name.
LPCSTR	<i>lpszAttributes</i>	Input	List of attributes in the form of keyword-value pairs. For information about the list structure, see "Comments."

Returns The function returns TRUE if it is successful. It returns FALSE if it fails.

Comments **ConfigDSN** receives connection information from the installer DLL as a list of attributes in the form of keyword-value pairs. Each pair is terminated with a null byte and the entire list is terminated with a null byte (that is, two null bytes mark the end of the list). The keywords used by **ConfigDSN** are the same as those used by **SQLBrowseConnect** and **SQLDriverConnect**, except that **ConfigDSN** does not accept the **DRIVER** keyword. As in **SQLBrowseConnect** and **SQLDriverConnect**, the keywords and their values should not contain the `[]{}(),;?*!=!@` characters, and the value of the **DSN** keyword cannot consist only of

blanks. Because of the registry grammar, keywords and data source names cannot contain the backslash (\) character.

For example, to configure a data source that requires a user ID, password, and database name, a setup application might pass the following keyword-value pairs:

DSN=Personnel Data\0UID=Smith\0PWD=Sesame\0DATABASE=Personnel\0\0

For more information about these keywords, see **SQLDriverConnect** and each driver's documentation.

In order to display a dialog box, *hwndParent* must not be null.

Adding a Data Source

If a data source name is passed to **ConfigDSN** in *lpszAttributes*, **ConfigDSN** checks that the name is valid. If the data source name matches an existing data source name and *hwndParent* is null, **ConfigDSN** overwrites the existing name. If it matches an existing name and *hwndParent* is not null, **ConfigDSN** prompts the user to overwrite the existing name.

If *lpszAttributes* contains enough information to connect to a data source, **ConfigDSN** can add the data source or display a dialog box with which the user can change the connection information. If *lpszAttributes* does not contain enough information to connect to a data source, **ConfigDSN** must determine the necessary information; if *hwndParent* is not null, it displays a dialog box to retrieve the information from the user.

If **ConfigDSN** displays a dialog box, it must display any connection information passed to it in *lpszAttributes*. In particular, if a data source name was passed to it, **ConfigDSN** displays that name but does not allow the user to change it. **ConfigDSN** can supply default values for connection information not passed to it in *lpszAttributes*.

If **ConfigDSN** cannot get complete connection information for a data source, it returns FALSE.

If **ConfigDSN** can get complete connection information for a data source, it calls **SQLWriteDSNTolni** in the installer DLL to add the new data source specification to the ODBC.INI file (or registry).

SQLWriteDSNTolni adds the data source name to the [ODBC Data Sources] section, creates the data source specification section, and adds the **Driver** keyword with the driver description as its value. **ConfigDSN** calls **SQLWritePrivateProfileString** in the installer DLL to add any additional keywords and values used by the driver.

Modifying a Data Source

To modify a data source, a data source name must be passed to **ConfigDSN** in *lpszAttributes*. **ConfigDSN** checks that the data source name is in the ODBC.INI file (or registry).

If *hwndParent* is null, **ConfigDSN** uses the information in *lpszAttributes* to modify the information in the ODBC.INI file (or registry). If *hwndParent* is not null, **ConfigDSN** displays a dialog box using the information in *lpszAttributes*; for information not in *lpszAttributes*, it uses information from the ODBC.INI file (or registry). The user can modify the information before **ConfigDSN** stores it in the ODBC.INI file (or registry).

If the data source name was changed, **ConfigDSN** first calls **SQLRemoveDSNFromIni** in the installer DLL to remove the existing data source specification from the ODBC.INI file (or registry). It then follows the steps in the previous section to add the new data source specification. If the data source name was not changed, **ConfigDSN** calls **SQLWritePrivateProfileString** in the installer DLL to make any other changes. **ConfigDSN** may not delete or change the value of the **Driver** keyword.

Deleting a Data Source

To delete a data source, a data source name must be passed to **ConfigDSN** in *lpszAttributes*. **ConfigDSN** checks that the data source name is in the ODBC.INI file (or registry). It then calls **SQLRemoveDSNFromIni** in the installer DLL to remove the data source.

Related Functions

For information about	See
Adding, modifying, or removing a data source	SQLConfigDataSource
Getting a value from the ODBC.INI file or the registry	SQLGetPrivateProfileString
Removing the default data source	SQLRemoveDefaultDataSource
Removing a data source name from ODBC.INI (or registry)	SQLRemoveDSNFromIni
Adding a data source name to ODBC.INI (or registry)	SQLWriteDSNToIni
Writing a value to the ODBC.INI file or the registry	SQLWritePrivateProfileString

ConfigTranslator



Purpose **ConfigTranslator** returns a default translation option for a translator. It can be in the translator DLL or a separate setup DLL.

Syntax **BOOL ConfigTranslator**(*hwndParent*, *pvOption*)

The **ConfigTranslator** function accepts the following arguments.

Type	Argument	Use	Description
HWND	<i>hwndParent</i>	Input	Parent window handle. The function will not display any dialog boxes if the handle is null.
DWORD FAR *	<i>pvOption</i>	Output	A 32-bit translation option.

Returns The function returns TRUE if it is successful. It returns FALSE if it fails.

Comments If the translator supports only a single translation option, **ConfigTranslator** returns TRUE and sets *pvOption* to the 32-bit option. Otherwise, it determines the default translation option to use. **ConfigTranslator** can display a dialog box with which a user selects a default translation option.

Related Functions	For information about	See
	Getting a translation option	SQLGetConnectOption
	Selecting a translator	SQLGetTranslator
	Setting a translation option	SQLSetConnectOption

!Unexpected End of Expression
C H A P T E R 2 4

- omitted -

Translation DLL Function Reference (Windows)

The following section describes the syntax of the translation DLL API, which consists of two functions: **SQLDriverToDataSource** and **SQLDataSourceToDriver**. These functions must be included in the DLL which performs translation for the driver.

For information on argument naming conventions, see Chapter 22, “ODBC Function Reference.”

SQLDataSourceToDriver

Extension Level 2 **SQLDataSourceToDriver** supports translations for ODBC drivers. This function is not called by ODBC-enabled applications; applications request translation through **SQLSetConnectOption**. The driver associated with the *hdbc* specified in **SQLSetConnectOption** calls the specified DLL to perform translations of all data flowing from the data source to the driver. A default translation DLL can be specified in the ODBC initialization file.

Syntax

BOOL **SQLDataSourceToDriver**(*fOption*, *fSqlType*, *rgbValueIn*, *cbValueIn*, *rgbValueOut*, *cbValueOutMax*, *pcbValueOut*, *szErrorMsg*, *cbErrorMsgMax*, *pcbErrorMsg*)

The **SQLDataSourceToDriver** function accepts the following arguments:

Type	Argument	Use	Description
UDWORD	<i>fOption</i>	Input	Option value.
SWORD	<i>fSqlType</i>	Input	<p>The SQL data type. This argument tells the driver how to convert <i>rgbValueIn</i> into a form acceptable by the application. This must be one of the following values:</p> <p>SQL_BIGINT SQL_BINARY SQL_BIT SQL_CHAR SQL_DATE SQL_DECIMAL SQL_DOUBLE SQL_FLOAT SQL_INTEGER SQL_LONGVARIABLE SQL_LONGVARIABLE SQL_NUMERIC SQL_REAL SQL_SMALLINT SQL_TIME SQL_TIMESTAMP SQL_TINYINT SQL_VARIABLE SQL_VARIABLE</p> <p>For information about SQL data types, see "SQL Data Types" in Appendix D, "Data Types."</p>
Type	Argument	Use	Description

Index

PTR	<i>rgbValueIn</i>	Input	Value to translate.
SDWORD	<i>cbValueIn</i>	Input	Length of <i>rgbValueIn</i> .
PTR	<i>rgbValueOut</i>	Output	Result of the translation.

Note: The translation DLL does not null-terminate this value.

SDWORD	<i>cbValueOutMax</i>	Input	Length of <i>rgbValueOut</i> .
SDWORD FAR * <i>pcbValueOut</i>		Output	The total number of bytes (excluding the null termination byte) available to return in <i>rgbValueOut</i> . For character or binary data, if this is greater than or equal to <i>cbValueOutMax</i> , the data in <i>rgbValueOut</i> is truncated to <i>cbValueOutMax</i> bytes. For all other data types, the value of <i>cbValueOutMax</i> is ignored and the translation DLL assumes the size of <i>rgbValueOut</i> is the size of the default C data type of the SQL data type specified with <i>fSqlType</i> .
UCHAR FAR * <i>szErrorMsg</i>		Output	Pointer to storage for an error message. This is an empty string unless the translation failed.
SWORD	<i>cbErrorMsgMax</i>	Input	Length of <i>szErrorMsg</i> .
SWORD FAR * <i>pcbErrorMsg</i>		Output	Pointer to the total number of bytes (excluding the null termination byte) available to return in <i>szErrorMsg</i> . If this is greater than or equal to <i>cbErrorMsg</i> , the data in <i>szErrorMsg</i> is truncated to <i>cbErrorMsgMax</i> – 1 bytes.

Returns TRUE if the translation was successful.

FALSE if the translation failed.

Comments The driver calls **SQLDataSourceToDriver** to translate *all* data (result set data, table names, row counts, error messages, and so on) passing from the data source to the driver. The translation DLL may not translate some data, depending on the data's type and the purpose of the translation DLL; for example, a DLL that translates character data from one code page to another ignores all numeric and binary data.

The value of *fOption* is set to the value of *vParam* specified by calling **SQLSetConnectOption** with the SQL_TRANSLATE_OPTION option. It is a 32-bit value which has a specific meaning for a given translation DLL. For example, it could specify a certain character set translation.

If the same buffer is specified for *rgbValueIn* and *rgbValueOut*, the translation of data in the buffer will be performed in place.

Note that, although *cbValueIn*, *cbValueOutMax*, and *pcbValueOut* are of the type SDWORD, **SQLDataSourceToDriver** does not necessarily support huge pointers.

If **SQLDataSourceToDriver** returns FALSE, data truncation may have occurred during translation. If *pcbValueOut*, the number of bytes available to return in the output buffer, is greater than *cbValueOutMax*, the length of the output buffer, then truncation occurred. The driver must determine whether the truncation was acceptable. If truncation did not occur, the **SQLDataSourceToDriver** returned FALSE due to another error. In either case, a specific error message is returned in *szErrorMsg*.

For more information about translating data, see “Translating Data” in Chapter 13, “Establishing Connections.”

Related Functions

For information about

See

Translating data being sent to the data source

SQLDriverToDataSource

Returning the setting of a connection option

SQLGetConnectOption
(extension)

Setting a connection option

SQLSetConnectOption
(extension)

SQLDriverToDataSource

Extension Level 2 **SQLDriverToDataSource** supports translations for ODBC drivers.

This function is not called by ODBC-enabled applications; applications request translation through **SQLSetConnectOption**. The driver associated with the *hdbc* specified in **SQLSetConnectOption** calls the specified DLL to perform translations of all data flowing from the driver to the data source. A default translation DLL can be specified in the ODBC initialization file.

Syntax

BOOL SQLDriverToDataSource(*fOption*, *fSqlType*, *rgbValueIn*, *cbValueIn*, *rgbValueOut*, *cbValueOutMax*, *pcbValueOut*, *szErrorMsg*, *cbErrorMsg*, *pcbErrorMsg*)

The **SQLDriverToDataSource** function accepts the following arguments:

Type	Argument	Use	Description
UDWORD	<i>fOption</i>	Input	Option value.
SWORD	<i>fSqlType</i>	Input	The ODBC SQL data type. This argument tells the driver how to convert <i>rgbValueIn</i> into a form acceptable by the data source. This must be one of the following values: SQL_BIGINT SQL_BINARY SQL_BIT SQL_CHAR SQL_DATE SQL_DECIMAL SQL_DOUBLE SQL_FLOAT SQL_INTEGER SQL_LONGVARBINARY SQL_LONGVARCHAR SQL_NUMERIC SQL_REAL SQL_SMALLINT SQL_TIME SQL_TIMESTAMP SQL_TINYINT SQL_VARBINARY SQL_VARCHAR For information about SQL data types, see "SQL Data Types" in Appendix D, "Data Types."

Contents

Type	Argument	Use	Description
PTR	<i>rgbValueIn</i>	Input	Value to translate.
SDWORD	<i>cbValueIn</i>	Input	Length of <i>rgbValueIn</i> .
PTR	<i>rgbValueOut</i>	Output	Result of the translation.
Note: The translation DLL does not null-terminate this value.			
SDWORD	<i>cbValueOutMax</i>	Input	Length of <i>rgbValueOut</i> .
SDWORD FAR *	<i>pcbValueOut</i>	Output	The total number of bytes (excluding the null termination byte) available to return in <i>rgbValueOut</i> . For character or binary data, if this is greater than or equal to <i>cbValueOutMax</i> , the data in <i>rgbValueOut</i> is truncated to <i>cbValueOutMax</i> bytes. For all other data types, the value of <i>cbValueOutMax</i> is ignored and the translation DLL assumes the size of <i>rgbValueOut</i> is the size of the default C data type of the SQL data type specified with <i>fSqlType</i> .
UCHAR FAR *	<i>szErrorMsg</i>	Output	Pointer to storage for an error message. This is an empty string unless the translation failed.
SWORD	<i>cbErrorMsgMax</i>	Input	Length of <i>szErrorMsg</i> .
SWORD FAR *	<i>pcbErrorMsg</i>	Output	Pointer to the total number of bytes (excluding the null termination byte) available to return in <i>szErrorMsg</i> . If this is greater than or equal to <i>cbErrorMsg</i> , the data in <i>szErrorMsg</i> is truncated to <i>cbErrorMsgMax</i> – 1 bytes.
Returns	TRUE if the translation was successful. FALSE if the translation failed.		
Comments	The driver calls SQLDriverToDataSource to translate <i>all</i> data (SQL statements, parameters, and so on) passing from the driver to the data source. The translation DLL may not translate some data, depending on the data's type and the purpose of the translation DLL. For example, a		

DLL that translates character data from one code page to another ignores all numeric and binary data.

The value of *fOption* is set to the value of *vParam* specified by calling **SQLSetConnectOption** with the SQL_TRANSLATE_OPTION option. It is a 32-bit value which has a specific meaning for a given translation DLL. For example, it could specify a certain character set translation.

If the same buffer is specified for *rgbValueIn* and *rgbValueOut*, the translation of data in the buffer will be performed in-place.

Note that, although *cbValueIn*, *cbValueOutMax*, and *pcbValueOut* are of the type SDWORD, **SQLDriverToDataSource** does not necessarily support huge pointers.

If **SQLDriverToDataSource** returns FALSE, data truncation may have occurred during translation. If *pcbValueOut*, the number of bytes available to return in the output buffer, is greater than *cbValueOutMax*, the length of the output buffer, then truncation occurred. The driver must determine whether or not the truncation was acceptable. If truncation did not occur, the **SQLDriverToDataSource** returned FALSE due to another error. In either case, a specific error message is returned in *szErrorMsg*.

For more information about translating data, see “Translating Data” in Chapter 13, “Establishing Connections.”

Related Functions

For information about	See
Translating data returned from the data source	SQLDataSourceToDriver
Returning the setting of a connection option	SQLGetConnectOption (extension)
Setting a connection option	SQLSetConnectOption (extension)

Appendixes

APPENDIX A

ODBC Error Codes

SQLError returns SQLSTATE values as defined by the X/Open and SQL Access Group SQL CAE specification (1992). SQLSTATE values are strings that contain five characters. The following table lists SQLSTATE values that a driver can return for **SQLError**.

The character string value returned for an SQLSTATE consists of a two character class value followed by a three character subclass value. A class value of "01" indicates a warning and is accompanied by a return code of SQL_SUCCESS_WITH_INFO. Class values other than "01", except for the class "IM", indicate an error and are accompanied by a return code of SQL_ERROR. The class "IM" is specific to warnings and errors that derive from the implementation of ODBC itself. The subclass value "000" in any class is for implementation defined conditions within the given class. The assignment of class and subclass values is defined by ANSI SQL-92.

Note: Although successful execution of a function is normally indicated by a return value of SQL_SUCCESS, the SQLSTATE 00000 also indicates success.

SQLSTATE	Error	Can be returned from
01000	General warning	All ODBC functions except: SQLAllocEnv SQLError
01002	Disconnect error	SQLDisconnect
SQLSTATE	Error	Can be returned from
01004	Data truncated	SQLBrowseConnect SQLColAttributes SQLDataSources SQLDescribeCol SQLDriverConnect SQLDrivers SQLExecDirect

		SQLExecute SQLExtendedFetch SQLFetch SQLGetCursorName SQLGetData SQLGetInfo SQLNativeSql SQLPutData SQLSetPos SQLExecDirect SQLExecute SQLBrowseConnect SQLDriverConnect SQLExtendedFetch SQLSetPos SQLSetConnectOption SQLSetStmtOption SQLExecDirect SQLExecute SQLSetPos SQLExecDirect SQLExecute SQLSetPos SQLExecDirect SQLExecute SQLBindParameter SQLExtendedFetch SQLFetch SQLGetData SQLBrowseConnect SQLConnect SQLDriverConnect SQLBrowseConnect SQLConnect SQLDriverConnect SQLSetConnectOption Can be returned from
01006	Privilege not revoked	
01S00	Invalid connection string attribute	
01S01	Error in row	
01S02	Option value changed	
01S03	No rows updated or deleted	
01S04	More than one row updated or deleted	
07001	Wrong number of parameters	
07006	Restricted data type attribute violation	
08001	Unable to connect to data source	
08002	Connection in use	
SQLSTATE	Error	
08003	Connection not open	SQLAllocStmt SQLDisconnect SQLGetConnectOption SQLGetInfo SQLNativeSql SQLSetConnectOption SQLTransact SQLBrowseConnect SQLConnect SQLDriverConnect
08004	Data source rejected establishment of connection	

08007	Connection failure during transaction	SQLTransact
08S01	Communication link failure	SQLBrowseConnect SQLColumnPrivileges SQLColumns SQLConnect SQLDriverConnect SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLForeignKeys SQLFreeConnect SQLGetData SQLGetTypeInfo SQLParamData SQLPrepare SQLPrimaryKeys SQLProcedureColumns SQLProcedures SQLPutData SQLSetConnectOption SQLSetStmtOption SQLSpecialColumns SQLStatistics SQLTablePrivileges SQLTables
21S01	Insert value list does not match column list	SQLExecDirect SQLPrepare
21S02	Degree of derived table does not match column list	SQLExecDirect SQLPrepare SQLSetPos

SQLSTATE	Error	Can be returned from
22001	String data right truncation	SQLPutData
22003	Numeric value out of range	SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLGetData SQLGetInfo SQLPutData SQLSetPos
22005	Error in assignment	SQLExecDirect

		SQLExecute SQLGetData SQLPrepare SQLPutData SQLSetPos
22008	Datetime field overflow	SQLExecDirect SQLExecute SQLGetData SQLPutData SQLSetPos
22012	Division by zero	SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch
22026	String data, length mismatch	SQLParamData
23000	Integrity constraint violation	SQLExecDirect SQLExecute SQLSetPos
SQLSTATE	Error	Can be returned from
24000	Invalid cursor state	SQLColAttributes SQLColumnPrivileges SQLColumns SQLDescribeCol SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLForeignKeys SQLGetData SQLGetStmtOption SQLGetTypeInfo SQLPrepare SQLPrimaryKeys SQLProcedureColumns SQLProcedures SQLSetCursorName SQLSetPos SQLSetStmtOption SQLSpecialColumns SQLStatistics SQLTablePrivileges SQLTables
25000	Invalid transaction state	SQLDisconnect
28000	Invalid authorization specification	SQLBrowseConnect SQLConnect

34000	Invalid cursor name	SQLDriverConnect SQLExecDirect SQLPrepare SQLSetCursorName
37000	Syntax error or access violation	SQLExecDirect SQLNativeSql SQLPrepare SQLSetCursorName
3C000	Duplicate cursor name	SQLSetCursorName
40001	Serialization failure	SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch
42000	Syntax error or access violation	SQLExecDirect SQLExecute SQLPrepare SQLSetPos

SQLSTATE	Error	Can be returned from
70100	Operation aborted	SQLCancel
IM001	Driver does not support this function	All ODBC functions except: SQLAllocConnect SQLAllocEnv SQLDataSources SQLDrivers SQLError SQLFreeConnect SQLFreeEnv SQLGetFunctions
IM002	Data source name not found and no default driver specified	SQLBrowseConnect SQLConnect SQLDriverConnect
IM003	Specified driver could not be loaded	SQLBrowseConnect SQLConnect SQLDriverConnect
IM004	Driver's SQLAllocEnv failed	SQLBrowseConnect SQLConnect SQLDriverConnect
IM005	Driver's SQLAllocConnect failed	SQLBrowseConnect SQLConnect SQLDriverConnect
IM006	Driver's SQLSetConnect-Option failed	SQLBrowseConnect SQLConnect SQLDriverConnect

IM007	No data source or driver specified; dialog prohibited	SQLDriverConnect
IM008	Dialog failed	SQLDriverConnect
IM009	Unable to load translation DLL	SQLBrowseConnect SQLConnect SQLDriverConnect SQLSetConnectOption
IM010	Data source name too long	SQLBrowseConnect SQLDriverConnect
IM011	Driver name too long	SQLBrowseConnect SQLDriverConnect
IM012	DRIVER keyword syntax error	SQLBrowseConnect SQLDriverConnect

SQLSTATE	Error	Can be returned from
IM013	Trace file error	All ODBC functions.
S0001	Base table or view already exists	SQLExecDirect SQLPrepare
S0002	Base table not found	SQLExecDirect SQLPrepare
S0011	Index already exists	SQLExecDirect SQLPrepare
S0012	Index not found	SQLExecDirect SQLPrepare
S0021	Column already exists	SQLExecDirect SQLPrepare
S0022	Column not found	SQLExecDirect SQLPrepare
S0023	No default for column	SQLSetPos
S1000	General error	All ODBC functions except: SQLAllocEnv SQLError
S1001	Memory allocation failure	All ODBC functions except: SQLAllocEnv SQLError SQLFreeConnect SQLFreeEnv
S1002	Invalid column number	SQLBindCol SQLColAttributes

		SQLDescribeCol SQLExtendedFetch SQLFetch SQLGetData
S1003	Program type out of range	SQLBindCol SQLBindParameter SQLGetData
S1004	SQL data type out of range	SQLBindParameter SQLGetTypeInfo
SQLSTATE	Error	Can be returned from
S1008	Operation canceled	All ODBC functions that can be processed asynchronously: SQLColAttributes SQLColumnPrivileges SQLColumns SQLDescribeCol SQLDescribeParam SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLForeignKeys SQLGetData SQLGetTypeInfo SQLMoreResults SQLNumParams SQLNumResultCols SQLParamData SQLPrepare SQLPrimaryKeys SQLProcedureColumns SQLProcedures SQLPutData SQLSetPos SQLSpecialColumns SQLStatistics SQLTablePrivileges SQLTables
S1009	Invalid argument value	SQLAllocConnect SQLAllocStmt SQLBindCol SQLBindParameter SQLExecDirect SQLForeignKeys SQLGetData SQLGetInfo

		SQLNativeSql SQLPrepare SQLPutData SQLSetConnectOption SQLSetCursorName SQLSetPos SQLSetStmtOption
SQLSTATE	Error	Can be returned from
S1010	Function sequence error	SQLBindCol SQLBindParameter SQLColAttributes SQLColumnPrivileges SQLColumns SQLDescribeCol SQLDescribeParam SQLDisconnect SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLForeignKeys SQLFreeConnect SQLFreeEnv SQLFreeStmt SQLGetConnectOption SQLGetCursorName SQLGetData SQLGetFunctions SQLGetStmtOption SQLGetTypeInfo SQLMoreResults SQLNumParams SQLNumResultCols SQLParamData SQLParamOptions SQLPrepare SQLPrimaryKeys SQLProcedureColumns SQLProcedures SQLPutData SQLRowCount SQLSetConnectOption SQLSetCursorName SQLSetPos SQLSetScrollOptions SQLSetStmtOption SQLSpecialColumns

SQLSTATE	Error	SQLStatistics SQLTablePrivileges SQLTables SQLTransact
		Can be returned from
S1011	Operation invalid at this time	SQLGetStmtOption SQLSetConnectOption SQLSetStmtOption
S1012	Invalid transaction operation code specified	SQLTransact
S1015	No cursor name available	SQLGetCursorName
S1090	Invalid string or buffer length	SQLBindCol SQLBindParameter SQLBrowseConnect SQLColAttributes SQLColumnPrivileges SQLColumns SQLConnect SQLDataSources SQLDescribeCol SQLDriverConnect SQLDrivers SQLExecDirect SQLExecute SQLForeignKeys SQLGetCursorName SQLGetData SQLGetInfo SQLNativeSql SQLPrepare SQLPrimaryKeys SQLProcedureColumns SQLProcedures SQLPutData SQLSetCursorName SQLSetPos SQLSpecialColumns SQLStatistics SQLTablePrivileges SQLTables
S1091	Descriptor type out of range	SQLColAttributes

SQLSTATE	Error	Can be returned from
S1092	Option type out of range	SQLFreeStmt SQLGetConnectOption SQLGetStmtOption SQLSetConnectOption SQLSetStmtOption
S1093	Invalid parameter number	SQLBindParameter SQLDescribeParam
S1094	Invalid scale value	SQLBindParameter
S1095	Function type out of range	SQLGetFunctions
S1096	Information type out of range	SQLGetInfo
S1097	Column type out of range	SQLSpecialColumns
S1098	Scope type out of range	SQLSpecialColumns
S1099	Nullable type out of range	SQLSpecialColumns
S1100	Uniqueness option type out of range	SQLStatistics
S1101	Accuracy option type out of range	SQLStatistics
S1103	Direction option out of range	SQLDataSources SQLDrivers
S1104	Invalid precision value	SQLBindParameter
S1105	Invalid parameter type	SQLBindParameter
S1106	Fetch type out of range	SQLExtendedFetch
S1107	Row value out of range	SQLExtendedFetch SQLParamOptions SQLSetPos SQLSetScrollOptions
S1108	Concurrency option out of range	SQLSetScrollOptions
SQLSTATE	Error	Can be returned from
S1109	Invalid cursor position	SQLExecute SQLExecDirect SQLGetData SQLGetStmtOption SQLSetPos
S1110	Invalid driver completion	SQLDriverConnect
S1111	Invalid bookmark value	SQLExtendedFetch

S1C00	Driver not capable	SQLBindCol SQLBindParameter SQLColAttributes SQLColumnPrivileges SQLColumns SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLForeignKeys SQLGetConnectOption SQLGetData SQLGetInfo SQLGetStmtOption SQLGetTypeInfo SQLPrimaryKeys SQLProcedureColumns SQLProcedures SQLSetConnectOption SQLSetPos SQLSetScrollOptions SQLSetStmtOption SQLSpecialColumns SQLStatistics SQLTablePrivileges SQLTables SQLTransact
-------	--------------------	--

SQLSTATE	Error	Can be returned from
----------	-------	----------------------

S1T00	Timeout expired	SQLBrowseConnect SQLColAttributes SQLColumnPrivileges SQLColumns SQLConnect SQLDescribeCol SQLDescribeParam SQLDriverConnect SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLForeignKeys SQLGetData SQLGetInfo SQLGetTypeInfo SQLMoreResults SQLNumParams SQLNumResultCols
-------	-----------------	--

SQLParamData
SQLPrepare
SQLPrimaryKeys
SQLProcedureColumns
SQLProcedures
SQLPutData
SQLSetPos
SQLSpecialColumns
SQLStatistics
SQLTablePrivileges
SQLTables

APPENDIX B

ODBC State Transition Tables

The tables in this appendix show how ODBC functions cause transitions of the environment, connection, and statement states. Generally speaking, the state of the environment, connection, or statement dictates when functions that use the corresponding type of handle (*henv*, *hdbc*, or *hstmt*) can be called. The environment, connection, and statement states overlap as follows, although the exact overlap of connection states C5 and C6 and statement states S1 through S12 is data source–dependent, since transactions begin at different times on different data sources. For a description of each state, see "Environment Transitions," "Connection Transitions," and "Statement Transitions," later in this appendix.

Environment: E0 E1 _____ E2
Connection: C0 C1 C2 C3 C4 C5 _____ C6
Statement: S0 S1 S2 S3 S4 S5 S6 S7 S8 S9 S10 S11 S12

Each entry in a transition table can be one of the following values:

- **--**. The state is unchanged after executing the function.
- **En**, **Cn**, or **Sn**. The environment, connection, or statement state moves to the specified state.
- **(IH)**. The function returned SQL_INVALID_HANDLE. Although this error is possible in any state, it is shown only when it is the only possible outcome of calling the function in the specified state. This error does not change the state and is always detected by the Driver Manager, as indicated by the parentheses.
- **NS**. Next State. The statement transition is the same as if the statement had not gone through the asynchronous states. For example, suppose a statement that creates a result set enters state S11 from state S1 because **SQLExecDirect** returned SQL_STILL_EXECUTING. The NS notation in state S11 means that the transitions for the statement are the same as those for a statement in state S1 that creates a result set: if **SQLExecDirect** returns an error; the statement remains in state S1; if it succeeds, the statement moves to state S5; if it needs data, the statement moves to state S8; and if it is still executing, it remains in state S11.

- **XXXXXX** or **(XXXXXX)**. An SQLSTATE that is related to the transition table; SQLSTATES detected by the Driver Manager are enclosed in parentheses. The function returned SQL_ERROR and the specified SQLSTATE, but the state does not change. For example, if **SQLExecute** is called before **SQLPrepare**, it returns SQLSTATE S1010 (Function sequence error).

Note: The tables do not show errors unrelated to the transition tables that do not change the state. For example, when **SQLAllocConnect** is called in environment state E1 and returns SQLSTATE S1001 (Memory allocation failure), the environment remains in state E1; this is not shown in the environment transition table for **SQLAllocConnect**.

Note: The Driver Manager does not exist on non-Windows systems. The Driver returns the errors directly.

If the environment, connection, or statement can move to more than one state, each possible state is shown and one or more footnotes explains the conditions under which each transition takes place. The following footnotes may appear in any table:

Footnote Meaning

b	Before or after. The cursor was positioned before the start of the result set or after the end of the result set.
c	Current function. The current function was executing asynchronously.
d	Need data. The function returned SQL_NEED_DATA.
e	Error. The function returned SQL_ERROR.
i	Invalid row. The cursor was positioned on a row in the result set and the value in the <i>rgfRowStatus</i> array in SQLExtendedFetch for the row was SQL_DELETED or SQL_ERROR.
nf	Not found. The function returned SQL_NO_DATA_FOUND.
np	Not prepared. The statement was not prepared.
nr	No results. The statement will not or did not create a result set.
o	Other function. Another function was executing asynchronously.
p	Prepared. The statement was prepared.
r	Results. The statement will or did create a (possibly empty) result set.
s	Success. The function returned SQL_SUCCESS_WITH_INFO or SQL_SUCCESS.
v	Valid row. The cursor was positioned on a row in the result set and the value in the <i>rgfRowStatus</i> array in SQLExtendedFetch for the row was SQL_ADDED, SQL_SUCCESS, or SQL_UPDATED.
x	Executing. The function returned SQL_STILL_EXECUTING.

For example, the environment state transition table for **SQLFreeEnv** is:

SQLFreeEnv

E0 Unallocated	E1 Allocated	E2 hdbc
(IH)	E0	(S1010)

If **SQLFreeEnv** is called in environment state E0, the Driver Manager returns SQL_INVALID_HANDLE. If it is called in state E1, the environment moves to state E0 if the function succeeds and remains in state E1 if the function fails. If it is called in state E2, the Driver Manager always returns SQL_ERROR and SQLSTATE S1010 (Function sequence error) and the environment remains in state E2.

Environment Transitions

The ODBC environment has the following three states:

State	Description
E0	Unallocated <i>henv</i>
E1	Allocated <i>henv</i> , unallocated <i>hdbc</i>
E2	Allocated <i>henv</i> , allocated <i>hdbc</i>

The following tables show how each ODBC function affects the environment state.

SQLAllocConnect

E0 Unallocated	E1 Allocated	E2 hdbc
(IH)	E2	-- ¹

¹ Calling **SQLAllocConnect** with a pointer to a valid *hdbc* overwrites that *hdbc*. This may be an application programming error.

SQLAllocEnv

E0 Unallocated	E1 Allocated	E2 hdbc
E1	-- ¹	E1 ¹

¹ Calling **SQLAllocEnv** with a pointer to a valid *henv* overwrites that *henv*. This may be an application programming error.

SQLDataSources and SQLDrivers

E0 Unallocated	E1 Allocated	E2 <i>hdbc</i>
(IH)	--	--

SQLError

E0 Unallocated	E1 Allocated	E2 <i>hdbc</i>
(IH) ¹	--	--

1 This row shows transitions when *henv* was non-null, *hdbc* was SQL_NULL_HDBC, and *hstmt* was SQL_NULL_HSTMT.

SQLFreeConnect

E0 Unallocated	E1 Allocated	E2 <i>hdbc</i>
(IH)	(IH)	-- ¹ E1 ²

1 There were other allocated *hdbcs*.

2 The *hdbc* was the only allocated *hdbc*.

SQLFreeEnv

E0 Unallocated	E1 Allocated	E2 <i>hdbc</i>
(IH)	E0	(S1010)

SQLTransact

E0 Unallocated	E1 Allocated	E2 <i>hdbc</i>
(IH)	-- ¹	-- ¹

1 The *hdbc* argument was SQL_NULL_HDBC.

All Other ODBC Functions

E0 Unallocated	E1 Allocated	E2 <i>hdbc</i>
(IH)	(IH)	--

Connection Transitions

ODBC connections have the following states:

State	Description
C0	Unallocated <i>henv</i> , unallocated <i>hdbc</i>
C1	Allocated <i>henv</i> , unallocated <i>hdbc</i>
C2	Allocated <i>henv</i> , allocated <i>hdbc</i>
C3	Connection function needs data
C4	Connected <i>hdbc</i>
C5	Connected <i>hdbc</i> , allocated <i>hstmt</i>
C6	Connected <i>hdbc</i> , transaction in progress

The following tables show how each ODBC function affects the connection state.

SQLAllocConnect

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	C2	-- ¹	C2 ¹	C2 ¹	C2 ¹	C2 ¹

¹ Calling **SQLAllocConnect** with a pointer to a valid *hdbc* overwrites that *hdbc*. This may be an application programming error.

SQLAllocEnv

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
C1	-- ¹	C1 ¹	C1 ¹	C1 ¹	C1 ¹	C1 ¹

¹ Calling **SQLAllocEnv** with a pointer to a valid *henv* overwrites that *henv*. This may be an application programming error.

SQLAllocStmt

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	(IH)	(08003)	(08003)	C5	-- ¹	C5 ¹

¹ Calling **SQLAllocStmt** with a pointer to a valid *hstmt* overwrites that *hstmt*. This may be an application programming error.

SQLBrowseConnect

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	(IH)	C3 ^d C4 ^s	-- ^d C2 ^e C4 ^s	(08002)	(08002)	(08002)

SQLColumnPrivileges, SQLColumns, SQLForeignKeys, SQLGetTypeInfo, SQLPrimaryKeys, SQLProcedureColumns, SQLProcedures, SQLSpecialColumns, SQLStatistics, SQLTablePrivileges, and SQLTables

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	(IH)	(IH)	(IH)	(IH)	-- ¹ C6 ²	--

1 The data source was in auto-commit mode or did not begin a transaction.

2 The data source was in manual-commit mode and began a transaction.

SQLColumns: see SQLColumnPrivileges**SQLConnect and SQLDriverConnect**

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	(IH)	C4	(08002)	(08002)	(08002)	(08002)

SQLDataSources and SQLDrivers

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	--	--	--	--	--	--

SQLDisconnect

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	(IH)	(08003)	C2	C2	C2	25000

SQLDriverConnect: see **SQLConnect**

SQLDrivers: see **SQLDataSources**

SQLError

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH) ¹	(IH)	--	--	--	--	--

1 This row shows transitions when *hdbc* was non-null and *hstmt* was SQL_NULL_HSTMT.

SQLExecDirect and SQLExecute

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	(IH)	(IH)	(IH)	(IH)	-- ¹ C6 ²	--

1 The data source was in auto-commit mode or did not begin a transaction.

2 The data source was in manual-commit mode and began a transaction.

SQLExecute: see **SQLExecDirect**

SQLForeignKeys: see **SQLColumnPrivileges**

SQLFreeConnect

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	(IH)	C1	(S1010)	(S1010)	(S1010)	(S1010)

SQLFreeEnv

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	C0 ¹ (S1010) ²	(S1010)	(S1010)	(S1010)	(S1010)	(S1010)

1 The *hdbc* was the only allocated *hdbc*.

2 There were other allocated *hdbcs*.

SQLFreeStmt

C0 No henv	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 hstmt	C6 Transaction
(IH)	(IH)	(IH)	(IH)	(IH)	-- ¹ C4 ²	-- ¹ C4 ²

¹ The *fOption* argument was SQL_CLOSE, SQL_UNBIND, or SQL_RESET_PARAMS.

² The *fOption* argument was SQL_DROP.

SQLGetConnectOption

C0 No henv	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 hstmt	C6 Transaction
(IH)	(IH)	-- ¹ (08003) ²	(S1010)	--	--	--

¹ The *fOption* argument was SQL_ACCESS_MODE or SQL_AUTOCOMMIT, or a value had been set for the connection option.

² The *fOption* argument was not SQL_ACCESS_MODE or SQL_AUTOCOMMIT, and a value had not been set for the connection option.

SQLGetFunctions

C0 No henv	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 hstmt	C6 Transaction
(IH)	(IH)	(S1010)	(S1010)	--	--	--

SQLGetInfo

C0 No henv	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 hstmt	C6 Transaction
(IH)	(IH)	-- ¹ (08003) ²	(08003)	--	--	--

¹ The *fInfoType* argument was SQL_ODBC_VER.

² The *fInfoType* argument was not SQL_ODBC_VER.

SQLGetTypeInfo: see **SQLColumnPrivileges**

SQLNativeSql

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	(IH)	(08003)	(08003)	--	--	--

SQLPrepare

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	(IH)	(IH)	(IH)	(IH)	-- ¹ C6 ²	--

1 The data source was in auto-commit mode or did not begin a transaction.

2 The data source was in manual commit mode and began a transaction.

SQLPrimaryKeys: see **SQLColumnPrivileges**

SQLProcedureColumns: see **SQLColumnPrivileges**

SQLProcedures: see **SQLColumnPrivileges**

SQLSetConnectOption

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	(IH)	-- ¹ (08003) ²	(S1010)	-- ³ (08002) ⁴	-- ³ (08002) ⁴	-- ³ and ⁵ C5 ⁶ (08002) ⁴ S1011 ⁷

¹ The *fOption* argument was not SQL_TRANSLATE_DLL or SQL_TRANSLATE_OPTION.

² The *fOption* argument was SQL_TRANSLATE_DLL or SQL_TRANSLATE_OPTION.

³ The *fOption* argument was not SQL_ODBC_CURSORS.

⁴ The *fOption* argument was SQL_ODBC_CURSORS.

⁵ If the *fOption* argument was SQL_AUTOCOMMIT, then the data source was in manual-commit mode or the *vParam* argument was SQL_AUTOCOMMIT_OFF.

⁶ The data source was in manual-commit mode, the *fOption* argument was SQL_AUTOCOMMIT, and the *vParam* argument was SQL_AUTOCOMMIT_ON.

⁷ The data source was in manual-commit mode and the *fOption* argument was SQL_TXN_ISOLATION.

SQLSpecialColumns: see **SQLColumnPrivileges**

SQLStatistics: see **SQLColumnPrivileges**

SQLTablePrivileges: see **SQLColumnPrivileges**

SQLTables: see **SQLColumnPrivileges**

SQLTransact

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH) ¹	(IH)	(IH)	(IH)	(IH)	(IH)	(IH)
(IH) ²	--	(08003)	(08003)	--	--	-- ^e and ⁴ C5 ^s or ⁵
(IH) ³	(IH)	(08003)	(08003)	--	--	-- ^e C5 ^s

¹ This row shows transitions when *henv* was SQL_NULL_HENV and *hdbc* was SQL_NULL_HDBC.

² This row shows transitions when *henv* was a valid environment handle and *hdbc* was

SQL_NULL_HDBC.

3 This row shows transitions when *hdbc* was a valid connection handle.

4 The commit or rollback failed on the connection.

5 The function returned SQL_ERROR but the commit or rollback succeeded on the connection.

All Other ODBC Functions

C0 No <i>henv</i>	C1 Unallocated	C2 Allocated	C3 Need Data	C4 Connected	C5 <i>hstmt</i>	C6 Transaction
(IH)	(IH)	(IH)	(IH)	(IH)	--	--

Statement Transitions

ODBC statements have the following states:

State	Description
S0	Unallocated <i>hstmt</i> . (The connection state must be C4. For more information, see "Connection Transitions.")
S1	Allocated <i>hstmt</i> .
S2	Prepared statement. No result set will be created.
S3	Prepared statement. A (possibly empty) result set will be created.
S4	Statement executed and no result set was created.
S5	Statement executed and a (possibly empty) result set was created. The cursor is open and positioned before the first row of the result set.
S6	Cursor positioned with SQLFetch .
S7	Cursor positioned with SQLExtendedFetch .
S8	Function needs data. SQLParamData has not been called.
S9	Function needs data. SQLPutData has not been called.
S10	Function needs data. SQLPutData has been called.
S11	Still executing.
S12	Asynchronous execution canceled. In S12, an application must call the canceled function until it returns a value other than SQL_STILL_EXECUTING. The function was canceled successfully only if the function returns SQL_ERROR and SQLSTATE S1008 (Operation canceled). If it returns any other value, such as SQL_SUCCESS, the cancel operation failed and the function executed normally.

States S2 and S3 are known as the prepared states, states S5 through S7 as the cursor states, states S8 through S10 as the need data states, and states S11 and S12 as the asynchronous states. In each of these groups, the transitions are shown separately only when they are different for each state in the group; generally, the transitions for each state in each a group are the same.

The following tables show how each ODBC function affects the statement state.

SQLAllocConnect

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
-- ¹	S0 ¹	S0 ¹	S0 ¹	S0 ¹	S0 ¹	S0 ¹

¹ Calling **SQLAllocConnect** with a pointer to a valid *hdbc* overwrites that *hdbc*. This may be an application programming error. Furthermore, this returns the connection state to C2; the connection state must be C4 before the statement state is S0.

SQLAllocEnv

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
-- ¹	S0 ¹	S0 ¹	S0 ¹	S0 ¹	S0 ¹	S0 ¹

¹ Calling **SQLAllocEnv** with a pointer to a valid *henv* overwrites that *henv*. This may be an application programming error. Furthermore, this returns the connection state to C1; the connection state must be C4 before the statement state is S0.

SQLAllocStmt

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
S1	-- ¹	S1 ¹	S1 ¹	S1 ¹	S1 ¹	S1 ¹

¹ Calling **SQLAllocStmt** with a pointer to a valid *hstmt* overwrites that *hstmt*. This may be an application programming error.

SQLBindCol

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
(IH)	--	--	--	--	(S1010)	(S1010)

SQLBindParameter

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	--	--	--	--	(S1010)	(S1010)

SQLBrowseConnect, SQLConnect, and SQLDriverConnect

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(08002)	(08002)	(08002)	(08002)	(08002)	(08002)	(08002)

SQLCancel¹

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	--	--	S1 ^{np} S2 ^p	S1 ^{np} S3 ^p	S1 ² S2 ^{nr and 3} S3 ^{r and 3} S7 ⁴	S12

1 This table does not cover cancellation of a function running synchronously on one thread when an application calls **SQLCancel** on a different thread with the same *hstmt*. In this case, the driver must note that **SQLCancel** was called and return the correct return code and SQLSTATE (if any) from the synchronous function. The statement transition when that function finishes is NS (Next State). That is, the statement transition is the same as if the function completed processing normally; the only difference is that it is possible for the function to return SQL_ERROR and SQLSTATE S1008 (Operation canceled).

2 **SQLExecDirect** returned SQL_NEED_DATA.

3 **SQLExecute** returned SQL_NEED_DATA.

4 **SQLSetPos** returned SQL_NEED_DATA.

SQLColAttributes

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	(S1010)	see below	24000	-- ^s S11 ^x	(S1010)	NS ^c (S1010) ^o

SQLColAttributes (Prepared states)

S2	S3
No Results	Results
24000	-- ^s S11 ^x

**SQLColumnPrivileges, SQLColumns, SQLForeignKeys, SQLGetTypeInfo,
SQLPrimaryKeys, SQLProcedureColumns, SQLProcedures,
SQLSpecialColumns, SQLStatistics, SQLTablePrivileges, and SQLTables**

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
(IH)	S5 ^s S11 ^x	S1 ^e S5 ^s S11 ^x	S1 ^e S5 ^s S11 ^x	see below	(S1010)	NS ^c (S1010) ^o

**SQLColumnPrivileges, SQLColumns, SQLForeignKeys, SQLGetTypeInfo,
SQLPrimaryKeys, SQLProcedureColumns, SQLProcedures,
SQLSpecialColumns, SQLStatistics, SQLTablePrivileges, and SQLTables** (Cursor
states)

S5	S6	S7
Opened	SQLFetch	SQLExtendedFetch
24000	(24000)	(24000)

SQLColumns: see **SQLColumnPrivileges**

SQLConnect: see **SQLBrowseConnect**

SQLDataSources and **SQLDrivers**

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
--	--	--	--	--	--	--

SQLDescribeCol

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
(IH)	(S1010)	see below	24000	-- ^s S11 ^x	(S1010)	NS ^c (S1010) ^o

SQLDescribeCol (Prepared states)

S2	S3
No Results	Results
24000	-- ^s S11 ^x

SQLDescribeParam

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
(IH)	(S1010)	-- ^s S11 ^x	-- ^s S11 ^x	-- ^s S11 ^x	(S1010)	NS ^c (S1010) ^o

SQLDisconnect

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
-- ¹	S0 ¹	S0 ¹	S0 ¹	S0 ¹	(S1010)	(S1010)

¹ Calling **SQLDisconnect** frees all *hstmts* associated with the *hdbc*. Furthermore, this returns the connection state to C2; the connection state must be C4 before the statement state is S0.

SQLDriverConnect: see **SQLBrowseConnect**

SQLDrivers: see **SQLDataSources**

SQLError

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
(IH) ¹	--	--	--	--	--	--

¹ This row shows transitions when *hstmt* was non-null.

SQLExecDirect

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
(IH)	S4 ^s and nr S5 ^s and r S8 ^d S11 ^x	S1 ^e S4 ^s and nr S5 ^s and r S8 ^d S11 ^x	S1 ^e S4 ^s and nr S5 ^s and r S8 ^d S11 ^x	see below	(S1010)	NS ^c (S1010) ^o

SQLExecDirect (Cursor states)

S5	S6	S7
Opened	SQLFetch	SQLExtendedFetch
24000	(24000)	(24000)

SQLExecute

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
(IH)	(S1010)	see below	S2 ^e and p S4 ^{s, p, and nr} S8 ^d and p S11 ^x and p (S1010) ^{np}	see below	(S1010)	NS ^c (S1010) ^o

SQLExecute (Prepared states)

S2	S3
No Results	Results
S4 ^s S8 ^d S11 ^x	S5 ^s S8 ^d S11 ^x

SQLExecute (Cursor states)

S5	S6	S7
Opened	SQLFetch	SQLExtendedFetch
24000 ^p (S1010) ^{np}	(24000) ^p (S1010) ^{np}	(24000) ^p (S1010) ^{np}

SQLExtendedFetch

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
(IH)	(S1010)	(S1010)	24000	see below	(S1010)	NS ^c (S1010) ^o

SQLExtendedFetch (Cursor states)

S5 Opened	S6 SQLFetch	S7 SQLExtendedFetch
S7 ^{s or nf} S11 ^x	(S1010)	-- ^{s or nf} S11 ^x

SQLFetch

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	(S1010)	(S1010)	24000	see below	(S1010)	NS ^c (S1010) ^o

SQLFetch (Cursor states)

S5 Opened	S6 SQLFetch	S7 SQLExtendedFetch
S6 ^{s or nf} S11 ^x	-- ^{s or nf} S11 ^x	(S1010)

SQLForeignKeys: see **SQLColumnPrivileges**

SQLFreeConnect

[illegible]

SQLFreeEnv

[illegible]

SQLFreeStmt

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH) ¹	--	--	S1 ^{np} S2 ^p	S1 ^{np} S3 ^p	(S1010)	(S1010)
(IH) ²	S0	S0	S0	S0	(S1010)	(S1010)
(IH) ³	--	--	--	--	(S1010)	(S1010)

1 This row shows transitions when *fOption* was SQL_CLOSE.

2 This row shows transitions when *fOption* was SQL_DROP.

3 This row shows transitions when *fOption* was SQL_UNBIND or SQL_RESET_PARAMS.

SQLGetConnectOption

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
--	--	--	--	--	--	--

SQLGetCursorName

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	-- ¹ (S1015) ²	-- ¹ (S1015) ²	-- ¹ (S1015) ²	--	(S1010)	(S1010)

1 A cursor name had been set by calling **SQLSetCursorName** or by creating a result set.

2 A cursor name had not been set by calling **SQLSetCursorName** or by creating a result set.

SQLGetData

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	(S1010)	(S1010)	(24000)	see below	(S1010)	NS ^c (S1010) ^o

SQLGetData (Cursor states)

S5 Opened	S6 SQLFetch	S7 SQLExtendedFetch
(24000)	-- s or nf S11 ^x 24000 ³	-- s or nf S11 ^x 24000 ^b S1109 ⁱ S1C00 ^{v and 1}

1 The rowset size was greater than 1 and the **SQLGetInfo** did not return the SQL_GD_BLOCK bit for the SQL_GETDATA_EXTENSIONS information type.

SQLGetFunctions

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
--	--	--	--	--	--	--

SQLGetInfo

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
--	--	--	--	--	--	--

SQLGetStmtOption

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	-- ¹ (24000) ²	-- ¹ (24000) ²	-- ¹ (24000) ²	see below	(S1010)	(S1010)

1 The statement option was not SQL_ROW_NUMBER or SQL_GET_BOOKMARK.

2 The statement option was SQL_ROW_NUMBER or SQL_GET_BOOKMARK.

SQLGetStmtOption (Cursor states)

S5 Opened	S6 SQLFetch	S7 SQLExtendedFetch
-- 1 (24000) ^{2 or 3}	-- 1 or (v and 3) 24000 ^{b and 3} S1011 ² S1109 ^{i and 3}	-- 1 or (v and (2 or 3)) 24000 ^{b and (2 or 3)} S1109 ^{i and (2 or 3)}

1 The *fOption* argument was not SQL_GET_BOOKMARK or SQL_ROW_NUMBER.

2 The *fOption* argument was SQL_GET_BOOKMARK.

3 The *fOption* argument was SQL_ROW_NUMBER.

SQLGetTypeInfo: see **SQLColumnPrivileges****SQLMoreResults**

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor ¹	S8 – S10 Need Data	S11 – S12 Async
(IH)	-- ²	-- ²	S1 ^{nf and np} S2 ^{nf and p} S11 ^x	-- ^s S1 ^{nf and np} S3 ^{nf and p} S11 ^x	(S1010)	NS ^c (S1010) ^o

1 For **SQLMoreResults**, the cursor states are somewhat modified. A batch statement, statement submitted with an array of parameters, or procedure can return result sets and **INSERT**, **UPDATE**, or **DELETE** row counts; these are collectively known as results. After the statement is executed, it moves to the S4 state if there are no results and the S5 state if there are results. In the S5 state, **SQLFetch** or **SQLExtendedFetch** returns SQLSTATE 24000 (Invalid cursor state) if the next result is a row count; the application must call **SQLMoreResults** to stop processing the current result and proceed to the next result. After the last result is processed, **SQLMoreResults** returns the statement to the allocated or prepared state.

2 The function always returns SQL_NO_DATA_FOUND in this state.

SQLNativeSql

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
--	--	--	--	--	--	--

SQLNumParams

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	(S1010)	-- ^s S11 ^x	-- ^s S11 ^x	-- ^s S11 ^x	(S1010)	NS ^c (S1010) ^o

SQLNumResultCols

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	(S1010)	-- ^s S11 ^x	-- ^s S11 ^x	-- ^s S11 ^x	(S1010)	NS ^c (S1010) ^o

SQLParamData

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	(S1010)	(S1010)	(S1010)	(S1010)	see below	NS ^c (S1010) ^o

SQLParamData (Need Data states)

S8 Need Data	S9 Must Put	S10 Can Put
S1 ^e and 1 S2 ^e , nr, and 2 S3 ^e , r, and 2 S7 ^e and 3 S9 ^s S11 ^x	S1010	S1 ^e and 1 S2 ^e , nr, and 2 S3 ^e , r, and 2 S4 ^s , nr, and (1 or 2) S5 ^s , r, and (1 or 2) S7 (s or e) and 3 S9 ^d S11 ^x

1 **SQLExecDirect** returned SQL_NEED_DATA.

2 **SQLExecute** returned SQL_NEED_DATA.

3 **SQLSetPos** returned SQL_NEED_DATA.

SQLParamOptions

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	--	--	--	--	(S1010)	(S1010)

SQLPrepare

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	S2 s and nr S3 s and r S11 x	-- s or (e and 1) S1 e and 2 S11 x	S1 e S2 s and nr S3 s and r S11 x	see below	(S1010)	NS ^c (S1010) ^o

1 The preparation fails for a reason other than validating the statement (in other words, the SQLSTATE was S1009 (Invalid argument value) or S1090 (Invalid string or buffer length)).

2 The preparation fails while validating the statement (in other words, the SQLSTATE was not S1009 (Invalid argument value) or S1090 (Invalid string or buffer length)).

SQLPrepare (Cursor states)

S5 Opened	S6 SQLFetch	S7 SQLExtendedFetch
24000	(24000)	(24000)

SQLPrimaryKeys: see **SQLColumnPrivileges**

SQLProcedureColumns: see **SQLColumnPrivileges**

SQLProcedures: see **SQLColumnPrivileges**

SQLPutData

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	(S1010)	(S1010)	(S1010)	(S1010)	see below	NS ^c (S1010) ^o

SQLPutData (Need Data states)

S8 Need Data	S9 Must Put	S10 Can Put
S1010	S1 ^e and 1 S2 ^e , nr, and 2 S3 ^e , r, and 2 S7 ^e and 3 S10 ^s S11 ^x	-- ^s S1 ^e and 1 S2 ^e , nr, and 2 S3 ^e , r, and 2 S7 ^e and 3 S11 ^x

1 **SQLExecDirect** returned SQL_NEED_DATA.

2 **SQLExecute** returned SQL_NEED_DATA.

3 **SQLSetPos** returned SQL_NEED_DATA.

SQLRowCount

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	(S1010)	(S1010)	--	--	(S1010)	(S1010)

SQLSetConnectOption

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
-- ¹	--	--	--	--	(S1010)	(S1010)

1 This row shows transitions when *fOption* was a connection option. For transitions when *fOption* was a statement option, see the statement transition table for **SQLSetStmtOption**.

SQLSetCursorName

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	--	--	(24000)	(24000)	(S1010)	(S1010)

SQLSetPos

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	(S1010)	(S1010)	(24000)	see below	(S1010)	NS ^c (S1010) ^o

SQLSetPos (Cursor states)

S5 Opened	S6 SQLFetch	S7 SQLExtendedFetch
(24000)	(S1010)	-- ^s S8 ^d S11 ^x 24000 ^b S1109 ⁱ

SQLSetScrollOptions

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	--	(S1010)	(S1010)	(S1010)	(S1010)	(S1010)

SQLSetStmtOption

S0 Unallocated	S1 Allocated	S2 – S3 Prepared	S4 Executed	S5 – S7 Cursor	S8 – S10 Need Data	S11 – S12 Async
(IH)	--	-- ¹ (S1011) ²	-- ¹ (24000) ²	-- ¹ (24000) ²	(S1010) ^{np or} ₁ (S1011) ^{p and} ₂	(S1010) ^{np or} ₁ (S1011) ^{p and} ₂

¹ The *fOption* argument was not SQL_CONCURRENCY, SQL_CURSOR_TYPE, SQL_SIMULATE_CURSOR, or SQL_USE_BOOKMARKS.

² The *fOption* argument was SQL_CONCURRENCY, SQL_CURSOR_TYPE, SQL_SIMULATE_CURSOR, or SQL_USE_BOOKMARKS.

SQLSpecialColumns: see **SQLColumnPrivileges**

SQLStatistics: see **SQLColumnPrivileges**

SQLTablePrivileges: see **SQLColumnPrivileges**

SQLTables: see **SQLColumnPrivileges**

SQLTransact

S0	S1	S2 – S3	S4	S5 – S7	S8 – S10	S11 – S12
Unallocated	Allocated	Prepared	Executed	Cursor	Need Data	Async
--	--	-- 2 or 3 S1 ¹	-- 3 S1 np and(1 or 2) S1 p and 1 S2 p and 2	-- 3 S1 np and(1 or 2) S1 p and 1 S3 p and 2	(S1010)	(S1010)

1 The *fType* argument is SQL_COMMIT and **SQLGetInfo** returns SQL_CB_DELETE for the SQL_CURSOR_COMMIT_BEHAVIOR information type, or the *fType* argument is SQL_ROLLBACK and **SQLGetInfo** returns SQL_CB_DELETE for the SQL_CURSOR_ROLLBACK_BEHAVIOR information type.

2 The *fType* argument is SQL_COMMIT and **SQLGetInfo** returns SQL_CB_CLOSE for the SQL_CURSOR_COMMIT_BEHAVIOR information type, or the *fType* argument is SQL_ROLLBACK and **SQLGetInfo** returns SQL_CB_CLOSE for the SQL_CURSOR_ROLLBACK_BEHAVIOR information type.

3 The *fType* argument is SQL_COMMIT and **SQLGetInfo** returns SQL_CB_PRESERVE for the SQL_CURSOR_COMMIT_BEHAVIOR information type, or the *fType* argument is SQL_ROLLBACK and **SQLGetInfo** returns SQL_CB_PRESERVE for the SQL_CURSOR_ROLLBACK_BEHAVIOR information type.

A P P E N D I X C

SQL Grammar

The following paragraphs list the recommended constructs to ensure interoperability in calls to **SQLPrepare**, **SQLExecute**, or **SQLExecDirect**. To the right of each construct is an indicator that tells whether the construct is part of the minimum grammar, the core grammar, or the extended grammar. ODBC does not prohibit the use of vendor-specific SQL grammar.

The Integrity Enhancement Facility (IEF) is included in the grammar but is optional. If drivers parse and execute SQL directly and wish to include referential integrity functionality, then we strongly recommend the SQL syntax used for this functionality conform to the grammar used here. The grammar for the IEF is taken directly from the X/Open and SQL Access Group SQL CAE specification (1992) and is a subset of the emerging ISO SQL-92 standard. Elements that are part of the IEF and are optional in the ANSI 1989 standard are presented in the following typeface and font, distinct from the rest of the grammar:

table-constraint-definition

A given driver and data source do not necessarily support all of the data types defined in this grammar. To determine which data types a driver supports, an application calls **SQLGetInfo** with the `SQL_ODBC_SQL_CONFORMANCE` flag. Drivers that support every core data type return 1 and drivers that support every core and every extended data type return 2. To determine whether a specific data type is supported, an application calls **SQLGetTypeInfo** with the *fSqlType* argument set to that data type.

If a driver supports data types that map to the ODBC SQL date, time, or timestamp data types, the driver must also support the extended SQL grammar for specifying date, time, or timestamp literals.

Note: In **CREATE TABLE** and **ALTER TABLE** statements, applications must use the data type name returned by **SQLGetTypeInfo** in the `TYPE_NAME` column.

Parameter Data Types

Even though each parameter specified with **SQLBindParameter** is defined using an SQL data type, the parameters in an SQL statement have no intrinsic data type. Therefore, parameter markers can be included in an SQL statement only if their data types can be inferred from another operand in the statement. For example, in an arithmetic expression such as **? + COLUMN1**, the data type of the parameter can be inferred from the data type of the named column represented by COLUMN1. An application cannot use a parameter marker if the data type cannot be determined.

The following table describes how a data type is determined for several types of parameters.

Location of Parameter	Assumed Data Type
One operand of a binary arithmetic or comparison operator	Same as the other operand
The first operand in a BETWEEN clause	Same as the other operand
The second or third operand in a BETWEEN clause	Same as the first operand
An expression used with IN	Same as the first value or the result column of the subquery
A value used with IN	Same as the expression
A pattern value used with LIKE	VARCHAR
An update value used with UPDATE	Same as the update column

Parameter Markers

An application cannot place parameter markers in the following locations:

- In a **SELECT** list.
- As both *expressions* in a *comparison-predicate*.
- As both operands of a binary operator.
- As both the first and second operands of a **BETWEEN** operation.
- As both the first and third operands of a **BETWEEN** operation.
- As both the expression and the first value of an **IN** operation.
- As the operand of a unary + or – operation.
- As the argument of a *set-function-reference*.

For more information, see the ANSI SQL-92 specification.

If an application includes parameter markers in the SQL statement, the application must call **SQLBindParameter** to associate storage locations with parameter markers before it calls **SQLExecute** or **SQLExecDirect**. If the application calls **SQLPrepare**, the application can call **SQLBindParameter** before or after it calls **SQLPrepare**.

The application can set parameter markers in any order. The driver buffers argument descriptors and sends the current values referenced by the **SQLBindParameter** argument *rgbValue* for the associated parameter marker when the application calls **SQLExecute** or **SQLExecDirect**. It is the application's responsibility to ensure that all pointer arguments are valid at execution time.

Note: The keyword **USER** in the following tables represents a character string containing the *user-name* of the current user.

SQL Statements

The following SQL statements define the base ODBC SQL grammar.

Statement	Mini- mum	Core	Ex- tende d
<i>alter-table-statement</i> ::= ALTER TABLE <i>base-table-name</i> { ADD <i>column-identifier data-type</i> ADD (<i>column-identifier data-type</i> [, <i>column-identifier data-type</i>]...) }		■	
Important: As a <i>data-type</i> in an <i>alter-table-statement</i> , applications must use a data type from the TYPE_NAME column of the result set returned by SQLGetTypeInfo .			
<i>alter-table-statement</i> ::= ALTER TABLE <i>base-table-name</i> { ADD <i>column-identifier data-type</i> ADD (<i>column-identifier data-type</i> [, <i>column-identifier data-type</i>]...) DROP [COLUMN] <i>column-identifier</i> [CASCADE RESTRICT] }			■
Important: As a <i>data-type</i> in an <i>alter-table-statement</i> , applications must use a data type from the TYPE_NAME column of the result set returned by SQLGetTypeInfo .			
<i>create-index-statement</i> ::= CREATE [UNIQUE] INDEX <i>index-name</i> ON <i>base-table-name</i> (<i>column-identifier</i> [ASC DESC] [, <i>column-identifier</i> [ASC DESC]]...)			■
Statement	Mini- mum	Core	Ex- tende d
<i>create-table-statement</i> ::= CREATE TABLE <i>base-table-name</i> (<i>column-element</i> [, <i>column-element</i>] ...) <i>column-element</i> ::= <i>column-definition</i> <i>table-constraint-definition</i> <i>column-definition</i> ::= <i>column-identifier data-type</i> [DEFAULT <i>default-value</i>]		■	

[*column-constraint-definition* [*column-constraint-definition*]...]

default-value ::= *literal* | NULL | USER

column-constraint-definition ::=

NOT NULL

| UNIQUE | PRIMARY KEY

| REFERENCES *ref-table-name* *referenced-columns*

| CHECK (*search-condition*)

table-constraint-definition ::=

UNIQUE (*column-identifier* [, *column-identifier*] ...)

| PRIMARY KEY (*column-identifier*

[, *column-identifier*] ...)

| CHECK (*search-condition*)

| FOREIGN KEY *referencing-columns* REFERENCES
ref-table-name *referenced-columns*

Important: As a *data-type* in a *create-table-statement*, applications must use a data type from the TYPE_NAME column of the result set returned by SQLGetTypeInfo.

create-view-statement ::=

CREATE VIEW *viewed-table-name*

[(*column-identifier* [, *column-identifier*] ...)]

AS *query-specification*

delete-statement-positioned ::=

DELETE FROM *table-name* WHERE CURRENT OF *cursor-name*

■ ■
ODBC ODBC
1.0 2.0

delete-statement-searched ::=

DELETE FROM *table-name* [WHERE *search-condition*]

drop-index-statement ::=

DROP INDEX *index-name*

drop-table-statement ::=

DROP TABLE *base-table-name*

[CASCADE | RESTRICT]

Statement

Mini-
Core
Ex-
tended

drop-view-statement ::=

DROP VIEW *viewed-table-name*

[CASCADE | RESTRICT]

grant-statement ::=

GRANT {ALL | *grant-privilege* [, *grant-privilege*]... }

ON *table-name*

<p>TO {PUBLIC <i>user-name</i> [, <i>user-name</i>]... }</p> <p><i>grant-privilege</i> ::=</p> <p>DELETE</p> <p> INSERT</p> <p> SELECT</p> <p> UPDATE [(<i>column-identifier</i> [, <i>column-identifier</i>]...)]</p> <p> REFERENCES [(<i>column-identifier</i></p> <p> [, <i>column-identifier</i>]...)]</p>				
<p><i>insert-statement</i> ::=</p> <p>INSERT INTO <i>table-name</i> [(<i>column-identifier</i> [, <i>column-identifier</i>]...)]</p> <p>VALUES (<i>insert-value</i> [, <i>insert-value</i>]...)</p>			■	
<p><i>insert-statement</i> ::=</p> <p>INSERT INTO <i>table-name</i> [(<i>column-identifier</i> [, <i>column-identifier</i>]...)]</p> <p>{ <i>query-specification</i> VALUES (<i>insert-value</i> [, <i>insert-value</i>]...)} </p>				■
<p><i>ODBC-procedure-extension</i> ::=</p> <p>ODBC-std-esc-initiator [?=] call procedure ODBC-std-esc-terminator</p> <p> ODBC-ext-esc-initiator [?=] call procedure ODBC-ext-esc-terminator</p>				■
<p><i>revoke-statement</i> ::=</p> <p>REVOKE {ALL <i>revoke-privilege</i> [, <i>revoke-privilege</i>]... }</p> <p>ON <i>table-name</i></p> <p>FROM {PUBLIC <i>user-name</i> [, <i>user-name</i>]... }</p> <p>[CASCADE RESTRICT]</p> <p><i>revoke-privilege</i> ::=</p> <p>DELETE</p> <p> INSERT</p> <p> SELECT</p> <p> UPDATE</p> <p> REFERENCES</p>				■
<p><i>select-statement</i> ::=</p> <p>SELECT [ALL DISTINCT] <i>select-list</i></p> <p>FROM <i>table-reference-list</i></p> <p>[WHERE <i>search-condition</i>]</p> <p>[<i>order-by-clause</i>]</p>			■	
Statement	Mini- mum	Core	Ex- tende d	
<p><i>select-statement</i> ::=</p> <p>SELECT [ALL DISTINCT] <i>select-list</i></p> <p>FROM <i>table-reference-list</i></p> <p>[WHERE <i>search-condition</i>]</p> <p>[GROUP BY <i>column-name</i> [, <i>column-name</i>]...]</p> <p>[HAVING <i>search-condition</i>]</p> <p>[<i>order-by-clause</i>]</p>			■	


```

select-statement ::=
    SELECT [ALL | DISTINCT] select-list
    FROM table-reference-list
    [WHERE search-condition]
    [GROUP BY column-name [, column-name]... ]
    [HAVING search-condition]
    [UNION [ALL] select-statement]...
    [order-by-clause]

```

(In ODBC 1.0, the **UNION** clause was in the Core SQL grammar and did not support the **ALL** keyword.)

```

select-for-update-statement ::=
    SELECT [ALL | DISTINCT] select-list
    FROM table-reference-list
    [WHERE search-condition]
    FOR UPDATE OF [column-name [, column-name]...]

```

```

statement ::= create-table-statement
            | delete-statement-searched
            | drop-table-statement
            | insert-statement
            | select-statement
            | update-statement-searched

```

```

statement ::= alter-table-statement
            | create-index-statement
            | create-table-statement
            | create-view-statement
            | delete-statement-searched
            | drop-index-statement
            | drop-table-statement
            | drop-view-statement
            | grant-statement
            | insert-statement
            | revoke-statement
            | select-statement
            | update-statement-searched

```

Statement	Mini- mum	Core	Ex- tende d
<pre> statement ::= alter-table-statement create-index-statement create-table-statement create-view-statement delete-statement-positioned delete-statement-searched drop-index-statement drop-table-statement </pre>			■

```

| drop-view-statement
| grant-statement
| insert-statement
| ODBC-procedure-extension
| revoke-statement
| select-statement
| select-for-update-statement
| statement-list
| update-statement-positioned
| update-statement-searched

```

(In ODBC 1.0, *select-for-update-statement*, *update-statement-positioned*, and *delete-statement-positioned* were in the Core SQL grammar.)

<i>statement-list</i> ::= <i>statement</i> <i>statement</i> ; <i>statement-list</i>			■
<i>update-statement-positioned</i> ::=		■	■
UPDATE <i>table-name</i>			
SET <i>column-identifier</i> = { <i>expression</i> NULL}		ODBC	ODBC
[, <i>column-identifier</i> = { <i>expression</i> NULL}]...		1.0	2.0
WHERE CURRENT OF <i>cursor-name</i>			
<i>update-statement-searched</i>			■
UPDATE <i>table-name</i>			
SET <i>column-identifier</i> = { <i>expression</i> NULL }			
[, <i>column-identifier</i> = { <i>expression</i> NULL}]...			
[WHERE <i>search-condition</i>]			

Elements Used in SQL Statements

The following elements are used in the SQL statements listed previously .

Element	Mini- mum	Core	Ex- tende d
<i>all-function</i> ::= {AVG MAX MIN SUM} (<i>expression</i>)		■	
<i>approximate-numeric-literal</i> ::= <i>mantissa</i> <i>E</i> <i>exponent</i>		■	
<i>approximate-numeric-type</i> ::= {approximate numeric types} (For example, FLOAT, DOUBLE PRECISION, or REAL. To determine the type name used by a data source, an application calls SQLGetTypeInfo .)		■	
<i>argument-list</i> ::= <i>expression</i> <i>expression</i> , <i>argument-list</i>	■		
<i>base-table-identifier</i> ::= <i>user-defined-name</i>	■		
<i>base-table-name</i> ::= <i>base-table-identifier</i>	■		
<i>base-table-name</i> ::= <i>base-table-identifier</i> <i>owner-name</i> . <i>base-table-identifier</i> <i>qualifier-name</i> <i>qualifier-separator</i> <i>base-table-identifier</i> <i>qualifier-name</i> <i>qualifier-separator</i> [<i>owner-name</i>]. <i>base-table-identifier</i> (The third syntax is valid only if the data source does not support owners.)		■	
<i>between-predicate</i> ::= <i>expression</i> [NOT] BETWEEN <i>expression</i> AND <i>expression</i>		■	
<i>binary-literal</i> ::= {implementation defined}			■
<i>binary-type</i> ::= {binary types} (For example, BINARY, VARBINARY, or LONG VARBINARY. To determine the type name used by a data source, an application calls SQLGetTypeInfo .)			■
<i>bit-literal</i> ::= 0 1			■
<i>bit-type</i> ::= {bit types} (For example, BIT. To determine the type name used by a data source, an application calls SQLGetTypeInfo .)			■
<i>boolean-factor</i> ::= [NOT] <i>boolean-primary</i>	■		
<i>boolean-primary</i> ::= <i>predicate</i> (<i>search-condition</i>)	■		
<i>boolean-term</i> ::= <i>boolean-factor</i> [AND <i>boolean-term</i>]	■		
<i>character</i> ::= {any character in the implementor's character set}	■		
<i>character-string-literal</i> ::= '{ <i>character</i> }...' (To include a single literal quote character (') in a <i>character-string-literal</i> , use two literal quote characters (' ').)	■		

Element	Mini- mum	Core	Ex- tende d
<i>character-string-type</i> ::= {character types} (The Minimum SQL conformance level requires at least one character data type. For example, CHAR, VARCHAR, or LONG VARCHAR. To determine the type name used by a data source, an application calls SQLGetTypeInfo .)	■		
<i>column-alias</i> ::= <i>user-defined-name</i>		■	
<i>column-identifier</i> ::= <i>user-defined-name</i>	■		
<i>column-name</i> ::= [<i>table-name</i>]. <i>column-identifier</i>	■		
<i>column-name</i> ::= [{ <i>table-name</i> <i>correlation-name</i> }.] <i>column-identifier</i>		■	
<i>comparison-operator</i> ::= < > <= >= = <>	■		
<i>comparison-predicate</i> ::= <i>expression comparison-operator expression</i>	■		
<i>comparison-predicate</i> ::= <i>expression comparison-operator {expression (sub-query)}</i>		■	
<i>correlation-name</i> ::= <i>user-defined-name</i>		■	
<i>cursor-name</i> ::= <i>user-defined-name</i>		■	
<i>data-type</i> ::= <i>character-string-type</i>	■		
<i>data-type</i> ::= <i>character-string-type</i> <i>exact-numeric-type</i> <i>approximate-numeric-type</i>		■	
<i>data-type</i> ::= <i>character-string-type</i> <i>exact-numeric-type</i> <i>approximate-numeric-type</i> <i>bit-type</i> <i>binary-type</i> <i>date-type</i> <i>time-type</i> <i>timestamp-type</i>			■
<i>date-separator</i> ::= -			■
<i>date-type</i> ::= {date types} (For example, DATE. To determine the type name used by a data source, an application calls SQLGetTypeInfo .)			■
<i>date-value</i> ::= <i>years-value date-separator months-value date-separator days-value</i>			■
<i>days-value</i> ::= <i>digit digit</i>			■
<i>digit</i> ::= 0 1 2 3 4 5 6 7 8 9	■		

ADABAS D		503
Element	Mini- mum	Core Ex- tende d
<i>distinct-function</i> ::= {AVG COUNT MAX MIN SUM} (DISTINCT <i>columnname</i>)		■
<i>dynamic-parameter</i> ::= ?	■	
<i>empty-string</i> ::=		■
<i>escape-character</i> ::= <i>character</i>		■
<i>exact-numeric-literal</i> ::= [+ -] { <i>unsigned-integer</i> [<i>unsigned-integer</i>] <i>unsigned-integer</i> . <i>unsigned-integer</i> }		■
<i>exact-numeric-type</i> ::= {exact numeric types} (For example, DECIMAL, NUMERIC, SMALLINT, or INTEGER. To determine the type name used by a data source, an application calls SQLGetTypeInfo .)		■
<i>exact-numeric-type</i> ::= {exact numeric types} (For example, DECIMAL, NUMERIC, SMALLINT, INTEGER, and BIGINT. To determine the type name used by a data source, an application calls SQLGetTypeInfo .)		■
<i>exists-predicate</i> ::= EXISTS (<i>sub-query</i>)		■
<i>exponent</i> ::= [+ -] <i>unsigned-integer</i>		■
<i>expression</i> ::= <i>term</i> <i>expression</i> {+ -} <i>term</i>	■	
<i>factor</i> ::= [+ -] <i>primary</i>	■	
<i>hours-value</i> ::= <i>digit digit</i>		■
<i>index-identifier</i> ::= <i>user-defined-name</i>		■
<i>index-name</i> ::= [<i>index-qualifier</i> .] <i>index-identifier</i>		■
<i>index-qualifier</i> ::= <i>user-defined-name</i>		■
<i>in-predicate</i> ::= <i>expression</i> [NOT] IN {(value {, value}...) (<i>subquery</i>)}		■
<i>insert-value</i> ::= <i>dynamic-parameter</i> <i>literal</i> NULL USER	■	
<i>keyword</i> ::= (see list of reserved keywords)	■	
<i>length</i> ::= <i>unsigned-integer</i>	■	
<i>letter</i> ::= <i>lower-case-letter</i> <i>upper-case-letter</i>	■	
	Mini-	Ex-

Element	mum	Core	tende d
<i>like-predicate</i> ::= <i>expression</i> [NOT] LIKE <i>pattern-value</i>	■		
<i>like-predicate</i> ::= <i>expression</i> [NOT] LIKE <i>pattern-value</i> [ODBC-like-escape-clause]			■
<i>literal</i> ::= <i>character-string-literal</i>	■		
<i>literal</i> ::= <i>character-string-literal</i> <i>numeric-literal</i>		■	
<i>literal</i> ::= <i>character-string-literal</i> <i>numeric-literal</i> <i>bit-literal</i> <i>binary-literal</i> <i>ODBC-date-time-extension</i>			■
<i>lower-case-letter</i> ::= a b c d e f g h i j k l m n o p q r s t u v w x y z	■		
<i>mantissa</i> ::= <i>exact-numeric-literal</i>		■	
<i>minutes-value</i> ::= <i>digit digit</i>			■
<i>months-value</i> ::= <i>digit digit</i>			■
<i>null-predicate</i> ::= <i>column-name</i> IS [NOT] NULL	■		
<i>numeric-literal</i> ::= <i>exact-numeric-literal</i> <i>approximate-numeric-literal</i>	■		
<i>ODBC-date-literal</i> ::= <i>ODBC-std-esc-initiator</i> d 'date-value' <i>ODBC-std-esc-terminator</i> <i>ODBC-ext-esc-initiator</i> d 'date-value' <i>ODBC-ext-esc-terminator</i>			■
<i>ODBC-date-time-extension</i> ::= <i>ODBC-date-literal</i> <i>ODBC-time-literal</i> <i>ODBC-timestamp-literal</i>			■
<i>ODBC-like-escape-clause</i> ::= <i>ODBC-std-esc-initiator</i> escape 'escape-character' <i>ODBC-std-esc-terminator</i> <i>ODBC-ext-esc-initiator</i> escape 'escape-character' <i>ODBC-ext-esc-terminator</i>			■
<i>ODBC-time-literal</i> ::= <i>ODBC-std-esc-initiator</i> t 'time-value' <i>ODBC-std-esc-terminator</i> <i>ODBC-ext-esc-initiator</i> t 'time-value' <i>ODBC-ext-esc-terminator</i>			■
<i>ODBC-timestamp-literal</i> ::= <i>ODBC-std-esc-initiator</i> ts 'timestamp-value' <i>ODBC-std-esc-terminator</i> <i>ODBC-ext-esc-initiator</i> ts 'timestamp-value' <i>ODBC-ext-esc-terminator</i>			■
	Mini-		Ex-

ADABAS D	505
Element	Minimum Core tendend
<i>ODBC-ext-esc-initiator</i> ::= {	■
<i>ODBC-ext-esc-terminator</i> ::= }	■
<i>ODBC-outer-join-extension</i> ::= <i>ODBC-std-esc-initiator</i> oj <i>outer-join</i> <i>ODBC-std-esc-terminator</i> <i>ODBC-ext-esc-initiator</i> oj <i>outer-join</i> <i>ODBC-ext-esc-terminator</i>	■
<i>ODBC-scalar-function-extension</i> ::= <i>ODBC-std-esc-initiator</i> fn <i>scalar-function</i> <i>ODBC-std-esc-terminator</i> <i>ODBC-ext-esc-initiator</i> fn <i>scalar-function</i> <i>ODBC-ext-esc-terminator</i>	■
<i>ODBC-std-esc-initiator</i> ::= <i>ODBC-std-esc-prefix</i> <i>SQL-esc-vendor-clause</i>	■
<i>ODBC-std-esc-prefix</i> ::= --(*	■
<i>ODBC-std-esc-terminator</i> ::= *)--	■
<i>order-by-clause</i> ::= ORDER BY <i>sort-specification</i> [, <i>sort-specification</i>]...	■
<i>outer-join</i> ::= <i>table-name</i> [<i>correlation-name</i>] LEFT OUTER JOIN { <i>table-name</i> [<i>correlation-name</i>] <i>outer-join</i> } ON <i>search-condition</i> (For outer joins, <i>search-condition</i> must contain only the join condition between the specified <i>table-names</i> .)	■
<i>owner-name</i> ::= <i>user-defined-name</i>	■
<i>pattern-value</i> ::= <i>character-string-literal</i> <i>dynamic-parameter</i> (In a <i>character-string-literal</i> , the percent character (%) matches 0 or more of any character; the underscore character (_) matches 1 character.)	■
<i>pattern-value</i> ::= <i>character-string-literal</i> <i>dynamic-parameter</i> USER (In a <i>character-string-literal</i> , the percent character (%) matches 0 or more of any character; the underscore character (_) matches 1 character.)	■
<i>precision</i> ::= <i>unsigned-integer</i>	■
<i>predicate</i> ::= <i>comparison-predicate</i> <i>like-predicate</i> <i>null-predicate</i>	■
<i>predicate</i> ::= <i>between-predicate</i> <i>comparison-predicate</i> <i>exists-predicate</i> <i>in-predicate</i> <i>like-predicate</i> <i>null-predicate</i> <i>quantified-predicate</i>	■
<i>primary</i> ::= <i>column-name</i> <i>dynamic-parameter</i> <i>literal</i> (<i>expression</i>)	■
Element	Minimum Core tendend

<i>primary</i> ::= <i>column-name</i> <i>dynamic-parameter</i> <i>literal</i> <i>set-function-reference</i> USER (<i>expression</i>)	■		
<i>primary</i> ::= <i>column-name</i> <i>dynamic-parameter</i> <i>literal</i> <i>ODBC-scalar-function-extension</i> <i>set-function-reference</i> USER (<i>expression</i>)			■
<i>procedure</i> ::= <i>procedure-name</i> <i>procedure-name</i> (<i>procedure-parameter-list</i>)			■
<i>procedure-identifier</i> ::= <i>user-defined-name</i>			■
<i>procedure-name</i> ::= <i>procedure-identifier</i> <i>owner-name.procedure-identifier</i> <i>qualifier-name qualifier-separator procedure-identifier</i> <i>qualifier-name qualifier-separator [owner-name].procedure-identifier</i> (The third syntax is valid only if the data source does not support owners.)			■
<i>procedure-parameter-list</i> ::= <i>procedure-parameter</i> <i>procedure-parameter</i> , <i>procedure-parameter-list</i>			■
<i>procedure-parameter</i> ::= <i>dynamic-parameter</i> <i>literal</i> <i>empty-string</i> (If a procedure parameter is an empty string, the procedure uses the default value for that parameter.)			■
<i>qualifier-name</i> ::= <i>user-defined-name</i>			■
<i>qualifier-separator</i> ::= {implementation-defined} (The qualifier separator is returned through SQLGetInfo with the SQL_QUALIFIER_NAME_SEPARATOR option.)			■
<i>quantified-predicate</i> ::= <i>expression comparison-operator</i> {ALL ANY} (<i>sub-query</i>)			■
<i>query-specification</i> ::= SELECT [ALL DISTINCT] <i>select-list</i> FROM <i>table-reference-list</i> [WHERE <i>search-condition</i>] [GROUP BY <i>column-name</i> , [<i>column-name</i>]...] [HAVING <i>search-condition</i>]			■
<i>ref-table-name</i> ::= <i>base-table-identifier</i>		■	
Element	Mini- mum	Core	Ex- tende d

ADABAS D	507
<i>ref-table-name</i> ::= <i>base-table-identifier</i> <i>owner-name.base-table-identifier</i> <i>qualifier-name qualifier-separator base-table-identifier</i> <i>qualifier-name qualifier-separator [owner-name].base-table-identifier</i> (The third syntax is valid only if the data source does not support owners.)	■
<i>referenced-columns</i> ::= (<i>column-identifier</i> [, <i>column-identifier</i>]...)	■
<i>referencing-columns</i> ::= (<i>column-identifier</i> [, <i>column-identifier</i>]...)	■
<i>scalar-function</i> ::= <i>function-name</i> (<i>argument-list</i>) (The definitions for the non-terminals <i>function-name</i> and <i>function-name</i> (<i>argument-list</i>) are derived from the list of scalar functions in Appendix F, "Scalar Functions.")	■
<i>scale</i> ::= <i>unsigned-integer</i>	■
<i>search-condition</i> ::= <i>boolean-term</i> [OR <i>search-condition</i>]	■
<i>seconds-fraction</i> ::= <i>unsigned-integer</i>	■
<i>seconds-value</i> ::= <i>digit digit</i>	■
<i>select-list</i> ::= * <i>select-sublist</i> [, <i>select-sublist</i>]...	■
<i>select-sublist</i> ::= <i>expression</i>	■
<i>select-sublist</i> ::= <i>expression</i> [[AS] <i>column-alias</i>] { <i>table-name</i> <i>correlation-name</i> }.*	■
<i>set-function-reference</i> ::= COUNT(*) <i>distinct-function</i> <i>all-function</i>	■
<i>sort-specification</i> ::= { <i>unsigned-integer</i> <i>column-name</i> } [ASC DESC]	■
<i>SQL-esc-vendor-clause</i> ::= VENDOR(Microsoft), PRODUCT(ODBC)	■
<i>sub-query</i> ::= SELECT [ALL DISTINCT] <i>select-list</i> FROM <i>table-reference-list</i> [WHERE <i>search-condition</i>] [GROUP BY <i>column-name</i> [, <i>column-name</i>]...] [HAVING <i>search-condition</i>]	■
<i>table-identifier</i> ::= <i>user-defined-name</i>	■
<i>table-name</i> ::= <i>table-identifier</i>	■
<i>table-name</i> ::= <i>table-identifier</i> <i>owner-name.table-identifier</i> <i>qualifier-name qualifier-separator table-identifier</i> <i>qualifier-name qualifier-separator [owner-name].table-identifier</i> (The third syntax is valid only if the data source does not support owners.)	■
<i>table-reference</i> ::= <i>table-name</i>	■
Element	Mini- mum
	Ex- Core
	tende

		d
<i>table-reference</i> ::= <i>table-name</i> [<i>correlation-name</i>]	■	
<i>table-reference</i> ::= <i>table-name</i> [<i>correlation-name</i>] <i>ODBC-outer-join-extension</i> (A SELECT statement can contain only one <i>table-reference</i> that is an <i>ODBC-outer-join-extension</i> .)		■
<i>table-reference-list</i> ::= <i>table-reference</i> [, <i>table-reference</i>]...	■	
<i>term</i> ::= <i>factor</i> <i>term</i> { <i>*</i> <i>/</i> } <i>factor</i>	■	
<i>time-separator</i> ::= :		■
<i>time-type</i> ::= {time types} (For example, TIME. To determine the type name used by a data source, an application calls SQLGetTypeInfo .)		■
<i>time-value</i> ::= <i>hours-value time-separator minutes-value time-separator seconds-value</i>		■
<i>timestamp-separator</i> ::= (The blank character.)		■
<i>timestamp-type</i> ::= {timestamp types} (For example, TIMESTAMP. To determine the type name used by a data source, an application calls SQLGetTypeInfo .)		■
<i>timestamp-value</i> ::= <i>date-value timestamp-separator time-value</i> [<i>seconds-fraction</i>]		■
<i>unsigned-integer</i> ::= { <i>digit</i> }...	■	
<i>upper-case-letter</i> ::= A B C D E F G H I J K L M N O P Q R S T U V W X Y Z	■	
<i>user-defined-name</i> ::= <i>letter</i> [<i>digit</i> <i>letter</i> <i>_</i>]...	■	
<i>user-name</i> ::= <i>user-defined-name</i>		■
<i>value</i> ::= <i>literal</i> USER <i>dynamic-parameter</i>		■
<i>viewed-table-identifier</i> ::= <i>user-defined-name</i>		■
<i>viewed-table-name</i> ::= <i>viewed-table-identifier</i> <i>owner-name.viewed-table-identifier</i> <i>qualifier-name qualifier-separator viewed-table-identifier</i> <i>qualifier-name qualifier-separator</i> [<i>owner-name</i>]. <i>viewed-table-identifier</i> (The third syntax is valid only if the data source does not support owners.)		■
Element	Mini- mum	Ex- Core tend

ADABAS D	509
	d
<i>years-value ::= digit digit digit digit</i>	■

List of Reserved Keywords

The following words are reserved for use in ODBC function calls. These words do not constrain the minimum SQL grammar; however, to ensure compatibility with drivers that support the core SQL grammar, applications should avoid using any of these keywords. The **#define** value `SQL_ODBC_KEYWORDS` contains a comma-separated list of these keywords.

ABSOLUTE	COLLATION
ADA	COLUMN
ADD	COMMIT
ALL	CONNECT
ALLOCATE	CONNECTION
ALTER	CONSTRAINT
AND	CONSTRAINTS
ANY	CONTINUE
ARE	CONVERT
AS	CORRESPONDING
ASC	COUNT
ASSERTION	CREATE
AT	CURRENT
AUTHORIZATION	CURRENT_DATE
AVG	CURRENT_TIME
BEGIN	CURRENT_TIMESTAMP
BETWEEN	CURSOR
BIT	DATE
BIT_LENGTH	DAY
BY	DEALLOCATE
CASCADE	DEC
CASCADDED	DECIMAL
CASE	DECLARE
CAST	DEFERRABLE
CATALOG	DEFERRED
CHAR	DELETE
CHAR_LENGTH	DESC
CHARACTER	DESCRIBE
CHARACTER_LENGTH	DESCRIPTOR
CHECK	DIAGNOSTICS
CLOSE	DICTIONARY
COALESCE	DISCONNECT
COBOL	DISPLACEMENT
COLLATE	DISTINCT

DOMAIN	INPUT
DOUBLE	INSENSITIVE
DROP	INSERT
ELSE	INTEGER
END	INTERSECT
END-EXEC	INTERVAL
ESCAPE	INTO
EXCEPT	IS
EXCEPTION	ISOLATION
EXEC	JOIN
EXECUTE	KEY
EXISTS	LANGUAGE
EXTERNAL	LAST
EXTRACT	LEFT
FALSE	LEVEL
FETCH	LIKE
FIRST	LOCAL
FLOAT	LOWER
FOR	MATCH
FOREIGN	MAX
FORTTRAN	MIN
FOUND	MINUTE
FROM	MODULE
FULL	MONTH
GET	MUMPS
GLOBAL	NAMES
GO	NATIONAL
GOTO	NCHAR
GRANT	NEXT
GROUP	NONE
HAVING	NOT
HOURL	NULL
IDENTITY	NULLIF
IGNORE	NUMERIC
IMMEDIATE	OCTET_LENGTH
IN	OF
INCLUDE	OFF
INDEX	ON
INDICATOR	ONLY
INITIALLY	OPEN
INNER	OPTION
OR	SQLSTATE
ORDER	SQLWARNING

OUTER	SUBSTRING
OUTPUT	SUM
OVERLAPS	SYSTEM
PARTIAL	TABLE
PASCAL	TEMPORARY
PLI	THEN
POSITION	TIME
PRECISION	TIMESTAMP
PREPARE	TIMEZONE_HOUR
PRESERVE	TIMEZONE_MINUTE
PRIMARY	TO
PRIOR	TRANSACTION
PRIVILEGES	TRANSLATE
PROCEDURE	TRANSLATION
PUBLIC	TRUE
RESTRICT	UNION
REVOKE	UNIQUE
RIGHT	UNKNOWN
ROLLBACK	UPDATE
ROWS	UPPER
SCHEMA	USAGE
SCROLL	USER
SECOND	USING
SECTION	VALUE
SELECT	VALUES
SEQUENCE	VARCHAR
SET	VARYING
SIZE	VIEW
SMALLINT	WHEN
SOME	WHENEVER
SQL	WHERE
SQLCA	WITH
SQLCODE	WORK
SQLERROR	YEAR

Note: For other reserved keywords see the ADABAS D Reference Manual.

A P P E N D I X D

Data Types

Data stored on a data source has an SQL data type, which may be specific to that data source. A driver maps data source–specific SQL data types to ODBC SQL data types and driver-specific SQL data types. (A driver returns these mappings through **SQLGetTypeInfo**. It also returns the SQL data types when describing the data types of columns and parameters in **SQLColAttributes**, **SQLColumns**, **SQLDescribeCol**, **SQLDescribeParam**, **SQLProcedureColumns**, and **SQLSpecialColumns**.)

Each SQL data type corresponds to an ODBC C data type. By default, the driver assumes that the C data type of a storage location corresponds to the SQL data type of the column or parameter to which the location is bound. If the C data type of a storage location is not the *default* C data type, the application can specify the correct C data type with the *fCType* argument in **SQLBindCol**, **SQLGetData**, or **SQLBindParameter**. Before returning data from the data source, the driver converts it to the specified C data type. Before sending data to the data source, the driver converts it from the specified C data type.

This appendix discusses the following:

- ODBC SQL data types
- ODBC C data types
- Default ODBC C data types
- Transferring data in its binary form
- Precision, scale, length, and display size of SQL data types
- Converting data from SQL to C data types
- Converting data from C to SQL data types

For information about driver-specific SQL data types, see the driver's documentation.

SQL Data Types

The ODBC SQL grammar defines three sets of SQL data types, each of which is a superset of the previous set.

- **Minimum** SQL data types provide a basic level of ODBC conformance.
- **Core** SQL data types are the data types in the X/Open and SQL Access Group SQL CAE specification (1992) and are supported by most SQL data sources.
- **Extended** SQL data types are additional data types supported by some SQL data sources.

A given driver and data source do not necessarily support all of the SQL data types defined in the ODBC grammar. Furthermore, they may support additional, driver-specific SQL data types. To determine which data types a driver supports, an application calls **SQLGetTypeInfo**. For information about driver-specific SQL data types, see the driver's documentation.

Minimum SQL Data Types

The following table lists valid values of *fSqlType* for the minimum SQL data types. These values are defined in SQL.H. The table also lists the name and description of the corresponding data type from the X/Open and SQL Access Group SQL CAE specification (1992).

Note: The minimum SQL grammar requires that a data source support at least one character SQL data type. This table is only a guideline and shows commonly used names and limits of these data types. For a given data source, the characteristics of these data types may differ from those listed below. For information about the data types in a specific data source, see the documentation for that data source.

To determine which data types are supported by a data source and the characteristics of those data types, an application calls **SQLGetTypeInfo**.

fSqlType	SQL Data Type	Description
SQL_CHAR	CHAR(<i>n</i>)	Character string of fixed string length <i>n</i> ($1 \leq n \leq 254$).
SQL_VARCHAR	VARCHAR(<i>n</i>)	Variable-length character string with a maximum string length <i>n</i> ($1 \leq n \leq 254$).
SQL_LONGVARCHAR	LONG VARCHAR	Variable length character data. Maximum length is data source–dependent.

Core SQL Data Types

The following table lists valid values of *fSqlType* for the core SQL data types. These values are defined in SQL.H. The table also lists the name and description of the corresponding data type from the X/Open and SQL Access Group SQL CAE specification (1992). In the table, precision refers to the total number of digits and scale refers to the number of digits to the right of the decimal point.

Note: This table is only a guideline and shows commonly used names, ranges, and limits of core SQL data types. A given data source may support only some of the listed data types and the characteristics of the supported data types may differ from those listed below. For example, some data sources support unsigned numeric data types. For information about the data types in a specific data source, see the documentation for that data source.

To determine which data types are supported by a data source and the characteristics of those data types, an application calls **SQLGetTypeInfo**.

fSqlType	SQL Data Type Description	
SQL_DECIMAL	DECIMAL(<i>p</i> , <i>s</i>)	Signed, exact, numeric value with a precision <i>p</i> and scale <i>s</i> ($1 \leq p \leq 15$; $0 \leq s \leq p$).
SQL_NUMERIC	NUMERIC(<i>p</i> , <i>s</i>)	Signed, exact, numeric value with a precision <i>p</i> and scale <i>s</i> ($1 \leq p \leq 15$; $0 \leq s \leq p$).
SQL_SMALLINT	SMALLINT	Exact numeric value with precision 5 and scale 0 (signed: $-32,768 \leq n \leq 32,767$, unsigned: $0 \leq n \leq 65,535$) ^a .
SQL_INTEGER	INTEGER	Exact numeric value with precision 10 and scale 0 (signed: $-2^{31} \leq n \leq 2^{31} - 1$, unsigned: $0 \leq n \leq 2^{32} - 1$) ^a .
SQL_REAL	REAL	Signed, approximate, numeric value with a mantissa precision 7 (zero or absolute value 10^{-38} to 10^{38}).
SQL_FLOAT	FLOAT	Signed, approximate, numeric value with a mantissa precision 15 (zero or absolute value 10^{-308} to 10^{308}).
SQL_DOUBLE	DOUBLE PRECISION	Signed, approximate, numeric value with a mantissa precision 15 (zero or absolute value 10^{-308} to 10^{308}).

^a An application uses **SQLGetTypeInfo** or **SQLColAttributes** to determine if a particular data type or a particular column in a result set is unsigned.

Extended SQL Data Types

The following table lists valid values of *fSqlType* for the extended SQL data types. These values are defined in SQLEXT.H. The table also lists the name and description of the corresponding data type. In the table, precision refers to the total number of digits and scale refers to the number of digits to the right of the decimal point.

Note: This table is only a guideline and shows commonly used names, ranges, and limits of extended SQL data types. A given data source may support only some of the listed data types and the characteristics of the supported data types may differ from those listed below. For example, some data sources support unsigned numeric data types. For information about the data types in a specific data source, see the documentation for that data source.

To determine which data types are supported by a data source and the characteristics of those data types, an application calls **SQLGetTypeInfo**.

fSqlType	Typical SQL Data Type Description	
SQL_BIT	BIT	Single bit binary data.
SQL_TINYINT	TINYINT	Exact numeric value with precision 3 and scale 0 (signed: $-128 \leq n \leq 127$, unsigned: $0 \leq n \leq 255$) ^a .
SQL_BIGINT	BIGINT	Exact numeric value with precision 19 (if signed) or 20 (if unsigned) and scale 0 (signed: $-2^{63} \leq n \leq 2^{63} - 1$, unsigned: $0 \leq n \leq 2^{64} - 1$) ^a .
SQL_BINARY	BINARY(<i>n</i>)	Binary data of fixed length <i>n</i> ($1 \leq n \leq 255$).
SQL_VARBINARY	VARBINARY(<i>n</i>)	Variable length binary data of maximum length <i>n</i> ($1 \leq n \leq 255$).
SQL_LONGVARBINARY	LONG VARBINARY	Variable length binary data. Maximum length is data source-dependent.
SQL_DATE	DATE	Date data.
SQL_TIME	TIME	Time data.
SQL_TIMESTAMP	TIMESTAMP	Date/time data.

^a An application uses **SQLGetTypeInfo** or **SQLColAttributes** to determine if a particular data type or a particular column in a result set is unsigned.

C Data Types

Data is stored in the application in ODBC C data types. The core C data types are those that support the minimum and core SQL data types. They also support some extended SQL data types. The extended C data types are those that only support extended SQL data types. The bookmark C data type is used only to retrieve bookmark values and should not be converted to other data types.

Note: Unsigned C data types for integers were added to ODBC 2.0. Drivers must support the integer C data types specified in both ODBC 1.0 and ODBC 2.0; ODBC 2.0 or later applications must use the ODBC 1.0 integer C data types with ODBC 1.0 drivers and the ODBC 2.0 integer C data types with ODBC 2.0 drivers..

The C data type is specified in the **SQLBindCol**, **SQLGetData**, and **SQLBindParameter** functions with the *fCType* argument.

Core C Data Types

The following table lists valid values of *fCType* for the core C data types. These values are defined in SQL.H. The table also lists the ODBC C data type that implements each value of *fCType* and the definition of this data type from SQL.H.

fCType	ODBC C Typedef	C Type
SQL_C_CHAR	UCHAR FAR *	unsigned char FAR *
SQL_C_SSHORT	SWORD	short int
SQL_C_USHORT	UWORD	unsigned short int
SQL_C_SLONG	SDWORD	long int
SQL_C_ULONG	UDWORD	unsigned long int
SQL_C_FLOAT	SFLOAT	float
SQL_C_DOUBLE	SDOUBLE	double

Note: Because objects of the CString class in Microsoft C++ are signed and string arguments in ODBC functions are unsigned, applications that pass CString objects to ODBC functions without casting them will receive compiler warnings.

Extended C Data Types

The following table lists valid values of *fCType* for the extended C data types. These values are defined in SQLEXT.H. The table also lists the ODBC C data type that implements each value of *fCType* and the definition of this data type from SQLEXT.H or SQL.H.

fCType	ODBC C Typedef	C Type
SQL_C_BIT	UCHAR	unsigned char
SQL_C_STINYINT	SCHAR	signed char
SQL_C_UTINYINT	UCHAR	unsigned char
SQL_C_BINARY	UCHAR FAR *	unsigned char FAR *
SQL_C_DATE	DATE_STRUCT	struct tagDATE_STRUCT { SWORD year; ^a UWORD month; ^b UWORD day; ^c }
SQL_C_TIME	TIME_STRUCT	struct tagTIME_STRUCT { UWORD hour; ^d UWORD minute; ^e UWORD second; ^f }
SQL_C_TIMESTAMP	TIMESTAMP_STRUCT	struct tagTIMESTAMP_STRUCT { SWORD year; ^a UWORD month; ^b UWORD day; ^c UWORD hour; ^d UWORD minute; ^e UWORD second; ^f UDWORD fraction; ^g }

^a The value of the year field must be in the range from 0 to 9,999. Years are measured from 0 A.D. Some data sources do not support the entire range of years.

^b The value of the month field must be in the range from 1 to 12.

^c The value of day field must be in the range from 1 to the number of days in the month. The number of days in the month is determined from the values of the year and month fields and is 28, 29, 30, or 31.

^d The value of the hour field must be in the range from 0 to 23.

^e The value of the minute field must be in the range from 0 to 59.

^f The value of the second field must be in the range from 0 to 59.

^g The value of the fraction field is the number of billionths of a second and ranges from 0 to 999,999,999 (1 less than 1 billion). For example, the value of the fraction field for a half-second is 500,000,000, for a thousandth of a second (one millisecond) is 1,000,000, for a millionth of a second (one microsecond) is 1,000, and for a billionth of a second (one nanosecond) is 1.

Bookmark C Data Type

Bookmarks are 32-bit values used by an application to return to a specific row; an application retrieves a bookmark either from column 0 of the result set with **SQLExtendedFetch** or **SQLGetData** or by calling **SQLGetStmtOption**. For more information, see "Using Bookmarks" in Chapter 7, "Retrieving Results."

The following table lists the value of *fCType* for the bookmark C data type, the ODBC C data type that implements the bookmark C data type, and the definition of this data type from SQL.H.

fCType	ODBC C Typedef	C Type
SQL_C_BOOKMARK	BOOKMARK	unsigned long int

ODBC 1.0 C Data Types

In ODBC 1.0, all integer C data types were signed. The following table lists values of *fCType* for the integer C data types that were valid in ODBC 1.0. To remain compatible with applications that use ODBC 1.0, all drivers must support these values of *fCType*. To remain compatible with drivers that use ODBC 1.0, ODBC 2.0 or later applications must pass these values of *fCType* to ODBC 1.0 drivers. However, ODBC 2.0 or later applications must not pass these values to ODBC 2.0 or later drivers.

fCType	ODBC C Typedef	C Type
SQL_C_TINYINT	SCHAR	signed char
SQL_C_SHORT	WORD	short int
SQL_C_LONG	DWORD	long int

Because the ODBC 1.0 integer C data types (SQL_C_TINYINT, SQL_C_SHORT, and SQL_C_LONG) are signed, and because the ODBC integer SQL data types can be signed or unsigned, ODBC 1.0 applications and drivers had to interpret signed integer C data as signed or unsigned.

ODBC 2.0 applications and drivers treat the ODBC 1.0 integer C data types as unsigned only when:

- The column from which data will be retrieved is unsigned, and
- The C data type of the storage location in which the data will be placed is the default C data type for that column. (For a list of default C data types, see "Default C Data Types" later in this chapter.)

In all other cases, these applications and drivers treat the ODBC 1.0 integer C data types as signed.

In other words, for any conversion except the default conversion, ODBC 2.0 drivers check the validity of the conversion based on the numeric data value. For the default conversion, the drivers simply pass the data value without attempting to validate it numerically and applications interpret the data value according to whether the column is signed. (Applications call **SQLGetTypeInfo** to determine whether a column is signed or unsigned.)

For example, the following table shows how an ODBC 2.0 driver interprets ODBC 1.0 integer C data sent to both signed and unsigned SQL_SMALLINT columns.

From C Data Type	To SQL Data Type	C Data Values	SQL Data Values
SQL_C_TINYINT	SQL_SMALLINT (signed)	–128 to 127	–128 to 127
	SQL_SMALLINT (unsigned)	< 0 0 to 127	--- ^a 0 to 127
SQL_C_SHORT (default conversion)	SQL_SMALLINT (signed)	–32,768 to 32,767	–32,768 to 32,767
	SQL_SMALLINT (unsigned)	–32,768 to –1 0 to 32,767	32,768 to 65,535 0 to 32,767
SQL_C_LONG	SQL_SMALLINT (signed)	< –32,768	--- ^a
		–32,768 to 32,767	–32,768 to 32,767
		> 32,767	--- ^a
	SQL_SMALLINT (unsigned)	< 0	--- ^a
		0 to 32,767	0 to 32,767
		> 32,767	--- ^a

^a The driver returns SQLSTATE 22003 (Numeric value out of range).

Default C Data Types

If an application specifies SQL_C_DEFAULT for the *fCType* argument in **SQLBindCol**, **SQLGetData**, or **SQLBindParameter**, the driver assumes that the C data type of the output or input buffer corresponds to the SQL data type of the column or parameter to which the buffer is bound. For each ODBC SQL data type, the following table shows the corresponding, or *default*, C data type. For information about driver-specific SQL data types, see the driver's documentation.

Note: For maximum interoperability, applications should specify a C data type other than SQL_C_DEFAULT. This allows drivers that promote SQL data types (and therefore cannot always determine default C data types) to return data. It also allows drivers that cannot determine whether an integer column is signed or unsigned to correctly return data.

Note: ODBC 2.0 drivers use the ODBC 2.0 default C data types for both ODBC 1.0 and ODBC 2.0 integer C data.

SQL Data Type	Default C Data Type
SQL_CHAR	SQL_C_CHAR
SQL_VARCHAR	SQL_C_CHAR
SQL_LONGVARCHAR	SQL_C_CHAR
SQL_DECIMAL	SQL_C_CHAR
SQL_NUMERIC	SQL_C_CHAR
SQL_BIT	SQL_C_BIT
SQL_TINYINT	SQL_C_STINYINT or SQL_C_UTINYINT ^a
SQL_SMALLINT	SQL_C_SSHORT or SQL_C_USHORT ^a
SQL_INTEGER	SQL_C_SLONG or SQL_C_ULONG ^a
SQL_BIGINT	SQL_C_CHAR
SQL_REAL	SQL_C_FLOAT
SQL_FLOAT	SQL_C_DOUBLE
SQL_DOUBLE	SQL_C_DOUBLE
SQL_BINARY	SQL_C_BINARY
SQL_VARBINARY	SQL_C_BINARY
SQL_LONGVARBINARY	SQL_C_BINARY
SQL_DATE	SQL_C_DATE
SQL_TIME	SQL_C_TIME
SQL_TIMESTAMP	SQL_C_TIMESTAMP

^a If the driver can determine whether the column is signed or unsigned, such as when the driver is fetching data from the data source or when the data source supports only a signed type or only an unsigned type, but not both, the driver uses the corresponding signed or unsigned C data type. If the driver cannot determine whether the column is signed or unsigned, it passes the data value without attempting to validate it numerically.

Transferring Data in its Binary Form

Among data sources that use the same DBMS, an application can safely transfer data in the internal form used by that DBMS. For a given piece of data, the SQL data types must be the same in the source and target data sources. The C data type is SQL_C_BINARY.

When the application calls **SQLFetch**, **SQLExtendedFetch**, or **SQLGetData** to retrieve the data from the source data source, the driver retrieves the data from

the data source and transfers it, without conversion, to a storage location of type `SQL_C_BINARY`. When the application calls **SQLExecute**, **SQLExecDirect**, or **SQLPutData** to send the data to the target data source, the driver retrieves the data from the storage location and transfers it, without conversion, to the target data source.

Note: Applications that transfer any data (except binary data) in this manner are not interoperable among DBMS's.

Precision, Scale, Length, and Display Size

SQLColAttributes, **SQLColumns**, and **SQLDescribeCol** return the precision, scale, length, and display size of a column in a table. **SQLProcedureColumns** returns the precision, scale, and length of a column in a procedure.

SQLDescribeParam returns the precision or scale of a parameter in an SQL statement; **SQLBindParameter** sets the precision or scale of a parameter in an SQL statement. **SQLGetTypeInfo** returns the maximum precision and the minimum and maximum scales of an SQL data type on a data source.

Due to limitations in the size of the arguments these functions use, precision, length, and display size are limited to the size of an `SDWORD`, or 2,147,483,647.

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 nonnumeric column or parameter generally refers to either the maximum length or the defined length of the column or parameter. To determine the maximum precision allowed for a data type, an application calls **SQLGetTypeInfo**. The following table defines the precision for each ODBC SQL data type.

fSqlType	Precision
<code>SQL_CHAR</code> <code>SQL_VARCHAR</code>	The defined length of the column or parameter. For example, the precision of a column defined as <code>CHAR(10)</code> is 10.
<code>SQL_LONGVARCHAR</code> ^{a, b}	The maximum length of the column or parameter.
<code>SQL_DECIMAL</code> <code>SQL_NUMERIC</code>	The defined number of digits. For example, the precision of a column defined as <code>NUMERIC(10,3)</code> is 10.
<code>SQL_BIT</code> ^c	1
<code>SQL_TINYINT</code> ^c	3
<code>SQL_SMALLINT</code> ^c	5

SQL_INTEGER ^c	10
SQL_BIGINT ^c	19 (if signed) or 20 (if unsigned)
fSqlType	Precision
SQL_REAL ^c	7
SQL_FLOAT ^c	15
SQL_DOUBLE ^c	15
SQL_BINARY SQL_VARBINARY	The defined length of the column or parameter. For example, the precision of a column defined as <code>BINARY(10)</code> is 10.
SQL_LONGVARBINARY ^{a,b}	The maximum length of the column or parameter.
SQL_DATE ^c	10 (the number of characters in the yyyy-mm-dd format).
SQL_TIME ^c	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[.f...]" format used by the <code>TIMESTAMP</code> 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 23 (the number of characters in the "yyyy-mm-dd hh:mm:ss.fff" format).

^a For an ODBC 1.0 application calling **SQLSetParam** in an ODBC 2.0 driver, and for an ODBC 2.0 application calling **SQLBindParameter** in an ODBC 1.0 driver, when *pcbValue* is `SQL_DATA_AT_EXEC`, *cbColDef* must be set to the total length of the data to be sent, not the precision as defined in this table.

^b If the driver cannot determine the column or parameter length, it returns `SQL_NO_TOTAL`.

^c The *cbColDef* argument of **SQLBindParameter** 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. For approximate floating point number columns or parameters, the scale is undefined, since the number of digits to the right of the decimal point is not fixed. (For the SQL_DECIMAL and SQL_NUMERIC data types, the maximum scale is generally the same as the maximum precision. However, some data sources impose a separate limit on the maximum scale. To determine the minimum and maximum scales allowed for a data type, an application calls **SQLGetTypeInfo**.) The following table defines the scale for each ODBC SQL data type.

fSqlType	Scale
SQL_CHAR ^a	Not applicable.
SQL_VARCHAR ^a	
SQL_LONGVARCHAR ^a	
SQL_DECIMAL	The defined number of digits to the right of the decimal point. For example, the scale of a column defined as NUMERIC(10,3) is 3.
SQL_NUMERIC	
SQL_BIT ^a	0
SQL_TINYINT ^a	
SQL_SMALLINT ^a	
SQL_INTEGER ^a	
SQL_BIGINT ^a	
SQL_REAL ^a	Not applicable.
SQL_FLOAT ^a	
SQL_DOUBLE ^a	
SQL_BINARY ^a	Not applicable.
SQL_VARBINARY ^a	
SQL_LONGVARBINARY ^a	
SQL_DATE ^a	Not applicable.
SQL_TIME ^a	
SQL_TIMESTAMP	The number of digits to the right of the decimal point in the "yyyy-mm-dd hh:mm:ss[.f...]" format. For example, if the TIMESTAMP data type uses the "yyyy-mm-dd hh:mm:ss.fff" format, the scale is 3.

^a The *ibScale* argument of **SQLBindParameter** is ignored for this data type.

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 may 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" earlier in this appendix.

The following table defines the length for each ODBC SQL data type.

fSqlType	Length
SQL_CHAR SQL_VARCHAR	The defined length of the column. For example, the length of a column defined as CHAR(10) is 10.
SQL_LONGVARCHAR ^a	The maximum length of the column.
SQL_DECIMAL SQL_NUMERIC	The maximum number of digits plus 2. 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_BIT SQL_TINYINT	1 (one byte).
SQL_SMALLINT	2 (two bytes).
SQL_INTEGER	4 (four bytes).
SQL_BIGINT	20 (since this data type is returned as a character string, characters are needed for 19 digits and a sign, if signed, or 20 digits, if unsigned).
SQL_REAL	4 (four bytes).
SQL_FLOAT	8 (eight bytes).
SQL_DOUBLE	8 (eight bytes).
SQL_BINARY SQL_VARBINARY	The defined length of the column. For example, the length of a column defined as BINARY(10) is 10.
SQL_LONGVARBINARY ^a	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).

^a If the driver cannot determine the column or parameter length, it returns SQL_NO_TOTAL.

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 ODBC SQL data type.

fSqlType	Display Size
SQL_CHAR	The defined length of the column. For example, the display size of a column defined as CHAR(10) is 10.
SQL_VARCHAR	
SQL_LONGVARCHAR ^a	The maximum length of the column.
SQL_DECIMAL	The precision of the column plus 2 (a sign, <i>precision</i> digits, and a decimal point). For example, the display size of a column defined as NUMERIC(10,3) is 12.
SQL_NUMERIC	
SQL_BIT	1 (1 digit).
SQL_TINYINT	4 if signed (a sign and 3 digits) or 3 if unsigned (3 digits).
SQL_SMALLINT	6 if signed (a sign and 5 digits) or 5 if unsigned (5 digits).
SQL_INTEGER	11 if signed (a sign and 10 digits) or 10 if unsigned (10 digits).
SQL_BIGINT	20 (a sign and 19 digits if signed or 20 digits if unsigned).
SQL_REAL	13 (a sign, 7 digits, a decimal point, the letter E, a sign, and 2 digits).
SQL_FLOAT	22 (a sign, 15 digits, a decimal point, the letter E, a sign, and 3 digits).
SQL_DOUBLE	
SQL_BINARY	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 BINARY(10) is 20.
SQL_VARBINARY	
SQL_LONGVARBINARY ^a	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[.f...]" 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.fff").

^a If the driver cannot determine the column or parameter length, it returns

SQL_NO_TOTAL.

Converting Data from SQL to C Data Types

When an application calls **SQLExtendedFetch**, **SQLFetch**, or **SQLGetData**, the driver retrieves the data from the data source. If necessary, it converts the data from the data type in which the driver retrieved it to the data type specified by the *fCType* argument in **SQLBindCol** or **SQLGetData**. Finally, it stores the data in the location pointed to by the *rgbValue* argument in **SQLBindCol** or **SQLGetData**.

Note: The word *convert* is used in this section in a broad sense, and includes the transfer of data, without a conversion in data type, from one storage location to another.

The following table shows the supported conversions from ODBC SQL data types to ODBC C data types. A solid circle indicates the default conversion for an SQL data type (the C data type to which the data will be converted when the value of *fCType* is **SQL_C_DEFAULT**). A hollow circle indicates a supported conversion.

SQL Data Type	C Data Type	SQL_C_CHAR	SQL_C_BIT	SQL_C_SIGNED	SQL_C_UNSIGNED	SQL_C_TINY	SQL_C_SHORT	SQL_C_USHORT	SQL_C_LONG	SQL_C_ULONG	SQL_C_FLOAT	SQL_C_DOUBLE	SQL_C_BINARY	SQL_C_DATE	SQL_C_TIME	SQL_C_TIMESTAMP	P
SQL_CHAR		●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_VARCHAR		●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_LONGVARCHAR		●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_DECIMAL		●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_NUMERIC		●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_BIT		○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_TINYINT (signed)		○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_TINYINT (unsigned)		○	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○
SQL_SMALLINT (signed)		○	○	○	○	○	●	○	○	○	○	○	○	○	○	○	○
SQL_SMALLINT (unsigned)		○	○	○	○	○	○	●	○	○	○	○	○	○	○	○	○
SQL_INTEGER (signed)		○	○	○	○	○	○	○	●	○	○	○	○	○	○	○	○
SQL_INTEGER (unsigned)		○	○	○	○	○	○	○	○	●	○	○	○	○	○	○	○
SQL_BIGINT (signed and unsigned)		●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_REAL		○	○	○	○	○	○	○	○	○	●	○	○	○	○	○	○
SQL_FLOAT		○	○	○	○	○	○	○	○	○	○	●	○	○	○	○	○
SQL_DOUBLE		○	○	○	○	○	○	○	○	○	○	○	●	○	○	○	○
SQL_BINARY		○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_VARBINARY		○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_LONGVARBINARY		○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_DATE		○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_TIME		○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_TIMESTAMP		○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○

● Default conversion
 ○ Supported conversion

The tables in the following sections describe how the driver or data source converts data retrieved from the data source; drivers are required to support conversions to all ODBC C data types from the ODBC SQL data types that they support. For a given ODBC 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. For each outcome, the third and fourth columns list the values of the *rgbValue* and *pcbValue* arguments specified in **SQLBindCol** or **SQLGetData** after the driver has attempted to convert the data. The last column lists the SQLSTATE returned for each outcome by **SQLExtendedFetch**, **SQLFetch**, or **SQLGetData**.

If the *fCType* argument in **SQLBindCol** or **SQLGetData** contains a value for an ODBC C data type not shown in the table for a given ODBC SQL data type, **SQLExtendedFetch**, **SQLFetch**, or **SQLGetData** returns SQLSTATE 07006 (Restricted data type attribute violation). If the *fCType* argument contains a value that specifies a conversion from a driver-specific SQL data type to an ODBC C data type and this conversion is not supported by the driver, **SQLExtendedFetch**, **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).

The following terms and conventions are used in the tables:

- **Length of data** is the number of bytes of C data available to return in *rgbValue*, regardless of whether the data will be truncated before it is returned to the application. For string data, this does not include the null termination byte.
- **Display size** is the total number of bytes needed to display the data in character format.
- Words in *italics* represent function arguments or elements of the ODBC SQL grammar. For the syntax of grammar elements, see Appendix C, "SQL Grammar."

SQL to C: Character

The character ODBC SQL data types are:

```
SQL_CHAR
SQL_VARCHAR
SQL_LONGVARCHAR
```

The following table shows the ODBC C data types to which character SQL data may be converted. For an explanation of the columns and terms in the table, see the list above.

fCType	Test	rgbValue	pcbValue	SQL-STATE
SQL_C_CHAR	Length of data < <i>cbValueMax</i>	Data	Length of data	N/A
	Length of data ≥ <i>cbValueMax</i>	Truncated data	Length of data	01004
SQL_C_STINYINT SQL_C_UTINYINT SQL_C_TINYINT ^a	Data converted without truncation ^b	Data	Size of the C data type	N/A
SQL_C_SSHORT SQL_C_USHORT SQL_C_SHORT ^a	Data converted with truncation of fractional digits ^b	Truncated data	Size of the C data type	01004
SQL_C_SLONG SQL_C_ULONG SQL_C_LONG ^a	Conversion of data would result in loss of whole (as opposed to fractional) digits ^b	Untouched	Untouched	22003
	Data is not a <i>numeric-literal</i> ^b	Untouched	Untouched	22005
SQL_C_FLOAT SQL_C_DOUBLE	Data is within the range of the data type to which the number is being converted ^b	Data	Size of the C data type	N/A
		Untouched	Untouched	22003
	Data is outside the range of the data type to which the number is being converted ^b	Untouched	Untouched	22005
	Data is not a <i>numeric-literal</i> ^b			
SQL_C_BIT	Data is 0 or 1 ^a	Data	1 ^c	N/A
	Data is greater than 0, less than 2, and not equal to 1 ^a	Truncated data	1 ^c	01004
	Data is less than 0 or greater than or equal to 2 ^a	Untouched	Untouched	22003
	Data is not a <i>numeric-literal</i> ^a	Untouched	Untouched	22005
SQL_C_BINARY	Length of data ≤ <i>cbValueMax</i>	Data	Length of data	N/A
	Length of data > <i>cbValueMax</i>	Truncated data	Length of data	01004
SQL_C_DATE	Data value is a valid <i>date-value</i> ^b	Data	6 ^c	N/A
	Data value is a valid <i>timestamp-value</i> ; time portion is zero ^b	Data	6 ^c	N/A
	Data value is a valid <i>timestamp-value</i> ; time portion is non-zero ^{b, d}	Truncated data	6 ^c	01004
	Data value is not a valid <i>date-value</i>	Untouched	Untouched	22008

or *timestamp-value*^b

fCType	Test	rgbValue	pcbValue	SQL-STATE
SQL_C_TIME	Data value is a valid <i>time-value</i> ^b	Data	6 ^c	N/A
	Data value is a valid <i>timestamp-value</i> ; fractional seconds portion is zero ^{b, e}	Data	6 ^c	N/A
	Data value is a valid <i>timestamp-value</i> ; fractional seconds portion is non-zero ^{b, e, f}	Truncated data	6 ^c	01004
	Data value is not a valid <i>time-value</i> or <i>timestamp-value</i> ^b	Untouched	Untouched	22008
SQL_C_TIMESTAMP	Data value is a valid <i>timestamp-value</i> ; fractional seconds portion not truncated ^b	Data	16 ^c	N/A
	Data value is a valid <i>timestamp-value</i> ; fractional seconds portion truncated ^b	Truncated data	16 ^c	N/A
	Data value is a valid <i>date-value</i> ^b	Data ^g	16 ^c	N/A
	Data value is a valid <i>time-value</i> ^b	Data ^h	16 ^c	N/A
	Data value is not a valid <i>date-value</i> , <i>time-value</i> , or <i>timestamp-value</i> ^b	Untouched	Untouched	22008

^a For more information, see "ODBC 1.0 C Data Types," earlier in this appendix.

^b 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.

^c This is the size of the corresponding C data type.

^d The time portion of the *timestamp-value* is truncated.

^e The date portion of the *timestamp-value* is ignored.

^f The fractional seconds portion of the timestamp is truncated.

^g The time fields of the timestamp structure are set to zero.

^h The date fields of the timestamp structure are set to the current date.

When character SQL data is converted to numeric, date, time, or timestamp C data, leading and trailing spaces are ignored.

All drivers that support date, time, and timestamp data can convert character SQL data to date, time, or timestamp C data as specified in the previous table. Drivers may be able to convert character SQL data from other, driver-specific formats to date, time, or timestamp C data. Such conversions are not interoperable among data sources.

SQL to C: Numeric

The numeric ODBC SQL data types are:

SQL_DECIMAL	SQL_BIGINT
SQL_NUMERIC	SQL_REAL
SQL_TINYINT	SQL_FLOAT
SQL_SMALLINT	SQL_DOUBLE
SQL_INTEGER	

The following table shows the ODBC C data types to which numeric SQL data may be converted.

fCType	Test	rgbValue	pcbValue	SQL-STATE
SQL_C_CHAR	Display size < <i>cbValueMax</i>	Data	Length of data	N/A
	Number of whole (as opposed to fractional) digits < <i>cbValueMax</i>	Truncated data	Length of data	01004
	Number of whole (as opposed to fractional) digits ≥ <i>cbValueMax</i>	Untouched	Untouched	22003
SQL_C_TINYINT	Data converted without truncation ^b	Data	Size of the C data type	N/A
SQL_C_UTINYINT				
SQL_C_TINYINT ^a	Data converted with truncation of fractional digits ^b	Truncated data	Size of the C data type	01004
SQL_C_SSHORT				
SQL_C_USHORT				
SQL_C_SHORT ^a	Conversion of data would result in loss of whole (as opposed to fractional) digits ^b	Untouched	Untouched	22003
SQL_C_SLONG				
SQL_C_ULONG				
SQL_C_LONG ^a				
SQL_C_FLOAT SQL_C_DOUBLE	Data is within the range of the data type to which the number is being converted ^b	Data	Size of the C data type	N/A
		Untouched	Untouched	22003
	Data is outside the range of the data type to which the number is being converted ^b			
SQL_C_BIT	Data is 0 or 1 ^b	Data	1 ^c	N/A
	Data is greater than 0, less than 2, and not equal to 1 ^b	Truncated data	1 ^c	01004
	Data is less than 0 or greater than or equal to 2 ^b	Untouched	Untouched	22003

SQL-

fCType	Test	rgbValue	pcbValue	STATE
SQL_C_BINARY	Length of data \leq <i>cbValueMax</i>	Data	Length of data	N/A
	Length of data $>$ <i>cbValueMax</i>	Untouched ^a	Untouched	22003

^b 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.

^c This is the size of the corresponding C data type.

SQL to C: Bit

The bit ODBC SQL data type is:

SQL_BIT

The following table shows the ODBC C data types to which bit SQL data may be converted.

fCType	Test	rgbValue	pcbValue	SQL-STATE
SQL_C_CHAR	$cbValueMax > 1$	Data	1	N/A
	$cbValueMax \leq 1$	Untouched	Untouched	22003
SQL_C_STINYINT SQL_C_UTINYINT SQL_C_TINYINT ^a SQL_C_SSHORT SQL_C_USHORT SQL_C_SHORT ^a SQL_C_SLONG SQL_C_ULONG SQL_C_LONG ^a SQL_C_FLOAT SQL_C_DOUBLE	None ^b	Data	Size of the C data type	N/A
SQL_C_BIT	None ^b	Data	1 ^c	N/A
SQL_C_BINARY	$cbValueMax \geq 1$	Data	1	N/A
	$cbValueMax < 1$	Untouched	Untouched	22003

^a For more information, see "ODBC 1.0 C Data Types," earlier in this appendix.

^b 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.

^c This is the size of the corresponding C data type.

When bit SQL data is converted to character C data, the possible values are "0" and "1".

SQL to C: Binary

The binary ODBC SQL data types are:

SQL_BINARY
SQL_VARBINARY
SQL_LONGVARBINARY

The following table shows the ODBC C data types to which binary SQL data may be converted.

fCType	Test	rgbValue	pcbValue	SQL-STATE
SQL_C_CHAR	$(\text{Length of data}) * 2 < cbValueMax$	Data	Length of data	N/A
	$(\text{Length of data}) * 2 \geq cbValueMax$	Truncated data	Length of data	01004
SQL_C_BINARY	$\text{Length of data} \leq cbValueMax$	Data	Length of data	N/A
	$\text{Length of data} > cbValueMax$	Truncated data	Length of data	01004

When binary SQL data is converted to character C data, each byte (8 bits) of source data is represented as two ASCII characters. These characters are the ASCII character representation of the number in its hexadecimal form. For example, a binary 00000001 is converted to "01" and a binary 11111111 is converted to "FF".

The driver always converts individual bytes to pairs of hexadecimal digits and terminates the character string with a null byte. Because of this, if *cbValueMax* is even and is less than the length of the converted data, the last byte of the *rgbValue* buffer is not used. (The converted data requires an even number of bytes, the next-to-last byte is a null byte, and the last byte cannot be used.)

SQL to C: Date

The date ODBC SQL data type is:

SQL_DATE

The following table shows the ODBC C data types to which date SQL data may be converted.

fCType	Test	rgbValue	pcbValue	SQL-STATE
SQL_C_CHAR	$cbValueMax \geq 11$	Data	10	N/A
	$cbValueMax < 11$	Untouched	Untouched	22003
SQL_C_BINARY	Length of data $\leq cbValueMax$	Data	Length of data	N/A
	Length of data $> cbValueMax$	Untouched	Untouched	22003
SQL_C_DATE	None ^a	Data	6 ^c	N/A
SQL_C_TIMESTAMP	None ^a	Data ^b	16 ^c	N/A

^b The time fields of the timestamp structure are set to zero.

^c This is the size of the corresponding C data type.

When date SQL data is converted to character C data, the resulting string is in the "yyyy-mm-dd" format.

SQL to C: Time

The time ODBC SQL data type is:

SQL_TIME

The following table shows the ODBC C data types to which time SQL data may be converted.

fCType	Test	rgbValue	pcbValue	SQL-STATE
SQL_C_CHAR	$cbValueMax \geq 9$	Data	8	N/A
	$cbValueMax < 9$	Untouched	Untouched	22003
SQL_C_BINARY	Length of data $\leq cbValueMax$	Data	Length of data	N/A
	Length of data $> cbValueMax$	Untouched	Untouched	22003
SQL_C_TIME	None ^a	Data	6 ^c	N/A
SQL_C_TIMESTAMP	None ^a	Data ^b	16 ^c	N/A

^b The date fields of the timestamp structure are set to the current date and the fractional seconds field of the timestamp structure is set to zero.

^c This is the size of the corresponding C data type.

When time SQL data is converted to character C data, the resulting string is in the "hh:mm:ss" format.

SQL to C: Timestamp

The timestamp ODBC SQL data type is:

SQL_TIMESTAMP

The following table shows the ODBC C data types to which timestamp SQL data may be converted.

fCType	Test	rgbValue	pcbValue	SQL-STATE
SQL_C_CHAR	$cbValueMax > \text{Display size}$	Data	Length of data	N/A
	$20 \leq cbValueMax \leq \text{Display size}$	Truncated data ^b	Length of data	01004
	$cbValueMax < 20$	Untouched	Untouched	22003
SQL_C_BINARY	Length of data $\leq cbValueMax$	Data	Length of data	N/A
	Length of data $> cbValueMax$	Untouched	Untouched	22003
SQL_C_DATE	Time portion of timestamp is zero ^a	Data	6 ^f	N/A
	Time portion of timestamp is non-zero ^a	Truncated data ^c	6 ^f	01004
SQL_C_TIME	Fractional seconds portion of timestamp is zero ^a	Data ^d	6 ^f	N/A
	Fractional seconds portion of timestamp is non-zero ^a	Truncated data ^{d,e}	6 ^f	01004
SQL_C_TIMESTAMP	Fractional seconds portion of timestamp is not truncated ^a	Data ^e	16 ^f	N/A
	Fractional seconds portion of timestamp is truncated ^a	Truncated data ^e	16 ^f	01004

^b The fractional seconds of the timestamp are truncated.

^c The time portion of the timestamp is truncated.

^d The date portion of the timestamp is ignored.

^e The fractional seconds portion of the timestamp is truncated.

^f This is the size of the corresponding C data type.

When timestamp SQL data is converted to character C data, the resulting string is in the "yyyy-mm-dd hh:mm:ss[.f...]" format, where up to nine digits may be used for fractional seconds. (Except for the decimal point and fractional seconds, the entire format must be used, regardless of the precision of the timestamp SQL data type.)

SQL to C Data Conversion Examples

The following examples illustrate how the driver converts SQL data to C data:

SQL Data Type	SQL Data Value	C Data Type	cbValueMax	rgbValue	SQL-STATE
SQL_CHAR	abcdef	SQL_C_CHAR	7	abcdef\0 ^a	N/A
SQL_CHAR	abcdef	SQL_C_CHAR	6	abcde\0 ^a	01004
SQL_DECIMAL	1234.56	SQL_C_CHAR	8	1234.56\0 ^a	N/A
SQL_DECIMAL	1234.56	SQL_C_CHAR	5	1234\0 ^a	01004
SQL_DECIMAL	1234.56	SQL_C_CHAR	4	----	22003
SQL_DECIMAL	1234.56	SQL_C_FLOAT	ignored	1234.56	N/A
SQL_DECIMAL	1234.56	SQL_C_SSHORT	ignored	1234	01004
SQL_DECIMAL	1234.56	SQL_C_STINYINT	ignored	----	22003
SQL_DOUBLE	1.2345678	SQL_C_DOUBLE	ignored	1.2345678	N/A
SQL_DOUBLE	1.2345678	SQL_C_FLOAT	ignored	1.234567	N/A
SQL_DOUBLE	1.2345678	SQL_C_STINYINT	ignored	1	N/A
SQL_DATE	1992-12-31	SQL_C_CHAR	11	1992-12-31\0 ^a	N/A
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 ^b	N/A
SQL_TIMESTAMP	1992-12-31 23:45:55.12	SQL_C_CHAR	23	1992-12-31 23:45:55.12\0 ^a	N/A
SQL_TIMESTAMP	1992-12-31 23:45:55.12	SQL_C_CHAR	22	1992-12-31 23:45:55.1\0 ^a	01004
SQL_TIMESTAMP	1992-12-31 23:45:55.12	SQL_C_CHAR	18	----	22003

^a "\0" represents a null-termination byte. The driver always null-terminates SQL_C_CHAR data.

^b The numbers in this list are the numbers stored in the fields of the TIMESTAMP_STRUCT structure.

Converting Data from C to SQL Data Types

When an application calls **SQLExecute** or **SQLExecDirect**, the driver retrieves the data for any parameters bound with **SQLBindParameter** from storage locations in the application. For data-at-execution parameters, the application sends the parameter data with **SQLPutData**. If necessary, the driver converts the data from the data type specified by the *fCType* argument in **SQLBindParameter** to the data type specified by the *fSqlType* argument in **SQLBindParameter**. Finally, the driver sends the data to the data source.

Note: The word *convert* is used in this section in a broad sense, and includes the transfer of data, without a conversion in data type, from one storage location to another.

The following table shows the supported conversions from ODBC C data types to ODBC SQL data types. A solid circle indicates the default conversion for an SQL data type (the C data type from which the data will be converted when the value of *fCType* is `SQL_C_DEFAULT`). A hollow circle indicates a supported conversion.

C Data Type	SQL Data Type	SQL_CHAR	SQL_VARCHAR	SQL_LONGVARCHAR	SQL_DECIMAL	SQL_NUMERIC	SQL_BIT	SQL_TINYINT (signed)	SQL_TINYINT (unsigned)	SQL_SMALLINT (signed)	SQL_SMALLINT (unsigned)	SQL_INTEGER (signed)	SQL_INTEGER (unsigned)	SQL_BIGINT (signed and unsigned)	SQL_REAL	SQL_FLOAT	SQL_DOUBLE	SQL_BINARY	SQL_VARBINARY	SQL_LONGVARBINARY	SQL_DATE	SQL_TIME	SQL_TIMESTAMP
SQL_C_CHAR		●	●	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_C_BIT		○	○	○	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_C_TINYINT		○	○	○	○	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_C_UTINYINT		○	○	○	○	○	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_C_TINYINT		○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_C_SSHORT		○	○	○	○	○	○	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_C_USHORT		○	○	○	○	○	○	○	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○
SQL_C_SHORT		○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_C_SLONG		○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_C_ULONG		○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_C_LONG		○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
SQL_C_FLOAT		○	○	○	○	○	○	○	○	○	○	○	○	○	○	●	○	○	○	○	○	○	○
SQL_C_DOUBLE		○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	●	○	○	○	○	○	○
SQL_C_BINARY		○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	●	●	○	○	○
SQL_C_DATE		○	○	○																	●	○	○
SQL_C_TIME		○	○	○																		●	○
SQL_C_TIMESTAMP		○	○	○																	○	○	●

● Default conversion
 ○ Supported conversion

The tables in the following sections describe how the driver or data source converts data sent to the data source; drivers are required to support conversions from all ODBC C data types to the ODBC SQL data types that they support. For a given ODBC C data type, the first column of the table lists the legal input values of the *fSqlType* argument in **SQLBindParameter**. The second column lists the outcomes of a test that the driver performs to

determine if it can convert the data. The third column lists the SQLSTATE returned for each outcome by **SQLExecDirect**, **SQLExecute**, or **SQLPutData**. Data is sent to the data source only if SQL_SUCCESS is returned.

If the *fSqlType* argument in **SQLBindParameter** contains a value for an ODBC SQL data type that is not shown in the table for a given C data type, **SQLBindParameter** returns SQLSTATE 07006 (Restricted data type attribute violation). If the *fSqlType* argument contains a driver-specific value and the driver does not support the conversion from the specific ODBC C data type to that driver-specific SQL data type, **SQLBindParameter** returns SQLSTATE S1C00 (Driver not capable).

If the *rgbValue* and *pcbValue* arguments specified in **SQLBindParameter** are both null pointers, that function returns SQLSTATE S1009 (Invalid argument value). Though it is not shown in the tables, an application sets the value pointed to by the *pcbValue* argument of **SQLBindParameter** or the value of the *cbValue* argument to SQL_NULL_DATA to specify a NULL SQL data value. The application sets these values to SQL_NTS to specify that the value in *rgbValue* is a null-terminated string.

The following terms are used in the tables:

- **Length of data** is the number of bytes of SQL data available to send to the data source, regardless of whether the data will be truncated before it is sent to the data source. For string data, this does not include the null termination byte.
- **Column length** and **display size** are defined for each SQL data type in the section "Precision, Scale, Length, and Display Size" earlier in this chapter.
- **Number of digits** is the number of characters used to represent a number, including the minus sign, decimal point, and exponent (if needed).
- Words in *italics* represent elements of the ODBC SQL grammar. For the syntax of grammar elements, see Appendix C, "SQL Grammar."

C to SQL: Character

The character ODBC C data type is:

SQL_C_CHAR

The following table shows the ODBC SQL data types to which C character data may be converted.

fSqlType	Test	SQL-STATE
SQL_CHAR	Length of data \leq Column length	N/A
SQL_VARCHAR	Length of data $>$ Column length	01004
SQL_LONGVARCHAR		
SQL_DECIMAL	Data converted without truncation	N/A
SQL_NUMERIC	Data converted with truncation of fractional digits	01004
SQL_TINYINT	Conversion of data would result in loss of whole (as opposed to fractional) digits	22003
SQL_SMALLINT		
SQL_INTEGER	Data value is not a <i>numeric-literal</i>	22005
SQL_BIGINT		
SQL_REAL	Data is within the range of the data type to which the number is being converted	N/A
SQL_FLOAT		
SQL_DOUBLE	Data is outside the range of the data type to which the number is being converted	22003
	Data value is not a <i>numeric-literal</i>	22005
SQL_BIT	Data is 0 or 1	N/A
	Data is greater than 0, less than 2, and not equal to 1	01004
	Data is less than 0 or greater than or equal to 2	22003
	Data is not a <i>numeric-literal</i>	22005
SQL_BINARY	(Length of data) / 2 \leq Column length	N/A
SQL_VARBINARY	(Length of data) / 2 $>$ Column length	01004
SQL_LONGVARBINARY	Data value is not a hexadecimal value	22005
SQL_DATE	Data value is a valid <i>ODBC-date-literal</i>	N/A
	Data value is a valid <i>ODBC-timestamp-literal</i> ; time portion is zero	N/A
	Data value is a valid <i>ODBC-timestamp-literal</i> ; time portion is non-zero ^a	01004
	Data value is not a valid <i>ODBC-date-literal</i> or <i>ODBC-timestamp-literal</i>	22008

fSqlType	Test	SQL-STATE
SQL_TIME	Data value is a valid <i>ODBC-time-literal</i>	N/A
	Data value is a valid <i>ODBC-timestamp-literal</i> ; fractional seconds portion is zero ^b	N/A
	Data value is a valid <i>ODBC-timestamp-literal</i> ; fractional seconds portion is non-zero ^{b, c}	01004
	Data value is not a valid <i>ODBC-time-literal</i> or <i>ODBC-timestamp-literal</i>	22008
SQL_TIMESTAMP	Data value is a valid <i>ODBC-timestamp-literal</i> ; fractional seconds portion not truncated	N/A
	Data value is a valid <i>ODBC-timestamp-literal</i> ; fractional seconds portion truncated	01004
	Data value is a valid <i>ODBC-date-literal</i> ^d	N/A
	Data value is a valid <i>ODBC-time-literal</i> ^e	N/A
	Data value is not a valid <i>ODBC-date-literal</i> , <i>ODBC-time-literal</i> , or <i>ODBC-timestamp-literal</i>	22008

^a The time portion of the timestamp is truncated.

^b The date portion of the timestamp is ignored.

^c The fractional seconds portion of the timestamp is truncated.

^d The time portion of the timestamp is set to zero.

^e The date portion of the timestamp is set to the current date.

When character C data is converted to numeric, date, time, or timestamp SQL data, leading and trailing blanks are ignored.

When character C data is converted to binary SQL data, each two bytes of character data are converted to a single byte (8 bits) of binary data. Each two bytes of character data represent a number in hexadecimal form. For example, "01" is converted to a binary 00000001 and "FF" is converted to a binary 11111111.

The driver always converts pairs of hexadecimal digits to individual bytes and ignores the null termination byte. Because of this, if the length of the character string is odd, the last byte of the string (excluding the null termination byte, if any) is not converted.

All drivers that support date, time, and timestamp data can convert character C data to date, time, or timestamp SQL data as specified in the previous table. Drivers may be able to convert character C data from other, driver-specific formats to date, time, or timestamp SQL data. Such conversions are not interoperable among data sources.

C to SQL: Numeric

The numeric ODBC C data types are:

SQL_C_STINYINT	SQL_C_SLONG
SQL_C_UTINYINT	SQL_C_ULONG
SQL_C_TINYINT	SQL_C_LONG
SQL_C_SSHORT	SQL_C_FLOAT
SQL_C_USHORT	SQL_C_DOUBLE
SQL_C_SHORT	

For more information about the SQL_C_TINYINT, SQL_C_SHORT, and SQL_C_LONG data types, see "ODBC 1.0 C Data Types," earlier in this appendix. The following table shows the ODBC SQL data types to which numeric C data may be converted.

fSqlType	Test	SQL-STATE
SQL_CHAR	Number of digits \leq Column length	N/A
SQL_VARCHAR	Number of whole (as opposed to fractional) digits \leq Column length	01004
SQL_LONGVARCHAR	Number of whole (as opposed to fractional) digits $>$ Column length	22003
SQL_DECIMAL	Data converted without truncation	N/A
SQL_NUMERIC	Data converted with truncation of fractional digits	01004
SQL_TINYINT	Conversion of data would result in loss of whole (as opposed to fractional) digits	22003
SQL_SMALLINT		
SQL_INTEGER		
SQL_BIGINT		
SQL_REAL	Data is within the range of the data type to which the number is being converted	N/A
SQL_FLOAT		
SQL_DOUBLE	Data is outside the range of the data type to which the number is being converted	22003
SQL_BIT	Data is 0 or 1	N/A
	Data is greater than 0, less than 2, and not equal to 1	01004
	Data is less than 0 or greater than or equal to 2	22003

The value pointed to by the *pcbValue* argument of **SQLBindParameter** and the value of the *cbValue* argument of **SQLPutData** are ignored when data is converted from the numeric C data types. The driver assumes that the size of *rgbValue* is the size of the numeric C data type.

C to SQL: Bit

The bit ODBC C data type is:

SQL_C_BIT

The following table shows the ODBC SQL data types to which bit C data may be converted.

fSqlType	Test	SQL-STATE
SQL_CHAR SQL_VARCHAR SQL_LONGVARCHAR	None	N/A
SQL_DECIMAL SQL_NUMERIC SQL_TINYINT SQL_SMALLINT SQL_INTEGER SQL_BIGINT SQL_REAL SQL_FLOAT SQL_DOUBLE	None	N/A
SQL_BIT	None	N/A

The value pointed to by the *pcbValue* argument of **SQLBindParameter** and the value of the *cbValue* argument of **SQLPutData** are ignored when data is converted from the bit C data type. The driver assumes that the size of *rgbValue* is the size of the bit C data type.

C to SQL: Binary

The binary ODBC C data type is:

SQL_C_BINARY

The following table shows the ODBC SQL data types to which binary C data may be converted.

fSqlType	Test	SQL-STATE
SQL_CHAR	Length of data \leq Column length	N/A
SQL_VARCHAR	Length of data $>$ Column length	01004
SQL_LONGVARCHAR		
SQL_DECIMAL	Length of data = SQL data length ^a	N/A
SQL_NUMERIC	Length of data \neq SQL data length ^a	22003
SQL_TINYINT		
SQL_SMALLINT		
SQL_INTEGER		
SQL_BIGINT		
SQL_REAL		
SQL_FLOAT		
SQL_DOUBLE		
SQL_BIT	Length of data = SQL data length ^a	N/A
	Length of data \neq SQL data length ^a	22003
SQL_BINARY	Length of data \leq Column length	N/A
SQL_VARBINARY	Length of data $>$ Column length	01004
SQL_LONGVARBINARY		
SQL_DATE	Length of data = SQL data length ^a	N/A
SQL_TIME	Length of data \neq SQL data length ^a	22003
SQL_TIMESTAMP		

^a The SQL data length is the number of bytes needed to store the data on the data source. (This may be different than the column length, as defined earlier in this appendix.)

C to SQL: Date

The date ODBC C data type is:

SQL_C_DATE

The following table shows the ODBC SQL data types to which date C data may be converted.

fSqlType	Test	SQL-STATE
SQL_CHAR	Column length \geq 10	N/A
SQL_VARCHAR	Column length < 10	22003
SQL_LONGVARCHAR	Data value is not a valid date	22008
SQL_DATE	Data value is a valid date	N/A
	Data value is not a valid date	22008
SQL_TIMESTAMP	Data value is a valid date ^a	N/A
	Data value is not a valid date ^a	22008

^a The time portion of the timestamp is set to zero.

For information about what values are valid in a SQL_C_DATE structure, see "Extended C Data Types," earlier in this appendix.

When date C data is converted to character SQL data, the resulting character data is in the "yyyy-mm-dd" format.

The value pointed to by the *pcbValue* argument of **SQLBindParameter** and the value of the *cbValue* argument of **SQLPutData** are ignored when data is converted from the date C data type. The driver assumes that the size of *rgbValue* is the size of the date C data type.

C to SQL: Time

The time ODBC C data type is:

SQL_C_TIME

The following table shows the ODBC SQL data types to which time C data may be converted.

fSqlType	Test	SQL-STATE
SQL_CHAR	Column length \geq 8	N/A
SQL_VARCHAR	Column length $<$ 8	22003
SQL_LONGVARCHAR	Data value is not a valid time	22008
SQL_TIME	Data value is a valid time	N/A
	Data value is not a valid time	22008
SQL_TIMESTAMP	Data value is a valid time ^a	N/A
	Data value is not a valid time	22008

^a The date portion of the timestamp is set to the current date and the fractional seconds portion of the timestamp is set to zero.

For information about what values are valid in a SQL_C_TIME structure, see "Extended C Data Types," earlier in this appendix.

When time C data is converted to character SQL data, the resulting character data is in the "hh:mm:ss" format.

The value pointed to by the *pcbValue* argument of **SQLBindParameter** and the value of the *cbValue* argument of **SQLPutData** are ignored when data is converted from the time C data type. The driver assumes that the size of *rgbValue* is the size of the time C data type.

C to SQL: Timestamp

The timestamp ODBC C data type is:

SQL_C_TIMESTAMP

The following table shows the ODBC SQL data types to which timestamp C data may be converted.

fSqlType	Test	SQL-STATE
SQL_CHAR	Column length \geq Display size	N/A
SQL_VARCHAR	$19 \leq$ Column length $<$ Display size ^a	01004
SQL_LONGVARCHAR	Column length < 19	22003
	Data value is not a valid date	22008
SQL_DATE	Time fields are zero	N/A
	Time fields are non-zero ^b	01004
	Data value does not contain a valid date	22008
SQL_TIME	Fractional seconds fields are zero ^c	N/A
	Fractional seconds fields are non-zero ^{c, d}	01004
	Data value does not contain a valid time	22008
SQL_TIMESTAMP	Fractional seconds fields are not truncated	N/A
	Fractional seconds fields are truncated ^d	01004
	Data value is not a valid timestamp ^a	22008

^b The time fields of the timestamp structure are truncated.

^c The date fields of the timestamp structure are ignored.

^d The fractional seconds fields of the timestamp structure are truncated.

For information about what values are valid in a SQL_C_TIMESTAMP structure, see "Extended C Data Types," earlier in this appendix.

When timestamp C data is converted to character SQL data, the resulting character data is in the "yyyy-mm-dd hh:mm:ss[.f...]" format.

The value pointed to by the *pcbValue* argument of **SQLBindParameter** and the value of the *cbValue* argument of **SQLPutData** are ignored when data is converted from the timestamp C data type. The driver assumes that the size of *rgbValue* is the size of the timestamp C data type.

C to SQL Data Conversion Examples

The following examples illustrate how the driver converts C data to SQL data:

C DataType	C Data Value	SQL Data Type	Column length	SQL Data Value	SQL-STATE
SQL_C_CHAR	abcdef\0 ^a	SQL_CHAR	6	abcdef	N/A
SQL_C_CHAR	abcdef\0 ^a	SQL_CHAR	5	abcde	01004
SQL_C_CHAR	1234.56\0 ^a	SQL_DECIMAL	8 ^b	1234.56	N/A
SQL_C_CHAR	1234.56\0 ^a	SQL_DECIMAL	7 ^b	1234.5	01004
SQL_C_CHAR	1234.56\0 ^a	SQL_DECIMAL	4	----	22003
SQL_C_FLOAT	1234.56	SQL_FLOAT	not applicable	1234.56	N/A
SQL_C_FLOAT	1234.56	SQL_INTEGER	not applicable	1234	01004
SQL_C_FLOAT	1234.56	SQL_TINYINT	not applicable	----	22003
SQL_C_DATE	1992,12,31 ^c	SQL_CHAR	10	1992-12-31	N/A
SQL_C_DATE	1992,12,31 ^c	SQL_CHAR	9	----	22003
SQL_C_DATE	1992,12,31 ^c	SQL_TIMESTAMP	not applicable	1992-12-31 00:00:00.0	N/A
SQL_C_TIMESTAMP	1992,12,31, 23,45,55, 120000000 ^d	SQL_CHAR	22	1992-12-31 23:45:55.12	N/A
SQL_C_TIMESTAMP	1992,12,31, 23,45,55, 120000000 ^d	SQL_CHAR	21	1992-12-31 23:45:55.1	01004
SQL_C_TIMESTAMP	1992,12,31, 23,45,55, 120000000 ^d	SQL_CHAR	18	----	22003

^a "\0" represents a null-termination byte. The null-termination byte is required only if the length of the data is SQL_NTS.

^b In addition to bytes for numbers, one byte is required for a sign and another byte is required for the decimal point.

^c The numbers in this list are the numbers stored in the fields of the DATE_STRUCT structure.

^d The numbers in this list are the numbers stored in the fields of the TIMESTAMP_STRUCT structure.

APPENDIX E

Comparison Between Embedded SQL and ODBC

This appendix compares ODBC and embedded SQL.

ODBC to Embedded SQL

The following table compares core ODBC functions to embedded SQL statements. This comparison is based on the X/Open and SQL Access Group SQL CAE specification (1992).

ODBC uses a parameter marker in place of a host variable, wherever a host variable would occur in embedded SQL.

The SQL language is based on the X/Open and SQL Access Group SQL CAE specification (1992).

ODBC Function	Statement	Comments
SQLAllocEnv	none	Driver Manager and driver memory allocation.
SQLAllocConnect	none	Driver Manager and driver memory allocation.
SQLConnect	CONNECT	Association management.
SQLAllocStmt	none	Driver Manager and driver memory allocation.
SQLPrepare	PREPARE	The prepared SQL string can contain any of the valid preparable functions as defined by the X/Open specification, including ALTER, CREATE, <i>cursor-specification</i> , searched DELETE, dynamic SQL positioned DELETE, DROP, GRANT, INSERT, REVOKE, searched UPDATE, or dynamic SQL positioned UPDATE.

ODBC Function	Statement	Comments
SQLBindParameter	USING DESCRIPTOR	Dynamic SQL INCLUDE SQLDA and dynamic SQL USING DESCRIPTOR. USING DESCRIPTOR would normally be issued on the first call to SQLBindParameter for an <i>hstmt</i> . Alternatively, USING DESCRIPTOR can be called during SQLAllocStmt , although this call would be unneeded by SQL statements containing no embedded parameters. The descriptor name is generated by the driver.
SQLSetCursorName	none	The specified cursor name is used in the DECLARE CURSOR statement generated by SQLExecute or SQLExecDirect .
SQLGetCursorName	none	Driver cursor name management.
SQLExecute	EXECUTE or DECLARE CURSOR and OPEN CURSOR	Dynamic SQL EXECUTE. If the SQL statement requires a cursor, then a dynamic SQL DECLARE CURSOR statement and a dynamic SQL OPEN are issued at this time.
SQLExecDirect	EXECUTE IMMEDIATE or DECLARE CURSOR and OPEN CURSOR	The ODBC function call provides for support for a <i>cursor specification</i> and statements allowed in an EXECUTE IMMEDIATE dynamic SQL statement. In the case of a <i>cursor specification</i> , the call corresponds to static SQL DECLARE CURSOR and OPEN statements.
SQLNumResultCols	USING DESCRIPTOR	COUNT from sqllda.sqln.
SQLColAttributes	USING DESCRIPTOR	COUNT from sqllda.sqln or VALUE from sqlvar[col].field-name with <i>field-name</i> any of {colname, coltype, collength, colfrac, colmode}.
SQLDescribeCol	USING DESCRIPTOR	VALUE from sqlvar[col].field-name with <i>field-name</i> any of {colname, coltype, collength, colfrac, colmode}.
SQLBindCol	none	This function establishes output buffers that correspond in usage to host variables for static SQL FETCH, and to an SQL DESCRIPTOR for dynamic SQL FETCH <i>cursor</i> USING DESCRIPTOR <i>descriptor</i> .
SQLFetch	FETCH	Static or dynamic SQL FETCH. If the call is a dynamic SQL FETCH, then the VALUE from sqllda.sqlvar[col].field-name with <i>field-name</i> any of {hostvaraddr, hostindicator}. hostvaraddr and hostindicator values are placed in output buffers specified in SQLBindCol .

ODBC Function	Statement	Comments
SQLFreeStmt (SQL_CLOSE option)	CLOSE	Dynamic SQL CLOSE.
SQLFreeStmt (SQL_DROP option)	none	Driver Manager and driver memory deallocation.
SQLTransact	COMMIT WORK or COMMIT ROLLBACK	None.
SQLDisconnect	DISCONNECT	Association management.
SQLFreeConnect	none	Driver Manager and driver memory deallocation.
SQLFreeEnv	none	Driver Manager and driver memory deallocation.
SQLCancel	none	None.
SQLError	WHENEVER	WHENEVER retrieves information from the SQL error area that pertains to the most recently executed SQL statement. This information can be retrieved following execution and preceding the deallocation of the statement.

Embedded SQL to ODBC

The following tables list the relationship between the X/Open Embedded SQL language and corresponding ODBC functions. The section number shown in the first column of each table refers to the section of the X/Open and SQL Access Group SQL CAE specification (1992).

Declarative Statements

The following table lists declarative statements.

Section	SQL Statement	ODBC Function	Comments
4.3.1	Static SQL DECLARE CURSOR	none	Issued implicitly by the driver if a <i>cursor specification</i> is passed to SQLExecDirect .
4.3.2	Dynamic SQL DECLARE CURSOR	none	Cursor is generated automatically by the driver. To set a name for the cursor, use SQLSetCursorName . To retrieve a cursor name, use SQLGetCursorName .

Data Definition Statements

The following table lists data definition statements.

Section	SQL Statement	ODBC Function	Comments
5.1.2	ALTER TABLE	SQLPrepare, SQLExecute, or SQLExecDirect	None.
5.1.3	CREATE INDEX		
5.1.4	CREATE TABLE		
5.1.5	CREATE VIEW		
5.1.6	DROP INDEX		
5.1.7	DROP TABLE		
5.1.8	DROP VIEW		
5.1.9	GRANT		
	REVOKE		

Data Manipulation Statements

The following table lists data manipulation statements.

Section	SQL Statement	ODBC Function	Comments
5.2.1	CLOSE	SQLFreeStmt (SQL_CLOSE option)	None.
5.2.2	Positioned DELETE	SQLExecDirect (..., "DELETE FROM <i>table-name</i> WHERE CURRENT OF <i>cursor-name</i> ")	Driver-generated <i>cursor-name</i> can be obtained by calling SQLGetCursorName .
5.2.3	Searched DELETE	SQLExecDirect (..., "DELETE FROM <i>table-name</i> WHERE <i>search-condition</i> ")	None.
5.2.4	FETCH	SQLFetch	None.
5.2.5	INSERT	SQLExecDirect (..., "INSERT INTO <i>table-name</i> ...")	Can also be invoked by SQLPrepare and SQLExecute .
5.2.6	OPEN	none	Cursor is OPENed implicitly by SQLExecute or SQLExecDirect when a SELECT statement is specified.
5.2.7	SELECT ...INTO	SQLExecDirect (..., "SELECT <i>-column-list</i> INTO parameter ...FROM <i>table-name</i> ")	Select statement is executed and the Parameters are moved directly to the paramters bind with SQLSetParam
5.2.8	Positioned UPDATE	SQLExecDirect (..., "UPDATE <i>table-name</i> SET <i>column-identifier</i> = <i>expression</i> ...WHERE CURRENT OF <i>cursor-name</i> ")	Driver-generated <i>cursor-name</i> can be obtained by calling SQLGetCursorName .
5.2.9	Searched UPDATE	SQLExecDirect (..., "UPDATE <i>table-name</i> SET <i>column-identifier</i> = <i>expression</i> ...WHERE <i>search-condition</i> ")	None.

Dynamic SQL Statements

The following table lists dynamic SQL statements.

Section	SQL Statement	ODBC Function	Comments
5.3 (see 5.2.1)	Dynamic SQL CLOSE	SQLFreeStmt (SQL_CLOSE option)	None.
5.3 (see 5.2.2)	Dynamic SQL Positioned DELETE	SQLExecDirect (..., "DELETE FROM <i>table-name</i> WHERE CURRENT OF <i>cursor-name</i> ")	Can also be invoked by SQLPrepare and SQLExecute .
5.3 (see 5.2.8)	Dynamic SQL Positioned UPDATE	SQLExecDirect (..., "UPDATE <i>table-name</i> SET <i>column-identifier</i> = <i>expression</i> ...WHERE CURRENT OF <i>cursor-name</i> ")	Can also be invoked by SQLPrepare and SQLExecute .
5.3.3	USING DESCRIPTOR	None	Descriptor information is implicitly allocated and attached to the <i>hstmt</i> by the driver. Allocation occurs at either the first call to SQLBindParameter or at SQLExecute or SQLExecDirect time.
5.3.4	none	SQLFreeStmt (SQL_DROP option)	None.
5.3.5	DESCRIBE	none	None.
5.3.6	EXECUTE	SQLExecute	None.
5.3.7	EXECUTE IMMEDIATE	SQLExecDirect	None.
5.3.8	Dynamic SQL FETCH	SQLFetch	None.
5.3.9	USING DESCRIPTOR	SQLNumResultCols SQLDescribeCol SQLColAttributes	COUNT from <i>sqlda.sqln</i> . VALUE from <i>sqlda.sqlvar[col].field-name</i> with <i>field-name</i> any of { <i>colname</i> , <i>coltype</i> , <i>collength</i> , <i>colfrac</i> , <i>colmode</i> }.
5.3.10	Dynamic SQL OPEN	SQLExecute	None.
5.3.11	PREPARE	SQLPrepare	None.

Transaction Control Statements

The following table lists transaction control statements.

Section	SQL Statement	ODBC Function	Comments
5.4.1	COMMIT WORK	SQLTransact (SQL_COMMIT option)	None.
5.4.2	ROLLBACK WORK	SQLTransact (SQL_ROLLBACK option)	None.

Association Management Statements

The following table lists association management statements.

Section	SQL Statement	ODBC Function	Comments
5.5.1	CONNECT	SQLConnect	None.
5.5.2	DISCONNECT	SQLDisconnect	ODBC does not support DISCONNECT ALL.

A P P E N D I X F

Scalar Functions

ODBC specifies five types of scalar functions:

- String functions
- Numeric functions
- Time and date functions
- System functions
- Data type conversion functions

The following sections list functions by function type. Descriptions include associated syntax.

String Functions

The following table lists string manipulation functions.

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* or *count* can be a numeric literal or the result of another scalar function, where the underlying data type can be represented as SQL_TINYINT, SQL_SMALLINT, or SQL_INTEGER.

The string functions listed here are 1-based, that is, the first character in the string is character 1.

Function	Description
ASCII (<i>string_exp</i>)	Returns the ASCII code value of the leftmost character of <i>string_exp</i> as an integer.
CHAR (<i>code</i>)	Returns the character that has the ASCII code value specified by <i>code</i> . The value of <i>code</i> should be between 0 and 255; otherwise, the return value is data source–dependent.
CONCAT (<i>string_exp1</i> , <i>string_exp2</i>)	Returns a character string that is the result of concatenating <i>string_exp2</i> to <i>string_exp1</i> . The resulting string is DBMS dependent. For example, if the column represented by <i>string_exp1</i> contained a NULL value, DB2 would return NULL, but SQL Server would return the non-NULL string.
DIFFERENCE (<i>string_exp1</i> , <i>string_exp2</i>)	Returns an integer value that indicates the difference between the values returned by the SOUNDEX function for <i>string_exp1</i> and <i>string_exp2</i> .
INSERT (<i>string_exp1</i> , <i>start</i> , <i>length</i> , <i>string_exp2</i>)	Returns a character string where <i>length</i> characters have been deleted from <i>string_exp1</i> beginning at <i>start</i> and where <i>string_exp2</i> has been inserted into <i>string_exp</i> , beginning at <i>start</i> .
LCASE (<i>string_exp</i>)	Converts all upper case characters in <i>string_exp</i> to lower case.
LEFT (<i>string_exp</i> , <i>count</i>)	Returns the leftmost <i>count</i> of characters of <i>string_exp</i> .
LENGTH (<i>string_exp</i>)	Returns the number of characters in <i>string_exp</i> , excluding trailing blanks and the string termination character.

Function	Description
LOCATE (<i>string_exp1</i> , <i>string_exp2</i> [, <i>start</i>])	Returns the starting position of the first occurrence of <i>string_exp1</i> within <i>string_exp2</i> . The search for the first occurrence of <i>string_exp1</i> begins with the first character position in <i>string_exp2</i> unless the optional argument, <i>start</i> , is specified. If <i>start</i> is specified, the search begins with the character position indicated by the value of <i>start</i> . The first character position in <i>string_exp2</i> is indicated by the value 1. If <i>string_exp1</i> is not found within <i>string_exp2</i> , the value 0 is returned.
LTRIM (<i>string_exp</i>)	Returns the characters of <i>string_exp</i> , with leading blanks removed.
REPEAT (<i>string_exp</i> , <i>count</i>)	Returns a character string composed of <i>string_exp</i> repeated <i>count</i> times.
REPLACE (<i>string_exp1</i> , <i>string_exp2</i> , <i>string_exp3</i>)	Replaces all occurrences of <i>string_exp2</i> in <i>string_exp1</i> with <i>string_exp3</i> .
RIGHT (<i>string_exp</i> , <i>count</i>)	Returns the rightmost <i>count</i> of characters of <i>string_exp</i> .
RTRIM (<i>string_exp</i>)	Returns the characters of <i>string_exp</i> with trailing blanks removed.
SOUNDEX (<i>string_exp</i>)	Returns a data source–dependent character string representing the sound of the words in <i>string_exp</i> . For example, SQL Server returns a four digit SOUNDEX code; Oracle returns a phonetic representation of each word.
SPACE (<i>count</i>)	Returns a character string consisting of <i>count</i> spaces.
SUBSTRING (<i>string_exp</i> , <i>start</i> , <i>length</i>)	Returns a character string that is derived from <i>string_exp</i> beginning at the character position specified by <i>start</i> for <i>length</i> characters.
UCASE (<i>string_exp</i>)	Converts all lower case characters in <i>string_exp</i> to upper case.

Numeric Functions

The following table describes numeric functions that are included in the ODBC scalar function set.

Arguments denoted as *numeric_exp* can be the name of a column, the result of another scalar function, or a numeric literal, where the underlying data type could be represented as SQL_NUMERIC, SQL_DECIMAL, SQL_TINYINT, SQL_SMALLINT, SQL_INTEGER, SQL_BIGINT, SQL_FLOAT, SQL_REAL, or SQL_DOUBLE.

Arguments denoted as *float_exp* can be the name of a column, the result of another scalar function, or a numeric literal, where the underlying data type can be represented as SQL_FLOAT.

Arguments denoted as *integer_exp* can be the name of a column, the result of another scalar function, or a numeric literal, where the underlying data type can be represented as SQL_TINYINT, SQL_SMALLINT, SQL_INTEGER, or SQL_BIGINT.

Function	Description
ABS (<i>numeric_exp</i>)	Returns the absolute value of <i>numeric_exp</i> .
ACOS (<i>float_exp</i>)	Returns the arccosine of <i>float_exp</i> as an angle, expressed in radians.
ASIN (<i>float_exp</i>)	Returns the arcsine of <i>float_exp</i> as an angle, expressed in radians.
ATAN (<i>float_exp</i>)	Returns the arctangent of <i>float_exp</i> as an angle, expressed in radians.
ATAN2 (<i>float_exp1</i> , <i>float_exp2</i>)	Returns the arctangent of the x and y coordinates, specified by <i>float_exp1</i> and <i>float_exp2</i> , respectively, as an angle, expressed in radians.
CEILING (<i>numeric_exp</i>)	Returns the smallest integer greater than or equal to <i>numeric_exp</i> .
COS (<i>float_exp</i>)	Returns the cosine of <i>float_exp</i> , where <i>float_exp</i> is an angle expressed in radians.
COT (<i>float_exp</i>)	Returns the cotangent of <i>float_exp</i> , where <i>float_exp</i> is an angle expressed in radians.
DEGREES (<i>numeric_exp</i>)	Returns the number of degrees converted from <i>numeric_exp</i> radians.
EXP (<i>float_exp</i>)	Returns the exponential value of <i>float_exp</i> .
FLOOR (<i>numeric_exp</i>)	Returns largest integer less than or equal to <i>numeric_exp</i> .

Function	Description
LOG (<i>float_exp</i>)	Returns the natural logarithm of <i>float_exp</i> .
LOG10 (<i>float_exp</i>)	Returns the base 10 logarithm of <i>float_exp</i> .
MOD (<i>integer_exp1</i> , <i>integer_exp2</i>)	Returns the remainder (modulus) of <i>integer_exp1</i> divided by <i>integer_exp2</i> .
PI ()	Returns the constant value of pi as a floating point value.
POWER (<i>numeric_exp</i> , <i>integer_exp</i>)	Returns the value of <i>numeric_exp</i> to the power of <i>integer_exp</i> .
RADIANS (<i>numeric_exp</i>)	Returns the number of radians converted from <i>numeric_exp</i> degrees.
RAND ([<i>integer_exp</i>])	Returns a random floating point value using <i>integer_exp</i> as the optional seed value.
ROUND (<i>numeric_exp</i> , <i>integer_exp</i>)	Returns <i>numeric_exp</i> rounded to <i>integer_exp</i> places right of the decimal point. If <i>integer_exp</i> is negative, <i>numeric_exp</i> is rounded to <i>integer_exp</i> places to the left of the decimal point.
SIGN (<i>numeric_exp</i>)	Returns an indicator or the sign of <i>numeric_exp</i> . If <i>numeric_exp</i> is less than zero, -1 is returned. If <i>numeric_exp</i> equals zero, 0 is returned. If <i>numeric_exp</i> is greater than zero, 1 is returned.
SIN (<i>float_exp</i>)	Returns the sine of <i>float_exp</i> , where <i>float_exp</i> is an angle expressed in radians.
SQRT (<i>float_exp</i>)	Returns the square root of <i>float_exp</i> .
TAN (<i>float_exp</i>)	Returns the tangent of <i>float_exp</i> , where <i>float_exp</i> is an angle expressed in radians.
TRUNCATE (<i>numeric_exp</i> , <i>integer_exp</i>)	Returns <i>numeric_exp</i> truncated to <i>integer_exp</i> places right of the decimal point. If <i>integer_exp</i> is negative, <i>numeric_exp</i> is truncated to <i>integer_exp</i> places to the left of the decimal point.

Time and Date Functions

The following table lists time and date functions that are included in the ODBC scalar function set.

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, where the underlying data type could be represented as SQL_CHAR, SQL_VARCHAR, SQL_TIME, SQL_DATE, or SQL_TIMESTAMP.

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 could be represented as SQL_CHAR, SQL_VARCHAR, SQL_DATE, or SQL_TIMESTAMP.

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 type could be represented as SQL_CHAR, SQL_VARCHAR, SQL_TIME, or SQL_TIMESTAMP.

Values returned are represented as ODBC data types.

Function	Description
CURDATE()	Returns the current date as a date value.
CURTIME()	Returns the current local time as a time value.
DAYNAME(<i>date_exp</i>)	Returns a character string containing the data source-specific name of the day (for example, Sunday, through Saturday or Sun. through Sat. for a data source that uses English, or Sonntag through Samstag for a data source that uses German) for the day portion of <i>date_exp</i> .
DAYOFMONTH(<i>date_exp</i>)	Returns the day of the month in <i>date_exp</i> as an integer value in the range of 1–31.
DAYOFWEEK(<i>date_exp</i>)	Returns the day to the week in <i>date_exp</i> as an integer value in the range of 1–7, where 1 represents Sunday.
DAYOFYEAR(<i>date_exp</i>)	Returns the day of the year in <i>date_exp</i> as an integer value in the range of 1–366.
HOUR(<i>time_exp</i>)	Returns the hour in <i>time_exp</i> as an integer value in the range of 0–23.
MINUTE(<i>time_exp</i>)	Returns the minute in <i>time_exp</i> as an integer value in the range of 0–59.
MONTH(<i>date_exp</i>)	Returns the month in <i>date_exp</i> as an integer value in the range of 1–12.
Function	Description

MONTHNAME (<i>date_exp</i>)	Returns a character string containing the data source–specific name of the month (for example, January through December or Jan. through Dec. for a data source that uses English, or Januar through Dezember for a data source that uses German) for the month portion of <i>date_exp</i> .
NOW ()	Returns current date and time as a timestamp value.
QUARTER (<i>date_exp</i>)	Returns the quarter in <i>date_exp</i> as an integer value in the range of 1–4, where 1 represents January 1 through March 31.
SECOND (<i>time_exp</i>)	Returns the second in <i>time_exp</i> as an integer value in the range of 0–59.
TIMESTAMPADD (<i>interval</i> , <i>integer_exp</i> , <i>timestamp_exp</i>)	<p>Returns the timestamp calculated by adding <i>integer_exp</i> intervals of type <i>interval</i> to <i>timestamp_exp</i>. Valid values of <i>interval</i> are the following keywords:</p> <p>SQL_TSI_FRAC_SECOND SQL_TSI_SECOND SQL_TSI_MINUTE SQL_TSI_HOUR SQL_TSI_DAY SQL_TSI_WEEK SQL_TSI_MONTH SQL_TSI_QUARTER SQL_TSI_YEAR</p> <p>where fractional seconds are expressed in billionths of a second. For example, the following SQL statement returns the name of each employee and their one-year anniversary dates:</p> <pre>SELECT NAME, {fn TIMESTAMPADD(SQL_TSI_YEAR, 1, HIRE_DATE)} FROM EMPLOYEES</pre> <p>7</p> <p>If <i>timestamp_exp</i> is a time value and <i>interval</i> specifies days, weeks, months, quarters, or years, the date portion of <i>timestamp_exp</i> is set to the current date before calculating the resulting timestamp.</p> <p>If <i>timestamp_exp</i> is a date value and <i>interval</i> specifies fractional seconds, seconds, minutes, or hours, the time portion of <i>timestamp_exp</i> is set to 0 before calculating the resulting timestamp.</p> <p>An application determines which intervals a data source supports by calling SQLGetInfo with the SQL_TIMEDATE_ADD_INTERVALS option.</p>

Function**Description**

TIMESTAMPDIFF (<i>interval</i> , <i>timestamp_exp1</i> , <i>timestamp_exp2</i>)	Returns the integer number of intervals of type <i>interval</i> by which <i>timestamp_exp2</i> is greater than <i>timestamp_exp1</i> . Valid values of <i>interval</i> are the following keywords:
--	--

SQL_TSI_FRAC_SECOND
SQL_TSI_SECOND
SQL_TSI_MINUTE
SQL_TSI_HOUR
SQL_TSI_DAY
SQL_TSI_WEEK
SQL_TSI_MONTH
SQL_TSI_QUARTER
SQL_TSI_YEAR

where fractional seconds are expressed in billionths of a second.
For example, the following SQL statement returns the name of each employee and the number of years they have been employed.

```
SELECT NAME,  
       {fn TIMESTAMPDIFF('SQL_TSI_YEAR',  
       {fn CURDATE()}, HIRE_DATE)}  
FROM EMPLOYEES
```

If either timestamp expression is a time value and *interval* specifies days, weeks, months, quarters, or years, the date portion of that timestamp is set to the current date before calculating the difference between the timestamps.

If either timestamp expression is a date value and *interval* specifies fractional seconds, seconds, minutes, or hours, the time portion of that timestamp is set to 0 before calculating the difference between the timestamps.

An application determines which intervals a data source supports by calling **SQLGetInfo** with the SQL_TIMEDATE_DIFF_INTERVALS option.

WEEK(*date_exp*)

Returns the week of the year in *date_exp* as an integer value in the range of 1–53.

YEAR(*date_exp*)

Returns the year in *date_exp* as an integer value. The range is data source–dependent.

System Functions

The following table lists system functions that are included in the ODBC scalar function set.

Arguments denoted as *exp* can be the name of a column, the result of another scalar function, or a literal, where the underlying data type could be represented as SQL_NUMERIC, SQL_DECIMAL, SQL_TINYINT, SQL_SMALLINT, SQL_INTEGER, SQL_BIGINT, SQL_FLOAT, SQL_REAL, SQL_DOUBLE, SQL_DATE, SQL_TIME, or SQL_TIMESTAMP.

Arguments denoted as *value* can be a literal constant, where the underlying data type can be represented as SQL_NUMERIC, SQL_DECIMAL, SQL_TINYINT, SQL_SMALLINT, SQL_INTEGER, SQL_BIGINT, SQL_FLOAT, SQL_REAL, SQL_DOUBLE, SQL_DATE, SQL_TIME, or SQL_TIMESTAMP.

Values returned are represented as ODBC data types.

Function	Description
DATABASE()	Returns the name of the database corresponding to the connection handle (<i>hdbc</i>). (The name of the database is also available by calling SQLGetConnectOption with the SQL_CURRENT_QUALIFIER connection option.)
IFNULL(<i>exp,value</i>)	If <i>exp</i> is null, <i>value</i> is returned. If <i>exp</i> is not null, <i>exp</i> is returned. The possible data type(s) of <i>value</i> must be compatible with the data type of <i>exp</i> .
USER()	Returns the user's authorization name. (The user's authorization name is also available via SQLGetInfo by specifying the information type: SQL_USER_NAME.)

Explicit Data Type Conversion

Explicit data type conversion is specified in terms of ODBC SQL data type definitions.

The ODBC syntax for the explicit data type conversion function does not restrict conversions. The validity of specific conversions of one data type to another data type will be determined by each driver-specific implementation. The driver will, as it translates the ODBC syntax into the native syntax, reject those conversions that, although legal in the ODBC syntax, are not supported by the data source. The ODBC function **SQLGetInfo** provides a way to inquire about conversions supported by the data source.

The format of the **CONVERT** function is:

CONVERT(*value_exp*, *data_type*)

The function returns the value specified by *value_exp* converted to the specified *data_type*, where *data_type* is one of the following keywords:

SQL_BIGINT	SQL_LONGVARBINARY
SQL_BINARY	SQL_LONGVARCHAR
SQL_BIT	SQL_REAL
SQL_CHAR	SQL_SMALLINT
SQL_DATE	SQL_TIME
SQL_DECIMAL	SQL_TIMESTAMP
SQL_DOUBLE	SQL_TINYINT
SQL_FLOAT	SQL_VARBINARY
SQL_INTEGER	SQL_VARCHAR

The ODBC syntax for the explicit data type conversion function does not support specification of conversion format. If specification of explicit formats is supported by the underlying data source, a driver must specify a default value or implement format specification.

The argument *value_exp* can be a column name, the result of another scalar function, or a numeric or string literal. For example:

```
{ fn CONVERT( { fn CURDATE() }, SQL_CHAR) }
```

converts the output of the CURDATE scalar function to a character string..

The following two examples illustrate the use of the **CONVERT** function. These examples assume the existence of a table called EMPLOYEES, with an EMPNO column of type SQL_SMALLINT and an EMPNAME column of type SQL_CHAR.

If an application specifies the following:

```
SELECT EMPNO FROM EMPLOYEES WHERE
--(*vendor(Microsoft),product(ODBC) fn CONVERT(EMPNO,SQL_CHAR)*)--
LIKE '1%'
```

or its equivalent in shorthand form:

```
SELECT EMPNO FROM EMPLOYEES WHERE
{fn CONVERT(EMPNO,SQL_CHAR)} LIKE '1%'
```

A driver that supports an ORACLE DBMS would translate the request to:

```
SELECT EMPNO FROM EMPLOYEES WHERE to_char(EMPNO) LIKE '1%'
```

A driver that supports a SQL Server DBMS would translate the request to:

```
SELECT EMPNO FROM EMPLOYEES WHERE convert(char,EMPNO) LIKE '1%'
```

If an application specifies the following:

```
SELECT
--(*vendor(Microsoft),product(ODBC) fn ABS(EMPNO)*)--,
--(*vendor(Microsoft),product(ODBC) fn CONVERT(EMPNAME,SQL_SMALLINT)*)--
FROM EMPLOYEES WHERE EMPNO <> 0
```

or its equivalent in shorthand form:

```
SELECT {fn ABS(EMPNO)}, {fn CONVERT(EMPNAME,SQL_SMALLINT)}
FROM EMPLOYEES WHERE EMPNO <> 0
```

A driver that supports an Oracle DBMS would translate the request to:

```
SELECT abs(EMPNO), to_number(EMPNAME) FROM EMPLOYEES WHERE EMPNO <> 0
```

A driver that supports a SQL Server DBMS would translate the request to:

```
SELECT abs(EMPNO), convert(smallint, EMPNAME) FROM EMPLOYEES
WHERE EMPNO != 0
```

A driver that supports an Ingres DBMS would translate the request to:

```
SELECT abs(EMPNO), int2(EMPNAME) FROM EMPLOYEES WHERE EMPNO <> 0
```

A P P E N D I X G

- omitted -

A P P E N D I X H

ODBC Library (Unix, OS/2)

To install the ODBC software, use the installation procedures described in the ADABAS D installation manuals of the particular platforms. This chapter only contains a description of the individual ODBC software components.

Installed Files

The following is an overview on the files created during installation and their functionality.

The files are specified with their relative paths, starting from DBROOT.

lib/odbcLib.a	ODBC driver library
lib/sql*.a	Libraries for the ADABAS D runtime environment (see the platform-specific installation manual).
bin/callpc	Shell script (UNIX only) for the translation of an ODBC application.
bin/odbcInk	Shell script (UNIX only) for the linking of an application.
incl/sql.h	Header file for the ODBC driver (Core).
incl/sqlext.h	Header file for the ODBC driver (Level1 and Level2).
incl/WINDOWS.H	Header file for non-MS WINDOWS platforms. On MS WINDOWS 3.1 or WINDOWS NT, this file can be replaced by windows.h if an SDK has been installed.

demo/eng/ODBC/sqlxamp.c C example program for an ODBC application
(see Chapter 10).

demo/eng/ODBC/sqladhoc.c C example program for an ODBC application
(see Chapter 10)

A P P E N D I X I

ODBC Library (Windows)

To install the ODBC software, use the installation procedures described in the ADABAS D installation manuals of the particular platforms. This chapter only contains a description of the individual ODBC software components.

Installed Files

The following is an overview on the files created during installation and their functionality.

The files are specified with their relative paths, starting from the respective directory.

<DBROOT> directory :

bin\sql*.dll	Libraries for the ADABAS D runtime environment (see the platform-specific installation manual)
---------------------	---

<Windows> directory :

odbc.ini	Information about data sources
-----------------	--------------------------------

odbcinst.ini	Information about installed drivers and translators
---------------------	---

<Windows System> directory :

odbc.dll	Driver Manager
-----------------	----------------

ctl3dv2.dll	3D Windows Controls
--------------------	---------------------

odbcinst.dll	Installer library
---------------------	-------------------

odbcinst.hlp	Installer library help
---------------------	------------------------

odbcadm.exe	Administrator program
--------------------	-----------------------

odbccurs.dll	ODBC cursor library
---------------------	---------------------

ODBC Run-Time Options

This chapter describes how information about the data flow and the executed SQL commands can be obtained.

Setting Options

Different options can be set for the ODBC driver at runtime. These options either influence the behavior of the application or provide information about the execution of SQL commands.

These options can be set by using the environment variable SQLOPT.

Connecting to a Database Session

The parameters passed with SQLConnect can be overridden at the runtime of the application. The parameters that can be overridden are SERVERDB, SERVERNODE (szDSN), user name, and password. Such a switch has only an effect on the first database session. The session that is the first to call the function SQLConnect after calling the function SQLAllocConnect always becomes the first database session.

The switch -n overrides the SERVERNODE in the parameter szDSN of the function SQLConnect; -d overrides the SERVERDB name; -u overrides the user name and password.

Example: This example overrides the parameter szDSN passed with SQLConnect.

```
UCHAR FAR szDSN[]="sqlberlin";
....
main ()
{
....
    retcode = SQLConnect(hdbc, szDSN, SQL_NTS, ...);
....
}
```

If the program is started with

 SQLOPT="-nsqlserver -dDATABASE" example

entered in the command line, the program is executed on the SERVERDB "DATABASE" and the SERVERNODE "sqlserver".

Summary of the Connect Options

-n <servername>	overriding the SERVERNODE
-d <serverdb>	overriding the SERVERDB
-u <name>,<passw>	user name Password

Trace of SQL Commands

The options can be used to generate a trace file of the SQL commands sent to the database. If the option -X is enabled, all SQL commands are recorded with time and date. The default filename is 'sqltrace.pct'. The file is stored in the working directory of the application. If this is not desired or not possible (missing write privilege), the trace filename can be rerouted by using the option -F <filename>.

With enabled option, a separate trace file is created for each ODBC Environment generated using the function SQLAllocEnv. The filename is kept and the file extension is increased (the last two characters only).

Example: Two trace files are generated which named trace.pct and trace.p01

```
main()
....
retcode =SQLAllocEnv(&henv);
retcode = SQLAllocEnv(&henv2);
retcode = SQLAllocConnect(henv, &hdbc)
retcode = SQLAllocConnect(henv2, &hdbc2)
retcode = SQLConnect(hdbc, szDSN, SQL_NTS, ...);
retcode = SQLConnect(hdbc2, "szDSN, SQL_NTS, ...);
```

If the program is started with:

 SQLOPT=-X -F trace.pct example

entered in the command line, the program is executed in two separate database sessions for each of which a trace file is written.

Further Trace File Options

As the trace file can become very large, options can be used to control the trace output.

Summary of the Trace Options

-X	enabling the trace file
-N	trace without date/time
-F <tracefn>	trace filename (if necessary.: directory)
-Y <statement count>	alternately overwriting
-L <seconds>	recording statements with a runtime exceeding <seconds> seconds

Index

Symbols

% (Percent sign) 85; 403
' (Single-quote) 108
* (Asterisk) 120
... (Ellipsis) x
? (Question mark)
 browse result connection string 121
 parameter markers 185; 309
[] (Brackets) x
_ (Underscore) 85
{ x; 120
{ } (Braces) 168
| (Vertical bar) x

—A—

Access mode 172
 specifying 344
Access plans 21
 deleting 310
ALIAS table type 405
Aliases, column 258
Allocating memory See Memory, allocating
ALTER TABLE statements
 data type names 285; 469
 interoperability 19
 modifying syntax 19
 owner usage in 270
 qualifier usage in 271
 supported clauses 257
Applications
 examples 57; 61
Architecture, ODBC

 error handling 50
 gateways 53
Arguments
 naming conventions 81
 null pointers 10
 pointers 10
 search patterns 85
 See also Parameters
 SQLTables 403
Arrays See Binding columns, column-wise See
 Parameters, arrays
Association management statements 533
Asterisk (*) 120
Asynchronous execution
 described 26
 specifying 376
 SQLCancel 127
 state transitions 453
 terminating 55
Attributes
 columns 132; 157
 data types 288
Authentication strings See Connection strings
Auto incrementing 132; 287
Auto-commit mode
 described 23
 See also Transactions 23
 specifying 344
 SQLTransact 407

—B—

Bar, vertical (|) x
Batch statements
 processing 290

- syntax 477
- Binary data
 - converting to C 510
 - converting to SQL 522
 - converting with CONVERT 259
 - literal length 265
 - null termination byte 110
 - retrieving in parts 237; 242
 - specifying conversions
 - SQLBindCol 94
 - SQLBindParameter 107
 - SQLGetData 236
 - transferring 497
- Binary large object (BLOB) See SQLGetData
- Bindind columns
 - code examples 204
- Binding columns
 - binding type 376
 - column attributes 36
 - column-wise 39; 196
 - null pointers 100
 - row-wise 39; 197
 - See also Converting data
 - See also Retrieving data
 - See also Rowsets
 - single row 35
 - SQLBindCol 98
 - unbinding 100; 226
- Binding parameters See Parameters, binding
- Bit data
 - converting to C 509
 - converting to SQL 521
 - converting with CONVERT 259
 - specifying conversions
 - SQLBindCol 94
 - SQLBindParameter 107
 - SQLGetData 236
- BLOB, Binary large object See SQLGetData
- Block cursors 40
- Bookmarks
 - described 43
 - enabling 380
 - getting 281
 - persistence 257; 258
 - SQLExtendedFetch 201

- Boundaries, segment 10
- Braces ({}) x; 120; 168
- Brackets ([]) x
- Browse request connection string 119
- Browse result connection string 120
- Buffers
 - allocating 9
 - input 10
 - interoperability 10
 - maintaining pointers 10
 - NULL data 10; 11
 - null pointers 10
 - null termination 11
 - output 10
 - See also NULL data
 - See also Null termination byte
 - segment boundaries 10
 - truncating data 11
- Bulk operations 302

—C—

- C data types
 - conversion examples 514; 526
 - converting from SQL data types 503; 505; 508; 509; 510; 511; 512; 513
 - converting to SQL data types 515; 518; 520; 521; 522; 523; 524; 525
 - defaults 496
 - defined 493
 - in ODBC 1.0 495
 - See also SQL data types 500
 - specifying conversions
 - SQLBindCol 94
 - SQLBindParameter 107
 - SQLGetData 236
 - using 12
- Call Level Interface (CLI)
 - support of 268
- Canceled
 - asynchronous execution 127
 - connection browsing 163
 - data-at-execution 127
 - on another thread 128
- Cardinality 394
- Cascading deletes 218

Cascading updates 218

Case sensitivity

- columns 132

- data types 286

- quoted SQL identifiers 271

- SQL identifiers 265

Catalog functions

- list of 24

- search patterns 85; 273

- summary 79

Character data

- case sensitivity 286

- converting to C 505

- converting to SQL 518

- converting with CONVERT 259

- empty string 10

- literal length 265

- literal prefix string 285

- literal suffix string 285

- maximum length 378

- NULL 10

- null termination byte 10; 110

- retrieving in parts 237; 242

- specifying conversions

 - SQLBindCol 94

 - SQLBindParameter 107

 - SQLGetData 236

- string functions 274

Characters, escape 85; 273

Characters, special See Special characters

CLI (Call Level Interface)

- support of 268

Closing cursors 226; 260

- SQLCancel 127

- SQLFreeStmt 55

Clustered indexes 393

Code examples

- ad hoc query 61

- parameter values 23

- See also specific function description 57

- static SQL 57

Codes, return 49

Collating 267; 394

Columns

- aliases 258

- attributes 132; 157

- binding See Binding columns

- foreign keys 217

- in GROUP BY clauses 266

- in indexes 266

- in ORDER BY clauses 266

- in select list 266

- in tables 266

- indexes 393

- labels 132

- listing 143

- maximum name length 266

- nullability 267

- number of 132; 296

- precision See Precision

- primary keys 217; 314

- privileges 135; 138

- procedure See Procedure columns

- pseudo 387

- scale See Scale

- See also Result sets

- See also Retrieving data

- See also SQL data types 490

- See also Tables

- signed 134

- titles 132

- unbinding 100; 226

- uniquely identifying row 386

COMMIT

- SQLExecDirect 185

COMMIT statements

- cursor behavior 260

- interoperability 24; 185

- SQLExecute 191

- SQLPrepare 309

Committing transactions 407

Compatibility, backwards See Version compatibility

Concurrency

- defined 43

- row locking 361

- serializability 43

- simulating transactions 361

- specifying 369; 377

supported types 272

ConfigDSN

- function description 413
- See also Driver setup DLL
- with SQLRemoveDSNFromIni 415
- with SQLWriteDSNToIni 414
- with SQLWritePrivateProfileString 414

ConfigTranslator

- function description 416
- See also Translator setup DLL

Configuring data sources See Data sources, configuring

Conformance levels

- by function 247
- determining 7; 268; 269
- in ODBC 1.0 46
- SQL 472; 490; 491; 492

Connecting to the ServerDB 17

Connection handles

- active 257
- allocating 16; 88
- defined 11
- driver 262
- freeing 56
- overwriting 88
- See also Handles
- SQLFreeConnect 222
- state transitions 447
- with threads 88

Connection options

- access mode 172; 344
- commit mode 344
- current qualifier 344
- cursor library 345
- dialog boxes 346
- login interval 172; 344
- maximum length 230; 343
- packet size 346
- releasing 343
- reserved 343
- retrieving 229
- See also Statement options 231
- setting 343
- trace file 345; 346
- tracing 345

Connection strings

- browse request 119
- browse result 120
- special characters 119; 120; 168
- SQLDriverConnect 168

Connections

- dialog boxes 169; 171; 172
- disconnecting 56
- Driver Manager 150
- function summary 76; 79
- in embedded SQL 533
- SQLBrowseConnect 121
- SQLConnect 150
- SQLDisconnect 163
- SQLDriverConnect 168
- terminating browsing 122

CONVERT function

- described 31; 544
- information types 256
- supported conversions 259

Converting data

- C to SQL 515; 518; 520; 521; 522; 523; 524; 525
- CONVERT function 31; 544
- converting with CONVERT 259
- default conversions 496
- examples 514; 526
- hexadecimal characters 510; 519; 522
- information types 256
- parameters 22
- See also Translation DLLs
- specifying conversions
 - SQLBindCol 94
 - SQLBindParameter 107
 - SQLGetData 236
- SQL to C 503; 505; 508; 509; 510; 511; 512; 513
- SQLExtendedFetch 196
- SQLFetch 99
- SQLGetData 242

Correlated subqueries 274

Correlation names 259

CREATE TABLE statements

- data type names 285; 469
- interoperability 19

- modifying syntax 19
- NOT NULL clauses 267
- owner usage in 270
- qualifier usage in 271

Currency columns 133

Currency data types 287

Current qualifier 344

Current row 281

Cursor library

- Driver Manager 345
- enabling 345

Cursor position

- bookmarks 43; 202
- current 281
- errors 200
- required 45; 242
- SQLExtendedFetch 43; 198; 201
- SQLFetch 210
- SQLSetPos 359
- supported operations 263

Cursors

- closing 260
 - SQLFreeStmt 55
- closing cursors
 - SQLCancel 127
- deleting 260
- described 37
- dynamic 41
- getting names 233
- holes in 41; 273
- keyset size 370; 372
- keyset-driven 41
- maximum name length 266
- mixed 42
- position See Cursor position
- positioned statements 45
- scrollable 40
- See also Concurrency 43
- sensitivity 273
- setting names 350
- simulated 380; 386
- specifying type 42
- SQLFreeStmt 226
- SQLSetScrollOptions 369; 372

- SQLSetStmtOption 377
- static 41
- supported types 272
- transaction behavior 260; 408

Cursors, block 40

—D—

Data

- converting See Converting data
- long See Long data values
- retrieving See Retrieving data
- transferring in binary form 497
- truncating See Truncating data

Data conversion See Converting data

Data Definition Language (DDL)

- in embedded SQL 530
- owner usage in 270
- qualifier usage in 271

Data Manipulation Language (DML)

- in embedded SQL 531
- owner usage in 270
- qualifier usage in 271

Data sources

- adding 414
- configuring 414
- connecting to 18
- default 18
- deleting 415
- information types 254
- name 260
- read only 260

Data types See SQL data types See C data types

Data-at-execution

- canceling 127
- columns 46; 47
- macro 105; 108; 110; 111; 363
- parameters 25
- SQLBindParameter 105; 108; 110; 111
- SQLSetPos 363

Databases

- current 344
- information types 254
- name 261

Date data

- converting to C 511
- converting to SQL 523
- converting with CONVERT 259
- intervals 275
- literals 29
- scalar functions 276; 540
- specifying conversions
 - SQLBindCol 94
 - SQLBindParameter 107
 - SQLGetData 236
- Date functions 276; 540
- DBMS product information 254; 261
- Declarative statements 530
- DELETE
 - privileges 399
- Delete rules 218
- DELETE statements
 - affected rows 338
 - cascade delete 218
 - owner usage in 270
 - qualifier usage in 271
 - restrictive delete 218
 - See also SQLSetPos
- Deletes, positioned See Positioned delete statements
- Deleting cursors 260
- Delimiter character, SQL identifiers 271
- Deprecated functions 353
- Descriptors
 - columns 132; 157
 - driver-specific 132
 - parameters 161
- Dialog boxes
 - disabling 346
 - SQLDriverConnect 169; 171
- Dirty reads 261; 277; 347
- Display size 132; 502
- Documentation x
- Driver keyword 168
 - browse request connection string 121
 - version compatibility 121
- Driver Manager
 - allocating handles 89
 - cursor library 345
 - described (UNIX) 8
 - described (WINDOWS) 9
 - errors 50; 54
 - functions supported 246
 - listing drivers 175
 - loading drivers 150
 - ODBC version supported 269
 - SQLBrowseConnect 121
 - SQLConnect 150
 - SQLDriverConnect 169
 - SQLError 177
 - SQLSTATE values 84
 - state transitions 444
 - tracing 224; 345; 346
 - transactions 407
 - unloading drivers 151
- Driver setup DLL
 - ConfigDSN 411
- Drivers
 - allocating handles 89
 - errors 49
 - file usage 263
 - functions supported 246
 - information types 253
 - keywords 175
 - listing installed 175
 - loading 150
 - setup DLL See Driver setup DLL
 - SQLError 177
 - SQLSTATE values 84
- DROP INDEX statements 393
- DSN keyword 121; 168; 413
- Dynamic cursors 41
- Dynamic SQL 532

—E—

- Elements, SQL statements 477
- Ellipsis (...) x
- Embedded SQL
 - executing statements 21
 - ODBC function equivalents 527; 530; 531
 - positioned statements 45
- Empty strings 10; 85
- Environment handles
 - allocating 90

- defined 11
- driver 262
- freeing 56
- initializing 16
- overwriting 90
- See also Handles
- SQLFreeEnv 224
- state transitions 445
- with threads 90

Errors

- application processing 54
- clearing 178
- format 51
- messages 50
- queues 178
- return codes 49
- rowsets 200; 365
- See also specific function description 176
- source 50
- SQLError 176
- SQLSTATE values 429
- with parameter arrays 303

Escape characters 85; 273

Escape clauses

- datetime literals 29
- described 31
- ESCAPE clause 31; 265
- outer joins 32
- procedures 33
- scalar functions 30
- scanning for 379
- support of 265
- syntax 28

Examples

- ad hoc query 61
- data conversion 514; 526
- parameter values 23
- static SQL 57

Executable objects See Procedures

Expressions

- data types 477
- in ORDER BY clauses 263

Extensions to SQL 28

F

Fat cursors 40

Fetching data See Retrieving data

Files

- trace 345; 346
- usage 263

Filtered indexes 394

Floating point data

- converting to C 508
- converting to SQL 520
- converting with CONVERT 259
- specifying conversions
 - SQLBindCol 94
 - SQLBindParameter 107
 - SQLGetData 236

Foreign keys 217

Freeing handles 222; 224; 227

Functionality, driver 246; 253

Functions, ODBC

- asynchronous execution 26
- buffers 9
- canceling 127
- deprecated 353
- list of 75
- return codes 49
- See also specific function 16
- SQL statement equivalents 527; 530; 531; 532; 533
- state transitions 443
- supported 247

Functions, scalar

- CONVERT 31; 544
- date 540
- numeric 538
- string manipulation 536
- system 543
- time 540

G

Gateways 53

GLOBAL TEMPORARY table type 405

Global transactions 407

Grammar, SQL 477

Granting privileges 139; 399
GROUP BY clauses 264; 266

—H—

Handles

- connection See Connection handles
- defined 11
- error queues 178
- library 262
- retrieving 262
- SQLAllocConnect 88
- SQLAllocEnv 90
- SQLAllocStmt 92
- SQLFreeConnect 222
- SQLFreeEnv 224
- SQLFreeStmt 227
- statement See Statement handles

Hashed indexes 393

hdbc See Connection handles

Header files

- required 7
- SQL.H
 - C data types 12
 - contents 84
 - SQL data type 490; 491
- SQLEXT.H
 - C data types 12
 - contents 84
 - macros 366
 - SQL data type 492

henv See Environment handles

Hexadecimal characters 510; 519; 522

Holes in cursors 41; 273

Host variables 527

hstmt See Statement handles

—I—

Identifiers

- case sensitive 265
- quote character 265

Identifiers, quoted

- case sensitive 271

Incrementing, auto 132

Indexes

- cardinality 394
- clustered 393
- collating 394
- filtered 394
- hashed indexes 393
- listing 392
- maximum columns in 266
- maximum length 266
- owner usage in 270
- pages 394
- qualifier usage in 271
- See also SQLStatistics
- sorting 394
- unique 393

Information types, returning 252

Information, status

- retrieving 50

Input buffers 10

Input parameters 105; 320

Input/output parameters 106; 320

INSERT statements

- affected rows 338
- bulk 26
- owner usage in 270
- privileges 140; 399
- qualifier usage in 271
- See also SQLSetPos
- SQLParamOptions 302

Installer DLL

- See also Driver setup DLL

Integer data

- converting to C 508
- converting to SQL 520
- converting with CONVERT 259
- ODBC 1.0 compatibility 495
- specifying conversions
 - SQLBindCol 94
 - SQLBindParameter 107
 - SQLGetData 236

Integrity enhancement facility (IEF) 269; 469

Interoperability

- affected rows 36
- buffer length 10

- cursor names 350
- default C data types 496
- functionality 253
- functions 253
- procedure parameters 33
- pseudo-columns 387
- SQL statements
 - ALTER TABLE 19; 282
 - COMMIT 24; 185; 191; 309
 - CREATE TABLE 19; 282
 - ODBC extensions 28
 - ROLLBACK 185; 191; 309
 - syntax 19; 469
- SQLGetData 38; 40; 242; 264
- transactions 24
- transferring data 497
- Intervals, datetime 275; 541; 542
- Isolation levels, transaction 261; 277; 347

—J—

- Joins, outer
 - described 32
 - support of 269

—K—

- Keys
 - foreign 216
 - primary 314
- Keyset size 370; 372; 378
- Keyset-driven cursors 41
- Keywords
 - data source-specific 265
 - in SQLBrowseConnect 119; 120
 - in SQLDriverConnect 168
 - ODBC 486

—L—

- Labels, column 132
- Laufzeittrace
 - Optionen
 - sqltrace.pct 553
- Length, available
 - SQLBindParameter 111

- SQLExtendedFetch 197
- SQLFetch 211
- SQLGetData 241
- Length, buffer
 - input 10
 - output 10
 - SQLBindCol 99
 - SQLBindParameter 109
 - SQLGetData 241
- Length, column
 - defined 501
 - procedures 321
 - result sets 133; 153
 - tables 145; 387
- Length, data-at-execution 110; 111; 363
- Length, maximum
 - buffers 10
 - column names length 266
 - columns 211; 241
 - connection options 230; 343
 - cursor names 266; 350
 - data 378
 - error messages 177
 - indexes 266
 - literals 265
 - owner names 266
 - procedure names 266
 - qualifiers 266
 - rows 266
 - statement options 281; 375
 - table names 267
 - user names 267
- Length, unknown
 - in length 501
 - in precision 498
 - SQLBindParameter 111
 - SQLExtendedFetch 197
 - SQLFetch 211
 - SQLGetData 241
- Levels, conformance See Conformance levels
- Library handles, driver 262
- LIKE predicates
 - described 31
 - support of 265

Limitations, SQL statements 255

Literals

- binary 265
- character 265
- date 29
- prefix string 285
- procedure parameters 33
- suffix string 285
- time 29
- timestamp 29

Loading drivers 150

LOCAL TEMPORARY table type 405

Locking

- concurrency 369; 377
- row 361
- supported locks 265

Login authorization

- connection strings
 - SQLBrowseConnect 119; 120; 171
- example 151
- interval period 172; 344
- SQLSetConnectOption 344
- timeout 172

Long data values

- in row length 267
- length required 267
- retrieving 38
- sending
 - from rowsets 46
 - in parameters 25
 - SQLBindParameter 111
 - SQLSetPos 363
- SQLGetData 242

—M—

Macros

- data-at-execution 105; 108; 110; 111; 363
- SQLSetPos 366

Manual-commit mode

- beginning transactions 408
- described 23
- See also Transactions 23
- specifying 344
- SQLTransact 408

Manuals x

Memory, allocating

- buffers 9
- connection handles 11
- environment handles 11
- result sets 36
- rowsets 39
- SQLAllocConnect 88
- SQLAllocEnv 90
- SQLAllocStmt 92
- SQLFreeConnect 222
- SQLFreeEnv 224
- SQLFreeStmt 227
- statement handles 11

Messages, error See Errors

Mixed cursors 42

Modes

- access 172; 344
- auto-commit See Auto-commit mode
- manual-commit See Manual-commit mode

Money columns 133

Money data types 287

Multiple-tier drivers

- error messages 52
- identifying data sources 51

Multi-threading

- canceling functions 128
- with connection handles 88
- with environment handles 90
- with statement handles 92

—N—

Names

- arguments 81
- column labels 132
- correlation 259
- cursors 233; 350
- data source 260
- data types 285
- database 261
- DBMS 261
- driver 262
- index 393
- localized data types 287

- procedure 326
 - procedure columns 319
 - See also Terms
 - server 273
 - table 403
 - tables 405
 - user 277
- Network traffic
- maximum data length 378
 - packet size 346
 - prepared statements 21
- Nonrepeatable reads 261; 277; 347
- NOT NULL clauses 267
- NULL data
- collating 267
 - concatenation behavior 258
 - input buffers 10
 - output buffers 11
 - retrieving 99
 - SQLBindParameter 110
 - SQLExtendedFetch 197
 - SQLFetch 211
 - SQLGetData 241
 - SQLPutData 330
 - SQLSetPos 361
- Null pointers
- input buffers 10
 - output buffers 10
 - parameter length 110
 - unbinding columns 100
- Null termination byte
- binary data 110
 - character data 110
 - embedded 10
 - examples 514; 526
 - input buffers 10
 - output buffers 10
 - parameters 23
- Nullability
- columns 133; 145; 155; 321
 - data types 286
 - parameters 159
- Numeric data
- auto incrementing 132
 - converting to C 508
 - converting to SQL 520
 - converting with CONVERT 259
 - currency data type 287
 - money data type 287
 - radix 145; 321
 - scalar functions 268
 - See also Floating point data 538
 - See also Integer data
 - specifying conversions
 - SQLBindCol 94
 - SQLBindParameter 107
 - SQLGetData 236
- Numeric functions 268; 538
- O—
- Objects, executable See Procedures
- ODBC
- described ix
 - functions See Functions, ODBC
 - version compatibility See Version compatibility
- ODBC architecture
- error handling 50
 - gateways 53
- ODBC.INF
- keywords
 - retrieving 175
- ODBC.INI
- keywords
 - trace 90
- ODBC functions using
- SQLAllocEnv 90
 - SQLBrowseConnect 121
 - SQLDriverConnect 169
 - SQLFreeEnv 224
 - SQLSetConnectOption 345; 346
- Setup DLL functions using
- ConfigDSN 414
- ODBCINST.INI
- listing drivers 175
- Optimistic concurrency control 369; 377 See Concurrency
- Options
- connection See Connection options

- fetch 201
- statement See Statement options
- ORDER BY clauses
 - columns in select list 269
 - expressions in 263
 - maximum columns in 266
- Outer joins
 - described 32
 - support of 269
- Output buffers 10
- Output parameters 106; 320
- Owner names
 - maximum length 266
 - procedure 319; 326
 - table 392; 399; 403
 - SQLColAttributes 133
 - SQLColumnPrivileges 139
 - SQLColumns 144
 - term, vendor-specific 269

—P—

- Packet size 346
- Pages, index or table 394
- Parameters
 - arrays 26; 109; 302
 - binding 23; 105
 - data types 158; 470
 - data-at-execution 25; 105; 108; 110; 111
 - default 107
 - descriptors 161
 - host variables 527
 - input 105; 320
 - input/output 106; 320
 - interoperability 33
 - long data values 25
 - markers 33; 470
 - nullability 159
 - number of 294
 - order 161
 - output 106; 320
 - preparing 309
 - procedure 33; 105; 318
 - return values 33; 319
 - See also Arguments

- single-quotes 108
- unbinding 23; 226
- values 22
- version compatibility 113; 353
- Patterns, search patterns 85
- Percent sign (%) 85; 403
- Phantoms 261; 277; 347
- Plans, access 21
- Pointers, maintaining 10
- Pointers, null See Null pointers
- Position, cursor See Cursor position
- Positioned delete statements
 - cursor name 233; 350
 - cursor position 198; 359
 - cursor simulation 380
 - executing 45
 - support of 270
 - version compatibility 46
- Positioned update statements
 - cursor name 233; 350
 - cursor position 198; 359
 - cursor simulation 380
 - executing 45
 - support of 270
 - version compatibility 46
- Precision
 - columns
 - procedures 320
 - result sets 133; 154
 - tables 144; 387
 - data types 285
 - defined 498
- Prepared statements, deleting 260
- Preparing statements 21; 309
- Preserving cursors 260
- Primary keys 217; 314
- Privileges
 - columns 135; 138
 - data source 260
 - grantable 140; 400
 - grantee 139; 399
 - grantor 139; 399
 - owner usage in 270
 - qualifier usage in 271

table user granting 139
tables 399

Procedure columns

data type 320
listing 318
name 319
owner name 319
parameters 105
qualifier 319
See also Columns
See also Procedures

Procedures

accessibility 257
defined 323
interoperability 33
name 326
name, maximum length 266
ODBC syntax 33
owner name 326
owner usage in 270
qualifier 326
qualifier usage in 271
return values 106; 320; 327
See also Procedure columns
support of 270
term, vendor-specific 270

Pseudo-columns 387

—Q—

Qualifiers

current 344
foreign key 217
index 393
location 271
maximum length 266
primary key 217; 314
procedure 319; 326
separator 271
table 133; 139; 144; 392; 399; 403; 405
term, vendor-specific 271
usage 271

Queries See SQL statements

Question mark (?)

browse result connection string 121

parameter markers 185; 309

Queues, error 178

Quote character 265

Quoted identifiers

case sensitivity 271

—R—

Radix 145; 321

Read committed isolation level 261; 277; 347

Read only

access mode 172; 344

columns 135

concurrency 369; 377

data sources 260

Read uncommitted isolation level 261; 277; 347

Read/write access mode 172; 344

Reads, dirty 261

Reads, nonrepeatable 261

REFERENCES statements 140; 399

Referential integrity 269; 469

Refreshing data

SQLExtendedFetch 201

SQLSetPos 360

Remarks 145; 405

Repeatable read isolation level 261; 277; 347

Restricted deletes 218

Restricted updates 218

Result sets

arrays See Rowsets

binding columns 36

column attributes 36

described 35

described2 13

fetching data 37

fetching rowsets 40

multiple 48; 267

number of columns 36

number of rows 36

retrieving data in parts 38

rowset size 39

See also Cursors 40

See also Rows

SQLBindCol 98

SQLColAttributes 132

- SQLDescribeCol 153
- SQLFreeStmt 226
- SQLMoreResults 290
- SQLNumResultCols 296
- SQLProcedures 325
- SQLRowCount 338
- Result states See State transitions
- Retrieving data
 - arrays See Rowsets
 - binding columns See Binding columns
 - cursor position 198; 210; 359
 - disabling 379
 - fetch directions 263
 - fetching data 37
 - fetching rowsets 40
 - in parts 38; 242
 - long data values 38
 - maximum length 378
 - multiple result sets 48; 267
 - NULL data 99; 211; 241
 - retrieving 379
 - rows See Rows
 - SQLExtendedFetch 196
 - SQLFetch 210
 - SQLGetData 242
 - truncating data 211; 242
 - unbound columns 38
- Return codes 49
- Return values, procedure 106; 320; 327
- ROLLBACK
 - SQLExecDirect 185
- ROLLBACK statements
 - cursor behavior 260
 - interoperability 24; 185; 191
 - SQLExecute 191
 - SQLPrepare 309
- Rolling back transactions 24; 407
- Row status array
 - errors 200; 365
 - SQLExtendedFetch 203
 - SQLSetPos 360; 365
 - updating 360
- Row versioning isolation level 261; 277; 347
- ROWID 386; 388

- Rows
 - adding 46
 - affected 338
 - after last row 201
 - concurrency 369; 377
 - current 281
 - deleting 45; 48
 - errors in 200; 365
 - interoperability 36
 - locking 265; 361
 - maximum 378
 - maximum length 266
 - refreshing 360
 - See also Cursors
 - See also Retrieving data
 - See also Rowsets
 - SQLExtendedFetch 203
 - SQLSetPos 359
 - updated 272
 - updating 45; 47
- Rowsets
 - allocating 39
 - binding 39
 - binding type 376
 - cursor position 198; 359
 - errors in 200; 365
 - modifying 46
 - retrieving 40
 - size 372; 379
 - SQLExtendedFetch 196
 - SQLSetPos 359
 - status 203

—S—

- SAG CLI compliance 268
- Scalar functions
 - data conversion 259; 544
 - date functions 276; 540
 - information types 256
 - numeric functions 268; 538
 - string manipulation 274; 536
 - syntax 30
 - system functions 275; 543
 - time functions 276; 540

- TIMESTAMPADD intervals 275
- TIMESTAMPDIFF intervals 275
- Scale
 - columns
 - procedures 321
 - result sets 134; 154
 - tables 145; 387
 - data types 287
 - defined 500
- Scanning for escape clauses 379
- Scope, rowid 386; 388
- Scrollable cursors 40
- Search patterns 85
 - escape character 273
- Searchability
 - columns 134
 - data types 286
- Segment boundaries 10
- SELECT FOR UPDATE statements 46; 270
- Select list
 - maximum columns in 266
 - ORDER BY columns in 269
- SELECT statements
 - affected rows 338
 - bulk 302
 - cursor name 185; 230; 350
 - maximum rows 378
 - maximum tables in 267
 - multiple result sets 45; 290
 - owner usage in 270
 - privileges 140; 399
 - qualifier usage in 271
 - reexecuting 191
 - See also Result sets
 - UNION clauses 277
- Sensitivity, cursor 273
- Separator, qualifier 271
- Serializability 43
- Serializable isolation level 261; 277; 347
- Server name 273
- Signed columns 134
- Signed data types 287
- Simulated cursors
 - by applications 386
 - controlling 380
- Simulating transactions 361
- Single-quote (') conversion 108
- Single-tier drivers
 - error messages" 52
 - file usage 263
- Size, display 132; 502
- Sorting 267; 394
- Special characters
 - in search patterns 85
 - in SQL identifiers 273
 - in SQLBrowseConnect 119; 120
 - in SQLDriverConnect 168
- SQL
 - data types See SQL data types
 - dynamic 532
 - embedded
 - executing statements 21
 - ODBC function equivalents 527; 530; 531
 - positioned statements 45
 - information types 254; 255
 - interoperability 19; 469
 - keywords 486
 - ODBC extensions to 28
 - ODBC grammar 469
 - referential integrity 469
 - See also SQL statements 472
 - static 57
- SQL Access Group 268
- SQL data types
 - columns
 - indexes 393
 - procedures 320
 - result sets 132; 134; 154
 - tables 144; 387
 - conformance levels 490; 491; 492
 - conversion examples 514; 526
 - converting from C data types 515; 518; 520; 521; 522; 523; 524; 525
 - converting to C data types 503; 505; 508; 509; 510; 511; 512; 513
 - converting with CONVERT 259
 - default C data types 496
 - defined 490

display size 502
in translation DLLs 420; 423
length 501
parameters 158; 470
precision 498
scale 500
See also C data types 522
See also Converting data 522
specifying conversions
 SQLBindCol 94
 SQLBindParameter 107
 SQLGetData 236
supported 284; 285
type names 285

SQL identifiers
 case sensitivity 265
 quote character 265
 special characters 273

SQL quoted identifiers
 case sensitivity 271

SQL statements
 batch 290; 477
 direct execution 21; 22
 embedded 527; 530
 ODBC function equivalents 530; 531
 information types 254
 interoperability 469
 keywords 486
 maximum length 267
 modifying syntax 19
 native 292
 owner usage in 270
 prepared execution 21
 qualifier usage in 271
 query timeout 379
 scanning 379
 See also Conformance levels 26
 See also specific statements 26
 See also SQL 26
 SQLCancel 127
 SQLExecDirect 185
 SQLExecute 191
 SQLPrepare 309
 terminating 55

SQL.H See Header files

SQL_ERROR 49

SQL_INVALID_HANDLE 49

SQL_NEED_DATA 49

SQL_NO_DATA_FOUND 49

SQL_STILL_EXECUTING 49

SQL_SUCCESS 49

SQL_SUCCESS_WITH_INFO 49

SQLAllocConnect
 allocating handles 16
 function description 87
 with SQLConnect 150

SQLAllocEnv
 function description 89
 initializing handles 16
 with SQLConnect 150

SQLAllocStmt 20; 91

SQLBindCol
 binding columns 36
 code example 100
 function description 93
 truncating data 99
 unbinding columns 55; 100
 with SQLFetch 210

SQLBindParameter
 code example 114
 function description 102
 replaces SQLSetParam 106; 107
 sending long data values 25
 unbinding parameters 55; 226
 version compatibility 113; 353
 with SQLParamData 108; 111
 with SQLParamOptions 108; 109
 with SQLPutData 111

SQLBrowseConnect
 code example 122
 Driver Manager 121
 function description 116
 with SQLDisconnect 122

SQLCancel
 function description 126
 terminating asynchronous execution 55

SQLColAttributes
 column attributes 36

- function description 129
 - versus SQLDescribeCol 135
 - versus SQLGetTypeInfo 288
- SQLColumnPrivileges 24; 136
- SQLColumns
- catalog 24
 - code example 146
 - function description 141
 - intended usage 143
- SQLConnect
- code example 151
 - connecting 18
 - Driver Manager 150
 - function description 148
 - with SQLAllocConnect 150
 - with SQLAllocEnv 150
 - with SQLSetConnectOption 150
- SQLDataSourceToDriver 420
- SQLDescribeCol 36; 153
- SQLDescribeParam 158
- SQLDisconnect
- disconnecting 56
 - function description 162
 - with SQLBrowseConnect 122
 - with SQLFreeConnect 222
- SQLDriverConnect
- Driver Manager 169
 - function description 164
 - login interval 172
- SQLDrivers
- function description 173
- SQLDriverToDataSource 423
- SQLError
- Driver Manager 177
 - function description 176
 - See also Errors 49
 - See also SQLSTATes 429
- SQLExecDirect
- direct execution 21
 - function description 179
- SQLExecute
- function description 187
 - prepared execution 21
 - with SQLPrepare 191
- SQLExt.H See Header files
- SQLExtendedFetch
- code examples 204
 - function description 193
 - retrieving rowsets 40
 - with SQLGetData 40
 - with SQLSetPos 198
 - with SQLSetStmtOption 196
- SQLFetch
- fetching data 37
 - function description 208
 - length transferred 133
 - positioning cursor 37
 - with SQLBindCol 210
 - with SQLGetData 211
- SQLForeignKeys 24; 213
- SQLFreeConnect
- freeing connection handles 56
 - function description 222
 - with SQLDisconnect 222
 - with SQLFreeEnv 224
- SQLFreeEnv
- freeing environment handles 56
 - function description 224
 - with SQLFreeConnect 224
- SQLFreeStmt
- closing cursors 37; 226
 - freeing statement handles 55; 227
 - function description 226
 - unbinding columns 55; 100; 226
 - unbinding parameters 23; 105; 226
- SQLGetConnectOption 229
- SQLGetCursorName 232
- SQLGetData
- code example 243
 - extensions 264
 - function description 235
 - interoperability 38; 40
 - length transferred 133
 - long data values 38; 98
 - unbound columns 38
 - with SQLExtendedFetch 40
 - with SQLFetch 211
- SQLGetFunctions

code example 248
function description 245
using 24

SQLGetInfo
code example 277
data conversion 256
data sources 254
DBMSs 254
drivers 253
function description 250
scalar functions 256
SQL statements 254

SQLGetStmtOption 279

SQLGetTypeInfo
ALTER TABLE statements 469
CREATE TABLE statements 469
function description 282
supported data types 12; 490
versus SQLColAttributes 288

SQLMoreResults 48; 289

SQLNativeSql 292

SQLNumParams 294

SQLNumResultCols 36; 296

SQLParamData
data-at-execution parameters 25
function description 299
with SQLBindParameter 108; 111
with SQLPutData 299
with SQLSetPos 363

SQLParamOptions
code example 303
function description 302
multiple parameter values 26; 302
with SQLBindParameter 108; 109

SQLPrepare
function description 306
preparing statements 21
with SQLExecute 191

SQLPrimaryKeys 24; 312

SQLProcedureColumns
catalog 25
function description 316
listing columns 34

SQLProcedures
catalog 25
code example 327
function description 323
listing procedures 34

SQLPutData
code example 334
data-at-execution parameters 25; 363
function description 330
with SQLBindParameter 111
with SQLParamData 299
with SQLSetPos 363

SQLRemoveDSNFromIni 415

SQLRowCount
affected rows 48
function description 338
interoperability 36

SQLSetConnectOption
commit mode 23; 344
function description 340
See also Connection options
with SQLConnect 150
with SQLTransact 344

SQLSetCursorName 349

SQLSetParam 113; 353

SQLSetPos
code example 366
function description 354
lock types supported 265
locking rows 43
macros 366
modifying rowsets 46
operations supported 270
with SQLExtendedFetch 198
with SQLParamData 363
with SQLPutData 363

SQLSetScrollOptions 369

SQLSetStmtOption
function description 373
See also Statement options 373
with SQLExtendedFetch 196

SQLSpecialColumns 25; 382

SQLSTATEs
guidelines 49
naming conventions 84

ADABAS D

- values 429

SQLStatistics 25; 389

SQLTablePrivileges 24; 396

SQLTables

- argument syntax 403
- catalog 24
- formatting 403
- function description 401

SQLTransact

- commit mode 344
- committing 23; 56
- function description 406
- rolling back 23; 56
- two-phase commit 407
- with SQLSetConnectOption 344

SQLWriteDSNToIni 414

SQLWritePrivateProfileString 414

State transitions

- connection handles 447
- defined 443
- environment handles 445
- statement handles 453

Statement handles

- active 257
- allocating 20
- defined 11
- driver 262
- freeing 55; 375
- overwriting 92
- See also Handles
- SQLAllocStmt 92
- SQLFreeStmt 227
- SQLMoreResults 290
- state transitions 453
- with threads 92

Statement options

- asynchronous execution 27; 376
- binding type 376
- concurrency 43; 377
- current row 281
- cursor type 377
- driver-specific 375
- getting bookmarks 281
- keyset size 378

- maximum data length 378

- maximum length 281; 375

- maximum rows 378

- query timeout 379

- releasing 375

- reserved 375

- retrieving 281

- retrieving data 379

- scanning SQL statements 379

- setting 375

- simulated 380

- substituting values 375

- using bookmarks 380

Statements See SQL statements

Static cursors

- described 41
- sensitivity 273

Static SQL 57

Statistics, listing 392

Status array

- errors 200; 365
- SQLExtendedFetch 203
- SQLSetPos 360; 365
- updating 360

Status information

- retrieving 50
- See also Errors 50

Stored procedures See Procedures

String data See Character data

String functions 274; 536

Subqueries 274

SYNONYM table type 405

Syntax

- connection strings 119; 120; 168
- ODBC SQL 28
- SQL 477

System functions 275; 543

SYSTEM TABLE table type 405

—T—

Table definition statements

- owner usage in 270
- qualifier usage in 271

TABLE table type 405

Tables

- accessibility 257
- cardinality 394
- columns See Columns
- correlation names 259
- defined 84
- foreign keys 217
- in SELECT statement 267
- indexes See Indexes
- maximum columns in 266
- names
 - maximum length 267
 - retrieving 392; 399
 - special characters 273
 - SQLColAttributes 132
 - SQLColumnPrivileges 139
 - SQLColumns 144
 - SQLTables 403; 405
- owner name See Owner names
- pages 394
- primary keys 217; 314
- privileges See Privileges
- qualifier See Qualifiers
- rows See Rows
- statistics 392
- term, vendor-specific 275
- types 405
- versus views 84

Terminating

- asynchronous execution 55
- transactions 56

Termination byte, null See Null termination byte

Terms, vendor-specific

- owner 269
- procedure 270
- qualifier 271
- table 275

Threads, multiple

- canceling functions 128
- with connection handles 88
- with environment handles 90
- with statement handles 92

Time data

- converting to C 512
- converting to SQL 524
- converting with CONVERT 259
- literals 29
- scalar functions 276; 540
- specifying conversions
 - SQLBindCol 94
 - SQLBindParameter 107
 - SQLGetData 236

Time functions 276; 540

Timeout

- login 172; 344
- query 379

Timestamp data

- converting to C 513
- converting to SQL 525
- converting with CONVERT 259
- literals 29
- scalar functions 540
- specifying conversions
 - SQLBindCol 94
 - SQLBindParameter 107
 - SQLGetData 236
- TIMESTAMPADD intervals 275
- TIMESTAMPDIFF intervals 275

Titles, column 132

Trace file 345; 346

Trace keyword

- SQLAllocEnv 90
- SQLFreeEnv 224

TraceAutoStop keyword 224

Tracing

- SQLAllocEnv 90
- SQLFreeEnv 224
- SQLSetConnectOption 345; 346

Traffic, network

- maximum data length 378
- packet size 346
- prepared statements 21

Transactions

- access plan behavior 21
- beginning 23; 408
- commit mode 23; 344
- committing 24; 407
- concurrency See Concurrency

- cursor behavior 38; 260; 408
- incomplete 163
- interoperability 24
- isolation levels 261; 277; 347
- multiple active 267
- ODBC function equivalents 533
- rolling back 24; 407
- serializability 43
- simulating 361
- SQLExecDirect 185
- SQLExecute 191
- SQLPrepare 310
- SQLTransact 407
- support of 276
- terminating 56
- two-phase commit 407

Transferring binary data 497

Transitions, state See State transitions

Translation DLLs

- default 172
- default option 416
- described 348
- See also Converting data 503
- setup DLL See Translation setup DLL
- specifying 346
- SQLDataSourceToDriver 422
- SQLDriverConnect 172
- SQLDriverToDataSource 425
- truncating data 422; 425

Translation options

- default 416
- described 348
- specifying 346
- SQLDriverConnect 172

Translation setup DLL

- ConfigTranslator 411
- See also Translation DLLs

Truncating data

- maximum data length 378
- output buffers 11
- See also Binary data 489
- See also Character data 489
- SQLBindCol 99
- SQLBindParameter 109

- SQLFetch 211
- SQLGetData 242
- translation DLLs 422; 425

Two-phase commit 407

Types, information 252

Typographic conventions x

—U—

Unbinding columns 100; 226

Unbinding parameters 226

Underscore (_) 85

UNION clauses 277

Unique indexes 393

Unloading drivers 151

Updatability 135

Update rules 218

UPDATE statements

- affected rows 338
- bulk 26; 302
- cascade update 218
- owner usage in 270
- privileges 140; 399
- qualifier usage in 271
- restrictive update 218
- See also SQLSetPos

Updates, positioned See Positioned update statements

User names

- maximum length 267
- retrieving 277

—V—

Values

- procedure return 106

Values, compare before update 369; 377

Variables, host 527

Version

- DBMS 261
- driver
 - number 263
 - ODBC version 262
- Driver Manager 269

Version compatibility

- C data types 495
- default C data types 497
- DRIVER keyword 121
- fetch options 201
- filtered indexes 394
- information types 253
- number of input parameters 326
- number of output parameters 326
- number of result sets 326
- parameter binding 113; 353
- positioned statements 46
- procedure type 327
- pseudo columns 388
- SQLBindParameter 113; 247; 353
- SQLForeignKeys 218
- SQLGetTypeInfo 288
- SQLPrimaryKeys 315
- SQLProcedureColumns 320
- SQLSetParam 113; 247; 353
- SQLSetScrollOptions 369

- Versioning isolation level 261; 277; 347
- Vertical bar (|) x
- VIEW table type 405
- Views 84

—W—

- Window handles
 - null 414
 - parent
 - ConfigDSN 414
 - quiet mode 346
 - See also SQLDriverConnect
- Windows NT registry
 - data sources
 - adding 414
 - configuring 414
 - removing 415
 - drivers 175
- Write privileges, column 135

Contents

About This Manual.....	
Organization of this Manual.....	
Document Conventions.....	

Part 1 Introduction to ODBC

- omitted -

Part 2 Developing Applications

Chapter 3 Guidelines for Calling ODBC Functions.....	
General Information.....	
Determining Driver Conformance Levels.....	
Determining API Conformance Levels.....	
Determining SQL Conformance Levels.....	
Using the Driver Manager (Unix, OS/2).....	
Using the Driver Manager (Windows).....	
Calling ODBC Functions.....	
Buffers.....	
Input Buffers.....	
Output Buffers.....	
Environment, Connection, and Statement Handles.....	
Using Data Types.....	
ODBC Function Return Codes.....	
Chapter 4 Basic Application Steps.....	
Chapter 5 Connecting to the SERVERDB.....	
About Data Sources (Windows).....	
Initializing the ODBC Environment.....	
Allocating a Connection Handle.....	
Connecting to the Database (Unix, OS/2).....	
Connecting to the Database (Windows).....	
Additional Functions.....	
Chapter 6 Executing SQL Statements.....	
Allocating a Statement Handle.....	
Executing an SQL Statement.....	
Prepared Execution.....	
Direct Execution.....	
Setting Parameter Values.....	
Performing Transactions.....	
ODBC Extensions for SQL Statements.....	

Retrieving Information About the Data Source's Catalog.....	
Sending Parameter Data at Execution Time.....	
Specifying Arrays of Parameter Values.....	
Executing Functions Asynchronously.....	
Using ODBC Extensions to SQL.....	
Date, Time, and Timestamp Data.....	
Scalar Functions.....	
Data Type Conversion Function.....	
LIKE Predicate Escape Characters.....	
Outer Joins.....	
Procedures.....	
Additional Extension Functions.....	
Chapter 7 Retrieving Results.....	
Assigning Storage for Results (Binding).....	
Determining the Characteristics of a Result Set.....	
Fetching Result Data.....	
Using Cursors.....	
ODBC Extensions for Results.....	
Retrieving Data from Unbound Columns.....	
Assigning Storage for Rowsets (Binding).....	
Column-Wise Binding.....	
Row-Wise Binding.....	
Retrieving Rowset Data.....	
Using Block and Scrollable Cursors.....	
Block Cursors.....	
Scrollable Cursors.....	
Static Cursors.....	
Dynamic Cursors.....	
Keyset-Driven Cursors.....	
Mixed (Keyset/Dynamic) Cursors.....	
Specifying the Cursor Type.....	
Specifying Cursor Concurrency.....	
Using Bookmarks.....	
Modifying Result Set Data.....	
Executing Positioned Update and Delete Statements.....	
Modifying Data with SQLSetPos.....	
Processing Multiple Results.....	
Chapter 8 Retrieving Status and Error Information.....	
Function Return Codes.....	
Retrieving Error Messages.....	
ODBC Error Messages.....	
Error Text Format.....	
Sample Error Messages.....	
Single-Tier Driver.....	

Multiple-Tier Driver.....	
Gateways.....	
Driver Manager.....	
Processing Error Messages.....	
Chapter 9 Terminating Transactions and Connections.....	
Terminating Statement Processing.....	
Terminating Transactions.....	
Terminating Connections.....	
Chapter 10 Constructing an ODBC Application.....	
Sample Application Code.....	
Static SQL Example.....	
Interactive Ad Hoc Query Example.....	

Part 3 Developing Drivers

- omitted -

Part 4 Installing and Configuring ODBC Software

- omitted -

Part 5 API Reference

Chapter 21 Function Summary.....	
ODBC Function Summary.....	
Chapter 22 ODBC Function Reference.....	
Arguments.....	
ODBC Include Files.....	
Diagnostics.....	
Tables and Views.....	
Catalog Functions.....	
Search Pattern Arguments.....	
SQLAllocConnect.....	
SQLAllocEnv.....	
SQLAllocStmt.....	
SQLBindCol.....	
SQLBindParameter.....	
fParamType Argument.....	
fCType Argument.....	
fSqlType Argument.....	
cbColDef Argument.....	
rgbValue Argument.....	

	cbValueMax Argument.....	
	pcbValue Argument.....	
	Passing Parameter Values.....	
	Conversion of Calls to and from SQLSetParam.....	
SQLBrowseConnect.....		
	Comments szConnStrIn Argument.....	
	szConnStrOut Argument.....	
	Using SQLBrowseConnect.....	
SQLCancel.....		
	Canceling Asynchronous Processing.....	
	Canceling Functions that Need Data.....	
	Canceling Functions in Multithreaded Applications.....	
SQLColAttributes.....		
SQLColumnPrivileges.....		
SQLColumns.....		
SQLConnect.....		
SQLDescribeCol.....		
SQLDescribeParam.....		
SQLDisconnect.....		
SQLDriverConnect (Windows).....		
	Comments Connection Strings.....	
	Driver Manager Guidelines.....	
	Driver Guidelines.....	
	Connection Options.....	
SQLDrivers (Windows).....		
SQLError.....		
SQLExecDirect.....		
SQLExecute.....		
SQLExtendedFetch.....		
	Binding.....	
	Column-Wise Binding.....	
	Row-Wise Binding.....	
	Positioning the Cursor.....	
	Processing Errors.....	
	<i>fFetchType</i> Argument.....	
	Moving by Row Position.....	
	Positioning to a Bookmark.....	
	<i>irow</i> Argument.....	
	<i>rgfRowStatus</i> Argument.....	
	Column-Wise Binding.....	
	Row-Wise Binding.....	
SQLFetch.....		
SQLForeignKeys.....		
SQLFreeConnect.....		
SQLFreeEnv.....		

SQLFreeStmt.....	
SQLGetConnectOption.....	
SQLGetCursorName.....	
SQLGetData.....	
SQLGetFunctions.....	
SQLGetInfo.....	
Information Types.....	
Driver Information.....	
DBMS Product Information.....	
Data Source Information.....	
Supported SQL.....	
SQL Limits.....	
Scalar Function Information.....	
Conversion Information.....	
Information Type Descriptions.....	
SQLGetStmtOption.....	
SQLGetTypeInfo.....	
SQLMoreResults.....	
SQLNativeSql.....	
SQLNumParams.....	
SQLNumResultCols.....	
SQLParamData.....	
SQLParamOptions.....	
SQLPrepare.....	
SQLPrimaryKeys.....	
SQLProcedureColumns.....	
SQLProcedures.....	
SQLPutData.....	
SQLRowCount.....	
SQLSetConnectOption.....	
Data Translation.....	
SQLSetCursorName.....	
SQLSetParam.....	
SQLSetPos.....	
Comments <i>irow</i> Argument.....	
<i>fOption</i> Argument.....	
<i>fLock</i> Argument.....	
Using SQLSetPos.....	
Performing Bulk Operations.....	
SQLSetPos Macros.....	
SQLSetScrollOptions.....	
SQLSetStmtOption.....	
SQLSpecialColumns.....	
SQLStatistics.....	
SQLTablePrivileges.....	

SQLTables.....	
SQLTransact.....	
Chapter 23 Setup DLL Function Reference (Windows).....	
ConfigDSN.....	
Adding a Data Source.....	
Modifying a Data Source.....	
Deleting a Data Source.....	
ConfigTranslator.....	
Chapter 24 - omitted -	
Chapter 25 Translation DLL Function Reference (Windows).....	
SQLDataSourceToDriver.....	
SQLDriverToDataSource.....	

Appendixes

Appendix A ODBC Error Codes.....	
Appendix B ODBC State Transition Tables.....	
Environment Transitions.....	
Connection Transitions.....	
Statement Transitions.....	
Appendix C SQL Grammar.....	
Parameter Data Types.....	
Parameter Markers.....	
SQL Statements.....	
Elements Used in SQL Statements.....	
List of Reserved Keywords.....	
Appendix D Data Types.....	
SQL Data Types.....	
Minimum SQL Data Types.....	
Core SQL Data Types.....	
Extended SQL Data Types.....	
C Data Types.....	
Core C Data Types.....	
Extended C Data Types.....	
Bookmark C Data Type.....	
ODBC 1.0 C Data Types.....	
Default C Data Types.....	
Transferring Data in its Binary Form.....	
Precision, Scale, Length, and Display Size.....	
Precision.....	
Scale.....	
Length.....	

Display Size.....	
Converting Data from SQL to C Data Types.....	
SQL to C: Character.....	
SQL to C: Numeric.....	
SQL to C: Bit.....	
SQL to C: Binary.....	
SQL to C: Date.....	
SQL to C: Time.....	
SQL to C: Timestamp.....	
SQL to C Data Conversion Examples.....	
Converting Data from C to SQL Data Types.....	
C to SQL: Character.....	
C to SQL: Numeric.....	
C to SQL: Bit.....	
C to SQL: Binary.....	
C to SQL: Date.....	
C to SQL: Time.....	
C to SQL: Timestamp.....	
C to SQL Data Conversion Examples.....	
Appendix E Comparison Between Embedded SQL and ODBC.....	
ODBC to Embedded SQL.....	
Embedded SQL to ODBC.....	
Declarative Statements.....	
Data Definition Statements.....	
Data Manipulation Statements.....	
Dynamic SQL Statements.....	
Transaction Control Statements.....	
Association Management Statements.....	
Appendix F Scalar Functions.....	
String Functions.....	
Numeric Functions.....	
Time and Date Functions.....	
System Functions.....	
Explicit Data Type Conversion.....	
Appendix G - omitted -	
Appendix H ODBC Library (Unix, OS/2).....549	
ODBC Library (Unix, OS/2).....	
Installed Files.....	
Appendix I ODBC Library (Windows).....551	
Installed Files.....	
Appendix J ODBC Run-Time Options.....553	
Setting Options.....	

Connecting to a Database Session.....	
Trace of SQL Commands.....	
Further Trace File Options.....	

Index.....	
-------------------	--

