

Application Programmer's Guide

Contents

μ	
Introduction	7
Audience	7
Document Conventions.....	8
Where to Find Additional Information.....	8
1 ODBC Theory of Operation.....	9
ODBC Components.....	10
Application.....	11
Driver Manager.....	11
Driver	11
Data Source.....	12
Types of Drivers.....	12
Single-Tier Configuration.....	12
Multiple-Tier Configuration.....	13
Network Example.....	14
Matching an Application to a Driver.....	15
ODBC Function Call Support.....	15
SQL Statement Support.....	15
How to Select a Set of Functionality.....	16
Connections and Transactions.....	16
2 A Short History of SQL.....	17
SQL Background Information.....	17
ANSI 1989 Standard.....	17
Embedded SQL.....	17
Future ANSI Specifications.....	18
Dynamic SQL.....	18
Call Level Interface.....	19
Interoperability.....	19
3 Guidelines for Calling ODBC Functions.....	21
Calling ODBC Functions.....	21
General Information.....	21
Variable Length Data in Function Arguments.....	21
Environment, Connection, and Statement Handles.....	22
Submitting SQL Statements.....	23
Data Type Support.....	23
Handling Results.....	24
Handling Errors.....	24
Processing Result Sets.....	25

4 Basic Application Steps.....	27
Additional Information.....	28

5	Establishing Connections.....	29
	Using the Driver Manager.....	29
	Initializing the ODBC Environment.....	29
	Establishing a Connection to a Data Source.....	29
	Accessing ODBC Functions.....	30
	Extensions for Establishing Connections.....	30
	SQLDriverConnect.....	31
6	Preparing and Executing an SQL Statement.....	33
	Allocating a Statement Handle.....	34
	Assigning Storage for Results (Binding).....	34
	Choosing Prepared or Direct Execution.....	34
	Prepared Execution.....	34
	Direct Execution.....	35
	Processing Positioned Updates and Deletes.....	35
	Example	36
	Extensions for Processing SQL Statements.....	37
	Obtaining Information about the Data.....	37
	Sending Large Data Values.....	38
	Specifying Arrays of Parameter Values.....	38
	Using Scrollable Cursors.....	39
	Basic Cursors.....	39
	Scrollable Cursors.....	39
	Requesting Asynchronous Processing.....	40
	Using Extended Data Types, Functions, and Outer Joins.....	40
	Date and Time Data Types.....	41
	Scalar Functions.....	42
	Data Type Conversion Function.....	42
	Requesting Outer Joins.....	43
7	Retrieving Results.....	45
	Determining Characteristics of a Result Set.....	45
	Fetching Result Data.....	45
	Retrieving Error and Status Information.....	46
	Related Extension Functions.....	47
	Retrieving Data in Large Columns.....	47
	Processing Multiple Result Sets.....	47
	Processing Blocks of Results.....	48
8	Terminating Transactions and Connections.....	49
	Terminating Statement Processing.....	49
	Terminating Transactions.....	49
	Terminating Connections.....	49
9	Constructing an ODBC Application.....	51
	Developer's Kit Contents.....	51

System Requirements.....	52
Hardware Requirements.....	52
Software Requirements.....	52
Environmental Requirements.....	52
Installing the Developer's Kit.....	53
Constructing an ODBC Environment.....	57
Data Source Specification.....	57
Default Data Source Specification.....	58
Sample Data Source Specifications.....	58
ODBC Data Source List.....	59
How ODBC Functions Use the ODBC.INI File.....	59
SQLConnect.....	59
SQLDataSources.....	60
SQLDriverConnect.....	60
The ODBC SETUP Routine.....	60
Sample Application Code.....	61
Static SQL Example.....	61
Interactive Ad-Hoc Query Example.....	64
Testing and Debugging an Application.....	67
Support	67
 Appendix A SQL Grammar.....	 69
Elements Used in SQL Statements.....	73
List of Reserved Keywords.....	79
Index	83

Introduction

The ODBC (Open Database Connectivity) interface is a C programming language interface for database connectivity. The *ODBC Application Programmer's Guide* is designed to address the following three questions:

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

The guide is organized into the following chapters:

- Chapter 1, “ODBC Theory of Operation,” provides conceptual information about the ODBC interface.
- Chapter 2, “A Short History of SQL,” contains a brief history of SQL.
- Chapters 3 through 8 describe how and when to call ODBC functions.
- Chapter 9, “Constructing an ODBC Application,” lists developer's kit contents, describes how to build an application, and includes sample code.

Appendix A contains SQL syntax information.

For information about the syntax and semantics of each ODBC function, refer to the *ODBC API Reference*. For information about driver development, refer to the *ODBC Driver Developer's Guide*.

Audience

The ODBC software development kit is available for use with the C language in a Windows environment. Use of the ODBC interface spans four areas of knowledge: SQL statements, ODBC function calls, C programming, and Windows programming. This manual assumes the following expertise:

- Working knowledge of the C programming language.
- General database knowledge and a familiarity with SQL.

Document Conventions

This manual uses the following typographic conventions.

Format	Used for
WIN.INI	Names of applications, programs, and other files.
RETCODE SQLFetch (hDBC)	Sample command lines and program code.
<i>argument</i>	Information that you or the application must provide, or word emphasis.
SQLTransact	Syntax that must be typed exactly as shown, including function names.
[]	Optional items or, if in bold text, brackets that must be included in the text string.
	Separates two mutually exclusive choices in a syntax line.
{ }	Delimits mutually exclusive choices in a syntax line.
...	Arguments that can be repeated several times.

Where to Find Additional Information

XE "SQL:references for additional information"§For more information about driver development, refer to the *ODBC Driver Developer's Guide*, included in the ODBC Developer's Kit.

For more information about SQL, the following standards are available:

- n Database Language - SQL with Integrity Enhancement, ANSI, 1989 ANSI

X3.135-1989.

- ⌘ X/Open and SQL Access Group SQL CAE draft specification (1991).
- ⌘ Database Language SQL: ANSI X3H2 and ISO/IEC JTC1,SC21,WG3 (draft international standard).

In addition to standards and vendor-specific SQL guides, there are many books that describe SQL, including:

- ⌘ Date, C. J.: *A Guide to the SQL Standard* (Addison-Wesley, 1989).
- ⌘ Emerson, Sandra L., Darnovsky, Marcy, and Bowman, Judith S.: *The Practical SQL Handbook* (Addison-Wesley, 1989).
- ⌘ Groff, James R. and Weinberg, Paul N.: *Using SQL* (Osborne McGraw-Hill, 1990).
- ⌘ Gruber, Martin: *Understanding SQL* (Sybex, 1990).
- ⌘ Hursch, Jack L. and Carolyn J.: *SQL, The Structured Query Language* (TAB Books, 1988).
- ⌘ Pascal, Fabian: *SQL and Relational Basics* (M & T Books, 1990).
- ⌘ Trimble, J. Harvey, Jr. and Chappell, David: *A Visual Introduction to SQL* (Wiley, 1989).
- ⌘ Van der Lans, Rick F.: *Introduction to SQL* (Addison-Wesley, 1988).
- ⌘ Vang, Soren: *SQL and Relational Databases* (Microtrend Books, 1990).
- ⌘ Viescas, John: *Quick Reference Guide to SQL* (Microsoft Corp., 1989).

1 ODBC Theory of Operation

XE "ODBC"§The Open Database Connectivity (ODBC) interface allows applications to access data from database management systems (DBMS).

XE "Interoperability"§The interface permits maximum *interoperability*—a single application can access diverse database management systems. You can develop, compile, and ship an application without targetting a specific DBMS. Users can then add modules called database *drivers* that link your application to their choice of database management systems.

The ODBC interface defines the following:

- A library of ODBC function calls that allow an application to connect to a DBMS, execute SQL statements, and retrieve results.
- SQL syntax. The syntax is based on the X/Open and SQL Access Group SQL CAE draft specification (1991).. To send an SQL statement, you include the statement as an argument in an ODBC function call. You need not customize the statement for a specific DBMS.

Appendix A, "SQL Grammar," contains a SQL syntax based on the X/Open and SQL Access Group SQL CAE draft specification (1991). To ensure maximum interoperability, ODBC applications are encouraged to use only the SQL syntax defined in Appendix A.

- A standardized set of error codes.
- A standard way to connect and log on to a DBMS.
- A standardized representation for data types.

The interface is flexible:

- You can construct SQL statements at compile time or at run time.
- You can use the same object code to access different DBMS products.
- You can connect to multiple instances of DBMS products.
- Your application can ignore underlying data communications protocols between your application and a DBMS product.
- You can send and retrieve data values in a format convenient to the application.

XE "ODBC:function call types"§The ODBC interface provides two types of function calls:

- Core functions based on the X/Open and SQL Access Group Call Level Interface specification.

- Extended functions support additional functionality, including scrollable cursors and asynchronous processing.

An application can check for availability of functions at run-time.

The following code fragment shows one way to use the ODBC interface to submit an SQL CREATE TABLE statement:

```
UCHAR sqltext[] ="CREATE TABLE NAMEID (ID integer, NAME varchar(50))";
if(SQLExecDirect(hstmt, sqltext, SQL_NTS) !=SUCCESS)
return(print_err(hDBC,hstmt));
```

Note that an application could obtain the value of sqltext from a dialog box at run time.

The following sections describe the ODBC architecture in more detail.

ODBC Components

The following paragraphs describe components within the ODBC architecture. This information is provided for reference purposes; the boundaries between the three underlying components are transparent to an application.

XE "ODBC:components"§The ODBC architecture has four components:

- Application Performs processing, possibly on behalf of a user, and calls ODBC functions to submit SQL statements and retrieve results.
- Driver Manager Loads drivers on behalf of an application.
- Driver Processes ODBC function calls, submits SQL requests to a specific data source, and returns results to the application. If necessary, the driver modifies an application's request so that the request conforms to syntax supported by the associated DBMS.
- Data source Consists of the data the user wants to access and its associated operating system, DBMS, and network platform (if any) used to access the DBMS.

The Driver Manager and driver appear to an application as one unit that processes ODBC function calls. The following diagram shows the relationship between the four components. The following paragraphs describe each component in more detail.

Application

XE "Application"§An application that uses the ODBC interface performs the following tasks:

- Requests a connection, or session, with a data source.
- Sends SQL requests to the data source.
- Defines storage areas and data formats for the results of SQL requests.
- Requests results.
- Processes errors.
- Reports results back to a user, if necessary.
- Requests commit or rollback operations for transaction control.
- Terminates the connection to the data source.

An application can provide a variety of features external to the ODBC interface, including spreadsheet capabilities, online transaction processing, and report generation; the application may or may not interact with users.

Driver Manager

XE "Driver Manager"§The Driver Manager, provided by Microsoft, is a dynamically-linked library (DLL) with an import library. The primary purpose of the Driver Manager is to load drivers. (The following section describes drivers.) In addition to loading drivers, the Driver Manager performs the following:

- Uses the ODBC.INI file to map a data source name to a specific driver dynamically-linked library (DLL).
- Processes several initialization-oriented ODBC calls.
- Provides entry points to ODBC functions for each driver.
- Provides parameter validation and sequence validation for ODBC calls.

Driver

XE "Driver"§A driver is a DLL that implements ODBC function calls and interacts with a data source

A driver is loaded by the Driver Manager when the application calls the **SQLConnect** or **SQLDriverConnect** function.

A driver performs the following tasks in response to ODBC function calls from an application:

- Establishes a connection to a data source.

- Submits requests to the data source.
- Translates data to or from other formats, if requested by the application.
- Returns results to the application.
- Formats errors into standard error codes and returns them to the application.
- Declares and manipulates cursors if necessary. (This operation is invisible to the application unless you request access to a cursor name.)
- Initiates transactions if the data source requires explicit transaction initiation. (This operation is invisible to the application.)

Data Source

XE "Database environment"§This manual refers to the general features and functionality provided by an SQL database management system as a DBMS. The manual refers to a specific instance of a combination of a DBMS product, remote operating system, and network as a data source.

An application establishes a connection with a particular vendor's DBMS product on a particular operating system, accessible by a particular network. For example, the application might establish connections to:

- An Oracle DBMS running on an OS2 operating system, accessed by Novell netware.
- A local xBase file, in which case the network and remote operating system are not part of the communication path.
- A Tandem NonStop SQL DBMS running on the Guardian 90 operating system, accessed via a gateway.

Types of Drivers

XE "Driver:types"§ODBC defines two types of drivers:

- Single-tier The driver processes both ODBC calls and SQL statements. (In this case, the driver performs part of the data source functionality.)
- Multiple-tier The driver processes ODBC calls and passes SQL statements to the data source.

One system can contain both types of configurations.

The following paragraphs describe single-tier and multiple-tier configurations in more detail.

Single-Tier Configuration

XE "Configuration:single-tier"§In a single-tier implementation, the database is a file and is processed directly by the driver. The driver processes SQL statements and retrieves information from the database. One example of a single-tier implementation is a driver that manipulates an xBase file.

The set of SQL statements that you can submit may be limited by a single-tier driver. The minimum set of SQL statements that must be supported by a single-tier driver is defined in Appendix A, "SQL Grammar."

The following diagram shows two types of single-tier configurations.

Multiple-Tier Configuration

XE "Configuration:multiple-tier"§In a multiple-tier configuration, the driver sends SQL requests to a server that processes SQL requests.

The application, driver, and Driver Manager reside on one system, typically called the client. The database and the software that controls access to the database typically reside on another system, typically called the server.

One type of multiple-tier configuration is a gateway architecture where the driver passes SQL requests to a gateway process. The gateway process sends the requests to the data source.

The following diagram shows three types of multiple-tier configurations. From the perspective of an application, all three configurations are identical.

Network Example

The following diagram shows how each of the preceding configurations could appear in a single network. The diagram includes examples of the types of DBMS that could reside in a network.

Applications can also communicate across wide-area networks:

Matching an Application to a Driver

§ One of the strengths of the ODBC interface is interoperability: you can create your ODBC application without targeting a specific data source. Users can add drivers to your application after you compile and ship the application.

From an application standpoint, it would be easiest if every driver and data source supported the same set of ODBC function calls and SQL statements. However, drivers and associated data sources provide a varying range of functionality. Therefore, the ODBC interface defines conformance levels, which determine the functions and SQL statements supported by a driver.

ODBC Function Call Support

Each ODBC driver supports a set of core ODBC functions and, optionally, one or more extended functions or data types, defined as extensions:

- Core functions and data types are based on the X/Open and SQL Access Group Call Level Interface specification. If a driver supports all core functions, it is said to conform to X/Open and SQL Access Group core functionality. If any core functions are not supported, the driver does not conform to X/Open and SQL Access Group core functionality. The specific set of core data types supported by a driver depends upon the set of data types supported by the data source.
- Extension functions and data types support additional features, including date, time, and datetime data type literals, scrollable cursors, retrieval of data dictionary information, and asynchronous execution of function calls. Extended functions may not be supported by a specific driver. The **SQLGetFunctions** function returns information about supported functions. **SQLGetFunctions** is always available to applications.

Extended functions are divided into two conformance levels, Level 1 and Level 2, each of which is a superset of the core functions.

For a list of functions and their conformance levels, refer to the *ODBC API Reference*, Chapter 1, "ODBC Function Summary."

§

Note Each function description in this manual indicates whether the function is a core function or a level 1 or level 2 extension.

§

SQL Statement Support

Each ODBC driver supports one of two sets of SQL statements:

- The Minimum set is a set of SQL statements that can be implemented by single-tier drivers.
- The Core set is based on the X/Open and SQL Access Group SQL draft CAE specification (1991).

In addition to the core and minimum sets, ODBC defines SQL syntax for data literals, outer joins, and SQL scalar functions. For more information about SQL statement sets, refer to Appendix A, "SQL Grammar."

The grammar listed in Appendix A is not intended to restrict the set of statements that can be supported. A driver can support additional syntax that is unique to the associated data source.

How to Select a Set of Functionality

The decision to use a set of functionality depends on:

- The set of features you want to access from your application, including features that may not be available from all data sources.
- How much interoperability you would like to provide.
- How much conditional code you want to include to determine whether a function or data type is supported by the driver.
- Performance requirements. The choice of performance options increases from the minimum to the extended set, depending on data source support.

To communicate with a specific driver and data source, select the appropriate functionality set for the driver. If you want additional interoperability, the following tables may help you select a functionality set.

If you want to communicate with	Choose
Single-tier and multiple-tier drivers, with maximum interoperability and the least amount of application work	Minimum functionality. All drivers support core ODBC functions and minimum SQL statements.
Single-tier and multiple-tier drivers, with maximum interoperability and maximum functionality	Check before you issue core or extended functions. If supported, use them. If not, perform equivalent work using minimum functions.
Single-tier drivers	Minimum functionality.
Multiple-tier drivers only, with maximum interoperability and the least amount of application work	Core functionality.
Multiple-tier drivers only, with maximum interoperability, maximum functionality, and maximum performance	Extended functionality. Check functions and, if not available, perform equivalent work using core functions.

Chapters 3 through 9 describe how to develop an application that uses ODBC functions. The *ODBC API Reference* lists all ODBC functions in alphabetic order.

Connections and Transactions

Before your application can communicate with a data source, you must establish a connection. If the connection is successful, the driver returns a connection handle (a pointer to a storage area) for use in subsequent ODBC calls.

The ODBC interface allows you to request multiple connections for one or more data sources. Each connection is considered a separate transaction space.

An active connection can have one or more statement processing streams.

XE "Transactions"§XE "Connection handles:and transactions"§Each active connection maintains a transaction in progress. You can request that the results of each SQL statement be automatically committed upon completion or you can choose to explicitly request commit or rollback operations. When an application commits or rolls back a transaction, the driver resets all statement requests associated with the connection.

2 A Short History of SQL

XE "SQL, overview"§This chapter provides a brief history of SQL and describes programmatic interfaces to SQL. For more information about SQL, refer to the references listed in the introduction.

SQL Background Information

SQL, or Structured Query Language, is a widely accepted industry standard for data definition, data manipulation, data management, access protection, and transaction control. SQL originated from the concept of relational databases and uses tables, indexes, keys, rows, and columns to identify storage locations.

Many types of applications use SQL statements to access data. Examples include ad-hoc query facilities, decision support applications, report generation utilities, and online transaction processing systems.

SQL is not a complete programming language in itself. For example, there are no provisions for flow control.

One of the challenges during the evolution of SQL has been to provide a standard access to SQL database management systems from traditional programming languages like C, COBOL, and PL/1.

ANSI 1989 Standard

XE "ANSI standards"§SQL was first standardized by the American National Standards Institute (ANSI) in 1986. The first ANSI standard defined a language that was independent of any programming language.

XE "Module language"§XE "Embedded SQL (ANSI)"§XE "SQL statements:embedded"§XE "Direct invocation (ANSI)"§The first ANSI standard has since been refined; the current standard is ANSI 1989. The ANSI 1989 standard defines three programmatic interfaces to SQL:

- Module language Allows you to define procedures within compiled programs (modules). You then call these procedures from traditional programming languages. The module language uses parameters to return values to the calling program.
- Embedded SQL Allows you to embed SQL statements within your program. The specification defines embedded statements for COBOL, FORTRAN, Pascal, and PL/1.
- Direct invocation Access is implementation-defined.

Neither the module language nor the direct invocation approach has been widely implemented; most implementations use the embedded approach.

Embedded SQL

Embedded SQL allows you to place SQL statements into a program that is written in a traditional programming language (for example, COBOL or Pascal). You delimit SQL statements with specific starting and ending statements defined by the host language. The resulting program contains source code from two languages—SQL and the host language.

When you compile a program with embedded SQL statements, you use a precompiler to compile the SQL statements. The precompiler replaces the SQL statements with equivalent host language source code. After you precompile the program, you use your host language compiler to compile the resulting source code.

XE "Static SQL"§XE "SQL statements:static"§The term *static SQL* encompasses the basic features of embedded SQL. Static SQL has the following characteristics:

- To use static SQL, you define each SQL statement within the source code of your program. You specify the number of result columns and their data types before you compile your program.
- Variables called *host variables* are accessible to both your host-language code and your SQL requests. You cannot, however, use host variables for column names or table names. In addition, host variables are fully defined (including length and data type) prior to compilation.
- If you submit an SQL request that returns more than one row of data, you define a *cursor* that points to one row of result data at a time.
- Each run of the associated program performs exactly the same SQL request, with possible variety in the values of host variables. All table names and column names must remain the same from one execution of the program to the next; otherwise, you must recompile the program.
- You use standard data storage areas for status and error information.

Static SQL is efficient; you can precompile SQL statements prior to execution and run them multiple times without recompiling the statements. The application is bound to a particular DBMS when it is compiled.

Static SQL cannot defer the definition of the SQL statement until run-time. Therefore, static SQL is not the best option for client-server configurations or for ad-hoc requests.

Future ANSI Specifications

SQL2 is the most recent ANSI specification, and is in the final stages of becoming an international standard. SQL2 defines three levels of functionality: entry, intermediate, and full. SQL2 adds many new features, including:

- Additional data types, including date and time.
- Connections to database environments, to address the needs of client-server architectures.
- Support for dynamic SQL (described in the following subsection).
- Scrollable cursors for access to result sets (full level).
- Outer joins (intermediate and full levels).

Dynamic SQL

Dynamic SQL allows an application to generate and execute SQL statements at run time.

You can prepare dynamic SQL statements. When you prepare a statement, the database environment generates an access plan and a description of the result set. You can then execute the statement multiple times with the previously-generated access plan, which minimizes processing overhead.

You can include parameters in dynamic SQL statements. Parameters function in much the same way as host variables in embedded SQL. Prior to execution, you assign values to the place held by each parameter. Unlike static SQL, parameters do not require length or data type definition prior to program compilation.

Dynamic SQL is not as efficient as static SQL, but is very useful if an application requires:

- Flexibility to construct SQL statements at run time.
- Flexibility to defer an association with a database until run time.

Call Level Interface

XE "Call level interface"§XE "CLI"§A Call Level Interface (CLI) for SQL consists of a library of function calls that support SQL statements. The ODBC interface is a CLI.

The ODBC interface is designed to be used directly by application programmers, and not as the target of a preprocessor for embedded SQL.

A CLI is very straightforward to programmers who are familiar with function libraries. The function call interface does not require host variables or other embedded SQL concepts.

A CLI does not require a precompiler. To submit an SQL request, you place an SQL command into a text buffer and pass the buffer as a parameter in a function call. CLI functions provide declarative capabilities and request management. You obtain error information as you would for any function call—by return code or error function call, depending on the CLI.

XE "Binding result columns"§A CLI allows you to specify result storage before or after the results are available. This allows you to determine what the results are and take appropriate action without being limited to a specific set of data structures that were defined prior to the request. Deferral of storage specification is called late binding of variables.

The concept of a CLI is very useful in a client/server environment; the interface between the application and the data source can be designed to minimize network traffic.

A CLI is typically used for dynamic access because applications that use a CLI are often driven by user input. The CLI defined by the X/Open and SQL Access Group—and therefore the ODBC interface—are similar to the dynamic embedded version of SQL described by in X/Open and SQL Access Group draft specification "Structured Query Language (SQL)" (1991).

For a comparison between embedded SQL statements and the ODBC call level interface, refer to the *ODBC API Reference*, Appendix E, "Comparison Between Embedded SQL and ODBC ."

Interoperability

XE "Interoperability"§Interoperability for call-level interfaces can be addressed in

the following ways:

- All clients and data sources adhere to a standard interface.
- All clients adhere to a standard interface; driver programs interpret the commands for a specific data source.

The second approach allows drivers to shield clients from database functionality differences, database protocol differences, and network differences. ODBC follows the second approach. ODBC can take advantage of standard database protocols and network protocols, but does not require the use of a standard database protocol or network protocol.

3 Guidelines for Calling ODBC Functions

This chapter describes characteristics of ODBC functions and discusses how to perform the following tasks:

- Calling ODBC functions.
- Submitting SQL statements from ODBC functions.
- Handling ODBC results.

Calling ODBC Functions

The following paragraphs describe general characteristics of ODBC functions.

General Information

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

The initialization file must contain driver location information. For additional information about initialization file contents, refer to Chapter 9, “Constructing an ODBC Application.”

C programs that call ODBC functions must use header files that define constants, type definitions, and function prototypes for all ODBC functions. To view the `SQL.H` and `SQLEXT.H` header files, refer to the *ODBC API Reference*, Appendix F, “C Header Files.”

For a list of valid data types, refer to the *ODBC API Reference*, Appendix D, “Data Type Definitions.”

Variable Length Data in Function Arguments

All function arguments that point to variable length data (for example, column names and parameter values) have an associated length argument.

You can specify one of the following lengths for each input argument:

- A length greater than or equal to zero specifies the actual length. A length of zero describes a zero length string, which is distinct from a `NULL` value.

Retrieving Results

- ⌘ A length equal to `SQL_NULL_DATA` specifies a null parameter value.
- ⌘ A length equal to `SQL_NTS` specifies that a value is a null terminated string.

Nulls are always valid for output pointers, unless otherwise noted in the syntax description for a function.

The application is responsible for allocating memory for output buffers. Therefore, the application must indicate the length of each buffer. On output, the driver returns the actual length of data that was stored. For each output argument there are two length arguments:

- An input argument that contains the buffer length as allocated by the application, including one byte for a null termination character that the driver returns for arguments that contain character data.

If a string argument is null for an input parameter, the driver ignores the argument unless it is required for proper operation of the function. If required, as in **SQLPrepare**, the driver returns an error. Nulls are always valid for output pointers, unless noted otherwise in the syntax description for a function.

- An output argument that contains the actual number of bytes written to the buffer by the driver (not including the null termination character), or `SQL_NULL_DATA`, if null.

If the data does not fit in the output buffer, the driver stores the number of bytes available and returns the value `SQL_SUCCESS_WITH_INFO`. If the application calls **SQLError**, the driver returns a truncation error. The application can compare the output length with the buffer size to determine which value was truncated.

When converting a binary (hexadecimal) SQL data type to a character format, the output length will always be an even number of bytes.

There may be instances—for example, if the buffer is very large—when you would like to use a null pointer for the output length, then search for the termination character in the buffer. If you use null terminated strings, you can pass a null pointer for the output length, in which case the driver does not return the length. This is not recommended, however, because the driver cannot return a truncation indicator if you use a null pointer.

For more information about error names and other predefined constants, refer to the header file listed in the *ODBC API Reference*, Appendix F, “C Header Files.”

Environment, Connection, and Statement Handles

XE "Handles:connection and statement"\$XE "Connection handles"\$XE "Statement handles"\$XE To communicate with a data source, an application establishes a connection with the driver. The driver returns handles that reference data structures that store information pertinent to the ODBC environment, a specific connection to an instance of a data source, or a statement being sent to an instance of a data source. These handles are required by most ODBC functions.

the ODBC interface defines three types of handles:

- Environment handles identify memory storage for global information, including valid connection handles and current active connection handle. ODBC defines environment handles as variables of type HENV. An application must request an environment 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 a connection to a specific instance of a data source.
- 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, refer to Chapter 5, “Establishing Connections.” For more information about requesting a statement handle, refer to Chapter 6, “Preparing and Executing an SQL Statement.”

Submitting SQL Statements

XE "SQLPrepare:submitting SQL statements"§XE "SQLExecDirect:submitting SQL statements"§XE "SQL statements:processing"§To submit an SQL statement, you pass it as an argument in an ODBC function call. For more information about submitting SQL statements, refer to Chapter 6, “Preparing and Executing an SQL Statement.”

The application is responsible for submitting correct SQL syntax.

For a description of grammar that is valid in ODBC function calls, refer to Appendix A, “SQL Grammar.” For a comparison between embedded SQL statements and ODBC function calls, refer to the *ODBC API Reference*, Appendix E, “Comparison Between Embedded SQL and ODBC.”

Data Type Support

XE "Data types"§The ODBC interface defines two sets of data types:

- ODBC data types define the SQL data type in the data source. This set is further divided into two subsets:
 - Core data types provide a standard set of data types. If you use the core set of functionality, the driver maps all data to a core data type. For example, the driver stores a date as a character string.

- Extended data types support data types in many current DBMS products.
- C data types describe how data is stored in your C program.

Each ODBC data type has a corresponding C data type. These data types are defined in the ODBC header files. For a list of ODBC data types, their meanings, and how they correspond to C data types, refer to the *ODBC API Reference*, Appendix D, "Data Type Definitions." For information about the header files, refer to the *ODBC API Reference*, Appendix F, "C Header Files."

XE "SQLDescribeCol:retrieving data types"§To retrieve the underlying data type for a column, call **SQLDescribeCol**.

XE "SQLGetTypeInfo:retrieving supported data types"§The ODBC interface provides support for any data type from any data source if you retrieve the type code for the data. If you access extended functions, call **SQLGetTypeInfo** to retrieve a description of data types supported by the data source.

Handling Results

XE "Return codes"§XE "SQL_SUCCESS return code"§XE "SQL_SUCCESS_WITH_INFO return code"§XE "SQL_NO_DATA_FOUND return code"§XE "SQL_ERROR return code"§XE "SQL_INVALID_HANDLE return code"§When you call an ODBC function, the driver returns a predefined status code that indicates success or failure. The status codes indicate success, warning, or failure status. The application can then call **SQLError**, if necessary, to retrieve additional information. The following table lists return constants.

#define name	Description
SQL_SUCCESS	Function completed successfully; no additional information is available.
SQL_SUCCESS_WITH_INFO	Function completed successfully. Call SQLError to retrieve a warning or additional information.
SQL_NO_DATA_FOUND	All rows from the result set have been fetched.
SQL_ERROR	Function failed. Call SQLError for more information.
SQL_INVALID_HANDLE	Function failed due to an invalid connection handle or statement handle. This indicates a programming error. No further information is available from SQLError .
SQL_STILL_EXECUTING	A function was called asynchronously and is still executing.
SQL_NEED_DATA	While processing an SQL statement, the driver determined that the application needs to send large data values.

Handling Errors

XE "SQLError"§If an ODBC function other than **SQLError** returns `SQL_SUCCESS_WITH_INFO` or `SQL_ERROR`, call **SQLError** to obtain additional information. 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 driver buffers errors or messages for only one ODBC call at a time; a subsequent call overwrites existing error information.

SQLError never returns error information about itself.

If you are familiar with `SQLSTATE` in the X/Open and SQL Access Group "Structured Query Language (SQL)" CAE draft specification (1991), note that the information provided by **SQLError** is in the same format as that provided by `SQLSTATE`.

For more information about error codes, refer to the *ODBC API Reference*, Appendix A, "ODBC Error Codes."

Processing Result Sets

You can use the following mechanisms to retrieve information about results:

- The return code.
- A call to **SQLRowCount**.
- A call to **SQLNumResultCols**.
- A call to **SQLDescribeCol**.
- A call to **SQLColAttributes**.
- A call to **SQLNumParams**.

XE "Processing results"§XE "Results:processing"§If the operation does not affect or return rows, such as an SQL **GRANT** or **REVOKE** operation, check the return code to determine the outcome of the operation. If the operation affected rows, obtain the row count to determine the outcome of the operation. If your request was a **SELECT** query, check the number of result columns and data descriptions to gain information about the result set.

For more information about retrieving data, refer to Chapter 7, “Retrieving Results.”

4 Basic Application Steps

To use ODBC functions, perform the following steps:

1. Establish a connection to the appropriate driver. You specify a data source name, defined in the initialization file, and additional connection information.
2. For each SQL request:
 - n Place the SQL text string into a buffer. Your request can include embedded parameter markers. Set parameter values if necessary. Associate a cursor name with the request if desired; otherwise, the driver assigns a cursor name. The driver requests cursor operations automatically.
 - n Submit the SQL string for execution. The ODBC interface provides separate function calls for prepared and immediate execution.
 - n Inquire about the results. For example, for a **SELECT** query you can request information about the columns in the result set. You can then allocate storage and define the data format for each result column.

For unsuccessful requests, process error information.
 - n Fetch data row by row.
3. When you finish a transaction, perform a commit or rollback command if you performed **INSERT**, **UPDATE**, or **DELETE** operations.
4. When you finish submitting statements to the data source, terminate the connection.

The following diagram shows an example of the basic command flow for connecting to a data source, processing SQL statements, and disconnecting from the data source. The words starting with SQL are ODBC function call names.

Depending on the types of requests your application makes, you may decide to use additional ODBC functions.

Additional Information

Chapters 5 through 8 describe how to use ODBC functions that provide these services.

The *ODBC API Reference* lists syntax and usage information for each ODBC function.

5 Establishing Connections

XE "Connections:establishing"§This chapter describes how to establish a connection to a target data source.

Using the Driver Manager

XE "Driver Manager:communicating with"§The Driver Manager is a DLL that provides access to ODBC drivers. You do not need to call an ODBC function to initiate communication with the Driver Manager; the Driver Manager is automatically involved in all ODBC calls.

Whenever you call an ODBC function, the Driver Manager performs one of the following actions:

- For **SQLAllocEnv**, **SQLAllocConnect**, **SQLDataSources**, **SQLFreeConnect**, or **SQLFreeEnv**, the Driver Manager processes the call.
- For **SQLConnect**, **SQLDriverConnect**, **SQLError**, or **SQLGetFunctions**, the Driver Manager performs initial processing then sends 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.

Initializing the ODBC Environment

The **SQLAllocEnv** function initializes the ODBC interface for use by an application. **SQLAllocEnv** must be called prior to any other ODBC function:

1. Declare a variable of type HENV. For example, you could declare a variable called henv1: "HENV henv1." HENV is defined in the SQL.H ODBC header file. For Windows, the HENV type is a memory handle. For more information about the header file, refer to the *ODBC API Reference*, Appendix F, "C Header Files."
2. Pass the address of this variable as an argument in a call to **SQLAllocEnv**. The driver allocates storage for environment information and places the address of the storage into the HENV variable.

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

Establishing a Connection to a Data Source

First, request a connection handle:

1. Declare a variable of type HDBC. For example, you could declare a variable called `hdbc1`: "HDBC `hdbc1`." HDBC is defined in the ODBC header file. For Windows, the HDBC type is a memory handle.
2. Pass the address of this variable as an argument in a call to **SQLAllocConnect**. The driver allocates storage for connection information and places the address of the storage into the HDBC variable.

Next, specify a specific driver and data source. Pass 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. For Windows, this corresponds to an entry in the ODBC initialization file (ODBC.INI). For more information, refer to Chapter 9, "Constructing an ODBC 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).

The Driver Manager establishes a connection with the specified data source and returns connection status to the application.

Accessing ODBC Functions

The ODBC interface defines two types of conformance:

- SAG core conformance, which is met if a driver supports all core functions and false if the driver does not support one or more core functions.
- ODBC conformance, which has two levels:
 - Level 1. The driver supports all core functions plus an additional set of functions that provides a basic level of support and optimization for an interactive query application.
 - Level 2. The driver supports all core functions and all ODBC functions.

If the driver does not support Level 1 or Level 2 functionality, it must return "None" for ODBC conformance, but may support one or more extended functions. An application can call **SQLGetFunctions** to determine if the driver supports a particular function.

The *ODBC API Reference*, Chapter 1, "ODBC Function Summary," lists

conformance levels for all functions. In addition, all function descriptions in the reference manual indicate whether a function is a core function or a level 1 or level 2 extension.

In addition, you can call **SQLGetInfo** to determine the conformance level supported by a driver.

Extensions for Establishing Connections

Several extended functions support the connection process. The following table lists these functions in alphabetic order. The paragraphs following the table describe **SQLDriverConnect** in more detail. XE "SQLDataSources" XE "Data source:listing" XE "Timeout values, setting" XE "Rowcount, setting maximum" XE "Autocommit, setting" XE "Transactions:setting autocommit option" XE "SQLSetConnectOption" XE "SQLSetStmtOption" XE "SQLGetInfo" XE "SQLGetTypeInfo" §

Function Name	Description
SQLDataSources	Requests a list of available data sources. The Driver Manager retrieves this information from the ODBC.INI file. You can present this information to a user or select a database and driver combination from within your application.
SQLDriverConnect	Ask the Driver Manager to present a connection dialog box to the user.
SQLGetFunctions	Returns 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 file names, versions, and the maximum length of names supported by the data source.

Function Name	Description
SQLGetTypeInfo	Determines the data types supported by a driver and data source.
SQLSetStmtOption SQLSetConnectOption SQLGetStmtOption SQLGetConnectOption	Set or retrieve operational parameters for the driver and data source. Options include access mode, timeout values, implicit commit operation, and asynchronous execution of ODBC functions.

SQLDriverConnect

XE "SQLConnect:compared to SQLDriverConnect"\$XE

"SQLDriverConnect"\$The **SQLDriverConnect** function allows you to request that the driver and Driver Manager obtain login information from the user prior to establishing a connection. **SQLDriverConnect** uses a connection string to connect to a driver and data source. This function is useful if the data source requires information that cannot be supplied in the standard **SQLConnect** function.

A connection string contains the following information:

- n Data source name
- n One or more user IDs
- n One or more passwords
- n One or more database-specific parameter values

The connection string is a more flexible interface than the data source name, user ID, and password used by **SQLConnect**. You can use the connection string for multiple levels of logon authorization or to convey other data source-specific connection information.

You can call **SQLDriverConnect** in two ways:

- n Specify a connection string that allows the driver to connect to the data source.
- n Specify a partial connection or no connection string. The Driver Manager displays a dialog that allows the user to select a data source name. The driver displays a login dialog, includes any partial information as default values, obtains necessary connection information from the user, and then establishes the connection.

Once the driver establishes a connection, **SQLDriverConnect** returns a connection string that you can use to call **SQLDriverConnect** again later, if

necessary.

The Driver Manager displays the following dialog if the application calls **SQLDriverConnect** and requests that the user be prompted for information.

Upon request from the application, the driver displays a dialog similar to the following to retrieve login information.

6 Preparing and Executing an SQL Statement

XE "SQL statements:preparing and executing"§Your application can submit the SQL statements listed in Appendix A, "SQL Grammar" (or data source-specific SQL statements) in ODBC function calls. The list in Appendix A is similar to the set of SQL statements that can be prepared in embedded SQL.

§

NoteTo request a commit or rollback operation, call **SQLTransact**.

§

The following diagram shows a sample sequence of ODBC commands that can be used for SQL statement processing. For more information about statement sequencing, refer to the *ODBC API Reference*, Appendix B, "ODBC State Transition Table."

Note that any valid SQL statement can be executed with either the **SQLPrepare** and **SQLExecute** sequence or the **SQLExecDirect** command, depending on whether the you plan to submit the SQL statement once or more than once. This functionality differs from embedded SQL, since both statements with and without cursors are executed the same way.

Note also that there are other valid calling sequences that include functions such as **SQLBindCol** and **SQLGetData**.

Sample Flow Control

This chapter describes ODBC functions that support SQL statement processing. For information about results, refer to Chapter 7, "Retrieving Results."

Allocating a Statement Handle

XE "Statement handles:requesting"\$XE "Handles:requesting statement"\$Allocate a statement handle as follows:

1. Declare a variable of type **HSTMT**. For example, you could declare “HSTMT hstmt1.” **HSTMT** is defined in the SQL.H file, listed in the *ODBC API Reference*, Appendix F, “C Header Files.” For Windows applications, the hstmt variable is a memory handle.
2. Pass the address of this variable and an existing connection handle in a call to **SQLAllocStmt**.

The driver allocates a statement handle and returns the handle to your application.

Assigning Storage for Results (Binding)

XE "Binding result columns"\$XE "Results:binding"\$XE "Results:assigning storage for"\$You can assign storage for results before or after you execute an SQL statement. XE "SQLBindCol"\$To allocate storage for a column of data, call **SQLBindCol** and include the following information:

- n Decide whether you want the driver to convert the results to a different data type. If so, include this request in your call to **SQLBindCol**.
- n Define a storage buffer for the data. This storage area must be large enough to hold the maximum number of bytes needed for the column for the specified data type.
- n Define a storage buffer for the data length.

Choosing Prepared or Direct Execution

XE "Prepared execution"\$XE "Direct execution"\$XE "SQL statements:executing"\$You have two execution options when you send an SQL request:

- n Prepared Use this option if you want to execute the same statement more than once without respecifying the SQL string or if you need information about the result set prior to execution.

- Direct Use this option if you want to submit the statement once and you do not need the result format prior to execution.

These two options differ from the prepared and immediate options in embedded SQL. For a comparison between ODBC functions and embedded SQL, refer to the *ODBC API Reference*, Appendix E, "Comparison Between Embedded SQL and ODBC."

Prepared Execution

If you plan to submit the SQL statement multiple times, possibly with intermediate changes to parameter values, prepare the request as follows.

```
XE "SQLPrepare"§XE "SQLSetParam"§XE "UPDATE (SQL
statement):positioned"§XE "DELETE (SQL statement):positioned"§XE "SQL
statements:positioned UPDATE and DELETE"§XE "SQLGetCursorName"§XE
"SQLSetCursorName"§Issue the following two calls in any order after
SQLAllocStmt and prior to SQLExecute:
```

- Call **SQLPrepare** and pass the SQL statement as an argument.
- If your SQL statement includes parameter markers, call **SQLSetParam** to associate storage areas with corresponding parameter markers. If this is not the first execution of an SQL statement, you can reuse previous storage areas.

If your request requires a cursor name, as in positioned update or delete (**UPDATE WHERE CURRENT OF** *cursor-name* or **DELETE WHERE CURRENT OF** *cursor-name*), you can allow the driver to generate the cursor name or you can call **SQLSetCursorName** to associate a cursor name with your prepared request.

After setting the cursor name with **SQLSetCursorName** or implicitly obtaining a cursor name by executing a **SELECT** statement, you can call **SQLGetCursorName** to retrieve the cursor name. (The following subsection, “Performing Positioned Updates and Deletes,” contains additional information about positioned updates and deletes.)

Set parameter values for all parameter markers, and then call **SQLExecute** to submit the request.

XE "Declare Cursor (SQL statement)"§If you do not set new parameter values prior to a subsequent call to **SQLExecute**, the driver reuses existing parameter values.

XE "Prepared SQL statements, advantages"§The prepare and execute approach provides the following advantages:

- If the data source supports statement preparation, this is the most efficient way to perform multiple iterations of the same request, especially for complex SQL statements. The data source minimizes processing time by compiling the SQL statements once, producing an access plan, then using the plan for each execution of the request. An access plan identifier allows the driver to send a tag instead of the full SQL statement for subsequent requests, thus minimizing network traffic on subsequent executions of a statement.
- You can retrieve information about the format of the result set prior to executing the SQL statement.

Direct Execution

XE "SQLExecDirect"§If you do not require information about the result set prior to completion of your SQL request and you plan to submit a statement only once, you can call **SQLExecDirect** to submit the SQL statement.

Processing Positioned Updates and Deletes

A positioned update or positioned delete performs an update or delete operation, respectively, based on cursor position.

XE "UPDATE, positioned"§XE "DELETE, positioned"§XE "SQL statements:positioned UPDATE and DELETE"§After you submit a **SELECT** statement that returns multiple rows and you fetch one or more result rows, you can perform a positioned **UPDATE** or **DELETE** to update or delete the row

referenced by the cursor. To request a positioned update or delete operation, use the following SQL syntax:

```
UPDATE {tablename | viewname} SET {columnname = expression} WHERE  
CURRENT OF cursorname
```

```
DELETE FROM {tablename | viewname } WHERE CURRENT OF  
cursorname
```

Next, execute the statement:

- n Prepared execution Call **SQLPrepare** with *szSqlStr* set to the text of the positioned **UPDATE** or **DELETE** statement. Use a different *hstmt* than that used for the **SELECT** statement. Include the cursor name associated with the earlier **SELECT** statement. Set parameter values as necessary, and then call **SQLExecute** to submit the statement.
- n Direct execution Call **SQLExecDirect** with *szSqlStr* set to the text of the **UPDATE** or **DELETE** statement. Use a different *hstmt* than that used for the **SELECT** statement. Include the cursor name associated with the earlier **SELECT** statement.

The driver associates the new *hstmt* with the existing *hstmt* , includes parameter values if necessary, submits the positioned **szSqlStr** or **DELETE** statement, and returns results to the application.

Example

The following diagram lists sample calling sequences for prepared and direct positioned update or delete operations. This sequence is an example; you could include calls to **SQLSetParam** in either sequence or combine prepared and direct requests in the same processing stream.

Extensions for Processing SQL Statements

The following table lists related extended functions. The paragraphs following this table describe the following topics in more detail: retrieving data dictionary information; sending arrays of parameters; sending large data values; using scrollable cursors; requesting asynchronous processing; and requesting scalar functions, extended data types, and outer joins.

To determine whether a particular driver supports these functions, call **SQLGetFunctions**.

Function or Operation	Description
Data dictionary functions	Return data dictionary information. These functions are useful if the data source does not support SQL system views.
SQLDescribeParam	Return information about prepared parameters.
SQLGetData	Return one column of one row of data to the application. SQLGetData is useful for returning large data values.
Sending large data values	Allow the application to send large data values to the data source.
SQLParamOptions	Allow the application to specify multiple sets of parameter values for a single SQL statement. This capability, if supported by the data source, minimizes network traffic..
SQLSetScrollOptions	Establish a scrollable cursor for the result set.
Asynchronous processing	Allow the application to request asynchronous processing of a subset of ODBC functions.

Scalar functions	Support the use of scalar functions in SQL statements.
Extended data types	Support embedded extended data type literals in SQL statements.
Outer joins	Support outer join requests in SQL statements.

Obtaining Information about the Data

The following functions return information about data:

- **SQLColumnPrivileges** returns a list of columns and associated privileges for one or more tables
- **SQLColumns** returns the list of columns names in a specified table
- **SQLForeignKeys** returns a list of column names that compose foreign keys for a specified table
- **SQLPrimaryKeys** returns the column name (or names) that comprise the primary key for a table
- **SQLSpecialColumns** returns information about the optimal set of columns that uniquely identifies a row in a table or the columns that are automatically updated when any value in the row is updated by a transaction
- **SQLStatistics** returns a list of statistics about a single table and the indexes associated with the table
- **SQLTablePrivileges** returns privileges associated with one or more tables
- **SQLTables** returns the list of table names stored in a specific data source

Each function returns the information as a result set. An application fetches these results in the same manner as it retrieves query results (through a call to **SQLFetch**).

Sending Large Data Values

XE "Large data values:sending"\$To send large data values, use the following three functions:

- **SQLSetParam**

- n **SQLParamData**
- n **SQLPutData**
- n **SQLParamOptions**

To indicate that you plan to send a large data value, call **SQLSetParam** to associate storage with the parameter—and set `pcbValue` to `SQL_LONG_DATA` for the parameter.

In the call to **SQLSetParam**, set `rgbValue` to a value that, at run time, references the location of the data. The driver returns this value to the application at statement execution time.

When the driver processes a call to **SQLExecute** or **SQLExecDirect**, the driver returns `SQL_NEED_DATA` as soon as it encounters a parameter that requires a large data value. The application then calls **SQLParamData** and **SQLPutData** to send data values:

- n **SQLParamData** searches for the next large data value parameter and returns the value referenced by `rgbValue` (in the earlier call to **SQLSetParam**).
- n **SQLPutData** transports the actual data value to the data source.

For additional information, refer to the description of **SQLSetParam** in Chapter 2 of the *ODBC API Reference*.

Specifying Arrays of Parameter Values

XE "Parameters:specifying arrays of"§XE "SQLParamOptions"§To specify multiple sets of parameter values for a single SQL statement, call **SQLParamOptions**. For example, if you have ten sets of column values to insert into a table—and you can use the same SQL statement for all ten operations—you can set up an array of values, then submit a single INSERT statement.

If you use **SQLParamOptions**, your application must allocate enough memory to handle the arrays of values.

Using Scrollable Cursors

SQL was originally designed to return one row at a time to an application. Scrollable cursors provide more flexible access to blocks of result data. XE "Scrollable cursors, overview"§The following paragraphs provide an overview of scrollable cursors and describe ODBC features that support scrollable cursors.

Basic Cursors

An SQL **SELECT** statement extracts data that meets a set of specifications. For example, `SELECT * FROM EMPLOYEE WHERE EMPNAME = "JONES"` returns all columns of all rows in `EMPLOYEE` where the employee's name is Jones. This set of information, called a result set, can contain zero, one, or more than one row.

Applications retrieve single rows as follows:

- A cursor, managed by the driver, points to the current row in the result set.
- A call to **SQLFetch** moves the cursor to the next row in the result set and retrieves the row.

This basic form of a cursor is called a forward-only scrolling cursor, and is supported by core ODBC functions. To fetch a previous row using a forward-only cursor, the driver closes the cursor, reopens the cursor for the same result set, and fetches rows until it retrieves the target row.

Scrollable Cursors

Scrollable cursors allow a user to scan results in a flexible manner without excessive support from the application. Users can view rows within a block of data and update, delete, refresh, or browse through the data. Scrollable cursors use the following concepts:

- A block of data is called a rowset.
- A set of keys that uniquely identifies the rows in a rowset is called a keyset. If a table does not contain unique key fields, the keyset may be the whole row.

XE "Concurrency control, setting"§As the size of a rowset increases, so does the possibility that another user may want to access or update one of the rows. You can request four types of locking for a keyset:

- Read only Read the data and build a keyset, but do not lock the data. This approach does not guarantee that the key will point to the same row at a later time.
- Locked Read the data with a lock. Other users cannot modify the data until you remove the lock. This approach guarantees that data is the same when a subsequent update or delete is performed.
- Optimistic concurrency control comparing timestamps Do not lock the data.

Instead, store the time the row was last modified (if available). If the user requests a positioned update or delete operation, check the timestamps to make sure the row was not modified since the keyset was built.

- Optimistic concurrency control comparing values Do not lock the data. Instead, include all row data values in the keyset. If the user requests a positioned update or delete operation, compare these values to values in the database to make sure the row was not modified since the keyset was built.

Call **SQLSetScrollOptions** to specify rowset, keyset, and concurrency control. If the application uses scrollable cursors, the application must call **SQLSetScrollOptions** before it calls **SQLPrepare** or **SQLExecDirect**.

SQLBindCol binds storage areas for result columns.

XE "SQLExtendedFetch"§To fetch a block of data, call **SQLExtendedFetch**. Following the extended fetch, the cursor points to the entire rowset for subsequent positioned operations. To position the cursor on specific rows within the rowset, call **SQLSetPos**.

For scrollable cursor operations, the rowset behaves as a single fat cursor.XE "SQLSetPos"§

Regardless of position, a fetch next or fetch previous operation moves the entire rowset as if it were one cursor.

The **SQLExtendedFetch** function supports forward, backwards, and arbitrary retrieval of blocks, so that there are two levels of movement within a result set: at the rowset level and at the row level.

Requesting Asynchronous Processing

XE "SQL statements:processing asynchronously"§By default, a driver processes ODBC functions synchronously; the driver does not return control to your application until a function call completes. If a driver supports asynchronous processing, however, you can request asynchronous processing for the functions listed below.

SQLColAttributes	SQLExecDirect	SQLGetTypeInfo	SQLSetPos
SQLColumns	SQLExecute	SQLMoreResults	SQLSpecialColumns
SQLColumnPrivileges	SQLExtendedFetch	SQLNumResultCols	SQLStatistics

SQLConnect	SQLFetch	SQLParamData	SQLTablePrivileges
SQLDescribeCol	SQLForeignKey s	SQLPrepare	SQLTables
SQLDescribeParam	SQLGetData	SQLPrimaryKeys	
SQLDriverConnect	SQLGetInfo	SQLPutData	

All of these functions either submit requests to a data source or retrieve data. Any of these functions can initiate extensive processing.

To enable or disable asynchronous processing for the above set of functions, call **SQLSetStmtOption** and specify ON or OFF for the `SQL_ASYNC_ENABLE` option. To check the setting of the `SQL_ASYNC_ENABLE` option, call **SQLGetStmtOption**. To enable or disable asynchronous processing for all `hstmts` associated with an `hdbc`, call **SQLSetConnectOption**.

Upon successful initiation of an asynchronously-initiated function, the driver returns `SQL_STILL_EXECUTING`. If you resubmit the function call, the driver returns `SQL_STILL_EXECUTING` until the function completes. Once the function completes, the driver returns a standard return code.

Using Extended Data Types, Functions, and Outer Joins

ODBC supports an escape mechanism that allows an application to embed database-specific data types, scalar functions, or outer join specifications within an SQL statement. XE "Escape sequence:data types"§The escape mechanism follows the format for vendor-specific commands as defined in Appendix A, "SQL Grammar." In addition, ODBC defines canonical forms for the following:

- n Date, time, and timestamp data types
- n Scalar functions such as string and numeric functions
- n Data type conversion

n Outer join request

The application can use one of two forms when specifying a data type or scalar function:

- n The canonical form as defined by ODBC. This approach provides database independence. The driver translates the canonical form into the database-specific form.
- n The form required by the data source. This approach does not provide database independence.

§

Note An application can submit literal data, without using an escape sequence, as long as it calls **SQLGetTypeInfo** to make sure the data type is supported and to obtain appropriate prefix and suffix information.

§

The following paragraphs list the syntax for each type of escape clause. In addition, ODBC defines a shorthand syntax for escape clauses, defined in Appendix A, "SQL Grammar."

Date and Time Data Types

The following escape clause allows an application to specify a date, time, or timestamp data type:

--*(vendor(Microsoft),product(ODBC), {d|t|ts} value --*)

The following table describes each element of the preceding clause.

Argument	Description
vendor(Microsoft), product(ODBC)	Identifies the vendor and product that support the escape clause. Vendor and product are not used with the shorthand notation.
d	Indicates value is in date format (yyyy-mm-dd).
t	Indicates value is in time format

(hh:mm:ss).

<i>ts</i>	Indicates value is in timestamp format (yyyy-mm-dd hh:mm:ss.[ffffff]), where the precision represented by fraction of a second ([ffffff]) depends on the data source).
<i>value</i>	The value of the date, time, or timestamp variable.

You can specify a date, time, or timestamp escape clause in place of a literal in an SQL statement. For detailed syntax information, refer to Appendix A, "SQL Grammar."

To submit a data type in the native form of the data source, submit the SQL statement with vendor-specific syntax.

Call **SQLGetFunctions** to determine if a driver supports a specific extended data type.

For a list of core and extended data types, refer to the *ODBC API Reference*, Appendix D, "Data Type Definitions."

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. ODBC supports a set of canonical scalar functions that may be a subset or superset of functions actually supported by a given DBMS. If the application specifies the canonical form of a scalar function, the driver translates the function to the syntax required by the data source. If the application specifies the native form of a scalar function, the driver does not translate the function, but sends it to the data source in the form specified by the application. In either case, the data source makes the final determination of the validity of the scalar function and its arguments.

The following escape clause allows an application to specify a scalar function:

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

The following table describes each element of the preceding clause.

Argument	Description
vendor(Microsoft), product(ODBC)	Identifies the vendor and product that support the escape clause. Vendor and product are not used with the shorthand notation.
fn	Indicates that the escape sequence requests a scalar function.
<i>function</i>	Is an expression that contains one or more function names and function arguments. For a list of canonical functions, refer to the <i>ODBC API Reference</i> , Appendix G, "Canonical Functions."

You can use an escape clause with a canonical function in place of a scalar function in an SQL statement. For information about the syntax of canonical functions, refer to Appendix A, "SQL Grammar," and the *ODBC API Reference*, Appendix G, "Canonical Functions."

An application can also include native DBMS functions in an escape clause. The following example shows how an application would specify a native function ("round," in this example) and two ODBC canonical functions (ABS and SQRT). This example uses shorthand escape clause syntax, as defined in Appendix A,

"SQL Grammar," and accesses three columns—EMPNO, EMPNAME, and EMPDIST—in a table called EMPLOYEE.

```
SELECT EMPNO, {fn abs( round ( {fn sqrt(EMPNAME)} ))}, EMPDIST  
FROM EMPLOYEE
```

An application can call **SQLGetInfo** to determine which functions are supported by a specific driver and associated data source.

Data Type Conversion Function

ODBC defines a special type of canonical function, called the CONVERT function, that allows applications to explicitly request a data type conversion when the database processes the SQL statement. The driver translates the canonical form into the database-specific form.

The canonical form of the CONVERT function does not restrict the range of data type conversions. Instead, each driver determines the valid set of conversions. Call **SQLGetInfo** to determine which conversions are supported by the data source and its associated driver.

The following example shows how an application would specify a conversion to a character data type:

```
SELECT EMPNO FROM EMPLOYEE WHERE
--
*(vendor(Microsoft),product(ODBC),fn(CONVERT(EMPNO,SQL_CHAR))--
*) LIKE '1%'
```

For more information, refer to the *ODBC API Reference*, Appendix G, "Canonical Functions."

Requesting Outer Joins

XE "Escape sequence:outer joins"§The ODBC interface uses an escape sequence to support outer joins. The format is:

```
--*(vendor(Microsoft),product(ODBC) oj join-type --*)
```

Join-type specifies a type of outer join. For example, to perform an outer join between two tables named EMPLOYEE and DEPT, using the DEPT.ID column for the join condition, use the following statement (shown in shorthand escape clause syntax):

```
SELECT employee.name, dept.name FROM
{oj employee LEFT OUTER JOIN dept ON employee.deptid=dept.deptid}
WHERE employee.projid=544
```

This syntax is a subset of ANSI SQL2 outer join syntax.

If you use an escape clause with a outer join request in an SQL statement, the escape clause must appear after the FROM clause and before the WHERE clause, if either exist in the statement. For detailed syntax information, refer to Appendix A, "SQL Grammar."

Call **SQLGetFunctions** to determine if a driver supports a specific join type and, if so, what join characters to use in *value*.

7 Retrieving Results

XE "Results:processing"§This chapter describes result-handling operations.

The steps you take to process results depend on your knowledge about the result set.

- Known result set You know the exact form of your SQL statement prior to execution. For example, the query **SELECT EMPNO, EMPNAME FROM EMPLOYEE** returns two specific columns. The statement **DELETE FROM EMPLOYEE WHERE EMPNAME = "Harry Jones"** returns a row count.
- Variably structured result set You do not know the exact form of your results at compile time. For example, the ad-hoc query **SELECT * FROM EMPLOYEE** returns all currently defined columns in the EMPLOYEE table. You may not be able to predict the format of these results prior to execution.

Determining Characteristics of a Result Set

XE "SQLNumResultCols:determining result characteristics"§XE

"SQLRowCount:determining result characteristics"§XE

"SQLDescribeCol:determining result characteristics"§XE "Results:determining characteristics"§To determine the characteristics of a result set, perform these steps:

- 1 Call **SQLNumResultCols** to see if your request returned a set of rows.
- 2 If **SQLNumResultCols** returns a nonzero result, your request was a **SELECT** statement. Call **SQLDescribeCol** and **SQLBindCol** for each column, and then call **SQLFetch** to retrieve each row.

If **SQLNumResultCols** returns a zero, call **SQLRowCount** to see if your request inserted, deleted, or modified rows:

- If **SQLRowCount** returns a nonzero result, that result is the number of rows modified.
- If **SQLRowCount** returns a zero, the statement performed other work: commit, rollback, or other operation that does not modify or return rows.

§

NoteIf you prefer to check for row modification first, call **SQLRowCount** before you call **SQLNumResultCols** .

§

Fetching Result Data

XE "Results:retrieving data"§If you requested an SQL statement that does not manipulate data (for example, **GRANT** or **REVOKE**), the return code for the ODBC function call indicates whether your request was successful or not.

If you requested an **INSERT**, **UPDATE**, or **DELETE** operation, **SQLRowCount** returns the number of rows affected by the operation.

XE "SQL statements:processing results"§XE "SELECT (SQL statement):storing results"§XE "SQLFetch:retrieving result rows"§If you requested a **SELECT** statement and you did not bind the columns to storage locations prior to execution, call **SQLBindCol** to perform bind operations prior to retrieving results. Call **SQLFetch** to retrieve each row of results.

The following diagram shows the flow of the preceding two operations:

Retrieving Error and Status Information

XE "SQLException"§XE "Errors, retrieving"§If your ODBC call returns `SQL_ERROR` or `SQL_SUCCESS_WITH_INFO`, call **SQLException** to find out more information. **SQLException** returns standardized messages and data source-specific messages. The application is responsible for reading these messages and taking appropriate action.

For information about how to handle standard error and status messages, refer to “Calling ODBC Functions” in Chapter 3. For more information about ODBC error codes, refer to the *ODBC API Reference*, Appendix A, “ODBC Error Codes.”

Related Extension Functions

XE "SQLNativeSql"§XE "SQLDescribeParam"§XE "Escape sequence:translating to SQL text"§XE "Parameters:obtaining data type information"§The following table lists related extended functions. The paragraphs following this table contains additional information about **SQLGetData**, **SQLExtendedFetch**, and retrieving multiple result sets.

Function Name	Description
SQLNativeSql	Retrieve the SQL statement as processed by the data source, with escape sequences translated to native SQL code.
SQLDescribeParam	Retrieve data type information for parameters in an SQL statement.
SQLParamOptions	Specify multiple values for a set of parameters.
SQLGetData	Fetch one column in a result row. This function is useful for retrieving large data values.
SQLExtendedFetch	Fetch a block of result data.
SQLSetPos	Position the cursor within a fetched block of data.
SQLScrollOptions	Request a scrollable cursor.

Retrieving Data in Large Columns

To retrieve data from large columns—or retrieve data one column at a time—call

SQLGetData to retrieve the column into a buffer. If you use **SQLGetData** to retrieve data, do not call **SQLBindCol** for the column.

SQLFetch does not retrieve data for columns that are not bound, but does move the cursor from row to row and retrieve data in bound columns. Prior to calling **SQLGetData**, call **SQLFetch** to position the cursor at the first row. Call **SQLFetch** to move the cursor to subsequent rows.

You can combine **SQLGetData** with **SQLBindCol** and **SQLFetch** access within a row. To combine these operations, bind as many contiguous columns as necessary. Call **SQLFetch** to retrieve all bound rows and leave the pointer at the end of the bound rows. Call **SQLGetData** to retrieve the unbound columns. (Note, however, that you cannot use **SQLGetData** to retrieve a column to the left of the rightmost bound column.)

You can call **SQLGetData** multiple times for a single column, if necessary.

Processing Multiple Result Sets

Call **SQLMoreResults** function to process multiple result sets. XE "Stored procedures, processing results"§XE "Batched SQL statements, results"§XE "SQL statements:batched processing"§The following SQL statements can return multiple result sets:

- Batched SQL statements You submit more than one **SELECT** statement in one request.
- Statements with arrays of parameters You submit more than one set of parameter values for a SQL statement (**SQLParamOptions**).
- Stored procedures You submit a stored procedure name as an SQL statement; such a procedure can return multiple result sets.

Processing Blocks of Results

If the driver supports scrollable cursors, your application can retrieve blocks of result rows. For more information about scrollable cursors and processing blocks of result data, refer to "Using Scrollable Cursors," in Chapter 6.

8 Terminating Transactions and Connections

XE "Transactions:terminating"§XE "Handles:releasing statement"§XE "Statement handles:releasing"§XE "Statements, terminating"§The ODBC interface provides functions that terminate SQL transactions, statement-processing connections (hstmts), connections (hdbc), and environment connections (henvs).

Terminating Statement Processing

Call **SQLFreeStmt** to free resources associated with a statement handle. The **SQLFreeStmt** function has four options:

- **SQL_CLOSE** Close the cursor, if one exists, and discard pending results. You can use the statement handle again later.
- **SQL_DROP** Close the cursor, if one exists, discard pending results, and free all resources associated with the statement handle.
- **SQL_UNBIND** Release all return buffers bound by **SQLBindCol** for the statement handle.
- **SQL_RESET_PARAMS** Release all parameter buffers requested by **SQLSetParam** for the statement handle.

Terminating Transactions

XE "SQLTransact:terminating a transaction"§At the transaction level, call **SQLTransact** to commit or roll back the current transaction.

XE "SQLCancel:terminating statement processing"§XE "SQLError:checking statement cancellation"§To cancel a specific statement, call **SQLCancel**, check the return code of the function that was canceled, and then call **SQLError** to determine the status of the cancellation. For additional information about **SQLCancel**, refer to the *ODBC API Reference*, Chapter 2, “ODBC API Reference.”

Terminating Connections

XE "Connections:terminating"§XE "SQLFreeConnect:releasing a connection handle"§XE "SQLDisconnect:closing a connection"§To terminate a connection to a driver and data source, perform the following steps:

1. Call **SQLDisconnect** to close the connection. You can then use the handle to reconnect to the same data source or to a different data source.

2. Call **SQLFreeConnect** to release the connection handle and free all resources associated with the handle.
3. Call **SQLFreeEnv** to release the environment handle and free all resources associated with the handle.

9 Constructing an ODBC Application

XE "Developer's Kit, contents"§The following paragraphs describe how to construct an application, starting with the contents of the developer's kit and including sample code, setup information, and references to testing and debugging information. This manual assumes that you are developing an application for the Windows environment.

Developer's Kit Contents

Your development kit should contain the following items:

- n SDK cover letter and license agreement
- n SDK installation instructions
- n *ODBC Application Programmer's Guide* (this manual), *ODBC Driver Developer's Guide*, *ODBC API Reference*, *ODBC Test Application User's Guide*, a binder, and tabbed dividers
- n One 3.5-inch, 1.44 MB disk and one 5.25-inch, 1.2 MB disk, both containing:
 - n Installation batch file and utilities (including INSTALL.BAT, INSTALL2.BAT, and CHKDIR.EXE)
 - n Driver Manager (DRVRMGR.DLL)
 - n Driver Manager Import Library (DRVRMGR.LIB)
 - n Test application (GATOR.EXE)
 - n Sample Driver DLL (SAMPLE.DLL)
 - n ODBC core functions header file (SQL.H)
 - n ODBC extension functions header file (SQLEXT.H)
 - n ODBC initialization file (ODBC.INI)
 - n Test application initialization file (GATOR.INI)
 - n Common dialogs DLL (COMMDLG.DLL)
 - n Test application source files, header files, object files, makefile, and other files
 - n Sample driver source files, header files, make file, and other files
 - n Other common header files
 - n Other common libraries

System Requirements

The following paragraphs list hardware, software, and environmental requirements for the ODBC environment.

Hardware Requirements

ODBC requires approximately two megabytes of disk space for installation of SDK files and for assembling, compiling and linking the test application and sample driver.

The SDK software has been tested on the following hardware, although it may be possible to use other configurations:

- Personal computer with an 80386 processor and at least five megabytes of RAM.

Software Requirements

The SDK software has been tested with the following system and development software, although it may be possible to use other configurations:

- MS-DOS 5.0
- Microsoft Windows 3.0 or Windows 3.1 Beta-3
- Microsoft C6.00A
- Microsoft Windows 3.0 SDK or Windows 3.1 Beta-3 SDK
- Microsoft MASM 5.3

Environmental Requirements

When using the supplied makefiles to build customized versions of the test application (gator.exe) or the sample driver (sample.dll), you must be sure that the \INCLUDE and \LIB directories in your ODBC SDK directory are specified in your PATH or environment variables as specified in the *Microsoft C Compiler Reference Manual*. These two directories should be searched first during the compile and link process. You must also make sure that the source, object and other files for the test application or sample driver can be found by the assembler, compiler and linker. Review the makefiles and any necessary assembler, compiler and linker documentation to ensure that you have your environment correctly defined for your configuration.

Installing the Developer's Kit

XE "Installing ODBC"§XE "ODBC:installing"§To install the software for the SDK, refer to the SDK installation instructions.

Upon successful completion of the installation process, your Windows directory should contain the ODBC initialization file (ODBC.INI). The directory that you specified for the installation of the rest of the SDK files should contain the following files (the default ODBC directory will be used for purposes of illustration):

Directory	File	File Description
C:\WINDOWS	ODBC.INI	The ODBC initialization file. Includes a data source specification for the SAMPLE.DLL and the sample driver.
C:\ODBC	GATOR.EXE	ODBC GATOR; driver test application.
	DRVRMGR.DLL	ODBC Driver Manager DLL.
	COMMDLG.DLL	Common Dialogs DLL.
	SAMPLE.DLL	ODBC Sample Driver DLL.
	GATOR.INI	ODBC GATOR initialization file.
C:\ODBC\GATOR	AUTOTEST.C	Source for selecting and initiating auto-tests. For example, CORQTEST.C is the source for an autotest and is initiated from here. Custom autotests can be hooked into GATOR via this file.
	AUTOTEST.H	Header file for AUTOTEST.C

CORQTEST.C	Source for the core conformance quick test autotest.
CUSTOM.C	An autotest prototype that can be used as the basis for creating customized auto-tests.
GATOR.DEF	Windows module definition file for GATOR.
GATOR.H	GATOR header file.
GATOR.RES	Windows resource file for GATOR.
L1QTEST.C	Source for the ODBC conformance level 1 quick test autotest.
L2QTEST.C	Source for the ODBC conformance level 2 quick test autotest.
MAKEFILE	GATOR makefile; a working model.
SQL_0001.AC T	Results from the query defined by the file SQL_0001.QRY.
SQL_0001.QR Y	Example of a query file that can be run against the sample driver.

Directory

File

File Description

C:\ODBC\GATOR
(continued)

SQL_0001.RST

The comparison file for the results of SQL_0001.QRY.

	TESTS.H	Contains references to global variables used in GATOR.
C:\ODBC\GATOR\OBJ	AUTO.OBJ	The object file that supports the "Auto" menu selections in GATOR.
	CONNECT.OBJ	The object file that supports the "Connect" menu selection in GATOR.
	DDA.OBJ	The object file that supports the "Dictionary" menu selection in GATOR.
	GATOR.OBJ	The object file for the basic GATOR application.
	MFQENG.OBJ	The object file that supports the execution of query files (.QRY) from the "Auto-Tests..." menu selection in GATOR.
	MISC.OBJ	The object file that supports the "Misc" and "Options" menu selections in GATOR.
	RECEIVE.OBJ	The object file that supports the "Results" menu selections in GATOR.
	SEND.OBJ	The object file that supports the "Statements" menu selections in GATOR.
	SERVERCN.OBJ	The object file that supports the "tools" section of the the "Connect" menu in GATOR.

C:\ODBC\INCLUDE	SQL.H	ODBC core functions header file.
	SQLEXT.H	ODBC extension functions header file.
	WINDOWS.H	Microsoft Windows header file.
C:\ODBC\LIB	COMMDLG.LIB	Common dialogs import library.
	DRVRMGR.LIB	ODBC Driver Manager import library.
	LIBW.LIB	Microsoft Windows library.
	MDLLCEW.LIB	Microsoft Windows library.
	MLIBCEW.LIB	Microsoft Windows library.
C:\ODBC\SAMPLE	CONNECT.C	Sample driver implementation of: SQLAllocEnv, SQLFreeEnv, SQLAllocConnect, SQLConnect, SQLDisconnect and SQLFreeConnect.
	DATABASE.C	Sample driver "database" functions..
	DATABASE.H	Header file for DATABASE.C.
Directory	File	File Description
C:\ODBC\SAMPLE	ECONNECT.C	Sample driver implementation of:

(continued)

SQLDriverConnect and
SQLGetInfo.

EDATA.C	Sample driver implementation of: SQLGetTypeInfo .
EDICT.C	Sample driver implementation of: SQLTables and SQLColumns .
EMISC.C	Sample driver implementation of: SQLSetConnectOption , SQLSetStmtOption , SQLGetConnectOption and SQLGetStmtOption .
ERESULTS.C	Sample driver implementation of: SQLGetData
ERR.H	Header file mapping error code constants for SQLSTATE values to error values.
EXECUTE.C	Sample driver implementation of: SQLAllocStmt , SQLPrepare , SQLSetParam , SQLSetCursorName , SQLGetCursorName , SQLExecute , and SQLExecDirect .
LIBSTART.ASM	Entry point startup routine for Windows sharable code libraries.
MAKEFILE	Sample driver makefile; a working model.
MEMORY.C	Memory and resource management cover functions.

MISC.C	Sample driver implementation of: SQLError and SQLCancel .
RESULTS.C	Sample driver implementation of: SQLNumResultsCols , SQLDescribeCol , SQLBindCol , SQLFetch , SQLRowCount and SQLFreeStmt
SAMPLE.DEF	Windows module definition file for the sample driver.
SAMPLE.H	Overall header file for the sample driver.
SAMPLE.RC	Windows resource file for the sample driver.
TRANSACT.C	Sample driver implementation of: SQLTransact .
WEP.C	Necessary for writing DLLs for Windows 3.0.

Directory	File	File Description
C: \ODBC\SAMPLE\DEB UG		The directory where object files for the sample driver are placed at compile time, if DEBUG is set as an environment variable.
C:\ODBC\SAMPLE \NODEBUG		The directory where object files for the sample driver are placed at compile time, if DEBUG is not set as an environment variable or if it is set to null.

Several components of the Windows SDK have been included with the ODBC SDK. These need only be used if you are using a beta version of the Windows 3.1 SDK. If you are using a beta version of the Windows 3.1 SDK, make sure that the INCLUDE and LIB directories of your ODBC SDK are searched first during compilation and link. The specific files of interest are WINDOWS.H, LIBW.LIB, MDLLCEW.LIB, MLIBCEW.LIB, COMMDLG.LIB and the COMMDLG.DLL.

Constructing an ODBC Environment

The ODBC environment uses an initialization file, called ODBC.INI, to store data source names and related information.

The ODBC.INI file stores information used by the ODBC Driver Manager, ODBC drivers and the ODBC SETUP routine. For example, all connection-related ODBC functions (**SQLConnect**, **SQLDataSources**, and **SQLDriverConnect**) accept a data source name (DSN) as an argument or as an element of an argument.

ODBC drivers can read the ODBC.INI and can update it in certain instances. The ODBC Driver Manager reads the ODBC.INI file, but does not update it. Applications should not read directly from ODBC.INI. ODBC functions supply information from the initialization file in a consistent and structured manner.

There are two types of sections in the ODBC.INI file. One type of section defines the data sources accessible through ODBC. These are called data source specifications. The ODBC.INI file can contain one or more data source specifications. The section name of each data source specification defines the data source name associated with the specification.

The other type of section is a list of the data source names of each of the data source specifications. This is a single section that contains all valid data source names.

Data Source Specification

Before accessing a driver or data source, the data source must be defined in the ODBC initialization file. Each data source definition resides in a separate section in the ODBC.INI file. The data source definition section is called a data source specification.

A data source specification, at a minimum, consists of:

- A data source name (this must be the name that appears in the section heading)
- The name of an ODBC driver
- A description of the data source

The data source name and description are defined by the user.

The following depicts a basic data source specification:

```
[data-source-name]
driver = driver-DLL-name
description = description-of-data-source
```

The specification can include driver-specific information, as well. Driver-specific information can be supplied by the user or by the driver. Some examples of driver-specific information are:

```
server=<server-name>
```

```
lastuid=<last user id used to logon>
database=<database name>
OemAnsi=<conversion indicator>
```

The ODBC SETUP routine allows users to add driver-specific information to the ODBC initialization file.

Default Data Source Specification

The ODBC.INI file can contain a single default data source specification. This default specification is optional. If the default specification exists, the data source name must be "default" and the driver can be any one of the set of installed drivers. When initially defined through ODBC.SETUP, the default data source specification consists of only the data source name and the driver DLL. The driver can add information at connection time.

Sample Data Source Specifications

The exact specification of a data source is dependent upon the implementation of the ODBC driver used to access the data source and the characteristics of the data source itself. Therefore, it is important to document your requirements for the ODBC.INI file.

A data source specification for a driver for DEC Rdb might contain the following information:

```
[personnel]
driver=rdb.dll
description=Personnel database: CURLY
lastuid=smithjo
server=curly
schema=declare schema personnel filename"sys$sysdevice:
[corpdata]personnel.rdb"
```

```
[inventory]
driver=rdb.dll
description=Western Region Inventory
lastuid=smithjo
server=larry
schema=declare schema filename "sys$sysdevice:
[regionw.inventory]inventory.rdb"
```

A given driver can be referenced in more than one data source specification.

An example for MS SQL Server might contain the following:

```
[payroll]
driver=sqlsrvr.dll
lastuid=sa
```

```
database=pubs  
OemAnsi=no
```

Note that for SQL Server, a data source specification in ODBC.INI maps to a server specification in the [SQLSERVER] section of WIN.INI. In this case, the data source name must be identical to the left side of the server specification entry in the WIN.INI file.

The data source specification need not contain all of the information necessary for completing a connection to a data source. Instead, the information can be used by a driver to obtain information from another source. For example, Microsoft SQL Server maintains a list of database server connections in the WIN.INI file. A WIN.INI entry for SQL Server might contain the following:

```
[sqlserver]  
payroll=dbnmp3,\\server\pipe\sqlquery
```

A data source specification in the ODBC.INI file for the PUBS database accessed by an ODBC driver—via a server called PAYROLL—might contain the following:

```
[payroll]
driver=sqlsrvr.dll
lastuid=sa
database=pubs
OemAnsi=no
```

In this case, the driver uses the data source name to locate a corresponding entry in the [sqlserver] section of the WIN.INI file.

ODBC Data Source List

Whenever a data source specification is defined, the data source name must be added to the list maintained in the section called "[ODBC Data Sources]". The ODBC Data Sources section permits the list of data source names and associated specifications to be easily enumerated.

Each entry in the data sources section consists of the data source name and a short description of the DBMS product associated with the driver referenced by the corresponding data source specification.

This section in the ODBC.INI file might appear as follows, given the entries in the preceding example and a default data source specification:

```
[ODBC Data Sources]
default=SQL Server
personnel=Rdb
inventory=Rdb
payroll=SQL Server
```

How ODBC Functions Use the ODBC.INI File

Three ODBC functions access data source specifications in the ODBC.INI file.

SQLConnect

SQLConnect accepts a data source name as one of its arguments. When **SQLConnect** is called, the Driver Manager reads the data source specification that matches the data source name (DSN) argument. The Driver Manager loads the driver DLL listed in the data source specification. Each of the **SQLConnect** arguments—data source name, user ID, and authentication ID—is passed to the driver. The driver can read the data source specification in the ODBC.INI file, if necessary, to obtain additional connection information.

If the application specifies a data source name in its call to **SQLConnect** but there is no corresponding data source specification in the ODBC.INI file, the Driver

Manager locates the default data source specification, listed under the data source name "[default]," and loads the corresponding driver DLL. The Driver Manager passes the application-specified data source name to the driver. If there is no default data source specification, the Driver Manager returns an error.

If the application does not specify a data source name, the Driver Manager attempts to locate a default data source specification in the ODBC.INI file. If there is a default data source specification, the Driver Manager loads the driver DLL named in the default specification and passes "default" to the driver as the data source name.

If the application does not specify a data source name and no default data source specification exists, the Driver Manager returns an error.

SQLDataSources

SQLDataSources reads the [ODBC Data Sources] section of the ODBC.INI file and returns the associated list of data source names. It also reads the data source specifications that correspond to the names in the data sources section. If there is a user-defined description associated with a data source specification, **SQLDataSources** returns the description.

SQLDriverConnect

SQLDriverConnect is used as an alternative to **SQLConnect** for data sources that require connection information other than the three arguments provided by **SQLConnect**. The application specifies a data source name as part of the connection string argument of **SQLDriverConnect**.

The connection string allows an application to pass all information required by a driver to establish a connection to a specific data source. **SQLDriverConnect** can, however, prompt the user for connection information. The Driver Manager provides an optional dialog to allow the user to select a data source from a list of the data source names from the ODBC.INI file. Once the Driver Manager has a specific data source name, the Driver Manager loads the driver DLL that is listed in the corresponding data source specification.

Once the driver is loaded, the driver can display a dialog to elicit implementation-specific logon information from the user. This information depends on the requirements of the data source; it typically consists of user ID and password. The driver uses this information to replace or supplement information from the data source specification in the ODBC.INI file, or to update the data source specification. For example, after a successful connection, a driver might save the user ID for later connections to the data source.

If the application supplies a data source name but there is no corresponding data source specification in the ODBC.INI file, the Driver Manager locates the default data source specification and loads the corresponding driver DLL. The Driver Manager passes the application-supplied data source name to the driver as part of

the connection string.

If the application supplies a data source name but there is no corresponding data source specification in the ODBC.INI file and no default data source specification exists, the Driver Manager returns an error.

The ODBC SETUP Routine

The ODBC SETUP routine creates the ODBC.INI file when ODBC is first installed. The SETUP routine prompts for information to create an initial set of user-defined data source specifications. Once ODBC is installed, the user can run the SETUP routine to add, modify and delete data source specification entries from the file.

The SETUP routine uses a two-layer architecture: a top layer for generic management tasks and a lower layer for driver-specific tasks. Microsoft supplies the top layer, which supports installation, configuration, and management of drivers.

A user can run SETUP.EXE to define a default driver or select one or more drivers to install. Once the user selects a driver, the SETUP routine loads a driver-specific setup DLL. This DLL, written by the driver developer, displays a configuration dialog box that prompts the user for all relevant connection information.

Sample Application Code

XE "Application:sample code"\$XE "Static SQL:example"\$The following subsections contain two ODBC examples that are written in the C language:

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

Static SQL Example

The following example constructs SQL statements within the application. The example includes embedded SQL calls for illustrative purposes.

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

#ifdef NULL
#define NULL 0
#endif

int print_err(HDBC hdbc, HSTMT hstmt);

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

  SDWORD id;
  UCHAR name[51];
  SDWORD namelen;
  UWORD scale;

  scale = 0;
  /* EXEC SQL CONNECT TO :server USER :uid USING :authentication_string; */
  SQLAllocEnv(&henv); /* allocate an environment handle */

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

  /* connect to database */
  if (SQLConnect(hdbc, server, SQL_NTS, uid, SQL_NTS, pwd, SQL_NTS)
      != SQL_SUCCESS)
    return( print_err(hdbc, SQL_NULL_HSTMT) );

  SQLAllocStmt(hdbc, &hstmt); /* allocate a statement handle */

  /* EXEC SQL CREATE TABLE NAMEID (ID integer, NAME varchar(50)); */
  {
  UCHAR create[] = "CREATE TABLE NAMEID (ID integer, NAME varchar(50))";

  /* execute the sql statement */
  if (SQLExecDirect(hstmt, create, SQL_NTS) != SQL_SUCCESS)
    return(print_err(hdbc, hstmt));
  }
  /* EXEC SQL COMMIT WORK; */
  SQLTransact(hdbc, SQL_COMMIT); /* commit create table */

  /* EXEC SQL INSERT INTO NAMEID VALUES ( :id, :name ); */
  {
  UCHAR insert[] = "INSERT INTO NAMEID VALUES (?, ?)";
```

```

/* show the use of SQLPrepare/SQLExecute method */
/* prepare the insert*/
if(SQLPrepare(hstmt, insert, SQL_NTS) != SQL_SUCCESS)
    return(print_err(hdbc, hstmt));
SQLSetParam(hstmt, 1,SQL_C_LONG, SQL_INTEGER, (UDWORD)sizeof(UDWORD),
    scale, (PTR)&id, (SDWORD *)NULL);
SQLSetParam(hstmt, 2,SQL_C_CHAR, SQL_VARCHAR, (UDWORD)sizeof(name),
    scale, (PTR)name, (SDWORD *)NULL);

/* now assign parameter values and execute the insert*/
id=500;
(void)strcpy(name, "Babbage");
if(SQLExecute(hstmt) != SQL_SUCCESS)
    return(print_err(hdbc, hstmt));
}
/* EXEC SQL COMMIT WORK; */
SQLTransact(hdbc, SQL_COMMIT); /* commit inserts */

/* EXEC SQL DECLARE c1 CURSOR FOR SELECT ID, NAME FROM NAMEID; */
/* EXEC SQL OPEN c1; */
/* note that "declare c1 cursor for" is NOT specified by the app.
ExecuteDBL */
{
    UCHAR 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; */
/* use column binding with SQLSetParam */
SQLBindCol(hstmt, 1,SQL_C_LONG, (PTR)&id, (SDWORD)sizeof(SDWORD),
    (SDWORD *)NULL);
SQLBindCol(hstmt, 2,SQL_C_CHAR, (PTR)name, (SDWORD)sizeof(name),
    &namelen);

SQLFetch(hstmt); /* now execute the fetch */

/* finally, we should commit, discard hstmt, disconnect */
/* EXEC SQL COMMIT WORK; */
SQLTransact(hdbc, SQL_COMMIT); /* commit the transaction */

/* EXEC SQL CLOSE c1; */
SQLFreeStmt(hstmt, SQL_DROP); /* free the statement handle */

/* EXEC SQL DISCONNECT; */
SQLDisconnect(hdbc); /* disconnect from the database */

SQLFreeConnect(hdbc); /* free the connection handle */
SQLFreeEnv(henv); /* free the environment handle */

return(0);
}

```

Interactive Ad-Hoc Query Example

XE "Interactive ODBC example"§XE "Results:determining (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);
int build_indicator_message(UCHAR * errmsg, PTR * data, SDWORD *len);
UDWORD display_length(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;
SWORD rc;

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

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

/* connect to database */
if (SQLConnect(hdbc, server, SQL_NTS, uid, SQL_NTS, pwd, SQL_NTS)
    != SQL_SUCCESS )
    return( print_err(hdbc,SQL_NULL_HSTMT) );

/* allocate a statement handle */
SQLAllocStmt(hdbc, &hstmt);

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

/* see what kind of statement it was */
SQLNumResultCols(hstmt, &nresultcols);
if (nresultcols == 0) {
    /* no result columns, so must be non-select */
    /* check rowcount */
    SQLRowCount(hstmt, &rowcount);
    if (rowcount > 0) {
        /* rowcount was affected, so must have been update, insert or delete */
        (void)printf("%ld rows affected\n", rowcount);
    }
    else {
        /* rowcount == 0, so assume it was not an update/delete and
```

```

        therefore a DDL, Grant/Revoke, or Commit/Rollback.
        Of course this isn't necessarily so--it could be that the
        where clause in the update/delete did not match any rows */
        (void)printf("Operation successful\n");
        SQLTransact(hdbc, SQL_COMMIT);
    }
}
else { /* must have result rows */
    /* display column names */
    for (i=0; i<nresultcols; i++) {
        SQLDescribeCol(hstmt, i+1, colname, sizeof(colname),
            &colnamelen, &coltype, &collen[i], &scale, &nullable);
        /* assume there is a display_length function which computes
        correct length given the data type */
        collen[i] = display_length(coltype, collen[i], colname);
        (void)printf("%*.*s", collen[i], collen[i], colname);
        /* allocate memory to bind column */
        data[i] = (UCHAR *) malloc(collen[i]);
        /* bind columns to program vars, converting all types to CHAR */
        SQLBindCol(hstmt, i+1, SQL_C_CHAR, data[i], collen[i], &outlen[i]);
    }
    /* display result rows */
    while((rc=SQLFetch(hstmt))!=SQL_ERROR) {
        errmsg[0] = '\0';
        if (rc == SQL_SUCCESS_WITH_INFO) {
            for (i=0; i<nresultcols; i++)
                if (outlen[i] == SQL_NULL_DATA || outlen[i] >= collen[i])
                    build_indicator_message(errmsg, (PTR *)&data[i], collen[i], &outlen[i], i);
            (void)printf("%*.*s ", outlen[i], outlen[i], data[i]);
        } /* for all columns in this row */
        (void)printf("\n%s", errmsg); /* print truncation messages, if any */
    } /* while rows to fetch */
} /* else select statement */

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

SQLFreeStmt(hstmt, SQL_DROP); /* free statement handle */
SQLDisconnect(hdbc); /* disconnect from database */
SQLFreeConnect(hdbc); /* free connection handle */
SQLFreeEnv(henv); /* free environment handle */

return(0);
}

```

```

/*****
The following functions are given for completeness, but are
not relevant for understanding the database processing
nature 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_length(coltype, collen, colname)
SWORD coltype;
UDWORD collen;
UCHAR * colname;
{
switch (coltype) {

```

```

    case SQL_VARCHAR:
    case SQL_CHAR:
        return(max(collen, strlen((char *)colname)));
        break;

```

```

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

case SQL_INTEGER:
    return(max(10,strlen((char *)colname)));
    break;

case SQL_SMALLINT:
    return(max(5,strlen((char *)colname)));
    break;

default:
    (void)printf("Unknown datatype, %d\n", coltype);
    return(0);
    break;
}
}

int build_indicator_message(errmsg, data, collen, outlen, colnum)
UCHAR * errmsg;
PTR * data;
UDWORD collen;
SDWORD * outlen;
UWORD colnum;
{
    if (*outlen == SQL_NULL_DATA) {
        (void)strcpy((char *)data, "NULL");
        *len=4;
    }
    else {
        sprintf((char *)errmsg+strlen((char *)errmsg),
            "%d chars truncated, col %d\n", *outlen-collen+1, colnum);
        *len=255;
    }
}
}

```

Testing and Debugging an Application

The ODBC developer's kit provides information for implementing Windows-based ODBC applications. Windows development and debugging tools are available for the development process. Refer to the Windows development kit for additional information about testing and debugging your ODBC application.

Support

If you need technical support regarding the Microsoft ODBC Software Development Kit, you have a wide choice of support offerings, including:

- n Developer Services on CompuServe The developer services area is dedicated to providing developers with access to peer support and information specific to Microsoft development products. The **Microsoft Developer Services** area (GO MSDS) offers:

- Developer Forums This set of public forums covers information on Windows SDKs, languages, tools, and utilities from a developer's perspective. For example, the Client Server forum provides information about ODBC development. Here, you can receive peer support and Microsoft Support Engineer support for general API and function questions.
- Developer Knowledge Base An up-to-date reference tool containing developer-specific technical information about Microsoft products, compiled by Microsoft Product Support Engineers.
- Software Library A collection of text and graphics files, sample code, and utilities. The entire library is keyword-searchable, and the files can be downloaded for use locally.
- Confidential Technical Service Requests Microsoft offers private (fee-based per incident) technical support to help solve more complex development problems.

For more information about signing up for a CompuServe account, call (800) 848-8199. Ask for Operator 230 to receive a \$15 free usage credit for sampling the information located in the Microsoft Developer Services area.

- Microsoft Fee-Based Offerings Microsoft offers a wide range of fee-based comprehensive support plans, one which meets your specific needs. Most plans offer telephone and electronic Service Request support from knowledgeable Microsoft engineers and access to the Knowledge Base/Software library, all with a Windows interface.

For more information about ODBC support options, please call Developer Services at **(800) 227-4679**.

Appendix A SQL Grammar

The following paragraphs list the constructs that are valid in a call to **SQLPrepare**, **SQLExecute**, or **SQLExecDirect**. This grammar is not intended to restrict the SQL syntax supported by a driver. Instead, it defines a base grammar. A driver can extend this grammar to include the syntax for a specific data source.

To the left of each construct is an indicator that tells whether the construct is part of the core grammar, the minimum grammar, or both.

Elements that are part of Integrity Enhancement Facility (IEF) and are separate from the ANSI 1989 standard are presented in the following typeface and font, distinct from the rest of the grammar:

table-constraint-definition

The set of data types defined in this grammar is not necessarily supported by a specific data source; the use and syntax of each data type is database-dependent.

Element	Core	Minimum
<i>alter-table-statement</i> ::= ALTER TABLE <i>base-table-name</i> { ADD <i>column-identifier</i> <i>data-type</i> ADD (<i>column-identifier</i> <i>data-type</i> [, <i>column-identifier</i> <i>data-type</i>]...) }	X	
<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]]...)	X	

Element	Core	Mini- mum
<p><i>create-table-statement</i> ::=</p> <pre>CREATE TABLE <i>base-table-name-1</i> (<i>column-element</i> [, <i>column-element</i>] ...)</pre> <p><i>column-element</i> ::= <i>column-definition</i> <i>table-constraint-definition</i></p> <p><i>column-definition</i> ::=</p> <pre><i>column-identifier</i> <i>data-type</i> [DEFAULT <i>default-value</i>] [<i>column-constraint-definition</i> [, <i>column-constraint-definition</i>]...]</pre> <p><i>column-constraint-definition</i> ::=</p> <pre>NOT NULL [UNIQUE PRIMARY KEY] (<i>column-identifier</i>[,<i>column-identifier</i>]...) [REFERENCES <i>base-table-name-2</i> <i>referenced-columns</i>] [CHECK (<i>search-condition</i>)]</pre> <p><i>default-value</i> ::= <i>literal</i> NULL USER</p> <p><i>table-constraint-definition</i> ::=</p> <pre>UNIQUE (<i>column-identifier</i> [, <i>column-identifier</i>] ...) PRIMARY KEY (<i>column-identifier</i> [, <i>column-identifier</i>] ...) CHECK (<i>search-condition</i>) FOREIGN KEY <i>referencing-columns</i> REFERENCES <i>base-table-name-2</i> <i>referenced-columns</i></pre>	X	
<p><i>create-view-statement</i> ::=</p> <pre>CREATE VIEW <i>viewed-table-name</i> [(<i>column-identifier</i> [, <i>column-identifier</i>]...)] AS <i>query-specification</i></pre>	X	
<p><i>delete-statement-positioned</i> ::=</p> <pre>DELETE FROM <i>table-name</i> WHERE CURRENT OF <i>cursor-name</i></pre>	X	
<p><i>delete-statement-searched</i> ::=</p> <pre>DELETE FROM <i>table-name</i> [WHERE <i>search-condition</i>]</pre>	X	X

drop-index-statement ::= X
DROP INDEX *index-name*

drop-table-statement ::= X
DROP TABLE *base-table-name*
[{ CASCADE | RESTRICT }]

Element	Core	Mini- mum
<p><i>drop-view-statement</i> ::= DROP VIEW <i>viewed-table-name</i> [{ CASCADE RESTRICT }]</p>	X	
<p><i>grant-statement</i> ::= GRANT {ALL <i>grant-privilege</i> [, <i>grant-privilege</i>]... } ON <i>table-name</i> TO {PUBLIC <i>user-name</i> [, <i>user-name</i>]... } <i>grant-privilege</i> ::= DELETE INSERT SELECT UPDATE [(<i>column-identifier</i> [, <i>column-identifier</i>]...)] REFERENCES [(<i>column-identifier</i> [, <i>column-identifier</i>]...)]</p>	X	
<p><i>insert-statement</i> ::= INSERT INTO <i>table-name</i> [(<i>column-identifier</i> [, <i>column-identifier</i>]...)] { <i>query-specification</i> VALUES (<i>insert-value</i> [, <i>insert-value</i>]...) }</p> <p><i>insert-value</i> ::= <i>dynamic-parameter</i> <i>literal</i> NULL USER</p>	X	
<p><i>revoke-statement</i> ::= REVOKE {ALL <i>revoke-privilege</i> [, <i>revoke-privilege</i>]... } ON <i>table-name</i> FROM {PUBLIC <i>user-name</i> [, <i>user-name</i>]... } [{ CASCADE RESTRICT }]</p> <p><i>revoke-privilege</i> ::= DELETE INSERT SELECT</p>	X	

| UPDATE
| REFERENCES

Element	Core	Minimum
<i>select-statement ::=</i> SELECT [ALL DISTINCT] <i>select-list</i> FROM <i>table-reference</i> [, <i>table-reference</i>]... [WHERE <i>search-condition</i>] [GROUP BY <i>column-name</i> [, <i>column-name</i>]...] [HAVING <i>search-condition</i>] [UNION <i>select-statement</i>]... [<i>order-by-clause</i>]	X	
<i>select-statement ::=</i> SELECT [ALL DISTINCT] <i>select-list</i> FROM <i>table-reference</i> [, <i>table-reference</i>]... [WHERE <i>search-condition</i>] [<i>order-by-clause</i>]		X
<i>select-for-update-statement ::=</i> SELECT [ALL DISTINCT] <i>select-list</i> FROM <i>table-reference</i> [, <i>table-reference</i>]... [WHERE <i>search-condition</i>] FOR UPDATE OF <i>column-name</i> [, <i>column-name</i>]...	X	
<i>update-statement-positioned ::=</i> UPDATE <i>table-name</i> SET <i>column-identifier</i> = { <i>expression</i> NULL} [, <i>column-identifier</i> = { <i>expression</i> NULL}]... WHERE CURRENT OF <i>cursor-name</i>	X	
<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>]	X	X

Elements Used in SQL Statements

The following elements are used in the SQL statements listed previously .

Element	Core	Minimum
<i>approximate-numeric-literal</i> ::= <i>mantissa</i> <i>E</i> <i>exponent</i> <i>mantissa</i> ::= <i>exact-numeric-literal</i> <i>exponent</i> ::= [+ -] <i>unsigned-integer</i>	X	X
<i>approximate-numeric-type</i> ::= FLOAT DOUBLE PRECISION REAL	X	X
<i>base-table-identifier</i> ::= <i>user-defined-name</i>	X	X
<i>base-table-name</i> ::= [<i>user-name</i> .] <i>base-table-identifier</i>	X	
<i>base-table-name</i> ::= <i>base-table-identifier</i>		X
<i>between-predicate</i> ::= <i>expression</i> [NOT] BETWEEN <i>expression</i> AND <i>expression</i>	X	
<i>binary-literal</i> ::= {implementation defined}	X	X
<i>binary-type</i> ::= BINARY (<i>length</i>) VARBINARY (<i>length</i>) LONG VARBINARY(<i>length</i>)	X	X
<i>character</i> ::= {any character in the implementor's character set except the newline indication}	X	X
<i>character-string-literal</i> ::= '{ <i>character</i> }...'	X	X

<i>character-string-type ::=</i> CHARACTER(<i>length</i>) CHAR(<i>length</i>) CHARACTER VARYING(<i>length</i>) VARCHAR (<i>length</i>) LONG VARCHAR(<i>length</i>)	X	X
<i>column-identifier ::= user-defined-name</i>	X	X
<i>column-name ::=</i> [{ <i>table-name</i> <i>correlation-name</i> }.] <i>column-identifier</i>	X	
<i>column-name ::=</i> [<i>table-name</i> .] <i>column-identifier</i>		X
<i>comparison-operator ::=</i> < > <= >= = <>	X	X
<i>comparison-predicate ::= expression comparison-operator</i> { <i>expression</i> (<i>sub-query</i>)}	X	
<i>comparison-predicate ::=</i> <i>expression comparison-operator expression</i>		X

Element	Core	Minimum
<i>correlation-name ::= user-defined-name</i>	X	
<i>cursor-name ::= user-defined-name</i>	X	
<i>data-type ::=</i> <i>binary-type</i> <i>character-string-type</i> <i>date-type</i> <i>exact-numeric-type</i> <i>approximate-numeric-type</i> <i>time-type</i> <i>timestamp-type</i>	X	X
<i>date-literal ::= 'date-value'</i>	X	X
<i>date-separator ::= -</i>		
<i>date-type ::= DATE</i>	X	X
<i>date-value ::=</i> <i>years-value date-separator months-value date-separator</i> <i>days-value</i>	X	X
<i>days-value ::= digit digit</i>	X	X
<i>digit ::= 0 1 2 3 4 5 6 7 8 9</i>	X	X
<i>dynamic-parameter ::= ?</i>	X	X
<i>exact-numeric-literal ::=</i> [+ -] { <i>unsigned-integer</i> [<i>unsigned-integer</i>] <i>unsigned-integer</i> .	X	X

| *.unsigned-integer* }

exact-numeric-type ::= X X
DECIMAL(*precision*, *scale*)
| INTEGER
| SMALLINT
| NUMERIC(*precision*, *scale*)
| TINYINT
| BIGINT
| BIT
precision ::= *unsigned-integer*
scale ::= *unsigned-integer*

exists-predicate ::= EXISTS (*sub-query*) X

Element	Core	Minimum
<i>expression</i> ::= <i>term</i> <i>expression</i> {+ -} <i>term</i>	X	X
<i>term</i> ::= <i>factor</i> <i>term</i> {*/\} <i>factor</i>	X	X
<i>factor</i> ::= [+ -] <i>primary</i>	X	X
<i>primary</i> ::= <i>column-name</i> <i>dynamic-parameter</i> <i>literal</i> <i>set-function-reference</i> USER (2) (<i>expression</i>)	X	
<i>primary</i> ::= <i>column-name</i> <i>dynamic-parameter</i> <i>literal</i> (<i>expression</i>)		X
<i>hours-value</i> ::= <i>digit digit</i>	X	X
<i>index-identifier</i> ::= <i>user-defined-name</i>	X	
<i>index-name</i> ::= [<i>user-name.</i>] <i>index-identifier</i>	X	
<i>in-predicate</i> ::= <i>expression</i> [NOT] IN {(<i>value</i> {, <i>value</i> }...) (<i>sub-query</i>)} <i>value</i> ::= <i>literal</i> USER <i>dynamic-parameter</i>	X	
<i>join-condition</i> ::= ON <i>search-condition</i>	X	X
<i>keyword</i> ::= (see list of reserved keywords)	X	X

<i>length ::= unsigned-integer</i>	X	X
<i>letter ::= lower-case-letter upper-case-letter</i>	X	X
<i>like-predicate ::= column-name [NOT] LIKE pattern-value</i>	X	X
<i>pattern-value ::= character-string-literal dynamic-parameter USER</i>	X	
<i>pattern-value ::= character-string-literal dynamic-parameter</i> (in character-string-literal, the percent character ('%') matches 0 or more of any character; the underscore character ('_') matches 1 or 0 characters)		X
<i>literal ::= character-string-literal numeric-literal</i>	X	X
<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	X	X
<i>minutes-value ::= digit digit</i>	X	X
<i>months-value ::= digit digit</i>	X	X
<i>null-predicate ::= column-name IS [NOT] NULL</i>	X	X

Element	Core	Minimum
<i>numeric-literal</i> ::= <i>exact-numeric-literal</i> <i>approximate-numeric-literal</i>	X	X
<i>order-by-clause</i> ::= ORDER BY <i>sort-specification</i> [, <i>sort-specification</i>] <i>sort-specification</i> ::= { <i>unsigned-integer</i> <i>column-name</i> } [ASC DESC]	X	X
<i>outer-join</i> ::= <i>table-reference</i> LEFT OUTER JOIN <i>table-reference</i> <i>join-condition</i> <i>join-condition</i> ::= ON <i>search-condition</i> (Notes: For outer joins, <i>search-condition</i> must contain only the join condition between the specified <i>table-references</i> . The outer-join syntax must be placed within an escape clause.)		
<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>	X	
<i>predicate</i> ::= <i>comparison-predicate</i> <i>like-predicate</i> <i>null-predicate</i>		X
<i>quantified-predicate</i> ::= <i>expression</i> <i>comparison-operator</i> {ALL ANY} (<i>sub-query</i>)	X	
<i>referenced-columns</i> ::= (<i>column-identifier</i> [, <i>column-identifier</i>]...)	X	
<i>referencing-columns</i> ::= (<i>column-identifier</i> [, <i>column-identifier</i>]...)	X	
<i>search-condition</i> ::= <i>boolean-term</i> [OR <i>search-condition</i>] <i>boolean-term</i> ::= <i>boolean-factor</i> [AND <i>boolean-term</i>] <i>boolean-factor</i> ::= [NOT] <i>boolean-primary</i>	X	X

<i>boolean-primary ::= predicate (search-condition)</i>		
<i>seconds-fraction ::= digit digit digit [digit digit digit]</i>	X	X
<i>seconds-value ::= digit digit</i>	X	X
<i>select-list ::= * select-sublist [, select-sublist]...</i>	X	X
<i>select-sublist ::= expression {table-name correlation-name}.*</i>	X	
<i>select-sublist ::= expression</i>		X
<i>separator ::=</i> The blank character or an implementation-defined end-of-line indicator.	X	X

Element	Core	Minimum
<p><i>set-function-reference</i> ::= COUNT(*) <i>distinct-function</i> <i>all-function</i> <i>distinct-function</i> ::= {AVG COUNT MAX MIN SUM} (DISTINCT <i>column-name</i>) <i>all-function</i> ::= {AVG MAX MIN SUM} (<i>expression</i>)</p>	X	
<p>SQL-<i>escape-clause</i> ::= <i>standard-SQL-escape-initiator</i> <i>extended-SQL-text</i> <i>standard-SQL-escape-terminator</i> <i>extended-SQL-escape-prefix</i> <i>extended-SQL-text</i> <i>extended-SQL-escape-terminator</i> <i>standard-SQL-escape-initiator</i> ::= <i>standard-SQL-escape-prefix</i> <i>escape-identification</i>, <i>standard-SQL-escape-prefix</i> ::= --*(<i>extended-SQL-escape-prefix</i> ::= { <i>standard-SQL-escape-terminator</i> ::= --*) <i>extended-SQL-escape-terminator</i> ::= } SQL-<i>escape-identification</i> ::= SQL-<i>escape-vendor-clause</i> SQL-<i>escape-vendor-clause</i> ::= VENDOR(Microsoft), PRODUCT(ODBC)</p>	X	X
<p><i>sub-query</i> ::= SELECT [ALL DISTINCT] <i>select-list</i> FROM <i>table-reference</i> [, <i>table-reference</i>]... [WHERE <i>search-condition</i>] [GROUP BY <i>column-name</i> [, <i>column-name</i>]...] [HAVING <i>search-condition</i>]</p>	X	
<p><i>table-identifier</i> ::= <i>user-defined-name</i></p>	X	X
<p><i>table-name</i> ::= [<i>user-name</i>.]<i>table-identifier</i></p>	X	

<i>table-name ::= table-identifier</i>		X
<i>table-reference ::= table-name [correlation-name]</i>	X	
<i>table-reference ::= table-name</i>		X
<i>time-literal ::= 'time-value'</i>	X	X

Element	Core	Minimum
<i>time-separator</i> ::= :		
<i>time-type</i> ::= TIME	X	X
<i>time-value</i> ::= <i>hours-value time-separator minutes-value time-separator</i> <i>seconds-value</i>	X	X
<i>timestamp-literal</i> ::= 'date-value:time-value.[seconds-fraction]'	X	X
<i>timestamp-type</i> ::= TIMESTAMP	X	X
<i>token</i> ::= <i>delimiter-token</i> <i>non-delimiter-token</i> <i>delimiter-token</i> ::= <i>character-string-literal</i> , () < > . : = * + - / <> >= <= ? <i>non-delimiter-token</i> ::= <i>keyword</i> <i>numeric-literal</i> <i>user-defined-name</i>	X	X
<i>unsigned-integer</i> ::= { <i>digit</i> }...	X	X
<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	X	X
<i>user-defined-name</i> ::= <i>letter</i> [<i>digit</i> <i>letter</i> _]...	X	X
<i>user-name</i> ::= <i>user-defined-name</i>	X	X

<i>viewed-table-identifier ::= user-defined-name</i>	X	
<i>viewed-table-name ::= [user-name.]viewed-table-identifier</i>	X	
<i>years-value ::= digit digit digit digit</i>	X	X

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.

ABSOLUTE	CREATE	FROM
ADA	CURRENT	FULL
ADD	CURRENT_DATE	GET
ALL	CURRENT_TIME	GLOBAL
ALLOCATE	CURRENT_TIMEST	GO
ALTER	AMP	GOTO
AND	CURSOR	GRANT
ANY	DATE	GROUP
ARE	DAY	HAVING
AS	DEALLOCATE	HOUR
ASC	DEC	IDENTITY
ASSERTION	DECIMAL	IGNORE
AT	DECLARE	IMMEDIATE
AUTHORIZATION	DEFERRABLE	IN
AVG	DEFERRED	INCLUDE
BEGIN	DELETE	INDEX
BETWEEN	DESC	INDICATOR
BIT	DESCRIBE	INITIALLY
BIT_LENGTH	DESCRIPTOR	INNER
BY	DIAGNOSTICS	INPUT
CASCADE	DICTIONARY	INSENSITIVE
CASCADED	DISCONNECT	INSERT
CASE	DISPLACEMENT	INTEGER
CAST	DISTINCT	INTERSECT
CATALOG	DOMAIN	INTERVAL
CHAR	DOUBLE	INTO
CHAR_LENGTH	DROP	IS
CHARACTER	ELSE	ISOLATION
CHARACTER_LEN	END	JOIN
GTH	END-EXEC	KEY
CHECK	ESCAPE	LANGUAGE
CLOSE	EXCEPT	LAST
COALESCE	EXCEPTION	LEFT
COBOL	EXEC	LEVEL
COLLATE	EXECUTE	LIKE
COLLATION	EXISTS	LOCAL
COLUMN	EXTERNAL	LOWER
COMMIT	EXTRACT	MATCH
CONNECT	FALSE	MAX
CONNECTION	FETCH	MIN
CONSTRAINT	FIRST	MINUTE
CONSTRAINTS	FLOAT	MODULE
CONTINUE	FOR	MONTH
CONVERT	FOREIGN	MUMPS
CORRESPONDING	FORTRAN	NAMES
COUNT	FOUND	NATIONAL

NCHAR	SQL
NEXT	SQLCA
NONE	SQLCODE
NOT	SQLERROR
NULL	SQLSTATE
NULLIF	SQLWARNING
NUMERIC	SUBSTRING
OCTET_LENGTH	SUM
OF	SYSTEM
OFF	TABLE
ON	TEMPORARY
ONLY	THEN
OPEN	TIME
OPTION	TIMESTAMP
OR	TIMEZONE_HOUR
ORDER	TIMEZONE_MINUTE
OUTER	E
OUTPUT	TO
OVERLAPS	TRANSACTION
PARTIAL	TRANSLATE
PASCAL	TRANSLATION
PLI	TRUE
POSITION	UNION
PRECISION	UNIQUE
PREPARE	UNKNOWN
PRESERVE	UPDATE
PRIMARY	UPPER
PRIOR	USAGE
PRIVILEGES	USER
PROCEDURE	USING
PUBLIC	VALUE
RESTRICT	VALUES
REVOKE	VARCHAR
RIGHT	VARYING
ROLLBACK	VIEW
ROWS	WHEN
SCHEMA	WHENEVER
SCROLL	WHERE
SECOND	WITH
SECTION	WORK
SELECT	YEAR
SEQUENCE	
SET	
SIZE	
SMALLINT	
SOME	

Index

- µANSI standards, 23
- Application, 17
 - sample code, 67
- Arguments
 - variable-length data and, 27
- Autocommit, setting, 36
- Batched SQL statements, results, 53
- Binding result columns, 25, 40
- Call level interface, 25
- CLI, 25
- Concurrency control, setting, 45
- Configuration
 - multiple-tier, 19
 - single-tier, 18
- Connection handles, 28
 - and transactions, 22
- Connections
 - establishing, 35
 - terminating, 55
- Data source
 - listing, 36
- Data types, 29
- Database environment, 18
- Declare Cursor (SQL statement), 41
- DELETE (SQL statement)
 - positioned, 40
 - DELETED, positioned, 41
- Developer's Kit, contents, 57
- Direct execution, 40
- Direct invocation (ANSI), 23
- Driver, 17
 - types, 18
- Driver Manager, 17
 - communicating with, 35
- Dynamic SQL, 24
- Embedded SQL (ANSI), 23
- Errors, retrieving, 52
- Escape sequence
 - data types, 46
 - outer joins, 49
 - translating to SQL text, 53

- Handles
 - connection and statement, 28
 - releasing statement, 55
 - requesting statement, 40
- Installing ODBC, 59
- Interactive ODBC example, 70
- Interoperability, 15, 25
- Large data values
 - sending, 44
- Module language, 23
- ODBC, 15
 - components, 16
 - function call types, 15
 - functionality sets, 21
 - functions, general information, 27
 - installing, 59
- Parameters
 - markers, 24
 - obtaining data type information, 53
 - specifying arrays of, 44
- Prepared execution, 40
- Prepared SQL statements, advantages, 41
- Processing results, 31
- Results
 - assigning storage for, 40
 - binding, 40
 - determining (example), 70
 - determining characteristics, 51
 - processing, 31, 51
 - retrieving data, 51
- Return codes, 30
- Rowcount, setting maximum, 36
- Scrollable cursors, overview, 45
- SELECT (SQL statement)
 - storing results, 51
- SQL
 - references for additional information, 14
- SQL statements
 - batched processing, 53
 - dynamic, 24
 - embedded, 23
 - executing, 40
 - positioned UPDATE and DELETE, 40, 41
 - preparing and executing, 39
 - processing, 29
 - processing asynchronously, 46

- processing results, 51
- static, 24
- SQL, overview, 23
- SQL_ERROR return code, 30
- SQL_INVALID_HANDLE return code, 30
- SQL_NO_DATA_FOUND return code, 30
- SQL_NTS, 27
- SQL_NULL_DATA, 27
- SQL_SUCCESS return code, 30
- SQL_SUCCESS_WITH_INFO return code, 30
- SQL2, overview, 24
- SQLBindCol, 40
- SQLCancel
- terminating statement processing, 55
- SQLConnect
 - compared to SQLDriverConnect, 37
 - SQLDataSources, 36
 - SQLDescribeCol
 - determining result characteristics, 51
 - retrieving data types, 29
 - SQLDescribeParam, 53
 - SQLDisconnect
 - closing a connection, 55
 - SQLDriverConnect, 37
 - SQLError, 30, 52
 - checking statement cancellation, 55
 - SQLExecDirect, 41
 - submitting SQL statements, 29
 - SQLExtendedFetch, 46
 - SQLFetch
 - retrieving result rows, 51
 - SQLFreeConnect
 - releasing a connection handle, 55
 - SQLGetCursorName, 40
 - SQLGetInfo, 36
 - SQLGetTypeInfo, 36
 - retrieving supported data types, 29
 - SQLNativeSql, 53
 - SQLNumResultCols
 - determining result characteristics, 51
 - SQLParamOptions, 44
 - SQLPrepare, 40
 - submitting SQL statements, 29
 - SQLRowCount
 - determining result characteristics, 51
 - SQLSetConnectOption, 36

- SQLSetCursorName, 40
- SQLSetParam, 40
- SQLSetPos, 46
- SQLSetStmtOption, 36
- SQLTransact
 - terminating a transaction, 55
- Statement handles, 28
 - releasing, 55
 - requesting, 40
- Statements, terminating, 55
- Static SQL, 24
 - example, 67
- Stored procedures, processing results, 53
- Timeout values, setting, 36
- Transactions, 22
 - setting autocommit option, 36
 - terminating, 55
- UPDATE (SQL statement)
 - positioned, 40
- UPDATE, positioned, 41
- Variable-length data, in arguments, 27