

User DSN Tab

Adds, deletes, or sets up data sources with User DSNs. These data sources are local to a computer, and can be used only by the current user.

Option	Description
User Data Sources list	A list of all user DSNs that includes the name of each DSN and the driver associated with the DSN. Double-clicking a user DSN displays the driver-specific data source setup dialog box.
Configure	Displays the driver-specific data source setup dialog box that enables you to change the configuration of an existing user data source. You must select the name of a user data source from the list before clicking Configure .
Add	Adds a new user data source. When you click Add , the Create New Data Source dialog box appears with a list of drivers. Choose the driver for which you are adding a user data source. After you click Finish , a driver-specific setup dialog box appears.
Remove	Removes an existing user data source. You must select the name of the user data source you want to remove from the list before clicking Remove .
OK	Closes the Administrator dialog box. You do not have to click OK to accept changes to the User Data Sources list. Changes to the list are accepted once you click OK in the data source setup dialog box.
Cancel	Closes the Administrator dialog box. Changes to the User Data Sources list will not be rejected if you click Cancel .
Help	Displays this Help screen.

System DSN Tab

Adds, deletes, or sets up data sources with system DSNs. These data sources are local to a computer, rather than dedicated to a user. The system, or any user having privileges, can use a data source set up with a system DSN.

Option	Description
System Data Sources list	A list of all system DSNs that includes the name of each DSN and the driver associated with the DSN. Double-clicking a system DSN displays the driver-specific data source setup dialog box.
Configure	Displays the driver-specific data source setup dialog box that enables you to change the configuration of an existing system data source. You must select the name of a system data source from the list before clicking Configure .
Add	Adds a new system data source. If you click this button, the Create New Data Source dialog box appears with a list of drivers. Choose the driver for which you are adding a system data source. After you click Finish , a driver-specific setup dialog box appears.
Remove	Removes an existing system data source. You must select the name of the system data source you want to remove from the list before clicking Remove .
OK	Closes the Administrator dialog box. You do not have to click OK to accept changes to the System Data Sources list. Changes to the list are accepted once you click OK in the data source setup dialog box.
Cancel	Closes the Administrator dialog box. Changes to the User Data Sources list will not be rejected if you click Cancel .
Help	Displays this Help screen.

File DSN Tab

Adds, deletes, or sets up data sources with file DSNs. These are file-based data sources that can be shared among all users that have the same drivers installed, so they have access to the database. These data sources need not be dedicated to a user or local to a computer.

Option	Description
File Data Sources list	Displays all file DSNs and subdirectories contained within the directory displayed in the Look In box. Double-clicking a file DSN displays the driver-specific data source setup dialog box.
New	Adds a new file data source. If you click this button, the Create New Data Source dialog box appears with a list of drivers. Choose the driver for which you are adding a file DSN. After you click Next , you can specify the keywords for the file DSN.
Remove	Removes an existing file data source. You must select the file data source you want to remove from the list before clicking Remove .
Configure	Displays the driver-specific data source setup dialog box that enables you to change the configuration of an existing file data source. You must select the name of a file data source from the list before clicking Configure .
Help	Displays this Help screen.
Set Directory	Establishes the displayed directory as the default directory to be displayed when the ODBC Administrator is executed.
Look In	Displays the current directory for which the subdirectories and file DSNs are displayed in the window below. Clicking the down arrow to the right of the text box displays the entire path to that directory. The default directory that is initially displayed when the ODBC Administrator is first executed is contained in the system information, but you can change it with the Set Directory button.
Up	Changes the directory displayed in the Look In box to be the directory directly above the current directory.
OK	Closes the Administrator dialog box. You do not have to press OK to accept changes to the File Data Sources list. Changes to the list are accepted once you click OK in the data source setup dialog box.
Cancel	Closes the Administrator dialog box. Changes to the File Data Sources list will not be rejected if you click Cancel .
Help	Displays this Help screen.

ODBC Drivers Tab

Displays information about the installed ODBC drivers. The **ODBC Drivers** list shows you which drivers are already installed on your computer.

Option	Description
ODBC Drivers list	The name, version, company, file name, and file creation date of each ODBC driver installed on the computer.
OK	Closes the Administrator dialog box.
Cancel	Closes the Administrator dialog box.
Help	Displays this Help screen.

This dialog box no longer includes **Add** and **Delete** buttons to add or delete a driver. You must add or delete an ODBC driver through its setup program.

Tracing Tab

Enables you to specify how the ODBC Driver Manager traces calls to ODBC functions. The Driver Manager can trace calls continuously, or for one connection only; can perform tracing dynamically; or can allow tracing to be performed by a custom trace DLL.

Option	Description
Start Tracing Now	Enables dynamic tracing that is performed as long as the ODBC Administrator dialog box is displayed. Dynamic tracing can be enabled whether or not a connection has been made. After you click Start Tracing Now , the button is replaced by Stop Tracing Now . Dynamic tracing remains enabled until you click Stop Tracing Now .
Start Visual Studio Analyzer Tracing	Enables Emitting Visual Studio Analyzer Events. After you click Start Visual Studio Analyzer Tracing , the button is replaced by Stop Visual Studio Analyzer Tracing , which disables Visual Studio Analyzer Events. Visual Studio Analyzer Tracing remains enabled until you click Stop Visual Studio Analyzer Tracing .
Log file path	Displays the path and file name for the file in which the tracing information will be stored. The default path and file name (sql.log) are taken from the system information, but you can specify a new file by either entering a new path and file name, or clicking Browse and selecting a directory and file.
Browse	Allows you to select the path and file name for the log file by browsing machine directories.
Custom Trace DLL	Allows the user to select a trace DLL other than odbctrac.dll to perform tracing. The odbctrac.dll file that is shipped with the SDK can be replaced by a custom DLL of your choice. Enter the path and file name of the custom DLL, or click Select DLL to browse the directories.
Select DLL	Allows the user to browse the directory structure for a custom trace DLL. When you choose a DLL, its path and file name will appear in the Custom Trace DLL text box.
OK	Accepts changes to tracing and closes the Administrator dialog box.
Cancel	Closes the Administrator dialog box without accepting changes to tracing.

Apply	Accepts any changes that have been made to tracing without closing the Administrator dialog box. The Apply button is shaded if no changes have been made.
Help	Displays this Help screen.

User Data Sources List

This list shows all user DSNs, including the name of each DSN and the driver associated with the DSN. To display the driver-specific data source setup dialog box for a user data source, double-click the user DSN.

Configure

To display the driver-specific data source setup dialog box that enables you to change the configuration of an existing user data source, click **Configure**. You must select the name of a user data source from the list before clicking this button.

Add

To add a new user data source, click **Add**. The **Create New Data Source** dialog box appears with a list of drivers. Choose the driver for which you are adding a user data source. After you click **Finish**, a driver-specific setup dialog box appears.

Remove

To remove an existing user data source from the list, click **Remove**. You must select the name of the user data source you want to remove from the list before clicking this button.

OK

To close the **Administrator** dialog box, click **OK**. You do not have to click **OK** to accept changes to the data sources list. Changes to the list are accepted once you click **OK** in the data source setup dialog box.

Cancel

To close the **Administrator** dialog box, click **Cancel**. Changes to the data sources list will not be rejected if you click **Cancel**.

Help

To display help for this tab, click **Help**.

System Data Sources List

This list shows all system DSNs, including the name of each DSN and the driver associated with the DSN. To display the driver-specific data source setup dialog box for a user data source, double-click the system DSN.

Configure

To display the driver-specific data source setup dialog box that enables you to change the configuration of an existing system data source, click **Configure**. You must select the name of a system data source from the list before clicking this button.

Add

To add a new system data source, click **Add**. The **Create New Data Source** dialog box appears with a list of drivers. Choose the driver for which you are adding a system data source. After you click **Finish**, a driver-specific setup dialog box appears.

Remove

To remove an existing system data source from the list, click **Remove**. You must select the name of the system data source you want to remove from the list before clicking this button.

File Data Sources List

This list shows all file DSNs and subdirectories contained within the directory displayed in the **Look In** box. Double-clicking a file DSN displays the driver-specific data source setup dialog box.

Configure

To display the driver-specific data source setup dialog box that enables you to change the configuration of an existing file data source, click **Configure**. You must select the name of a file data source from the list before clicking this button.

New

To add a new file data source, click **New**. The **Create New Data Source** dialog box appears with a list of drivers. Choose the driver for which you are adding a file data source. After you choose **Next**, you can specify the keywords for the file DSN.

Remove

To remove an existing file data source from the list, click **Remove**. You must select the name of the file data source you want to remove from the list before clicking this button.

Set Directory

To establish the displayed directory as the default directory to be displayed when the ODBC Administrator is executed, click **Set Directory**.

Look In

The **Look In** box displays the current directory for which the subdirectories and file DSNs are displayed in the window below. The default directory contained in the system information is initially displayed.

Up

To change the directory displayed in the **Look In** box to be the directory directly above the current directory, click **Up**.

Directory

To specify the path of the directory whose contents are displayed, enter text in the **Directory** text box. You can enter an entire path, or a path relative to the current directory.

ODBC Drivers List

This list shows the name, version, company, file name, and release date of each ODBC driver installed on the computer.

Start Tracing Now

To enable tracing to be performed as long as the **ODBC Administrator** dialog box is displayed, click **Start Tracing Now**. Dynamic tracing can be enabled whether or not a connection has been made. After you click **Start Tracing Now**, the button is replaced by **Stop Tracing Now**. Dynamic tracing is disabled when you click **Stop Tracing Now** or close the **ODBC Administrator** dialog box.

Enable Visual Studio Analyzer Tracing

To enable Emitting Visual Studio Analyzer Event tracing to be performed continuously, click **Start Visual Studio Analyzer Tracing**. After you click **Start Visual Studio Analyzer Tracing**, the button is replaced by **Stop Visual Studio Analyzer Tracing**. Visual Studio Analyzer tracing is disabled when you click **Stop Visual Studio Analyzer Tracing**.

Log File Path

To change the path and file name of the log file, enter text into the **Log File Path** text box. This box displays the path and file name for the file in which the tracing information will be stored. The default path and file name (sql.log) are taken from the system information, but you can specify a new file by either entering a new path and file name, or clicking **Browse** and selecting a directory and file.

Browse

Allows you to select the path and file name for the log file by browsing machine directories.

Custom Trace DLL

To select a trace DLL other than odbctrac.dll to perform tracing, enter a path and file name in the **Custom Trace DLL** text box. The odbctrac.dll file that is shipped with the SDK can be replaced by a custom DLL of your choice. Enter the path and file name of the custom DLL or click **Select DLL** to browse the directories.

Select DLL

To browse the directory structure for a custom trace DLL, click **Select DLL**. When you choose a DLL, its path and file name will appear in the **Custom Trace DLL** text box.

OK

Click **OK** to accept changes to tracing and close the **Administrator** dialog box.

Cancel

Click **Cancel** to close the **Administrator** dialog box without accepting changes to tracing.

Apply

Click **Apply** to accept any changes that have been made to tracing without closing the **Administrator** dialog box. The **Apply** button is shaded if no changes have been made.

File Data Source Tab

Allows you to connect with a data source that has file DSNs. These are file-based data sources that can be shared among all users that have the same drivers installed, so they have access to the database. These data sources need not be dedicated to a user or local to a computer.

Option	Description
File Data Sources list	Displays all file DSNs and subdirectories contained within the directory displayed in the Look In box. Double-clicking a DSN connects to the data source.
Look In	Displays the current directory for which the subdirectories and file DSNs are displayed in the window below. Clicking the down arrow to the right of the text box displays the entire directory structure.
Up	Changes the directory displayed in the Look In box to be the directory directly above the current directory.
DSN Name	Displays the file DSN name selected in the File Data Sources list, or you can enter a new file DSN name.
New	Adds a new file data source. If you click this button, the Create New Data Source dialog box appears with a list of drivers. Choose the driver for which you are adding a file DSN. After you click Next , you can specify the keywords for the file DSN.
OK	<p>Closes the Administrator dialog box and connects to the file data source that is highlighted in the list or entered in the DSN Name text box.</p> <p>You do not have to click OK to accept changes to the File Data Sources list. Changes to the list are accepted once you click OK in the Data Source Setup dialog box.</p>
Cancel	Closes the Administrator dialog box without connecting to the file data source. Changes to the File Data Sources list will not be rejected if you click Cancel .
Help	Displays this Help screen.

Machine Data Source Tab

Allows you to connect with a data source that has a user DSN or a system DSN. Machine data sources are specific to this machine and cannot be shared. User data sources are specific to a user on this machine. System data sources can be used by all users on this machine, or by a system-wide service.

Option	Description
Machine Data Sources list	A list of all user and system DSNs that includes the name and type of each DSN. Double-clicking a DSN connects to the data source.
New	Adds a new machine data source. If you click this button, the Create New Data Source dialog box appears with a list of drivers. Choose the driver for which you are adding a user or system DSN. After you click Next , you can verify that the keywords for the DSN are correct.
OK	Closes the Administrator dialog box and connects to the file data source that is highlighted in the list. You do not have to click OK to accept changes to the Machine Data Sources list. Changes to the list are accepted once you click OK in the Data Source Setup dialog box.
Cancel	Closes the Administrator dialog box without connecting to the file data source. Changes to the Machine Data Sources list will not be rejected if you click Cancel .
Help	Displays this Help screen.

File Data Sources List

Displays all file DSNs and subdirectories contained within the directory displayed in the **Look In** box. Double-clicking a DSN connects to the data source.

Look In

Displays the current directory for which the subdirectories and file DSNs are displayed in the window below. Clicking the down arrow to the right of the text box displays the entire directory structure.

Up

Changes the directory displayed in the **Look In** box to be the directory directly above the current directory.

DSN Name

Displays the file DSN name selected in the **File Data Sources** list, or you can enter a new file DSN name.

Connect

Connects to the file DSN that is highlighted in the list or entered in the **DSN Name** text box.

New

Adds a new file data source. If you click this button, the **Create New Data Source** dialog box appears with a list of drivers. Choose the driver for which you are adding a file DSN. After you click **Next**, you can specify the keywords for the file DSN.

OK

Closes the **Administrator** dialog box and connects to the file data source that is highlighted in the list or entered in the **DSN Name** text box.

You do not have to click **OK** to accept changes to the **File Data Sources** list. Changes to the list are accepted once you click **OK** in the **Data Source Setup** dialog box.

Cancel

Closes the **Administrator** dialog box without connecting to the file data source. Changes to the **Machine Data Sources** list will not be rejected if you click **Cancel**.

Machine Data Sources List

A list of all user and system DSNs that includes the name and type of each DSN. Double-clicking a DSN connects to the data source.

New

Adds a new machine data source. If you click this button, the **Create New Data Source** dialog box appears with a list of drivers. Choose the driver for which you are adding a user or system DSN. After you click **Next**, you can verify that the keywords for the DSN are correct.

OK

Closes the **Administrator** dialog box and connects to the file data source that is highlighted in the list.

You do not have to click **OK** to accept changes to the **Machine Data Sources** list. Changes to the list are accepted once you click **OK** in the **Data Source Setup** dialog box.

Cancel

Closes the **Administrator** dialog box without connecting to the file data source. Changes to the **Machine Data Sources** list will not be rejected if you click **Cancel**.

Connection Pooling Tab

Allows you to alter the connection retry wait time and connection timeout period for a selected driver when using connection pooling. It also allows you to enable and disable Performance Monitoring, which records a number of connection statistics.

Option	Description
ODBC Drivers list	The name, version, company, file name, and file-creation date of each ODBC driver installed on the computer.
Timeout for selected driver	Sets connection pooling timeout in seconds for the selected driver.
Enable	Enables performance monitoring.
Disable	Disables performance monitoring.
Retry Wait Time	When the ODBC Driver Manager detects that a database server is unavailable, it waits for a certain amount of time before it attempts to connect again. You can adjust the waiting time by setting the Retry Wait Time value in seconds.
Help	Displays this screen.

Connection Pooling Timeout

Specifies the timeout period in seconds for a selected driver.

PerfMon

Enables or disables Performance Monitoring of the connection.

Retry Wait Time

Specifies the wait time for connection retries in seconds.

OK

Closes the **Administrator** dialog box and activates any Connection Pooling options selected.

Cancel

Closes the **Administrator** dialog box without activating any options selected.

About Tab

Displays information about the ODBC core components, which include the Driver Manager, the cursor library, the installer DLL, and any other files that make up the core components.

Option	Description
Core Component list	The description, version, file name, and location of each ODBC core component.
OK	Closes the Administrator dialog box.
Cancel	Closes the Administrator dialog box.
Help	Displays this screen.

Glossary

A

Access plan A plan generated by the database engine to execute an SQL statement. Equivalent to executable code compiled from a third-generation language such as C.

Aggregate function A function that generates a single value from a group of values, often used with GROUP BY and HAVING clauses. Aggregate functions include AVG, COUNT, MAX, MIN, and SUM. Also known as *set functions*. See also scalar function.

ANSI American National Standards Institute. The ODBC API is based on the ANSI Call-Level Interface.

APD Application Parameter Descriptor.

API Application Programming Interface. A set of routines that an application uses to request and carry out lower-level services. The ODBC API is composed of the ODBC functions.

Application An executable program that calls functions in the ODBC API.

Application Parameter Descriptor (APD) A descriptor that describes the dynamic parameters used in an SQL statement before any conversion specified by the application.

Application Row Descriptor (ARD) A descriptor that represents the column metadata and data in the application's buffers, describing a row of data following any data conversion specified by the application.

ARD Application Row Descriptor.

Auto-commit mode A transaction commit mode in which transactions are committed immediately after they are executed.

B

Behavioral change A change in certain functionality from ODBC 3.x behavior to ODBC 2.x behavior, or vice versa. Caused by changing the SQL_ATTR_ODBC_VERSION environment attribute.

Binary large object (BLOB) Any binary data over a certain number of bytes, such as 255. Typically much longer. Such data is generally sent to and retrieved from the data source in parts. Also known as *long data*.

Binding As a verb, the act of associating a column in a result set or a parameter in an SQL statement with an application variable. As a noun, the association.

Binding offset A value added to the data buffer addresses and length/indicator buffer addresses for all bound column or parameter data, producing new addresses.

Block cursor A cursor capable of fetching more than one row of data at a time.

Buffer A piece of application memory used to pass data between the application and driver. Buffers often come in pairs: a *data buffer* and a *data length buffer*.

Byte Eight bits or one octet. See also octet.

C

C data type The data type of a variable in a C program, in this case the application.

Catalog The set of system tables in a database that describe the shape of the database. Also known as a schema or data dictionary.

Catalog function An ODBC function used to retrieve information from the database's catalog.

CLI See API.

Client/server A database access strategy in which one or more clients access data through a server. The clients usually implement the user interface while the server controls database access.

Column The container for a single item of information in a row. Also known as a *field*.

Commit To make the changes in a transaction permanent.

Concurrency The ability of more than one transaction to access the same data at the same time.

Conformance level A discrete set of functionality supported by a driver or data source. ODBC defines API conformance levels and SQL conformance levels.

Connection A particular instance of a driver and data source.

Connection browsing Searching the network for data sources to connect to. Connection browsing might involve several steps. For example, the user might first browse the network for servers, then browse a particular server for a database.

Connection handle A handle to a data structure that contains information about a connection.

Current row The row currently pointed to by the cursor. Positioned operations act on the current row.

Cursor A piece of software that returns rows of data to the application. Probably named after the blinking cursor on a computer terminal; just as that cursor indicates the current position on the screen, a cursor on a result set indicates the current position in the result set.

D

Data buffer A buffer used to pass data. Often associated with a data buffer is a *data length buffer*.

Data dictionary See catalog.

Data length buffer A buffer used to pass the length of the value in a corresponding *data buffer*. The data length buffer is also used to store indicators, such as whether the data value is null terminated.

Data source The data that the user wants to access and its associated operating system, DBMS, and network platform (if any).

Data type The type of a piece of data. ODBC defines C and SQL data types. See *also* type indicator.

Data-at-execution column A column for which data is sent after **SQLSetPos** is called. So named because the data is sent at execution time rather than being placed in a rowset buffer. Long data is usually sent in parts at execution time.

Data-at-execution parameter A parameter for which data is sent after **SQLExecute** or **SQLExecDirect** is called. So named because the data is sent when the SQL statement is executed rather than being placed in a parameter buffer. Long data is usually sent in parts at execution time.

Database A discrete collection of data in a DBMS. Also a DBMS.

Database engine The software in a DBMS that parses and executes SQL statements and accesses the physical data.

DBMS Database Management System. A layer of software between the physical database and the user. The DBMS manages all access to the database.

DBMS-based driver A driver that accesses physical data through a standalone database engine.

DDL Data Definition Language. Those statements in SQL that define, as opposed to manipulate, data. For example, **CREATE TABLE**, **CREATE INDEX**, **GRANT**, and **REVOKE**.

Delimited Identifier An identifier that is enclosed in identifier quote characters so it can contain special characters or match keywords (also known as a quoted identifier).

Descriptor A data structure that holds information about either column data or dynamic parameters. The physical representation of the descriptor is not defined; applications gain direct access to a descriptor only by manipulating its fields by calling ODBC functions with the descriptor handle.

Desktop database A DBMS designed to run on a personal computer. In most cases, these DBMSs do not provide a standalone database engine and must be accessed through a file-based driver. The engines in these drivers usually have reduced support for SQL and transactions—for example, dBASE, Paradox, Btrieve, or FoxPro.

Diagnostic A record containing diagnostic information about the last function called that used a particular handle. Diagnostic records are associated with environment, connection, statement, and descriptor handles.

DML Data Manipulation Language. Those statements in SQL that manipulate, as opposed to define, data—for example, **INSERT**, **UPDATE**, **DELETE**, and **SELECT**.

Driver A routine library that exposes the functions in the ODBC API. Drivers are specific to a single DBMS.

Driver Manager A routine library that manages access to drivers for the application. The Driver Manager loads and unloads drivers, and passes calls to ODBC functions to the correct driver. Another way to say this is that the Driver Manager connects (and disconnects) to the drivers while passing calls to ODBC functions.

Driver setup DLL A DLL that contains driver-specific installation and configuration functions.

Dynamic cursor A scrollable cursor capable of detecting updated, deleted, or inserted rows in the result set.

Dynamic SQL A type of embedded SQL in which SQL statements are created and compiled at run time. *See also* static SQL.

E

Embedded SQL SQL statements that are included directly in a program written in another language, such as COBOL or C. ODBC does not use embedded SQL. *See also* static SQL *and* dynamic SQL.

Environment A global context in which to access data; associated with the environment is any information that is global in nature, such as a list of all connections in that environment.

Environment handle A handle to a data structure that contains information about the environment.

Escape clause A clause in an SQL statement.

Execute To run an SQL statement.

F

Fat cursor *See* block cursor.

Fetch To retrieve one or more rows from a result set.

Field *See* column.

File-based driver A driver that accesses physical data directly. In this case, the driver contains a database engine and acts as both driver and data source.

File data source A data source for which connection information is stored in a .DSN file.

Foreign key A column or columns in a table that match the primary key in another table.

Forward-only cursor A cursor that can only move forward through the result set and fetch one row at a time. Most relational databases support only forward-only cursors.

H

Handle A value that uniquely identifies something such as a file or data structure. Handles are meaningful only to the software that creates and uses them, but are passed by other software to identify things. ODBC defines handles for environments, connections, statements, and descriptors.

I

Implementation Parameter Descriptor (IPD) A descriptor that describes the dynamic parameters used in an SQL statement after any conversion specified by the application.

Implementation Row Descriptor (IRD) A descriptor that describes a row of data before any conversion specified by the application.

Installer DLL A DLL that installs ODBC components and configures data sources.

Integrity Enhancement Facility A subset of SQL designed to maintain the integrity of a database.

Interface conformance level The level of the ODBC 3.7 interface supported by a driver; can be Core, Level 1, or Level 2.

Interoperability The ability of one application to use the same code when accessing data in different DBMSs.

IPD Implementation Parameter Descriptor.

IRD Implementation Row Descriptor.

ISO/IEC International Standards Organization/International Electrotechnical Commission. The ODBC API is based on the ISO/IEC Call-Level Interface.

J

Join An operation in a relational database that links the rows in two or more tables by matching values in specified columns.

K

Key A column or columns whose values identify a row. *See also* primary key *and* foreign key.

Keyset A set of keys used by a mixed or keyset-driven cursor to refetch rows.

Keyset-driven cursor A scrollable cursor that detects updated and deleted rows by using a keyset.

L

Literal A character representation of an actual data value in an SQL statement.

Locking The process by which a DBMS restricts access to a row in a multiuser environment. The DBMS usually sets a bit on a row or the physical page containing a row that indicates the row or page is locked.

Long data Any binary or character data over a certain length, such as 255 bytes or characters. Typically much longer. Such data is usually sent to and retrieved from the data source in parts. Also known as *BLOBs* or *CLOBs*.

M

Machine data source A data source for which connection information is stored on the system (for example, the registry).

Manual-commit mode A transaction commit mode in which transactions must be explicitly committed by calling **SQLTransact**.

Metadata Data that describes a parameter in an SQL statement or a column in a result set. For example, the data type, byte length, and precision of a parameter.

Multiple-tier driver See DBMS-based driver.

N

NULL value Having no explicitly assigned value. In particular, a NULL value is different from a zero or a blank.

O

Octet Eight bits or one byte. See *also* byte.

Octet length The length in octets of a buffer or the data it contains.

ODBC Open Database Connectivity. A specification for an API that defines a standard set of routines with which an application can access data in a data source.

ODBC Administrator An executable program that calls the installer DLL to configure data sources.

ODBC SDK ODBC Software Development Kit. A product used to help develop ODBC applications and drivers.

Optimistic concurrency A strategy to increase concurrency in which rows are not locked. Instead, before rows are updated or deleted, a cursor checks to see if they have been changed since they were last read. If so, the update or delete fails. See *also* pessimistic concurrency.

Outer join A join in which both matching and nonmatching rows are returned. The values of all columns from the unmatched table in nonmatching rows are set to NULL.

Owner The owner of a table.

P

Parameter A variable in an SQL statement, marked with a parameter marker or question mark (?). Parameters are bound to application variables and their values retrieved when the statement is executed.

Parameter descriptor A descriptor that describes the run-time parameters used in an SQL statement, either before any conversion specified by the application (an application parameter descriptor, or APD) or after any conversion specified by the application (an implementation parameter descriptor, or IPD).

Parameter operation array An array containing values that an application can set to indicate that the corresponding parameter should be ignored in an **SQLExecDirect** or **SQLExecute** operation.

Parameter status array An array containing the status of a parameter after a call to **SQLExecDirect** or **SQLExecute**.

Pessimistic concurrency A strategy for implementing serializability in which rows are locked so that other transactions cannot change them. See *also* optimistic concurrency.

Positioned operation Any operation that acts on the current row—for example, positioned update and delete statements, **SQLGetData**, and **SQLSetPos**.

Positioned update statement An SQL statement used to update the values in the current row.

Positioned delete statement An SQL statement used to delete the current row.

Prepare To compile an SQL statement. You create an access plan by preparing an SQL statement.

Primary key A column or columns that uniquely identifies a row in a table.

Procedure A group of one or more precompiled SQL statements that are stored as a named object in a database.

Procedure column An argument in a procedure call, the value returned by a procedure, or a column in a result set created by a procedure.

Q

Qualifier A database that contains one or more tables.

Query An SQL statement. Sometimes used to mean a **SELECT** statement.

Quoted Identifier An identifier that is enclosed in identifier quote characters so it can contain special characters or match keywords (also known in SQL92 as a delimited identifier).

R

Radix The base of a number system. Usually 2 or 10.

Record See row.

Result set The set of rows created by executing a **SELECT** statement.

Return code The value returned by an ODBC function.

Roll back To return the values changed by a transaction to their original state.

Row A set of related columns that describe a specific entity. Also known as a *record*.

Row descriptor A descriptor that describes the columns of a result set, either before any conversion specified by the application (an implementation row descriptor, or IRD) or after any conversion specified by the application (an application row descriptor, or ARD).

Row operation array An array containing values that an application can set to indicate that the corresponding row should be ignored in a **SQLSetPos** operation.

Row status array An array containing the status of a row after a call to **SQLFetch**, **SQLFetchScroll**, or **SQLSetPos**.

Rowset The set of rows returned in a single fetch by a block cursor.

Rowset buffers The buffers bound to the columns of a result set and in which the data for an entire rowset is returned.

S

SAG SQL Access Group. An industry consortium of companies concerned with SQL DBMSs. The X/Open Call-Level Interface is based on work originally done by the SQL Access Group.

Scalar function A function that generates a single value from a single value—for example, a function that changes the case of character data.

Schema See catalog.

Scrollable cursor A cursor that can move forward or backward through the result set.

Serializability Whether two transactions executing simultaneously produce a result that is the same as the serial (or sequential) execution of those transactions. Serializable transactions are required to maintain database integrity.

Server database A DBMS designed to be run in a client/server environment. These DBMSs provide a standalone database engine that provides rich support for SQL and transactions. They are accessed through DBMS-based drivers. For example, Oracle, Informix, DB/2, or Microsoft SQL Server.

Set function See aggregate function.

Setup DLL See driver setup DLL *and* translator setup DLL.

Single-tier driver See file-based driver.

SQL Structured Query Language. A language used by relational databases to query, update, and manage data.

SQL conformance level The level of SQL92 grammar supported by a driver; can be Entry, FIPS

Transitional, Intermediate, or Full.

SQL data type The data type of a column or parameter as it is stored in the data source.

SQLSTATE A five-character value that indicates a particular error.

SQL statement A complete phrase in SQL that begins with a keyword and completely describes an action to be taken—for example, **SELECT * FROM Orders**. SQL statements should not be confused with statements.

State A well-defined condition of an item. For example, a connection has seven states, including unallocated, allocated, connected, and needing data. Certain operations can be done only when an item is in a particular state. For example, a connection can be freed only when it is in an allocated state and not, for example, when it is in a connected state.

State transition The movement of an item from one state to another. ODBC defines rigorous state transitions for environments, connections, and statements.

Statement A container for all the information related to an SQL statement. Statements should not be confused with SQL statements.

Statement handle A handle to a data structure that contains information about a statement.

Static cursor A scrollable cursor that cannot detect updates, deletes, or inserts in the result set. Usually implemented by making a copy of the result set.

Static SQL A type of embedded SQL in which SQL statements are hard-coded and compiled when the rest of the program is compiled. See *also* dynamic SQL.

Stored procedure See procedure.

T

Table A collection of rows.

Thunking The conversion of 16-bit addresses to 32-bit addresses, or vice versa, when 16-bit applications are used with 32-bit ODBC drivers.

Transaction An atomic unit of work. The work in a transaction must be completed as a whole; if any part of the transaction fails, the entire transaction fails.

Transaction isolation The act of isolating one transaction from the effects of all other transactions.

Transaction isolation level A measure of how well a transaction is isolated. There are five transaction isolation levels: Read Uncommitted, Read Committed, Repeatable Read, Serializable, and Versioning.

Translator DLL A DLL used to translate data from one character set to another.

Translator setup DLL A DLL that contains translator-specific installation and configuration functions.

Two-phase commit The process of committing a distributed transaction in two phases. In the first phase, the transaction processor verifies that all parts of the transaction can be committed. In the second phase, all parts of the transaction are committed. If any part of the transaction indicates in the first phase that it cannot be committed, the second phase does not occur. ODBC does not support two-phase commits.

Type indicator An integer value passed to or returned from an ODBC function to indicate the data type of an application variable, a parameter, or a column. ODBC defines type indicators for both C and SQL data types.

V

View An alternative way of looking at the data in one or more tables. A view is usually created as a subset of the columns from one or more tables. In ODBC, views are generally equivalent to tables.

X

X/Open A company that publishes standards. In particular, it publishes SAG standards.

