

Overview

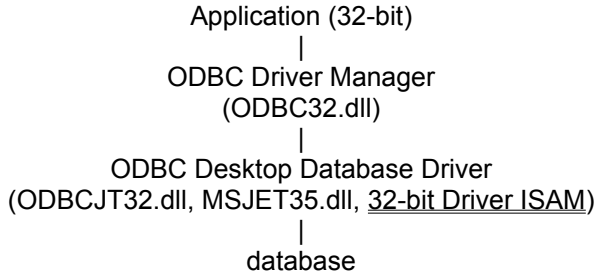
See Also

Microsoft ODBC Desktop Database Drivers 3.5 allow you to open and query a database through the Open Database Connectivity (ODBC) interface. These drivers work with ODBC 2.5, but are shipped with ODBC 3.x. The drivers are designed for use on Microsoft Windows 95 or later, or Windows NT 3.51 or later. Only 32-bit applications are supported on Windows 95 or later; 16-bit and 32-bit applications are supported on Windows NT 3.51 or later.

Note For information about the version of ODBC to be used with these drivers, refer to the ODBC 2.0 Programmer's Reference and SDK Guide, and release notes up to ODBC version 2.5. The ODBC 3.0 Programmer's Reference should not be used as reference material for issues with these drivers.

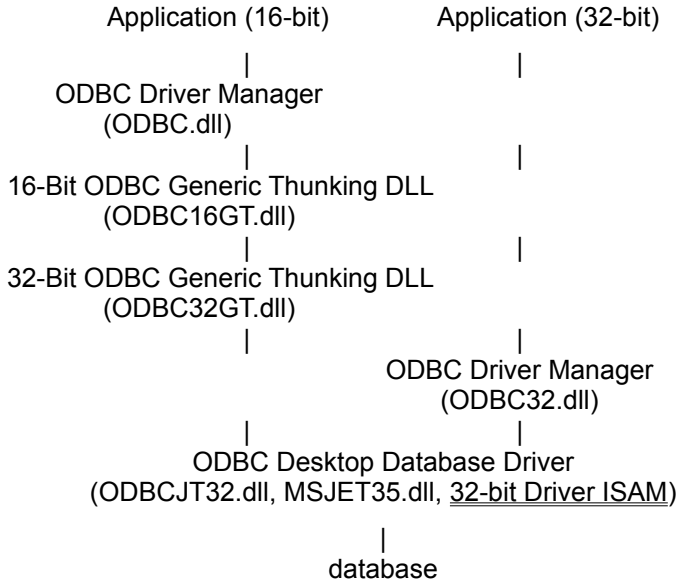
The ODBC Desktop Database Drivers include 32-bit drivers for Microsoft Access, dBASE, Microsoft Excel, Microsoft FoxPro, Paradox, and Text. No 16-bit drivers are included. (A driver for Microsoft FoxPro 3.0 is available separately.)

The application/driver architecture on Windows 95 or later is:



The use of these drivers by 16-bit applications on Windows 95 is not supported.

The application/driver architecture on Windows NT 3.51 or later is:



Driver ISAM Files

[See Also](#)

ODBC Desktop Database Drivers 3.5 include ISAM files. These files are as follows:

Driver	32-bit ISAM	Versions
Microsoft Access	None (Included in MSJET35.dll)	Microsoft Access versions 7.0 and 97
Microsoft Access dBase	MSRD2X35.dll MSXBSE35.dll	Microsoft Access versions 1.0, 1.1, 2.0 dBase versions III, IV, and V (dBASE for Windows)
Microsoft Excel	MSEXCL35.dll	Microsoft Excel versions 3.0, 4.0, 5.0, 7.0, 97 (1)
Microsoft FoxPro Paradox	MSXBSE35.dll MSPDOX35.dll	Microsoft FoxPro versions 2.0, 2.5, 2.6 Paradox versions 3.5, 4.0, 4.5, 5.x (Paradox for Windows)
Text	MSTEXT35.dll	ASCII Text

- 1 The Microsoft Excel driver is not able to read Microsoft Excel files whose format is earlier than version 3.0. These files can be opened in a later release of Microsoft Excel and saved in version 3.0, 4.0, 5.0, 7.0, or 97 format that the driver is able to read.

Microsoft FoxPro 3.0

The ODBC Desktop Database Drivers do not include a driver for Microsoft FoxPro 3.0, and the Desktop Database Drivers do not support access to Microsoft FoxPro 3.0 files. The Microsoft FoxPro 3.0 driver should be used for access to these files.

ODBC Component Files

See Also

The following table lists the files required by each component of ODBC Desktop Database Drivers 3.5. These files are installed in the \Windows\System directory for Windows 95 or later, or the \Windows\System32 directory for Windows NT 3.51 or later. If ODBC files were previously installed in a different directory, make sure that you use the newer files in the \Windows\System(or System32) directory.

Some files are required by more than one component. These files must be redistributed if you intend to redistribute any of the ODBC Desktop Database Drivers with your commercial application.

The following files are common for each of the ODBC Desktop Database Drivers 3.5:

DS16GT.dll	ODBCINST.hlp
DS32GT.dll	ODBCINT.dll
MSJINT35.dll	ODBCJET.hlp
MSJET35.dll	ODBCJET.cnt
MSJTER35.dll	ODBCJI32.dll
MSVCRT40.dll	ODBCJT32.dll
ODBC16GT.dll	ODBCJTNW.hlp
ODBC32.dll	ODBCJTNW.cnt
ODBC32GT.dll	ODBCTL32.dll
ODBCCP32.cpl	VBAJET32.dll
ODBCCP32.dll	VBAR32.dll
ODBCCR32.dll	ODBCTRAC.dll
ODBCINST.cnt	

The following files are unique to each driver:

Driver	File
Microsoft Access	MSRD2X35.dll
dBASE	MSXBSE35.dll ODDBSE32.dll
Microsoft Excel	MSEXCL35.dll ODEXL32.dll
Microsoft FoxPro	MSXBSE35.dll ODFOX32.dll
Paradox	MSPDOX35.dll ODPDX32.dll
Text	MSTEXT35.dll ODTEXT32.dll

Hardware and Software Requirements

Hardware Requirements

To use the ODBC Desktop Database Drivers, you must have:

- An IBM-compatible personal computer.
- A hard disk with 6 MB of free disk space.
- At least 16 MB of random-access memory (RAM).

Software Requirements

To access data with the ODBC driver, you must have:

- The ODBC driver.
- The 32-bit ODBC Driver Manager, version 3.0 or later (ODBC32.dll). Note that this file is included with the ODBC Desktop Database Drivers.
- Microsoft Windows 95 or later, or Windows NT 3.51 or later.
- The stack size for an application using a Microsoft ODBC driver should be at least 20K.

When using Microsoft Windows NT 3.51 or later, the 32-bit driver is thread-safe, but only through the use of a global semaphore that controls access to the driver. Concurrent use of the driver will be very limited under Windows NT. All access to the Jet ISAM layer will be single-threaded for all applications using the Microsoft Jet engine.

When running multiple 16-bit applications on Windows on Windows (WOW) on Microsoft Windows NT 3.51 or 4.0, the applications must be run in separate memory spaces. (The same memory space cannot be used because ODBC does not support multiple environments in the same process.) To run an application in a separate memory space, select the application's icon in the Microsoft Windows NT Program Manager, choose the File-Properties command, and then choose Run In Separate Memory Space.

The use of these drivers by 16-bit applications on Windows 95 is not supported.

Driver-Specific Hardware and Software Requirements

- The **Microsoft Access**, **dBASE**, and **Microsoft FoxPro** driver may require changes in the AUTOEXEC.bat or CONFIG.sys files.

CONFIG.sys and AUTOEXEC.bat Changes

If you are working with the **Microsoft Access** ODBC driver on Microsoft Windows 95 (not on Microsoft Windows NT) and do not have Microsoft Access installed on your computer, add the line "FILES= 50" to your CONFIG.sys file.

Setting Up Drivers

- 1 Double-click the Control Panel icon, and in the Control Panel window, double-click the ODBC icon (or open the ODBC Administrator directly).
- 2 Choose the **ODBC Drivers** button.

Drivers can no longer be added or deleted from the Control Panel or Administrator. The Add and Delete buttons have been removed from the **ODBC Drivers** tab. Adding or deleting drivers is only supported during installation of the ODBC Desktop Database Drivers.

Adding a Data Source

See Also

- 1 Double-click the Control Panel icon. In the Control Panel window, double-click the ODBC icon.
- 2 Choose the **User DSN**, **System DSN**, or **File DSN** tab.
- 3 Choose the **Add** button.
- 4 In the **Create New Data Source** dialog box, select a driver's name and follow the instructions of the wizard.

Modifying a Data Source

- 1 Double-click the Control Panel icon. In the Control Panel window, double-click the ODBC icon.
- 2 Select the **User DSN**, **System DSN**, or **File DSN** tab, depending on the type of the data source to be modified.
- 3 Select the data source from the list. If modifying a file DSN, find the appropriate directory in the **Look In** box (using the **Up** control if necessary), then select the data source from the directory.
- 4 Choose the **Configure** button, and set up the data source as necessary.

Deleting a Data Source

- 1 Double-click the Control Panel icon. In the Control Panel window, double-click the ODBC icon.
- 2 Select the **User DSN**, **System DSN**, or **File DSN** tab, depending on the type of the data source to be modified.
- 3 Select the data source from the list. If modifying a file DSN, find the appropriate directory in the **Look In** box (using the **Up** control if necessary), then select the data source from the directory.
- 4 Click the **Remove** button, and then click the **Yes** button to confirm the deletion.

Data Source Conversion

The 32-bit data sources used with the ODBC Desktop Database Drivers version 2.0 are automatically converted to new 32-bit data sources for version 3.5 drivers. No conversion is provided for 16-bit data sources. To change a 16-bit data source to a 32-bit data source, create a new 32-bit data source, then (optionally) delete the old 16-bit data source. A 32-bit data source and a 16-bit data source cannot share the same name.

Adding, Modifying, or Deleting a Data Source Dynamically

See Also

You can call the **SQLConfigDataSource** function to add, modify, or delete a data source dynamically. This function uses keywords to set connect options that when connecting through the Control Panel are set through the **Setup** dialog box. This function should be used when you want to add, modify, or delete a data source without displaying the **Setup** dialog box.

Product Support

Product support for ODBC is provided by Microsoft Product Support Services (PSS). Because many Microsoft products utilize ODBC as a core component, and also redistribute ODBC drivers, PSS considers ODBC to be a part of whatever product the user has that incorporates ODBC as a component. Support for ODBC is provided in accordance with the support agreement of that product. Each of these Microsoft products has its own support offering, as described in the documentation that comes with the product. Please refer to this documentation to determine which support options are available.

ODBC Setup Dialog Box

The **ODBC Setup** dialog box contains the following fields. Note that most of these options can also be set up dynamically by calling **SQLConfigDataSource** using keywords.

- >> Approximate Row Count
- >> Buffer Size
- >> Collating Sequence
- >> Data Source Name
- >> Database
- >> Define Format
- >> Description
- >> Directory
- >> Exclusive
- >> Extensions List
- >> ImplicitCommitSync
- >> Net Style
- >> Page Timeout
- >> Read Only
- >> Rows to Scan
- >> Select Directory
- >> Select Indexes
- >> Select Network Directory
- >> SelectWorkbook
- >> Show Deleted Rows
- >> Sort Order
- >> System Database
- >> Threads
- >> Use Current Directory
- >> UserCommitSync
- >> User Name
- >> Version

Approximate Row Count

(This control is displayed only for the **dBASE** and **Microsoft FoxPro** drivers.)

Determines whether table size statistics are approximated. Note that this option applies to all data sources that use the ODBC driver. To set this option dynamically, use the **STATISTICS** keyword in a call to **SQLConfigDataSource**.

Buffer Size

(This control is displayed only for the **Microsoft Access** driver.)

The size of the internal buffer, in kilobytes, that is used by Microsoft Access to transfer data to and from the disk. The default buffer size is 512K (displayed as 512). Any integer value divisible by 256 can be entered. To set this option dynamically, use the MAXBUFFERSIZE keyword in a call to **SQLConfigDataSource**.

Collating Sequence

(This control is displayed only for the **dBASE**, **Microsoft FoxPro**, and **Paradox** drivers.)

The sequence in which the fields are sorted.

When the **dBASE** or **Microsoft FoxPro** driver is used, the sequence can be: ASCII (the default) or International.

When the **Paradox** driver is used, the sequence can be: ASCII (the default), International, Swedish-Finnish, or Norwegian-Danish.

To set this option dynamically, use the **COLLATINGSEQUENCE** keyword in a call to **SQLConfigDataSource**.

Data Source Name

A name that identifies the data source, such as Payroll or Personnel.

To set this option dynamically, use the **DSN** keyword in a call to **SQLConfigDataSource**.

Database

See Also

(This control is displayed only for the **Microsoft Access** driver.)

The full path of the Microsoft Access database you want to access. Use the Select button to select an existing database. Use the Create button to create and select a new Microsoft Access database.

A Microsoft Access data source can be set up without selecting or creating a database. If no database is provided upon setup, the user will be prompted to choose a database file when connecting to the data source.

Use the Repair button to repair a damaged database. Use the Compact button to compact a database.

To set this option dynamically, use the **DBQ** keyword in a call to SQLConfigDataSource.

Define Format

[See Also](#)

(This control is displayed only for the **Text** driver.)

Displays the **Define Text Format** dialog box and enables you to specify the schema for individual tables in the data source directory.

This option cannot be set dynamically by a call to **SQLConfigDataSource**.

Description

An optional description of the data in the data source; for example, “Hire date, salary history, and current review of all employees.”

To set this option dynamically, use the **DESCRIPTION** keyword in a call to **SQLConfigDataSource**.

Directory

Displays the currently selected directory. Before you add the data source, you must either use the Select Directory button to select a directory, or select the Use Current Directory check box to use the application's current working directory. (This does **not** apply to the **Microsoft Access** driver.)

For **Microsoft Excel** 3.0/4.0 files, the path display is labeled "Directory", while for Microsoft Excel 5.0, 7.0, or 97 files, the path display is labeled "Workbook".

To set this option dynamically, use the **DEFAULTDIR** keyword in a call to SQLConfigDataSource.

Exclusive

(This control is displayed only for the **dBASE**, **Microsoft Access**, **Microsoft FoxPro**, and **Paradox** drivers.)

If the **Exclusive** box is selected, the database will be opened in Exclusive mode and can be accessed by only one user at a time. If the **Exclusive** box is cleared, the database will be opened in Shared mode and can be accessed by more than one user at a time. Performance is enhanced when running in Exclusive mode.

To set this option dynamically, use the **EXCLUSIVE** keyword in a call to **SQLConfigDataSource**.

Extensions List

(This control is displayed only for the **Text** driver.)

Lists the file name extensions of the text files on the data source. To use all files in the directory, select the **Default (*.*)** check box. To use only those files with certain extensions, clear the **Default (*.*)** check box and add each extension you want to use. When **Default (*.*)** is unchecked, the list of extensions defaults to *.asc, *.csv, *.tab, and *.txt.

To add an extension, type the extension in the **Extension** box and click the **Add** button. The extension must use the format *.xxx. For example, to use .dat files, type the extension *.dat. To remove an extension, select the extension in the **Extensions** list and click the Remove button. Removing all extensions is like checking the **Default (*.*)** check box: all files will be displayed.

When the **Text** driver is used, a file with no extension is created when the CREATE TABLE statement is executed with a name that has no extension. Other drivers create a file with a default extension when no extension is provided. To create a file with a .TXT extension, the extension must be included in the name. To display files without extensions in the Define Text Format dialog box, "*" must be added to the Extensions List.

To set this option dynamically, use the **EXTENSIONS** keyword in a call to [SQLConfigDataSource](#).

ImplicitCommitSync

(This control is displayed only for the **Microsoft Access** driver.)

Determines how changes made outside of a transaction are written to the database. This value is initially set to "Yes", which means that the **Microsoft Access** driver will wait for commits in an internal/implicit transaction to be completed.

This option is included in the Advanced Options dialog box for the **Microsoft Access** driver.

Net Style

(This control is displayed only for the **Paradox** driver.)

The network access style to use when accessing Paradox data: either "3.x" for Paradox 3.x or "4.x" for Paradox 4.x or 5.x. Can be set to "3.x" or "4.x" if the version is Paradox 4.x or 5.x; if the version is Paradox 3.x, the style must be "3.x".

To set this option dynamically, use the **PARADOXNETSTYLE** keyword in a call to **SQLConfigDataSource**.

Page Timeout

(This control is displayed only for the **dBASE**, **Microsoft Access**, **Microsoft FoxPro**, and **Paradox** drivers.)

Specifies the period of time, in tenths of a second, that a page (if not used) remains in the buffer before being removed. For the **Microsoft Access** driver, the default is five-tenths of a second (0.5 seconds). For the **dBASE**, **Microsoft FoxPro**, and **Paradox** drivers, the default is 600 tenths of a second (60 seconds). Note that this option applies to all data sources that use the ODBC driver.

The page timeout cannot be 0 because of an inherent delay. The page timeout cannot be less than the inherent delay, even if the page timeout option is set below that value.

To set this option dynamically, use the **PAGETIMEOUT** keyword in a call to **SQLConfigDataSource**.

Read-Only

Designates the database as read-only.

To set this option dynamically, use the **READONLY** keyword in a call to [SQLConfigDataSource](#).

Rows to Scan

(This control is displayed only for the **Microsoft Excel** and **Text** drivers.)

The number of rows to scan to determine the data type of each column. The data type is determined given the maximum number of kinds of data found. If data is encountered that does not match the data type guessed for the column, the data type will be returned as a NULL value.

For the **Microsoft Excel** driver, you may enter a number from 1 to 16 for the rows to scan. The value defaults to 8; if it is set to 0, all rows are scanned. (A number outside the limit will return an error.)

For the **Text** driver, you may enter a number from 1 to 32767 for the number of rows to scan; however, the value will always default to 25. (A number outside the limit will return an error.)

To set this option dynamically, use the **MAXSCANROWS** keyword in a call to **SQLConfigDataSource**.

Select Directory

(This control is displayed only for the **dBASE**, **Microsoft Excel 3.0/4.0**, **Microsoft FoxPro**, **Paradox**, and **Text** drivers.)

Displays a dialog box where you can select a directory containing the files you want to access.

When defining a data source directory (for all drivers except **Microsoft Access**), specify the directory where your most commonly used files are located. The ODBC driver uses this directory as the default directory. Copy other files into this directory if they are used frequently. Alternatively, you can qualify filenames in a SELECT statement with the directory name:

```
SELECT * FROM C:\MYDIR\EMP
```

Or, you can specify a new default directory by using the **SQLSetConnectOption** function with the **SQL_CURRENT_QUALIFIER** option.

When using the **Microsoft Access** driver, you can use the **DefaultDir** setting in the Advanced Option dialog box to define a default directory, since you can create a data source without a corresponding .mdb file.

Directory/Workbook Path Displays (Microsoft Excel)

For Microsoft Excel 3.0 or 4.0 files, the path display is labeled "Directory", and the path selection button is labeled "Select Directory". For Microsoft Excel 5.0, 7.0, or 97 files, the path display is labeled "Workbook", and the path selection button is labeled "Select Workbook". When defining a data source directory, specify the directory where your most commonly used Microsoft Excel files are located for Microsoft Excel 3.0/4.0, or the directory where the workbook file is located for Microsoft Excel 5.0, 7.0, or 97. The **Use Current Directory** control is disabled for Microsoft Excel 5.0, 7.0, or 97.

To set this option dynamically, use the **DEFAULTDIR** keyword in a call to **SQLConfigDataSource**.

Select Indexes

See Also

(This control is displayed only for the **dBASE** and **Microsoft FoxPro** drivers.)

Displays the **Select Indexes** dialog box where you can associate dBASE or Microsoft FoxPro files with index files. The **Select Indexes** dialog box contains the following fields.

Tables

Displays a list of the files in the currently selected data source.

Indexes

Displays the indexes assigned to the currently selected file in the Tables list.

List Files of Type

Displays and allows you to choose the types of files to display in the Indexes list.

To associate an index with a table

- 1 From the **Tables** list, select a file.
- 2 From the **Indexes** list, select an index.
- 3 Click **OK** to save the table/index associations.

Note dBASE III indexes or Microsoft FoxPro 2.0 indexes must be assigned using this dialog box for the driver to recognize them.

Note The ODBC dBASE and Microsoft FoxPro drivers do not support Clipper .NTX index files. This option cannot be set dynamically by a call to **SQLConfigDataSource**.

Select Network Directory

(This control is displayed only for the **Paradox** driver.)

The full path of the directory containing a Paradox lock database, because it contains either the PDOXUSRS.net file (in Paradox 4.x) or the PARADOX.net file (in Paradox 5.x). If the directory does not contain one of these files, the Paradox driver creates one. For information about these files, see the Paradox documentation.

Before you can select a network directory, you must enter your Paradox user name in the **User Name** text box. Use the **Select Network Directory** button to select a network directory.

To set this option dynamically, use the **PARADOXNETPATH** keyword in a call to **SQLConfigDataSource**.

Select Workbook

(This control is displayed only for **Microsoft Excel** 5.0, 7.0, or 97.)

For **Microsoft Excel** 5.0, 7.0, or 97 files, this option displays the currently selected workbook. Before you add the data source, you must use the **Select Workbook** button to select a directory.

Worksheets within workbooks are treated as SYSTEM TABLES. Defined or named ranges within a worksheet are treated as TABLES.

Show Deleted Rows

(This control is displayed only for the **dBASE** and **Microsoft FoxPro** drivers.)

Specifies whether or not rows that have been marked as deleted can be retrieved or positioned on. If unchecked, deleted rows are not displayed; if checked, deleted rows are treated the same as non-deleted rows. The default is unchecked.

To set this option dynamically, use the **DELETED** keyword in a call to **SQLConfigDataSource**.

Sort Order

(This control is displayed only for the **Microsoft Access** driver.)

Sets a default sort order for the database. Can be:

General (English, French, German, Portuguese, Italian), traditional Spanish, Dutch, Swedish/Finnish, Norwegian/Danish, Icelandic, Czech, Hungarian, Polish, Russian, Turkish, Arabic, Hebrew, Greek, Japanese, Korean, Taiwanese, PRC (People's Republic of China).

This option cannot be set dynamically by a call to **SQLConfigDataSource**.

System Database

(This control is displayed only for the **Microsoft Access** driver.)

The full path of the Microsoft Access system database to be used with the Microsoft Access database you want to access.

Select the **System Database** option button to select the system database to be used. The ODBC Microsoft Access driver prompts the user for a name and password. The default name is Admin and the default password in Microsoft Access for the Admin user is an empty string.

To increase the security of your Microsoft Access database, create a new user to replace the Admin user and delete the Admin user, or change the objects to which the Admin user has access.

Select the **None** option button if no system database will be used. Select the **None** option button to log into the Microsoft Access database as the Admin user; no system database will be used.

To set this option dynamically, use the **SYSTEMDB** keyword in a call to [SQLConfigDataSource](#).

For more information about system databases and the Admin user, see the Microsoft Access documentation.

Threads

(This control is displayed only for the **Microsoft Access** driver.)

The number of background threads for the engine to use. For the Microsoft Access driver, this value defaults to 3, but can be changed. The user may want to increase the number of threads if there is a large amount of activity in the database.

This option is included in the Advanced Options dialog box for the **Microsoft Access** driver.

To set this option dynamically, use the **THREADS** keyword in a call to [SQLConfigDataSource](#).

Use Current Directory

(This control is displayed only for the **dBASE**, **Microsoft Excel**, **Microsoft FoxPro**, **Paradox**, and **Text** drivers.)

When selected, makes the application's current working directory the data source directory and disables the **Select Directory** option. When cleared, enables you to select the data source directory using the **Select Directory** option.

The **Use Current Directory** control is disabled for **Microsoft Excel** 5.0, 7.0, or 97.

UserCommitSync

(This control is displayed only for the **Microsoft Access** driver.)

Determines whether the **Microsoft Access** driver will perform an explicit user-defined transactions asynchronously. This value is initially set to “Yes”, which means that the **Microsoft Access** driver will wait for commits in a user-defined transaction to be completed.

Setting this option to False can have unpredictable consequences in a multi-user environment.

This option is included in the Advanced Options dialog box for the **Microsoft Access** driver.

To set this option dynamically, use the **USERCOMMITSYNC** keyword in a call to **SQLConfigDataSource**.

User Name

(This control is displayed only for the **Paradox** driver.)

The Paradox user name. This is the name displayed to other users of Paradox files when a lock is encountered.

To set this option dynamically, use the **PARADOXUSERNAME** keyword in a call to **SQLConfigDataSource**.

Version

A pull-down control that enables you to select the version of the files used, for example, dBASE III, dBASE IV, or dBASE 5.0 for the dBASE driver.

Select Database

Enter the full path of an existing **Microsoft Access** database you want to access.

Create Database

Enter the full path of the new **Microsoft Access** database you want to create.

Defining Text Format

[See Also](#)

When the **Text** driver is used, you can use the **Define Text Format** dialog box to define the format for columns in a selected file. This dialog box enables you to specify the schema for each data table. This information is written to a [SCHEMA.ini](#) file in the data source directory. A separate SCHEMA.ini is created for each text data source directory.

Note The same default file format applies to all new text data tables. All files created by the CREATE TABLE statement inherit those same default format values, which are set by selecting file format values in the Define Text Format dialog box with <default> chosen in the Tables list box. The **Text** driver does not change the format of an existing text file to match the format defined in this dialog box, but returns an error when it uses the format, such as when it attempts to retrieve data from the text file.

Control	Information
Add	Adds a column using the values in the Data Type, Name, and Width fields from the dialog box, and if applicable, the Date Separator value from SCHEMA.ini.
Characters	ANSI or OEM. OEM specifies a non-ANSI character set. This defaults to OEM if the format of the item selected in the Tables list has not been previously defined by this dialog box
Column NameHeader	Indicates whether the columns of the first row of the selected table are to be used as column names. Either TRUE or FALSE. This defaults to FALSE if the format of the item selected in the Tables list has not been previously defined by this dialog box.
Columns	Contains a list of the column names for each column in the selected table. The order of the columns reflects the order of the columns in the table. This list is enabled if a file has been selected in the Tables list.
Data Type	Can be BIT, BYTE, CHAR, CURRENCY, DATE, FLOAT, INTEGER, LONGCHAR, SHORT, or SINGLE. Date data types can be in the following formats: "dd-mmm-yy", "mm-dd-yy", "mmm-dd-yy", "yyyy-mm-dd", or "yyyy-mmm-dd". "mm" denotes numbers for months; "mmm" denotes letters for months.
Delimiter	Specifies the custom delimiter character to be used to separate columns. Enabled when the "Custom Delimited" format is selected. The delimiter can only be one character in length, and double quotations marks (") cannot be used as the delimiter character. (Note that the delimiter cannot be specified in hexadecimal or decimal format.)
Format	Either delimited or fixed length. If delimited, indicates the type of delimiter used: comma (CSV), tab, or special character (custom). This defaults to CSV Delimited if the format of the item selected in the Tables list has not been previously defined by this dialog box. If Format is fixed-length and Column Name Header is TRUE, the first line must be comma-delimited.
Guess	Automatically generates the column's data type, name, and width values for the columns in the selected table by scanning the table's contents according to the Format list box selection. Enabled when the table format is delimited. Any previously

defined columns in the Columns list are cleared and replaced with new entries. If Column Name Header is not checked, column names are generated automatically as "F1", "F2", and so on. Note that no default value is shown in the Data Type box.

This functionality only works on columns that are less than 64,513 bytes.

Modify	Modifies the selected column using the values in the Data Type, Name, and Width fields.
Name	Displays the name of the selected column. May be used to specify a new column name for either an existing column or a new column. If Column Name Header is TRUE, the column name displayed is ignored.
Remove	Deletes the selected column.
Rows to Scan	The number of rows that Setup or the driver will scan when setting the columns and column data types based upon existing data. You may enter a number from 1 to 32767 for the number of rows to scan. This defaults to 25 if the format of the item selected in the Tables list has not been previously defined by this dialog box. (A number outside the limit will return an error.)
Tables	Contains a list of all files in the directory selected in the Text Setup dialog box that match the list of extensions specified. When <default> is selected, and one of the following is true, then the values of the table attributes in the Tables group are written to schema.ini. No other entries in schema.ini are touched. <ul style="list-style-type: none">• There is no schema.ini in the specified directory.• The schema.ini file exists, but there is no section in schema.ini for one of the Text files (with the specified extension) in the directory.• The section for a Text file exists in schema.ini, but the body is empty. When <default> is selected, the "Columns" group is disabled.
Width	The width of the column may be changed for CHAR or LONGCHAR columns. The width defaults to 1 if the format of the item selected in the Tables list has not been previously defined by this dialog box. For other data types, the width control is disabled, and no value is displayed.

Schema

The schema includes information about each table (text file) in a data source, including the table's format, the number of rows to scan to determine column types, whether the first row of the table contains column names, whether the source file is written using an OEM or ANSI codepage, and each column's name, data type, and width.

SCHEMA.ini File

[See Also](#)

When the **Text** driver is used, the format of the text file is determined by using a schema information file. The schema information file, which is always named Schema.ini and always kept in the same directory as the text data source, provides the IISAM with information about the general format of the file, the column name and data type information, and a number of other data characteristics. A Schema.ini file is always required for accessing fixed-length data; you should use a Schema.ini file when your text table contains DateTime, Currency, or Decimal data or any time you want more control over the handling of the data in the table.

Note The Text ISAM will obtain initial values from the registry, not from SCHEMA.ini. The same default file format applies to all new text data tables. All files created by the CREATE TABLE statement inherit those same default format values, which are set by selecting file format values in the **Define Text Format** dialog box with <default> chosen in the Tables list box. If the values in the registry are different from the values in SCHEMA.ini, the values in the registry will be overwritten by the values from SCHEMA.ini.

Understanding Schema.ini Files

Schema.ini files provide schema information about the records in a text file. Each Schema.ini entry specifies one of five characteristics of the table:

- The text file name
- The file format
- The field names, widths, and types
- The character set
- Special data type conversions

The following sections discuss these characteristics.

Specifying the File Name

The first entry in Schema.ini is always the name of the text source file enclosed in square brackets. The following example illustrates the entry for the file Sample.txt:

```
[Sample.txt]
```

Specifying the File Format

The **Format** option in Schema.ini specifies the format of the text file. The Text IISAM can read the format automatically from most character-delimited files. You can use any single character as a delimiter in the file except the double quotation mark ("). The **Format** setting in Schema.ini overrides the setting in the Windows Registry on a file-by-file basis. The following table lists the valid values for the **Format** option.

Format specifier	Table format
TabDelimited	Fields in the file are delimited by tabs.
CSVDelimited	Fields in the file are delimited by commas (comma-separated values).
Delimited(*)	Fields in the file are delimited by asterisks. You can substitute any character for the asterisk except the double quotation mark.
FixedLength	Fields in the file are of a fixed-length.

For example, to specify a comma-delimited format, you would add the following line to Schema.ini:

Format=CSVDelimited

Specifying the Fields

You can specify field names in a character-delimited text file in two ways:

- Include the field names in the first row of the table and set **ColNameHeader** to **True**.
- Specify each column by number and designate the column name and data type.

You must specify each column by number and designate the column name, data type, and width for fixed-length files.

Note The **ColNameHeader** setting in Schema.ini overrides the **FirstRowHasNames** setting in the Windows Registry on a file-by-file basis.

The data types of the fields can also be determined. Use the **MaxScanRows** option to indicate how many rows should be scanned when determining the column types. If you set **MaxScanRows** to 0, the entire file is scanned. The **MaxScanRows** setting in Schema.ini overrides the setting in the Windows Registry on a file-by-file basis.

The following entry indicates that Microsoft Jet should use the data in the first row of the table to determine field names and should examine the entire file to determine the data types used:

```
ColNameHeader=True  
MaxScanRows=0
```

The next entry designates fields in a table by using the column number (**Coln**) option, which is optional for character-delimited files and required for fixed-length files. The example shows the Schema.ini entries for two fields, a 10-character CustomerNumber text field and a 30-character CustomerName text field:

```
Col1=CustomerNumber Text Width 10  
Col2=CustomerName Text Width 30
```

The syntax of **Coln** is:

Coln=*ColumnName type [Width #]*

The following table describes each part of the **Coln** entry.

Parameter	Description
<i>ColumnName</i>	The text name of the column. If the column name contains embedded spaces, you must enclose it in double quotation marks.
<i>type</i>	Data types are: Microsoft Jet data types Bit Byte Short Long Currency Single Double DateTime Text Memo ODBC data types Char (same as Text) Float (same as Double) Integer (same as Short)

	LongChar (same as Memo)
	Date <i>date format</i>
Width	The literal string value <code>Width</code> . Indicates that the following number designates the width of the column (optional for character-delimited files, required for fixed-length files).
#	The integer value that designates the width of the column (required if Width is specified).

Selecting a Character Set

You can select from two character sets: ANSI and OEM. The following example shows the Schema.ini entry for an OEM character set. The **CharacterSet** setting in Schema.ini overrides the setting in the Windows Registry on a file-by-file basis. The following example shows the Schema.ini entry that sets the character set to ANSI:

```
CharacterSet=ANSI
```

Specifying Data Type Formats and Conversions

The Schema.ini file contains a number of options that you can use to specify how data is converted or displayed. The following table lists each of these options.

Option	Description
DateTimeFormat	Can be set to a format string indicating dates and times. You should specify this entry if all date/time fields in the import/export are handled with the same format. All of the Microsoft Jet formats except A.M. and P.M. are supported. In the absence of a format string, the Windows Control Panel short date picture and time options are used.
DecimalSymbol	Can be set to any single character that is used to separate the integer from the fractional part of a number.
NumberDigits	Indicates the number of decimal digits in the fractional portion of a number.
NumberLeadingZeros	Specifies whether a decimal value less than 1 and greater than -1 should contain leading zeros; this value can either be False (no leading zeros) or True .
CurrencySymbol	Indicates the currency symbol to be used for currency values in the text file. Examples include the dollar sign (\$) and Dm.
CurrencyPosFormat	Can be set to any of the following values: <ul style="list-style-type: none"> • Currency symbol prefix with no separation (\$1) • Currency symbol suffix with no separation (1\$) • Currency symbol prefix with one character separation (\$ 1) • Currency symbol suffix with one

character separation (1 \$)

CurrencyDigits

Specifies the number of digits used for the fractional part of a currency amount.

CurrencyNegFormat

Can be one of the following values:

- (\$1)
- -\$1
- \$-1
- \$1-
- (1\$)
- -1\$
- 1-\$
- 1\$-
- -1 \$
- -\$ 1
- 1 \$-
- \$ 1-
- \$ -1
- 1- \$
- (\$ 1)
- (1 \$)

This example shows the dollar sign, but you should replace it with the appropriate **CurrencySymbol** value in the actual program.

CurrencyThousandSymbol

Indicates the single-character symbol to be used for separating currency values in the text file by thousands.

CurrencyDecimalSymbol

Can be set to any single character that is used to separate the whole from the fractional part of a currency amount.

Note If you omit an entry, the default value in the Windows Control Panel is used.

Text File Format (Advanced)

[See Also](#)

The ODBC **Text** driver supports both delimited and fixed-width text files. A text file consists of an optional header line and zero or more text lines.

Although the header line uses the same format as the other lines in the text file, the ODBC text driver interprets the header line entries as column names, not data.

A delimited text line contains one or more data values separated by delimiters: commas, tabs, or a custom delimiter. The same delimiter must be used throughout the file. Null data values are denoted by two delimiters in a row with no data between them. Character strings in a delimited text line can be enclosed in double quotation marks (""). No blanks may occur before or after delimited values.

The width of each data entry in a fixed-width text line is specified in a [schema](#). Null data values are denoted by blanks.

Tables are limited to a maximum of 255 fields. Field names are limited to 64 characters, and field widths are limited to 32,766 characters. Records are limited to 65,000 bytes.

A text file can be opened only for a single user. Multiple users are not supported.

The following grammar, written for programmers, defines the format of a text file that can be read by the ODBC text driver. Non-italics represent characters that must be entered as shown, italics represent arguments that are defined elsewhere in the grammar, brackets ([]) represent optional items, braces ({}), braces ({}), vertical bars (|) separate these choices, and ellipses (...) represent items that can be repeated one or more times.

The format of a text file is:

```
text-file ::=  
    [delimited-header-line] [delimited-text-line]... end-of-file |  
    [fixed-width-header-line] [fixed-width-text-line]... end-of-file
```

```
delimited-header-line ::= delimited-text-line
```

```
delimited-text-line ::=  
    blank-line |  
    delimited-data [delimiter delimited-data]... end-of-line
```

```
fixed-width-header-line ::= fixed-width-text-line
```

```
fixed-width-text-line ::=  
    blank-line |  
    fixed-width-data [fixed-width-data]... end-of-line
```

```
end-of-file ::= <EOF>
```

```
blank-line ::= end-of-line
```

```
delimited-data ::= delimited-string | number | date | delimited-null
```

```
fixed-width-data ::= fixed-width-string | number | date | fixed-width-null
```

The width of each column in a fixed width text file is specified in the SCHEMA.INI file.

```
end-of-line ::= <CR> | <LF> | <CR><LF>
```

```
delimited-string ::= unquoted-string | quoted-string
```

```
unquoted-string ::= [character | digit] [character | digit | quote-character]...
```

```
quoted-string ::=  
    quote-character  
    [character | digit | delimiter | end-of-line | embedded-quoted-string]...  
    quote-character
```

embedded-quoted-string ::=
 quote-character quote-character
 [*character | digit | delimiter | end-of-line*]
 quote-character quote-character

fixed-width-string ::= [*character | digit | delimiter | quote-character*] ...

character ::= any character except:
 delimiter
 digit
 end-of-file
 end-of-line
 quote-character

digit ::= 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9

delimiter ::= , | <TAB> | *custom-delimiter*

custom-delimiter ::= any character except:
 end-of-file
 end-of-line
 quote-character

The delimiter in a custom-delimited text file is specified in the SCHEMA.INI file.

quote-character ::= "

number ::= *exact-number* | *approximate-number*

exact-number ::= [+ | -] {*unsigned-integer*[.*unsigned-integer*] |
 unsigned-integer. |
 .*unsigned-integer*}

approximate-number ::= *exact-number*{e | E}[+ | -]*unsigned-integer*

unsigned-integer ::= {*digit*}...

date ::=
 mm date-separator dd date-separator yy |
 mmm date-separator dd date-separator yy |
 dd date-separator mmm date-separator yy |
 yyyy date-separator mm date-separator dd |
 yyyy date-separator mmm date-separator dd

mm ::= *digit* [*digit*]

dd ::= *digit* [*digit*]

yy ::= *digit digit*

yyyy ::= *digit digit digit digit*

mmm ::= Jan | Feb | Mar | Apr | May | Jun | Jul | Aug | Sep | Oct | Nov | Dec

date-separator ::= - | / | .

delimited-null ::=

For delimited files, a NULL is represented by no data between two delimiters.

fixed-width-null ::= <SPACE>...

For fixed width files, a NULL is represented by spaces.

Select Unique Record Identifier

Select the column or columns that uniquely identify a record for a table.

Data Source

When the **Microsoft Access** driver is used, the data source specifies the Microsoft Access database (.mdb) you want to use and the information needed to get to that database.

When the **dBASE, Microsoft Excel 3.0/4.0, Microsoft FoxPro, Paradox, or Text** driver is used, the data source specifies the default data directory in which the ODBC driver searches for data files you want to access, as well as other information.

For **Microsoft Excel 5.0, 7.0, or 97**, the data source specifies the default data workbook in which the ODBC driver searches for data files you want to access, as well as other information.

Password Entry

[See Also](#)

Passwords must be entered at connection time, either by password in the connection string (UID and PWD for **Microsoft Access** or PWD for **Paradox**), or with the *szAuthStr* argument in the **SQLConnect** function.

When the **Microsoft Access** driver is used, if the data source uses a system database, an application may prompt you for your name and password. The default name is Admin and the default password in Microsoft Access for the Admin user is an empty string.

When the **Paradox** driver is used, a password is only supported if entered in the connection string. That password is valid whenever a table is opened. If no password is passed in the connection string, no password is established for a table. If tables have different passwords, they cannot both be opened in the same session, nor can the tables be joined.

Heterogeneous Joins/Attached Tables

When the **Microsoft Access** driver is used, attached tables and heterogeneous joins are supported.

Creating links to attached tables is not supported. The **Microsoft Access** driver should not be used to attach a table to another table from within the Microsoft Access database management system. The table should be attached directly.

Note that the use of the ODBC Desktop Database Drivers from within Data Access Objects (DAO) code is not supported.

Opening Microsoft Excel Tables

The **Microsoft Excel** driver is not able to read Microsoft Excel files whose format is earlier than version 3.0. These files can be opened in a later release of Microsoft Excel and saved in version 3.0, 4.0, 5.0, 7.0, or 97 format that the driver is able to read.

A BIFF3-version table is created if the version of the data source through which the connection was made is "Excel", which means either Microsoft Excel 3.0 or 4.0. A worksheet is created in the workbook that is connected to if the version of the data source was "Excel 5.0", "Excel 7.0", or "Excel 97".

Microsoft Excel 3.0 and Microsoft Excel 4.0 file formats must have a database range defined. The Microsoft Excel driver recognizes named ranges as tables. Databases specified by a named range will increase as new data is added. If any new data would overwrite existing data in a worksheet (outside the named range that represents the data), an error will be returned.

All Microsoft Excel tables (spreadsheets) that are created and opened for inserting are opened exclusive by default, and can be opened by only one user at a time. The user must explicitly choose to open Microsoft Excel tables as shared. "Read Only" allows multiple users to view .xls files.

Password-Protected Files

When the **Paradox** driver is used, after a password-protected file has been opened by a user, other users are not allowed to open the same file.

Failing after Repeated Connections

If an ODBC application working on Microsoft Windows 95 connects to and disconnects from the server repeatedly (over 50-60 times), the connection may fail with the driver returning SQLSTATE 01000 (General warning) and the error message "Failed to get the expression service". This may particularly be a problem for Internet web server applications. This failure is resolved by updating the rpctr4.dll file in the \\WINDOWS\\SYSTEM directory to release QFE 324. For more information, contact Microsoft Product Support Services.

Referential Integrity

Referential integrity is supported only for a Microsoft Access data source, and only if the referential integrity is set up for the data source by the Microsoft Access database management system.

Lotus 1-2-3 Data/EMS Data

The ODBC Desktop Database Drivers do not support Lotus 1-2-3 data or EMS data.

Visual Basic Support

The **Connect** and **Attributes** properties of a **TableDef** object in Visual Basic can be used to connect to an ODBC data source using the ODBC Desktop Database Drivers. For more information, see the Visual Basic documentation.

Connection Strings (Advanced)

[See Also](#)

The connection string for the ODBC drivers uses the following keywords:

Keyword	Description
DSN	Name of the data source
DBQ	Name of the directory
DRIVERID	An integer ID for the driver.
FIL	File type (MS Access for Microsoft Access; dBase III, dBase IV, or dBase5; Excel 3.0, Excel 4.0, Excel 5.0, Excel 7.0, or Excel 97 for Microsoft Excel; FoxPro 2.0, 2.5, or 2.6 for Microsoft FoxPro; Paradox 3.x, 4.x, or 5.x; or Text) This keyword is supported only for backward compatibility.

Driver-Specific Keywords

In addition to the keywords shown in the previous table, the **Microsoft Access** driver uses UID and PWD keywords and the **Paradox** driver uses the PWD keyword.

Keyword	Description
UID	The user login ID (Microsoft Access only). Initially set to "Admin". For more information on the Admin user ID, see "Creating a New Admin User" in Chapter 10, "Managing Security," of the Microsoft Jet Database Engine Programmer's Guide.
PWD	The user-specified password.

Other keywords are supported for use in connection strings, in addition to the ones documented above. For more information on these keys, see [SQLConfigDataSource](#). The **Microsoft Access**-specific keywords compacting a database (COMPACT_DB) and creating a database (CREATE_DB) are not supported in connection strings.

Note When the **Microsoft Access** or **Paradox** driver is used, the PWD keyword should not include any of the special characters (see SQL_SPECIAL_CHARACTERS in [SQLGetInfo Returned Values](#)).

Select Database (Advanced)

When the **Microsoft Access** driver is used, the **Select** button may be selected in the Microsoft Access Setup dialog box to select an existing database.

The **Open Database** dialog box contains the following fields:

Control	Description
Database Name	Name of the existing database file with an .mdb extension.
Exclusive	If the Exclusive box is selected, the database will be opened in Exclusive mode and can be accessed by only one user at a time. If the Exclusive box is cleared, the database will be opened in Shared mode and can be accessed by more than one user at a time. Performance is enhanced when running in Exclusive mode.
Read Only	Designates the database as read-only.

Create Database (Advanced)

When the Microsoft Access driver is used, the Create button may be pressed to create a new database.

Note A database created using the Create button will work only with Microsoft Access 7.0.

The **New Database** dialog box contains the following fields:

Control	Description
Database Name	Name of the database file with an .mdb extension.
Old format (2.x)	Creates an ODBC 2.x-compatible database.
Sort Order	Sets a default sort order for the database. Can be: General (English, French, German, Portuguese, Italian), traditional Spanish, Dutch, Swedish/Finnish, Norwegian/Danish, Icelandic, Czech, Hungarian, Polish, Russian, Turkish, Arabic, Hebrew, Greek, Japanese, Korean, Taiwanese, PRC (People's Republic of China).
System Database	Creates a system database.

Database Repair (Advanced)

When the **Microsoft Access** driver is used, if a client failure occurs during the commit process (after **SQLTransact** has been issued, but before the function returns), the affected database can be corrupted and need to be repaired. Repair can be performed either upon setup through the Microsoft Access driver setup dialog box, or programmatically by using the **SQLConfigDataSource** function.

To repair the database through the Microsoft Access driver setup dialog box, click the Repair button. Select the file and path, and then click OK.

To repair the database programmatically, include "REPAIR_DB=<path name>" in the list of attributes (*pszAttributes*) of the **SQLConfigDataSource** function. In this function, <path name> is the full path to the Microsoft Access database.

When repairing a password-protected file, a password and user id must be entered in the advanced options dialog box (or in the list of attributes for **SQLConfigDataSource**, if repairing programmatically). The password and user id should be removed after the repair.

These operations will run the standard Microsoft Access repair process on the database. Data in Microsoft Access databases could be affected by either of these actions.

Note that a database indicated as corrupted may in fact not be corrupted. The indication may be prompted by an improper shutdown that does not clear the appropriate bits. In this and other cases, a programmer may want to keep the corruption message box transparent to the end user. This can be accomplished by passing `SQL_DRIVER_NOPROMPT` with the **SQLDriverConnect** function.

For more information about database repair in Microsoft Access, refer to the Microsoft Access help file.

Database Compaction (Advanced)

When the **Microsoft Access** driver is used, data compaction can be performed on a database file either upon setup through the Microsoft Access driver setup dialog box, or programmatically by using the **SQLConfigDataSource** function.

To perform data compaction on a database file through the Microsoft Access driver setup dialog box, click the **Compact** button. From the **Database to Compact From** dialog box, select the file and path, and then click **OK**. From the **Database to Compact Into** dialog box, select the file and path, select the version of the resulting compacted file and a sort order (which can be changed for the destination database), and then click **OK**. The name of the file to compact into can be the same as the file to be compacted. An error is returned if the name of the file to compact into is the same as an existing file.

To perform data compaction on a database file programmatically, include "COMPACT_DB=<source db> <dest db> <sort order>" in the list of attributes (*pszAttributes*) of the **SQLConfigDataSource** function. In this function, <source db> is the full path to the Microsoft Access database to be compacted and <dest db> is the full path for the compacted Microsoft Access database.

When compacting a password-protected file, a password and user id must be entered in the advanced options dialog box (or in the list of attributes for **SQLConfigDataSource**, if compacting programmatically). The password and user id should be removed after the compaction. Compacting password-protected files may have unpredicted results.

These operations will run the standard Microsoft Access compaction process on the database.

For more information about database compaction in Microsoft Access, refer to the Microsoft Access Help file.

SQL-92 Compliance (Advanced)

The ODBC Desktop Database Drivers and the underlying Microsoft Jet engine are not SQL-92 compliant. They support many features that have been defined in SQL-92. Some features supported in the driver are not supported in SQL-92. For more information, see the Microsoft Jet Database Engine Programmers's Guide. The following are the major differences between the two:

- The SQL used by the Desktop Database Drivers supports more powerful expressions than those specified by SQL-92.
- Different rules apply to the BETWEEN predicate.
- The SQL used by the Desktop Database Drivers and ANSI SQL supports different keywords.

The following SQL-92 features are not supported by Microsoft Jet SQL:

- Security statements, such as GRANT and LOCK.
- DISTINCT with aggregate function references.

The following features are enhancements in the SQL used by the Desktop Database Drivers that are not specified by SQL-92:

- The TRANSFORM statement providing support for crosstab queries.
- Additional aggregate functions (**StDev** and **VarP**).

Additional Supported ODBC SQL Grammar (Advanced)

The ODBC drivers completely support the Minimum ODBC 2.5 grammar and the following clauses in the Core and Extended ODBC 2.5 grammar. Additional information may be displayed by jumping to the ALTER TABLE statement, correlation names, and procedure invocation topics.

- Approximate numeric literal (Core)
- ALTER TABLE statement (Core)
- Binary literals (Extended)
- BETWEEN predicate (Core)
- Correlation names (Core)
- EXISTS predicate (Core)
- IN (*valuelist*) (Core)
- Outer joins (Extended)
- Procedure invocation (Extended)
- SELECT statement (Extended) (including GROUP BY, HAVING, and UNION clauses)
- Subqueries (SUBSELECTs) (Core)
- Topic

ALTER TABLE Statement (Core)

[See Also](#)

When the **Microsoft Access**, **dBASE**, **Microsoft FoxPro**, or **Paradox** driver is used, the ALTER TABLE statement is supported with add and drop table constraint definition clauses. (ALTER TABLE statements are not supported for the **Microsoft Excel** or **Text** driver.)

Correlation Names (Core)

Correlation names are fully supported, including within the table list. For example, in the following string, E1 is the correlation name for the table named Emp:

```
SELECT * FROM Emp E1  
WHERE E1.LastName = 'Smith'
```

Procedure Invocation (Extended)

When the **Microsoft Access** driver is used, procedures can be invoked from the driver by using the **SQLExecDirect** or **SQLPrepare** function with the following syntax: {CALL *procedure-name* [(*parameter*[,*parameter*]...)}]. Note that expressions are not supported as parameters to a called procedure.

If a procedure name includes a dash, the name must be delimited with backquotes (`).

A parameterized query can be called using the previous statement.

BETWEEN Predicate

The syntax:

expression1 BETWEEN *expression2* AND *expression3*

returns true only if *expression1* is greater than or equal to *expression2* and *expression1* is less than or equal to *expression3*.

The semantics of this syntax are different for the Desktop Database Drivers and the Microsoft Jet engine. In Microsoft Jet SQL, *expression2* can be greater than *expression3* so that the statement will return TRUE only if *expression1* is greater than or equal to *expression3*, and *expression1* is less than or equal to *expression2*.

CREATE INDEX Statement

[See Also](#)

The syntax of the CREATE INDEX statement is:

```
CREATE [UNIQUE] INDEX index-name ON table-name (column-identifier [ASC][DESC][, column-identifier [ASC][DESC]...]) WITH <index option list>
```

where <*index option list*> can be: PRIMARY | DISALLOW NULL | IGNORE NULL

Only the **Microsoft Access** driver uses the DISALLOW NULL and IGNORE NULL index options. The **dBASE**, **Microsoft FoxPro**, and **Paradox** drivers accept the syntax, but ignore the presence of either option.

When the **Paradox** driver is used, the CREATE INDEX statement creates Paradox primary key files and secondary files.

This statement is not supported by the Microsoft Excel or Text drivers.

Date Arithmetic

The driver supports adding and subtracting an integer from a DATE, TIME, or TIMESTAMP column. For a DATE column, the integer specifies the number of days to add or subtract. For a TIME or TIMESTAMP column, the integer specified the number of seconds to add or subtract.

Date, Time, and Timestamp Literals

For maximum interoperability, applications should pass date literals in the ODBC canonical format using escape-clause syntax:

- For date literals, {d '*value*'}, where *value* is in the form “yyyy-mm-dd”
- For time literals, {t '*value*'}, where *value* is in the form “hh:mm:ss”
- For timestamp literals {ts '*value*'}, where *value* is in the form “yyyy-mm-dd hh:mm:ss[.f...]”.

DROP INDEX Statement

When the **Microsoft Access**, **dBASE**, **Microsoft FoxPro**, or **Paradox** driver is used, the syntax of the DROP INDEX statement is "DROP INDEX a on b" where "a" is the name of the index and "b" is the name of the table (not DROP INDEX *index-name*).

When the **Paradox** driver is used, the DROP INDEX statement deletes Paradox secondary index files.

The DROP INDEX statement is not supported for the **Microsoft Excel** or **Text** driver.

Fixed-Width Text File

When the **Text** driver is used, the last column of a fixed-width text file can be variable length.

GROUP BY *expression-list*

GROUP BY supports an expression list as well as a column name. The select list of a SELECT statement that has a GROUP BY clause can only include expressions from the GROUP BY clause or set functions.

ORDER BY *expression-list*

Expressions can be used in the ORDER BY clause. For example, in the following clauses the table is ordered by three key expressions: a+b, c+d, and e.

```
SELECT * FROM emp  
ORDER BY a+b,c+d,e
```

No ordering is allowed on set functions or an expression that contains a set function.

ORDER BY with GROUP BY

ORDER BY can be performed on any expression in the GROUP BY *expression-list* or any column in the result set.

Outer Joins

A `SELECT` statement can contain a list of `OUTER JOIN` clauses. Nested `OUTER JOINS` are supported. For more information on `OUTER JOINS`, see the Microsoft Jet Database Engine Programmers's Guide.

Scalar Functions

The Microsoft ODBC Desktop Database drivers support the following scalar functions:

CONCAT	LCASE	RIGHT
CONVERT	LEFT	RTRIM
CURDATE	LENGTH	SUBSTRING
CURTIME	LOCATE	UCASE
DATABASE	LTRIM	USER
DAYOFMONTH	MOD	YEAR
DAYOFWEEK	MONTH	

For information about the arguments and return values of scalar functions, see Appendix F of the *Microsoft ODBC SDK Programmer's Reference*.

Table Names

When the **dBASE**, **Microsoft Excel**, **Microsoft FoxPro**, **Paradox**, or **Text** driver is used, table names that occur in the FROM clause of SELECT or DELETE, after the INTO clause in INSERT, and after UPDATE, CREATE TABLE, and DROP TABLE can contain a valid path, primary name, and file name extension.

Use of a table name elsewhere in an SQL statement does not support the use of paths or extensions but will accept only the primary name (for example, EMP FROM C:\ABC\EMP).

Correlation names (aliases) can be used. For example:

```
SELECT *  
FROM C:\ABC\EMP T1  
WHERE T1.COL1 = 'aaa'
```

Aggregate Function Limitations

An aggregate function and a non-aggregate column reference cannot both be used as arguments to a single SQL statement.

ALTER TABLE Statement Limitations

See Also

When the **dBASE**, **Microsoft FoxPro**, or **Paradox** driver is used, once an index has been created and a new record added, the structure of the table cannot be changed by the ALTER TABLE statement unless the index is dropped and the contents of the table are deleted.

ALTER TABLE statements are not supported for the **Microsoft Excel** or **Text** driver.

AND Predicate Limitations

A maximum of 40 is supported.

CALL Statement Limitations

(This limitation applies to the **Microsoft Access** driver only.)

Expressions are not supported as parameters to a called procedure.

Column Name Limitations

Column names can contain any valid characters (for example, spaces). If column names contain any characters except letters, numbers, and underscores, the name must be delimited by enclosing it in back quotes (`).`).

When the **Microsoft Access** or **Microsoft Excel** driver is used, column names are limited to 64 characters, and longer names generate an error. When the **Paradox** driver is used, the maximum column name is 25 characters. When the **Text** driver is used, the maximum column name is 64 characters, and longer names are truncated.

When the **dBASE** or **Microsoft FoxPro** driver is used, characters with an ASCII value greater than 127 are converted to underscores.

When the **Microsoft Excel** driver is used, if column names are present, they must be in the first row. A name that in Microsoft Excel would use the "!" character must be enclosed in back quotes (""). The "!" character is converted to the "\$" character, because the "!" character is not legal in an ODBC name, even when the name is enclosed in back quotes. All other valid Microsoft Excel characters (except the pipe character (|)) can be used in a column name, including spaces. A delimited identifier must be used for a Microsoft Excel column name to include a space. Unspecified column names will be replaced with driver-generated names, for example, "Col1" for the first column.

The pipe character (|) cannot be used in a column name, whether the name is enclosed in back quotes or not.

When the **Text** driver is used, the driver provides a default name if a column name is not specified. For example, the driver calls the first column F1, the second column F2, and so on.

CONVERT Function Limitations

Type conversion failures result in the affected column being set to NULL.

Neither the DATE nor TIMESTAMP data type can be converted to another data type (or itself) by the CONVERT function.

CREATE INDEX Statement Limitations

See Also

The CREATE INDEX statement is not supported for the **Microsoft Excel** or **Text** driver.

An index can be defined on a maximum of 10 columns. If more than 10 columns are included in a CREATE INDEX statement, the index will not be recognized and the table will be treated as though no index were created.

The **dBASE** driver cannot create an index on a LOGICAL column.

When the **dBASE** or **Microsoft FoxPro** driver is used, response time on large files can be improved by building an .MDX (or .NDX) index on the column (field) specified in the WHERE clauses of a SELECT statement. Existing .MDX indexes will automatically be applied for =, >, <, >=, =<, and BETWEEN operators in a WHERE clause, and LIKE predicates, as well as in join predicates.

When the **dBASE** or **Microsoft FoxPro** driver is used, the index created by a CREATE UNIQUE INDEX statement is actually non-unique, and duplicate values can be inserted into the indexed column. Only one record from a set with identical key values can be added to the index.

When the **Paradox** driver is used, a unique index must be defined upon a contiguous subset of the columns in a table, including the first column. A table cannot be updated by the **Paradox** driver if a unique index is not defined on the table.

CREATE TABLE Statement Limitations

When the **Microsoft Access**, **Microsoft Excel**, or **Paradox** driver is used, and the length of a text or binary column is not specified (or is specified as 0), the column length will be set to 255.

When the **dBASE** or **Microsoft FoxPro** driver is used, and the length of a text or binary column is not specified (or is specified as 0), the column length will be set to 254.

A maximum of 255 columns is supported.

When the **Microsoft Excel** driver is used on a **Microsoft Excel** 5.0, 7.0, or 97 data source, a worksheet cannot be created with the same name as a worksheet that was previously dropped. When the **Microsoft Excel** driver is used to access a version 5.0, 7.0, or 97 worksheet, a DROP TABLE statement clears the worksheet, but does not delete the worksheet name.

When the **Paradox** driver is used, columns cannot be added once an index has been defined on a table. If the first column of the argument list of a CREATE TABLE statement creates an index, a second column cannot be included in the argument list.

Date Arithmetic Limitations

Date arithmetic is not supported for subtracting a DATE data type from a DATE data type.

DELETE Statement Limitations

The DELETE statement is not supported for the **Microsoft Excel** or **Text** driver. Note that the INSERT statement is supported for the **Text** driver.

Neither the **dBASE** nor the **Microsoft FoxPro** driver support packing a table to remove “deleted” values.

For the **Paradox** driver to delete a row from a table, the table must have a unique index (Paradox primary key).

DISTINCT Keyword Limitations

Not supported for Long Text fields (**Microsoft Access**) or Memo fields (**dBASE** or **FoxPro**).

DROP INDEX Statement Limitations

The DROP INDEX statement is not supported for the **Microsoft Excel** or **Text** driver.

DROP TABLE Statement Limitations

When the **Microsoft Excel** 5.0, 7.0, or 97 driver is used, the DROP TABLE statement clears the worksheet, but does not delete the worksheet name. Because the worksheet name still exists in the workbook, another worksheet cannot be created with the same name.

FROM Clause Limitations

The maximum number of tables in a FROM clause is 16.

HAVING Clause Limitations

The maximum number of search conditions in a HAVING clause is 40.

Identifiers Limitations

If an identifier contains a space or a special symbol, the identifier must be enclosed in back quotes. A valid name is a string of no more than 64 characters, of which the first character must not be a space. Valid names cannot include control characters or the following special characters: ` | # * ? [] . ! \$.

Do not use the reserved words listed in the SQL grammar in Appendix C of the *Microsoft ODBC Programmer's Reference* (or the shorthand form of these reserved words) as identifiers (that is, table or column names), unless you surround the word in back quotes (`).

Index Name Limitations

When the **Paradox** driver is used, a primary index must have the same name as the table upon which it is defined. Other unique or non-unique indexes must have the same name as the table upon which they are defined.

INSERT Statement Limitations

Inserted data is truncated on the right without warning if it is too long to fit into the column.

Attempting to insert a value that is out of the range of a column's data type causes a NULL to be inserted into the column.

When a **dBASE**, **Microsoft Excel**, **Microsoft FoxPro**, **Paradox**, or **Text** driver is used, inserting a zero-length string into a column actually inserts a NULL instead.

When the **Microsoft Excel** driver is used, if an empty string is inserted into a column, the empty string is converted to a NULL; a searched SELECT statement that is executed with an empty string in the WHERE clause will not succeed on that column.

A table is not updatable by the **Paradox** driver if a unique index is not defined on the table. This is not true for an empty table, which can be updated with a single row even if a unique index is not defined on the table. If a single row is inserted in an empty table that does not have a unique index, an application cannot create a unique index or insert additional data after the single row has been inserted.

When the **Text** driver is used, NULL values are represented by a blank-padded string in fixed-length files, but are represented by no spaces in delimited files. For example, in the following row containing three fields, the second field is a NULL value:

```
"Smith:., 123
```

When the **Text** driver is used, all column values can be padded with leading spaces. The length of any row must be less than or equal to 65,543 bytes.

LIKE Predicate Limitations

If data in a column is longer than 255 characters, the LIKE comparison will be based only on the first 255 characters.

A LIKE used in a procedure is supported only with constant patterns. The Desktop Database Drivers support SQL 92 LIKE pattern matching.

Use of an escape clause in a LIKE predicate is not supported.

A LIKE comparison should not be performed on a column containing data of a numeric or float data type. The results may be unpredictable. For more information, see the Microsoft Jet Database Engine Programmers's Guide.

NOT NULL Limitations

The NOT NULL constraint in the CREATE TABLE statement is not supported.

ORDER BY Clause Limitations

If a SELECT statement contains a GROUP BY clause and an ORDER BY clause, the ORDER BY clause can contain only a column in the result set or an expression in the GROUP BY clause.

Parameterized Query Limitations

When the **Microsoft Access** driver is used, a parameterized query can be called using the following syntax: `CALL query-name [(parameter[,parameter]...)]`.

Reserved Word Limitations

Do not use the reserved words listed in the SQL grammar in Appendix C of the *Microsoft ODBC Programmer's Reference* (or the shorthand form of these reserved words) as identifiers (that is, table or column names), unless you surround the word in back quotes (` `).

Scalar Function Limitations

Scalar functions are supported only by using the ODBC canonical format.

SELECT DISTINCT Limitations

The DISTINCT keyword does not apply to binary data.

SELECT Statement Limitations

An aggregate-function column cannot be mixed with a non-aggregate column in a SELECT statement.

The select list of a SELECT statement that has a GROUP BY clause can only have expressions from the GROUP BY clause or set functions.

The use of an asterisk (to select all columns) in a SELECT statement containing a GROUP BY clause is not supported. The names of the columns to be selected must be specified.

The use of a vertical bar in a SELECT statement is not supported. Use a parameter in the SELECT statement if you need to refer to a data value that contains a vertical bar.

When using a column alias in a SELECT statement, the word "as" must precede the alias. For example, "SELECT col1 as a from b." Without the "as", the statement will return an error.

If an incorrect column name is entered into a SELECT statement, a SQLSTATE 07001 error, "Wrong Number of Parameters," is returned instead of a SQLSTATE S0022 error, "Column Not Found."

When the **Microsoft Excel** driver is used, if an empty string is inserted into a column, the empty string is converted to a NULL; a searched SELECT statement that is executed with an empty string in the WHERE clause will not succeed on that column.

Set Functions Limitations

The set functions (AVG, MAX, MIN, and SUM) do not support the DISTINCT keyword.

Sorting Limitations

The maximum length of a sort key in a GROUP BY clause, ORDER BY clause, SELECT DISTINCT statement, or outer join is 255 bytes; the maximum length of all sort keys in a sort row is 65,500 bytes.

String Limitations

The maximum length of an SQL statement string is 65,000 characters.

When the **Microsoft Access** driver is used, only SQL 92 string constants (with single quotation marks, not double quotation marks) are supported.

The pipe character (|) cannot be used in a string, whether the character is enclosed in back quotes or not.

For maximum interoperability, applications should pass strings in parameters, rather than passing quoted strings.

Table Name Limitations

Table names can contain any valid characters (for example, spaces). If table names contain any characters except letters, numbers, and underscores, the name must be delimited by enclosing it in back quotes (`).

When the **Microsoft Excel** driver is used, and a table name is not qualified by a database reference, the default database is implied. If a name in Microsoft Excel includes the “!” character, it will automatically be translated to the ‘\$’ character instead.

The Microsoft Excel table name that references <filename> is supported for Microsoft Excel 3.0 and 4.0 files. The Microsoft Excel table name that references <workbook-name> is supported for Microsoft Excel 5.0, 7.0, or 97 files.

When the **dBASE** or **Microsoft FoxPro** driver is used, characters with an ASCII value greater than 127 are converted to underscores.

When the **Microsoft Access** driver is used, the table name is limited to 64 characters.

When the **dBASE**, **Microsoft Excel 3.0** or **4.0**, **Microsoft FoxPro**, **Paradox**, or **Text** driver is used, special MS-DOS keywords CON, AUX, LPT1, and LPT2 should not be used as table names.

Table References Limitations

A maximum of 16 table references can be included in any query statement.

UPDATE Statement Limitations

For the **Paradox** driver to update a table, the table must have a unique index (Paradox primary key).

Not supported by the **Text** driver.

When the **Microsoft Excel** driver is used, it is possible to update values, but a row cannot be deleted from a table based on a **Microsoft Excel** spreadsheet. As a result, the UPDATE statement is not considered officially supported by the **Microsoft Excel** driver. Only the INSERT statement is considered supported.

Views Limitations

Not supported by the **dBASE**, **Microsoft Excel**, **Microsoft FoxPro**, **Paradox**, or **Text** driver.

WHERE Clause Limitations

The maximum number of clauses in a WHERE clause is 40.

LONGVARBINARY and LONGVARCHAR columns can be compared to literals of up to 255 characters in length, but cannot be compared using parameters.

WHERE CURRENT OF Clause Limitations

Not supported.

dBASE Indexes (Advanced)

[See Also](#)

The ODBC **dBASE** driver automatically opens and updates dBASE IV index files. You must use the Select Indexes dialog box displayed through the ODBC Control Panel option to associate dBASE III .NDX files with dBASE files.

The following limitations apply to the creation of dBASE indexes:

- All column names must be valid.
- All columns must be in the same ascending or descending order.
- The length of any single text column must be less than 100 bytes.
- If more than one column exists, all of the columns must be text columns and the sum of the column sizes must be less than 100 bytes.
- Memo fields cannot be indexed.
- An index must not be specified for the current set of fields (that is, duplicate indexes are not allowed).
- The index name must match the dBASE index naming convention. dBASE III requires that each index be in a separate file, each having an .NDX extension. In dBASE IV, indexes are created as tag names that are stored in a single .MDX file. The .MDX file has the same base name as the database file (for example, EMP.MDX is the index file for the EMP.DBF database).

Microsoft FoxPro Indexes (Advanced)

[See Also](#)

The ODBC **Microsoft FoxPro** driver automatically opens and updates Microsoft FoxPro 2.5 or 2.6 index files. You must use the ODBC Control Panel option to associate Microsoft FoxPro 2.0 .IDX files with Microsoft FoxPro files.

The following limitations apply to the creation of Microsoft FoxPro indexes:

- All column names must be valid.
- All columns must be in the same ascending or descending order.
- The length of any single text column must be less than 100 bytes.
- If more than one column exists, all columns must be text columns and the sum of the column sizes must be less than 100 bytes.
- Memo fields cannot be indexed.
- An index must not be specified for the current set of fields (that is, duplicate indexes are not allowed).
- The index name must match the Microsoft FoxPro index naming convention. Microsoft FoxPro 2.0 requires that each index be in a separate file, each having an .IDX extension. In Microsoft FoxPro 2.5/2.6, indexes are created as tag names that are stored in a single .CDX file. The .CDX file has the same base name as the database file (for example, EMP.CDX is the index file for the EMP.DBF database).

CREATE INDEX for Paradox (Advanced)

[See Also](#)

The syntax of the CREATE INDEX statement for the ODBC **Paradox** driver is:

```
CREATE [UNIQUE] INDEX index-name
  ON table-name
  (column-identifier [ASC]
  [, column-identifier [ASC]...])
```

The ODBC **Paradox** driver does not support the **DESC** keyword in the ODBC SQL grammar for the CREATE INDEX statement. The *table-name* argument can specify the full path of the table.

If the keyword **UNIQUE** is specified, the ODBC **Paradox** driver will create a unique index. The first unique index is created as a primary index. This is a Paradox primary key file named *table-name.PX*. Primary indexes are subject to the following restrictions:

- The primary index must be created before any rows are added to the table.
- A primary index must be defined upon the first “n” columns in a table.
- Only one primary index is allowed per table.
- A table cannot be updated by the **Paradox** driver if a primary index is not defined on the table. (Note that this is not true for an empty table, which can be updated even if a unique index is not defined on the table.)
- The *index-name* argument for a primary index must be the same as the base name of the table, as required by Paradox.

If the keyword **UNIQUE** is omitted, the ODBC **Paradox** driver will create a non-unique index. This consists of two Paradox secondary index files named *table-name.Xnn* and *table-name.Ynn*, where *nn* is the number of the column in the table. Non-unique indexes are subject to the following restrictions:

- Before a non-unique index can be created for a table, a primary index must exist for that table.
- For Paradox 3.x, the *index-name* argument for any index other than a primary index (unique or non-unique) must be the same as the column name. For Paradox 4.x and 5.x, the name of such an index can be, but doesn't have to be, the same as the column name.
- Only one column can be specified for a non-unique index.

Columns cannot be added once an index has been defined on a table. If the first column of the argument list of a CREATE TABLE statement creates an index, a second column cannot be included in the argument list.

For example, to use the sales order number and line number columns as the unique index on the SO_LINES table, use the statement:

```
CREATE UNIQUE INDEX SO_LINES
  ON SO_LINES (SONum, LineNum)
```

To use the part number column as a non-unique index on the SO_LINES table, use the statement:

```
CREATE INDEX PartNum
  ON SO_LINES (PartNum)
```

Note that when two CREATE INDEX statements are performed, the first statement will always create a primary index with the same name as the table and the second statement will always create a non-unique index with the same name as the column. These indexes will be named this way even if different names are entered in the CREATE INDEX statements and even if the index is labeled **UNIQUE** in the second CREATE INDEX statement.

Microsoft Access Data Types (Advanced)

The following table shows the **Microsoft Access** data types, data types used to create tables, and ODBC SQL data types.

Microsoft Access data type	Data type (CREATETABLE)	ODBC SQL data type
BINARY	BINARY	SQL_BINARY
BIT	BIT	SQL_BIT
COUNTER	COUNTER	SQL_INTEGER
CURRENCY	CURRENCY	SQL_NUMERIC
DATE/TIME	DATETIME	SQL_TIMESTAMP
LONG BINARY	LONGBINARY	SQL_LONGVARBINARY
LONG TEXT	LONGTEXT	SQL_LONGVARCHAR
MEMO	LONGTEXT	SQL_LONGVARCHAR
NUMBER (FieldSize=SINGLE)	SINGLE	SQL_REAL
NUMBER (FieldSize=DOUBLE)	DOUBLE	SQL_DOUBLE
NUMBER (FieldSize=BYTE)	UNSIGNED BYTE	SQL_TINYINT
NUMBER (FieldSize=INTEGER)	SHORT	SQL_SMALLINT
NUMBER (FieldSize=LONG INTEGER)	LONG	SQL_INTEGER
OLE	LONGBINARY	SQL_LONGVARBINARY
TEXT	VARCHAR	SQL_VARCHAR
VARBINARY	VARBINARY	SQL_VARBINARY

Note **SQLGetTypeInfo** returns ODBC data types. It will not return all Microsoft Access data types if more than one **Microsoft Access** type is mapped to the same ODBC SQL data type. All conversions in Appendix D of the *Microsoft ODBC Programmer's Reference* are supported for the SQL data types listed in the previous table.

The following table shows limitations on Microsoft Access data types.

Data type	Description
BINARY, VARBINARY, and VARCHAR	Creating a BINARY, VARBINARY, or VARCHAR column of zero or unspecified length actually returns a 255-byte column.
BYTE	Even though a Microsoft Access NUMBER field with a FieldSize equal to BYTE is unsigned, a negative number can be inserted into the field when using the Microsoft Access driver.

CHAR,
LONGVARCHAR,
and VARCHAR

A character string literal can contain any ANSI character (1-255 decimal). Use two consecutive single quotation marks (") to represent one single quotation mark (').

Procedures should be used to pass character data when using any special character in a character data type column.

DATE

Date values must be either delimited according to the ODBC canonical date format or delimited by the datetime delimiter ("#"). Otherwise, **Microsoft Access** will treat the value as an arithmetic expression and will not raise a warning or error.

For example, the date "March 5, 1996" must be represented as {d '1996-03-05'} or #03/05/1996#; otherwise, if only 03/05/1993 is submitted, **Microsoft Access** will evaluate this as 3 divided by 5 divided by 1996. This value rounds up to the integer 0, and since the zero day maps to 1899-12-31, this is the date used.

A pipe character (|) cannot be used in a date value, even if enclosed in back quotes.

More limitations on data types can be found in [Data Type Limitations](#).

dBASE Data Types (Advanced)

The following table shows how **dBASE** data types are mapped to ODBC SQL data types. Note that not all ODBC SQL data types are supported.

dBASE data type	ODBC data type
CHAR	SQL_VARCHAR
DATE	SQL_DATE
FLOAT *	SQL_DOUBLE
LOGICAL	SQL_BIT
MEMO	SQL_LONGVARCHAR
NUMERIC (BCD)	SQL_DOUBLE
OLEOBJECT *	SQL_LONGBINARY

* Only valid for dBASE version 5.x

Precision in dBASE III allows numbers with up to two-digit exponents and in dBASE IV numbers with up to three-digit exponents. Because numbers are stored as text, they are converted to numbers. If the number to convert does not fit in a field, unexplained results may occur.

While **dBASE** allows a precision and a scale to be specified with a NUMERIC data type, it is not supported by the ODBC **dBASE** driver. The ODBC **dBASE** driver always returns a precision of 15 and a scale of 0 for a NUMERIC data type.

A column created with the Numeric data type using the ODBC **dBASE** driver maps to the SQL_DOUBLE ODBC data type. Thus the data in this column is subject to rounding. This behavior is not the same as that of the NUMERIC data type in **dBASE** (type N), which is Binary Coded Decimal (BCD).

Note **SQLGetTypeInfo** returns ODBC SQL data types. All conversions in Appendix D of the *Microsoft ODBC SDK Programmer's Reference* are supported for the ODBC SQL data types listed earlier in this topic.

The following table shows limitations on dBASE data types.

Data type	Description
CHAR	Creating a CHAR column of zero or unspecified length actually returns a 254-byte column.
Encrypted data	The dBASE driver does not support encrypted dBASE tables.
LOGICAL	The dBASE driver cannot create an index on a LOGICAL column.
MEMO	The maximum length of a MEMO column is 65,500 bytes.

More limitations on data types can be found in [Data Type Limitations](#).

Microsoft Excel Data Types (Advanced)

The following table shows how **Microsoft Excel** driver data types are mapped to ODBC SQL data types. The **Microsoft Excel** driver assigns these data types to columns in **Microsoft Excel** tables based on the data in the column.

Microsoft Excel data type	ODBC data type
CURRENCY	SQL_NUMERIC
DATETIME	SQL_TIMESTAMP
LOGICAL	SQL_BIT
NUMBER	SQL_DOUBLE
TEXT	SQL_VARCHAR

Note **SQLGetTypeInfo** returns ODBC SQL data types. All conversions in Appendix D of the *Microsoft ODBC SDK Programmer's Reference* are supported for the ODBC SQL data types listed earlier in this topic.

The following table shows limitations on **Microsoft Excel** data types.

Data type	Description
Encrypted data	The Microsoft Excel driver cannot read encrypted data.
Error Strings	The Microsoft Excel driver cannot return a character string for the Microsoft Excel error values (#N/A!, #VALUE!, #REF!, #DIV/0!, #NUM!, #NAME?, and #NULL!), but returns a NULL instead.
LOGICAL	The value in a LOGICAL column is returned in a SQL_C_CHAR buffer as either 0 or 1.
NUMBER	If an integer column is created, numbers that are too big for the integer data type can be entered, and data containing non-integer values can be inserted, with the result that the column may be converted to SQL_DOUBLE.
TEXT	<p>When the rows of a column contain more than one Microsoft Excel data type, the ODBC Microsoft Excel driver assigns the SQL_VARCHAR data type to the column. There is one exception to this: if the column contains only two or three of the datetime data types (DATE, TIME, and DATETIME), the ODBC Microsoft Excel driver assigns the SQL_TIMESTAMP data type to the column.</p> <p>Creating a TEXT column of zero or unspecified length actually returns a 255-byte column.</p> <p>A character string literal can contain any ANSI character (1-255 decimal). Use two consecutive single quotation marks (") to represent one single quotation mark (').</p> <p>Inserting a NULL into a column with a data type other than SQL_VARCHAR will cause the data type of the column to change to SQL_VARCHAR.</p>

More limitations on data types can be found in [Data Type Limitations](#).

Microsoft FoxPro Data Types (Advanced)

The following table shows how **Microsoft FoxPro** data types are mapped to ODBC SQL data types. Note that not all ODBC SQL data types are supported.

Microsoft FoxPro data type	ODBC data type
CHAR	SQL_VARCHAR
DATE	SQL_DATE
GENERAL	SQL_LONGVARBINARY
LOGICAL	SQL_BIT
MEMO	SQL_LONGVARCHAR
NUMERIC (BCD)	SQL_DOUBLE

Precision in Microsoft FoxPro 2.0 allows numbers with up to two-digit exponents and in Microsoft FoxPro 2.5 and 2.6 numbers with up to three-digit exponents. Because numbers are stored as text, they are converted to numbers. If the number to convert does not fit in a field, unexplained results may occur.

While Microsoft FoxPro allows a precision and a scale to be specified with a NUMERIC data type, it is not supported by the ODBC **Microsoft FoxPro** driver. The ODBC **Microsoft FoxPro** driver always returns a precision of 15 and a scale of 0 for a NUMERIC data type.

A column created with the NUMERIC data type using the ODBC **Microsoft FoxPro** driver maps to the SQL_DOUBLE ODBC data type. Thus the data in this column is subject to rounding. This behavior is not the same as that of the NUMERIC data type in Microsoft FoxPro (type N), which is Binary Coded Decimal (BCD).

Note **SQLGetTypeInfo** returns ODBC SQL data types. All conversions in Appendix D of the *Microsoft ODBC SDK Programmer's Reference* are supported for the ODBC SQL data types listed earlier in this topic.

The following table shows limitations on **Microsoft FoxPro** data types.

Data type	Description
CHAR	Creating a CHAR column of zero or unspecified length actually returns a 254-byte column.
GENERAL	The GENERAL data type is supported only for Microsoft FoxPro version 2.5 and 2.6 files.
MEMO	The maximum length of a MEMO column is 65,500 bytes.

More limitations on data types can be found in [Data Type Limitations](#).

Paradox Data Types (Advanced)

The ODBC **Paradox** driver maps Paradox data types to ODBC SQL data types. The following table lists all Paradox data types and shows the ODBC SQL data types they are mapped to.

Paradox data type	ODBC data type
ALPHANUMERIC	SQL_VARCHAR
AUTOINCREMENT #	SQL_INTEGER
BCD #	SQL_DOUBLE
BYTES #	SQL_BINARY
DATE	SQL_DATE
IMAGE *	SQL_LONGVARBINARY
LOGICAL #	SQL_BIT
LONG #	SQL_INTEGER
MEMO *	SQL_LONGVARCHAR
MONEY #	SQL_DOUBLE
NUMBER	SQL_DOUBLE
SHORT	SQL_SMALLINT
TIME #	SQL_TIMESTAMP
TIMESTAMP #	SQL_TIMESTAMP

* Only valid for Paradox versions 4.x and 5.x.

Only valid for Paradox versions 5.x.

Note **SQLGetTypeInfo** returns ODBC SQL data types. All conversions in Appendix D of the *Microsoft ODBC SDK Programmer's Reference* are supported for the ODBC SQL data types listed earlier in this topic.

The following table shows limitations on Paradox data types.

Data type	Description
ALPHANUMERIC	Creating an ALPHANUMERIC column of zero or unspecified length actually returns a 255-byte column.
BYTES	If you insert NULL into a binary column with the Paradox5 driver, it is changed to 0.
LONG	The maximum negative value supported by the Paradox driver for the Long data type in Paradox 5.x is not -2^{31} (-2147483648), as it should be since Long maps to the ODBC data type SQL_INTEGER. The maximum negative value supported for Long is actually $-2^{31} + 1$ (-2147483647).
TIMESTAMP	When a value is inserted into a TIMESTAMP column by the Paradox driver, then subsequently retrieved from the column, the retrieved value may differ from the inserted value by as much as 1 second because of rounding.

More limitations on data types can be found in [Data Type Limitations](#).

Text Data Types (Advanced)

The following table shows how text data types are mapped to ODBC SQL data types. Note that not all ODBC SQL data types are supported by the ODBC **Text** driver.

Text data type	ODBC data type
CHAR	SQL_VARCHAR
DATETIME	SQL_TIMESTAMP
FLOAT	SQL_DOUBLE
INTEGER	SQL_INTEGER
LONGCHAR	SQL_LONGVARCHAR

Note **SQLGetTypeInfo** returns ODBC data types. All conversions in Appendix D of the *Microsoft ODBC Programmer's Reference* are supported for the SQL data types listed in the previous table.

The following table shows limitations on **Text** data types.

Data type	Description
CHAR	<p>Creating a CHAR column of zero or unspecified length actually returns a 255-bit column.</p> <p>In delimited files, a CHAR column may or may not have double quotation mark delimiters at the beginning and the end; in fixed-length files, double quotation marks are not used as delimiters.</p>
DATETIME	<p>MM-DD-YY (for example, 01-17-92) MMM-DD-YY (for example, Jan-17-92) DD-MMM-YY (for example, 17-Jan-92) YYYY-MM-DD (for example, 1992-01-17) YYYY-MMM-DD (for example, 1992-Jan-17)</p> <p>Mixed date separators are not allowed within a table.</p> <p>The Text ISAM formats a DATETIME field in the American or European format, depending upon the International setting in the Windows Control Panel.</p>
FLOAT	<p>The maximum width includes the sign and decimal point. In SCHEMA.INI, the width is denoted as follows:</p> <p>14.083 is FLOAT Width 6 -14.083 is FLOAT Width 7 +14.083 is FLOAT Width 7 14083. is FLOAT Width 6</p> <p>ODBC always returns 8 for FLOAT columns.</p> <p>FLOAT columns can also be in scientific notation, for example:</p> <p>-3.04E+2 is Float Width 8 25E4 is Float Width 4</p> <p>Note Decimal and scientific notation cannot be mixed in a column.</p> <p>NULL values are represented by a blank padded string in fixed-length files, and are omitted in delimited files.</p> <p>Float data may be padded with leading blanks.</p>
INTEGER	<p>Valid values for INTEGER columns are 32767 to -32766.</p> <p>In SCHEMA.INI, the width is denoted as follows:</p> <p>14083 is INTEGER Width 5</p>

0 is INTEGER Width 1

ODBC always returns 4 for INTEGER columns.

The maximum width includes a sign. The maximum width of an INTEGER column is 11, although the width can be greater due to blanks that are allowed in fixed-format tables.

LONGCHAR

The theoretical limit on the width of a LONGCHAR column in either a fixed-length or delimited table is 65500K. The **Text** ISAM is more likely to provide reliable support up to about 32K.

More limitations on data types can be found in [Data Type Limitations](#).

Data Type Limitations (Advanced)

The Microsoft ODBC Desktop Database Drivers impose the following limitations on the data types:

Data type	Description
All data types	Type conversion failures may result in the affected column being set to NULL.
BINARY	Creating a zero-length BINARY column actually returns a 255-byte BINARY column.
DATE	The DATE data type cannot be converted to another data type (or itself) by the CONVERT function.
DECIMAL (Exact Numeric)	Not supported.
Floating-Point Data Types	The number of decimal places in a floating-point number may be limited by the number format set in the International section of the Windows Control Panel.
TIMESTAMP	The TIMESTAMP data type cannot be converted to itself by the CONVERT function.
TINYINT	TINYINT values are always unsigned.
Zero-Length Strings	When a dBASE , Microsoft Excel , Microsoft FoxPro , Paradox , or Text driver is used, inserting a zero-length string into a column actually inserts a NULL instead.

Error Messages (Advanced)

When an error occurs, the Microsoft ODBC Desktop Database drivers return the native error number, the SQLSTATE (an ODBC error code), and an error message.

Native Error

Positive native error codes are generated by the Desktop Database Drivers. Negative error codes are generated by Microsoft Jet. For more information, see the Microsoft Jet Database Engine Programmers's Guide.

SQLSTATE

For errors that occur in the data source, the ODBC driver maps the returned native error to the appropriate SQLSTATE. For errors that are detected by the driver or the Driver Manager, the ODBC driver or Driver Manager generates the appropriate SQLSTATE.

Error Message

For errors that occur in the data source, the ODBC driver returns an error message returned to it by the ODBC File Library. For errors that occur in the ODBC driver or the Driver Manager, the driver returns an error message based on the text associated with the SQLSTATE.

Error messages have the following format:

[vendor][ODBC-component][data-source]message-text

where the prefixes in brackets ([]) identify the location of the error. When the error occurs in the Driver Manager, *data-source* is not given. When the error occurs in the data source, the *[vendor]* and *[ODBC-component]* prefixes identify the vendor and name of the ODBC component that received the error from the data source.

The following table shows the error messages returned by the Driver Manager and driver ISAM:

Error message	Error location
[Microsoft][ODBC Driver Manager] <i>message-text</i>	Driver Manager (ODBC32.DLL)
[Microsoft][ODBC <i>driver-name</i>] <i>message-text</i>	Driver ISAM (see Driver ISAMs Table)

SQLGetInfo Returned Values (Programming)

The following table lists the C-language #defines for the *flInfoType* argument and the corresponding values returned by **SQLGetInfo**. This information can be retrieved by passing the listed C-language #defines to **SQLGetInfo** in the *flInfoType* argument. Where **SQLGetInfo** returns a 32-bit bitmask, a vertical bar (|) represents a bitwise OR. For more information about the values returned by **SQLGetInfo**, see the *Microsoft ODBC SDK Programmer's Reference, Version 2.0*.

>>	SQL_ACCESSIBLE_PROCEDURES
>>	SQL_ACCESSIBLE_TABLES
>>	SQL_ACTIVE_CONNECTIONS
>>	SQL_ACTIVE_STATEMENTS
>>	SQL_ALTER_TABLE
>>	SQL_BOOKMARK_PERSISTENCE
>>	SQL_COLUMN_ALIAS
>>	SQL_CONCAT_NULL_BEHAVIOR
>>	SQL_CONVERT_BIGINT
>>	SQL_CONVERT_BINARY
>>	SQL_CONVERT_BIT
>>	SQL_CONVERT_CHAR
>>	SQL_CONVERT_DATE
>>	SQL_CONVERT_DECIMAL
>>	SQL_CONVERT_DOUBLE
>>	SQL_CONVERT_FLOAT
>>	SQL_CONVERT_FUNCTIONS
>>	SQL_CONVERT_INTEGER
>>	SQL_CONVERT_LONGVARBINARY
>>	SQL_CONVERT_LONGVARCHAR
>>	SQL_CONVERT_NUMERIC
>>	SQL_CONVERT_REAL
>>	SQL_CONVERT_SMALLINT
>>	SQL_CONVERT_TIME
>>	SQL_CONVERT_TIMESTAMP
>>	SQL_CONVERT_TINYINT
>>	SQL_CONVERT_VARBINARY
>>	SQL_CONVERT_VARCHAR
>>	SQL_CORRELATION_NAME
>>	SQL_CURSOR_COMMIT_BEHAVIOR
>>	SQL_CURSOR_ROLLBACK_BEHAVIOR
>>	SQL_DATA_SOURCE_NAME
>>	SQL_DATA_SOURCE_READ_ONLY
>>	SQL_DATABASE_NAME
>>	SQL_DBMS_NAME
>>	SQL_DBMS_VER
>>	SQL_DEFAULT_TXN_ISOLATION
>>	SQL_DRIVER_HDBC
>>	SQL_DRIVER_HENV
>>	SQL_DRIVER_HLIB
>>	SQL_DRIVER_HSTMT
>>	SQL_DRIVER_NAME
>>	SQL_DRIVER_ODBC_VER
>>	SQL_DRIVER_VER
>>	SQL_EXPRESSIONS_IN_ORDERBY
>>	SQL_FETCH_DIRECTION
>>	SQL_FILE_USAGE

» SQL_GETDATA_EXTENSIONS
» SQL_GROUP_BY
» SQL_IDENTIFIER_CASE
» SQL_IDENTIFIER_QUOTE_CHAR
» SQL_KEYWORDS
» SQL_LIKE_ESCAPE_CLAUSE
» SQL_LOCK_TYPES
» SQL_MAX_BINARY_LITERAL_LEN
» SQL_MAX_CHAR_LITERAL_LEN
» SQL_MAX_COLUMN_NAME_LEN
» SQL_MAX_COLUMNS_IN_GROUP_BY
» SQL_MAX_COLUMNS_IN_INDEX
» SQL_MAX_COLUMNS_IN_ORDER_BY
» SQL_MAX_COLUMNS_IN_SELECT
» SQL_MAX_COLUMNS_IN_TABLE
» SQL_MAX_CURSOR_NAME_LEN
» SQL_MAX_INDEX_SIZE
» SQL_MAX_OWNER_NAME_LEN
» SQL_MAX_PROCEDURE_NAME_LEN
» SQL_MAX_QUALIFIER_NAME_LEN
» SQL_MAX_ROW_SIZE
» SQL_MAX_ROW_SIZE_INCLUDES_LONG
» SQL_MAX_STATEMENT_LEN
» SQL_MAX_TABLE_NAME_LEN
» SQL_MAX_TABLES_IN_SELECT
» SQL_MAX_USER_NAME_LEN
» SQL_MULT_RESULT_SETS
» SQL_MULTIPLE_ACTIVE_TXN
» SQL_NEED_LONG_DATA_LEN
» SQL_NON_NULLABLE_COLUMNS
» SQL_NULL_COLLATION
» SQL_NUMERIC_FUNCTIONS
» SQL_ODBC_API_CONFORMANCE
» SQL_ODBC_SAG_CLI_CONFORMANCE
» SQL_ODBC_SQL_CONFORMANCE
» SQL_ODBC_SQL_OPT_IEF
» SQL_ODBC_VER
» SQL_OJ_CAPABILITIES
» SQL_ORDER_BY_COLUMNS_IN_SELECT
» SQL_OUTER_JOINS
» SQL_OWNER_TERM
» SQL_OWNER_USAGE
» SQL_POS_OPERATIONS
» SQL_POSITIONED_STATEMENTS
» SQL_PROCEDURE_TERM
» SQL_PROCEDURES
» SQL_QUALIFIER_LOCATION
» SQL_QUALIFIER_NAME_SEPARATOR
» SQL_QUALIFIER_TERM
» SQL_QUALIFIER_USAGE
» SQL_QUOTED_IDENTIFIER_CASE
» SQL_ROW_UPDATES
» SQL_SCROLL_CONCURRENCY
» SQL_SCROLL_OPTIONS
» SQL_SEARCH_PATTERN_ESCAPE

- » SQL_SERVER_NAME
- » SQL_SPECIAL_CHARACTERS
- » SQL_STATIC_SENSITIVITY
- » SQL_STRING_FUNCTIONS
- » SQL_SUBQUERIES
- » SQL_SYSTEM_FUNCTIONS
- » SQL_TABLE_TERM
- » SQL_TIMEDATE_ADD_INTERVALS
- » SQL_TIMEDATE_DIFF_INTERVALS
- » SQL_TIMEDATE_FUNCTIONS
- » SQL_TXN_CAPABLE
- » SQL_TXN_ISOLATION_OPTION
- » SQL_UNION
- » SQL_USER_NAME

SQL_ACCESSIBLE_PROCEDURES

"Y"

(Microsoft Access)

"N"

(dBASE, Microsoft Excel, Microsoft FoxPro, Paradox, Text)

SQL_ACCESSIBLE_TABLES

"Y"

SQL_ACTIVE_CONNECTIONS

64

SQL_ACTIVE_STATEMENTS

0

SQL ALTER TABLE

SQL_AT_ADD_COLUMN |

SQL_AT_DROP_COLUMN

(Microsoft Access, dBASE, Microsoft FoxPro, Paradox)

0

(Microsoft Excel, Text)

SQL_BOOKMARK_PERSISTENCE

SQL_BP_SCROLL |
SQL_BP_UPDATE (1)

1 Bookmarks persist after a commit, but do not persist after a rollback.

SQL_COLUMN_ALIAS

“Y”

SQL_CONCAT_NULL_BEHAVIOR
SQL_CB_NON_NULL

SQL_CONVERT_BIGINT

0

SQL_CONVERT_BINARY

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_BIT

0

SQL_CONVERT_CHAR

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_DATE

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_DECIMAL

0

SQL_CONVERT_DOUBLE

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_FLOAT

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_FUNCTIONS

SQL_FN_CVT_CONVERT

SQL_CONVERT_INTEGER

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_LONGVARBINARY

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_LONGVARCHAR

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_NUMERIC

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_REAL

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_SMALLINT

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_TIME

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_TIMESTAMP

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_TINYINT

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_VARBINARY

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CONVERT_VARCHAR

SQL_CVT_DOUBLE |
SQL_CVT_FLOAT |
SQL_CVT_INTEGER |
SQL_CVT_NUMERIC |
SQL_CVT_REAL |
SQL_CVT_SMALLINT |
SQL_CVT_VARCHAR

SQL_CORRELATION_NAME

SQL_CN_DIFFERENT

SQL_CURSOR_COMMIT_BEHAVIOR
SQL_CB_CLOSE

SQL_CURSOR_ROLLBACK_BEHAVIOR
SQL_CB_CLOSE

SQL_DATA_SOURCE_NAME

The DSN from ODBC.INI, or "" if DRIVER keyword is used in ODBC.INI

SQL_DATA_SOURCE_READ_ONLY

"N"

(Microsoft Access, dBASE, Microsoft FoxPro, Paradox)

"Y"

(Microsoft Excel, Text)

SQL_DATABASE_NAME

File name (Microsoft Access, Microsoft Excel 5.0/7.0//97)

Current database directory (dBASE, Microsoft Excel 3.0/4.0, Microsoft FoxPro, Paradox, Text)

SQL_DBMS_NAME

"ACCESS" (Microsoft Access)

"DBASE" (dBASE)

"EXCEL" (Microsoft Excel)

"FOXPRO" (Microsoft FoxPro)

"PARADOX" (Paradox)

"TEXT" (Text)

SQL_DBMS_VER

"1.0", "1.1", "2.0", or "3.0" (Microsoft Access)

(There is no difference in the data format in Microsoft Access versions 3.0, 7.0, and 97.)

"3.0", "4.0", or "5.0" (dBASE)

"3.0", "4.0", "5.0", "7.0", or "97" (Microsoft Excel)

"2.0", "2.5", or "2.6" (Microsoft FoxPro)

"3.x", "4.x", or "5.x" (Paradox)

"1.0" (Text)

SQL_DEFAULT_TXN_ISOLATION

SQL_TXN_READ_COMMITTED [Microsoft Access]

0 [dBASE, Microsoft Excel, Microsoft FoxPro, Paradox, Text]

SQL_DRIVER_HDBC

Handled by the Driver Manager.

SQL_DRIVER_HENV

Handled by the Driver Manager.

SQL_DRIVER_HLIB

Handled by the Driver Manager.

SQL_DRIVER_HSTMT

Handled by the Driver Manager.

SQL_DRIVER_NAME

"ODBCJT32.DLL"

SQL_DRIVER_ODBC_VER

"2.50"

SQL_DRIVER_VER

" 3.50.*nnnn*" (*nnnn* specifies the build date.)

SQL_EXPRESSIONS_IN_ORDERBY

"Y"

SQL_FETCH_DIRECTION

SQL_FD_FETCH_NEXT |
SQL_FD_FETCH_FIRST |
SQL_FD_FETCH_LAST |
SQL_FD_FETCH_PRIOR |
SQL_FD_FETCH_ABSOLUTE |
SQL_FD_FETCH_RELATIVE |
SQL_FD_FETCH_BOOKMARK

SQL_FILE_USAGE

SQL_FILE_QUALIFIER

(Microsoft Access, Microsoft Excel 5.0, 7.0, or 97)

SQL_FILE_TABLE

(dBASE, Microsoft Excel 3.0 or 4.0, Microsoft FoxPro, Paradox, Text)

SQL_GETDATA_EXTENSIONS

SQL_GD_ANY_ORDER |
SQL_GD_ANY_COLUMN |
SQL_GD_BLOCK |
SQL_GD_BOUND

SQL_GROUP_BY

SQL_GB_GROUP_BY_CONTAINS_SELECT

SQL_IDENTIFIER_CASE

SQL_IC_MIXED

(Microsoft Access, Microsoft Excel, Microsoft FoxPro, Paradox, Text)

SQL_IC_UPPER (1)

(dBASE)

1 The qualifier is returned in mixed case so that Windows NT can locate the directory.

SQL_IDENTIFIER_QUOTE_CHAR

"" (backquote)

SQL_KEYWORDS

ALPHANUMERIC, AUTOINCREMENT, BINARY, BOOLEAN, BYTE, COUNTER, CURRENCY, DATABASE, DATABASENAME, DATETIME, DISALLOW, DISTINCTROW, DOUBLEFLOAT, FLOAT4, FLOAT8, GENERAL, IEEEDOUBLE, IEEE SINGLE, IGNORE, IMAGE, INTEGER1, INTEGER2, INTEGER4, LOGICAL, LOGICAL1, LONG, LONGBINARY, LONGCHAR, LONGTEXT, MEMO, MONEY, NOTE, NUMBER, OLEOBJECT, OWNERACCESS, PARAMETERS, PERCENT, PIVOT, SHORT, SINGLE, SINGLEFLOAT, STDEV, STDEVP, STRING, TABLEID, TEXT, TOP, TRANSFORM, UNSIGNEDBYTE, VAR, VARBINARY, VARP, YESNO

SQL_LIKE_ESCAPE_CLAUSE

"N"

SQL_LOCK_TYPES

SQL_LCK_NO_CHANGE

SQL_MAX_BINARY_LITERAL_LEN

255

SQL_MAX_CHAR_LITERAL_LEN

255 (Microsoft Access)

254 (dBASE)

255 (Microsoft Excel 3.0, 4.0, 5.0, or 7.0)

65535 (Microsoft Excel 97)

254 (Microsoft FoxPro)

255 (Paradox)

255 (Text)

SQL_MAX_COLUMN_NAME_LEN

- 64 (Microsoft Access)
- 10 (dBASE)
- 64 (Microsoft Excel 5.0, 7.0, or 97)
- 30 (Microsoft Excel 3.0 or 4.0)
- 10 (Microsoft FoxPro)
- 25 (Paradox)
- 64 (Text)

SQL_MAX_COLUMNS_IN_GROUP_BY

10

SQL_MAX_COLUMNS_IN_INDEX

32 (Microsoft Access)

0 (Limit Unknown or Not Applicable) (dBase)

0 (Microsoft Excel)

0 (Limit Unknown or Not Applicable) (Microsoft FoxPro)

0 (Limit Unknown or Not Applicable) (Paradox)

0 (Text)

SQL_MAX_COLUMNS_IN_ORDER_BY

10

SQL_MAX_COLUMNS_IN_SELECT

255

SQL_MAX_COLUMNS_IN_TABLE

255

When using the **Microsoft Excel** driver, a CREATE TABLE statement may allow 256 columns, but the 255 column limit is still valid, and an insert into column number 256 will fail.

SQL_MAX_CURSOR_NAME_LEN

64

SQL_MAX_INDEX_SIZE

255 (Microsoft Access)
220 (dBASE)
0 (Microsoft Excel)
512 (Microsoft FoxPro)
1350 (Paradox)
0 (Text)

SQL_MAX_OWNER_NAME_LEN

0

SQL_MAX_PROCEDURE_NAME_LEN

64 (Microsoft Access)
0 (dBASE)
0 (Microsoft Excel)
0 (Microsoft FoxPro)
0 (Paradox)
0 (Text)

SQL_MAX_QUALIFIER_NAME_LEN

66

SQL_MAX_ROW_SIZE

2096 (Microsoft Access)
4000 (dBASE)
65535 (Microsoft Excel)
65500 (Microsoft FoxPro 2.5 and above)
4000 (Microsoft FoxPro 2.0)
1350 (Paradox)
65535 (Text)

SQL_MAX_ROW_SIZE_INCLUDES_LONG

"N" (Microsoft Access)

"N" (dBASE)

"Y" (Microsoft Excel)

"N" (Microsoft FoxPro)

"N" (Paradox)

"Y" (Text)

SQL_MAX_STATEMENT_LEN

65000

SQL_MAX_TABLE_NAME_LEN

- 64 (Microsoft Access)
- 12 (dBASE)
- 31 (Microsoft Excel 5.0, 7.0, or 97)
- 12 (Microsoft Excel 3.0 or 4.0)
- 12 (Microsoft FoxPro)
- 12 (Paradox)
- 12 (Text)

SQL_MAX_TABLES_IN_SELECT

16

SQL_MAX_USER_NAME_LEN

0

SQL_MULT_RESULT_SETS

"N"

SQL_MULTIPLE_ACTIVE_TXN

"Y"

SQL_NEED_LONG_DATA_LEN

"N"

SQL_NON_NULLABLE_COLUMNS

SQL_NNC_NON_NULL

SQL_NULL_COLLATION

SQL_NC_LOW

SQL_NUMERIC_FUNCTIONS

SQL_FN_NUM_ABS |
SQL_FN_NUM_ATAN |
SQL_FN_NUM_CEILING |
SQL_FN_NUM_COS |
SQL_FN_NUM_EXP |
SQL_FN_NUM_FLOOR |
SQL_FN_NUM_LOG |
SQL_FN_NUM_MOD |
SQL_FN_NUM_POWER |
SQL_FN_NUM_RAND |
SQL_FN_NUM_SIGN |
SQL_FN_NUM_SIN |
SQL_FN_NUM_SQRT |
SQL_FN_NUM_TAN

SQL_ODBC_API_CONFORMANCE
SQL_OAC_LEVEL1

SQL_ODBC_SAG_CLI_CONFORMANCE
SQL_OSCC_COMPLIANT

SQL_ODBC_SQL_CONFORMANCE
SQL_OSC_MINIMUM

SQL_ODBC_SQL_OPT_IEF

"N"

SQL_ODBC_VER

From Driver Manager

SQL_OJ_CAPABILITIES

SQL_OJ_LEFT

SQL_OJ_RIGHT

SQL_OJ_NOT_ORDERED

SQL_OJ_INNER

SQL_OJ_ALL_COMPARISON_OPS

SQL_ORDER_BY_COLUMNS_IN_SELECT

"N"

SQL_OUTER_JOINS

"Y"

SQL_OWNER_TERM

'''

SQL_OWNER_USAGE

0

SQL_POS_OPERATIONS

SQL_POS_POSITION |
SQL_POS_REFRESH |
SQL_POS_UPDATE |
SQL_POS_DELETE |
SQL_POS_ADD

(Microsoft Access, dBASE, Microsoft FoxPro, Paradox)

SQL_POS_POSITION |
SQL_POS_REFRESH

(Microsoft Excel, Text)

SQL_POSITIONED_STATEMENTS

0

SQL_PROCEDURE_TERM

"QUERY"

(Microsoft Access)

""

(dBASE, Microsoft Excel, Microsoft FoxPro, Paradox, Text)

SQL_PROCEDURES

"Y"

(Microsoft Access)

"N"

(dBASE, Microsoft Excel, Microsoft FoxPro, Paradox, Text)

SQL_QUALIFIER_LOCATION
SQL_QL_START

SQL_QUALIFIER_NAME_SEPARATOR

."

(Microsoft Access, Microsoft Excel 5.0, 7.0, or 97)

"\"

(dBASE, Microsoft Excel 3.0 or 4.0, Microsoft FoxPro, Paradox, Text)

SQL_QUALIFIER_TERM

"Database" (Microsoft Access)

"Directory" (dBASE)

"Workbook" (Microsoft Excel 5.0, 7.0, or 97)

"Directory" (Microsoft Excel 3.0 or 4.0)

"Directory" (Microsoft FoxPro)

"Directory" (Paradox)

"Directory" (Text)

SQL_QUALIFIER_USAGE

SQL_QU_DML_STATEMENTS |
SQL_QU_TABLE_DEFINITION |
SQL_QU_INDEX_DEFINITION |
SQL_QU_PROCEDURE_INVOCATION
(Microsoft Access)

SQL_QU_DML_STATEMENTS |
SQL_QU_TABLE_DEFINITION |
SQL_QU_INDEX_DEFINITION
(dBASE, Microsoft FoxPro, Paradox)

SQL_QU_DML_STATEMENTS |
SQL_QU_TABLE_DEFINITION
(Microsoft Excel, Text)

SQL_QUOTED_IDENTIFIER_CASE
SQL_IC_MIXED

SQL_ROW_UPDATES

"N"

SQL_SCROLL_CONCURRENCY

SQL_SCCO_READ_ONLY |

SQL_SCCO_LOCK |

SQL_SCCO_OPT_VALUES

(Microsoft Access, dBASE, Microsoft FoxPro, Paradox)

SQL_SCCO_READ_ONLY

(Microsoft Excel, Text)

SQL_SCROLL_OPTIONS

SQL_SO_FORWARD_ONLY |

SQL_SO_STATIC |

SQL_SO_KEYSET_DRIVEN

SQL_SEARCH_PATTERN_ESCAPE

"\"

SQL_SERVER_NAME

"ACCESS" (Microsoft Access)

"DBASE" (dBASE)

"EXCEL" (Microsoft Excel)

"FOXPRO" (Microsoft FoxPro)

"PARADOX" (Paradox)

"TEXT" (Text)

SQL_SPECIAL_CHARACTERS

"~`@\$%^&* _+=\}|";:~/><,!'[]"

SQL_STATIC_SENSITIVITY

SQL_SS_ADDITIONS |
SQL_SS_DELETIONS |
SQL_SS_UPDATES

SQL_STRING_FUNCTIONS

SQL_FN_STR_ASCII |
SQL_FN_STR_CHAR |
SQL_FN_STR_CONCAT |
SQL_FN_STR_LCASE |
SQL_FN_STR_LEFT |
SQL_FN_STR_LENGTH |
SQL_FN_STR_LOCATE |
SQL_FN_STR_LOCATE_2
SQL_FN_STR_LTRIM |
SQL_FN_STR_RIGHT |
SQL_FN_STR_RTRIM |
SQL_FN_STR_SPACE |
SQL_FN_STR_SUBSTRING |
SQL_FN_STR_UCASE

SQL_SUBQUERIES

SQL_SQ_COMPARISON |

SQL_SQ_EXISTS |

SQL_SQ_IN |

SQL_SQ_QUANTIFIED |

SQL_SQ_CORRELATED_SUBQUERIES

SQL_SYSTEM_FUNCTIONS

0

SQL_TABLE_TERM

"TABLE"

SQL_TIMEDATE_ADD_INTERVALS

0

SQL_TIMEDATE_DIFF_INTERVALS

0

SQL_TIMEDATE_FUNCTIONS

SQL_FN_TD_CURDATE |
SQL_FN_TD_CURTIME |
SQL_FN_TD_DAYOFMONTH |
SQL_FN_TD_DAYOFWEEK |
SQL_FN_TD_DAYOFYEAR |
SQL_FN_TD_HOUR |
SQL_FN_TD_MINUTE |
SQL_FN_TD_MONTH |
SQL_FN_TD_NOW |
SQL_FN_TD_SECOND |
SQL_FN_TD_WEEK |
SQL_FN_TD_YEAR
(Microsoft Access, Microsoft Excel, Text)

SQL_FN_TD_DAYOFMONTH |
SQL_FN_TD_DAYOFWEEK |
SQL_FN_TD_DAYOFYEAR |
SQL_FN_TD_HOUR |
SQL_FN_TD_MINUTE |
SQL_FN_TD_MONTH |
SQL_FN_TD_SECOND |
SQL_FN_TD_WEEK |
SQL_FN_TD_YEAR
(dBASE, Microsoft FoxPro, Paradox)

SQL_TXN_CAPABLE

SQL_TC_ALL
(Microsoft Access)

SQL_TC_NONE
(dBASE, Microsoft Excel, Microsoft FoxPro, Paradox, Text)

SQL_TXN_ISOLATION_OPTION

SQL_TXN_READ_COMMITTED
(Microsoft Access)

0
(dBASE, Microsoft Excel, Microsoft FoxPro, Paradox, Text)

SQL_UNION

SQL_U_UNION_ALL |

SQL_U_UNION

SQL_USER_NAME

Not supported

Supported ODBC API Functions

The Microsoft ODBC drivers support all Core and Level 1 functions and the following Level 2 functions:

- **SQLDataSources**
- **SQLDrivers**
- **SQLExtendedFetch**
- **SQLMoreResults**
- **SQLNativeSQL**
- **SQLNumParams**
- **SQLProcedures**
- **SQLProcedureColumns**
- **SQLSetPos**
- **SQLSetScrollOptions**

SQLBindParameter

When the **Microsoft Excel** driver is used, executing an INSERT statement that uses a parameter to insert a NULL into a SQL_CHAR column will return SQL_SUCCESS_WITH_INFO with SQLSTATE 01004, "Data Truncated."

SQLColAttributes

Attribute	Comments
SQL_COLUMN_DISPLAY_SIZE	For LONGVARBINARY data, SQL_COLUMN_DISPLAY_SIZE is the maximum length of the column, not the maximum length of the column times 2.
SQL_COLUMN_MONEY	FALSE is returned for a column of Paradox data type CURRENCY, which maps to the ODBC SQL data type SQL_DOUBLE.
SQL_OWNER_NAME	An empty string ("") is returned in this column, since owner name is not supported.
SQL_QUALIFIER_NAME	The path to a database file is returned for Microsoft Access ; the path to a directory is returned for dBASE , Microsoft Excel , Microsoft FoxPro , Paradox , or Text .
SQL_COLUMN_SEARCHABLE	LONGVARBINARY and LONGVARCHAR columns are reported as SQL_UNSEARCHABLE. Fixed-length and variable-length binary and character data types are seachable, even though LONGVARBINARY and LONGVARCHAR are not.

Note The above is not a complete list of the attributes returned by **SQLColAttributes**.

SQLColumns

Column	Comments
TABLE_QUALIFIER	The path to a database file is returned for Microsoft Access ; the path to a directory is returned for dBASE , Microsoft Excel , Microsoft FoxPro , Paradox , or Text .
TABLE_OWNER	NULL is returned in this column, since owner name is not supported.
NULLABLE	SQL_NO_NULLS is returned for columns that participate in a primary key or unique index.

SQLConfigDataSource

[See Also](#)

The **SQLConfigDataSource** function that is used to add, modify, or delete a data source dynamically uses the following keywords.

Keyword	Description
CHARACTERSET	For the Text driver, OEM or ANSI.
COLLATINGSEQUENCE	The sequence in which the fields are sorted. When the dBASE or Microsoft FoxPro driver is used, the sequence can be: ASCII (the default) or International. When the Paradox driver is used, the sequence can be: ASCII (default), International, Swedish-Finnish, or Norwegian-Danish. This sets the same option as the Collating Sequence control in the setup dialog box.
COLNAMEHEADER	For the Text driver, indicates whether the first record of data will specify the column names. Either TRUE or FALSE.
COMPACT_DB	For the Microsoft Access driver, performs <u>data compaction</u> on a database file.
CREATE_DB	For the Microsoft Access driver, creates a database file. Has the following format: CREATE_DB=<path-name><optional-sort-order>, where the path name is the full path to a Microsoft Access database. An error will be returned if the path name specifies an existing database. The sort order will be as set up in the New Database dialog box displayed when the Create button is pressed in the Microsoft Access Setup dialog box. If no sort order is specified, General is used. When using the CREATE_DB keyword, if the pathname of the Microsoft Access database to be created contains one or more spaces, then the entire pathname must be enclosed by double quotation marks, as shown in the following examples: "C:\PROGRAM FILES\COMMON FILES\ MyAccess.mdb" "C:\PROGRAM FILES\Access2.mdb" CREATE_DB=C:\TEMP\test.mdb (no quotation marks needed)
CREATE_SYSDB	For the Microsoft Access driver, creates a system database file. Has the following format: CREATE_SYSDB=<path-name><optional-sort-order>, where the path name is the full path to a Microsoft Access database. An error will be returned if the path name specifies an existing database. The sort order will be as set up in the New Database dialog box displayed when the Create button is pressed in the Microsoft Access Setup dialog box. If no sort order is specified, General is used.
CREATE_V2DB	For the Microsoft Access driver, creates a database

file that is compatible with Microsoft Access 2.0. Has the following format: CREATE_V2DB=<path-name><optional-sort-order>, where the path name is the full path to a Microsoft Access database. An error will be returned if the path name specifies an existing database. The sort order will be as set up in the New Database dialog box displayed when the Create button is pressed in the Microsoft Access Setup dialog box. If no sort order is specified, General is used.

When using the CREATE_V2DB keyword, if the pathname of the Microsoft Access database to be created contains one or more spaces, then the entire pathname must be enclosed by double quotation marks, as shown in the following examples:

“C:\PROGRAM FILES\COMMON FILES\
MyAccess.mdb”

“C:\PROGRAM FILES\Access2.mdb”

CREATE_V2DB=C:\TEMP\test.mdb (no quotation marks needed)

DBQ

For the **Microsoft Access** or **Paradox** driver, the name of the database file.

For the **Microsoft Excel** driver when accessing **Microsoft Excel** 5.0, 7.0, or 97 files, the name of the workbook file.

This sets the same option as the **Database** control in the setup dialog box.

DEFAULTDIR

The path specification to the database file (for the **Microsoft Access** driver) or directory (for the **dBASE**, **Microsoft Excel**, **Microsoft FoxPro**, **Paradox**, or **Text** driver).

This sets the same option as the **Select Directory control** in the setup dialog box, or for **Microsoft Excel**, the **Select Workbook** control in the setup dialog box.

DELETED

For the **dBASE** or **Microsoft FoxPro** driver, specifies whether or not rows that have been marked as deleted can be retrieved or positioned on. If set to 1, deleted rows are not displayed; if set to 0, deleted rows are treated the same as non-deleted rows. The default is unchecked.

This sets the same option as the **Show Deleted Rows** control in the setup dialog box.

DESCRIPTION

A description of the data in the data source.

This sets the same option as the **Description** control in the setup dialog box.

DRIVER

The path specification to the driver DLL.

DRIVERID

An integer ID for the driver.

25 (Microsoft Access)

21 (dBASE III)

277 (dBASE IV)

533 (dBASE 5.0)

534 (Microsoft Excel 3.0)

	278 (Microsoft Excel 4.0) 22 (Microsoft Excel 5.0/7.0) 790 (Microsoft Excel 97) 24 (Microsoft FoxPro 2.0) 280 (Microsoft FoxPro 2.5) 536 (Microsoft FoxPro 2.6) 26 (Paradox 3.x) 282 (Paradox 4.x) 538 (Paradox 5.x) 27 (Text)
EXCLUSIVE	<p>Determines whether the database will be opened in exclusive mode (accessed by only one user at a time) or shared mode (accessed by more than one user at a time). Can be true (exclusive mode) or false (shared mode). (This does not apply to the Microsoft Excel or Text driver.)</p> <p>This sets the same option as the Exclusive control in the setup dialog box.</p>
EXTENSIONS	<p>Lists the file name extensions of the Text files on the data source.</p> <p>This sets the same option as the Extensions List control in the setup dialog box.</p>
FIL	<p>File type</p> <ul style="list-style-type: none"> MS Access for Microsoft Access dBase III, dBase IV, or dBase 5 Excel 3.0, 4.0, 5.0, 7.0, or 97 for Microsoft Excel FoxPro 2.0, 2.5, or 2.6 for Microsoft FoxPro Paradox 3.x, 4.x, or 5.x Text.
FILETYPE	File type for the Text driver (Text).
FIRSTROWHASNAMES	For the Microsoft Excel driver, indicates whether the cells of the first row of the range contain the column names for the table (1) or not (0).
FORMAT	For the Text driver, can be FIXEDLENGTH, TABDELIMITED, CSVDELIMITED (by a comma), or DELIMITED() (by the special character specified in the parentheses). The special character is one character in length and can be in character, decimal, or hexadecimal format.
IMPLICITCOMMITSYNC	<p>Determines whether the Microsoft Access driver will perform internal or implicit commits asynchronously. This value is initially set to "Yes", which means that the Microsoft Access driver will wait for commits in an internal/implicit transaction to be completed.</p> <p>The value of this option should not be changed without careful consideration of the consequences. For more information on the option, see the the Microsoft Jet Database Engine Programmers's Guide.</p> <p>This sets the same option as the ImplicitCommitSync control in the setup dialog box.</p>
MAXBUFFERSIZE	The size of the internal buffer, in kilobytes, that is used by Microsoft Access to transfer data to and from the disk. The default buffer size is 512K (displayed as

	512). Any integer value divisible by 256 can be used. This sets the same option as the Buffer Size control in the setup dialog box.
MAXSCANROWS	<p>For the Microsoft Access, Microsoft Excel, or Text driver, the number of rows to be scanned when setting a column's data type based upon existing data.</p> <p>A number from 1 to 16 can be entered for the rows to scan. The value defaults to 8; if it is set to 0, all rows are scanned. (A number outside the limit will return an error.)</p> <p>For the Text driver, you can enter a number from 1 to 32767 for the number of rows to scan; however, the value will always default to 25. (A number outside the limit will return an error.)</p> <p>This sets the same option as the Rows to Scan Control in the setup dialog box.</p>
PAGETIMEOUT	<p>Specifies the period of time, in tenths of a second, that a page (if not used) remains in the buffer before being removed. For the Microsoft Access driver, the default is 5 tenths of a second (0.5 seconds). For the dBASE, Microsoft FoxPro, and Paradox drivers, the default is 600 tenths of a second (60 seconds). Note that this option applies to all data sources that use the ODBC driver.</p> <p>This sets the same option as the Page Timeout control in the setup dialog box.</p>
PARADOXNETPATH	<p>For the Paradox driver, the full path of the directory containing a Paradox lock database, because it contains either the PDOXUSRS.net file (in Paradox 4.x) or the PARADOX.net file (in Paradox 5.x). If the directory does not contain one of these files, the Paradox driver creates one. For information about these files, see the Paradox documentation.</p> <p>Before a network directory can be selected, a Paradox user name must be entered.</p> <p>This sets the same option as the Select Network Directory control in the Paradox setup dialog box.</p>
PARADOXNETSTYLE	<p>For the Paradox driver, the network access style to use when accessing Paradox data: either "3.x" for Paradox 3.x or "4.x" for Paradox 4.x or 5.x. Can be set to "3.x" or "4.x" if the version is Paradox 4.x or 5.x; if the version is Paradox 3.x, the style must be "3.x".</p> <p>This sets the same option as the Net Style control in the Paradox setup dialog box.</p>
PARADOXUSERNAME	<p>For the Paradox driver, the Paradox user name.</p> <p>This sets the same option as the User Name control in the Paradox setup dialog box.</p>
PWD	<p>For the Microsoft Access or Paradox driver, the password.</p> <p>For the Paradox driver, this is an optional keyword and will never be written to the file by the driver. It is used in a call to SQLDriverConnect against</p>

READONLY	<p>password-secured Paradox files.</p> <p>TRUE to make file read-only; FALSE to make file not read-only.</p> <p>This sets the same option as the Read-Only control in the setup dialog box.</p>
REPAIR_DB	<p>For the Microsoft Access driver, <u>repairs a database</u> damaged by a failure that occurs during the commit process.</p>
STATISTICS	<p>For the dBASE or Microsoft FoxPro driver, determines whether table size statistics are approximated. Note that this option applies to all data sources that use the ODBC driver.</p> <p>This sets the same option as the Approximate Row Count control in the setup dialog box.</p>
SYSTEMDB	<p>For the Microsoft Access driver, the path specification to system database file.</p> <p>This sets the same option as the System Database control in the setup dialog box.</p>
THREADS	<p>The number of background threads for the engine to use. For the Microsoft Access driver, this value defaults to 3, but can be changed. For the dBASE, Microsoft Excel, Microsoft FoxPro, Paradox, and Text drivers, this value is 3, and cannot be changed.</p> <p>This sets the same option as the Threads control in the setup dialog box.</p>
UID	<p>For the Microsoft Access driver, the user ID name used for login.</p>
USERCOMMITSYNC	<p>Determines whether the Microsoft Access driver will perform user-defined transactions asynchronously. This value is initially set to "Yes", which means that the Microsoft Access driver will wait for commits in a user-defined transaction to be completed.</p> <p>The value of this option should not be changed without careful consideration of the consequences. For more information on the option, see the the Microsoft Jet Database Engine Programmers's Guide.</p> <p>This sets the same option as the UserCommitSync control in the setup dialog box.</p>

SQLDriverConnect

[See Also](#)

SQLDriverConnect enables you to connect to a driver without creating a data source (DSN).

The following keywords are supported in the [connection string](#) for all drivers: **DSN**, **DBQ**, and **FIL**.

For the **Microsoft Access** driver, the **UID** and **PWD** keywords are also supported. For the **Paradox** driver, the **PWD** keyword is also supported.

When the **Microsoft Access** or **Paradox** driver is used, the PWD keyword should not include any of the special characters (see SQL_SPECIAL_CHARACTERS in [SQLGetInfo Returned Values](#)).

When the **Paradox** driver is used, after a password-protected file has been opened by a user, other users are not allowed to open the same file.

The following table shows the minimum keywords required to connect to each driver, and provides an example of keyword/value pairs used with **SQLDriverConnect**. For a full list of DRIVERID values, see [SQLConfigDataSource implementation](#).

Note If DBQ or DefaultDir is not specified for the **dBASE**, **Microsoft Excel 3.0/4.0**, **Microsoft FoxPro**, **Paradox**, or **Text** driver, the driver will connect to the current directory.

Driver	Keywords Required	Examples
Microsoft Access	Driver, DBQ	Driver={Microsoft Access Driver (*.mdb)}; DBQ=c:\temp\sample.mdb
dBASE	Driver, DriverID	Driver={Microsoft dBASE Driver (*.dbf)}; DBQ=c:\temp; DriverID=277
Microsoft Excel 3.0 or 4.0	Driver, DriverID	Driver={Microsoft Excel Driver (*.xls)}; DBQ=c:\temp; DriverID=278
Microsoft Excel 5.0/7.0	Driver, DriverID, DBQ	Driver={Microsoft Excel Driver (*.xls)}; DBQ=c:\temp\sample.xls; DriverID=22
Microsoft Excel 97	Driver, DriverID, DBQ	Driver={Microsoft Excel Driver (*.xls)}; DBQ=c:\temp\sample.xls; DriverID=790
Microsoft FoxPro	Driver, DriverID	Driver={Microsoft FoxPro Driver (*.dbf)}; DBQ=c:\temp; DriverID=536
Paradox	Driver, DriverID	Driver={Microsoft Paradox Driver (*.db)}; DBQ=c:\temp; DriverID=26
Text	Driver	Driver={Microsoft Text Driver (*.txt;*.csv)}; DefaultDir=c:\temp

SQLGetCursorName

SQLGetCursorName is supported, but can only be used when the Cursor Library is used, because positioned operations are not supported in the driver.

SQLGetData

This function can retrieve data from any column, whether or not there are bound columns after it and regardless of the order in which the columns are retrieved.

SQLGetInfo

SQLGetInfo supports the SQL_FILE_USAGE information type. The returned value is a 16-bit integer that indicates how the driver directly treats files in a data source:

SQL_FILE_NOT_SUPPORTED = The driver is not a single-tier driver.

SQL_FILE_TABLE = A single-tier driver treats files in a data source as tables.

SQL_FILE_QUALIFIER = A single-tier driver treats files in a data source as a qualifier.

The ODBC driver returns SQL_FILE_QUALIFIER for the **Microsoft Access** driver, because each file is a complete database. The ODBC driver returns SQL_FILE_TABLE for the **dBASE, Microsoft Excel, Microsoft FoxPro, Paradox, or Text** driver, because each file is a table.

SQLGetStmtOption

The bookmarks returned by an *fOption* of SQL_GETBOOKMARK are only valid while the query is open and are invalidated when the query is reissued. Persistent bookmarks are not supported.

SQLGetTypeInfo

The name of the type (TYPE_NAME) returned in the table produced by **SQLGetTypeInfo** will be the name most commonly used by the data source.

SQL_ALL_EXCEPT_LIKE will be returned in the SEARCHABLE column for the Byte, Counter, Double, Single, Long, and Short data types. (The LIKE capability can be achieved by converting the value to a character using the ODBC canonical conversion functions, then performing the comparison.)

When the **Microsoft Excel** driver is used, the ODBC type names are returned in the TYPE_NAME column that is returned by **SQLGetTypeInfo**.

When the **Text** driver is used, **SQLGetTypeInfo** returns a CASE_SENSITIVE value of FALSE for the text data types (CHAR and LONGCHAR), when the data types actually are case-sensitive.

SQLMoreResults

This function always returns SQL_NO_DATA_FOUND.

SQLPrepare

A statement prepared with an incorrect number of columns returns an error at execution time, not upon statement preparation.

SQLProcedureColumns

Supported for **Microsoft Access** DSNs only.

Application developers should look for driver-defined columns starting at the end of the result set and proceeding backward.

Column	Comments
COLUMN_TYPE	SQL_PARAM_INPUT or SQL_RESULT_COL
ORDINAL	This is a driver-specific column that is returned at the end of the result set. The SQL type of the column is an integer.

SQLProcedures

SQLProcedures will only return rows for those procedures that have at least one argument. Procedures that have no arguments are treated as views.

Column	Comments
PROCEDURE_QUALIFIER	The path to the database file.
PROCEDURE_OWNER	NULL
PROCEDURE_NAME	Undelimited procedure name
PROCEDURE_TYPE	SQL_PT_PROCEDURE

SQLSetConnectOption

<i>fOption</i>	Comment
SQL_ACCESS_MODE	The SQL_ACCESS_MODE <i>fOption</i> can be set to either SQL_MODE_READ_ONLY or SQL_MODE_READ_WRTE. However, the driver does not prevent updates if SQL_ACCESS_MODE is set to SQL_MODE_READ_ONLY.
SQL_AUTOCOMMIT	<p>When the Microsoft Access driver is used, the SQL_AUTOCOMMIT option may be set to either SQL_AUTOCOMMIT_ON or SQL_AUTOCOMMIT_OFF, because the Microsoft Access driver supports transactions(1).</p> <p>The dBASE, Microsoft Excel, Microsoft FoxPro, Paradox, and Text drivers only support SQL_AUTOCOMMIT being set to ON (the default state), because they do not support transactions.</p>
SQL_CURRENT_QUALIFIER	Supported.
SQL_LOGIN_TIMEOUT	Not supported.
SQL_OPT_TRACE	Supported.
SQL_OPT_TRACEFILE	Supported.
SQL_PACKET_SIZE	Not supported.
SQL_QUIET_MODE	Not supported.
SQL_TRANSLATE_DLL	Not supported.
SQL_TRANSLATION_OPTION	Not supported.
SQL_TXN_ISOLATION	When the Microsoft Access driver is used, SQL_TXN_ISOLATION is always SQL_TXN_READ_COMMITTED. This <i>fOption</i> is not supported for the other drivers.

1 Atomic transactions are not supported by the **Microsoft Access** driver. When committing a transaction using the **Microsoft Access** driver, a finite delay exists between the time the transaction is committed and the time the values are written to disk. This delay is determined by a delay inherent in the Microsoft Jet engine. The page timeout will not be less than a minimum value, even if the PageTimeout option is set below that value. As a result, there is no guarantee that committed data is stable, since changes may be made during the delay.

SQLSetCursorName

Because the driver does not support a positioned update or delete by the WHERE CURRENT OF *cursorname* syntax, **SQLSetCursorName** is supported, but cannot be used for positioned updates. It can only be used when the Cursor Library is enabled and the application is using **SQLExtendedFetch**.

SQLSetPos

The bulk-model semantics for **SQLSetPos** calls with the *row* argument equal to 0 are supported.

SQL_LOCK_NO_CHANGE is supported for *fLock*. SQL_LOCK_EXCLUSIVE and SQL_LOCK_UNLOCK are not supported.

SQLSetPos supports updatable joins. (For more information, see the Microsoft Jet Database Engine Programmer's Guide.)

SQLSetScrollOptions

Forward and static cursors are supported for SQL_CONCUR_READ_ONLY.

Only keyset-driven cursors are supported for an *fConcurrency* argument of SQL_CONCUR_LOCK.

An *fConcurrency* argument of SQL_CONCUR_ROWVER is not supported.

Dynamic cursors and mixed cursors are not supported.

SQLSetStmtOption

<i>fOption</i>	Comments
SQL_ASYNC_ENABLE	Asynchronous processing is not supported. The SQL_ASYNC_ENABLE <i>fOption</i> will return SQLSTATE S1C00 (Driver not capable).
SQL_KEYSET_SIZE	The only valid keyset size is 0, because mixed and dynamic cursors are not supported. If this value is set to any other number, it will be changed to 0 and the call will return SQL_SUCCESS_WITH_INFO and SQLSTATE 01S02 (Option value changed).
SQL_MAX_ROWS	The only valid rowset size is 0, because the Desktop Database Drivers do not support limiting the number of rows that are returned. If this value is set to any other number, it will be changed to 0 and the call will return SQL_SUCCESS_WITH_INFO and SQLSTATE 01S02 (Option value changed).
SQL_QUERY_TIMEOUT	Not supported.
SQL_ROW_NUMBER	Not supported.
SQL_SIMULATE_CURSOR	Not supported.

SQLSpecialColumns

A unique index will be returned (if one exists) for the SQL_BEST_ROWID flag in *fColType*. No result set will be returned for the SQL_ROWVER flag.

All row IDs have a scope of SQL_SCOPE_CURROW.

Pattern matching is not supported for either the *szTableQualifier* or *szTableName* argument.

SQLStatistics

Column	Comments
TABLE_QUALIFIER	The path to a database file is returned for Microsoft Access; the path to a directory is returned for dBASE, Microsoft Excel, Microsoft FoxPro, Paradox, or Text. Pattern matching is not supported in the <i>szTableQualifier</i> argument.
TABLE_OWNER	NULL is returned in this column, since owner name is not supported.
TABLE_NAME	Undelimited table name. Pattern matching is not supported in the <i>szTableName</i> argument.
INDEX_QUALIFIER	NULL is always returned.
INDEX_NAME	Index-dependent.
TYPE	Only SQL_TABLE_STAT or SQL_INDEX_OTHER will be returned for TYPE.
SEQ_IN_INDEX	Index-dependent.
COLUMN_NAME	Index-dependent.
COLLATION	Index-dependent.
CARDINALITY	Returned for Microsoft Access only.
PAGES	NULL is always returned.

Filtering is based on uniqueness (the *fUnique* argument). The *fAccuracy* parameter is ignored.

SQLTables

Argument	Comments
<i>szTableOwner</i>	The only valid argument for <i>szTableOwner</i> is NULL, since none of the drivers support owner names. With <i>szTableOwner</i> set to NULL, all tables are returned. NULL is returned in the TABLE_OWNER column.
<i>szTableQualifier</i>	<p>When the Microsoft Excel 3.0 or 4.0 driver is used, if you call SQLTables with a value for <i>szTableQualifier</i> that is not the name of an existing table, the driver will create a table with that name.</p> <p>In the TABLE_QUALIFIER column, SQLTables will return either the path to a database file (for Microsoft Access or Microsoft Excel 5.0, 7.0, or 97) or the path to a directory (dBASE, Microsoft Excel 3.0 or 4.0, Microsoft FoxPro, Paradox, or Text).</p>
<i>szTableType</i>	<p>When the Microsoft Access driver is used, "SYSTEM TABLE" is supported for <i>szTableType</i> for system tables, "SYNONYM" is supported for attached tables, and "VIEW" is supported for row-returning queries.</p> <p>For dBASE, Microsoft Excel 3.0 or 4.0, Microsoft FoxPro, Paradox, or Text files, "TABLE" is the only table type supported.</p> <p>For Microsoft Excel 5.0, 7.0, or 97 files, "SYSTEM TABLE" is returned for sheet names (tables with a "\$" on the end), and "TABLE" is returned for tables within worksheets.</p> <p>When the Text driver is used, the list of files returned by SQLTables is determined by the file extensions in the Extensions List box in the ODBC Text Setup dialog box.</p>

SQLTransact

When the **Microsoft Access** driver is used, SQL_COMMIT and SQL_ROLLBACK are supported for the *fType* argument in a call to **SQLTransact**.

When the **dBASE, Microsoft Excel, Microsoft FoxPro, Paradox, or Text** driver is used, this function supports SQL_COMMIT, but not SQL_ROLLBACK.

When the **Microsoft Access** driver is used, if a failure occurs during the commit process, the affected database can be repaired using the [Repair Database](#) option in the Microsoft Access driver setup, or through the use of the REPAIR_DB keyword in the SQLConfigDataSource function.

Arithmetic Errors

The ODBC driver evaluates the WHERE clause in a SELECT statement as it fetches each row. If a row contains a value that causes an arithmetic error, such as divide-by-zero or numeric overflow, the driver returns all rows, but returns errors for columns with arithmetic errors. When inserting or updating, however, the ODBC driver stops inserting or updating data when the first arithmetic error is encountered.

Base Address of Drivers

Each of the ODBC Desktop Database Drivers has been linked with a base address of 0x04c0000.

Creating and Opening Tables

When the **Text** driver is used, a new table is created using the format specified in ODBCINST.INI. If not specified, tables are created in CSVDELIMITED format. By default, INTEGER columns default to 11 characters and FLOAT columns default to 22 characters. DATE columns use the YYYY-MM-DD format. CHAR and LONGCHAR columns are the width specified in the CREATE statement.

Multiple *hstmts*

When the ODBC **Paradox** driver is used, if you want to use more than one *hstmt* to execute queries on a table, the table must have a unique index (Paradox primary key).

Read-Only Status

When the **Microsoft Excel** or **Text** driver is used, data source tables are read-only. Even though tables have read-only status, however, applications can perform insertions and updates for Microsoft Excel tables, or insertions for Text tables.

When an application performs a Save As command on Microsoft Excel data through the **Microsoft Excel** driver, the application should create a new table and insert the data to be saved into the new table. Inserts result in an append to the table. No other operations can be performed on the table until it is closed and reopened. Once the table is closed, no subsequent insert can be performed, since the table is then a read-only table.

It is possible to update values when using the **Microsoft Excel** driver, but a row cannot be deleted from a table based on a Microsoft Excel spreadsheet, so updates are not considered officially supported by the **Microsoft Excel** driver.

Setup DLL

The ODBC Administrator calls the function **ConfigDSN** when users configure data sources.

Translation DLLs

Not supported.

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.

