



DataDirect ODBC Paradox Driver

[About the Paradox Driver](#)

[System Requirements](#)

[Multuser Access to Tables](#)

[Locking](#)

[Configuring Data Sources](#)

[Connecting to Paradox Using a Connection String](#)

[Data Types](#)

[Select Statement](#)

[Password Protection](#)

[Index Files](#)

[Create and Drop Index Statements](#)

[Transactions](#)

[Isolation and Lock Levels Supported](#)

[ODBC Conformance Level](#)

[Number of Connections and Statements Supported](#)

[Copyright](#)

Copyright 1996 INTERSOLV Inc. All rights reserved. INTERSOLV is a registered trademark, and DataDirect is a trademark of INTERSOLV, Inc. Other company or product names mentioned herein may be trademarks or registered trademarks of their respective companies.

About the Paradox Driver

The Paradox driver supports Paradox 3.0, 3.5, 4.0, 4.5, 5.0, and 7.0 tables.

The driver filename is `IVIDPnn.DLL`.

System Requirements

To use the Paradox driver, you must install the Borland Database Engine, which conforms to the IDAPI programming interface. The Borland Database Engine can be found in any of the following software packages:

- Borland C++ for Windows NT/95, version 5
- Delphi for Windows NT/95
- Paradox 7 for Windows NT/95

If you attempt to configure a data source and you do not have IDAPI32.DLL on your path or in your Windows 95 SYSTEM directory or Windows NT SYSTEM32 directory, the following message appears:

The setup routines for the INTERSOLV ParadoxFile (*.db) ODBC driver could not be loaded. You may be low on memory and need to quit a few applications.

Multiusers Access to Tables

You can query Paradox tables that reside in a shared directory on a network or that are to be shared among applications running on a local workstation. If the tables are on a network server, multiple users can query these tables simultaneously. To share Paradox tables among multiple users, the tables must be located in a shared directory on your network server.

Two connection attributes identify the Paradox database you are accessing: Database (database directory) and NetDir (network directory). The Database setting specifies the directory of Paradox tables that is the database. If the Database setting you specify is a shared network directory, then Paradox requires a NetDir specification as well. This value identifies the directory containing the PARADOX.NET file that corresponds to the Database setting you have specified.

Every user who accesses the same shared directory of tables must set the NetDir value to point to the same PARADOX.NET file. If your connection string does not specify a NetDir value, then Paradox uses the NetDir value specified in the Paradox section of the IDAPI configuration file. This makes it important that the NetDir specification in each user's IDAPI configuration file be set correctly.

Whenever you open a Paradox table that another user opens at the same time, the consistency of the data becomes an issue if both individuals are updating the table.

Paradox Driver Locking

The Paradox driver has two locking levels: record locking and table locking. Tables that have no primary key always have a prevent write lock placed on the table when the table is opened. The table lock is escalated to a write lock when an operation that changes the table is attempted.

Tables that have primary keys use record locking. The locking level is escalated from record locking to a table write lock if the transaction runs out of record locks.

Primary keys provide the greatest concurrency for tables that are accessed and modified by multiple users.

Note If a table lock is placed on a Paradox table, the Paradox driver prevents users from updating and deleting records but does not prevent them from inserting records into the locked table.

Configuring Data Sources

To configure a Paradox data source, do the following:

- 1 Start the ODBC Administrator. A list of data sources appears.
- 2 If you are configuring a new data source, click **Add**. A list of installed drivers appears. Select INTERSOLV ParadoxFile and click **OK**.

If you are configuring an existing data source, select the data source name and click **Setup**.

The [Paradox ODBC Setup](#) dialog box appears.

- 3 Specify a data source name, a database directory and optionally, a description.
- 4 Click the **Advanced** button to configure optional data source settings, such as server list and database list.

The [ODBC Advanced Driver Setup](#) dialog box appears.

- 5 Click **Translate** to perform a translation of your data from one character set to another. The Select Translator dialog box appears in which you can select a translator to perform the translation. INTERSOLV provides a translator named INTERSOLV OEM ANSI that translates your data from the IBM PC character set to the ANSI character set. Click **OK** to leave this dialog box and perform the translation.

The translators that are listed in this dialog box are determined by the values listed in the ODBC Translators section of your ODBCINST.INI file.

- 6 Click **OK** to write these values to ODBC.INI. These values are now the defaults when you connect to the data source. You can change the defaults by configuring your data source again. You can override the defaults by connecting to the data source using a connection string with alternate values.

Connecting to Paradox Using a Connection String

If your application requires a connection string to connect to a data source, you must specify the data source name that tells the driver which ODBC.INI section to use for the default connection information. Optionally, you may specify *attribute=value* pairs in the connection string to override the default values stored in ODBC.INI. These values are not written to ODBC.INI.

You can specify either long or short names in the connection string. The connection string has the form:

```
DSN=data_source_name[;attribute=value[;attribute=value]...]
```

An example of a connection string for Paradox is

```
DSN=PARADOX TABLES;DB=C:\ODBC\EMP;PW=ABC,DEF,GHI
```

The following table gives the long and short names for each attribute, as well as a description.

The defaults listed in the table are initial defaults that apply when no value is specified in either the connection string or in the data source definition in ODBC.INI. If you specified a value for the attribute when configuring the data source, that value is your default.

Attribute	Description
DataSourceName (DSN)	A string that identifies a Paradox data source configuration in ODBC.INI. Examples include "Accounting" or "Paradox5-Serv1."
Database (DB)	The directory in which the Paradox files are stored. For this attribute, you can also specify aliases that are defined in your IDAPI configuration file, if you have one. To do this, enclose the alias name in colons. For example, to use the alias MYDATA, specify "Database=:MYDATA:"
NetDir (ND)	The directory containing the PARADOX.NET file that corresponds to the database you have specified. If the Paradox database you are using is shared on a network, then every user who accesses it must set this value to point to the same PARADOX.NET file. If not specified, this value is determined by the NetDir setting in the Paradox section of the IDAPI configuration file. If you are not sure how to set this value, contact your network administrator.
FileOpenCache (FOC)	The maximum number of unused table opens to cache. For example, when FileOpenCache=4, and a user opens and closes 4 tables, the tables are not actually closed. The driver keeps them open so that if another query uses one of these tables, the driver does not have to perform another open, which is expensive. The advantage of using file open caching is increased performance. The disadvantage is that a user who tries

to open the table exclusively may get a locking conflict even though no one appears to have the table open. The initial default is 0.

CreateType (CT)

CreateType={3 | 4 | 5 | 7 | null (blank)}. The table version for Create Table statements. There are five valid values for this connection string: 3, 4, 5, 7 and null (blank). The numeric values map to the major revision numbers of the Paradox family of products. To override another CreateType setting chosen during data source configuration with the default create type determined by the Level setting in the Paradox section of the IDAPI configuration file, set CreateType=(null).

Note: When CreateType is set to 7, the Paradox driver supports table names up to 128 characters long. For all other Create Type settings, the driver supports table names up to 8 characters long.

IntlSort (IS)

IntlSort={0 | 1}. A value that determines the order that records are retrieved when you issue a Select statement. If IntlSort=1, the driver uses the international sort order as defined by your operating system. The sort is case-insensitive (*a* precedes *B*); the sorting of accented characters is also affected (check your operating system documentation). If IntlSort=0 (the default), the driver uses the ASCII sort order, where uppercase letters precede lowercase letters (*B* precedes *a*).

UltraSafeCommit (USF)

UltraSafeCommit={0 | 1}. A value that determines when the driver flushes its changes to disk. If UltraSafeCommit=1, the driver does this at each COMMIT. This decreases performance. The default is 0. This means that the driver flushes its changes to disk when the table is closed or when internal buffers are full. In this case, a machine "crash" before closing a table may cause recent changes to be lost.

Passwords (PW)

A password or list of passwords. You can add 1 to 50 passwords into the system using a comma-separated list of passwords. Passwords are case-sensitive. For example, Passwords=psw1, psw2, psw3.

DeferQueryEvaluation
(DQ)

DeferQueryEvaluation={0 | 1}. A value that determines when a query is evaluated-- after all records are read or each time a record is fetched.

If DeferQueryEvaluation=0, the driver generates a result set when the first record is fetched. The driver reads all records, evaluates each one against the Where clause, and compiles a result set containing the records that satisfy the search criteria. This process slows performance when the first record is fetched, but activity performed on the result set after this point is much faster, since the result set has already been created. You do not see any additions, deletions, or changes in the database that occur while working from this result set.

If DeferQueryEvaluation=1 (the default), the driver evaluates the query each time a record is fetched, and stops reading through the records when it finds one that matches the search criteria. This setting avoids the slowdown while fetching the first record, but each fetch takes longer because of the evaluation taking place. The data you retrieve reflects the latest changes to the database. However, a result set is still generated if the query is a Union of multiple Select statements, if it contains the Distinct keyword, or if it has an Order By or Group By clause.

Data Types

The following table shows how the Paradox data types are mapped to the standard ODBC data types. It also identifies the types supported by Paradox versions 3.x and 4.x. These Paradox data types can be used in a [Create Table statement](#).

Paradox	ODBC	3.x Support	4.x Support
Alpha	SQL_VARCHAR	Yes	Yes
AutoIncrement	SQL_INTEGER	No	No
BCD	SQL_DECIMAL	No	No
Binary*	SQL_LONGVARBINARY	No	Yes
Bytes*	SQL_BINARY	No	No
Date	SQL_DATE	Yes	Yes
Formatted Memo*	SQL_LONGVARBINARY	No	Yes
Graphic*	SQL_LONGVARBINARY	No	Yes
Logical*	SQL_BIT	No	No
Long Integer	SQL_INTEGER	No	No
Memo*	SQL_LONGVARCHAR	No	Yes
Money	SQL_DOUBLE	Yes	Yes
Number	SQL_DOUBLE	Yes	Yes
OLE*	SQL_LONGVARBINARY	No	Yes
Short	SQL_SMALLINT	Yes	No
Time	SQL_TIME	No	No
TimeStamp	SQL_TIMESTAMP	No	No

* Cannot be used in an index. Of these types, only Logical can be used in a Where clause.

SQL Statements for the Paradox Driver

You use a SQL Select statement to specify the columns and records to be read. The Paradox driver supports all of the [Select statement](#) clauses. The following topics provide additional information about using SQL with the Paradox driver:

[Column Names](#)

[Alter Table Statement](#)

Column Names

Paradox column names are case-insensitive and their maximum length is 25 characters. If a column name contains a special character, does not begin with a letter, or is a reserved word, surround it with the grave character (`) (ASCII 96). For example,

```
SELECT `last name` FROM emp
```

Alter Table Statement

The dBASE driver supports the Alter Table statement to add one or more columns to the table, or to delete (drop) a single column.

The Alter Table statement has the form:

```
ALTER TABLE table_name {ADD column_name data_type | ADD (column_name data_type [, column_name
data_type] . . . ) |
DROP [COLUMN] column_name [CASCADE | RESTRICT]}
```

table_name is the name of the table for which you are adding or dropping columns.

column_name assigns a name to the column you are adding or specifies the column you are dropping. See [Column Names](#) for more information.

data_type specifies the native data type of each column you add. See [Paradox Data Types](#) for more information.

For example, to add two columns to the emp table,

```
ALTER TABLE emp {ADD startdate date, dept char 10}
```

Dropping Columns

You cannot add and drop columns in a single statement, and you can drop only one column at a time.

When dropping a column, use the CASCADE keyword to drop the column while removing references from any dependent objects, such as indexes or views. Use RESTRICT to cause the Alter Table statement to fail if other objects are dependent upon the column you are dropping. For example, to drop a column and remove its references from dependent objects:

```
ALTER TABLE emp DROP startdate CASCADE
```

If the Alter Table statement contains neither CASCADE nor RESTRICT, it fails when you attempt to drop a column upon which other objects are dependent.

Password Protection

Paradox supports two types of passwords: master and auxiliary. The Paradox driver supports master passwords only and can manage up to 50 passwords.

Paradox tables may be encrypted to provide limited access to users who do not know the password. The driver maintains a list of passwords for each connection. The driver can access only encrypted tables for which a password appears in this list. You can supply a password in three ways: by typing it in the Password dialog box (which appears when the driver does not have the password to open an encrypted table), by including it in a connection string (with the Passwords attribute), or by using the Add Password statement.

Paradox provides five statements that manage passwords for Paradox tables. These statements are specific to the Paradox driver:

```
ENCRYPT filename USING PASSWORD password
ADD PASSWORD password
DECRYPT filename USING PASSWORD password
REMOVE PASSWORD password
REMOVE ALL PASSWORDS
```

filename can be a simple filename or a full pathname. If a simple filename is given, the file must be in the directory specified with the Database connection string attribute. The .DB extension is not required.

password is a case-sensitive text string up to 15 characters in length, including blanks. If your password includes lower-case letters or nonalphanumeric characters, enclose it in single quotation marks (').

Encrypting a Paradox Table

The Encrypt statement associates a password with a table. For example,

```
ENCRYPT emp USING PASSWORD test
```

Accessing an Encrypted Paradox Table

To access an encrypted Paradox table, add the password to the list of passwords Paradox maintains for that connection. To do so, you can

- Issue an Add Password statement before you access the table. For example,
ADD PASSWORD test
SELECT * FROM emp
- Specify the passwords using the Passwords attribute at connection time.

If you do not add the password, the driver prompts you for it when you access the table.

Decrypting a Paradox Table

The Decrypt statement disassociates a password from a table. You no longer need to enter the password to open the table. For example,

```
DECRYPT emp USING PASSWORD test
```

Removing a Password from Paradox

The Remove Password statement removes a password from the list Paradox maintains for the connection. For example,

```
REMOVE PASSWORD test
```

Removing All Passwords from Paradox

The Remove All Passwords statement removes the list of passwords Paradox maintains.

If you remove a password from Paradox and do not decrypt the table, you must continue entering the password to open the table.

Index Files

An index is used to read records in sorted order and to improve performance when selecting records and joining tables. Paradox indexes are stored in separate files and are either *primary* or *non-primary*.

Primary Index

A primary index is made up of one or more fields from the Paradox table. The primary key fields of a primary index consist of one or more consecutive fields in the table, beginning with the first field in the table. A table can have only one primary index.

Collectively, the primary key fields uniquely identify each record in the Paradox table. Thus, no two records in a Paradox table can share the same values in their primary key fields. Once a primary index is created for a Paradox table, the table's records are re-ordered based on the primary key fields. At the time a primary index is created, if any records have matching primary key field values, Paradox deletes all but the first record. Paradox creates this index as maintained; that is, if you modify, add, or delete records in the table, the primary index is updated automatically to reflect these changes.

A primary index is a single file with the same name as the table on which it is based but with a .PX extension.

To lock records, you must have a primary index.

Non-Primary Index

Unlike a primary key field of a primary index, a non-primary key field does not uniquely identify each record. Thus, two or more records in a Paradox table can share the same value in their non-primary key field. A Paradox table can have more than one non-primary index as long as each one is based on different fields.

A non-primary index is defined by specifying one or more fields in the Paradox table that constitute the non-primary key field. It allows Paradox to sort each record in the table according to the values of the non-primary key fields.

There are two kinds of non-primary indexes: maintained and non-maintained. A maintained index is automatically updated when the table is changed, whereas a non-maintained index is not. Instead, a non-maintained index is tagged out-of-date and is updated when the index is used again.

You must have a primary index on a table before you create a maintained, non-primary index.

The Paradox driver uses non-maintained indexes only for read-only queries on locked tables. A primary index is not required for the non-maintained index to be used.

For Paradox 3.x, a single non-primary index consists of a pair of files with the same name as the table on which the non-primary index is based; one of these files has an .Xnn extension while the other has a .Ynn extension (where the hexadecimal number *nn* corresponds to the field number of the non-primary key field for the non-primary index).

For Paradox 4.x, 5, and 7, single-field non-primary indexes that are case sensitive have the same name as their associated table and are assigned file extensions .X01 through .XFF, depending on the number of the field on which the index is based. Single-field non-primary indexes that are case insensitive and composite indexes have the same name as the table on which they are based. They are assigned file extensions sequentially starting with .XG0 (with hexadecimal increments).

Create and Drop Index Statements

The Paradox driver supports SQL statements to create and delete indexes. The Create Index Primary statement is used to create primary indexes. The Create Index statement is used to create non-primary indexes. The Drop Index statement is used to delete indexes.

Create Index Primary Statement

The syntax for creating a primary index is

```
CREATE [UNIQUE] INDEX PRIMARY
      ON table_name (column [,column...])
```

The UNIQUE keyword is optional; the index is unique whether or not you include this keyword.

table_name is the name of the table on which the index is to be based.

column is the name of a column that is included as a the key field for the index. The column list must contain one or more consecutive fields in the table, beginning with the first field in the table.

For example,

```
CREATE UNIQUE INDEX PRIMARY ON emp (emp_id)
```

Be careful when you create a primary key because any rows that have a primary key duplication are deleted when you execute the Create Index Primary statement.

Create Index Statement

The syntax for creating a non-primary index is

```
CREATE [UNIQUE] INDEX index_name [/NON_MAINTAINED] [/CASE_INSENSITIVE] ON table_name (column
[DESC] [, column...])
```

The optional UNIQUE keyword prevents duplicate values in the non-primary index. The UNIQUE keyword is supported for non-primary indexes only for Paradox 7 tables (when the Create Type is 7).

index_name identifies the index. If the name contains blanks or special characters, or does not begin with a letter, surround it with the grave character (`) (ASCII 96).

The NON_MAINTAINED switch makes the index non-maintained. The default is to create a maintained index.

The CASE_INSENSITIVE switch makes the index case-insensitive. The default is to create a case-sensitive index.

table_name is the name of the table on which the index is to be based.

column is the name of a column that is included as a the key field for the index.

The DESC keyword creates a non-primary index that uses descending keys.

Paradox 3.0 and 3.5 tables cannot have composite or case-insensitive indexes. When you create a non-primary index for Paradox 3.0 and 3.5 tables, follow these rules:

- Specify only one column name.
- Do not use the CASE_INSENSITIVE switch.
- Use the column name as the index name.

For example,

```
CREATE INDEX last_name ON emp (last_name)
```

Drop Index Statement

The syntax for dropping a primary index is

```
DROP INDEX path_name.PRIMARY
```

For example,

```
DROP INDEX emp.PRIMARY
```

The syntax for dropping a non-primary index is

```
DROP INDEX path_name.index_name
```

path_name is the name of the table from which the index is being dropped. The pathname can be either the fully qualified pathname or, if the table is specified with the Database attribute of the connection string, a simple table name.

index_name is the name that was given to the index when it was created. If the name contains blanks or special characters, or does not begin with a letter, surround it with the grave character (`) (ASCII 96). Use the column name as the index name when dropping indexes from Paradox 3.5 or 3.0 tables.

For example,

```
DROP INDEX emp.last_name
```

Transactions

The Paradox driver supports transactions. A transaction is a series of database changes that is treated as a single unit. In applications that don't use transactions, the Paradox driver immediately executes Insert, Update, and Delete statements on the database tables and the changes are automatically committed when the SQL statement is executed. There is no way to undo such changes. In applications that use transactions, inserts, updates, and deletes are held until a COMMIT or ROLLBACK is specified. A COMMIT saves the changes to the database file; a ROLLBACK discards the changes.

Transactions affect the removal of record locking. All locks are removed when SQLTransact is called with the Commit or Rollback option to end the active transaction.

Isolation and Lock Levels Supported

Paradox supports isolation levels 1 (read committed) and 3, (serializable). It supports record- and table-level locking.

ODBC Conformance Level

The Paradox driver supports the Core, Level 1, and Level 2 API functions listed in [Supported ODBC Functions](#). In addition, it supports the following Level 2 functions:

- SQLSetPos
- SQLPrimaryKeys

When used with Paradox 5 tables, the Paradox driver also supports the SQLForeignKeys function.

Number of Connections and Statements Supported

The Paradox database system supports multiple connections and multiple statements per connection.

Paradox ODBC Setup Dialog

Use the Paradox ODBC Setup dialog to create new Paradox data sources or configure existing data sources.

Data Source Name: A string that identifies this Paradox data source configuration in ODBC.INI. Examples include "Accounting" or "Paradox5-Serv1."

Description: An optional long description of a data source name. For example, "My Accounting Database" or "Paradox Files on Server number 1."

Database Directory: The directory in which the Paradox files are stored. If a directory is not specified, the current working directory is used. You can also specify aliases that are defined in your IDAPI configuration file, if you have one. To do this, enclose the alias name in colons. For example, to use the alias MYDATA, specify ":MYDATA:"

Advanced

Displays the [ODBC Advanced Driver Setup](#) dialog box to configure optional data source settings, such as server name.

OK

Creates or modifies the current data source using the options you specify.

Cancel

Exits the ODBC Setup dialog without creating or modifying a data source.

Advanced Dialog

To configure optional settings for a Paradox data source, specify values as follows:

Network Directory: The directory containing the PARADOX.NET file that corresponds to the database you have specified. If the Paradox database you are using is shared on a network, then every user who accesses it must set this value to point to the same PARADOX.NET file. If not set here, this value is determined by the NetDir setting in the Paradox section of the IDAPI configuration file. If you are not sure how to set this value, contact your network administrator.

File Open Cache: A numeric value to specify the maximum number of unused file opens to cache. For example, the value 4 specifies that when a user opens and closes four tables, the tables are not actually closed. The driver keeps them open so that if another query uses one of these tables, the driver does not have to perform another open, which is expensive. The advantage of file open caching is increased performance. The disadvantage is that a user who specifies file locking on open may get a locking conflict even though no one appears to have the file open. The default is 0, which means no file open caching.

Create Type: The Paradox table version desired for any Create Table statements that you execute. You can specify 3, 4, 5, 7 or leave the box blank and use the default, which is determined by the Level setting in the Paradox section of the IDAPI configuration file. The numeric values map to the major revision numbers of the Paradox family of products.

Note: If you select Create Type 7, the Paradox driver supports table names up to 128 characters long. For all other Create Type settings, the driver supports table names up to 8 characters long.

Translate

Displays the Select Translator dialog box to allow you to perform a translation of your data from one character set to another. Choose the INTERSOLV OEM ANSI translator to translate your data from the IBM PC character set to the ANSI character set.

Close

Returns to the [Paradox ODBC Setup](#) dialog box where you can click the **OK** button to write these settings to the ODBC.INI file.

Entering Paradox Passwords

Paradox supports two types of passwords: master and auxiliary. The Paradox driver supports master passwords only and can manage up to 50 passwords.

Paradox tables may be encrypted to provide limited access to users who do not know the password. The driver maintains a list of passwords for each connection. The driver can access only encrypted tables for which a password appears in this list. You can supply a password in three ways: by typing it in the Password dialog box (which appears when the driver does not have the password to open an encrypted table), by including it in a connection string (with the Passwords attribute), or by using the Add Password statement.

Entering Passwords with a Dialog Box

The password dialog box appears when the driver does not have the password to open an encrypted table. Entering the password into the dialog box will enable the driver to access that table for the rest of that connection.

Entering Passwords During the Connection

You may enter passwords that may be needed during a connection as a connection option. Use connection attribute "Passwords" to enter a comma-separated list of password names.

Entering Passwords with the Add Password Statement

You may enter passwords by executing the ADD PASSWORD statement. For further help with ADD PASSWORD, consult the *INTERSOLV DataDirect ODBC Drivers Reference* or double-click the Driver help icon.

