ALTER TABLE statement

Adds or deletes a column from a table.

```
ALTER TABLE table_reference DROP [COLUMN] column_reference | ADD [COLUMN]
  column_reference [,reference DROP [COLUMN] column_reference | ADD
  [COLUMN] column_reference...]
```

Description

Use the ALTER TABLE statement to add a column to or delete a column from an existing table. It is possible to delete one column and add another in the same ALTER TABLE statement.

The DROP keyword requires only the name of the column to be deleted. The ADD keyword requires the same combination of column name, type, and possibly dimension definition as CREATE TABLE when defining new columns. See the section on defining column types for the specific syntax of all supported column types.

The statement below deletes the column FULLNAME and adds the column LASTNAME.

```
ALTER TABLE "Names.db"
DROP FullName, ADD LastName CHAR(25)
```

It is possible to delete and add a column of the same name in the same ALTER TABLE statement, however any data in the column is lost in the process. This allows quick redefinition of columns while still in the database design stages.

```
ALTER TABLE "Names.db"
DROP LastName, ADD LastName CHAR(30)
```

If a column to be deleted is part of a primary key, the primary index is deleted. ALTER TABLE fails on an attempt to delete a column that is the target of a foreign key constraint (referential integrity).

To reference columns with non-alphanumeric characters or spaces embedded in the column name, you must enclose the column name in quotation marks and prefix the quoted column name with the table name in quotes.

```
ALTER TABLE "Customer.db"
ADD "Customer.db"."#ID" CHAR(3)
```

Arithmetic operators

Perform arithmetic operations.

```
numeric_value1 + numeric_value2
numeric_value1 – numeric_value2
numeric_value1 * numeric_value2
numeric_value1 / numeric_value2
```

Description

Use arithmetic operators to perform arithmetic calculations on data in SELECT queries. Calculations can be performed wherever non-aggregated data values are allowed, such as in a SELECT or WHERE clause. In the statement below, a column value is multiplied by a numeric literal.

```
SELECT (ItemsTotal * 0.0825) AS Tax
FROM Orders
```

Arithmetic calculations are performed in the normal order of precedence: multiplication, division, addition, and then subtraction. To cause a calculation to be performed out of the normal order of precedence, use parentheses around the operation to be performed first. In the statement below the addition is performed before the multiplication.

```
SELECT (N.Numbers * (N.Multiple + 1)) AS Result
FROM Numbertable N
```

Applicability

Arithmetic operators operate only on numeric values. To use arithmetic operators on non-numeric values, first use the CAST function to convert the column to a numeric type.

AVG function

Returns the average of the values in a specified column or an expression.

```
AVG([ALL] column_reference | DISTINCT column_reference)
```

Description

Use AVG to calculate the average value for a numeric value (column, literal, parameter, or caculated values). The average is a total of the column values from all rows in the dataset divided by the number of rows. The dataset may be the entire table, a filtered dataset, or a logical group produced by a GROUP BY clause.

```
SELECT AVG(ItemsTotal)
FROM Orders
```

Column values of zero are included in the averaging, so values of 1, 2, 3, 0, 0, and 0 result in an average of 1. NULL column values are not counted in the calculation (values of 1, 2, 3, 0, NULL, and 0 result in an average of 1.2).

Constant values (such as literals and parameter values) can only be used with the AVG function if part of a calculation involving a table column. If AVG is applied to a calculation, the calculation is performed for each row, the result of the calculation added together, and the whole divided by the number of rows in the dataset.

```
SELECT AVG(ItemsTotal), AVG(ItemsTotal * 0.0825) AS AverageTax
FROM Orders
```

ALL returns the average for all rows. When DISTINCT is not specified, ALL is the implied default.

DISTINCT ignores duplicate values when averaging values in the specified column.

When used with a GROUP BY clause, AVG calculates one value for each group. This value is the aggregation of the specified column for all rows in each group. The statement below aggregates the average value for the order totals column in the ORDERS table, producing a subtotal for each company in the COMPANY table.

```
SELECT C."Company", AVG(O."ItemsTotal") AS Average,
  MAX(O."ItemsTotal") AS Biggest,
  MIN(O."ItemsTotal") AS Smallest
FROM "Customer.db" C, "Orders.db" O
WHERE (C."CustNo" = O."CustNo")
GROUP BY C."Company"
ORDER BY C."Company"
```

Applicability

SELECT statements. AVG operates only on numeric values. To use AVG on non-numeric values, first use the CAST function to convert the column to a numeric type.

Note: the AVG function cannot be used with memo or BLOB columns.

BETWEEN predicate

Determines whether a value falls inside a range.

```
value1 [NOT] BETWEEN value2 AND value3
```

Description

Use the BETWEEN comparison predicate to compare a value to a value range. If the value is greater than or equal to the low end of the range and less than or equal to the high end of the range, BETWEEN returns a TRUE value. If the value is less than the low end value or greater than the high end value, BETWEEN returns a FALSE value. For example, the expression below returns a FALSE value because 10 is not between 1 and 5.

```
10 BETWEEN 1 AND 5
```

Use NOT to return the converse of a BETWEEN comparison. For example, the expression below returns a TRUE value.

```
10 NOT BETWEEN 1 AND 5
```

BETWEEN can be used with all non-BLOB data types, but all values compared must be of the same or a compatible data type. If one value is of an incompatible data type, convert that value with the CAST function to a compatible data type. Values used in a BETWEEN comparison may be column, literal, or calculated values.

```
SELECT Saledate
FROM Orders
WHERE (Saledate BETWEEN "1/1/1988" AND "12/31/1988")
```

Hint BETWEEN is useful when filtering to retrieve rows with contiguous values that fall within the specified range. For filtering to retrieve rows with noncontiguous values, use the IN predicate.

Cartesian join

Joins two tables in a non-relational manner.

```
SELECT *
FROM table_reference, table_reference [,table_reference...]
```

Description

Use the Cartesian join to join the column of two tables into one result set, but without correlation between the rows from the tables. Cartesian joins match **each** row of the source table with **each** row of the joining table. No column comparisons are used, just simple association. If the source table has 10 rows and the joining table has 10, the result set will contain 100 rows as each row from the source table is joined with each row from the joined table.

```
SELECT *
FROM "Employee.dbf", "Items.db"
```

CAST function

Converts specified value to the specified data type.

```
CAST(column_reference AS data_type)
```

Description

Use CAST to convert the value in the specified column to the data type specified. CAST can also be applied to literal and calculated values. CAST can be used in the columns list of a SELECT statement, in the predicate for a WHERE clause, or to modify the update atom of an UPDATE statement.

The Data_Type parameter may be one of most column data type applicable to the table type used: CHAR, INTEGER, NUMERIC, and so on. Certain column types cannot be used as the source or target data types: BLOB, MEMO, and BYTES.

The statement below converts a Paradox TIMESTAMP column value to DATE.

```
SELECT CAST(SaleDate AS DATE)
FROM ORDERS
```

Converting a column value with CAST allows use of other functions or predicates on an otherwise incompatible data type, such as using the SUBSTRING function on a DATE column.

```
SELECT SaleDate,
  SUBSTRING(CAST(CAST(SaleDate AS DATE) AS CHAR(10)) FROM 1 FOR 1)
FROM Orders
```

When applied to retrieved data of a SELECT statement, the effect is transient and does not affect stored data. When applied to the update atoms of an UPDATE statement, the effect is persistent and permanently converts the case of the stored values.

Note: the CAST function cannot be used with memo or BLOB columns.

Clause list

Local SQL supports the following SQL statement clauses:

Clause	Description
FROM	Specifies the tables used for the statement.
WHERE	Specifies filter criteria to limit rows retrieved.
ORDER BY	Specifies the columns on which to sort the result set.
GROUP BY	Specifies the columns used to group rows with aggregation.
HAVING	Specifies filter criteria using aggregated data.

Note: Not all of the above clauses are supported by all statements. See the topics for individual statements for specific indications of what clauses are and are not supported.

Column correlation names

Use the AS keyword to assign a correlation name to a column, aggregated value, or literal. Column correlation names cannot be enclosed in quotation marks and so cannot contain embedded spaces. In the statement below, the tokens Sub and Word are column correlation names.

```
SELECT SUBSTRING(Company FROM 1 FOR 1) AS Sub, "Text" AS Word
FROM Customer
```

Note: Column correlation names cannot be used in all contexts. For instance, column correlation names cannot be used in the SELECT clause other than the original use (such as in other calculations), in the column comparisons that join tables, in comparisons in a WHERE clause, or in a GROUP BY clause. Column correlation names can be used in an ORDER BY clause.

Column names

ANSI-standard SQL confines each column name to a single word comprised of alphanumeric characters and the underscore symbol, "_". Local SQL, however, is enhanced to support the Paradox and dBASE feature of non-alpha characters and spaces embedded in column names. Local SQL supports Paradox multi-word column names and column names that duplicate SQL keywords as long as those column names are

- Enclosed in single or double quotation marks
- Prefaced with an SQL table name or table correlation name

For example, the following column name consists of two words:

```
SELECT E."Emp Id"
FROM Employee E
```

In the next example, the column name is the same as the SQL keyword DATE:

```
SELECT DateLog."Date"
FROM DateLog
```

Comparison predicates

Compare two values.

```
value1 < value2        less than
value1 > value2        greater than
value1 = value2        equal to
value1 <> value2       not equal to
value1 != value2       not equal to (alternate syntax)
value1 >= value2       greater than or equal to
value1 <= value2       less than or equal to
```

Description

Use comparison predicates to compare two like values. Values compared can be: column values, literals, or calculations. The result of the comparison is a boolean value that is used in contexts like a WHERE clause to determine on a row-by-row basis whether a row meets the filtering criteria.

```
SELECT *
FROM Orders
WHERE (ItemsTotal >= 1000)
```

Comparisons must be between two values of the same or a compatible data type. If one value is of an incompatible data type, convert that value with the CAST function to a compatible data type.

The result of a comparison predicate can be modified by a logical operator, such as NOT.

```
SELECT *
FROM Orders
WHERE NOT (ItemsTotal >= 1000)
```

Note: comparison predicates can only be used in a WHERE or HAVING clause, or in the ON clause of a join; they cannot be used in the SELECT clause.

Concatenation function

Concatenates two character values.

```
value1 || value2
```

Description

Use the concatenation function to concatenate two character values (column, literal, parameter, or caculated values) into a single string.

The expression below returns the string "ABCdef".

```
"ABC" || "def"
```

The statement below uses the concatenation function to combine column values with a character literal.

```
SELECT LastName || ", " || FirstName
FROM Names
```

Applicability

The concatenation function can only be used with character columns or literals. To use on values of other data types, the values must first be converted to CHAR using the CAST function.

Note: the concatenation function cannot be used with memo or BLOB columns.

Conventions

Certain rules apply during use of local SQL. These rules govern the naming of columns and tables, the format for DATE values, the use of boolean values, and other aspects of the local SQL language. Failure to abide by these rules when composing SQL statements will result in the SQL statements not working and errors incurred.

Table names

Column names

Date formats

Time formats

Boolean values

Table correlation names

Column correlation names

Embedded comments

COUNT function

Returns the number of rows that satisfy a query's search condition.

```
COUNT(* | [ALL] column_reference | DISTINCT column_reference)
```

Description

Use COUNT to count the number of rows retrieved by a SELECT statement. The SELECT statement may be a single- or multi-table query.

```
SELECT COUNT(Amount)
FROM Averaging
```

ALL returns the count for all rows. When DISTINCT is not specified, ALL is the implied default.

DISTINCT ignores duplicate values in the specified column when counting rows.

If the dataset is filtered using a WHERE clause, the COUNT function reflects the number of rows in the filtered version of the dataset, not the unfiltered version.

```
SELECT COUNT(*)
FROM "Clients.dbf"
WHERE State = "CA"
```

CREATE INDEX statement

Creates a secondary index.

```
CREATE [UNIQUE] [ASC | DESC] INDEX index_reference ON table_reference
  (column_reference [,column_reference...])
```

Description

Use the CREATE INDEX statement to create a secondary index for an existing table. Index names may not have embedded spaces. Paradox indexes may be based on multiple columns. Due to the distinctive nature of dBASE expression indexes, only single-column indexes can be created with CREATE INDEX.

Use UNIQUE to create an index that raises an error if rows with duplicate column values are inserted. By default, indexes are not unique.

Use ASC (or ASCENDING) to create an index that orders data in an ascending direction (smallest to largest). DESC (or DESCENDING) creates a descending ordering (largest to smallest). When a direction imperative is not specified, ASC is the implied default.

The following statement creates a multi-column (compound) Paradox secondary index.

```
CREATE INDEX CustDate ON "Orders.db" (CustNo, SaleDate)
```

The following statement creates a unique dBASE secondary index.

```
CREATE UNIQUE INDEX Namex ON "Employee.dbf" (Last_Name)
```

The existence of indexes may affect the updatability of queries. See the section on updatable queries for more information.

CREATE TABLE statement

Creates a table.

```
CREATE TABLE table_reference (column_definition [, column_definition,...]
  [, primary_key_constraint])
```

Description

Use the CREATE TABLE statement to create a dBASE or Paradox table, define its columns, and define a primary key constraint.

The table name reference for CREATE TABLE must comply with the rules described in the section on naming conventions. Table names with embedded spaces must be enclosed in quotation marks.

Column definitions consist of a comma-separated list of combinations of column name, data type, and (if applicable) dimensions. The list of column definitions must be enclosed in parentheses. The number and type of dimensions that must be specified varies with column type. See the section on defining column types for specific syntax of all supported column types.

Use the PRIMARY KEY (or CONSTRAINT) keyword to create a primary index for the new table.The following statement creates a Paradox table with a PRIMARY KEY constraint on the LAST_NAME and FIRST_NAME columns:

```
CREATE TABLE "employee.db"
(
  Last_Name CHAR(20),
  First_Name CHAR(15),
  Salary NUMERIC(10,2),
  Dept_No SMALLINT,
  PRIMARY KEY (Last_Name, First_Name)
)
```

An alternate syntax for creating a primary key constraint is using the CONSTRAINT keyword. While Paradox primary indexes do not have names, an arbitrary name needs to be provided to satisfy the CONSTRAINT keyword need for a token name.

```
CREATE TABLE "employee.db"
(
  Last_Name CHAR(20),
  First_Name CHAR(15),
  Salary NUMERIC(10,2),
  Dept_No SMALLINT,
  CONSTRAINT z PRIMARY KEY (Last_Name, First_Name)
)
```

Indicate whether the table is a Paradox or dBASE table by specifying the file extension when naming the table:

- ".DB" for Paradox tables
- ".DBF" for dBASE tables

If you omit the file extension for a local table name, the table created is the table type specified in the Default Driver setting in the System INIT page of the BDE Administrator utility. When specifying a file extension, the table name reference for CREATE TABLE must be enclosed in quotation marks.

Column definitions based on domains are not supported. Primary keys are the only form of constraint that can be defined with CREATE TABLE.

Note To create a table with columns that have non-alphanumeric characters or spaces in the column name, you must enclose the column name in quotation marks and prefix the quoted column name with the table name in quotes (if the filename extension is included) or just the table name. Any one of the three variations below are acceptable.

```
CREATE TABLE "abc.db"
```

```
(
  ID CHAR(3),
  "ABC.db"."Funny Name" CHAR(10)
)
CREATE TABLE abc
(
  ID CHAR(3),
  ABC."Funny Name" CHAR(10)
)
CREATE TABLE "abc.db"
(
  ID CHAR(3),
  ABC."Funny Name" CHAR(10)
)
```

Data manipulation language overview

The Data Manipulation Language (DML) part of local SQL consists of the following language elements.

Category	Description
Statements	Statements, or commands, that retrieve, modify, and delete data.
Clauses	Clauses of statements that affect how a statement operates.
Functions	Functions like UPPER that alter data and SUM that aggregate data.
Operators	Arithmetic, comparison, logical, and string concatenation operators.
Predicates	Keywords used in WHERE clauses to qualify rows returned by a statement.
Relational operators	Operators that allow joining multiple tables.
Updatable queries	Conditions under which queries are updateable or read-only.
Parameters in queries	Using parameters to make queries dynamic.
Local SQL VIEWs	The local SQL emulation of SQL VIEWs.

Date formats

There are two acceptable formats for the date parts of DATE and TIMESTAMP literals in local SQL statements: the US date format and the European date format. Which one of the two you use is predicated on the Windows configuration settings for Date Separator and Short Date Style. If the Date Separator is a slash (/) and Short Date Style puts the month before the day (such as MM-dd-yy), the date must be in a US date format, MM/DD/YY or MM/DD/YYYY. If the Date Separator is a period (.) and Short Date Style puts the day before the month (such as dd.MM.yy), the date must be in a European date format, DD.MM.YY or DD.MM.YYYY. A combination of Short Date Style of month before day and a Date Separator of anything but a period is not supported in local SQL. DATE and TIMESTAMP values passed to an SQL statement via parameters and the values in columns are unaffected by these restrictions on format. Only literal values are affected.

To prevent date literals from being mistaken by the SQL parser for arithmetic calculations, enclose them in quotation marks. This keeps 1/23/1998 from being mistaken for 1 divided by 23 divided by 1998.

Example of a date literal in a US format:

```
SELECT *
FROM Orders
WHERE (SaleDate <= "1/23/1998")
```

Example of the same date literal in a European format:

```
SELECT *
FROM Orders
WHERE (SaleDate <= "23.1.1998")
```

Leading zeros for the month and day fields are optional.

If the century is not specified for the year in a date value, the BDE settings FOURDIGITYEAR and YEARBIASED control the century of the date.

Data definition overview

Local SQL supports data definition language (DDL) for creating, altering, and deleting tables and indexes.

Local SQL does not permit the metadata object names to be represented by parameters in DDL statements.

Local SQL supports the following DDL statements:

DDL Statement	Description
CREATE TABLE	Creates a new table.
ALTER TABLE	Adds columns to and deletes columns from an existing table.
DROP TABLE	Deletes an existing table.
CREATE INDEX	Creates a new secondary index for an existing table.
DROP INDEX	Deletes an existing primary or secondary index.

DELETE statement

Deletes one or more rows from a table.

```
DELETE FROM table_reference
[WHERE predicates]
```

Description

Use DELETE to delete one or more rows from one existing table per statement.

```
DELETE FROM "Employee.db"
```

The optional WHERE clause restricts row deletions to a subset of rows in the table that meet a logical criteria. If no WHERE clause is specified, all rows in the table are deleted. The statement below deletes all of the rows from the EMPLOYEE table if the value in its EmpNo column is found in the result set produced by a SELECT subquery against the OLD_EMPLOYEE table.

```
DELETE FROM "Employee.db"
WHERE (Empno IN (SELECT Empno FROM "Old_Employee.db"))
```

The DELETE statement only supports SELECT subqueries in the WHERE clause. References to tables other than the one from which rows are deleted or columns in such tables are only possible in SELECT subqueries.

DML statement list

Local SQL supports the following data manipulation language (DML) statements:

DML Statements	Description
SELECT	Retrieves existing data from a table.
DELETE	Deletes existing data from a table.
INSERT	Adds new data to a table.
UPDATE	Modifies existing data in a table.

DROP INDEX

Deletes an index.

```
DROP INDEX table_reference.index_reference | PRIMARY
```

Description

Use the DROP INDEX statement to delete a primary or secondary index.

To delete a dBASE primary or secondary index or a Paradox secondary index, identify the index using the table name and index name separated by an identifier connector symbol (.).

```
DROP INDEX "Employee.dbf".Namex
```

To delete a Paradox primary index, identify the index with the keyword PRIMARY.

```
DROP INDEX Orders.PRIMARY
```

The existence of indexes may affect the updatability of queries. See the section on updatable queries for more information.

DROP TABLE statement

Deletes a table.

```
DROP TABLE table_reference
```

Description

Use the DROP TABLE statement to delete an existing table. The statement below drops a Paradox table:

```
DROP TABLE "Employee.db"
```

Embedded comments

Comments, or remarks, can be embedded in SQL statements to add clarity or explanation. Text is designated as a comment and not treated as SQL by enclosing it within the beginning **/*** and ending ***/** comment symbols. The symbols and comments need not be on the same line.

```
/*
  This is a comment
*/
SELECT SUBSTRING(Company FROM 1 FOR 1) AS sub, "Text" AS word
FROM Customer
```

Comments can also be embedded within an SQL statement. This is useful when debugging an SQL statement, such as removing one clause for testing.

```
SELECT Company
FROM Customer
/* WHERE (State = "TX") */
ORDER BY Company
```

Equi-join

Joins two tables based on column values common between the two, excluding non-matches.

```
SELECT column_list
FROM table_reference, table_reference [, table_reference...]
WHERE predicate [AND predicate...]
```

Description

Use equi-join to join two tables, a source and joining table, that have values from one or more columns in common. One or more columns from each table are compared in the WHERE clause for equal values. For rows in the source table that have a match in the joining table, the data for the source table rows and matching joining table rows are included in the result set. Rows in the source table without matches in the joining table are excluded from the joined result set. In the statement below, the CUSTOMER and ORDERS tables are joined based on values in the CUSTNO column, which each table contains.

```
SELECT *
FROM Customer C, Orders O
WHERE (C.CustNo = O.CustNo)
```

More that one table may be joined with an equi-join. One column comparison predicate in the WHERE clause is required for each column compared to join each two tables. The statement below joins the CUSTOMER table to ORDERS, and then ORDERS to ITEMS. In this case, the joining table ORDERS acts as a source table for the joining table ITEMS.

```
SELECT *
FROM Customer C, Orders O, Items I
WHERE (C.CustNo = O.CustNo) AND
  (O.OrderNo = I.OrderNo)
```

Tables may also be joined using a concatenation of multiple column values to produce a single value for the join comparison predicate. Here, the ID1 and ID2 columns in JOINING are concatenated and compared with the values in the single column ID in SOURCE.

```
SELECT *
FROM Source S, Joining J
WHERE (S.ID = J.ID1 || J.ID2)
```

An ORDER BY clause in equi-join statements can use columns from any table specified in the FROM clause to sort the result set.

EXISTS predicate

Indicates whether values exist in a subquery.

```
EXISTS subquery
```

Description

Use the EXISTS comparison predicate to filter a table based on the existence of column values from the table in a subquery. The subquery is filtered using a WHERE clause comparing one or more columns in the filtered table to corresponding columns in the subquery. EXISTS returns a true value if the subquery has at least one row in its result set, false if zero rows are retrieved. The subquery is executed once for each row in the filtered table and the existence of rows in the subquery is used to include or exclude the rows in the filtered table.

```
SELECT O.OrderNo, O.CustNo
FROM Orders O
WHERE EXISTS
  (SELECT C.CustNo
   FROM Customer C
   WHERE (C.CustNo = O.CustNo))
```

The subquery may be further filtered with other conditions. For example, the statement below returns the rows pertaining to all customers who have placed Orders the totals for which exceed $1000.

```
SELECT C.Company, C.CustNo
FROM Customer C
WHERE EXISTS
  (SELECT O.CustNo
   FROM Orders O
   WHERE (O.CustNo = C.CustNo) AND
     (O.ItemsTotal > 1000))
```

Use NOT to return the converse of an EXISTS comparison.

EXTRACT function

Returns one field from a date value.

```
EXTRACT(extract_field FROM column_reference)
```

Description

Use EXTRACT to return the year, month, or day field from a DATE or TIMESTAMP column. If the column used with the EXTRACT function contains a NULL value, the return value of EXTRACT will be NULL. If the value is not NULL, EXTRACT returns the value for the specified element in the date, expressed as a SMALLINT.

The Extract_Field parameter may contain any one of the specifiers: YEAR, MONTH, DAY, HOUR, MINUTE, or SECOND. The specifiers YEAR, MONTH, and DAY can only be used with DATE and TIMESTAMP columns. The specifiers HOUR, MINUTE, and SECOND can only be used with TIMESTAMP and TIME columns.

```
SELECT SaleDate,
  EXTRACT(YEAR FROM saledate) AS YY,
  EXTRACT(MONTH FROM saledate) AS MM,
  EXTRACT(DAY FROM saledate) AS DD
FROM Orders
```

The statement below uses a DOB column (containing birthdates) to filter to those rows where the date is in the month of May. The month field from the DOB column is retrieved using the EXTRACT function and compared to 5, May being the fifth month.

```
SELECT DOB, LastName, FirstName
FROM People
WHERE (EXTRACT(MONTH FROM DOB) = 5)
```

Applicability

EXTRACT operates only on DATE, TIME, and TIMESTAMP values. To use EXTRACT on non-date values, first use the CAST function to convert the column to a date type.

Note: while SQL-92 provides the EXTRACT function specifiers TIMEZONE_HOUR and TIMEZONE_MINUTE, these specifiers are not supported in local SQL.

FROM clause

Specifies the tables from which a SELECT statement retrieves data.

```
FROM table_reference [, table_reference...]
```

Description

Use a FROM clause to specify the table or tables from which a SELECT statement retrieves data. The value for a FROM clause is a comma-separated list of table names. Specified table names must follow local SQL naming conventions for tables. In statements that support the FROM clause, the clause is required.

In the following example of a SELECT statement, a single Paradox table is specified in the statement's FROM clause.

```
SELECT *
FROM "Customer.db"
```

Define table correlation names for tables in the FROM clause. These table correlation names can then be used elsewhere in the statement. In the SELECT statement below, the table correlation name Res is defined for the table RERSERVAT.DB:

```
SELECT Res.ResNo, Res.Amt_Paid
FROM "Reservat.db" Res
WHERE (Res.Pay_Method = "VISA")
```

See the section Relational Operators for more information on retrieving data from multiple tables in a single SELECT query.

Local SQL also supports retrieving data from saved queries (SQL or QBE) by specifying the saved file in the FROM clause of a statement. See local SQLVIEWs for more information on this.

Applicability

SELECT statements

Function list

Local SQL supports the following data manipulation language functions:

String functions	Description
Concatenation	Concatenates two string values.
LOWER	Forces a string to lowercase.
UPPER	Forces a string to uppercase.
SUBSTRING	Extracts a portion of a string value.
TRIM	Removes repetitions of a specified character from the left, right, or both sides of a string.

Aggregate functions	Description
AVG	Averages all non-NULL numeric values in a column.
COUNT	Counts the number of rows in a result set.
MAX	Determines the maximum value in a column.
MIN	Determines the minimum value in a column.
SUM	Totals all numeric values in a column.

Data functions	Description
CAST	Converts values from one data type to another.
EXTRACT	Extracts the year, month, or day field of a date.

GROUP BY clause

Combines rows with column values in common into single rows.

```
GROUP BY column_reference [, column reference...]
```

Description

Use a GROUP BY clause to cause an aggregation process to be repeated once for each group of similar rows. Similarity between rows is determined by the distinct values (or combination of values) in the columns specified in the GROUP BY. For instance, a query with a SUM function produces a result set with a single row with the total of all the values for the column used in the SUM function. But when a GROUP BY clause is added, the SUM function performs its summing action once for each group of rows, a subtotalling. In statements that support a GROUP BY clause, the use of a GROUP BY clause is optional. A GROUP BY clause becomes necessary when both aggregated and non-aggregated columns are included in the same SELECT statement.

In the statement below, the SUM function produces one subtotal of the ItemsTotal column for each distinct value in the CustNo column (i.e., one subtotal for each different customer).

```
SELECT CustNo, SUM(ItemsTotal)
FROM Orders
GROUP BY CustNo
```

The value for the GROUP BY clause is a comma-separated list of columns. Each column in this list must meet the following criteria:

- Be in one of the tables specified in the FROM clause of the query.
- Also be in the SELECT clause of the query.
- Cannot have an aggregate function applied to it (in the SELECT clause).

When a GROUP BY clause is used, all table columns in the SELECT clause of the query must meet at least one of the following criteria, or it cannot be included in the SELECT clause:

- Be in the GROUP BY clause of the query.
- Be in the subject of an aggregate function.

Literal values in the SELECT clause are not subject to the preceding criteria and are not required to be in the GROUP BY clause in addition to the SELECT clause.

The distinctness of rows is based on the columns in the column list specified. All rows with the same values in these columns are combined into a single row (or logical group). Columns that are the subject of an aggregate function have their values across all rows in the group combined. All columns not the subject of an aggregate function retain their value and serve to distinctly identify the group. For example, in the SELECT statement below, the values in the Sales column are aggregated (totalled) into groups based on distinct values in the Company column. This produces total sales for each company.

```
SELECT C.Company, SUM(O.ItemsTotal) AS TotalSales
FROM Customer C, Orders O
WHERE C.CustNo = O.CustNo
GROUP BY C.Company
ORDER BY C.Company
```

A column may be referenced in a GROUP BY clause by a column correlation name, instead of actual column names. The statement below forms groups using the first column, Company, represented by the column correlation name Co.

```
SELECT C.Company Co, SUM(O.ItemsTotal) AS TotalSales
FROM Customer C, Orders O
WHERE C.CustNo = O.CustNo
GROUP BY Co
ORDER BY 1
```

Note Derived values (calculated values) cannot be used as the basis for a GROUP BY clause.

Applicability
SELECT statements (when data aggregation is performed)

HAVING clause

Specifies filtering conditions for a SELECT statement.

```
HAVING predicates
```

Description

Use a HAVING clause to limit the rows retrieved by a SELECT statement to a subset of rows where aggregated column values meet the specified criteria. A HAVING clause can only be used in a SELECT statement when:

- The statement also has a GROUP BY clause.
- One or more columns are the subjects of aggregate functions.

The value for a HAVING clause is one or more logical expressions, or predicates, that evaluate to true or false for each aggregate row retrieved from the table. Only those rows where the predicates evaluate to true are retrieved by a SELECT statement. For example, the SELECT statement below retrieves all rows where the total sales for individual companies exceed $1,000.

```
SELECT Company, SUM(sales) AS TOTALSALES
FROM Sales1998
GROUP BY Company
HAVING (SUM(sales) >= 1000)
ORDER BY Company
```

Multiple predicates must be separated by one of the logical operators OR or AND. Each predicate can be negated with the NOT operator. Parentheses can be used to isolate logical comparisons and groups of comparisons to produce different row evaluation criteria.

A SELECT statement can include both a WHERE clause and a HAVING clause. The WHERE clause filters the data to be aggregated, using columns not the subject of aggregate functions. The HAVING clause then further filters the data after the aggregation, using columns that are the subject of aggregate functions. The SELECT query below performs the same operation as that above, but data limited to those rows where the State column is "CA".

```
SELECT Company, SUM(sales) AS TOTALSALES
FROM Sales1998
WHERE (State = "CA")
GROUP BY Company
HAVING (SUM(Sales) >= 1000)
ORDER BY Company
```

Subqueries are supported in the HAVING clause. A subquery works like a search condition to restrict the number of rows returned by the outer, or "parent" query. The following statement uses a subquery to total all the sales for California and calculates a value half that total. This calculated value is then used in the parent statement to restrict the result set to aggregated rows where the grouped subtotal (based on customer number) is at least as much as the calculated value.

```
SELECT O.CustNo, SUM(O.ItemsTotal)
FROM Orders O
GROUP BY O.CustNo
HAVING SUM(O.ItemsTotal) >=
  (SELECT SUM(O2.ItemsTotal) / 2
  FROM Customer C,Orders O2
  WHERE (C.CustNo = O2.CustNo) AND
    (C.State = "CA"))
```

Note A HAVING clause filters data after the aggregation of a GROUP BY clause. For filtering based on row values prior to aggregation, use a WHERE clause.

Applicability

SELECT statements (with GROUP BY clauses and data aggregation)

Heterogeneous joins

Joins two tables from different databases.

```
SELECT column_list
FROM ":database_reference:table_reference",
  ":database_reference:table_reference"
  [,":database_reference:table_reference"...]
WHERE predicate [AND predicate...]
```

Description

Use a heterogeneous join to join two tables that reside in different databases. The joined tables may be of different types (like dBASE to Paradox or Paradox to InterBase), but you can only join tables whose database types are accessible through the BDE (local, ODBC, or SQL Links). A hetergeneous join may be any of the joins supported by local SQL. The difference is in the syntax for the table reference: the database containing each table is specified in the table reference, surrounded by colons and the whole reference enclosed in quotation marks. The database specified as part of the table reference may be a drive and directory reference (for local tables) or a BDE alias.

```
SELECT *
FROM ":DBDEMOS:Customer.db" C, ":BCDEMOS:Orders.db" O
WHERE (C.CustNo = O.CustNo)
```

IN predicate

Indicates whether a value exists in a set of values.

```
value [NOT] IN (value_set)
```

Description

Use the IN comparison predicate to filter a table based on the existence of a column value in a specified set of comparison values. The set of comparison values may be either static using a comma-separated list of literals or dynamic using the result set from a subquery.

The value to compare with the values set can be any or a combination of: a column value, a literal value, or a calculated value.

The comparison set can be a static comma-separated list of literal values.

```
SELECT C.Company, C.State
FROM Customer C
WHERE (C.State IN ("CA", "HI"))
```

The comparison set can also be the result set from a subquery. The subquery may return multiple rows, but must only return a single column for comparison.

```
SELECT C.Company, C.State
FROM Customer C
WHERE (C.State IN
  (SELECT R.State
  FROM Regions R
  WHERE (R.Region = "Pacific")))
```

Use NOT to return the converse of an IN comparison.

IN can be used with all non-BLOB data types, but all values compared must be of the same or a compatible data type. If one value is of an incompatible data type, convert that value with the CAST function to a compatible data type.

Hint IN is useful when filtering to retrieve rows with noncontiguous values. For filtering to retrieve rows with contiguous values that fall within a specified range, use the BETWEEN predicate.

INNER join

Joins two tables based on column values common between the two, excluding non-matches.

```
SELECT column_list
FROM table_reference
  [INNER] JOIN table_reference
    ON predicate
  [[INNER] JOIN table_reference
    ON predicate...]
```

Description

Use an INNER JOIN to join two tables, a source and joining table, that have values from one or more columns in common. One or more columns from each table are compared in the ON clause for equal values. For rows in the source table that have a match in the joining table, the data for the source table rows and matching joining table rows are included in the result set. Rows in the source table without matches in the joining table are excluded from the joined result set. In the statement below, the CUSTOMER and ORDERS tables are joined based on values in the CUSTNO column, which each table contains.

```
SELECT *
FROM Customer C
  INNER JOIN Orders O
    ON (C.CustNo = O.CustNo)
```

More than one table may be joined with an INNER JOIN. One use of the INNER JOIN operator and corresponding ON clause is required for each each set of two tables joined. One columns comparison predicate in an ON clause is required for each column compared to join each two tables. The statement below joins the CUSTOMER table to ORDERS, and then ORDERS to ITEMS. In this case, the joining table ORDERS acts as a source table for the joining table ITEMS. (The statement below appears without the optional INNER keyword.)

```
SELECT *
FROM Customer C
  JOIN Orders O
    ON (C.CustNo = O.CustNo)
  JOIN items I
    ON (O.OrderNo = I.OrderNo)
```

Tables may also be joined using a concatenation of multiple column values to produce a single value for the join comparison predicate. Here, the ID1 and ID2 columns in JOINING are concatenated and compared with the values in the single column ID in SOURCE.

```
SELECT *
FROM Source S
  INNER JOIN Joining J
    ON (S.ID = J.ID1 || J.ID2)
```

An ORDER BY clause in INNER JOIN statements can use columns from any table specified in the FROM clause to sort the result set.

INSERT statement

Adds one or more new rows of data in a table

```
INSERT INTO table_reference
[(columns_list)]
VALUES (update_atoms)
```

Description

Use the INSERT statement to add new rows of data to a single table.

Use a table reference in the INTO clause to specify the table to receive the incoming data.

The columns list is a comma-separated list, enclosed in parentheses, of columns in the table and is optional. The VALUES clause is a comma-separated list of update atoms, enclosed in parentheses. Unless the source of new rows is a SELECT subquery, the VALUES clause is required.

If no columns list is specified, incoming update values (update atoms) are stored in fields as they are defined sequentially in the table structure. Update atoms are applied to columns in the order the update atoms are listed in the VALUES clause. There can be fewer update atoms than there are columns in the table, but there cannot be more.

```
INSERT INTO "Holdings.dbf"
VALUES (4094095, "INPR", 5000, 10.500, "1/2/1998")
```

If an explicit columns list is stated, incoming update atoms (in the order they appear in the VALUES clause) are stored in the listed columns (in the order they appear in the columns list). NULL values are stored in any columns that are not in a columns list. When a columns list is explicitly described, there must be exactly the same number of update atoms in the VALUES clause as there are columns in the list.

```
INSERT INTO "Customer.db"
(CustNo, Company)
VALUES (9842, "Inprise Corporation")
```

To add rows to one table that are retrieved from another table, omit the VALUES keyword and use a subquery as the source for the new rows.

```
INSERT INTO "Customer.db"
(CustNo, Company)
SELECT CustNo, Company
FROM "OldCustomer.db"
```

The INSERT statement only supports SELECT subqueries in the VALUES clause. References to tables other than the one to which rows are added or columns in such tables are only possible in SELECT subqueries.

Note: The INSERT statement can use a single SELECT statement as the source for the new rows, but not multiple statements joined with UNION. To do that, save the UNION join statement to a .SQL file and use that file in the INSERT statement. See Local SQL VIEWs for more information on using .SQL files.

IS NULL predicate

Indicates whether a column contains a NULL value.

```
column_reference IS [NOT] NULL
```

Description

Use the IS NULL comparison predicate to filter a table based on the specified column containing a NULL (empty) value.

```
SELECT *
FROM Customer
WHERE (InvoiceDate IS NULL)
```

Use NOT to return the converse of a IS NULL comparison.

Note For a numeric column, a zero value is not the same as a NULL value.

Legend

Typographical conventions

To ensure clarity in the SQL syntax examples and pseudo-code, the typographical conventions that follow are used throughout the local SQL help.

SQL statements

Each topic contains a pseudo-code prototype of an SQL statement or part of a statement that demonstrates the language element discussed. These prototype examples appear at the beginning of topics. Actual SQL statements appear within topics that demonstrate actual use of the language element discussed. Both the prototype and example statements appear in `Courier New` font.

While the local SQL language itself is case-insensitive (language elements and metadata object names), examples in this help file use the following convention to differentiate between language and metadata objects. All language elements appear in uppercase. Metadata names appear in lowercase. Correlation names appear in mixed case.

Optional elements

Language elements that are available, but that do not have to be used with an SQL statement for the statement to be valid appear in prototype syntax examples in brakets (`[` and `]`). For example, in the line below, the DISTINCT keyword is optional.

```
SELECT [DISTINCT] *
```

Syntax choices

When there is a choice between one of a number of possible syntax elements, such choices will be listed in prototype syntax examples separated by the vertical bar character (`|`). Unless also enclosed in brackets to make the group of choices optional, one of the group of choices must be used in the statement. In the prototype syntax example below, the SQL statement may include either the ASC or the DESC keyword, but not both. Because the list of choices (ASC and DESC) is enclosed in brackets, use of either keyword is optional.

```
ORDER BY column_reference [ASC | DESC]
```

Note Do not mistake the choice typographical symbol for the concatenation function, defined as two vertical bar characters together (`||`).

LIKE predicate

Indicates the similarity of one value as compared to another.

```
value [NOT] LIKE [substitution_char] comparison_value [substitution_char]
  [ESCAPE escape_char]
```

Description

Use the LIKE comparison predicate to filter a table based on the similarity of a column value to a comparison value. Use of substitution characters allows the comparison to be based on the whole column value or just a portion.

```
SELECT *
FROM Customer
WHERE (Company LIKE "Adventure Undersea")
```

The wildcard substitution character ("%") may be used in the comparison to represent an unknown number of characters. LIKE returns a TRUE when the portion of the column value matches that portion of the comparison value not corresponding to the position of the wildcard character. The wildcard character can appear at the beginning, middle, or end of the comparison value (or multiple combinations of these positions). For example, the statement below retrieves rows where the column value begins with "A" and is followed by any number of any characters. Matching values could include "Action Club" and "Adventure Undersea", but not "Blue Sports".

```
SELECT *
FROM Customer
WHERE (Company LIKE "A%")
```

The single-character substitution character ("_") may be used in the comparison to represent a single character. LIKE returns a TRUE when the portion of the column value matches that portion of the comparison value not corresponding to the position of the single-character substitution character. The single-character substitution character can appear at the beginning, middle, or end of the comparison value (or multiple combinations of these positions). Use one single-character substitution character for each character to be wild in the filter pattern For example, the statement below retrieves rows where the column value begins with "b" ends with "n", with one character of any value between. Matching values could include "bin" and "ban", but not "barn".

```
SELECT Words
FROM Dictionary
WHERE (Words LIKE "b_n")
```

Use NOT to return the converse of a LIKE comparison.

Use ESCAPE when the wildcard character "%" or "_" appear as data in the column. The ESCAPE keyword designates an escape character. In the comparison value for the LIKE predicate, the character that follows the escape character is treated as a data character and not a wildcard character. Other wildcard characters in the comparison value are unaffected.

In the example below, the "^" character is designated as the escape character. In the comparison value for the LIKE predicate ("%10^%%"), the "%" that immediately follows the escape character is treated as data in the PercentValue. This allows filtering based on the string "10%".

```
SELECT *
FROM Sales
WHERE (PercentValue LIKE "%10^%%" ESCAPE "^")
```

LIKE can be used only with CHAR or compatible data types. If one value is of an incompatible data type, convert that value with the CAST function to a compatible data type. The comparison performed by the LIKE predicate is case-sensitive.

Logical operators

Connect multiple predicates.

```
[NOT] predicate OR [NOT] predicate
[NOT] predicate AND [NOT] predicate
```

Description

Use the logical operators OR and AND to connect two predicates in a single WHERE clause. This allows the table to be filtered based on multiple conditions. Logical operators compare the boolean result of two predicate comparisons, each producing a boolean result. If OR is used, **either** of the two predicate comparisons can result on a TRUE value for the whole expression to evaluate to TRUE. If AND is used, **both** predicate comparisons must evaluate to TRUE for the whole expression to be TRUE; if **either** is FALSE, the whole is FALSE. In the statement below, if only one of the two predicate comparisons is TRUE (`reservdate < "1/31/1998"` or `paid = TRUE`), the row will be included in the query result set.

```
SELECT *
FROM Reservations
WHERE ((ReserveDate < "1/31/1998") OR (Paid = TRUE))
```

Logical operator comparisons are performed in the order of precedence: OR and then AND. To perform a comparison out of the normal order of precedence, use parentheses around the comparison to be performed first. The SELECT statement below retrieves all rows where the SHAPE column is "round" and the COLOR "blue". It also returns those rows where the COLOR column is "red", regardless of the value in the SHAPE column (such as "triangle"). It would not return rows where the SHAPE is "round" and the COLOR anything but "blue" or where the COLOR is "blue" and the SHAPE anything but "round".

```
SELECT Shape, Color, Cost
FROM Objects
WHERE ((Shape = "round") AND (Color = "blue")) OR
  (Color = "red")
```

Without the parentheses, the default order of precedence is used and the logic changes. The statement below, a variation on the above statement, would return rows where the SHAPE is "square" and the COLOR is "blue". It would also return rows where the SHAPE is "square" and the COLOR is "red". But unlike the preceding statement, it would not return rows where the COLOR is "red" and the SHAPE "triangle".

```
SELECT Shape, Color, Cost
FROM Objects
WHERE Shape = "round" AND Color = "blue" OR
 Color = "red"
```

Use the NOT operator to negate the boolean result of a comparison. In the statement below, only those rows where the PAID column contains a FALSE value are retrieved.

```
SELECT *
FROM Reservations
WHERE (NOT (Paid = "TRUE"))
```

LOWER function

Converts all characters to lowercase.

```
LOWER(column_reference)
```

Description

Use LOWER to convert all of the characters in a character value (column, literal, parameter, or caculated values) to lowercase. For example, in the SELECT statement below the values in the Name column appear all in lowercase.

```
SELECT LOWER(Name)
FROM Country
```

When applied to retrieved data of a SELECT statement, the effect is transient and does not affect stored data. When applied to the update atoms of an UPDATE statement, the effect is persistent and permanently converts the case of the stored values.

The LOWER function can be used in WHERE clause string comparisons to effect a case-insensitive comparison. Apply LOWER to the values on both sides of the comparison operator (if one of the comparison values is a literal, simply enter it all in lower case).

```
SELECT *
FROM Cames
WHERE LOWER(LastName) = "smith"
```

Applicability

LOWER can only be used with character columns or literals. To use on values of other data types, the values must first be converted to CHAR using the CAST function.

Note: The LOWER function cannot be used with memo or BLOB columns.

Note: The lower case version of a given character is determined by the language driver used.

MAX function

Returns the largest value in the specified column.

```
MAX([ALL] column_reference | DISTINCT column_reference)
```

Description

Use MAX to calculate the largest value for a value (column, literal, parameter, or caculated values) using a dataset. The dataset may be the entire table, a filtered dataset, or a logical group produced by a GROUP BY clause. Column values of zero are included in the aggregation. NULL column values are not counted in the calculation. If the number of qualifying rows is zero, MAX returns a NULL value.

```
SELECT MAX(ItemsTotal)
FROM Orders
```

Constant values (such as literals and parameter values) can only be used with the MAX function if part of a calculation involving a table column. If MAX is applied to a calculation, the calculation is performed for each row and then the maximum value calculated based on the data as modified by the calculation.

```
SELECT MAX(UPPER(LastName))
FROM "Employee.db"
```

ALL returns the largest value for all rows. When DISTINCT is not specified, ALL is the implied default.

DISTINCT ignores duplicate values when calculating the largest value in the specified column.

MAX returns the largest value in a column or a calculation using a column performed for each row (a calculated field).

```
SELECT MAX(ItemsTotal), MAX(ItemsTotal * 0.0825) AS HighestTax
FROM Orders
```

When used with a GROUP BY clause, MAX returns one row for each group with the maximum value for that group. The statement below aggregates the largest value for the order totals column in the Orders table, producing a subtotal for each company in the Company table.

```
SELECT C."Company", AVG(O."ItemsTotal") AS Average,
  MAX(O."ItemsTotal") AS Biggest,
  MIN(O."ItemsTotal") AS Smallest
FROM "Customer.db" C, "Orders.db" O
WHERE (C."CustNo" = O."CustNo")
GROUP BY C."Company"
ORDER BY C."Company"
```

Applicability

SELECT statements. MAX can be used with all non-BLOB columns. When used with numeric columns, the return value is of the same type as the column (such as INTEGER or FLOAT). When used with a CHAR column, the largest value returned will depend on the Borland Database Engine (BDE) language driver used.

Note: the MAX function cannot be used with memo or BLOB columns.

MIN function

Returns the smallest value in the specified column.

```
MIN([ALL] column_reference | DISTINCT column_reference)
```

Description

Use MIN to calculate the smallest value (column, literal, parameter, or caculated values) using a dataset. The dataset may be the entire table, a filtered dataset, or a logical group produced by a GROUP BY clause. Column values of zero are included in the aggregation. NULL column values are not counted in the calculation. If the number of qualifying rows is zero, MIN returns a NULL value.

```
SELECT MIN(ItemsTotal)
FROM Orders
```

Constant values (such as literals and parameter values) can only be used with the MIN function if part of a calculation involving a table column. If MIN is applied to a calculation, the calculation is performed for each row and then the minimum value calculated based on the data as modified by the calculation.

```
SELECT MAX(UPPER(LastName))
FROM "Employee.db"
```

ALL returns the smallest value for all rows. When DISTINCT is not specified, ALL is the implied default.

DISTINCT ignores duplicate values when calculating the smallest value in the specified column.

MIN returns the smallest value in a column or a calculation using a column performed for each row (a calculated field).

```
SELECT MIN(ItemsTotal), MIN(ItemsTotal * 0.0825) AS LowestTax
FROM Orders
```

When used with a GROUP BY clause, MIN returns one row for each group with the minimum value for that group. The statement below aggregates the smallest value for the order totals column in the Orders table, producing a subtotal for each company in the Company table.

```
SELECT C."Company", AVG(O."ItemsTotal") AS Average,
  MAX(O."ItemsTotal") AS Biggest,
  MIN(O."ItemsTotal") AS Smallest
FROM "Customer.db" C, "Orders.db" O
WHERE (C."CustNo" = O."CustNo")
GROUP BY C."Company"
ORDER BY C."Company"
```

Applicability

SELECT statements. MIN can be used with all non-BLOB columns. When used with numeric columns, the return value is of the same type as the column (such as INTEGER or FLOAT). When used with a CHAR column, the smallest value returned will depend on the Borland Database Engine (BDE) language driver used.

Note: the MIN function cannot be used with memo or BLOB columns.

Operators list

Local SQL supports the following operators:

Type	Operators
Arithmetic	`+ – * /`
Logical	`AND  OR  NOT`

ORDER BY clause

Sorts the rows retrieved by a SELECT statement.

```
ORDER BY column_reference [, column_reference...] [ASC|DESC]
```

Description

Use an ORDER BY clause to sort the rows retrieved by a SELECT statement based on the values from one or more columns. In statements that support an ORDER BY clause, use of the clause is optional.

The value for the ORDER BY clause is a comma-separated list of column names. The columns in this list must also be in the SELECT clause of the query statement. Columns in the ORDER BY list can be from one or multiple tables. If the columns used for an ORDER BY clause come from multiple tables, the tables must all be those that are part of a join. They cannot be a table included in the statement only through a SELECT subquery.

A column may be specified in an ORDER BY clause using a number representing the relative position of the column in the SELECT of the statement. Column correlation names can also be used in an ORDER BY clause columns list. Calculations cannot be used directly in an ORDER BY clause. Instead, assign a column correlation name to the calculation and use that name in the ORDER BY clause.

Use ASC (or ASCENDING) to force the sort to be in ascending order (smallest to largest), or DESC (or DESCENDING) for a descending sort order (largest to smallest). When not specified, ASC is the implied default.

The statement below sorts the result set ascending by the year extracted from the LastInvoiceDate column, then descending by the State column, and then ascending by the uppercase conversion of the Company column.

```
SELECT EXTRACT(YEAR FROM LastInvoiceDate) AS YY, State, UPPER(Company)
FROM Customer
ORDER BY YY DESC, State ASC, 3
```

See the section Relational Operators for more information on retrieving data from multiple tables in a single SELECT query.

Applicability

SELECT statements

OUTER join

Joins two tables based on column values common between the two, including non-matches.

```
SELECT column_list
FROM table_reference
   LEFT | RIGHT | FULL [OUTER] JOIN table_reference
    ON predicate
  [LEFT | RIGHT | FULL [OUTER] JOIN table_reference
    ON predicate...]
```

Description

Use an OUTER JOIN to join two tables, a source and joining table, that have one or more columns in common. One or more columns from each table are compared in the ON clause for equal values. The primary difference between inner and outer joins is that, in outer joins rows from the source table that do not have a match in the joining table are **not** excluded from the result set. Columns from the joining table for rows in the source table without matches have NULL values.

In the statement below, the CUSTOMER and ORDERS tables are joined based on values in the CUSTNO column, which each table contains. For rows from CUSTOMER that do not have a matching value between CUSTOMER.CUSTNO and ORDERS.CUSTNO, the columns from ORDERS contain NULL values.

```
SELECT *
FROM Customer C
  LEFT OUTER JOIN Orders O
    ON (C.CustNo = O.CustNo)
```

The LEFT modifier causes all rows from the table on the left of the OUTER JOIN operator to be included in the result set, with or without matches in the table to the right. If there is no matching row from the table on the right, its columns contain NULL values. The RIGHT modifier causes all rows from the table on the right of the OUTER JOIN operator to be included in the result set, with or without matches. If there is no matching row from the table on the left, its columns contain NULL values. The FULL modifier causes all rows from the all tables specified in the FROM clause to be included in the result set, with or without matches. If there is no matching row from one of the tables, its columns contain NULL values.

More than one table may be joined with an INNER JOIN. One use of the INNER JOIN operator and corresponding ON clause is required for each each set of two tables joined. One column comparison predicate in an ON clause is required for each column compared to join each two tables. The statement below joins the CUSTOMER table to ORDERS, and then ORDERS to ITEMS. In this case, the joining table ORDERS acts as a source table for the joining table ITEMS.

```
SELECT *
FROM Customer C
  FULL OUTER JOIN Orders O
    ON (C.CustNo = O.CustNo)
  FULL OUTER JOIN items I
    ON (O.OrderNo = I.OrderNo)
```

Tables may also be joined using expressions to produce a single value for the join comparison predicate. Here, the ID1 and ID2 columns in JOINING are separately compared with two values produced by the SUBSTRING function using the single column ID in SOURCE.

```
SELECT *
FROM Source S
  RIGHT OUTER JOIN Joining J
    ON (SUBSTRING(S.ID FROM 1 FOR 2) = J.ID1) AND
    (SUBSTRING(S.ID FROM 3 FOR 1) = J.ID2)
```

An ORDER BY clause in OUTER JOIN statements can use columns from any table specified in the FROM clause to sort the result set.

Overview

What is local SQL?

Local SQL is the subset of the SQL-92 specification used to access dBASE, Paradox, and FoxPro tables. On receiving local SQL statements from front-end applications, the Borland Database Engine (BDE) translates the statements into BDE API functions.

The local SQL language set

The SQL statements fall into two categories: Data Manipulation Language (DML) and Data Definition Language (DDL).

DML consists of SQL statements used for retrieving, inserting, updating, and deleting table data. SELECT is a DML statement.

DDL consists of SQL statements used for creating, altering, and deleting tables, and for creating and deleting indexes. CREATE TABLE and DROP INDEX are DDL statements.

For a complete introduction to ANSI-standard SQL, see one of the many available third-party books.

Typographical conventions

See the Legend for definitions of the typographical conventions used in this help file for SQL syntax examples.

Note The information covered in this help file pertains to the 32-bit version of the BDE. Certain features discussed are not available in the 16-bit BDE, like nested SELECT queries.

Parameter substitutions in DML statements

Parameter markers can be used in DML statements in place of data values. Parameters are identified by a preceding colon (:). For example:

```
SELECT last_name, first_name
FROM "Customer.db"
WHERE (last_name > :lname) AND (first_name < :fname)
```

Parameters allow the same SQL statement to be used with different data values to be used for comparisons. Parameters are placeholders for data values. At runtime, the front-end application fills the parameter with a value, before the query is executed. When the query is executed, the data values passed into the parameters are substituted for the parameter placeholder and the SQL statement is applied.

Parameters are used to pass data values to be used in WHERE clause comparison and as update atoms in updating statements. Parameters cannot be used to pass values for metadata object names (table and column names).

```
UPDATE Orders
SET ItemsTotal = :TotalParam
WHERE (OrderNo = 1014)
```

Data values passed to SQL statements as parameters are enclosed in quotation marks (where applicable). Thus, when a front-end application supplies CHAR and DATE values for parameters, quotation marks need not be included when the parameter is populated with the values.

Predicate list

Local SQL supports the following predicates:

Predicate	Description
Comparison	Compares two values.
BETWEEN	Compares a value to a range formed by two values.
EXISTS	Determines whether a value exists in a look-up table.
IN	Determines whether a value exists in a list of values or a table.
LIKE	Compares, in part or in whole, one value with another.
IS NULL	Compares a value with an empty, or NULL, value.
SOME/ANY/ALL	Performs quantified comparisons.

Relational operators list

Local SQL supports the following join types:

Join operator	Description
Equi-join	Joins two tables, filtering out non-matching rows.
INNER	Joins two tables, filtering out non-matching rows.
OUTER	Joins two tables, retaining non-matching rows.
Cartesian	Joins two tables, matching each row of one table with each row from the other.
UNION	Concatenates the result set of one query with the result set of another query.
Heterogeneous	Joins two tables in different databases, including differing database types.

Reserved words

Below is an alphabetical list of words reserved by local SQL. Avoid using these reserved words for the names of metadata objects (tables, columns, and indexes). An "Invalid use of keyword error" occurs when reserved words are used as names for metadata objects. If a metadata object must have a reserved word as it name, prevent the error by enclosing the name in quotation marks and prefixing the reference with the table name.

```
ACTIVE                     DO                       MANUAL                  ROLLBACK
ADD                        DOMAIN                   MAX                     SECOND
ALL                        DOUBLE                   MAXIMUM_SEGMENT         SEGMENT
AFTER                      DROP                     MERGE                   SELECT
ALTER                      ELSE                     MESSAGE                 SET
AND                        END                      MIN                     SHARED
ANY                        ENTRY_POINT              MINUTE                  SHADOW
AS                         ESCAPE                   MODULE_NAME             SCHEMA
ASC                        EXCEPTION                MONEY                   SINGULAR
ASCENDING                  EXECUTE                  MONTH                   SIZE
AT                         EXISTS                   NAMES                   SMALLINT
AUTO                       EXIT                     NATIONAL                SNAPSHOT
AUTOINC                    EXTERNAL                 NATURAL                 SOME
AVG                        EXTRACT                  NCHAR                   SORT
BASE_NAME                  FILE                     NO                      SQLCODE
BEFORE                     FILTER                   NOT                     STABILITY
BEGIN                      FLOAT                    NULL                    STARTING
BETWEEN                    FOR                      NUM_LOG_BUFFERS         STARTS
BLOB                       FOREIGN                  NUMERIC                 STATISTICS
BOOLEAN                    FROM                     OF                      SUB_TYPE
BOTH                       FULL                     ON                      SUBSTRING
BY                         FUNCTION                 ONLY                    SUM
BYTES                      GDSCODE                  OPTION                  SUSPEND
CACHE                      GENERATOR                OR                      TABLE
CAST                       GEN_ID                   ORDER                   THEN
CHAR                       GRANT                    OUTER                   TIME
CHARACTER                  GROUP                    OUTPUT_TYPE             TIMESTAMP
CHECK                      GROUP_COMMIT_WAIT_TIME   OVERFLOW                TIMEZONE_HOUR
CHECK_POINT_LENGTH         HAVING                   PAGE_SIZE               TIMEZONE_MINUTE
COLLATE                    HOUR                     PAGE                    TO
COLUMN                     IF                       PAGES                   TRAILING
COMMIT                     IN                       PARAMETER               TRANSACTION
COMMITTED                  INT                      PASSWORD                TRIGGER
COMPUTED                   INACTIVE                 PLAN                    TRIM
CONDITIONAL                INDEX                    POSITION                UNCOMMITTED
CONSTRAINT                 INNER                    POST_EVENT              UNION
CONTAINING                 INPUT_TYPE               PRECISION               UNIQUE
COUNT                      INSERT                   PROCEDURE               UPDATE
CREATE                     INTEGER                  PROTECTED               UPPER
CSTRING                    INTO                     PRIMARY                 USER
CURRENT                    IS                       PRIVILEGES              VALUE
CURSOR                     ISOLATION                RAW_PARTITIONS          VALUES
DATABASE                   JOIN                     RDB$DB_KEY              VARCHAR
DATE                       KEY                      READ                    VARIABLE
DAY                        LONG                     REAL                    VARYING
DEBUG                      LENGTH                   RECORD_VERSION          VIEW
DEC                        LOGFILE                  REFERENCES              WAIT
DECIMAL                    LOWER                    RESERV                  WHEN
DECLARE                    LEADING                  RESERVING               WHERE
DEFAULT                    LEFT                     RETAIN                  WHILE
DELETE                     LEVEL                    RETURNING_VALUES        WITH
DESC                       LIKE                     RETURNS                 WORK
DESCENDING                 LOG_BUFFER_SIZE          REVOKE                  WRITE
DISTINCT                                            RIGHT                   YEAR
```

The following are operators used in local SQL. Avoid using these characters in the names of metadata objects.

```
||, -, *, /, <>, <, >, ,(comma), =, <=, >=, ~=, !=, ^=, (, )
```

SELECT statement

Retrieves data from tables.

```
SELECT [DISTINCT] * | column_list
FROM table_reference
[WHERE predicates]
[GROUP BY group_list]
[HAVING having_condition]
[ORDER BY order_list]
```

Description

Use the SELECT statement to retrieve zero, one, or multiple rows into a dynamic result set. The rows retrieved may come from one or multiple tables (the latter referred to as a "join"). The statement may retrieve data from all or only some of the columns in the source table(s).

The SELECT clause indidcates the columns from which data is retrieved and that compose the result set. Specify the columns of the result set with a comma-separated list comprising: table columns, literal values, and calculated values. Literal values in the columns list may be passed to the SELECT statement via parameters. You cannot use parameters to pass column names to the SQL statement. Use an asterisk to retrieve values from all columns. Columns in the columns list in the SELECT clause may come from more than one table, but can only come from those tables listed in the FROM clause. See Relational Operators for more information on using the SELECT statement to retrieve data from multiple tables (joining tables).

The FROM clause identifies the table(s) from which data is retrieved.

The following statement retrieves data from two columns (CustNo and Company), plus a literal value ("Current") and a calculated value (the date), for all rows of a single table (Customer).

```
SELECT CustNo, Company, "Current", CAST(LastInvoiceDate AS DATE)
FROM Customer
```

Use DISTINCT to limit the retrieved data to only distinct rows. The distinctness of rows is based on the combination of all of the columns in the SELECT clause columns list. DISTINCT can only be used with simple column types like CHAR and INTEGER; it cannot be used with complex column types like BLOB and memo.

```
SELECT DISTINCT CustNo, EXTRACT(YEAR FROM SaleDate) AS Yr
FROM Orders
```

In lieu of a table, a SELECT statement may retrieve rows from an SQL statement (or Paradox-style .QBE query) that has been saved to a file. See the topic Local SQL VIEWs for more information on this.

To limit the rows retrieved by a SELECT statement to only those meeting a specified logical criteria, use a WHERE clause in the statement. To cause aggregation to be performed once for each group of like rows, use a GROUP BY clause. To limit rows retrieved based on agregated values, use a HAVING clause. And to force the retrieved rows to be ordered a certain way in the result set, use an ORDER BY clause. Not all clauses must be in a given SELECT statement, but when they are they must appear in this specific order: SELECT (required), FROM (required), WHERE (optional), GROUP BY (optional), HAVING (optional), and ORDER BY (optional).

SELECT statements can be used as subqueries in INSERT, DELETE, and UPDATE statements. See the topics for these statements for specific requirements for using SELECT subqueries. SELECT subqueries can be used in SELECT statements, but are limited to the WHERE and HAVING clauses. They cannot be used in the SELECT, FROM, GROUP BY, or ORDER BY clauses of a SELECT statement.

Note: There are special limitations on what columns can appear in the SELECT clause of a grouped query. See the GROUP BY clause for more information.

SOME/ANY/ALL predicates

Compares a column value to a column value in multiple rows in a subquery.

```
column_reference comparison_predicate SOME | ANY | ALL (subquery)
```

Description

Use the quantified comparison predicates SOME, ANY, and ALL to filter a table by comparing a column value with multiple comparison values. The quantified comparison predicates are used with comparison predicates to compare a column value to the multiple values in a column of a subquery.

The ANY predicate evaluates TRUE when the accompanying comparison predicate evaluates TRUE for **any** value from the subquery. The SOME predicate operates functionally the same as ANY. For example, using the statement below, for any row to be retrieved from the HOLDINGS table, the value in the PUR_PRICE column need only be greater than any **one** value returned in the subquery's PRICE column.

```
SELECT *
FROM "Holdings.dbf" H
WHERE (H."Pur_Price" > ANY
  (SELECT O."Price"
  FROM "Old_Sales.dbf"))
```

The ALL predicate evaluates TRUE when the accompanying comparison predicate evaluates TRUE for **all** values from the subquery. For example, using the statement below, for any row to be retrieved from the HOLDINGS table, the value in the PUR_PRICE column needs to be greater than **every** value returned in the subquery's PRICE column.

```
SELECT *
FROM "Holdings.dbf" H
WHERE (H." Pur_Price" > ALL
  (SELECT O."Price"
  FROM "Old_Sales.dbf"))
```

Note The subquery providing the comparison values for the quantified comparison predicates may retrieve multiple rows, but can only have one column.

SUBSTRING function

Extracts a substring from a string.

```
SUBSTRING(column_reference FROM start_index [FOR length])
```

Description

Use SUBSTRING to extract a substring from a character value (column, literal, parameter, or caculated values), specified in the column reference.

FROM is the character position at which the extracted substring starts within the original string. The index for FROM is based on the first character in the source value being 1.

FOR is optional, and specifies the length of the extracted substring. If FOR is omitted, the substring goes from the position specified by FROM to the end of the string.

The example below, applied to the literal string "ABCDE" returns the value "BCD".

```
SELECT SUBSTRING("ABCDE" FROM 2 FOR 3) AS Sub
FROM Country
```

In the SELECT statement below only the second and subsequent characters of the NAME column are retrieved. For a column value of "Belgium" that would be a "elgium" result.

```
SELECT SUBSTRING(Name FROM 2)
FROM Country
```

When applied to retrieved data of a SELECT statement, the effect is transient and does not affect stored data. When applied to the update atoms of an UPDATE statement, the effect is persistent and permanently converts the case of the stored values.

Applicability

SUBSTRING can only be used with character columns or literals. To use on values of other data types, the values must first be converted to CHAR using the CAST function.

Note: the SUBSTRING function cannot be used with memo or BLOB columns.

SUM function

Calculates the sum of values for a column.

```
SUM([ALL] column_reference | DISTINCT column_reference)
```

Description

Use SUM to calculate the total of a numeric value (column, literal, parameter, or caculated values) using a dataset. The dataset may be the entire table, a filtered dataset, or a logical group produced by a GROUP BY clause. Column values of zero are included in the aggregation. NULL column values are not counted in the calculation. If the number of qualifying rows is zero, SUM returns a NULL value.

```
SELECT SUM(ItemsTotal)
FROM Orders
```

Constant values (such as literals and parameter values) can only be used with the SUM function if part of a calculation involving a table column. If SUM is applied to a calculation, the calculation is performed for each row and then the total calculated based on the data as modified by the calculation.

```
SELECT SUM(Salary + (Sarary * 0.1))
FROM "Employee.db"
```

ALL returns the smallest value for all rows. When DISTINCT is not specified, ALL is the implied default.

DISTINCT ignores duplicate values when calculating the smallest value in the specified column.

MIN returns the smallest value in a column or a calculation using a column performed for each row (a calculated field).

```
SELECT SUM(ItemsTotal), SUM(ItemsTotal * 0.0825) AS TotalTax
FROM Orders
```

When used with a GROUP BY clause, SUM returns one row for each group with a total of the column's value for that group (one subtotal for each group). The statement below aggregates the total value for the order totals column in the Orders table, producing a subtotal for each company in the Company table.

```
SELECT C."Company", SUM(O."ItemsTotal") AS SubTotal
FROM "Customer.db" C, "Orders.db" O
WHERE (C."CustNo" = O."CustNo")
GROUP BY C."Company"
ORDER BY C."Company"
```

Applicability

SELECT statements. SUM operates only on numeric values. To use SUM on non-numeric values, first use the CAST function to convert the column to a numeric type.

System INIT settings

The System INIT settings are the settings the BDE uses to start an application. This information is stored in the Windows Registry.

To display these settings:

1. Click the Configuration tab in the left pane and expand Configuration in the settings tree.
2. Expand System in the tree.
3. Expand INIT in the tree.

   The settings appear on the Definition tab in the right pane.

In the following table, **Parameter** lists all the configurable system and network parameters, with their current values. When your BDE application is first installed, all values are set to their defaults.

Description briefly notes the purpose of the selected parameter.

To change a setting, highlight the desired parameter. Then replace its old value with a new value. When edits are complete, choose Object|Apply or click the Apply button to save the changes.

Parameter	Description
AUTO ODBC	Set to TRUE to import installed ODBC drivers and datasources every time BDE is initialized. Default: FALSE. **Note:** Setting AUTO ODBC to TRUE is no longer recommended because of the Virtual configuration mode setting. For more information, see Options dialog box.
DATA REPOSITORY	The name of the active data dictionary.
DEFAULT DRIVER	Driver first tried when TYPE is FILE and the table name has no extension. Also, for files with a .DBF extension, specifies whether to use the dBASE or FoxPro driver.
LANGDRIVER	System language driver that defaults to the OEM driver appropriate for a country's version of Windows; for example, 'ascii' ANSI (DBWINUS0) for US workstations.
LOCAL SHARE	The ability to share access to local data between an active BDE application and an active non-BDE application. Set to TRUE if you need to work with the same files through both a BDE and a non-BDE application at the same time. (It is not necessary to set LOCAL SHARE to TRUE if you do not need to have both applications open at the same time.) Default: FALSE.
LOW MEMORY USAGE LIMIT	Maximum amount of low memory the BDE will attempt to use, in kilobytes. Default: 32.
MAXBUFSIZE	Maximum amount of memory for database data cache, in kilobytes. Can be any integer between MINBUFSIZE and 65,535 (inclusive). Must be a multiple of 128. Default: 2048.
MAXFILEHANDLES	Maximum number of file handles the BDE uses. Can be any integer ranging from 5 to 4096. High values improve performance but use more Windows resources. Default: 48.
MEMSIZE	Maximum amount of memory that BDE will use, in megabytes. Default: 16 MB. Maximum 205. Any number specified over 205 sets MEMSIZE to 205.
MINBUFSIZE	Minimum amount of memory for database data cache, in kilobytes. Can be any integer between 32 and 65535. Must be less than the total amount of RAM available to Windows. Default: 128.
MTS POOLING	Set to TRUE to enable MTS pooling, thus improving the initial

	connection time of opening a database. Default: FALSE.
SHAREDMEMLOCATION	Preferred address of shared memory manager. The shared memory manager and shared buffer manager load at a preferred address. If this address is used, the system can find a suitable address (first process only). If this address conflicts with another application, you might want to modify it. If the shared managers cannot be loaded into the default, or specified memory address, one attempt to allow the operating system to find a memory address is made. If a memory address if specified outside the boundry (below), the default is used. Default (for Windows 95): 0xEBDE0000 Minimum Value (for Windows 95): 0x90000000 Maximum Value (for Windows 95): 0xFFFF0000 Default (for Windows NT): 0x6BDE0000 Minimum Value (for Windows NT): 0x10000000 Maximum Value (for Windows NT): 0x7F000000 The range setting for SHAREDMEMLOCATION is 9000 to FFFF for Windows 95 and 1000 to 7F00 for Windows NT. Only the upper word is used to set SHAREDMEMLOCATION.
SHAREDMEMSIZE	Maximum amount of memory BDE will use for shared resources, in kilobytes. Default: 2048 KB. Minimum: 2048 KB. If a value below 2048 is specified, 2048 will be used. File handles, drivers, and system, client and table objects are shared resources and SHAREDMEMSIZE should be increased when using many of these objects.
SQLQRYMODE	Method for handling queries to SQL data. Can be NULL, SERVER, or LOCAL. The SQLQRYMODE parameter appears only if a Borland SQL Link driver is installed. Default: NULL.
SYSFLAGS	Internal BDE setting. Do not modify.
VERSION	An internal setting that describes the current version of the BDE. Do not modify.

Table correlation names

Table correlation names are used to explicitly associate a column with the table from which it comes. This is especially useful when multiple columns of the same name appear in the same query, typically in multi-table queries. A table correlation name is defined by following the table reference in the FROM clause of a SELECT query with a unique identifier. This identifier, or table correlation name, can then be used to prefix a column name.

If the table name is not a quoted string, the table name is the default implicit correlation name. An explicit correlation name the same as the table name need not be specified in the FROM clause and the table name can prefix column names in other parts of the statement.

```
SELECT *
FROM Customer
  LEFT OUTER JOIN Orders
    ON (Customer.CustNo = Orders.CustNo)
```

If the table name is a quoted string, you need to do one of the following:

Prefix column names with the exact quoted string used for the table in the FROM clause.

```
SELECT *
FROM "Customer.db"
  LEFT OUTER JOIN "Orders.db"
    ON ("Customer.db".CustNo = "Orders.db".CustNo)
```

Use the full table name as a correlation name in the FROM clause (and prefix all column references with the same correlation name).

```
SELECT *
FROM "Customer.db" CUSTOMER
  LEFT OUTER JOIN "Orders.db" ORDERS
    ON (CUSTOMER.CustNo = ORDERS.CustNo)
```

Use a distinctive token as a correlation name in the FROM clause (and prefix all column references with the same correlation name).

```
SELECT *
FROM "Customer.db" C
  LEFT OUTER JOIN "Orders.db" O
    ON (C.CustNo = O.CustNo)
```

Table names

ANSI-standard SQL confines each table name to a single word comprised of alphanumeric characters and the underscore symbol, "_". Local SQL, however, is enhanced to support the Paradox and dBASE feature of non-alpha characters and spaces embedded in table names. To use table names with these embedded characters, you must enclose the table name in quotation marks (double or single).

```
SELECT *
FROM "Customer Revenue.db"
```

Local SQL supports full file and path specifications in table references. Table references with path or filename extensions must be enclosed in single or double quotation marks. For example:

```
SELECT *
FROM 'Parts.dbf'
```

```
SELECT *
FROM "c:\sample\parts.dbf"
```

Local SQL also supports BDE aliases in table references. To do this, enclose the BDE alias in colons, prefix the table name with that colon-enclosed alias name, and then enclose the whole reference in quotation marks. For example:

```
SELECT *
FROM ":DBDEMOS:Customer.db"
```

If you omit the file extension for a local table name, the table is assumed to be the table type specified in the BDE configuration. The default table type is specified either in the default driver setting or in the default driver type for the standard alias associated with the query.

Finally, local SQL permits table names to duplicate SQL keywords as long as those table names are enclosed in single or double quotation marks. For example:

```
SELECT Passid
FROM "Password"
```

Time formats

Local SQL expects time literals to be in the format hh:mm:ss AM/PM; where hh are the hours, mm the minutes, and ss the seconds. When inserting new data with a time value, the AM/PM designator is optional and is case-insensitive ("AM" is the same as "am"). The time literal must be enclosed in quotation marks.

```
INSERT INTO WorkOrder
(ID, StartTime)
VALUES ("B00120","10:30:00 PM")
```

Indicate which half of the day (morning or after noon) a time literal falls under in one of two ways. If an AM or PM marker is specified, that determines the half of the day. If no AM/PM designator is specified, the hour field is compared to 12. If the hour is less than twelve, the time is in the AM; if greater than 12, after noon. The hour field overrides an AM/PM designator. For example, the time literal "15:03:22 AM" is translated as "3:03:22 PM". For example, the statement below only retrieves rows where the Event_Time column contains a value of "1:00:00 PM" and excludes those of "1:00:00 AM".

```
SELECT *
FROM "Events.db"
WHERE (Event_Time = "13:00:00")
```

TRIM function

Removes the trailing or leading character, or both, from a string.

```
TRIM([LEADING|TRAILING|BOTH] [trimmed_char] FROM column_reference)
```

Description

Use TRIM to delete the leading or trailing character, or both, from a character value (column, literal, parameter, or caculated values). The TRIM function only deletes characters located in the specified position.

The first parameter indicates the position of the character to be deleted, and has one of the following values:

Value	Description
LEADING	Deletes the character at the left end of the string.
TRAILING	Deletes the character at the right end of the string.
BOTH	Deletes the character at both ends of the string.

The trimmed character parameter specifies the character to be deleted, if present. Case-sensitivity is applied for this parameter. To make TRIM case-insensitive, use the UPPER function.

FROM specifies the character value from which to delete the character. The column reference for FROM can be a table column or a character literal.

Example variations:

TRIM syntax	Result
TRIM(LEADING "_" FROM "_ABC_")	"ABC_"
TRIM(TRAILING "_" FROM "_ABC_")	"_ABC"
TRIM(BOTH "_" FROM "_ABC_")	"ABC"
TRIM(BOTH "A" FROM "ABC")	"BC"

When applied to retrieved data of a SELECT statement, the effect is transient and does not affect stored data. When applied to the update atoms of an UPDATE statement, the effect is persistent and permanently converts the case of the stored values.

Applicability

TRIM can only be used with character columns or literals. To use on values of other data types, the values must first be converted to CHAR using the CAST function.

Note: the TRIM function cannot be used with memo or BLOB columns.

UNION join

Concatenates the rows of one table to the end of another table.

```
SELECT col_1 [, col_2, ... col_n]
FROM table_reference
UNION [ALL]
SELECT col_1 [, col_2, ... col_n]
FROM table_reference
```

Description

Use the UNION join to add the rows of one table to the end of another similarly structured SELECT query result sets. The SELECT statement for the source and joining tables must include the same number of columns for them to be UNION compatible. The table structures themselve need not be the same as long as those column included in the SELECT statements are.

```
SELECT CustNo, Company
FROM Customers
UNION
SELECT CustNo, Company
FROM Old_Customers
```

Matching data types for a column is not always mandatory for data retrieved by the UNION across the multiple tables. If there is a data type difference between two tables for a given column, an error occurs if the same column from the second (or subsequent) table would lose data. For example, if the first table 's column is of type DATE and the second table's of type TIMESTAMP, part of the TIMESTAMP value would be lost when put into a lesser DATE type column. A "Type mismatch in expression" error is generated for these situations. In general, when there are column differences between the tables, use the CAST function to convert the columns to a compatible type.

```
SELECT S.ID, CAST(S.Date_Field AS TIMESTAMP)
FROM Source S
UNION ALL
SELECT J.ID, J.Timestamp_Field
FROM Joiner J
```

Matching names is not mandatory for result set columns retrieved by the UNION across the multiple tables. Column name differences between the multiple source tables are automatically handled. If the first column of two tables has a different name, the first column in the UNION result set will use that from the first SELECT statement.

By default, non-distinct rows are aggregated into single rows in a UNION join. Use ALL to retain non-distinct rows.

To join two tables with UNION where one table does not have a column included by another, a compatible literal or expression may be used instead in the SELECT statement missing the column. For example, if there is no column in the JOINING table corresponding to the NAME column in SOURCE an expression is used to provide a value for a pseudo JOINING.NAME column. Assuming SOURCE.NAME is of type CHAR(10), the CAST function is used to convert an empty character string to CHAR(10).

```
SELECT S.ID, S.Name
FROM Source S
UNION ALL
SELECT J.ID, CAST("" AS CHAR(10))
FROM Joiner J
```

Unsupported language

The following SQL-92 language elements are not used in local SQL.

ALLOCATE CURSOR (Command)
ALLOCATE DESCRIPTOR (Command)
ALTER DOMAIN (Command)
CASE (Expression)
CHECK (Constraint)
CLOSE (Command)
COALESCE (Expression)
COMMIT (Command)
CONNECT (Command)
CONVERT (Function)
CORRESPONDING BY (Expression)
CREATE ASSERTION (Command)
CREATE CHARACTER SET (Command)
CREATE COLLATION (Command)
CREATE DOMAIN (Command)
CREATE SCHEMA (Command)
CREATE TRANSLATION (Command)
CREATE VIEW (Command)
CROSS JOIN (Relational operator)
CURRENT_DATE (Function)
CURRENT_TIME (Function)
CURRENT_TIMESTAMP (Function)
DEALLOCATE DESCRIPTOR (Command)
DEALLOCATE PREPARE (Command)
DECLARE CURSOR (Command)
DECLARE LOCAL TEMPORARY TABLE (Command)
DESCRIBE (Command)
DISCONNECT (Command)
DROP ASSERTION (Command)
DROP CHARACTER SET (Command)
DROP COLLATION (Command)
DROP DOMAIN (Command)
DROP SCHEMA (Command)

DROP TRANSLATION (Command)
DROP VIEW (Command)
EXCEPT (Relational operator)
EXECUTE (Command)
EXECUTE IMMEDIATE (Command)
FETCH (Command)
FOREIGN KEY (Constraint)
GET DESCRIPTOR (Command)
GET DIAGNOSTICS (Command)
GRANT (Command)
INTERSECT (Relational operator)
MATCH (Predicate)
NATURAL (Relational operator)
NULLIF (Expression)
OPEN (Command)
OVERLAPS (Predicate)
PREPARE (Command)
REFERENCES (Constraint)
REVOKE (Command)
ROLLBACK (Command)
Row value constructors
SET CATALOG (Command)
SET CONNECTION (Command)
SET CONSTRAINTS MODE (Command)
SET DESCRIPTOR (Command)
SET NAMES (Command)
SET SCHEMA (Command)
SET SESSION AUTHORIZATION (Command)
SET TIME ZONE (Command)
SET TRANSACTION (Command)
TRANSLATE (Function)
UNIQUE (Predicate)
USING (Relational operator)

Updatable queries

Definition...

Single-table queries

Queries that retrieve data from a single table are updatable provided that:

- There is no DISTINCT key word in the SELECT.
- Everything in the SELECT clause is a simple column reference or a calculated column, no aggregation is allowed. Calculated columns remain read-only.
- The table referenced in the FROM clause is an updatable base table.
- There is no GROUP BY or HAVING clause.
- There are no subqueries in the statement.
- There is no ORDER BY clause.

The read-only effect of an ORDER BY clause is negated and the query updatable if the ORDER BY clause uses a single column and there is a dBASE single-column primary or secondary index based on that same field. dBASE compound (expression) indexes will not negate the read-only effect of an ORDER BY clause. A Paradox single- or multi-field primary index will make the query updatable if the ORDER BY uses exactly the same columns (in the same order) as the index. Paradox secondary indexes will not negate the read-only effect of an ORDER BY clause.

Multi-table queries

All queries that join two or more tables will produce a read-only result set.

Calculated fields

For updateable queries with calculated fields, an additional field property identifies a result field as both read-only and calculated. Every call to the BDE function DbiPutField causes recalculation of any dependent fields.

UPDATE statement

Modifies one or more existing rows in a table.

```
UPDATE table_reference
SET column_ref = update_atom [, column_ref = update_atom...]
[WHERE predicates]
```

Description

Use the UPDATE statement to modify one or more column values in one or more existing rows in a single table per statement.

Use a table reference in the UPDATE clause to specify the table to receive the data changes.

The SET clause is a comma-separated list of update expressions. Each expression comprises the name of a column, the assignment operator (=), and the update value (update atom) for that column. The update atoms in any one update expression may be literal values, singleton return values from a SELECT subquery, or calculated values. Subqueries supplying an update atom for an update expression must return a singleton result set (one row) and return only a single column.

```
UPDATE SalesInfo
SET TaxRate = 0.0825
```

When update atoms come from a SELECT subquery, one SELECT statement must be provided for each column being updated in the target table. Each SELECT subquery providing an update atom must be enclosed in parentheses. The following statement updates the two columns OnHand and InventoryDate in the table INVENTORY, each with a separate SELECT subquery.

```
UPDATE Inventory
SET OnHand = OnHand –
  (SELECT SUM(Orders.QtySold)
  FROM Orders
  WHERE (Orders.PartNo = Inventory.PartNo) AND
    (Orders.OrderDate BETWEEN "10/01/1999" AND "10/31/1999")),
  InventoryDate = OnHand –
  (SELECT MAX(Orders.OrderDate)
  FROM Orders
  WHERE (Orders.PartNo = Inventory.PartNo) AND
    (Orders.OrderDate BETWEEN "10/01/1999" AND "10/31/1999")),
```

The optional WHERE clause restricts updates to a subset of rows in the table. If no WHERE clause is specified, all rows in the table are updated using the SET clause update expressions.

```
UPDATE SalesInfo
SET TaxRate = 0.0825
WHERE (State = "CA")
```

The UPDATE statement only supports SELECT subqueries in the WHERE clause. References to tables other than the one in which rows are update or columns in such tables are only possible in SELECT subqueries.

UPPER function

Converts all characters to uppercase.

```
UPPER(column_reference)
```

Description

Use UPPER to convert all of the characters in a character value (column, literal, parameter, or caculated values) to uppercase. For example, in the SELECT statement below the values in the NAME column are treated as all in uppercase. Because the same conversion is applied to both the filter column and comparison value in the WHERE clause, the filtering is effectively case-insensitive.

```
SELECT Name, Capital, Continent
FROM Country
WHERE UPPER(Name) LIKE UPPER("Pe%")
```

When applied to retrieved data of a SELECT statement, the effect is transient and does not affect stored data. When applied to the update atoms of an UPDATE statement, the effect is persistent and permanently converts the case of the stored values.

Applicability

UPPER can only be used with character columns or literals. To use on values of other data types, the values must first be converted to CHAR using the CAST function.

Note: The UPPER function cannot be used with memo or BLOB columns.

Note: The lower case version of a given character is determined by the language driver used.

WHERE clause

Specifies filtering conditions for a SELECT or UPDATE statement.

```
WHERE predicates
```

Description

Use a WHERE clause to limit the effect of a SELECT or UPDATE statement to a subset of rows in the table. In statements that support a WHERE clause, the clause is optional. (An exception where the WHERE clause is required is the equi-join.)

The value for a WHERE clause is one or more logical expressions, or predicates, that evaluate to TRUE or FALSE for each row in the table. Only those rows where the predicates evaluate to TRUE are retrieved by a SELECT statement or modified by an UPDATE statement. For example, the SELECT statement below retrieves all rows where the STATE column contains a value of "CA".

```
SELECT Company, State
FROM Customer
WHERE State = "CA"
```

A column used in the WHERE clause of a statement is not required to also appear in the SELECT clause of that statement. In the preceding statement, the State column could be used in the WHERE clause even if it was not also in the SELECT clause.

Multiple predicates must be separated by one of the logical operators OR or AND. Each predicate can be negated with the NOT operator. Parentheses can be used to isolate logical comparisons and groups of comparisons to produce different row evaluation criteria. For example, the SELECT statement below retrieves all rows where the STATE column contains a value of "CA" and those with a value of "HI".

```
SELECT Company, State
FROM Customer
WHERE (State = "CA") OR (State = "HI")
```

The SELECT statement below retrieves all rows where the SHAPE column is "round" or "square", but only if the the COLOR column also contains "red". It would not retrieve rows where, for example, the SHAPE is "round" and the COLOR "blue".

```
SELECT Shape, Color, Cost
FROM Objects
WHERE ((Shape = "round") OR (Shape = "square")) AND
  (Color = "red")
```

But without the parentheses to override the order of precedence of the logical operators, as in the statement that follows, the results are very different. This statement retrieves the rows where the SHAPE is "round", regardless of the value in the COLOR column. It also retrieves rows where the SHAPE column is "square", but only when the COLOR column contains "red". Unlike the preceding variation of this statement, this one would retrieve rows where the SHAPE is "round" and the COLOR "blue".

```
SELECT Shape, Color, Cost
FROM Objects
WHERE Shape = "round" OR Shape = "square" AND
  Color = "red"
```

Subqueries are supported in the WHERE clause. A subquery works like a search condition to restrict the number of rows returned by the outer, or "parent" query. Such subqueries must be valid SELECT statements. SELECT subqueries can be correlated by referring to columns in the outer (or "parent") statement. In the following statement, the subquery is said to be correlated because it refers to the C.Acct_Nbr column of the parent statement.

```
UPDATE "Clients.dbf" C
SET C.Interests = "Buys lotsa stock"
WHERE EXISTS
  (SELECT H.Acct_Nbr, SUM(H.Shares)
```

```
FROM "Holdings.dbf" H
WHERE (H.Acct_Nbr = C.Acct_Nbr)
GROUP BY H.Acct_Nbr
HAVING (SUM(H.Shares) >= 50000))
```

Note: Column correlation names cannot be used in filter comparisons in the WHERE clause. Use the actual column name.

Note: A WHERE clause filters data prior to the aggregation of a GROUP BY clause. For filtering based on aggregated values, use a HAVING clause.

Hint: When filtering using most column types, the simple comparison predicates suffice. But when filtering based on the contents of a MEMO column, you must use the LIKE predicate.

Hint: Columns devoid of data contain NULL values. To filter using such column values, use the IS NULL predicate.

Applicability
SELECT (filtering on non-aggregated columns) and UPDATE statements