

# **SQL Reference**

## **ExtenDB Parallel Server**

**Version 1.1**

**October 2006**



# ExtenDB Parallel Server SQL Reference

## Table of Contents

Table of Contents .....	2
1 Introduction .....	4
1.1 Overview .....	4
2 Data Types .....	5
3 Expressions .....	6
3.1 Expressions .....	6
3.2 Operators .....	6
3.2.1 Logical Operators .....	6
3.2.2 Mathematical Operators .....	8
4 Functions .....	9
4.1 Aggregate Functions .....	9
4.2 Mathematical Functions .....	11
4.3 Date & Time Functions .....	13
4.4 String and Character Functions .....	14
4.5 Other Functions .....	16
5 SQL Commands .....	17
5.1 Data Definition Statements .....	17
5.1.1 CREATE TABLESPACE .....	17
5.1.2 CREATE TABLE .....	18
5.1.3 DROP TABLE .....	23
5.1.4 TRUNCATE TABLE .....	23
5.1.5 ALTER TABLE ... ADD COLUMN .....	23
5.1.6 ALTER TABLE ... DROP COLUMN .....	23
5.1.7 ALTER TABLE ... ADD PRIMARY KEY .....	24
5.1.8 ALTER TABLE ... ADD FOREIGN KEY .....	24
5.1.9 ALTER TABLE ... DROP CONSTRAINT .....	25
5.1.10 ALTER TABLE ... DROP PRIMARY KEY .....	25
5.1.11 ALTER TABLE ... ALTER COLUMN .....	26
5.1.12 ALTER TABLE ... OWNER .....	26
5.1.13 ALTER TABLE ... SET TABLESPACE .....	26
5.1.14 RENAME TABLE .....	26
5.1.15 CREATE INDEX .....	27
5.1.16 DROP INDEX .....	27
5.1.17 CLUSTER .....	27
5.1.18 CREATE VIEW .....	28
5.1.19 DROP VIEW .....	28
5.2 Data Manipulation Statements .....	29
5.2.1 INSERT .....	29
5.2.2 UPDATE .....	29
5.2.3 DELETE .....	30
5.2.4 SELECT .....	31
5.3 Users and Privileges .....	32
5.3.1 CREATE USER .....	32
5.3.2 ALTER USER .....	33
5.3.3 DROP USER .....	33
5.3.4 GRANT .....	33
5.3.5 REVOKE .....	34
5.4 Other Commands .....	35

# ExtenDB Parallel Server SQL Reference

5.4.1	SHOW DATABASES.....	35
5.4.2	SHOW TABLES.....	36
5.4.3	SHOW VIEWS.....	36
5.4.4	DESCRIBE <table>.....	36
5.4.5	DESCRIBE <view>.....	36
5.4.6	SHOW INDEXES ON <table>.....	37
5.4.7	SHOW CONSTRAINTS ON <table>.....	37
5.4.8	SHOW USERS.....	37
5.4.9	SHOW STATEMENTS.....	38
5.4.10	KILL.....	38
5.4.11	UPDATE STATISTICS.....	38
5.4.12	ANALYZE.....	39
5.4.13	VACUUM.....	39

# ExtenDB Parallel Server SQL Reference

## 1 Introduction

---

### 1.1 Overview

The ExtenDB Parallel Server's supported SQL incorporates ANSI-92 and an overlap of the features in more popular open source databases like PostgreSQL, MySQL, and SAP-DB/MaxDB. To be clear, using ExtenDB in conjunction with the underlying database does not mean that you will have access to the full functionality of that particular database. Nonetheless, ExtenDB allows for a lot of customization in the xdb.config file to allow the DBA to define additional functions as well as map ExtenDB commands to the underlying database.

**This is not intended to be a comprehensive analysis of the SQL language, but is intended to provide information regarding the supported SQL and its syntax, and to point out noteworthy ExtenDB implementation details to bear in mind.**

# ExtenDB Parallel Server SQL Reference

## 2 Data Types

The following data types are supported:

Data Type	Comments
BIGINT, INT8	
BIT, VARBIT	
BLOB	Maps to BYTEA in PostgreSQL
BOOLEAN	(May not be supported by all underlying databases)
CHAR[ACTER] (length)	Width is fixed to length.
DATE	Accepts format YYYY-MM-DD or YYYYMMDD
DATETIME	Combination of date and time
DOUBLE PRECISION	8 byte floating point number
DECIMAL (length, decimals)	Mapped to NUMERIC
FLOAT[ (length, decimals)]	
INT[EGER], INT4	
INTERVAL [YEAR MONTH DAY HOUR MINUTE TO YEAR MONTH DAY HOUR MINUTE]	For time intervals
NCHAR[ACTER] (length)	For multi-language support, like CHAR
NUMERIC (length, decimals)	
NVARCHAR[ACTER] (length)	For multi-language support, like VARCHAR
REAL[(length, decimals)]	
SERIAL, BIGSERIAL	4 and 8 byte serial
SMALLINT, INT2	
TEXT	Acts as a CLOB
TIME	Accepts hh:mm:ss or hhmmss format
TIMESTAMP	Combination of date and time
VARCHAR (length) or CHAR VARYING (length)	Varying number of characters, with a maximum of specified length

## 3 Expressions

### 3.1 Expressions

Logical expressions are typically found as part of the WHERE clause of various statements in determining the rows that will be effected by the statement. The operands of a logical expression to be evaluated by a logical operator may in turn be a logical expression, or an SQL expression of any of the supported data types.

Operators are discussed in the following section. In addition, available functions are covered in a later chapter.

### 3.2 Operators

The logical and mathematical operators that are used in expressions that ExtenDB recognizes appear below, in ascending order of precedence by line.

Parentheses can also be used in expression to determine precedence.

#### 3.2.1 Logical Operators

```
OR
AND
NOT
BETWEEN
CASE, WHEN, THEN, ELSE
=, !=, <>, >=, >, <=, <,
IS, LIKE, IN
```

#### BETWEEN

The between operator allows for the comparison of a range of values.

Example:

```
SELECT *
FROM customer
WHERE est_income BETWEEN 100000 and 200000
```

#### CASE

```
CASE expression
  WHEN compare_expression THEN result_value
  [WHEN compare_expression THEN result_value ...]
  [ELSE result_value]
```

# ExtenDB Parallel Server SQL Reference

END

```
CASE
  WHEN condition THEN result_value
  [WHEN condition THEN result_value ...]
  [ELSE result_value]
END
```

There are two forms of CASE. The first form allows for comparing a single expression against a list of possible values, while the second form allows for various conditions to be evaluated.

In either case, an ELSE clause may appear to specify a default value.

Example:

```
SELECT CASE custtype
  WHEN 'P' THEN 'Platinum'
  WHEN 'G' THEN 'Gold'
  WHEN 'S' THEN 'Silver'
  ELSE 'Standard'
END
FROM customer
WHERE state = 'CA'

SELECT CASE
  WHEN custtype = 'P' THEN 'Red Carpet'
  WHEN ordercount <= 1 THEN 'New'
  ELSE 'Standard'
END
```

## IS

Example:

```
SELECT *
FROM customer
WHERE lastname IS NULL
```

## LIKE

LIKE is used to match part of a string. Wildcard characters are used as part of a string pattern.

Example: Select all customers whose names begin with the letter B.

```
SELECT *
FROM customer
WHERE lastname LIKE 'B%'
```

## IN

IN is used to compare an expression to a set of expressions

# ExtenDB Parallel Server SQL Reference

Example: Select all customers in New England states

```
SELECT *  
FROM customer  
WHERE state IN ('NH', 'VT', 'CT', 'MA', 'ME')
```

## 3.2.2 Mathematical Operators

```
- , +  
* , / , DIV , MOD  
^  
- (numeric negation)
```

Standard mathematical operators '-', '+', '\*', and '/' are available. In addition the following are available: '^' (raise to a power), DIV (divisor—no remainder), and MOD (Modulo- remainder).



## 4 Functions

The following functions are supported as part of SQL Expressions. One caveat is that for the most of these, the underlying database used must support them. Please consult your database documentation. Commonly used functions appear below.

You may also map functions from one to another or create templates that are substituted at runtime. Please see the Administrator's Guide for more details on how to set that up in the xdb.config file.

### 4.1 Aggregate Functions

Aggregate functions differ from regular functions in that regular functions are applied on a row by row basis, where as aggregate functions apply to a group of rows, either the entire set of rows being returned based on the specified conditions, or a subset of that, as determined by the GROUP BY clause.

#### **AVG(n)**

Calculates the average or mean of a numeric expression.

#### **COUNT(\*)** **COUNT([DISTINCT] expr)**

COUNT(\*) counts the number of rows that make up a group of rows.  
COUNT(DISTINCT expr) counts the number of unique appearances of the expression in the projected results.

#### **MAX(expr)**

Calculates the maximum value for a group of rows, whether the expression is of a numeric or string type.

#### **MIN(expr)**

Calculates the minimum value for a group of rows, whether the expression is of a numeric or string type.

#### **SUM(n)**

Calculates the SUM of an expression for a group of rows.

#### **STDDEV(n)**

Calculates the standard deviation of a group of rows.

# ExtenDB Parallel Server SQL Reference

## **VARIANCE(n)**

Calculates the statistical variance of a group of rows.

## **4.2 Mathematical Functions**

### **ABS(n)**

Calculates the absolute value of the numeric expression.

### **CEIL[ING](n)**

Calculates the ceiling of the numeric expression, that is, the next highest integer. CEIL(1.2) is equal to 2, and CEIL(-1.2) is equal to -1.

### **LN(n)**

Calculates the natural logarithm of numeric expression n.

### **LOG(n1)**

### **LOG(n1, n2)**

In the first form with one parameter, it is equivalent of LN(), the natural logarithm function.

In the second form, it calculates the logarithm of n1 for base n2.

### **MOD(n1, n2)**

Modulo. Calculates the remainder of n1 divided by n2. This is equivalent to n1 mod n2 or n1 % n2.

### **PI()**

Returns the value of Pi.

### **POWER(n1, n2)**

Calculates n1 to the power of n2.

### **ROUND(n)**

Calculates n rounded to the nearest integer.

### **SIGN(n)**

Calculates the sign of numeric expression n. The return value is -1 if n is negative, 0, if n is 0, or 1 if n is positive.

### **SQRT(n)**

Calculates the square root of n.

### **TRUNC(n1)**

# ExtenDB Parallel Server SQL Reference

## **TRUNC(n1,n2)**

In the first form, with a single parameter, the decimal places are truncated.

In the second form, with  $n1 > 0$ , the number  $n1$  is truncated to  $n2$  decimal places. If  $n2 < 0$  then  $n2$  places in front of the decimal point are set to 0.

## ***Trigonometric Functions***

### **ACOS(n)**

Returns the arc cosine of  $n$ .

### **ASIN(n)**

Returns the arc sin of  $n$ .

### **ATAN(n)**

Returns the arc tangent of  $n$ .

### **COS(n)**

Returns the cosine of  $n$ .

### **COT(n)**

Returns the cotangent of  $n$ .

### **DEGREES(n)**

Returns the radians value of  $n$  converted to degrees.

### **RADIANS(n)**

Returns the degrees value of  $n$  converted to radians.

### **SIN(n)**

Returns the sine of  $n$ .

### **TAN(n)**

Returns the tangent of  $n$ .

## 4.3 Date & Time Functions

### **AGE (timestamp [, timestamp])**

Returns the age or age difference of the timestamp.

### **CURRENT\_DATE**

Returns the current date.

### **CURRENT\_TIME**

Returns the current time.

### **CURRENT\_TIMESTAMP**

Returns the current date and time as a timestamp (datetime) type.

### **DATE\_PART(s, timestamp)**

Returns the part element specified in the first argument from the timestamp.

Example:

```
select DATE_PART('month', '2001-12-31') from table1
```

Returns

12

### **DATE\_TRUNC(s, timestamp)**

Returns the timestamp truncated to the specified precision.

### **EXTRACT (field from [timestamp|interval])**

Extracts the specified field from the timestamp

```
select EXTRACT (day from timestamp '1999-12-31')
```

Returns

31

### **ISFINITE (timestamp/interval)**

Tests for finite timestamp

# ExtenDB Parallel Server SQL Reference

## **LOCALTIME [(int)]**

Returns time of day

## **LOCALTIMESTAMP [(int)]**

Returns current timestamp

## **NOW()**

Returns current timestamp with time zone.

## **TIMEOFDAY()**

Returns the current date and time as a timestamp.

## **4.4 String and Character Functions**

### **ASCII(s)**

Returns the ASCII value of string s.

### **BIT\_LENGTH(s)**

Returns the number of bits in the string.

### **CHAR\_LENGTH(s)**

### **CHARACTER\_LENGTH(s)**

Returns the number of characters in the string.

### **CONVERT(str using conversion\_name)**

Change encoding using specified conversion name.

### **INITCAP(s)**

Returns the string passed in transformed such that the first letter in upper case, and the other letters are in lower case.

### **LENGTH(s)**

Returns the number of characters in string s.

### **LOWER(s)**

Returns string s with all characters converted to lower case.

# ExtenDB Parallel Server SQL Reference

## **LPAD(s1, n, s2)**

Returns string based on string s1 set to length n, with any extra padding needed taken from s2.

## **LTRIM(s1 [, s2])**

Remove the longest string containing only characters from s2 (a space by default) from the start of string s1.

## **OCTET\_LENGTH(s)**

Returns the number of bytes in the string

## **OVERLAY(s1 PLACING s2 FROM int [FOR int])**

Replaces characters in string s1 with those from s2 starting at the position specified.

## **POSITION(s1 IN s2)**

Returns the character position in which string s1 is found in string s2.

## **RPAD(s1,n,s2)**

Returns string based on string s set to length n, with any extra padding needed taken from s2.

## **RTRIM(s1)**

Returns string s1 with trailing spaces removed.

## **SUBSTR(s1, n1, n2)**

## **SUBSTRING(s1 FROM n1 [ FOR n2])**

Returns the substring of string s1, starting at position n1, continuing for n2 characters.

## **TRIM(leading | trailing | both c from s1)**

Remove the longest string containing only the characters c (a space by default) from the start, end, or both ends of the string s1.

## **UPPER(s)**

Returns the string s converted into all upper case letters.

## ***4.5 Other Functions***

### **CAST(expr1 as datatype)**

Casts expression to datatype.

### **DATABASE()**

Returns the current database being used.

### **VERSION()**

Returns the version of the database server being used.



## 5 SQL Commands

### 5.1 Data Definition Statements

#### 5.1.1 CREATE TABLESPACE

PostgreSQL and Bizgres Only

Syntax:

```
CREATE TABLESPACE tablespace
    LOCATION 'path' ON NODE[S] nodenum [,nodenum...]
    [,LOCATION 'path' ON NODE[S] nodenum [,nodenum...]]
```

Bizgres and PostgreSQL support the creation of tablespaces. This gives database administrators more flexibility in specifying locations of tables on physical disks. For example, a DBA may want to put a particularly large and often-used fact table on its own set of RAID hard drives, while keeping other tables in another location. Or, the DBA may want to put indexes on a table in a different location than the table itself.

Tablespaces are not required. If not used, all tables will be created in the default tablespace on the underlying database on each node.

ExtenDB allows you to conveniently assign a logical tablespace name that groups together and corresponds to tablespaces on the individual Bizgres or PostgreSQL instances on the individual nodes. Ideally, nodes should be equal in processing capability and configuration, so you could specify the same file system path location for all nodes. There is also the capability to specify different locations on different nodes, if desired.

Once a tablespace has been defined, it can be referenced when creating tables or indexes.

A table can be moved to a separate tablespace by using the ALTER TABLE SET TABLESPACE command.

**Note: the paths specified for the tablespaces must be empty directories that have already been created, and the Bizgres or PostgreSQL process must have permission to write to this directory.**

# ExtenDB Parallel Server SQL Reference

## 5.1.2 CREATE TABLE

Syntax:

```
CREATE [TEMP] TABLE table_name (create_definition, ...)
    [partitioning_options]
    [INHERITS (table_name)]
    [WITH XROWID | WITHOUT XROWID]
    [TABLESPACE tablespace]

create_definition:
    column_name data_type [NOT NULL | NULL] [DEFAULT default_value]
        [PRIMARY KEY]
    or
    constraint_definition

constraint_definition: [CONSTRAINT constraint_name]
    PRIMARY KEY (index_column_name,...)
    or
    FOREIGN KEY [CONSTRAINT symbol] (index_column_name,...)
        [reference_definition]
    or
    CHECK (expression)

data_type: (one of the data types that appear in the Data Type section)

:
reference_definition:
    REFERENCES table_name [(index_column_name,...)]

partitioning_options:
    [PARTITIONING KEY column_name] ON ALL
    | [PARTITIONING KEY column_name] ON [NODE[S]] node_num[,node_num ...]
    | REPLICATED
    | ROUND ROBIN ON (ALL | NODES node_num[,node_num ...])
```

The CREATE TABLE statement is used to create tables in databases. It can be executed only by a user of type DBA or RESOURCE.

Table names may not start with the temporary table character sequence defined in `xdb.tempTablePrefix` in the `xdb.config` file.

### **Tablespaces**

The optional TABLESPACE clause allows a tablespace to be specified, to determine the location of the table data. See the CREATE TABLESPACE command for more details.

### **Internode Partitioning**

The partitioning\_options allow the DBA to specify a partitioning strategy, which is very important to the performance of the system. **Additional information appears in the ExtenDB Planning Guide, and should be read carefully.**

# ExtenDB Parallel Server SQL Reference

Partitioning allows the DBA to distribute the data amongst multiple nodes, either based on round robin partitioning, or a partitioning column. The value in this column is used to calculate a hash value, which is mapped to a destination node. (ExtenDB does not support range partitioning or round-robin, just hash partitioning.)

This distribution of data allows the ExtenDB Parallel Server (EPS) to parallelize queries. In choosing a partitioning key, it is important to take into consideration what other tables this table will likely join with. That way, if these other tables are partitioned on the corresponding join column, the ExtenDB Optimizer will recognize that local joins can occur without having to resort to any row shipping.

It is recommended to use all available nodes when choosing which nodes to use for the partitioned table.

Normally, one would probably want to select a column to designate as the partitioning column. Designating a partitioning column allows local joins to occur for other tables that are similarly partitioned, in cases the tables have a parent-child relationship. ROUND ROBIN partitioning may be useful, too, however, in such cases where a table typically does not join with any other tables (or just replicated ones), and there is no natural column to select as a partitioning column.

Another important partitioning\_option is REPLICATED. This is appropriate for "lookup" tables, such as a state code table. Replicated tables appear on all nodes, with each node containing the exact same data. This also allows joins to occur on all nodes without having to ship any data. Depending on your database schema and queries, a DBA may even consider making other larger tables replicated, but caution is urged.

A table may also simply appear on a single node or subset of the available nodes, via the ON NODE clause, but it is recommended to use all nodes.

**\*\*Note that if no partitioning table options are specified, the table is partitioned on the on the first element of the primary key across all nodes. If none such key exists, the first column in the table is chosen as the partitioning key automatically.**

Examples:

```
CREATE TABLE part ( p_partkey    INTEGER NOT NULL,
                    p_name        VARCHAR(55) NOT NULL,
                    p_mfgr        CHAR(25) NOT NULL,
                    p_brand       CHAR(10) NOT NULL,
                    p_type        VARCHAR(25) NOT NULL,
                    p_size        INTEGER NOT NULL,
                    p_container   CHAR(10) NOT NULL,
                    p_retailprice DECIMAL(15,2) NOT NULL,
                    p_comment     VARCHAR(23) NOT NULL )
PARTITIONING KEY p_partkey ON ALL;

CREATE TABLE nation ( n_nationkey INTEGER NOT NULL,
```

# ExtenDB Parallel Server SQL Reference

```

n_name          CHAR(25) NOT NULL,
n_regionkey     INTEGER NOT NULL,
n_comment       VARCHAR(152) )

REPLICATED;
```

## **Constraint Exclusion Partitioning**

### PostgreSQL and Bizgres Only

PostgreSQL 8.1 and Bizgres 0.7 introduce the notion of “partitioning” within a database instance via check constraints. This allows the DBA to create segments for a table that contain ranges of values, for example. A table named orders could be partitioned into monthly subtables, allowing queries that include a condition based on order date to scan with a smaller set of data, and therefore have a faster query time.

This is a powerful feature that should be taken advantage of. Constraint exclusion partitioning coupled with ExtenDB’s partitioning across multiple nodes will result in significantly faster query response times; a large table can be broken into multiple subtables, each of which is partitioned across multiple nodes in the cluster.

Note that the postgresql.conf parameter `constraint_exclusion` is off by default and must be set to on in order to take advantage of this feature.

An example appears below.

```
CREATE TABLE orders (
  o_orderkey      INTEGER NOT NULL,
  o_custkey       INTEGER NOT NULL,
  o_orderstatus   CHAR(1) NOT NULL,
  o_totalprice    DECIMAL(15,2) NOT NULL,
  o_orderdate     DATE NOT NULL,
  o_orderpriority CHAR(15) NOT NULL,
  o_clerk         CHAR(15) NOT NULL,
  o_shippriority  INTEGER NOT NULL,
  o_comment       VARCHAR(79) NOT NULL)
PARTITIONING KEY o_orderkey ON ALL;

CREATE TABLE orders_199201
( CHECK (o_orderdate BETWEEN '19920101'::DATE AND '19920131'::DATE) )
INHERITS (orders);

CREATE TABLE orders_199202
( CHECK (o_orderdate BETWEEN '19920201'::DATE AND '19920228'::DATE) )
INHERITS (orders);
:
```

A query like “SELECT o\_orderdate, count(\*) from orders where o\_orderdate between ‘1992-01-01’ and ‘1992-01-15’ GROUP BY o\_orderdate” will only use tuples found from the orders\_199201 subtable (and the orders table, which should be left empty).

**Note that when loading data, you must insert data into the proper subtable. Using the above example, in the current implementation, loading into orders directly will not automatically just insert the data into the correct subtable.**

# ExtenDB Parallel Server SQL Reference

Another important consideration when creating subtables is to be aware that the current PostgreSQL implementation is a bit datatype sensitive, and you might find that the PostgreSQL executor is not taking full advantage of eliminating subtables.

In the case of dates, we recommend using the above syntax to cast the date type, as in

```
CHECK (o_orderdate BETWEEN '19920101'::DATE AND '19920131'::DATE)
```

Leaving it as either just a date, or as a quoted string may cause queries in PostgreSQL to not be executed optimally. This depends on how the CHECK constraints are formulated and how the WHERE conditions are formulated. The above check constraint syntax appears to handle various date constructs (quoted, cast) in SELECT WHERE clauses properly.

## **Temporary Tables**

Temporary tables (or temp tables) may be created using the CREATE TEMP TABLE command. A temporary table is accessible only for the session that created it. After the session ends, the temp table will be dropped automatically.

When a temporary table is created in ExtenDB, one or more temp tables will also be created on the underlying database. The xdb.config configuration file has a few of parameters that need to be set correctly for this to work properly, the defaults of which are set for PostgreSQL.

xdb.config parameter	Default	Description
xdb.tempTablePrefix	TMP	Temporary table prefix to use in underlying database. Various databases have different conventions, like "TEMP." or "#".
xdb.sqlcommand.createTempTable.start	CREATE TABLE	Start of command for CREATE TABLE statement for creating temp table on the underlying. Allows for alternate syntaxes like CREATE TEMP TABLE. For PostgreSQL, real (not temporary) tables are used, to take advantage of the COPY command when loading data.
xdb.sqlcommand.createTempTable.suffix		Suffix to add at the end of CREATE statements for temp tables. <b>This can be used to allow disabling of logging information on the underlying database and greatly improve performance, since temporary tables are used internally by ExtenDB.</b> For example, in SAP-DB, this value should be set to <b>IGNORE ROLLBACK</b> .

Note that for PostgreSQL, ExtenDB currently uses non-temporary tables in the temporary table implementation by default.

## **Column Definitions**

Columns may be of any of the data types listed in the data type chapter. In addition, the user may specify whether or not nulls are allowed with NULL or NOT NULL, and

# ExtenDB Parallel Server SQL Reference

include a DEFAULT clause to specify any default value for the column if none is specified as part of an insert. A column may also be designated as a primary key, or reference a foreign key.

## ***Constraint Definitions***

The CREATE TABLE statement may also include primary key or foreign key constraint definitions. Alternatively, these can also be specified as separate statements as part of ALTER TABLE. Please see ALTER TABLE for a more detailed discussion about constraints.

## ***Xrowid***

In previous versions of ExtenDB, an internal row identifier was always created for each row to aid with distributed constraint checking. This is no longer required if you have a primary key or unique index on a table. If you do not use distributed constraints (e.g., a tuple in one table references a tuple in another), you do not need to worry about a unique key for each row. For compatibility, one can still create XROWID by including the WITH XROWID clause. The implicit default is WITHOUT XROWID.

# ExtenDB Parallel Server SQL Reference

## 5.1.3 DROP TABLE

Syntax:

```
DROP TABLE table_name
```

The DROP TABLE command is used to drop a table from the database.

If any foreign key constraints exist and other tables reference this one, the DROP will not succeed. The foreign key constraint from the other table must be dropped first.

## 5.1.4 TRUNCATE TABLE

Syntax:

```
TRUNCATE table_name
```

The TRUNCATE command effectively deletes all the rows in a table. It is much faster than DELETE without any WHERE condition since it will not scan the entire table.

## 5.1.5 ALTER TABLE ... ADD COLUMN

Syntax:

```
ALTER TABLE table_name  
    ADD [COLUMN] create_definition
```

This command is used to add columns to an existing table.

## 5.1.6 ALTER TABLE ... DROP COLUMN

Syntax:

```
ALTER TABLE table_name  
    DROP COLUMN column_name
```

This command is used to drop a column from a table in the database.

# ExtenDB Parallel Server SQL Reference

## 5.1.7 ALTER TABLE ... ADD PRIMARY KEY

Syntax:

```
ALTER TABLE table_name
    ADD PRIMARY KEY (column_name [,column_name...])
```

A primary key may be added either as part of the CREATE TABLE statement, or as part of ALTER TABLE.

The ExtenDB Parallel Server will try and have constraints be enforced locally on the nodes if possible. That means, if the first column in the primary key is also the partitioning column for a table, ExtenDB will not enforce the primary key itself, and will allow the underlying database to do it, since we know that no two rows with the same key can appear on different nodes. In addition, if the table is a lookup table and replicated to all nodes, or it just appears on a single node, ExtenDB will also leave it to the underlying database to be enforced.

If, however, the table is partitioned, and the first column in the primary key is not the partitioning column, ExtenDB will enforce the primary key in case of INSERTs or UPDATES. **Distributed constraints like this are expensive to enforce, so if your database does a considerable amount of such write operations, you may not want to create any distributed constraints.** If just doing periodic loads, it should not be an issue.

If a unique index already exists on the primary key columns, it will be used. Otherwise, the command is issued on the underlying nodes.

## 5.1.8 ALTER TABLE ... ADD FOREIGN KEY

Syntax:

```
ALTER TABLE table_name
    ADD [CONSTRAINT constraint_name]
    FOREIGN KEY (column_name [, column_name...])
    REFERENCES table_name (column_name [, column_name...])
```

Foreign key constraints may be added by either the CREATE TABLE or ALTER TABLE command.

Foreign keys help to guarantee referential integrity in your database. The referring table's corresponding values must exist as a primary key or unique index in the referenced table.

The ExtenDB Parallel Server will try and have constraints be enforced locally on the nodes if possible. If we are referencing a replicated lookup table for example,



# ExtenDB Parallel Server SQL Reference

ExtenDB will leave it to the underlying nodes to enforce, and will create a foreign key constraint on each node locally.

In addition, if the table is partitioned and we are referencing another partitioned table, if the first column in the referenced and referring key in each of the tables is also the partitioning key for each, we will also rely on the underlying databases to enforce the constraint.

In other cases, local enforcement is not possible. In those cases, ExtenDB will provide enforcement of foreign keys for INSERT, UPDATE and DELETE operations, as a distributed constraint.

Distributed constraints like this are expensive to enforce, so if your database does a considerable amount of such write operations, you may not want to create any distributed constraints. If just doing periodic loads, it should not be an issue.

## 5.1.9 ALTER TABLE ... DROP CONSTRAINT

Syntax:

```
ALTER TABLE table_name
    DROP CONSTRAINT constraint_name
```

This command is used to drop an existing primary key or foreign key constraint. If the constraint had been created after an index on which it was based, only the constraint definition is dropped, and not the previously created index. If an internal index needed to be created for the constraint, it will be dropped as well, however.

If no constraint name was specified at the time of its creation, the internally generated name may be determined by issuing a SHOW CONSTRAINTS command.

## 5.1.10 ALTER TABLE ... DROP PRIMARY KEY

Syntax:

```
ALTER TABLE table_name
    DROP PRIMARY KEY
```

The DROP PRIMARY KEY clause of the ALTER TABLE command is used to remove the primary key for a table. If the primary key constraint was created after a unique index on which it was based, the index will not be dropped. If an internally generated index was created, it will be dropped along with the primary key.

# ExtenDB Parallel Server SQL Reference

## 5.1.11 ALTER TABLE ... ALTER COLUMN

Syntax:

```
ALTER [ COLUMN ] column TYPE type [ USING expression ]
ALTER [ COLUMN ] column SET DEFAULT expression
ALTER [ COLUMN ] column DROP DEFAULT
ALTER [ COLUMN ] column { SET | DROP } NOT NULL
```

This command is used to modify the definition of a column, changing its type, its default, and whether or not it can be null. These can also be combined into a single ALTER command.

## 5.1.12 ALTER TABLE ... OWNER

Syntax:

```
ALTER TABLE table OWNER TO new_owner
```

This variant of ALTER TABLE is used to modify the owner of the table.

## 5.1.13 ALTER TABLE ... SET TABLESPACE

**Bizgres and PostgreSQL Only**

Syntax:

```
ALTER TABLE table_name SET TABLESPACE tablespace
```

This variant of ALTER TABLE is used to move a table from one tablespace to another. More information about tablespaces can be found in the topic of CREATE TABLESPACE.

## 5.1.14 RENAME TABLE

Syntax:

```
RENAME TABLE table_name TO new_table_name
```

This command allows the DBA or owner of a table to rename a table.

# ExtenDB Parallel Server SQL Reference

## 5.1.15 CREATE INDEX

Syntax:

```
CREATE [UNIQUE] INDEX index_name
ON table_name [USING index_type]
(column_name [[ASC|DESC|operator_class],... ])
[TABLESPACE tablespace]
[WHERE predicate]
```

This command is used to create indexes for tables on the specified columns. The optional TABLESPACE clause allows a tablespace to be specified.

If the UNIQUE modifier is specified, the set of values in the columns must be unique. If the table is partitioned and the first column in the index is the partitioning column, ExtenDB will rely on the underlying nodes to enforce the unique index. If that is not the case, ExtenDB will enforce the unique index itself. **Note that enforcing a unique distributed index will significantly slow down INSERT and UPDATE operations.**

The USING clause allows for other types of indexes to be used, like Bitmap indexes, if the underlying database supports it. Bitmap indexes are not yet part of PostgreSQL 8.1, but can be added via a patch or by using the Bizgres distribution of PostgreSQL.

The WHERE predicate allows for the support of partial indexes.

## 5.1.16 DROP INDEX

Syntax:

```
DROP INDEX index_name [ON table_name]
```

This command is used to drop indexes on tables. It can be executed by the DBA or owner of the table.

## 5.1.17 CLUSTER

Syntax:

```
CLUSTER index_name ON table_name
CLUSTER table_name
CLUSTER
```

Clustering a table on an existing index causes the table to be reordered based on the specified index, to allow for faster access.

# ExtenDB Parallel Server SQL Reference

Note that when new rows are added to the table they will not be in clustered order. Issuing a CLUSTER command with just the table\_name parameter will recluster the table on the specified index. Issuing a CLUSTER command without any parameters will recluster all clustered tables.

As a practical matter, if your tables get quite large, and this command may execute for a long time. Using CLUSTER may just be practical if you have a loading strategy where you create subtables (see CREATE TABLE) based on a window of time of data, and where this data is fairly static. Then, you can CLUSTER the particular subtable just once and leave it alone after that.

## 5.1.18 CREATE VIEW

Syntax:

```
CREATE VIEW view_name [(column1[,column...])] AS select_statement
```

This statement allows views to be defined from one or more tables as SELECT statements.

ExtenDB views are not updatable; one cannot update data in a view, or insert data into a view.

Views are not materialized, and are evaluated as part of a query at execution time.

## 5.1.19 DROP VIEW

Syntax:

```
DROP VIEW view_name
```

This statement is used to drop a view.

## 5.2 Data Manipulation Statements

### 5.2.1 INSERT

Syntax:

```
INSERT [INTO] table_name [(column_name,...)]  
VALUES ((expression),...)  
  
INSERT [INTO] table_name [(column_name,...)]  
SELECT ...
```

The INSERT command is used to insert data into the table. There are two forms of the command, one for single row inserts, and one for multiple row inserts by taking the results of a SELECT statement.

If the table is replicated, inserted rows will be inserted into the replicated table on all nodes.

If the table is on a single node, the rows will just be inserted into the table on that node.

If the table is partitioned, a hash value is calculated based on the partitioning column of the table, and the row to be inserted will be inserted the appropriate target node.

For the INSERT to succeed, no unique index, primary key or referenced foreign key constraint violations may occur.

### 5.2.2 UPDATE

Syntax:

```
UPDATE table_name  
SET [table_name.]column_name1=expression1  
[, [table_name.]column_name2=expression2 ...]  
[WHERE where_definition]
```

The UPDATE command is used to update data within a table.

If a table is replicated, the table will be updated the same on all of the nodes.

# ExtenDB Parallel Server SQL Reference

If the table is on a single node, the rows will just be update in the table on that node.

For the UPDATE to succeed, no unique index, primary key or referenced foreign key constraint violations may occur.

If any distributed foreign key or primary key constraints exist, the execution of UPDATE may slow down considerably while these are checked cross-node. If the UPDATE affects the partitioning column of the table, the updated row may physically move from one node to another on the underlying database.

## 5.2.3 DELETE

Syntax:

```
DELETE FROM table_name
    [WHERE where_definition]
```

The DELETE command is used to delete data within a table.

If a table is replicated, the rows from the table will be deleted in the same manner on all of the nodes.

If the table is on a single node, the rows will just be deleted in the table on that node.

For the DELETE to succeed, no foreign key constraint references to this table may be violated.

If any distributed foreign key references exist, the execution of DELETE may slow down considerably while these are checked cross-node.

# ExtenDB Parallel Server SQL Reference

## 5.2.4 SELECT

Syntax:

```
SELECT
    [DISTINCT | UNIQUE | ALL]
    select_expression,...
    [FROM table_references
    [WHERE condition]
    [GROUP BY {unsigned_integer | column_name | expression}...]
    [HAVING where_definition]
    [UNION select.....]
    [ORDER BY {unsigned_integer | column_name | expression} [ASC |
DESC] ,...]]
    [LIMIT n] [OFFSET m]
```

table\_references may be

```
:
table_reference, table_reference
table_reference [CROSS] JOIN table_reference
table_reference [INNER] JOIN table_reference join_condition
table_reference NATURAL JOIN table_reference ON (column[,column_name...])
table_reference LEFT [OUTER] JOIN table_reference join_condition
```

Where table\_reference is defined as:

```
table_name [[AS] alias]
and join_condition is defined as:
```

```
ON conditional_expr | USING (column_list)
```

When using USING, column\_list must be column names that exist in both tables.

WHERE condition\_expr:

```
[NOT] condition_expr [ {AND | OR} condition_expr]
```

condition\_expr:

```
expression operator [ANY|ALL] expression
expression [NOT] BETWEEN expression AND expression
expression [NOT] {IN|EXISTS} ({select_statement|expression_list})
[{{table|alias}.]column IS [NOT] NULL
[{{table|alias}.]column [NOT] {LIKE|MATCHES} string
```

An expression itself can be a string literal, mathematical expression, the result of a SELECT, etc. More information on expressions and operators can be found in the Expressions chapter earlier in this document.

# ExtenDB Parallel Server SQL Reference

## 5.3 Users and Privileges

This section contains the SQL commands for creating and manipulating users and granting and revoking privileges.

A valid user must be specified when connecting to the ExtenDB database. When a database is first created with createdb, the user is required to assign a DBA username and password.

A user must be granted access to a table before being able to access it. By default, a user who creates a table has all privileges on that table.

There are 3 classes of users: DBA, RESOURCE, and STANDARD. DBA users have Database Administration privileges. RESOURCE users can create tables. STANDARD users cannot create tables, but can access the database.

### 5.3.1 CREATE USER

Syntax:

```
CREATE USER username PASSWORD password [user_class]
```

```
user_class = [DBA|RESOURCE|STANDARD]
```

By default, user\_class is RESOURCE.

This command may only be executed by a DBA user.

STANDARD users may only access tables for which they have been granted permission (or via PUBLIC, of course). They may also create temp tables.

RESOURCE users are the same as STANDARD, only they also have the ability to create tables.

DBA users have all of the permissions of a RESOURCE user, but can additionally perform other tasks: starting and stopping a database, and executing ANALYZE, CREATE USER, DROP USER, and ALTER USER.



# ExtenDB Parallel Server SQL Reference

## 5.3.2 ALTER USER

Syntax:

```
ALTER USER username PASSWORD password
ALTER USER username user_class
ALTER USER username PASSWORD password user_class
```

This command may only be executed by a DBA user, or by the same user wanting to change his or her own password.

A user cannot be demoted to the user class STANDARD if they own one or more tables. In such a case, this will fail. The DBA should instead first either drop the tables or reassign ownership via ALTER TABLE using the OWNER clause.

## 5.3.3 DROP USER

Syntax:

```
DROP USER username
```

This command may only be executed by a DBA user, and removes the user from the database.

## 5.3.4 GRANT

Syntax:

```
GRANT privilege_list ON [TABLE] table_list TO grantee_list

privilege_list ::= ALL | privilege[,privilege...]
privilege ::= SELECT | INSERT | UPDATE | DELETE |
            REFERENCES | INDEX | ALTER

table_list ::= * | table[,table...]

grantee_list ::= grantee[,grantee]
grantee ::= PUBLIC | user_name
```

The GRANT command is used to grant privileges on a table, and may only be executed by a DBA user, or the owner of a table.

The privileges for SELECT, INSERT, UPDATE and DELETE are for those respective commands.

# ExtenDB Parallel Server SQL Reference

The REFERENCES privilege refers to the ability to create foreign key constraints, which requires REFERENCES privilege in both the referrer and referee tables.

The INDEX privilege allows the user to create and drop indexes.

The ALTER privilege allows the user to execute the ALTER table command.

## 5.3.5 REVOKE

Syntax:

```
REVOKE privilege_list ON [TABLE] table_list FROM grantee_list
```

```
privilege_list ::= ALL | privilege[,privilege...]  
privilege ::= SELECT | INSERT | UPDATE | DELETE |  
REFERENCES | INDEX | ALTER
```

```
table_list ::= * | *.* | table[,table...]
```

```
grantee_list ::= grantee[,grantee]  
grantee ::= PUBLIC | user_name
```

REVOKE is used to revoke privileges on a table.

Note that we only revoke what has previously been granted. If a table has been granted permission to PUBLIC, revoking for user1 will still allow user1 to access the table because the table is still accessible to the PUBLIC.

# ExtenDB Parallel Server SQL Reference

## 5.4 Other Commands

ExtenDB also includes some administrative commands for administering the database and finding out information.

Command	Description
SHOW DATABASES	Lists all of the user-created ExtenDB databases
SHOW TABLES	Lists all of the tables that exist in the current database
SHOW VIEWS	Lists all of the views that exist in the current database
DESCRIBE <table>	Lists the columns and their definitions of the specified table
DESCRIBE <view>	Displays the view definition for the specified view
SHOW INDEXES ON <table>	Lists all indexes for <table>
SHOW CONSTRAINTS ON <table>	Lists the following types of constraints for <table>: primary keys, foreign keys, foreign key references
SHOW USERS	Lists all defined users and their class
SHOW STATEMENTS	Lists all of the currently executing SQL statements
KILL <request_id>	Kills execution of the request id specified. Request ids can be obtained by executing the SHOW STATEMENTS command
UPDATE STATISTICS	Updates the internal statistics in the database for creating better execution plans
ANALYZE	Updates the internal statistics in the database for creating better execution plans
VACUUM	For PostgreSQL and Bizgres, potentially frees up space in the database

### 5.4.1 SHOW DATABASES

Lists all of the user-created ExtenDB databases. Columns:

Column	Description
Database	Database name
Status	The status of the database, can be either Started or Down.
Nodes	A comma separated list of all of the node id numbers that the database is on.

# ExtenDB Parallel Server SQL Reference

## 5.4.2 SHOW TABLES

Lists the columns and their definitions of the specified table.

Column	Description
Table_name	The name of the table
Table_partitioning_column	If the table is partitioned, the column used in partitioning, otherwise it is null for single-node or replicated tables.
Table_nodes	A comma separated list of all of the node id numbers that the table is on.

## 5.4.3 SHOW VIEWS

Lists the views defined in the database.

Column	Description
View_name	The name of the view

## 5.4.4 DESCRIBE <table>

Lists the columns and their definitions of the specified table.

Column	Description
Column_name	The name of the column
Sql_data_type	Data type of column
Type_name	The data type name
Is_nullable	If YES, allows nulls, otherwise NO
Key	If when the table was defined, this single column was denoted as the primary key
Default	The default value of the column

## 5.4.5 DESCRIBE <view>

Displays the view definition and any named view columns.

Column	Description
View_text	The text defining the view
View_column	The comma separated list of the column names defined by the creator of the view, if any.

# ExtenDB Parallel Server SQL Reference

## 5.4.6 SHOW INDEXES ON <table>

Lists all indexes for specified table.

Column	Description
Indexname	The name of the index
Isunique	If YES, the index only allows unique values.
Columns	A comma-separated list of the columns that make up the index.

## 5.4.7 SHOW CONSTRAINTS ON <table>

Lists all constraints for specified table, including primary keys, foreign keys, foreign key references.

Column	Description
Constname	Constraint name
Type	The constraint type
Sourcetable	If a foreign key constraint, the referencing table
Sourcecolumns	If a foreign key constraint, the referencing columns
Desttable	If a foreign key constraint, the referenced table.
Destcolumns	If a foreign key constraint, the referenced columns.

## 5.4.8 SHOW USERS

Lists all defined users.

Column	Description
User_name	User name
User_class	Class of user

# ExtenDB Parallel Server SQL Reference

## 5.4.9 SHOW STATEMENTS

Shows all of the currently executing statements.

Column	Description
Request_id	The id of the request
Session_id	The session id number
Submit_time	The time at which the request was submitted
Status	P=processing, Q=queued
Statement	The actual request submitted
Nodes	For future use
Current_step	For future use

## 5.4.10 KILL

Syntax:

```
KILL statement_number
```

Kill is used to abort a currently executing or queued request. Kill is followed by the request id number, which can be determined from the SHOW STATEMENTS command.

## 5.4.11 UPDATE STATISTICS

Syntax:

```
UPDATE STAT[ISTICS] table_name
UPDATE STAT[ISTICS] *
UPDATE STAT[ISTICS] COLUMN (*) FOR table name
UPDATE STAT[ISTICS] COLUMN (column_name, ...) FOR table name
```

UPDATE STATISTICS is used to update the internal statistics in the database that are used in optimizing queries. Executing this command will result in a similar command being executed on the underlying database so that individual nodes have reasonable statistics, in addition to updating ExtenDB's statistics.

It is critical in the performance of the database that UPDATE STATISTICS (or ANALYZE) is run periodically.

UPDATE STATISTICS is a command that may need to be mapped for the underlying database. The two most common usages are UPDATE STATISTICS and ANALYZE.

# ExtenDB Parallel Server SQL Reference

The DBA must configure this in xdb.config file. The current defaults are for PostgreSQL. Please see Administration Guide for more details.

See also: ANALYZE, VACUUM ANALYZE

## 5.4.12 ANALYZE

Syntax:

```
ANALYZE [ table [ (column [, ...] ) ] ]
```

The ANALYZE command is a synonym for UPDATE STATISTICS. Both commands are supported to make administration easier for DBAs accustomed to one or the other command from other DBMS systems.

ANALYZE updates internal statistics within node databases level as well as within ExtenDB's metadata database, to allow better query plans to be created when processing queries.

ANALYZE by itself will analyze all tables. If a table name is included, it will analyze just the particular table. For finer granularity, individual column names may be specified as well.

See also: UPDATE STATISTICS, VACUUM.

## 5.4.13 VACUUM

### PostgreSQL and Bizgres Only

Syntax:

```
VACUUM [ FULL | FREEZE ] [ANALYZE] [ table [ (column [, ...] ) ] ]
```

The VACUUM command is used to clean up the database internally, freeing up space due to deleted and updated rows. It is also a good idea to run it periodically due to PostgreSQL transaction id wrap-around (more information can be found in the PostgreSQL documentation).

It is a good idea to run the variant VACUUM ANALYZE, achieving both the results of VACUUM and ANALYZE at the same time.

VACUUM can optionally accept a single table name, as well as a subset of columns to update.

# ExtenDB Parallel Server SQL Reference

The FULL parameter does a more thorough VACUUM, which may free up more space, but it takes longer to run.

FREEZE is used to mark tuples so that they will not be subject to the transaction id wrap-around issue. It is recommended to read more about FREEZE in the PostgreSQL documentation before using.

See also: ANALYZE