



[Home](#) → [Documentation](#) → [Manuals](#) → [PostgreSQL 9.1](#)

This page in other versions: [9.2](#) / [9.3](#) / [9.4](#) / [9.5](#) / [current \(9.6\)](#) | Development versions: [devel](#) / [10](#) |

Unsupported versions: [7.1](#) / [7.2](#) / [7.3](#) / [7.4](#) / [8.0](#) / [8.1](#) / [8.2](#) / [8.3](#) / [8.4](#) / [9.0](#) / [9.1](#)

[PostgreSQL 9.1.24 Documentation](#)

[Prev](#)

[Up](#)

[Next](#)

CREATE TABLE

Name

CREATE TABLE -- define a new table

Synopsis

```
CREATE [ [ GLOBAL | LOCAL ] { TEMPORARY | TEMP } | UNLOGGED ] TABLE [ IF NOT EXISTS
] table_name ( [
    { column_name data_type [ COLLATE collation ] [ column_constraint [ ... ] ]
      | table_constraint
      | LIKE parent_table [ like_option ... ] }
    [, ... ]
] )
[ INHERITS ( parent_table [, ... ] ) ]
[ WITH ( storage_parameter [= value] [, ... ] ) | WITH OIDS | WITHOUT OIDS ]
[ ON COMMIT { PRESERVE ROWS | DELETE ROWS | DROP } ]
[ TABLESPACE tablespace ]
```

```
CREATE [ [ GLOBAL | LOCAL ] { TEMPORARY | TEMP } | UNLOGGED ] TABLE [ IF NOT EXISTS
] table_name
    OF type_name [ (
    { column_name WITH OPTIONS [ column_constraint [ ... ] ]
      | table_constraint }
    [, ... ]
) ]
[ WITH ( storage_parameter [= value] [, ... ] ) | WITH OIDS | WITHOUT OIDS ]
[ ON COMMIT { PRESERVE ROWS | DELETE ROWS | DROP } ]
[ TABLESPACE tablespace ]
```

where column_constraint is:

```
[ CONSTRAINT constraint_name ]
{ NOT NULL |
  NULL |
  CHECK ( expression ) |
  DEFAULT default_expr |
  UNIQUE index_parameters |
```

```

PRIMARY KEY index_parameters |
REFERENCES reftable [ ( refcolumn ) ] [ MATCH FULL | MATCH PARTIAL | MATCH SIMPLE
]
[ ON DELETE action ] [ ON UPDATE action ] }
[ DEFERRABLE | NOT DEFERRABLE ] [ INITIALLY DEFERRED | INITIALLY IMMEDIATE ]

```

and table_constraint is:

```

[ CONSTRAINT constraint_name ]
{ CHECK ( expression ) |
  UNIQUE ( column_name [, ... ] ) index_parameters |
  PRIMARY KEY ( column_name [, ... ] ) index_parameters |
  EXCLUDE [ USING index_method ] ( exclude_element WITH operator [, ... ] )
index_parameters [ WHERE ( predicate ) ] |
  FOREIGN KEY ( column_name [, ... ] ) REFERENCES reftable [ ( refcolumn [, ... ] )
]
[ MATCH FULL | MATCH PARTIAL | MATCH SIMPLE ] [ ON DELETE action ] [ ON UPDATE
action ] }
[ DEFERRABLE | NOT DEFERRABLE ] [ INITIALLY DEFERRED | INITIALLY IMMEDIATE ]

```

and like_option is:

```

{ INCLUDING | EXCLUDING } { DEFAULTS | CONSTRAINTS | INDEXES | STORAGE | COMMENTS |
ALL }

```

index_parameters in UNIQUE, PRIMARY KEY, and EXCLUDE constraints are:

```

[ WITH ( storage_parameter [= value] [, ... ] ) ]
[ USING INDEX TABLESPACE tablespace ]

```

exclude_element in an EXCLUDE constraint is:

```

{ column | ( expression ) } [ opclass ] [ ASC | DESC ] [ NULLS { FIRST | LAST } ]

```

Description

CREATE TABLE will create a new, initially empty table in the current database. The table will be owned by the user issuing the command.

If a schema name is given (for example, **CREATE TABLE myschema.mytable ...**) then the table is created in the specified schema. Otherwise it is created in the current schema. Temporary tables exist in a special schema, so a schema name cannot be given when creating a temporary table. The name of the table must be distinct from the name of any other table, sequence, index, view, or foreign table in the same schema.

CREATE TABLE also automatically creates a data type that represents the composite type corresponding to one row of the table. Therefore, tables cannot have the same name as any existing data type in the same schema.

The optional constraint clauses specify constraints (tests) that new or updated rows must satisfy for an insert or update operation to succeed. A constraint is an SQL object that helps define the set of valid values in the table in various ways.

There are two ways to define constraints: table constraints and column constraints. A column constraint is defined as part of a column definition. A table constraint definition is not tied to a particular column, and it can encompass more than one column. Every column constraint can also be written as a table constraint; a column constraint is only a notational convenience for use when the constraint only affects one column.

Parameters

TEMPORARY or TEMP

If specified, the table is created as a temporary table. Temporary tables are automatically dropped at the end of a session, or optionally at the end of the current transaction (see `ON COMMIT` below). Existing permanent tables with the same name are not visible to the current session while the temporary table exists, unless they are referenced with schema-qualified names. Any indexes created on a temporary table are automatically temporary as well.

The [autovacuum daemon](#) cannot access and therefore cannot vacuum or analyze temporary tables. For this reason, appropriate vacuum and analyze operations should be performed via session SQL commands. For example, if a temporary table is going to be used in complex queries, it is wise to run `ANALYZE` on the temporary table after it is populated.

Optionally, `GLOBAL` or `LOCAL` can be written before `TEMPORARY` or `TEMP`. This makes no difference in PostgreSQL, but see [Compatibility](#).

UNLOGGED

If specified, the table is created as an unlogged table. Data written to unlogged tables is not written to the write-ahead log (see [Chapter 29](#)), which makes them considerably faster than ordinary tables. However, they are not crash-safe: an unlogged table is automatically truncated after a crash or unclean shutdown. The contents of an unlogged table are also not replicated to standby servers. Any indexes created on an unlogged table are automatically unlogged as well; however, unlogged [GiST indexes](#) are currently not supported and cannot be created on an unlogged table.

IF NOT EXISTS

Do not throw an error if a relation with the same name already exists. A notice is issued in this case. Note that there is no guarantee that the existing relation is anything like the one that would have been created.

table_name

The name (optionally schema-qualified) of the table to be created.

OF type_name

Creates a *typed table*, which takes its structure from the specified composite type (name optionally schema-qualified). A typed table is tied to its type; for example the table will be dropped if the type is dropped (with `DROP TYPE . . . CASCADE`).

When a typed table is created, then the data types of the columns are determined by the underlying composite type and are not specified by the `CREATE TABLE` command. But the `CREATE TABLE` command can add defaults and constraints to the table and can specify storage parameters.

`column_name`

The name of a column to be created in the new table.

`data_type`

The data type of the column. This can include array specifiers. For more information on the data types supported by PostgreSQL, refer to [Chapter 8](#).

`COLLATE collation`

The `COLLATE` clause assigns a collation to the column (which must be of a collatable data type). If not specified, the column data type's default collation is used.

`INHERITS (parent_table [, ...])`

The optional `INHERITS` clause specifies a list of tables from which the new table automatically inherits all columns.

Use of `INHERITS` creates a persistent relationship between the new child table and its parent table(s). Schema modifications to the parent(s) normally propagate to children as well, and by default the data of the child table is included in scans of the parent(s).

If the same column name exists in more than one parent table, an error is reported unless the data types of the columns match in each of the parent tables. If there is no conflict, then the duplicate columns are merged to form a single column in the new table. If the column name list of the new table contains a column name that is also inherited, the data type must likewise match the inherited column(s), and the column definitions are merged into one. If the new table explicitly specifies a default value for the column, this default overrides any defaults from inherited declarations of the column. Otherwise, any parents that specify default values for the column must all specify the same default, or an error will be reported.

`CHECK` constraints are merged in essentially the same way as columns: if multiple parent tables and/or the new table definition contain identically-named `CHECK` constraints, these constraints must all have the same check expression, or an error will be reported. Constraints having the same name and expression will be merged into one copy. Notice that an unnamed `CHECK` constraint in the new table will never be merged, since a unique name will always be chosen for it.

Column **STORAGE** settings are also copied from parent tables.

LIKE parent_table [like_option ...]

The **LIKE** clause specifies a table from which the new table automatically copies all column names, their data types, and their not-null constraints.

Unlike **INHERITS**, the new table and original table are completely decoupled after creation is complete. Changes to the original table will not be applied to the new table, and it is not possible to include data of the new table in scans of the original table.

Default expressions for the copied column definitions will be copied only if **INCLUDING DEFAULTS** is specified. The default behavior is to exclude default expressions, resulting in the copied columns in the new table having null defaults. Note that copying defaults that call database-modification functions, such as `nextval`, may create a functional linkage between the original and new tables.

Not-null constraints are always copied to the new table. **CHECK** constraints will be copied only if **INCLUDING CONSTRAINTS** is specified. No distinction is made between column constraints and table constraints.

Indexes, **PRIMARY KEY**, **UNIQUE**, and **EXCLUDE** constraints on the original table will be created on the new table only if **INCLUDING INDEXES** is specified. Names for the new indexes and constraints are chosen according to the default rules, regardless of how the originals were named. (This behavior avoids possible duplicate-name failures for the new indexes.)

STORAGE settings for the copied column definitions will be copied only if **INCLUDING STORAGE** is specified. The default behavior is to exclude **STORAGE** settings, resulting in the copied columns in the new table having type-specific default settings. For more on **STORAGE** settings, see [Section 55.2](#).

Comments for the copied columns, constraints, and indexes will be copied only if **INCLUDING COMMENTS** is specified. The default behavior is to exclude comments, resulting in the copied columns and constraints in the new table having no comments.

INCLUDING ALL is an abbreviated form of **INCLUDING DEFAULTS INCLUDING CONSTRAINTS INCLUDING INDEXES INCLUDING STORAGE INCLUDING COMMENTS**.

Note that unlike **INHERITS**, columns and constraints copied by **LIKE** are not merged with similarly named columns and constraints. If the same name is specified explicitly or in another **LIKE** clause, an error is signaled.

CONSTRAINT constraint_name

An optional name for a column or table constraint. If the constraint is violated, the constraint name is present in error messages, so constraint names like `col must be positive` can be

used to communicate helpful constraint information to client applications. (Double-quotes are needed to specify constraint names that contain spaces.) If a constraint name is not specified, the system generates a name.

NOT NULL

The column is not allowed to contain null values.

NULL

The column is allowed to contain null values. This is the default.

This clause is only provided for compatibility with non-standard SQL databases. Its use is discouraged in new applications.

CHECK (expression)

The **CHECK** clause specifies an expression producing a Boolean result which new or updated rows must satisfy for an insert or update operation to succeed. Expressions evaluating to **TRUE** or **UNKNOWN** succeed. Should any row of an insert or update operation produce a **FALSE** result an error exception is raised and the insert or update does not alter the database. A check constraint specified as a column constraint should reference that column's value only, while an expression appearing in a table constraint can reference multiple columns.

Currently, **CHECK** expressions cannot contain subqueries nor refer to variables other than columns of the current row.

DEFAULT default_expr

The **DEFAULT** clause assigns a default data value for the column whose column definition it appears within. The value is any variable-free expression (subqueries and cross-references to other columns in the current table are not allowed). The data type of the default expression must match the data type of the column.

The default expression will be used in any insert operation that does not specify a value for the column. If there is no default for a column, then the default is null.

UNIQUE (column constraint)

UNIQUE (column_name [, ...]) (table constraint)

The **UNIQUE** constraint specifies that a group of one or more columns of a table can contain only unique values. The behavior of the unique table constraint is the same as that for column constraints, with the additional capability to span multiple columns.

For the purpose of a unique constraint, null values are not considered equal.

Each unique table constraint must name a set of columns that is different from the set of columns named by any other unique or primary key constraint defined for the table. (Otherwise it would just be the same constraint listed twice.)

PRIMARY KEY (column constraint)

PRIMARY KEY (column_name [, ...]) (table constraint)

The **PRIMARY KEY** constraint specifies that a column or columns of a table can contain only unique (non-duplicate), nonnull values. Only one primary key can be specified for a table, whether as a column constraint or a table constraint.

The primary key constraint should name a set of columns that is different from the set of columns named by any unique constraint defined for the same table. (Otherwise, the unique constraint is redundant and will be discarded.)

PRIMARY KEY enforces the same data constraints as a combination of **UNIQUE** and **NOT NULL**, but identifying a set of columns as the primary key also provides metadata about the design of the schema, since a primary key implies that other tables can rely on this set of columns as a unique identifier for rows.

EXCLUDE [**USING** index_method] (exclude_element **WITH** operator [, ...]) index_parameters [**WHERE** (predicate)]

The **EXCLUDE** clause defines an exclusion constraint, which guarantees that if any two rows are compared on the specified column(s) or expression(s) using the specified operator(s), not all of these comparisons will return **TRUE**. If all of the specified operators test for equality, this is equivalent to a **UNIQUE** constraint, although an ordinary unique constraint will be faster. However, exclusion constraints can specify constraints that are more general than simple equality. For example, you can specify a constraint that no two rows in the table contain overlapping circles (see [Section 8.8](#)) by using the **&&** operator.

Exclusion constraints are implemented using an index, so each specified operator must be associated with an appropriate operator class (see [Section 11.9](#)) for the index access method **index_method**. The operators are required to be commutative. Each **exclude_element** can optionally specify an operator class and/or ordering options; these are described fully under [CREATE INDEX](#).

The access method must support **amgettuple** (see [Chapter 52](#)); at present this means GIN cannot be used. Although it's allowed, there is little point in using B-tree or hash indexes with an exclusion constraint, because this does nothing that an ordinary unique constraint doesn't do better. So in practice the access method will always be GiST.

The **predicate** allows you to specify an exclusion constraint on a subset of the table; internally this creates a partial index. Note that parentheses are required around the predicate.

REFERENCES reftable [(refcolumn)] [**MATCH** matchtype] [**ON DELETE** action] [**ON UPDATE** action] (column constraint)

```
FOREIGN KEY ( column [, ... ] ) REFERENCES reftable [ ( refcolumn  
[, ... ] ) ] [ MATCH matchtype ] [ ON DELETE action ] [ ON UPDATE  
action ] (table constraint)
```

These clauses specify a foreign key constraint, which requires that a group of one or more columns of the new table must only contain values that match values in the referenced column(s) of some row of the referenced table. If `refcolumn` is omitted, the primary key of the `reftable` is used. The referenced columns must be the columns of a non-deferrable unique or primary key constraint in the referenced table. Note that foreign key constraints cannot be defined between temporary tables and permanent tables.

A value inserted into the referencing column(s) is matched against the values of the referenced table and referenced columns using the given match type. There are three match types: `MATCH FULL`, `MATCH PARTIAL`, and `MATCH SIMPLE`, which is also the default. `MATCH FULL` will not allow one column of a multicolumn foreign key to be null unless all foreign key columns are null. `MATCH SIMPLE` allows some foreign key columns to be null while other parts of the foreign key are not null. `MATCH PARTIAL` is not yet implemented.

In addition, when the data in the referenced columns is changed, certain actions are performed on the data in this table's columns. The `ON DELETE` clause specifies the action to perform when a referenced row in the referenced table is being deleted. Likewise, the `ON UPDATE` clause specifies the action to perform when a referenced column in the referenced table is being updated to a new value. If the row is updated, but the referenced column is not actually changed, no action is done. Referential actions other than the `NO ACTION` check cannot be deferred, even if the constraint is declared deferrable. There are the following possible actions for each clause:

`NO ACTION`

Produce an error indicating that the deletion or update would create a foreign key constraint violation. If the constraint is deferred, this error will be produced at constraint check time if there still exist any referencing rows. This is the default action.

`RESTRICT`

Produce an error indicating that the deletion or update would create a foreign key constraint violation. This is the same as `NO ACTION` except that the check is not deferrable.

`CASCADE`

Delete any rows referencing the deleted row, or update the value of the referencing column to the new value of the referenced column, respectively.

`SET NULL`

Set the referencing column(s) to null.

`SET DEFAULT`

Set the referencing column(s) to their default values.

If the referenced column(s) are changed frequently, it might be wise to add an index to the foreign key column so that referential actions associated with the foreign key column can be performed more efficiently.

DEFERRABLE
NOT DEFERRABLE

This controls whether the constraint can be deferred. A constraint that is not deferrable will be checked immediately after every command. Checking of constraints that are deferrable can be postponed until the end of the transaction (using the [SET CONSTRAINTS](#) command). NOT DEFERRABLE is the default. Currently, only UNIQUE, PRIMARY KEY, EXCLUDE, and REFERENCES (foreign key) constraints accept this clause. NOT NULL and CHECK constraints are not deferrable.

INITIALLY IMMEDIATE
INITIALLY DEFERRED

If a constraint is deferrable, this clause specifies the default time to check the constraint. If the constraint is INITIALLY IMMEDIATE, it is checked after each statement. This is the default. If the constraint is INITIALLY DEFERRED, it is checked only at the end of the transaction. The constraint check time can be altered with the [SET CONSTRAINTS](#) command.

WITH (storage_parameter [= value] [, ...])

This clause specifies optional storage parameters for a table or index; see [Storage Parameters](#) for more information. The WITH clause for a table can also include OIDS=TRUE (or just OIDS) to specify that rows of the new table should have OIDs (object identifiers) assigned to them, or OIDS=FALSE to specify that the rows should not have OIDs. If OIDS is not specified, the default setting depends upon the [default_with_oids](#) configuration parameter. (If the new table inherits from any tables that have OIDs, then OIDS=TRUE is forced even if the command says OIDS=FALSE.)

If OIDS=FALSE is specified or implied, the new table does not store OIDs and no OID will be assigned for a row inserted into it. This is generally considered worthwhile, since it will reduce OID consumption and thereby postpone the wraparound of the 32-bit OID counter. Once the counter wraps around, OIDs can no longer be assumed to be unique, which makes them considerably less useful. In addition, excluding OIDs from a table reduces the space required to store the table on disk by 4 bytes per row (on most machines), slightly improving performance.

To remove OIDs from a table after it has been created, use [ALTER TABLE](#).

WITH OIDS
WITHOUT OIDS

These are obsolescent syntaxes equivalent to `WITH (OIDS)` and `WITH (OIDS=FALSE)`, respectively. If you wish to give both an `OIDS` setting and storage parameters, you must use the `WITH (. . .)` syntax; see above.

ON COMMIT

The behavior of temporary tables at the end of a transaction block can be controlled using `ON COMMIT`. The three options are:

PRESERVE ROWS

No special action is taken at the ends of transactions. This is the default behavior.

DELETE ROWS

All rows in the temporary table will be deleted at the end of each transaction block. Essentially, an automatic [TRUNCATE](#) is done at each commit.

DROP

The temporary table will be dropped at the end of the current transaction block.

TABLESPACE tablespace

The `tablespace` is the name of the tablespace in which the new table is to be created. If not specified, [default tablespace](#) is consulted, or [temp tablespaces](#) if the table is temporary.

USING INDEX TABLESPACE tablespace

This clause allows selection of the tablespace in which the index associated with a `UNIQUE`, `PRIMARY KEY`, or `EXCLUDE` constraint will be created. If not specified, [default tablespace](#) is consulted, or [temp tablespaces](#) if the table is temporary.

Storage Parameters

The `WITH` clause can specify *storage parameters* for tables, and for indexes associated with a `UNIQUE`, `PRIMARY KEY`, or `EXCLUDE` constraint. Storage parameters for indexes are documented in [CREATE INDEX](#). The storage parameters currently available for tables are listed below. For each parameter, unless noted, there is an additional parameter with the same name prefixed with `toast .`, which can be used to control the behavior of the table's secondary TOAST table, if any (see [Section 55.2](#) for more information about TOAST). Note that the TOAST table inherits the `autovacuum_*` values from its parent table, if there are no `toast .autovacuum_*` settings set.

`fillfactor (integer)`

The fillfactor for a table is a percentage between 10 and 100. 100 (complete packing) is the default. When a smaller fillfactor is specified, INSERT operations pack table pages only to the indicated percentage; the remaining space on each page is reserved for updating rows on that page. This gives UPDATE a chance to place the updated copy of a row on the same page as the original, which is more efficient than placing it on a different page. For a table whose entries are never updated, complete packing is the best choice, but in heavily updated tables smaller fillfactors are appropriate. This parameter cannot be set for TOAST tables.

`autovacuum_enabled, toast.autovacuum_enabled` (boolean)

Enables or disables the autovacuum daemon on a particular table. If true, the autovacuum daemon will initiate a VACUUM operation on a particular table when the number of updated or deleted tuples exceeds `autovacuum_vacuum_threshold` plus `autovacuum_vacuum_scale_factor` times the number of live tuples currently estimated to be in the relation. Similarly, it will initiate an ANALYZE operation when the number of inserted, updated or deleted tuples exceeds `autovacuum_analyze_threshold` plus `autovacuum_analyze_scale_factor` times the number of live tuples currently estimated to be in the relation. If false, this table will not be autovacuumed, except to prevent transaction Id wraparound. See [Section 23.1.4](#) for more about wraparound prevention. Observe that this variable inherits its value from the [autovacuum](#) setting.

`autovacuum_vacuum_threshold, toast.autovacuum_vacuum_threshold`
(integer)

Minimum number of updated or deleted tuples before initiate a VACUUM operation on a particular table.

`autovacuum_vacuum_scale_factor, toast.autovacuum_vacuum_scale_factor`
(float4)

Multiplier for `reltuples` to add to `autovacuum_vacuum_threshold`.

`autovacuum_analyze_threshold` (integer)

Minimum number of inserted, updated, or deleted tuples before initiate an ANALYZE operation on a particular table.

`autovacuum_analyze_scale_factor` (float4)

Multiplier for `reltuples` to add to `autovacuum_analyze_threshold`.

`autovacuum_vacuum_cost_delay, toast.autovacuum_vacuum_cost_delay`
(integer)

Custom [autovacuum vacuum cost delay](#) parameter.

`autovacuum_vacuum_cost_limit, toast.autovacuum_vacuum_cost_limit`
(integer)

Custom [autovacuum_vacuum_cost_limit](#) parameter.

`autovacuum_freeze_min_age, toast.autovacuum_freeze_min_age` (integer)

Custom [vacuum_freeze_min_age](#) parameter. Note that autovacuum will ignore attempts to set a per-table `autovacuum_freeze_min_age` larger than the half system-wide [autovacuum_freeze_max_age](#) setting.

`autovacuum_freeze_max_age, toast.autovacuum_freeze_max_age` (integer)

Custom [autovacuum_freeze_max_age](#) parameter. Note that autovacuum will ignore attempts to set a per-table `autovacuum_freeze_max_age` larger than the system-wide setting (it can only be set smaller).

`autovacuum_freeze_table_age, toast.autovacuum_freeze_table_age`
(integer)

Custom [vacuum_freeze_table_age](#) parameter.

Notes

Using OIDs in new applications is not recommended: where possible, using a `SERIAL` or other sequence generator as the table's primary key is preferred. However, if your application does make use of OIDs to identify specific rows of a table, it is recommended to create a unique constraint on the `oid` column of that table, to ensure that OIDs in the table will indeed uniquely identify rows even after counter wraparound. Avoid assuming that OIDs are unique across tables; if you need a database-wide unique identifier, use the combination of `tableoid` and row OID for the purpose.

Tip: The use of `OIDS=FALSE` is not recommended for tables with no primary key, since without either an OID or a unique data key, it is difficult to identify specific rows.

PostgreSQL automatically creates an index for each unique constraint and primary key constraint to enforce uniqueness. Thus, it is not necessary to create an index explicitly for primary key columns. (See [CREATE INDEX](#) for more information.)

Unique constraints and primary keys are not inherited in the current implementation. This makes the combination of inheritance and unique constraints rather dysfunctional.

A table cannot have more than 1600 columns. (In practice, the effective limit is usually lower because of tuple-length constraints.)

Examples

Create table `films` and table `distributors`:

```
CREATE TABLE films (  
    code          char(5) CONSTRAINT firstkey PRIMARY KEY,  
    title         varchar(40) NOT NULL,  
    did           integer NOT NULL,  
    date_prod     date,  
    kind          varchar(10),  
    len           interval hour to minute  
);  
  
CREATE TABLE distributors (  
    did           integer PRIMARY KEY DEFAULT nextval('serial'),  
    name          varchar(40) NOT NULL CHECK (name <> '')  
);
```

Create a table with a 2-dimensional array:

```
CREATE TABLE array_int (  
    vector        int[][]  
);
```

Define a unique table constraint for the table `films`. Unique table constraints can be defined on one or more columns of the table:

```
CREATE TABLE films (  
    code          char(5),  
    title         varchar(40),  
    did           integer,  
    date_prod     date,  
    kind          varchar(10),  
    len           interval hour to minute,  
    CONSTRAINT production UNIQUE(date_prod)  
);
```

Define a check column constraint:

```
CREATE TABLE distributors (  
    did           integer CHECK (did > 100),  
    name          varchar(40)  
);
```

Define a check table constraint:

```
CREATE TABLE distributors (  
    did           integer,  
    name          varchar(40)  
    CONSTRAINT con1 CHECK (did > 100 AND name <> '')  
);
```

Define a primary key table constraint for the table `films`:

```
CREATE TABLE films (  
    code          char(5),
```

```

        title      varchar(40),
        did        integer,
        date_prod  date,
        kind       varchar(10),
        len        interval hour to minute,
        CONSTRAINT code_title PRIMARY KEY(code,title)
    );

```

Define a primary key constraint for table `distributors`. The following two examples are equivalent, the first using the table constraint syntax, the second the column constraint syntax:

```

CREATE TABLE distributors (
    did        integer,
    name       varchar(40),
    PRIMARY KEY(did)
);

```

```

CREATE TABLE distributors (
    did        integer PRIMARY KEY,
    name       varchar(40)
);

```

Assign a literal constant default value for the column `name`, arrange for the default value of column `did` to be generated by selecting the next value of a sequence object, and make the default value of `modtime` be the time at which the row is inserted:

```

CREATE TABLE distributors (
    name       varchar(40) DEFAULT 'Luso Films',
    did        integer DEFAULT nextval('distributors_serial'),
    modtime    timestamp DEFAULT current_timestamp
);

```

Define two `NOT NULL` column constraints on the table `distributors`, one of which is explicitly given a name:

```

CREATE TABLE distributors (
    did        integer CONSTRAINT no_null NOT NULL,
    name       varchar(40) NOT NULL
);

```

Define a unique constraint for the `name` column:

```

CREATE TABLE distributors (
    did        integer,
    name       varchar(40) UNIQUE
);

```

The same, specified as a table constraint:

```

CREATE TABLE distributors (
    did        integer,
    name       varchar(40),
    UNIQUE(name)
);

```

Create the same table, specifying 70% fill factor for both the table and its unique index:

```
CREATE TABLE distributors (  
    did      integer,  
    name     varchar(40),  
    UNIQUE(name) WITH (fillfactor=70)  
)  
WITH (fillfactor=70);
```

Create table `circles` with an exclusion constraint that prevents any two circles from overlapping:

```
CREATE TABLE circles (  
    c circle,  
    EXCLUDE USING gist (c WITH &&)  
);
```

Create table `cinemas` in tablespace `diskvol1`:

```
CREATE TABLE cinemas (  
    id serial,  
    name text,  
    location text  
) TABLESPACE diskvol1;
```

Create a composite type and a typed table:

```
CREATE TYPE employee_type AS (name text, salary numeric);  
  
CREATE TABLE employees OF employee_type (  
    PRIMARY KEY (name),  
    salary WITH OPTIONS DEFAULT 1000  
);
```

Compatibility

The `CREATE TABLE` command conforms to the SQL standard, with exceptions listed below.

Temporary Tables

Although the syntax of `CREATE TEMPORARY TABLE` resembles that of the SQL standard, the effect is not the same. In the standard, temporary tables are defined just once and automatically exist (starting with empty contents) in every session that needs them. PostgreSQL instead requires each session to issue its own `CREATE TEMPORARY TABLE` command for each temporary table to be used. This allows different sessions to use the same temporary table name for different purposes, whereas the standard's approach constrains all instances of a given temporary table name to have the same table structure.

The standard's definition of the behavior of temporary tables is widely ignored. PostgreSQL's behavior on this point is similar to that of several other SQL databases.

The standard's distinction between global and local temporary tables is not in PostgreSQL, since that distinction depends on the concept of modules, which PostgreSQL does not have. For compatibility's

sake, PostgreSQL will accept the `GLOBAL` and `LOCAL` keywords in a temporary table declaration, but they have no effect.

The `ON COMMIT` clause for temporary tables also resembles the SQL standard, but has some differences. If the `ON COMMIT` clause is omitted, SQL specifies that the default behavior is `ON COMMIT DELETE ROWS`. However, the default behavior in PostgreSQL is `ON COMMIT PRESERVE ROWS`. The `ON COMMIT DROP` option does not exist in SQL.

Non-deferred Uniqueness Constraints

When a `UNIQUE` or `PRIMARY KEY` constraint is not deferrable, PostgreSQL checks for uniqueness immediately whenever a row is inserted or modified. The SQL standard says that uniqueness should be enforced only at the end of the statement; this makes a difference when, for example, a single command updates multiple key values. To obtain standard-compliant behavior, declare the constraint as `DEFERRABLE` but not deferred (i.e., `INITIALLY IMMEDIATE`). Be aware that this can be significantly slower than immediate uniqueness checking.

Column Check Constraints

The SQL standard says that `CHECK` column constraints can only refer to the column they apply to; only `CHECK` table constraints can refer to multiple columns. PostgreSQL does not enforce this restriction; it treats column and table check constraints alike.

EXCLUDE Constraint

The `EXCLUDE` constraint type is a PostgreSQL extension.

NULL "Constraint"

The `NULL` "constraint" (actually a non-constraint) is a PostgreSQL extension to the SQL standard that is included for compatibility with some other database systems (and for symmetry with the `NOT NULL` constraint). Since it is the default for any column, its presence is simply noise.

Inheritance

Multiple inheritance via the `INHERITS` clause is a PostgreSQL language extension. SQL:1999 and later define single inheritance using a different syntax and different semantics. SQL:1999-style inheritance is not yet supported by PostgreSQL.

Zero-column Tables

PostgreSQL allows a table of no columns to be created (for example, `CREATE TABLE foo();`). This is an extension from the SQL standard, which does not allow zero-column tables. Zero-column tables are not in themselves very useful, but disallowing them creates odd special cases for `ALTER TABLE DROP COLUMN`, so it seems cleaner to ignore this spec restriction.

LIKE Clause

While a LIKE clause exists in the SQL standard, many of the options that PostgreSQL accepts for it are not in the standard, and some of the standard's options are not implemented by PostgreSQL.

WITH Clause

The WITH clause is a PostgreSQL extension; neither storage parameters nor OIDs are in the standard.

Tablespaces

The PostgreSQL concept of tablespaces is not part of the standard. Hence, the clauses TABLESPACE and USING INDEX TABLESPACE are extensions.

Typed Tables

Typed tables implement a subset of the SQL standard. According to the standard, a typed table has columns corresponding to the underlying composite type as well as one other column that is the "self-referencing column". PostgreSQL does not support these self-referencing columns explicitly, but the same effect can be had using the OID feature.

See Also

[ALTER TABLE](#), [DROP TABLE](#), [CREATE TABLESPACE](#), [CREATE TYPE](#)

[Prev](#)

CREATE SERVER

[Home](#)

[Up](#)

[Next](#)

CREATE TABLE AS

Submit correction

If you see anything in the documentation that is not correct, does not match your experience with the particular feature or requires further clarification, please use [this form](#) to report a documentation issue.

[Privacy Policy](#) | [About PostgreSQL](#)

Copyright © 1996-2017 The PostgreSQL Global Development Group