

# SQL Functions



## Copyright

© Postgres Professional, 2017–2021

Authors: Egor Rogov, Pavel Luzanov

Translated by Liudmila Mantrova

## Usage of Course Materials

Non-commercial use of course materials (presentations, demonstrations) is allowed on an unrestricted basis. Commercial use is only possible with prior written permission of Postgres Professional company. Modification of course materials is forbidden.

## Contact Us

Please send your feedback to: [edu@postgrespro.ru](mailto:edu@postgrespro.ru)

## Disclaimer

In no event shall Postgres Professional company be liable for any damages or loss, including loss of profits, that arise from direct or indirect, special or incidental use of course materials. Postgres Professional company specifically disclaims any warranties on course materials. Course materials are provided “as is,” and Postgres Professional company has no obligations to provide maintenance, support, updates, enhancements, or modifications.

# Agenda



Functions and their specifics in the context of databases

Parameters and return values

Passing parameters when calling a function

Volatility categories and query planning

# Functions in Databases

The main goal is simplifying development tasks

interface (parameters) and implementation (function body)

abstracting from other tasks when implementing a particular function

	<i>Traditional languages</i>	<i>PostgreSQL</i>
side effects	global variables	the whole database (volatility categories)
modules	have their own interface and implementation	namespaces, client and server
challenges	function call overhead (inlining)	hiding the query from the planner (inlining, subqueries, views)

The main goal of introducing functions in programming is simplifying development tasks by decomposing them into smaller subtasks. Such simplification is possible because you can abstract from the “big” task when thinking of a function. For this purpose, the function provides a precise interface to the outside world (parameters and the return value). Its implementation (the function body) can change; the caller does not see these changes and does not depend on them. This ideal situation can be messed up by the global state (global variables), and you have to keep in mind that in the DB context the whole database constitutes such a state.

In traditional programming languages, functions are often grouped into modules (packages, classes for OOP, etc.), which have their own interface and implementation. This separation into modules can be more or less arbitrary. In PostgreSQL, there is a fixed boundary between client and server parts: the server code deals with the database, while the client code manages transactions. There are no modules (or packages), only namespaces are present.

The only disadvantage of extensive use of functions in traditional languages is function call overhead. It is sometimes overcome by inlining function code into the calling program. In databases, the consequences can be more serious: if some part of the query is moved into a function, the planner stops seeing the “big picture” and cannot build a good query plan. In some cases, PostgreSQL can also perform inlining; alternatively, subqueries or views can be used.

## A database object

function declaration is stored in the system catalog

## The structure of function declaration

name

parameters

return data type

body

## Can be written in various languages, including SQL

the code is stored as a string literal

a function is interpreted when it is called

## Is called in the context of an expression

4

Functions are regular database objects, just like tables or indexes. Function declarations are stored in the system catalog; that's why database functions are called *stored* functions.

PostgreSQL provides a lot of standard functions. Some of them are listed in the "Basic Data Types and Functions" handout.

You can also write your own functions in various programming languages. The information provided in this lecture applies to functions in any programming language, but we will use SQL in all examples.

Predictably, a function declaration consists of a name, optional parameters, a return data type, and a body. What may seem unexpected is that the body is written as a string literal, which contains the code written in the programming language of your choice. It makes function declarations look the same regardless of the used programming language. The body string is stored in the system catalog and is interpreted each time the function is called. At the moment, the only way to avoid interpretation is to write a function in the C language, but it's required not too often, so we are not going to discuss it in this course.

A function is always called within the context of an expression. For example, in the list of expressions of the SELECT statement, in the WHERE clause, in CHECK constraints, etc.

<https://postgrespro.com/docs/postgresql/12/sql-createfunction>

<https://postgrespro.com/docs/postgresql/12/sql-syntax-calling-funcs>

## Functions without Parameters

Here is a simple example of a function with no parameters:

```
=> CREATE FUNCTION hello_world() -- function name and an empty list of parameters
RETURNS text                    -- the type of the return value
AS $$ SELECT 'Hello, world!'; $$ -- function body
LANGUAGE sql;                  -- language specification
```

```
CREATE FUNCTION
```

It is convenient to write the body as a dollar-quoted string, as shown in the example above. Otherwise, you have to take care of escaping quotes, which are sure to appear in the function body. Compare the following strings:

```
=> SELECT ' SELECT ''Hello, world!''; ';
```

```
      ?column?
-----
SELECT 'Hello, world!';
(1 row)
```

```
=> SELECT $$ SELECT 'Hello, world!'; $$;
```

```
      ?column?
-----
SELECT 'Hello, world!';
(1 row)
```

If required, dollar quoting can be nested. It is achieved by using different text strings between dollars in each pair of quotes:

```
=> SELECT $func$ SELECT $$Hello, world!$$; $func$;
```

```
      ?column?
-----
SELECT $$Hello, world!$$;
(1 row)
```

A function is called in the context of an expression. For example:

```
=> SELECT hello_world(); -- empty brackets are mandatory
```

```
      hello_world
-----
Hello, world!
(1 row)
```

In general, a function body can contain several SQL operators. The return value is taken from the first row returned by the last operator.

Not all SQL operators can be used in a function. The following ones are forbidden:

- transaction control commands (BEGIN, COMMIT, ROLLBACK, etc.);
- service commands (such as VACUUM or CREATE INDEX).

Here is an example of an invalid function. We have used the void pseudotype, which indicates that the function returns nothing.

```
=> CREATE FUNCTION do_commit() RETURNS void AS $$
COMMIT;
$$ LANGUAGE sql;
```

```
CREATE FUNCTION
```

```
=> SELECT do_commit();
```

```
ERROR:  COMMIT is not allowed in a SQL function
CONTEXT:  SQL function "do_commit" during startup
```

You can use procedures to manage transactions; we will cover this topic in the next lecture.

---

## Functions with Input Parameters

Here is a function with a single parameter:

```
=> CREATE FUNCTION hello(name text) -- a formal parameter
RETURNS text AS $$
SELECT 'Hello, ' || name || '!';
$$ LANGUAGE sql;
```

CREATE FUNCTION

When calling this function, we have to specify the actual value that corresponds to the formal parameter:

```
=> SELECT hello('Alice');
```

```
      hello
-----
Hello, Alice!
(1 row)
```

When specifying parameter types, you can add a modifier (such as varchar(10)), but it will be ignored.

You can define a function parameter without a name; then the function body will have to refer to it by its position number. Let's delete this function and create a new one:

```
=> DROP FUNCTION hello(text); -- it is enough to specify the parameter type
```

DROP FUNCTION

```
=> CREATE FUNCTION hello(text)
RETURNS text AS $$
SELECT 'Hello, ' || $1 || '!'; -- a number instead of the name
$$ LANGUAGE sql;
```

CREATE FUNCTION

```
=> SELECT hello('Alice');
```

```
      hello
-----
Hello, Alice!
(1 row)
```

But this approach is inconvenient and should be avoided.

Let's delete and recreate the function again, now adding one more parameter to include the title of a person.

```
=> DROP FUNCTION hello(text);
```

DROP FUNCTION

Here we have used an optional IN keyword, which means the input parameter. The DEFAULT clause is used to define the default parameter value:

```
=> CREATE FUNCTION hello(IN name text, IN title text DEFAULT 'Mr.')
RETURNS text AS $$
SELECT 'Hello, ' || title || ' ' || name || '!';
$$ LANGUAGE sql;
```

CREATE FUNCTION

```
=> SELECT hello('Alice', 'Dr.');
```

```
      hello
-----
Hello, Dr. Alice!
(1 row)
```

```
=> SELECT hello('Bob'); -- the parameter with the default value is omitted
```

```
      hello
-----
Hello, Mr. Bob!
(1 row)
```

So far, we have provided function parameters as positional ones, in the order they were specified in the function declaration. In many standard functions, parameter names are not set, so it is the only way possible.

But if the formal parameters are named, you can use these names when providing their actual values. In this case, parameters can be specified in any order:

```
=> SELECT hello(title => 'Dr.', name => 'Alice');
```

```

      hello
-----
Hello, Dr. Alice!
(1 row)

```

```
=> SELECT hello(name => 'Bob');
```

```

      hello
-----
Hello, Mr. Bob!
(1 row)

```

This approach is convenient if the order of parameters is not quite obvious, especially if there are a lot of them.

You can combine both conventions: provide some parameters by position (starting from the first one) and specify the rest by name:

```
=> SELECT hello('Alice', title => 'Dr.');
```

```

      hello
-----
Hello, Dr. Alice!
(1 row)

```

If the function must return NULL when at least one of its input parameters is NULL, it can be declared STRICT. In this case, the function body will not be executed at all.

```
=> DROP FUNCTION hello(text, text);
```

```
DROP FUNCTION
```

```
=> CREATE FUNCTION hello(IN name text, IN title text DEFAULT 'NULL')
RETURNS text AS $$
SELECT 'Hello, ' || title || ' ' || name || '!';
$$ LANGUAGE sql STRICT;
```

```
CREATE FUNCTION
```

```
=> SELECT hello('Alice', NULL);
```

```

      hello
-----
(1 row)

```

## Input values

are defined by parameters with IN or INOUT modes

## Output value

is defined either by the RETURNS clause

or by parameters with INOUT or OUT modes

if both forms are specified, they must not contradict each other

Formal parameters that have IN or INOUT modes are considered to be *input* parameters. Their actual values must be specified in the function call (or the default values must be defined).

There are two ways to define the return value:

- use the RETURNS clause to specify the return data type
- define *output* parameters using INOUT or OUT modes

These two approaches are equivalent. For example, a function with the RETURNS integer clause and a function with the OUT integer parameter both return an integer number.

You can combine these two approaches. In this case, the function will also return *one* integer number. But note that the types of the output parameters and the RETURNS clause must not contradict each other.

Thus, you cannot write a function that returns one value, but passes another value into the OUT parameter (which is allowed in many traditional programming languages). In PostgreSQL, such functions *return both values*.



## Functions with Output Parameters

An alternative way to return a value is to use an output parameter.

```
=> DROP FUNCTION hello(text, text);
```

DROP FUNCTION

```
=> CREATE FUNCTION hello(  
    IN name text,  
    OUT text -- you can omit the parameter name if it is not required  
)  
AS $$  
SELECT 'Hello, ' || name || '!';  
$$ LANGUAGE sql;
```

CREATE FUNCTION

```
=> SELECT hello('Alice');
```

```
      hello  
-----  
Hello, Alice!  
(1 row)
```

The result is the same.

You can use the RETURNS clause and the OUT parameter together: the result will be the same anyway:

```
=> DROP FUNCTION hello(text); -- OUT parameters are omitted
```

DROP FUNCTION

```
=> CREATE FUNCTION hello(IN name text, OUT text)  
RETURNS text AS $$  
SELECT 'Hello, ' || name || '!';  
$$ LANGUAGE sql;
```

CREATE FUNCTION

```
=> SELECT hello('Alice');
```

```
      hello  
-----  
Hello, Alice!  
(1 row)
```

Or even use an INOUT parameter:

```
=> DROP FUNCTION hello(text);
```

DROP FUNCTION

```
=> CREATE FUNCTION hello(INOUT name text)  
AS $$  
SELECT 'Hello, ' || name || '!';  
$$ LANGUAGE sql;
```

CREATE FUNCTION

```
=> SELECT hello('Alice');
```

```
      hello  
-----  
Hello, Alice!  
(1 row)
```

Note that the actual value passed to the SQL function in an INOUT parameter is not modified: we pass an input value, and the output value is returned as a result (SQL differs from many other programming languages in this respect). That's why we can specify a constant, although other languages would require a variable.

While the RETURNS clause can take only one value, there can be several output parameters. For example:

```
=> DROP FUNCTION hello(text);
```

DROP FUNCTION

```

=> CREATE FUNCTION hello(
    IN name text,
    OUT greeting text,
    OUT clock timetz)
AS $$
SELECT 'Hello, ' || name || '!', current_time;
$$ LANGUAGE sql;

CREATE FUNCTION

=> SELECT hello('Alice');

             hello
-----
("Hello, Alice!",17:01:28.935637+03)
(1 row)

```

Indeed, our function has returned not just one, but several values at once.

We will provide more details about this feature and composite types in the “SQL. Composite Types” lecture.

# Volatility Categories

## Volatile

- returns different values for the same input arguments
- is used by default

## Stable

- the value cannot change within a single SQL operator
- the function cannot change the database state

## Immutable

- the value cannot change, the function is deterministic
- the function cannot change the database state

Each function is mapped to a particular volatility category, which defines the properties of the return value for the same input arguments.

The volatile category means that the return value can change randomly. Such functions will be executed each time they are called. If the function is declared without a category specification, it is assumed to be volatile.

The stable category is used for functions that always return the same value within a single SQL operator. In particular, such functions cannot change the state of the database. It is *possible* to execute such a function only once during the query and then use the computed value.

The immutable category is even more strict: the return value always remains the same. Such a function *can* be executed at the planning stage, before the query is actually executed.

It does not mean that it happens so all the time, but the planner has the right to perform such optimizations. In some (simple) cases, the planner makes its own assumptions about function volatility, regardless of the explicitly provided category.

<https://postgrespro.com/docs/postgresql/12/xfunc-volatility>

## Volatility Categories and Isolation

In general, using functions within queries does not violate the isolation level of the transaction, but there are two points worth knowing.

First, volatile functions can cause data inconsistency within the query when used at the Read Committed level.

Let's create a function that returns the number of rows in a table:

```
=> CREATE TABLE t(n integer);
```

CREATE TABLE

```
=> CREATE FUNCTION cnt() RETURNS bigint
AS $$
    SELECT count(*) FROM t;
$$ VOLATILE LANGUAGE sql;
```

CREATE FUNCTION

Now let's call it several times with a delay and insert a row into the table in a parallel session.

```
=> BEGIN ISOLATION LEVEL READ COMMITTED;
```

BEGIN

```
=> SELECT (SELECT count(*) FROM t), cnt(), pg_sleep(1)
FROM generate_series(1,4);
```

```
| => INSERT INTO t VALUES (1);
```

```
| INSERT 0 1
```

count	cnt	pg_sleep
0	0	
0	0	
0	1	
0	1	

(4 rows)

```
=> END;
```

COMMIT

It won't happen at stricter isolation levels, or if the function is stable or immutable.

```
=> ALTER FUNCTION cnt() STABLE;
```

ALTER FUNCTION

```
=> TRUNCATE t;
```

TRUNCATE TABLE

```
=> BEGIN ISOLATION LEVEL READ COMMITTED;
```

BEGIN

```
=> SELECT (SELECT count(*) FROM t), cnt(), pg_sleep(1)
FROM generate_series(1,4);
```

```
| => INSERT INTO t VALUES (1);
```

```
| INSERT 0 1
```

count	cnt	pg_sleep
0	0	
0	0	
0	0	
0	0	

(4 rows)

```
=> END;
```

COMMIT

---

Another point is the visibility of changes made by the same transaction.

Volatile functions can see all the changes, even those made by the current SQL operator that has not been completed yet.

```
=> ALTER FUNCTION cnt() VOLATILE;

ALTER FUNCTION

=> TRUNCATE t;

TRUNCATE TABLE

=> INSERT INTO t SELECT cnt() FROM generate_series(1,5);

INSERT 0 5

=> SELECT * FROM t;

 n
---
 0
 1
 2
 3
 4
(5 rows)
```

It is true for any isolation level.

Stable and immutable functions see only the changes performed by an already completed operator.

```
=> ALTER FUNCTION cnt() STABLE;

ALTER FUNCTION

=> TRUNCATE t;

TRUNCATE TABLE

=> INSERT INTO t SELECT cnt() FROM generate_series(1,5);

INSERT 0 5

=> SELECT * FROM t;

 n
---
 0
 0
 0
 0
 0
(5 rows)
```

---

## Volatility Categories and Performance Tuning

Thanks to the volatility labels that provide additional information about the function behavior, the optimizer can spare some function calls.

To try it out, let's create a function that returns a random number:

```
=> CREATE FUNCTION rnd() RETURNS float
AS $$
    SELECT random();
$$ VOLATILE LANGUAGE sql;
```

CREATE FUNCTION

Let's check the execution plan of the following query:

```
=> EXPLAIN (costs off)
SELECT * FROM generate_series(1,10) WHERE rnd() > 0.5;
```

```
          QUERY PLAN
-----
Function Scan on generate_series
  Filter: (random() > '0.5'::double precision)
(2 rows)
```

The query plan shows that the generate\_series function is honestly called several times; each result is compared with a random number and is filtered out, if required.

You can see it for yourself (we expect to get about 5 rows on average):

```
=> SELECT * FROM generate_series(1,10) WHERE rnd() > 0.5;

generate_series
-----
              1
              4
              5
              6
(4 rows)
```

A stable function will be called only once, because we have virtually specified that its value cannot change within a single operator:

```
=> ALTER FUNCTION rnd() STABLE;

ALTER FUNCTION

=> EXPLAIN (costs off)
SELECT * FROM generate_series(1,10) WHERE rnd() > 0.5;

              QUERY PLAN
-----
Result
One-Time Filter: (rnd() > '0.5'::double precision)
-> Function Scan on generate_series
(3 rows)
```

Finally, immutable functions are computed at the planning stage, so we do not need any filters during execution:

```
=> ALTER FUNCTION rnd() IMMUTABLE;

ALTER FUNCTION

=> EXPLAIN (costs off)
SELECT * FROM generate_series(1,10) WHERE rnd() > 0.5;

              QUERY PLAN
-----
Function Scan on generate_series
(1 row)
```

It is the developer's responsibility to provide the correct information.

---

## Function Inlining

In some (very simple) cases, a function can be inlined: the function body written in SQL can be inserted right into the main SQL operator while the query is being parsed. In this case, we can save some time on the function call.

Roughly speaking, the following conditions should be met:

- The function body contains only one SELECT operator.
- There are no table lookups, subqueries, grouping operations, etc.
- There must be only one return value.
- The called functions must not violate the specified volatility category.

We have already seen such an example: the rnd() function, which is declared volatile.

Let's take another look.

```
=> ALTER FUNCTION rnd() VOLATILE;

ALTER FUNCTION

=> EXPLAIN (costs off)
SELECT * FROM generate_series(1,10) WHERE rnd() > 0.5;

              QUERY PLAN
-----
Function Scan on generate_series
Filter: (random() > '0.5'::double precision)
(2 rows)
```

The Filter does not mention the rnd() function, only random() is present; it will be called directly, without using the rnd() wrapper.

You can create your own function and use it in the same manner as the built-in ones

Functions can be written in various languages, including SQL

Volatility categories affect optimization opportunities

An SQL function can sometimes be inlined



1. Create the `author_name` function to construct author names. The function takes three parameters (`last_name`, `first_name`, and `middle_name`) and returns the full name, with the middle name abbreviated to its initial. Use this function in the `authors_v` view.
  2. Create the `book_name` function to construct book titles. The function takes two parameters (book ID and the title) and returns a concatenation of the book title and the list of authors, in the `seq_num` order. The name of each author is produced by the `author_name` function. Use this function in the `catalog_v` view.
- Check the changes in the application.

Just a friendly reminder: all the required functions are listed in the “Basic Data Types and Functions” handout.

**Task 1.** `FUNCTION author_name(  
    first_name text, middle_name text, last_name text  
)  
RETURNS text`

For example: `author_name('Alexander', 'Sergeyevich', 'Pushkin')`  
→ 'Alexander S. Pushkin'

**Task 2.** `FUNCTION book_name(book_id integer, title text)  
RETURNS text`

For example: `book_name(3, 'Good Omens')` →  
→ 'Good Omens. Terry Pratchett, Neil Gaiman'

Stored functions can be edited directly in any tool. For example, `psql` provides the `\ef` command that opens the function body in an editor and saves the changes in the database.

You should avoid using this capability (or at least do not overuse it). A properly set up development process requires that all the code is stored in files under version control. If a function has to be changed, the file is modified and executed (using `psql` or an IDE). Function modifications made directly in the database can be easily lost. (In fact, setting up development processes is much more complex, but we are not going to cover it in this course.)



## Task 1. The author\_name Function

```
=> CREATE OR REPLACE FUNCTION author_name(  
    last_name text,  
    first_name text,  
    middle_name text  
) RETURNS text  
AS $$  
SELECT first_name ||  
    CASE WHEN middle_name != '' -- NOT NULL is implied  
        THEN ' ' || left(middle_name, 1) || '.'  
        ELSE ''  
    END || ' ' ||  
    last_name;  
$$ IMMUTABLE LANGUAGE sql;
```

CREATE FUNCTION

Volatility category: immutable. The function always returns the same value given the same input arguments.

```
=> CREATE OR REPLACE VIEW authors_v AS  
SELECT a.author_id,  
    author_name(a.last_name, a.first_name, a.middle_name) AS display_name  
FROM authors a  
ORDER BY display_name;
```

CREATE VIEW

## Task 2. The book\_name Function

```
=> CREATE OR REPLACE FUNCTION book_name(book_id integer, title text)  
RETURNS text  
AS $$  
SELECT title || '. ' ||  
    string_agg(  
        author_name(a.last_name, a.first_name, a.middle_name), ', '  
        ORDER BY ash.seq_num  
    )  
FROM authors a  
    JOIN authorship ash ON a.author_id = ash.author_id  
WHERE ash.book_id = book_name.book_id;  
$$ STABLE LANGUAGE sql;
```

CREATE FUNCTION

Volatility category: stable. The function returns the same value given the same input arguments, but only within a single SQL query.

```
=> DROP VIEW IF EXISTS catalog_v;
```

DROP VIEW

```
=> CREATE VIEW catalog_v AS  
SELECT b.book_id,  
    book_name(b.book_id, b.title) AS display_name  
FROM books b  
ORDER BY display_name;
```

CREATE VIEW

## Practice

1. Write a function that returns random timestamps, equally distributed within the specified time range. The lower bound of the range is set by a `timestampz`, while the upper bound is either a `timestampz` or an `interval`.
2. Sometimes phonewords are used as a mnemonic equivalent of telephone numbers, like 1-800-TAXICAB. A phoneword is constructed by replacing some digits with the corresponding letters, as seen on the keypad (even on modern phones). Write a function that returns number representation of a phoneword, no dashes.
3. Write a function that finds square roots of an equation.

12

In all tasks, make sure to pay attention to volatility categories of the functions.

**Task 2.** The correspondence between digits and letters is as follows:

2 – abc, 3 – def, 4 – ghi, 5 – jkl, 6 – mno, 7 – pqrs, 8 – tuv, 9 – wxyz

Note that a phoneword can be written in lowercase as well as in uppercase.

**Task 3.** For the equation  $y = ax^2 + bx + c$ :

discriminant  $D = b^2 - 4ac$ .

if  $D > 0$ , then there are two roots  $x_{1,2} = (-b \pm \sqrt{D}) / 2a$

if  $D = 0$ , then there is one root  $x = -b / 2a$  (null can be returned as  $x_2$ )

if  $D < 0$ , then there are no roots (both roots are null).

## Task 1. A Random Timestamp

Here is a function with two timestamps:

```
=> CREATE FUNCTION rnd_timestamp(t_start timestamptz, t_end timestamptz)
RETURNS timestamptz
AS $$
    SELECT t_start + (t_end - t_start) * random();
$$ VOLATILE LANGUAGE sql;
```

CREATE FUNCTION

This function is volatile. It will return different values given the same input arguments since the random() function is used.

```
=> SELECT current_timestamp,
        rnd_timestamp(
            current_timestamp,
            current_timestamp + interval '1 hour'
        )
FROM generate_series(1,10);
```

current_timestamp		rnd_timestamp
2021-10-19 17:05:26.981294+03		2021-10-19 17:36:20.432004+03
2021-10-19 17:05:26.981294+03		2021-10-19 18:01:25.771207+03
2021-10-19 17:05:26.981294+03		2021-10-19 17:29:37.688731+03
2021-10-19 17:05:26.981294+03		2021-10-19 17:20:28.820816+03
2021-10-19 17:05:26.981294+03		2021-10-19 17:51:26.45336+03
2021-10-19 17:05:26.981294+03		2021-10-19 17:54:06.357517+03
2021-10-19 17:05:26.981294+03		2021-10-19 17:59:15.308904+03
2021-10-19 17:05:26.981294+03		2021-10-19 17:19:25.796948+03
2021-10-19 17:05:26.981294+03		2021-10-19 17:06:41.498541+03
2021-10-19 17:05:26.981294+03		2021-10-19 17:10:44.500444+03

(10 rows)

The second function (that has an interval parameter) can be defined via the first one:

```
=> CREATE FUNCTION rnd_timestamp(t_start timestamptz, t_delta interval)
RETURNS timestamptz
AS $$
    SELECT rnd_timestamp(t_start, t_start + t_delta);
$$ VOLATILE LANGUAGE sql;
```

CREATE FUNCTION

```
=> SELECT rnd_timestamp(current_timestamp, interval '1 hour');
```

rnd_timestamp
2021-10-19 17:24:00.635935+03

(1 row)

## Task 2. Phonewords

```
=> CREATE FUNCTION phoneword_to_number(phoneword text) RETURNS numeric
AS $$
    SELECT translate(
        lower(phoneword),
        'abcdefghijklmnopqrstuvwxyz-',
        '2223334445556667778889999'
    )::numeric;
$$ IMMUTABLE LANGUAGE sql;
```

CREATE FUNCTION

Volatility category: immutable. The function always returns the same value given the same input arguments.

```
=> SELECT phoneword_to_number('1-800-TAXICAB');
```

phoneword_to_number
18008294222

(1 row)

### Task 3. Square Roots

```
=> CREATE FUNCTION square_roots(  
    a float,  
    b float,  
    c float,  
    x1 OUT float,  
    x2 OUT float  
)  
AS $$  
WITH discriminant(d) AS (  
    SELECT b*b - 4*a*c  
)  
SELECT CASE WHEN d >= 0.0 THEN (-b + sqrt(d))/2/a END,  
       CASE WHEN d > 0.0 THEN (-b - sqrt(d))/2/a END  
FROM discriminant;  
$$ IMMUTABLE LANGUAGE sql;
```

CREATE FUNCTION

Volatility category: immutable. The function always returns the same value given the same input arguments.

```
=> SELECT square_roots(1, 0, -4);
```

```
square_roots  
-----  
(2,-2)  
(1 row)
```

```
=> SELECT square_roots(1, -4, 4);
```

```
square_roots  
-----  
(2,)  
(1 row)
```

```
=> SELECT square_roots(1, 1, 1);
```

```
square_roots  
-----  
(,)  
(1 row)
```