

---

# PostgREST Documentation

*Release 5.0.0*

**Joe Nelson**

**Oct 22, 2019**



<b>1</b>	<b>Motivation</b>	<b>3</b>
<b>2</b>	<b>Declarative Programming</b>	<b>5</b>
<b>3</b>	<b>Leak-proof Abstraction</b>	<b>7</b>
<b>4</b>	<b>One Thing Well</b>	<b>9</b>
<b>5</b>	<b>Getting Support</b>	<b>11</b>
5.1	v5.2.0 . . . . .	11
5.2	v6.0.2 . . . . .	12
<b>6</b>	<b>Tutorials</b>	<b>15</b>
6.1	Tutorial 0 - Get it Running . . . . .	15
6.2	Tutorial 1 - The Golden Key . . . . .	18
<b>7</b>	<b>Reference guides</b>	<b>25</b>
7.1	Tables and Views . . . . .	25
7.2	Resource Embedding . . . . .	32
7.3	Custom Queries . . . . .	35
7.4	Stored Procedures . . . . .	36
7.5	Insertions / Updates . . . . .	41
7.6	Deletions . . . . .	43
7.7	Binary Output . . . . .	43
7.8	OpenAPI Support . . . . .	44
7.9	HTTP Status Codes . . . . .	45
7.10	Configuration . . . . .	46
<b>8</b>	<b>How-to guides</b>	<b>51</b>
8.1	Casting a type to a custom JSON object . . . . .	51
8.2	Embedding a table from another schema . . . . .	52
<b>9</b>	<b>Topic guides</b>	<b>55</b>
9.1	Overview of Role System . . . . .	55
9.2	Client Auth . . . . .	59
9.3	Schema Isolation . . . . .	62
9.4	SQL User Management . . . . .	62
9.5	Binary Release . . . . .	65

9.6	PostgreSQL dependency . . . . .	65
9.7	Configuration . . . . .	66
9.8	Running the Server . . . . .	66
9.9	Docker . . . . .	66
9.10	Deploying to Heroku . . . . .	68
9.11	Build from Source . . . . .	68
9.12	Hardening PostgREST . . . . .	71
9.13	Debugging . . . . .	73
9.14	Daemonizing . . . . .	75
9.15	Alternate URL Structure . . . . .	76
<b>10</b>	<b>Ecosystem</b>	<b>77</b>
10.1	Example Apps . . . . .	77
10.2	External Notification . . . . .	78
10.3	Extensions . . . . .	78
10.4	Client-Side Libraries . . . . .	78
10.5	Commercial . . . . .	79
<b>11</b>	<b>Release Notes</b>	<b>81</b>
<b>12</b>	<b>In Production</b>	<b>83</b>
<b>13</b>	<b>Testimonials</b>	<b>85</b>
<b>14</b>	<b>Translations</b>	<b>87</b>



PostgREST is a standalone web server that turns your PostgreSQL database directly into a RESTful API. The structural constraints and permissions in the database determine the API endpoints and operations.



PostgREST for the backend.  
Retool for the frontend.



# CHAPTER 1

---

## Motivation

---

Using PostgREST is an alternative to manual CRUD programming. Custom API servers suffer problems. Writing business logic often duplicates, ignores or hobbles database structure. Object-relational mapping is a leaky abstraction leading to slow imperative code. The PostgREST philosophy establishes a single declarative source of truth: the data itself.



## CHAPTER 2

---

### Declarative Programming

---

It's easier to ask PostgreSQL to join data for you and let its query planner figure out the details than to loop through rows yourself. It's easier to assign permissions to db objects than to add guards in controllers. (This is especially true for cascading permissions in data dependencies.) It's easier to set constraints than to litter code with sanity checks.



## CHAPTER 3

---

### Leak-proof Abstraction

---

There is no ORM involved. Creating new views happens in SQL with known performance implications. A database administrator can now create an API from scratch with no custom programming.



## CHAPTER 4

---

### One Thing Well

---

PostgREST has a focused scope. It works well with other tools like Nginx. This forces you to cleanly separate the data-centric CRUD operations from other concerns. Use a collection of sharp tools rather than building a big ball of mud.



The project has a friendly and growing community. Join our [chat room](#) for discussion and help. You can also report or search for bugs/features on the Github [issues](#) page.

### 5.1 v5.2.0

- [Explicit qualification](#) introduced in v5.0 is no longer necessary, this section will not be included from this version onwards. A *db-extra-search-path* configuration parameter was introduced to avoid the need to explicitly qualify database objects. If you install PostgreSQL extensions on the `public` schema, they'll work normally from now on.
- Now you can filter *Table / Columns with spaces*.
- Included the ability to quote columns that have *Reserved characters*.
- Thanks to [Zhou Feng](#), now is possible to reference an external file in *db-uri*.
- Thanks to [Russell Davies](#), Json Web Key Sets are now accepted by *jwt-secret*.

#### 5.1.1 Thanks

This release was made possible thanks to:

- [Daniel Babiak](#)
- [Michel Pelletier](#)
- [Tsingson Qin](#)
- [Jay Hannah](#)
- [Victor Adossi](#)
- [Petr Beles](#)

If you like to join them please consider [supporting PostgREST development](#).

## 5.2 v6.0.2

Full changelog is available at [PostgREST releases page](#).

### 5.2.1 Added

- Ignoring payload keys for insert/update can be now done with the `?columns` query parameter. See *Specifying Columns*. – @steve-chavez
- `websearch_to_tsquery` can now be used through the `wfts` operator. See *Full-Text Search*. – @herulume
- Resource Embedding on materialized views is now possible. See *Embedding Views*. – @vitorbaptista
- Bulk calling an RPC is now allowed. See *Bulk Call*. – @steve-chavez
- It's now possible to request a `text/plain` output. See *Plain Text Output*. – @steve-chavez
- Config option for specifying PostgREST database pool timeout. See *db-pool-timeout*. – @Qu4tro
- Config option for binding the PostgREST web server to an unix socket. See *server-unix-socket*. – @Dansvidania
- Config option for extending the supported media types. See *raw-media-types*. – @Dansvidania
- We now offer an statically linked binary for Linux. Look for `postgrest-<version>-linux-x64-static.tar.xz` on the [releases page](#). – @clojurians-org
- A *How-to guides* section was added to the documentation.

### 5.2.2 Changed

- SIGHUP support was removed. You should use SIGUSR1 instead. See *Schema Reloading*.
- `server-host` default of `127.0.0.1` was changed to `!4`. See *server-host*.

### 5.2.3 Thanks

This release is sponsored by:



PostgREST for the backend.  
Retool for the frontend.

- Daniel Babiak
- Evans Fernandes

- Tsingson Qin
- Michel Pelletier
- Jay Hannah
- Robert Stolarz
- Kofi Gumbs
- Nicholas DiBiase
- Christopher Reid
- Nathan Bouscal

If you like to join them please consider [supporting PostgREST development](#).




Are you new to PostgREST? This is the place to start!

## 6.1 Tutorial 0 - Get it Running

**author** begriffs

Welcome to PostgREST! In this pre-tutorial we're going to get things running so you can create your first simple API.

PostgREST is a standalone web server which turns a PostgreSQL database into a RESTful API. It serves an API that is customized based on the structure of the underlying database.



To make an API we'll simply be building a database. All the endpoints and permissions come from database objects like tables, views, roles, and stored procedures. These tutorials will cover a number of common scenarios and how to model them in the database.

By the end of this tutorial you'll have a working database, PostgREST server, and a simple single-user todo list API.

### 6.1.1 Step 1. Relax, we'll help

As you begin the tutorial, pop open the project [chat room](#) in another tab. There are a nice group of people active in the project and we'll help you out if you get stuck.

### 6.1.2 Step 2. Install PostgreSQL

You'll need a modern copy of the database running on your system, either natively or in a Docker instance. We require PostgreSQL 9.3 or greater, but recommend at least 9.5 for row-level security features that we'll use in future tutorials.

If you're already familiar with using PostgreSQL and have it installed on your system you can use the existing installation. For this tutorial we'll describe how to use the database in Docker because database configuration is otherwise too complicated for a simple tutorial.

If Docker is not installed, you can get it [here](#). Next, let's pull and start the database image:

```
sudo docker run --name tutorial -p 5433:5432 \  
    -e POSTGRES_PASSWORD=mysecretpassword \  
    -d postgres
```

This will run the Docker instance as a daemon and expose port 5433 to the host system so that it looks like an ordinary PostgreSQL server to the rest of the system.

### 6.1.3 Step 3. Install PostgREST

PostgREST is distributed as a single binary, with versions compiled for major distributions of Linux/BSD/Windows. Visit the [latest release](#) for a list of downloads. In the event that your platform is not among those already pre-built, see [Build from Source](#) for instructions how to build it yourself. Also let us know to add your platform in the next release.

The pre-built binaries for download are `.tar.xz` compressed files (except Windows which is a zip file). To extract the binary, go into the terminal and run

```
# download from https://github.com/PostgREST/postgrest/releases/latest  
  
tar xfJ postgrest-<version>-<platform>.tar.xz
```

The result will be a file named simply `postgrest` (or `postgrest.exe` on Windows). At this point try running it with

```
./postgrest
```

If everything is working correctly it will print out its version and information about configuration. You can continue to run this binary from where you downloaded it, or copy it to a system directory like `/usr/local/bin` on Linux so that you will be able to run it from any directory.

---

**Note:** PostgREST requires `libpq`, the PostgreSQL C library, to be installed on your system. Without the library you'll get an error like "error while loading shared libraries: libpq.so.5." Here's how to fix it:

---

### 6.1.4 Step 4. Create Database for API

Connect to the SQL console (`psql`) inside the container. To do so, run this from your command line:

```
sudo docker exec -it tutorial psql -U postgres
```

You should see the `psql` command prompt:

```
psql (9.6.3)
Type "help" for help.

postgres=#
```

The first thing we'll do is create a `named schema` for the database objects which will be exposed in the API. We can choose any name we like, so how about "api." Execute this and the other SQL statements inside the `psql` prompt you started.

```
create schema api;
```

Our API will have one endpoint, `/todos`, which will come from a table.

```
create table api.todos (
  id serial primary key,
  done boolean not null default false,
  task text not null,
  due timestamptz
);

insert into api.todos (task) values
('finish tutorial 0'), ('pat self on back');
```

Next make a role to use for anonymous web requests. When a request comes in, PostgREST will switch into this role in the database to run queries.

```
create role web_anon nologin;

grant usage on schema api to web_anon;
grant select on api.todos to web_anon;
```

The `web_anon` role has permission to access things in the `api` schema, and to read rows in the `todos` table.

It's a good practice to create a dedicated role for connecting to the database, instead of using the highly privileged `postgres` role. So we'll do that, name the role `authenticator` and also grant him the ability to switch to the `web_anon` role:

```
create role authenticator noinherit login password 'mysecretpassword';
grant web_anon to authenticator;
```

Now quit out of `psql`; it's time to start the API!

```
\q
```

## 6.1.5 Step 5. Run PostgREST

PostgREST uses a configuration file to tell it how to connect to the database. Create a file `tutorial.conf` with this inside:

```
db-uri = "postgres://authenticator:mysecretpassword@localhost:5433/postgres"
db-schema = "api"
db-anon-role = "web_anon"
```

The configuration file has other *options*, but this is all we need. Now run the server:

```
./postgrest tutorial.conf
```

You should see

```
Listening on port 3000
Attempting to connect to the database...
Connection successful
```

It's now ready to serve web requests. There are many nice graphical API exploration tools you can use, but for this tutorial we'll use `curl` because it's likely to be installed on your system already. Open a new terminal (leaving the one open that PostgREST is running inside). Try doing an HTTP request for the todos.

```
curl http://localhost:3000/todos
```

The API replies:

```
[
  {
    "id": 1,
    "done": false,
    "task": "finish tutorial 0",
    "due": null
  },
  {
    "id": 2,
    "done": false,
    "task": "pat self on back",
    "due": null
  }
]
```

With the current role permissions, anonymous requests have read-only access to the `todos` table. If we try to add a new todo we are not able.

```
curl http://localhost:3000/todos -X POST \
  -H "Content-Type: application/json" \
  -d '{"task": "do bad thing"}'
```

Response is 401 Unauthorized:

```
{
  "hint": null,
  "details": null,
  "code": "42501",
  "message": "permission denied for relation todos"
}
```

There we have it, a basic API on top of the database! In the next tutorials we will see how to extend the example with more sophisticated user access controls, and more tables and queries.

Now that you have PostgREST running, try the next tutorial, *Tutorial 1 - The Golden Key*

## 6.2 Tutorial 1 - The Golden Key

author `begriffs`

In *Tutorial 0 - Get it Running* we created a read-only API with a single endpoint to list todos. There are many directions we can go to make this API more interesting, but one good place to start would be allowing some users to change data in addition to reading it.

### 6.2.1 Step 1. Add a Trusted User

The previous tutorial created a `web_anon` role in the database with which to execute anonymous web requests. Let's make a role called `todo_user` for users who authenticate with the API. This role will have the authority to do anything to the todo list.

```
-- run this in psql using the database created
-- in the previous tutorial

create role todo_user nologin;
grant todo_user to authenticator;

grant usage on schema api to todo_user;
grant all on api.todos to todo_user;
grant usage, select on sequence api.todos_id_seq to todo_user;
```

### 6.2.2 Step 2. Make a Secret

Clients authenticate with the API using JSON Web Tokens. These are JSON objects which are cryptographically signed using a password known to only us and the server. Because clients do not know the password, they cannot tamper with the contents of their tokens. PostgREST will detect counterfeit tokens and will reject them.

Let's create a password and provide it to PostgREST. Think of a nice long one, or use a tool to generate it. **Your password must be at least 32 characters long.**

**Note:** Unix tools can generate a nice password for you:

```
# Allow "tr" to process non-utf8 byte sequences
export LC_CTYPE=C

# read random bytes and keep only alphanumerics
< /dev/urandom tr -dc A-Za-z0-9 | head -c32
```

Open the `tutorial.conf` (created in the previous tutorial) and add a line with the password:

```
# PASSWORD MUST BE AT LEAST 32 CHARS LONG
# add this line to tutorial.conf:

jwt-secret = "<the password you made>"
```

If the PostgREST server is still running from the previous tutorial, restart it to load the updated configuration file.

### 6.2.3 Step 3. Sign a Token

Ordinarily your own code in the database or in another server will create and sign authentication tokens, but for this tutorial we will make one “by hand.” Go to [jwt.io](https://jwt.io) and fill in the fields like this:

**Remember to fill in the password you generated rather than the word “secret”.** After you have filled in the password and payload, the encoded data on the left will update. Copy the encoded token.

**Encoded** PASTE A TOKEN HERE

```
eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJyY2x1IjoiaWoiG9kb191c2VyIn0.IF9PKgLZ_XA70Uz5vm00xrqYTCQdXxBA00z4uK81BqM
```

**Decoded** EDIT THE PAYLOAD AND SECRET (ONLY HS256 SUPPORTED)

HEADER: ALGORITHM & TOKEN TYPE

```
{
  "alg": "HS256",
  "typ": "JWT"
}
```

PAYLOAD: DATA

```
{
  "role": "todo_user"
}
```

VERIFY SIGNATURE

```
HMACSHA256(
  base64UrlEncode(header) + "." +
  base64UrlEncode(payload),
  secret
) secret base64 encoded
```

2. Enter this JSON

3. Copy the resulting token

1. Enter password

Fig. 1: How to create a token at <https://jwt.io>

**Note:** While the token may look well obscured, it's easy to reverse engineer the payload. The token is merely signed, not encrypted, so don't put things inside that you don't want a determined client to see.

### 6.2.4 Step 4. Make a Request

Back in the terminal, let's use `curl` to add a todo. The request will include an HTTP header containing the authentication token.

```
export TOKEN="<paste token here>"

curl http://localhost:3000/todos -X POST \
  -H "Authorization: Bearer $TOKEN" \
  -H "Content-Type: application/json" \
  -d '{"task": "learn how to auth"}'
```

And now we have completed all three items in our todo list, so let's set done to true for them all with a PATCH request.

```
curl http://localhost:3000/todos -X PATCH \
  -H "Authorization: Bearer $TOKEN" \
  -H "Content-Type: application/json" \
  -d '{"done": true}'
```

A request for the todos shows three of them, and all completed.

```
curl http://localhost:3000/todos
```

```
[
  {
    "id": 1,
    "done": true,
    "task": "finish tutorial 0",
    "due": null
  },
  {
    "id": 2,
    "done": true,
    "task": "pat self on back",
    "due": null
  },
  {
    "id": 3,
    "done": true,
    "task": "learn how to auth",
    "due": null
  }
]
```

### 6.2.5 Step 4. Add Expiration

Currently our authentication token is valid for all eternity. The server, as long as it continues using the same JWT password, will honor the token.

It's better policy to include an expiration timestamp for tokens using the `exp` claim. This is one of two JWT claims that PostgREST treats specially.

Claim	Interpretation
<code>role</code>	The database role under which to execute SQL for API request
<code>exp</code>	Expiration timestamp for token, expressed in "Unix epoch time"

**Note:** Epoch time is defined as the number of seconds that have elapsed since 00:00:00 Coordinated Universal Time (UTC), January 1st 1970, minus the number of leap seconds that have taken place since then.

To observe expiration in action, we'll add an `exp` claim of five minutes in the future to our previous token. First find the epoch value of five minutes from now. In `psql` run this:

```
select extract(epoch from now() + '5 minutes'::interval) :: integer;
```

Go back to `jwt.io` and change the payload to

```
{
  "role": "todo_user",
  "exp": 123456789
}
```

**NOTE:** Don't forget to change the dummy epoch value `123456789` in the snippet above to the epoch value returned by the `psql` command.

Copy the updated token as before, and save it as a new environment variable.

```
export NEW_TOKEN="<paste new token>"
```

Try issuing this request in curl before and after the expiration time:

```
curl http://localhost:3000/todos \  
-H "Authorization: Bearer $NEW_TOKEN"
```

After expiration, the API returns HTTP 401 Unauthorized:

```
{"message": "JWT expired"}
```

### 6.2.6 Bonus Topic: Immediate Revocation

Even with token expiration there are times when you may want to immediately revoke access for a specific token. For instance, suppose you learn that a disgruntled employee is up to no good and his token is still valid.

To revoke a specific token we need a way to tell it apart from others. Let's add a custom `email` claim that matches the email of the client issued the token.

Go ahead and make a new token with the payload

```
{  
  "role": "todo_user",  
  "email": "disgruntled@mycompany.com"  
}
```

Save it to an environment variable:

```
export WAYWARD_TOKEN="<paste new token>"
```

PostgreSQL allows us to specify a stored procedure to run during attempted authentication. The function can do whatever it likes, including raising an exception to terminate the request.

First make a new schema and add the function:

```
create schema auth;  
grant usage on schema auth to web_anon, todo_user;  
  
create or replace function auth.check_token() returns void  
language plpgsql  
as $$  
begin  
  if current_setting('request.jwt.claim.email', true) =  
    'disgruntled@mycompany.com' then  
    raise insufficient_privilege  
      using hint = 'Nope, we are on to you';  
  end if;  
end  
$$;
```

Next update `tutorial.conf` and specify the new function:

```
# add this line to tutorial.conf  
pre-request = "auth.check_token"
```

Restart PostgREST for the change to take effect. Next try making a request with our original token and then with the revoked one.

```
# this request still works

curl http://localhost:3000/todos -X PATCH \
  -H "Authorization: Bearer $TOKEN" \
  -H "Content-Type: application/json" \
  -d '{"done": true}'

# this one is rejected

curl http://localhost:3000/todos -X PATCH \
  -H "Authorization: Bearer $WAYWARD_TOKEN" \
  -H "Content-Type: application/json" \
  -d '{"task": "AAAHHHH!", "done": false}'
```

The server responds with 403 Forbidden:

```
{
  "hint": "Nope, we are on to you",
  "details": null,
  "code": "42501",
  "message": "insufficient_privilege"
}
```

- [Tutorial 0 - Get it Running](#)
- [Tutorial 1 - The Golden Key](#)

Also have a look at [Installation](#).



Technical references for PostgREST's functionality.

## 7.1 Tables and Views

All views and tables in the exposed schema and accessible by the active database role for a request are available for querying. They are exposed in one-level deep routes. For instance the full contents of a table *people* is returned at

```
GET /people HTTP/1.1
```

There are no deeply/nested/routes. Each route provides OPTIONS, GET, POST, PATCH, and DELETE verbs depending entirely on database permissions.

---

**Note:** Why not provide nested routes? Many APIs allow nesting to retrieve related information, such as `/films/1/director`. We offer a more flexible mechanism (inspired by GraphQL) to embed related information. It can handle one-to-many and many-to-many relationships. This is covered in the section about *Resource Embedding*.

---

### 7.1.1 Horizontal Filtering (Rows)

You can filter result rows by adding conditions on columns, each condition a query string parameter. For instance, to return people aged under 13 years old:

```
GET /people?age<13 HTTP/1.1
```

Multiple parameters can be logically conjoined by:

```
GET /people?age>=18&student=is.true HTTP/1.1
```

Multiple parameters can be logically disjoined by:

```
GET /people?or=(age.gte.14,age.lte.18) HTTP/1.1
```

Complex logic can also be applied:

```
GET /people?and=(grade.gte.90,student.is.true,or(age.gte.14,age.is.null)) HTTP/1.1
```

## Operators

These operators are available:

Ab- brevi- ation	In Post- greSQL	Meaning
eq	=	equals
gt	>	greater than
gte	>=	greater than or equal
lt	<	less than
lte	<=	less than or equal
neq	<> or !=	not equal
like	LIKE	LIKE operator (use * in place of %)
ilike	ILIKE	ILIKE operator (use * in place of %)
in	IN	one of a list of values, e.g. ?a=in.(1,2,3) – also supports commas in quoted strings like ?a=in.("hi,there","yes,you")
is	IS	checking for exact equality (null,true,false)
fts	@@	<i>Full-Text Search</i> using to_tsquery
plfts	@@	<i>Full-Text Search</i> using plainto_tsquery
phfts	@@	<i>Full-Text Search</i> using phraseto_tsquery
wfts	@@	<i>Full-Text Search</i> using websearch_to_tsquery
cs	@>	contains e.g. ?tags=cs.{example,new}
cd	<@	contained in e.g. ?values=cd.{1,2,3}
ov	&&	overlap (have points in common), e.g. ?period=ov.[2017-01-01,2017-06-30] – also supports array types, use curly braces instead of square brackets e.g. :code: ?arr=ov.{1,3}
sl	<<	strictly left of, e.g. ?range=sl.(1,10)
sr	>>	strictly right of
nrx	&<	does not extend to the right of, e.g. ?range=nrx.(1,10)
nxl	&>	does not extend to the left of
adj	- -	is adjacent to, e.g. ?range=adj.(1,10)
not	NOT	negates another operator, see below

To negate any operator, prefix it with not like ?a=not.eq.2 or ?not.and=(a.gte.0,a.lte.100) .

For more complicated filters you will have to create a new view in the database, or use a stored procedure. For instance, here’s a view to show “today’s stories” including possibly older pinned stories:

```
CREATE VIEW fresh_stories AS
SELECT *
FROM stories
WHERE pinned = true
OR published > now() - interval '1 day'
ORDER BY pinned DESC, published DESC;
```

The view will provide a new endpoint:

```
GET /fresh_stories HTTP/1.1
```

**Important:** Views are invoked with the privileges of the view owner, much like stored procedures with the SECURITY DEFINER option. When created by a SUPERUSER role, all row-level security will be bypassed unless a different, non-SUPERUSER owner is specified.

```
-- Workaround:
-- non-SUPERUSER role to be used as the owner of the views
CREATE ROLE api_views_owner;
-- alter the view owner so RLS can work normally
ALTER VIEW sample_view OWNER TO api_views_owner;
```

## Full-Text Search

The `fts` filter mentioned above has a number of options to support flexible textual queries, namely the choice of plain vs phrase search and the language used for stemming. Suppose that `tsearch` is a table with column `my_tsv`, of type `tsvector`. The following examples illustrate the possibilities.

```
GET /tsearch?my_tsv=fts(french).amusant HTTP/1.1
```

```
GET /tsearch?my_tsv=plfts.The%20Fat%20Cats HTTP/1.1
```

```
GET /tsearch?my_tsv=not.phfts(english).The%20Fat%20Cats HTTP/1.1
```

```
GET /tsearch?my_tsv=not.wfts(french).amusant HTTP/1.1
```

Using phrase search mode requires PostgreSQL of version at least 9.6 and will raise an error in earlier versions of the database.

Using `websearch_to_tsquery` requires PostgreSQL of version at least 11.0 and will raise an error in earlier versions of the database.

### 7.1.2 Vertical Filtering (Columns)

When certain columns are wide (such as those holding binary data), it is more efficient for the server to withhold them in a response. The client can specify which columns are required using the `select` parameter.

```
GET /people?select=first_name,age HTTP/1.1

[
  {"first_name": "John", "age": 30},
  {"first_name": "Jane", "age": 20}
]
```

The default is `*`, meaning all columns. This value will become more important below in *Resource Embedding*.

### Renaming Columns

You can rename the columns by prefixing them with an alias followed by the colon `:` operator.

```
GET /people?select=fullName:full_name,birthDate:birth_date HTTP/1.1
```

```
[
  {"fullName": "John Doe", "birthDate": "04/25/1988"},
  {"fullName": "Jane Doe", "birthDate": "01/12/1998"}
]
```

## Casting Columns

Casting the columns is possible by suffixing them with the double colon :: plus the desired type.

```
GET /people?select=full_name,salary::text HTTP/1.1
```

```
[
  {"full_name": "John Doe", "salary": "90000.00"},
  {"full_name": "Jane Doe", "salary": "120000.00"}
]
```

## JSON Columns

You can specify a path for a json or jsonb column using the arrow operators(-> or ->>) as per the PostgreSQL docs.

```
GET /people?select=id,json_data->>blood_type,json_data->phones HTTP/1.1
```

```
[
  { "id": 1, "blood_type": "A+", "phones": [{"country_code": "61", "number": "917-929-
  ↪5745"}] },
  { "id": 2, "blood_type": "O+", "phones": [{"country_code": "43", "number": "512-446-
  ↪4988"}, {"country_code": "43", "number": "213-891-5979"}] }
]
```

```
GET /people?select=id,json_data->phones->0->>number HTTP/1.1
```

```
[
  { "id": 1, "number": "917-929-5745"},
  { "id": 2, "number": "512-446-4988"}
]
```

## Computed Columns

Filters may be applied to computed columns(a.k.a. **virtual columns**) as well as actual table/view columns, even though the computed columns will not appear in the output. For example, to search first and last names at once we can create a computed column that will not appear in the output but can be used in a filter:

```
CREATE TABLE people (
  fname text,
  lname text
);

CREATE FUNCTION full_name(people) RETURNS text AS $$
  SELECT $1.fname || ' ' || $1.lname;
```

(continues on next page)

(continued from previous page)

```

$$ LANGUAGE SQL;

-- (optional) add an index to speed up anticipated query
CREATE INDEX people_full_name_idx ON people
  USING GIN (to_tsvector('english', full_name(people)));

```

A full-text search on the computed column:

```
GET /people?full_name=fts.Beckett HTTP/1.1
```

As mentioned, computed columns do not appear in the output by default. However you can include them by listing them in the vertical filtering `select` param:

```
GET /people?select=*,full_name HTTP/1.1
```

---

**Important:** Computed columns must be created under the *exposed schema* to be used in this way.

---

### 7.1.3 Unicode support

PostgREST supports unicode in schemas, tables, columns and values. To access a table with unicode name, use percent encoding.

To request this:

```
GET / HTTP/1.1
```

Do this:

```
GET /%D9%85%D9%88%D8%A7%D8%B1%D8%AF HTTP/1.1
```

### Table / Columns with spaces

You can request table/columns with spaces in them by percent encoding the spaces with `%20`:

```
GET /Order%20Items?Unit%20Price=1t.200 HTTP/1.1
```

### Reserved characters

If filters include PostgREST reserved characters(, , . , : , ( ) ) you'll have to surround them in percent encoded double quotes `%22` for correct processing.

Here `Hebdon, John` and `Williams, Mary` are values.

```
GET /employees?name=in. (%22Hebdon, John%22, %22Williams, Mary%22) HTTP/1.1
```

Here `information.cpe` is a column name.

```
GET /vulnerabilities?%22information.cpe%22=like.*MS* HTTP/1.1
```

**Note:** Some http libraries might encode URLs automatically(e.g. `axios`). In these cases you should use double quotes `" "` directly instead of `%22`.

---

### 7.1.4 Ordering

The reserved word `order` reorders the response rows. It uses a comma-separated list of columns and directions:

```
GET /people?order=age.desc,height.asc HTTP/1.1
```

If no direction is specified it defaults to ascending order:

```
GET /people?order=age HTTP/1.1
```

If you care where nulls are sorted, add `nullsfirst` or `nullslast`:

```
GET /people?order=age.nullsfirst HTTP/1.1
```

```
GET /people?order=age.desc.nullslast HTTP/1.1
```

You can also use *Computed Columns* to order the results, even though the computed columns will not appear in the output.

### 7.1.5 Limits and Pagination

PostgreSQL uses HTTP range headers to describe the size of results. Every response contains the current range and, if requested, the total number of results:

```
HTTP/1.1 200 OK
Range-Unit: items
Content-Range: 0-14/*
```

Here items zero through fourteen are returned. This information is available in every response and can help you render pagination controls on the client. This is an RFC7233-compliant solution that keeps the response JSON cleaner.

There are two ways to apply a limit and offset rows: through request headers or query params. When using headers you specify the range of rows desired. This request gets the first twenty people.

```
GET /people HTTP/1.1
Range-Unit: items
Range: 0-19
```

Note that the server may respond with fewer if unable to meet your request:

```
HTTP/1.1 200 OK
Range-Unit: items
Content-Range: 0-17/*
```

You may also request open-ended ranges for an offset with no limit, e.g. `Range: 10-`.

The other way to request a limit or offset is with query parameters. For example

```
GET /people?limit=15&offset=30 HTTP/1.1
```

This method is also useful for embedded resources, which we will cover in another section. The server always responds with range headers even if you use query parameters to limit the query.

In order to obtain the total size of the table or view (such as when rendering the last page link in a pagination control), specify your preference in a request header:

```
GET /bigtable HTTP/1.1
Range-Unit: items
Range: 0-24
Prefer: count=exact
```

Note that the larger the table the slower this query runs in the database. The server will respond with the selected range and total

```
HTTP/1.1 206 Partial Content
Range-Unit: items
Content-Range: 0-24/3573458
```

## 7.1.6 Response Format

PostgREST uses proper HTTP content negotiation ([RFC7231](#)) to deliver the desired representation of a resource. That is to say the same API endpoint can respond in different formats like JSON or CSV depending on the client request.

Use the Accept request header to specify the acceptable format (or formats) for the response:

```
GET /people HTTP/1.1
Accept: application/json
```

The current possibilities are

- \*/\*
- text/csv
- application/json
- application/openapi+json
- application/octet-stream

The server will default to JSON for API endpoints and OpenAPI on the root.

## 7.1.7 Singular or Plural

By default PostgREST returns all JSON results in an array, even when there is only one item. For example, requesting `/items?id=eq.1` returns

```
[
  { "id": 1 }
]
```

This can be inconvenient for client code. To return the first result as an object unenclosed by an array, specify `vnd.pgrst.object` as part of the Accept header

```
GET /items?id=eq.1 HTTP/1.1
Accept: application/vnd.pgrst.object+json
```

This returns

```
{ "id": 1 }
```

When a singular response is requested but no entries are found, the server responds with an error message and 406 Not Acceptable status code rather than the usual empty array and 200 status:

```
{
  "message": "JSON object requested, multiple (or no) rows returned",
  "details": "Results contain 0 rows, application/vnd.pgrst.object+json requires 1 row
  ↪"
}
```

---

**Note:** Many APIs distinguish plural and singular resources using a special nested URL convention e.g. */stories* vs */stories/1*. Why do we use */stories?id=eq.1*? The answer is because a singular resource is (for us) a row determined by a primary key, and primary keys can be compound (meaning defined across more than one column). The more familiar nested urls consider only a degenerate case of simple and overwhelmingly numeric primary keys. These so-called artificial keys are often introduced automatically by Object Relational Mapping libraries.

Admittedly PostgREST could detect when there is an equality condition holding on all columns constituting the primary key and automatically convert to singular. However this could lead to a surprising change of format that breaks unwary client code just by filtering on an extra column. Instead we allow manually specifying singular vs plural to decouple that choice from the URL format.

---


## 7.2 Resource Embedding

In addition to providing RESTful routes for each table and view, PostgREST allows related resources to be included together in a single API call. This reduces the need for multiple API requests. The server uses **foreign keys** to determine which tables and views can be returned together. For example, consider a database of films and their awards:

---

**Important:** PostgREST needs **FOREIGN KEY constraints** to be able to do Resource Embedding.

---



As seen above in *Vertical Filtering (Columns)* we can request the titles of all films like this:

```
GET /films?select=title HTTP/1.1
```

This might return something like

```
[
  { "title": "Workers Leaving The Lumière Factory In Lyon" },
  { "title": "The Dickson Experimental Sound Film" },
  { "title": "The Haunted Castle" }
]
```

However because a foreign key constraint exists between Films and Directors, we can request this information be included:

```
GET /films?select=title,directors(id,last_name) HTTP/1.1
```

Which would return

```
[
  { "title": "Workers Leaving The Lumière Factory In Lyon",
    "directors": {
      "id": 2,
      "last_name": "Lumière"
    }
  },
  { "title": "The Dickson Experimental Sound Film",
    "directors": {
      "id": 1,
      "last_name": "Dickson"
    }
  },
  { "title": "The Haunted Castle",
    "directors": {
      "id": 3,
      "last_name": "Méliès"
    }
  }
]
```

In this example, since the relationship is a forward relationship, there is only one director associated with a film. As the table name is plural it might be preferable for it to be singular instead. An table name alias can accomplish this:

```
GET /films?select=title,director:directors(id,last_name) HTTP/1.1
```

**Important:** Whenever FOREIGN KEY constraints change in the database schema you must refresh PostgreSQL's schema cache for Resource Embedding to work properly. See the section *Schema Reloading*.

## 7.2.1 Embedding through join tables

PostgreSQL can also detect relationships going through join tables. Thus you can request the Actors for Films (which in this case finds the information through Roles). You can also reverse the direction of inclusion, asking for all Directors with each including the list of their Films:

```
GET /directors?select=films(title,year) HTTP/1.1
```

## 7.2.2 Embedded Filters

Embedded resources can be shaped similarly to their top-level counterparts. To do so, prefix the query parameters with the name of the embedded resource. For instance, to order the actors in each film:

```
GET /films?select=*,actors(*)&actors.order=last_name,first_name HTTP/1.1
```

This sorts the list of actors in each film but does *not* change the order of the films themselves. To filter the roles returned with each film:

```
GET /films?select=*,roles(*)&roles.character=in.(Chico,Harpo,Groucho) HTTP/1.1
```

Once again, this restricts the roles included to certain characters but does not filter the films in any way. Films without any of those characters would be included along with empty character lists.

An `OR` filter can be used for a similar operation:

```
GET /films?select=*,roles(*)&roles.or=(character.eq.Gummo,character.eq.Zeppo) HTTP/1.1
```

Limit and offset operations are possible:

```
GET /films?select=*,actors(*)&actors.limit=10&actors.offset=2 HTTP/1.1
```

Embedded resources can be aliased and filters can be applied on these aliases:

```
GET /films?select=*,90_comps:competitions(name),91_comps:competitions(name)&90_comps.
↪year=eq.1990&91_comps.year=eq.1991 HTTP/1.1
```

### 7.2.3 Embedding Views

Embedding a view is possible if the view contains columns that have **foreign keys** defined in their source tables.

As an example, let's create a view called `nominations_view` based on the `nominations` table.

```
CREATE VIEW nominations_view AS
SELECT
  rank
  , competition_id
  , film_id
FROM
  nominations;
```

Since it contains `competition_id` and `film_id`—and each one has a **foreign key** defined in its source table—we can embed `competitions` and `films`:

```
GET /nominations_view?select=rank,competitions(name,year),films(title)&rank=eq.5 HTTP/
↪1.1
```

It's also possible to embed [Materialized Views](#).

**Warning:** Is not guaranteed that all kinds of views will be embeddable. In particular, views that contain UNIONS will not be made embeddable.

Why? PostgREST detects source table foreign keys in the view by querying and parsing `pg_rewrite`. This may fail depending on the complexity of the view.

[Report an issue](#) if your view is not made embeddable so we can keep continue improving foreign key detection.

In the future we'll include include a way to manually specify views source foreign keys to address this limitation.

**Important:** If view definitions change you must refresh PostgREST's schema cache for this to work properly. See the section [Schema Reloading](#).

## 7.3 Custom Queries

The PostgREST URL grammar limits the kinds of queries clients can perform. It prevents arbitrary, potentially poorly constructed and slow client queries. It's good for quality of service, but means database administrators must create custom views and stored procedures to provide richer endpoints. The most common causes for custom endpoints are

- Table unions
- More complicated joins than those provided by *Resource Embedding*
- Geo-spatial queries that require an argument, like “points near (lat,lon)”
- More sophisticated full-text search than a simple use of the `fts` filter

## 7.4 Stored Procedures

Every stored procedure in the API-exposed database schema is accessible under the `/rpc` prefix. The API endpoint supports POST (and in some cases GET) to execute the function.

```
POST /rpc/function_name HTTP/1.1
```

Such functions can perform any operations allowed by PostgreSQL (read data, modify data, and even DDL operations).

To supply arguments in an API call, include a JSON object in the request payload and each key/value of the object will become an argument.

For instance, assume we have created this function in the database.

```
CREATE FUNCTION add_them(a integer, b integer)
RETURNS integer AS $$
  SELECT a + b;
$$ LANGUAGE SQL IMMUTABLE;
```

The client can call it by posting an object like

```
POST /rpc/add_them HTTP/1.1

{ "a": 1, "b": 2 }

3
```

---

**Important:** Whenever you create or change a function you must refresh PostgREST’s schema cache. See the section *Schema Reloading*.

If the schema cache is not refreshed, PostgREST will assume `text` as the default type for function arguments. This could lead to getting error responses like:

```
{
  "hint": "No function matches the given name and argument types. You might need to add
  ↳ explicit type casts.",
  "details": null,
  "code": "42883",
  "message": "function test.add_them(a => text, b => text) does not exist"
}
```

You can also call a function that takes a single parameter of type `json` by sending the header `Prefer: params=single-object` with your request. That way the JSON request body will be used as the single argument.

```
CREATE FUNCTION mult_them(param json) RETURNS int AS $$
  SELECT (param->>'x')::int * (param->>'y')::int
$$ LANGUAGE SQL;
```

```
POST /rpc/mult_them HTTP/1.1
Prefer: params=single-object

{ "x": 4, "y": 2 }

8
```

Procedures must be declared with named parameters, procedures declared like:

```
CREATE FUNCTION non_named_args(integer, text, integer) ...
```

Can not be called with PostgREST, since we use `named notation` internally.

Note that PostgreSQL converts identifier names to lowercase unless you quote them like:

```
CREATE FUNCTION "someFunc"("someParam" text) ...
```

PostgreSQL has four procedural languages that are part of the core distribution: PL/pgSQL, PL/Tcl, PL/Perl, and PL/Python. There are many other procedural languages distributed as additional extensions. Also, plain SQL can be used to write functions (as shown in the example above).

**Note:** For versions prior to PostgreSQL 10, to pass a PostgreSQL native array you need to quote it as a string:

```
POST /rpc/native_array_func HTTP/1.1

{ "arg": "{1,2,3}" }
```

In these versions we recommend using function arguments of type json to accept arrays from the client:

```
POST /rpc/json_array_func HTTP/1.1

{ "arg": [1,2,3] }
```

Starting from PostgreSQL 10, a json array from the client gets mapped normally to a PostgreSQL native array.

**Note:** Why the `/rpc` prefix? One reason is to avoid name collisions between views and procedures. It also helps emphasize to API consumers that these functions are not normal restful things. The functions can have arbitrary and surprising behavior, not the standard “post creates a resource” thing that users expect from the other routes.

### 7.4.1 Immutable and stable functions

Procedures in PostgreSQL marked with `stable` or `immutable` `volatility` can only read, not modify, the database and PostgREST executes them in a read-only transaction compatible for read-replicas. Stable and immutable functions can be called with the HTTP GET verb if desired.

**Note:** The volatility marker is a promise about the behavior of the function. PostgreSQL will let you mark a function that modifies the database as `immutable/stable` without failure. However the function will fail when called through PostgREST since it executes it in a read-only transaction.

Because `add_them` was declared IMMUTABLE, we can alternately call the function with a GET request:

```
GET /rpc/add_them?a=1&b=2 HTTP/1.1
```

The function parameter names match the JSON object keys in the POST case, for the GET case they match the query parameters `?a=1&b=2`.

## 7.4.2 Scalar functions

PostgreSQL will detect if the function is scalar or table-valued and will shape the response format accordingly:

```
GET /rpc/add_them?a=1&b=2 HTTP/1.1
```

```
3
```

```
GET /rpc/best_films_2017 HTTP/1.1
```

```
[
  { "title": "Okja", "rating": 7.4},
  { "title": "Call me by your name", "rating": 8},
  { "title": "Blade Runner 2049", "rating": 8.1}
]
```

## 7.4.3 Bulk Call

It's possible to call a function in a bulk way, analogously to *Bulk Insert*.

```
POST /rpc/add_them HTTP/1.1
Content-Type: application/json
```

```
[
  { "a": 1, "b": 2},
  { "a": 3, "b": 4}
]
```

Result:

```
[ 3, 7 ]
```

It's also possible to *Specify Columns* on functions calls.

## 7.4.4 Function filters

A function that returns a table type response can be shaped using the same filters as the ones used for tables and views:

```
CREATE FUNCTION best_films_2017() RETURNS SETOF films ..
```

```
GET /rpc/best_films_2017?select=title,director:directors(*) HTTP/1.1
```

```
GET /rpc/best_films_2017?rating=gt.8&order=title.desc HTTP/1.1
```

## 7.4.5 Function privileges

By default, a function is executed with the privileges of the user who calls it. This means that the user has to have all permissions to do the operations the procedure performs.

Another option is to define the function with the `SECURITY DEFINER` option. Then only one permission check will take place, the permission to call the function, and the operations in the function will have the authority of the user who owns the function itself. See [PostgreSQL documentation](#) for more details.

**Warning:** Unlike tables/views, functions privileges work as a blacklist, so they're executable for all the roles by default. You can workaround this by revoking the PUBLIC privileges of the function and then granting privileges to specific roles:

```
REVOKE ALL PRIVILEGES ON FUNCTION private_func() FROM PUBLIC;
GRANT EXECUTE ON FUNCTION private_func() TO a_role;
```

Also to avoid doing REVOKE on every function you can enable this behavior by default with:

```
ALTER DEFAULT PRIVILEGES REVOKE EXECUTE ON FUNCTIONS FROM PUBLIC;
```

See [PostgreSQL alter default privileges](#) for more details.

## 7.4.6 Overloaded functions

You can call overloaded functions with different number of arguments.

```
CREATE FUNCTION rental_duration(customer_id integer) ..
CREATE FUNCTION rental_duration(customer_id integer, from_date date) ..
```

```
GET /rpc/rental_duration?customer_id=232 HTTP/1.1
```

```
GET /rpc/rental_duration?customer_id=232&from_date=2018-07-01 HTTP/1.1
```

## 7.4.7 Accessing Request Headers, Cookies and JWT claims

Stored procedures can access request headers, cookies and jwt claims by reading GUC variables set by PostgREST per request. They are named `request.header.XYZ`, `request.cookie.XYZ` and `request.jwt.claim.XYZ`.

```
-- To read the value of the Origin request header:
SELECT current_setting('request.header.origin', true);
-- To read the value of sessionId in a cookie:
SELECT current_setting('request.cookie.sessionId', true);
-- To read the value of the email claim in a jwt:
SELECT current_setting('request.jwt.claim.email', true);
```

**Note:** `request.jwt.claim.role` defaults to the value of `db-anon-role`.

## 7.4.8 Setting Response Headers

PostgreSQL reads the `response.headers` SQL variable to add extra headers to the HTTP response. Stored procedures can modify this variable. For instance, this statement would add caching headers to the response:

```
-- tell client to cache response for two days

SET LOCAL "response.headers" =
  '[{"Cache-Control": "public"}, {"Cache-Control": "max-age=259200"}]';
```

Notice that the variable should be set to an *array* of single-key objects rather than a single multiple-key object. This is because headers such as `Cache-Control` or `Set-Cookie` need to be repeated when setting multiple values and an object would not allow the repeated key.

## 7.4.9 Errors and HTTP Status Codes

Stored procedures can return non-200 HTTP status codes by raising SQL exceptions. For instance, here's a saucy function that always responds with an error:

```
CREATE OR REPLACE FUNCTION just_fail() RETURNS void
LANGUAGE plpgsql
AS $$
BEGIN
  RAISE EXCEPTION 'I refuse!'
  USING DETAIL = 'Pretty simple',
  HINT = 'There is nothing you can do.';
END
$$;
```

Calling the function returns HTTP 400 with the body

```
{
  "message": "I refuse!",
  "details": "Pretty simple",
  "hint": "There is nothing you can do.",
  "code": "P0001"
}
```

One way to customize the HTTP status code is by raising particular exceptions according to the PostgreSQL *error to status code mapping*. For example, `RAISE insufficient_privilege` will respond with HTTP 401/403 as appropriate.

For even greater control of the HTTP status code, raise an exception of the `PTxyz` type. For instance to respond with HTTP 402, raise `'PT402'`:

```
RAISE sqlstate 'PT402' using
  message = 'Payment Required',
  detail = 'Quota exceeded',
  hint = 'Upgrade your plan';
```

Returns:

```
HTTP/1.1 402 Payment Required
Content-Type: application/json; charset=utf-8

{"hint": "Upgrade your plan", "details": "Quota exceeded"}
```

## 7.5 Insertions / Updates

All tables and [auto-updatable views](#) can be modified through the API, subject to permissions of the requester's database role.

To create a row in a database table post a JSON object whose keys are the names of the columns you would like to create. Missing properties will be set to default values when applicable.

```
POST /table_name HTTP/1.1
{ "col1": "value1", "col2": "value2" }
```

The response will include a `Location` header describing where to find the new object. If the table is write-only then constructing the `Location` header will cause a permissions error. To successfully insert an item to a write-only table you will need to suppress the `Location` response header by including the request header `Prefer: return=minimal`.

On the other end of the spectrum you can get the full created object back in the response to your request by including the header `Prefer: return=representation`. That way you won't have to make another HTTP call to discover properties that may have been filled in on the server side. You can also apply the standard [Vertical Filtering \(Columns\)](#) to these results.

URL encoded payloads can be posted with `Content-Type: application/x-www-form-urlencoded`.

```
POST /people HTTP/1.1
Content-Type: application/x-www-form-urlencoded
name=John+Doe&age=50&weight=80
```

**Note:** When inserting a row you must post a JSON object, not quoted JSON.

```
Yes
{ "a": 1, "b": 2 }

No
"{ \"a\": 1, \"b\": 2 }"
```

Some javascript libraries will post the data incorrectly if you're not careful. For best results try one of the [Client-Side Libraries](#) built for PostgreSQL.

To update a row or rows in a table, use the PATCH verb. Use [Horizontal Filtering \(Rows\)](#) to specify which record(s) to update. Here is an example query setting the `category` column to `child` for all people below a certain age.

```
PATCH /people?age=1t.13 HTTP/1.1
{ "category": "child" }
```

Updates also support `Prefer: return=representation` plus [Vertical Filtering \(Columns\)](#).

**Warning:** Beware of accidentally updating every row in a table. To learn to prevent that see [Block Full-Table Operations](#).

**Warning:** Insertion on VIEWS with complex RULEs might not work out of the box with PostgREST. It's recommended that you use [triggers](#) instead of RULEs. If you want to keep using RULEs, a workaround is to wrap the VIEW insertion in a stored procedure and call it through the *Stored Procedures* interface.

### 7.5.1 Bulk Insert

Bulk insert works exactly like single row insert except that you provide either a JSON array of objects having uniform keys, or lines in CSV format. This not only minimizes the HTTP requests required but uses a single INSERT statement on the back-end for efficiency. Note that using CSV requires less parsing on the server and is much faster.

To bulk insert CSV simply post to a table route with `Content-Type: text/csv` and include the names of the columns as the first row. For instance

```
POST /people HTTP/1.1
Content-Type: text/csv

name,age,height
J Doe,62,70
Jonas,10,55
```

An empty field (, ) is coerced to an empty string and the reserved word NULL is mapped to the SQL null value. Note that there should be no spaces between the column names and commas.

To bulk insert JSON post an array of objects having all-matching keys

```
POST /people HTTP/1.1
Content-Type: application/json

[
  { "name": "J Doe", "age": 62, "height": 70 },
  { "name": "Janus", "age": 10, "height": 55 }
]
```

### 7.5.2 Specifying Columns

By using the `columns` query parameter it's possible to specify the payload keys that will be inserted/updated and ignore the rest of the payload.

```
POST /datasets?columns=source,publication_date,figure HTTP/1.1
Content-Type: application/json

{
  "source": "Natural Disaster Prevention and Control",
  "publication_date": "2015-09-11",
  "figure": 1100,
  "location": "...",
  "comment": "...",
  "extra": "...",
  "stuff": "..."
}
```

In this case, only `source`, `publication_date` and `figure` will be inserted. The rest of the JSON keys will be ignored.

Using this also has the side-effect of being more efficient for *Bulk Insert* since PostgREST will not process the JSON and it'll send it directly to PostgreSQL.

### 7.5.3 Upsert

You can make an UPSERT with POST and the `Prefer: resolution=merge-duplicates` header:

```
POST /employees HTTP/1.1
Prefer: resolution=merge-duplicates

[
  { "id": 1, "name": "Old employee 1", "salary": 30000 },
  { "id": 2, "name": "Old employee 2", "salary": 42000 },
  { "id": 3, "name": "New employee 3", "salary": 50000 }
]
```

UPSERT operates based on the primary key columns, you must specify all of them. You can also choose to ignore the duplicates with `Prefer: resolution=ignore-duplicates`. UPSERT works best when the primary key is natural, but it's also possible to use it if the primary key is surrogate (example: "id serial primary key"). For more details read [this issue](#).

---

**Important:** After creating a table or changing its primary key, you must refresh PostgREST schema cache for UPSERT to work properly. To learn how to refresh the cache see [Schema Reloading](#).

---

A single row UPSERT can be done by using PUT and filtering the primary key columns with `eq`:

```
PUT /employees?id=eq.4 HTTP/1.1

{ "id": 4, "name": "Sara B.", "salary": 60000 }
```

All the columns must be specified in the request body, including the primary key columns.

---

**Note:** This feature is only available starting from PostgreSQL 9.5 since it uses the `ON CONFLICT` clause.

---

## 7.6 Deletions

To delete rows in a table, use the DELETE verb plus *Horizontal Filtering (Rows)*. For instance deleting inactive users:

```
DELETE /user?active=is.false HTTP/1.1
```

**Warning:** Beware of accidentally deleting all rows in a table. To learn to prevent that see [Block Full-Table Operations](#).

## 7.7 Binary Output

If you want to return raw binary data from a bytea column, you must specify `application/octet-stream` as part of the `Accept` header and select a single column `?select=bin_data`.

```
GET /items?select=bin_data&id=eq.1 HTTP/1.1
Accept: application/octet-stream
```

You can also request binary output when calling *Stored Procedures* and since they can return a scalar value you are not forced to use `select` for this case.

```
CREATE FUNCTION closest_point(..) RETURNS bytea ..
```

```
POST /rpc/closest_point HTTP/1.1
Accept: application/octet-stream
```

If the stored procedure returns non-scalar values, you need to do a `select` in the same way as for GET binary output.

```
CREATE FUNCTION overlapping_regions(..) RETURNS SETOF TABLE (geom_twkb bytea, ..) ..
```

```
POST /rpc/overlapping_regions?select=geom_twkb HTTP/1.1
Accept: application/octet-stream
```

---

**Note:** If more than one row would be returned the binary results will be concatenated with no delimiter.

---

### 7.7.1 Plain Text Output

You can get raw output from a `text` column by using `Accept: text/plain`.

```
GET /workers?select=custom_psv_format HTTP/1.1
Accept: text/plain
```

```
09310817|JOHN|DOE|15/04/88|
42152780|FRED|BLOGGS|20/02/85|
43006541|OTTO|NORMALVERBRAUCHER|01/07/90|
02452492|ERIKA|MUSTERMANN|11/01/80|
```

This follows the same rules as *Binary Output*.

## 7.8 OpenAPI Support

Every API hosted by PostgreSQL automatically serves a full [OpenAPI](#) description on the root path. This provides a list of all endpoints (tables, foreign tables, views, functions), along with supported HTTP verbs and example payloads. For extra customization, the OpenAPI output contains a “description” field for every [SQL comment](#) on any database object. For instance,

```
COMMENT ON SCHEMA mammals IS
  'A warm-blooded vertebrate animal of a class that is distinguished by the secretion_
  ↳of milk by females for the nourishment of the young';

COMMENT ON TABLE monotremes IS
  'Freakish mammals lay the best eggs for breakfast';

COMMENT ON COLUMN monotremes.has_venomous_claw IS
  'Sometimes breakfast is not worth it';
```

These unsavory comments will appear in the generated JSON as the fields `info.description`, `definitions.monotremes.description` and `definitions.monotremes.properties.has_venomous_claw.description`.

Also if you wish to generate a `summary` field you can do it by having a multiple line comment, the `summary` will be the first line and the `description` the lines that follow it:

```
COMMENT ON TABLE entities IS
  $$Entities summary

  Entities description that
  spans
  multiple lines$$;
```

You can use a tool like [Swagger UI](#) to create beautiful documentation from the description and to host an interactive web-based dashboard. The dashboard allows developers to make requests against a live PostgREST server, and provides guidance with request headers and example request bodies.

---

**Important:** The OpenAPI information can go out of date as the schema changes under a running server. To learn how to refresh the cache see [Schema Reloading](#).

---

## 7.9 HTTP Status Codes

PostgREST translates PostgreSQL error codes into HTTP status as follows:

PostgreSQL error code(s)	HTTP status	Error description
08*	503	pg connection err
09*	500	triggered action exception
0L*	403	invalid grantor
0P*	403	invalid role specification
23503	409	foreign key violation
23505	409	uniqueness violation
25*	500	invalid transaction state
28*	403	invalid auth specification
2D*	500	invalid transaction termination
38*	500	external routine exception
39*	500	external routine invocation
3B*	500	savepoint exception
40*	500	transaction rollback
53*	503	insufficient resources
54*	413	too complex
55*	500	obj not in prerequisite state
57*	500	operator intervention
58*	500	system error
F0*	500	conf file error
HV*	500	foreign data wrapper error
P0001	400	default code for "raise"
P0*	500	PL/pgSQL error
XX*	500	internal error
42883	404	undefined function
42P01	404	undefined table
42501	if authenticated 403, else 401	insufficient privileges
other	400	

## 7.10 Configuration

PostgREST reads a configuration file to determine information about the database and how to serve client requests. There is no predefined location for this file, you must specify the file path as the one and only argument to the server:

```
./postgrest /path/to/postgrest.conf
```

The configuration file must contain a set of key value pairs. At minimum you must include these keys:

```
# postgrest.conf

# The standard connection URI format, documented at
# https://www.postgresql.org/docs/current/static/libpq-connect.html#AEN45347
db-uri      = "postgres://user:pass@host:5432/dbname"

# The name of which database schema to expose to REST clients
db-schema   = "api"
```

(continues on next page)

(continued from previous page)

```
# The database role to use when no client authentication is provided.
# Can (and should) differ from user in db-uri
db-anon-role = "anon"
```

The user specified in the db-uri is also known as the authenticator role. For more information about the anonymous vs authenticator roles see the *Overview of Role System*.

Here is the full list of configuration parameters.

Name	Type	Default	Required
db-uri	String		Y
db-schema	String		Y
db-anon-role	String		Y
db-pool	Int	10	
db-pool-timeout	Int	10	
db-extra-search-path	String	public	
server-host	String	!4	
server-port	Int	3000	
server-unix-socket	String		
server-proxy-uri	String		
jwt-secret	String		
jwt-aud	String		
secret-is-base64	Bool	False	
max-rows	Int	∞	
pre-request	String		
app.settings.*	String		
role-claim-key	String	.role	
raw-media-types	String		

### 7.10.1 db-uri

The standard connection PostgreSQL [URI format](#). Symbols and unusual characters in the password or other fields should be percent encoded to avoid a parse error. If enforcing an SSL connection to the database is required you can use [sslmode](#) in the URI, for example `postgres://user:pass@host:5432/dbname?sslmode=require`.

When running PostgREST on the same machine as PostgreSQL, it is also possible to connect to the database using a [Unix socket](#) and the [Peer Authentication method](#) as an alternative to TCP/IP communication and authentication with a password, this also grants higher performance. To do this you can omit the host and the password, e.g. `postgres://user@/dbname`, see the [libpq connection string](#) documentation for more details.

On older systems like Centos 6, with older versions of libpq, a different db-uri syntax has to be used. In this case the URI is a string of space separated key-value pairs (key=value), so the example above would be `"host=host user=user port=5432 dbname=dbname password=pass"`.

Choosing a value for this parameter beginning with the at sign such as `@filename` (e.g. `@./configs/my-config`) loads the secret out of an external file.

### 7.10.2 db-schema

The database schema to expose to REST clients. Tables, views and stored procedures in this schema will get API endpoints.

This schema gets added to the `search_path` of every request.

### 7.10.3 db-anon-role

The database role to use when executing commands on behalf of unauthenticated clients. For more information, see *Overview of Role System*.

### 7.10.4 db-pool

Number of connections to keep open in PostgREST's database pool. Having enough here for the maximum expected simultaneous client connections can improve performance. Note it's pointless to set this higher than the `max_connections` GUC in your database.

### 7.10.5 db-pool-timeout

Time to live for an idle database pool connection. If the timeout is reached the connection will be closed. Once a new request arrives a new connection will be started.

### 7.10.6 db-extra-search-path

Extra schemas to add to the `search_path` of every request. These schemas tables, views and stored procedures **don't get API endpoints**, they can only be referred from the database objects inside your *db-schema*.

This parameter was meant to make it easier to use **PostgreSQL extensions** (like PostGIS) that are outside of the *db-schema*.

Multiple schemas can be added in a comma-separated string, e.g. `public, extensions`.

### 7.10.7 server-host

Where to bind the PostgREST web server. In addition to the usual address options, PostgREST interprets these reserved addresses with special meanings:

- `*` - any IPv4 or IPv6 hostname
- `*4` - any IPv4 or IPv6 hostname, IPv4 preferred
- `!4` - any IPv4 hostname
- `*6` - any IPv4 or IPv6 hostname, IPv6 preferred
- `!6` - any IPv6 hostname

### 7.10.8 server-port

The TCP port to bind the web server.

### 7.10.9 server-unix-socket

Unix domain socket where to bind the PostgREST web server. If specified, this takes precedence over *server-port*. Example:

```
server-unix-socket = "/tmp/pgrst.sock"
```

### 7.10.10 server-proxy-uri

Overrides the base URL used within the OpenAPI self-documentation hosted at the API root path. Use a complete URI syntax scheme: `[[//[user:password@]host[:port]][/]path[?query][#fragment]`. Ex. `https://postgrest.com`

```
{
  "swagger": "2.0",
  "info": {
    "version": "0.4.3.0",
    "title": "PostgREST API",
    "description": "This is a dynamic API generated by PostgREST"
  },
  "host": "postgrest.com:443",
  "basePath": "/",
  "schemes": [
    "https"
  ]
}
```

### 7.10.11 jwt-secret

The secret or [JSON Web Key \(JWK\)](#) (or *set*) used to decode JWT tokens clients provide for authentication. For security the key must be **at least 32 characters long**. If this parameter is not specified then PostgREST refuses authentication requests. Choosing a value for this parameter beginning with the at sign such as `@filename` loads the secret out of an external file. This is useful for automating deployments. Note that any binary secrets must be base64 encoded. Both symmetric and asymmetric cryptography are supported. For more info see [Asymmetric Keys](#).

### 7.10.12 jwt-aud

Specifies the [JWT audience claim](#). If this claim is present in the client provided JWT then you must set this to the same value as in the JWT, otherwise verifying the JWT will fail.

### 7.10.13 secret-is-base64

When this is set to `true`, the value derived from `jwt-secret` will be treated as a base64 encoded secret.

### 7.10.14 max-rows

A hard limit to the number of rows PostgREST will fetch from a view, table, or stored procedure. Limits payload size for accidental or malicious requests.

### 7.10.15 pre-request

A schema-qualified stored procedure name to call right after switching roles for a client request. This provides an opportunity to modify SQL variables or raise an exception to prevent the request from completing.

### 7.10.16 app.settings.\*

Arbitrary settings that can be used to pass in secret keys directly as strings, or via OS environment variables. For instance: `app.settings.jwt_secret = "$(MYAPP_JWT_SECRET)"` will take `MYAPP_JWT_SECRET` from the environment and make it available to postgresql functions as `current_setting('app.settings.jwt_secret')`.

### 7.10.17 role-claim-key

A JSPath DSL that specifies the location of the `role` key in the JWT claims. This can be used to consume a JWT provided by a third party service like Auth0, Okta or Keycloak. Usage examples:

```
# {"postgrest":{"roles": ["other", "author"]}}
# the DSL accepts characters that are alphanumerical or one of "_$@" as keys
role-claim-key = ".postgrest.roles[1]"

# {"https://www.example.com/role": { "key": "author" }}
# non-alphanumerical characters can go inside quotes(escaped in the config_
↳value)
role-claim-key = ".\"https://www.example.com/role\".key"
```

### 7.10.18 raw-media-types

This serves to extend the [Media Types](#) that PostgREST currently accepts through an `Accept` header.

These media types can be requested by following the same rules as the ones defined in [Binary Output](#).

As an example, the below config would allow you to request an **image** and an **xml** by doing a request with `Accept: image/png` and a request with `Accept: text/xml`, respectively.

```
raw-media-types="image/png, text/xml"
```

- [API](#)
- [Configuration](#)

These are recipes that'll help you address specific use-cases.

## 8.1 Casting a type to a custom JSON object

**author** `steve-chavez`

While using PostgREST you might have noticed that certain PostgreSQL types translate to JSON strings when you would have expected a JSON object or array. For example, let's see the case of `range` types.

```
-- example taken from https://www.postgresql.org/docs/11/rangetypes.html#RANGETYPES-  
→EXAMPLES  
create table reservations (  
  room int  
, during tsrange  
);  
  
insert into  
  reservations  
values  
  (1108, tsrange('2010-01-01 14:30', '2010-01-01 15:30'));
```

Here we have a column named `during` as a `tsrange` type, we would like to get it as JSON through PostgREST.

```
curl "http://localhost:3000/reservations"
```

Result:

```
[  
  {  
    "room":1108,  
    "during":["2010-01-01 14:30:00","2010-01-01 15:30:00"]  
  }  
]
```

The **during** value is probably not the in the format you want. We get a JSON string because by default PostgreSQL casts the type to JSON by using its `text` representation. We can change this representation to a custom JSON object by creating a CAST.

To do this, first we'll define the function that will do the conversion from `tsrange` to `json`.

```
create or replace function tsrange_to_json(tsrange) returns json as $$
select json_build_object (
  'lower', lower($1)
, 'upper', upper($1)
, 'lower_inc', lower_inc($1)
, 'upper_inc', upper_inc($1)
);
$$ language sql;
```

Using this function we'll create the CAST.

```
create cast (tsrange as json) with function tsrange_to_json(tsrange) as assignment;
```

And we'll do the request and *cast the column*.

```
curl "http://localhost:3000/reservations?select=room,during::json"
```

The result now is:

```
[
  {
    "room":1108,
    "during":{
      "lower" : "2010-01-01T14:30:00",
      "upper" : "2010-01-01T15:30:00",
      "lower_inc" : true,
      "upper_inc" : false
    }
  }
]
```

You can use the same idea for creating custom CASTs for different types.

---

**Note:** If you don't want to modify CASTs for built-in types, an option would be to create a custom type for your own `tsrange` and add its own CAST.

```
create type mytsrange as range (subtype = timestamp, subtype_diff = tsrange_subdiff);
-- define column types and casting function analogously to the above example
-- ...

create cast (mytsrange as json) with function mytsrange_to_json(mytsrange) as
↪assignment;
```

---

## 8.2 Embedding a table from another schema

**author** steve-chavez

Suppose you have a **people** table in the `public` schema and this schema is exposed through PostgREST's *db-schema*.

```
create table public.people(
  id          int  primary key
, full_name text
);
```

And you want to *embed* the **people** table with a **details** table that's in another schema named `private`.

```
create schema if not exists private;

-- For simplicity's sake the table is devoid of constraints/domains on email, phone,
↪etc.
create table private.details(
  id          int  primary key references public.people
, email      text
, phone      text
, birthday   date
, occupation text
, company    text
);

-- other database objects in this schema
-- ...
-- ...
```

To solve this, you can create a view of **details** in the `public` schema. We'll call it **public\_details**.

```
create view public.public_details as
select
  id
, occupation
, company
from
  private.details;
```

Since PostgreSQL supports *Embedding Views*, you can embed **people** with **public\_details**.

Let's insert some data to test this:

```
insert into
  public.people
values
  (1, 'John Doe'), (2, 'Jane Doe');

insert into
  private.details
values
  (1, 'jhon@fake.com', '772-323-5433', '1990-02-01', 'Transportation attendant',
↪'Body Fate'),
  (2, 'jane@fake.com', '480-474-6571', '1980-04-21', 'Geotechnical engineer',
↪'Earthworks Garden Kare');
```

---

**Important:** Make sure PostgreSQL's schema cache is up-to-date. See *Schema Reloading*.

---

Now, make the following request:

```
curl "http://localhost:3000/people?select=full_name,public_details(occupation,company)
↳"
```

The result should be:

```
[
  { "full_name": "John Doe", "public_details": [{"occupation": "Transportation attendant",
↳ "company": "Body Fate"}]},
  { "full_name": "Jane Doe", "public_details": [{"occupation": "Geotechnical engineer",
↳ "company": "Earthworks Garden Kare"}]}
]
```

- *Casting a type to a custom JSON object*
- *Embedding a table from another schema*

Explanations of some key concepts in PostgREST.

## 9.1 Overview of Role System


PostgREST is designed to keep the database at the center of API security. All authorization happens through database roles and permissions. It is PostgREST's job to **authenticate** requests – i.e. verify that a client is who they say they are – and then let the database **authorize** client actions.

### 9.1.1 Authentication Sequence

There are three types of roles used by PostgREST, the **authenticator**, **anonymous** and **user** roles. The database administrator creates these roles and configures PostgREST to use them.



The authenticator should be created `NOINHERIT` and configured in the database to have very limited access. It is a chameleon whose job is to “become” other users to service authenticated HTTP requests. The picture below shows how the server handles authentication. If auth succeeds, it switches into the user role specified by the request, otherwise it switches into the anonymous role.



Here are the technical details. We use [JSON Web Tokens](#) to authenticate API requests. As you’ll recall a JWT contains a list of cryptographically signed claims. All claims are allowed but PostgreSQL cares specifically about a claim called `role`.

```
{
  "role": "user123"
}
```

When a request contains a valid JWT with a role claim PostgREST will switch to the database role with that name for the duration of the HTTP request.

```
SET LOCAL ROLE user123;
```

Note that the database administrator must allow the authenticator role to switch into this user by previously executing

```
GRANT user123 TO authenticator;
```

If the client included no JWT (or one without a role claim) then PostgREST switches into the anonymous role whose actual database-specific name, like that of with the authenticator role, is specified in the PostgREST server configuration file. The database administrator must set anonymous role permissions correctly to prevent anonymous users from seeing or changing things they shouldn't.

## 9.1.2 Users and Groups

PostgreSQL manages database access permissions using the concept of roles. A role can be thought of as either a database user, or a group of database users, depending on how the role is set up.

### Roles for Each Web User

PostgREST can accommodate either viewpoint. If you treat a role as a single user then the the JWT-based role switching described above does most of what you need. When an authenticated user makes a request PostgREST will switch into the role for that user, which in addition to restricting queries, is available to SQL through the `current_user` variable.

You can use row-level security to flexibly restrict visibility and access for the current user. Here is an [example](#) from Tomas Vondra, a chat table storing messages sent between users. Users can insert rows into it to send messages to other users, and query it to see messages sent to them by other users.

```
CREATE TABLE chat (
  message_uuid  UUID PRIMARY KEY DEFAULT uuid_generate_v4(),
  message_time  TIMESTAMP NOT NULL DEFAULT now(),
  message_from  NAME      NOT NULL DEFAULT current_user,
  message_to    NAME      NOT NULL,
  message_subject VARCHAR(64) NOT NULL,
  message_body  TEXT
);
```

We want to enforce a policy that ensures a user can see only those messages sent by him or intended for him. Also we want to prevent a user from forging the `message_from` column with another person's name.

PostgreSQL (9.5 and later) allows us to set this policy with row-level security:

```
CREATE POLICY chat_policy ON chat
  USING ((message_to = current_user) OR (message_from = current_user))
  WITH CHECK (message_from = current_user)
```

Anyone accessing the generated API endpoint for the chat table will see exactly the rows they should, without our needing custom imperative server-side coding.

**Warning:** Roles are namespaced per-cluster rather than per-database so they may be prone to collision.

## Web Users Sharing Role

Alternately database roles can represent groups instead of (or in addition to) individual users. You may choose that all signed-in users for a web app share the role webuser. You can distinguish individual users by including extra claims in the JWT such as email.

```
{
  "role": "webuser",
  "email": "john@doe.com"
}
```

SQL code can access claims through GUC variables set by PostgREST per request. For instance to get the email claim, call this function:

```
current_setting('request.jwt.claim.email', true)
```

This allows JWT generation services to include extra information and your database code to react to it. For instance the RLS example could be modified to use this `current_setting` rather than `current_user`. The second 'true' argument tells `current_setting` to return NULL if the setting is missing from the current configuration.

## Hybrid User-Group Roles

You can mix the group and individual role policies. For instance we could still have a webuser role and individual users which inherit from it:

```
CREATE ROLE webuser NOLOGIN;
-- grant this role access to certain tables etc

CREATE ROLE user000 NOLOGIN;
GRANT webuser TO user000;
-- now user000 can do whatever webuser can

GRANT user000 TO authenticator;
-- allow authenticator to switch into user000 role
-- (the role itself has nologin)
```

### 9.1.3 Custom Validation

PostgREST honors the `exp` claim for token expiration, rejecting expired tokens. However it does not enforce any extra constraints. An example of an extra constraint would be to immediately revoke access for a certain user. The configuration file parameter `pre-request` specifies a stored procedure to call immediately after the authenticator switches into a new role and before the main query itself runs.

Here's an example. In the config file specify a stored procedure:

```
pre-request = "public.check_user"
```

In the function you can run arbitrary code to check the request and raise an exception to block it if desired.

```
CREATE OR REPLACE FUNCTION check_user() RETURNS void AS $$
BEGIN
  IF current_user = 'evil_user' THEN
    RAISE EXCEPTION 'No, you are evil'
    USING HINT = 'Stop being so evil and maybe you can log in';
  END IF;
END
$$ LANGUAGE plpgsql;
```

## 9.2 Client Auth

To make an authenticated request the client must include an Authorization HTTP header with the value Bearer <jwt>. For instance:

```
GET /foo HTTP/1.1
Authorization: Bearer eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.
→eyJyY2x1IjoiamRvZSIsImV4cCI6MTQ3NTUxNjI1MH0.GYDZV3yM0gqvuEtJmfppLbXSGYnke_
→Pvn10tbKAjB4
```

### 9.2.1 JWT Generation

You can create a valid JWT either from inside your database or via an external service. Each token is cryptographically signed with a secret key. In the case of symmetric cryptography the signer and verifier share the same secret passphrase. In asymmetric cryptography the signer uses the private key and the verifier the public key. PostgREST supports both symmetric and asymmetric cryptography.

#### JWT from SQL

You can create JWT tokens in SQL using the [pgjwt extension](#). It's simple and requires only `pgcrypto`. If you're on an environment like Amazon RDS which doesn't support installing new extensions, you can still manually run the [SQL inside pgjwt](#) (you'll need to replace `@extschema@` with another schema or just delete it) which creates the functions you will need.

Next write a stored procedure that returns the token. The one below returns a token with a hard-coded role, which expires five minutes after it was issued. Note this function has a hard-coded secret as well.

```
CREATE TYPE jwt_token AS (
  token text
);

CREATE FUNCTION jwt_test() RETURNS public.jwt_token AS $$
SELECT public.sign(
  row_to_json(r), 'reallyreallyreallyreallyverysafe'
) AS token
FROM (
  SELECT
    'my_role'::text AS role,
    extract(epoch from now())::integer + 300 AS exp
) r;
$$ LANGUAGE sql;
```

PostgREST exposes this function to clients via a POST request to `/rpc/jwt_test`.

**Note:** To avoid hard-coding the secret in stored procedures, save it as a property of the database.

```
-- run this once
ALTER DATABASE mydb SET "app.jwt_secret" TO 'reallyreallyreallyreallyverysafe';

-- then all functions can refer to app.jwt_secret
SELECT sign(
  row_to_json(r), current_setting('app.jwt_secret')
) AS token
FROM ...
```

### JWT from Auth0

An external service like [Auth0](#) can do the hard work transforming OAuth from Github, Twitter, Google etc into a JWT suitable for PostgREST. Auth0 can also handle email signup and password reset flows.

To use Auth0, copy its client secret into your PostgREST configuration file as the `jwt-secret`. (Old-style Auth0 secrets are Base64 encoded. For these secrets set `secret-is-base64` to `true`, or just refresh the Auth0 secret.) You can find the secret in the client settings of the Auth0 management console.

**Note:** Make sure OIDC-conformant is toggled off.

A recent Auth0 change sets it on by default. Turn it *off* here:

Clients > *Your App* > Settings > Show Advanced Settings > OAuth > OIDC Conformant

Ensure also that your client application does not pass in any *audience* configuration.

Our code requires a database role in the JWT. To add it you need to save the database role in Auth0 [app metadata](#). Then, you will need to write a rule that will extract the role from the user metadata and include a `role` claim in the payload of our user object. Afterwards, in your Auth0Lock code, include the `role` claim in your `scope param`.

```
// Example Auth0 rule
function (user, context, callback) {
  user.app_metadata = user.app_metadata || {};
  user.role = user.app_metadata.role;
  callback(null, user, context);
}
```

```
// Example using Auth0Lock with role claim in scope
new Auth0Lock ( AUTH0_CLIENTID, AUTH0_DOMAIN, {
  container: 'lock-container',
  auth: {
    params: { scope: 'openid role' },
    redirectUrl: FQDN + '/login', // Replace with your redirect url
    responseType: 'token'
  }
})
```

## Asymmetric Keys

As described in the *Configuration* section, PostgREST accepts a `jwt-secret` config file parameter. If it is set to a simple string value like “reallyreallyreallyreallyverysafe” then PostgREST interprets it as an HMAC-SHA256 passphrase. However you can also specify a literal JSON Web Key (JWK) or set. For example, you can use an RSA-256 public key encoded as a JWK:

```
{
  "alg": "RS256",
  "e": "AQAB",
  "key_ops": ["verify"],
  "kty": "RSA",
  "n": "9zKNYTaYGfGmltBMpRT6FxOYrM720GhXdettc02uyakYSEHU2IJz90G_MLlE14-WWWYoS_
↪QKFupw3s7aPY1aAjamG22rAnvWu-rRkP5sSSkKvud_IgKL4iE6Y2WJx2Bk11XUFkdZ8w1EUR6O1ft3TS4uA-
↪qKifSZ43CahzAJyUezOH9shI--tirC0281Ng7671dEki3WnVr3zokSujc9YJ_9XXjw2hFBfmJUrnB0-
↪wldvxQbFU8RPXip-GQ_JPTrCTZhrzGFvPvhA6Rqmc3b1PhM9jY7Dur1sjYWYVyXlFNCK3c-
↪6feo5W1RfelaCWmwZQh6O18eTmLeT4nWYkDzQ"
}
```

**Note:** This could also be a JSON Web Key Set (JWKS) if it was contained within an array assigned to a *keys* member, e.g. `{ keys: [jwk1, jwk2] }`.

Just pass it in as a single line string, escaping the quotes:

```
jwt-secret = "{ \"alg\": \"RS256\", ... }"
```

To generate such a public/private key pair use a utility like `latchset/jose`.

```
jose jwk gen -i '{"alg": "RS256"}' -o rsa.jwk
jose jwk pub -i rsa.jwk -o rsa.jwk.pub

# now rsa.jwk.pub contains the desired JSON object
```

You can specify the literal value as we saw earlier, or reference a filename to load the JWK from a file:

```
jwt-secret = "@rsa.jwk.pub"
```

## JWT security

There are at least three types of common critiques against using JWT: 1) against the standard itself, 2) against using libraries with known security vulnerabilities, and 3) against using JWT for web sessions. We’ll briefly explain each critique, how PostgREST deals with it, and give recommendations for appropriate user action.

The critique against the *JWT standard* is voiced in detail [elsewhere on the web](#). The most relevant part for PostgREST is the so-called `alg=none` issue. Some servers implementing JWT allow clients to choose the algorithm used to sign the JWT. In this case, an attacker could set the algorithm to `none`, remove the need for any signature at all and gain unauthorized access. The current implementation of PostgREST, however, does not allow clients to set the signature algorithm in the HTTP request, making this attack irrelevant. The critique against the standard is that it requires the implementation of the `alg=none` at all.

Critiques against JWT libraries are only relevant to PostgREST via the library it uses. As mentioned above, not allowing clients to choose the signature algorithm in HTTP requests removes the greatest risk. Another more subtle attack is possible where servers use asymmetric algorithms like RSA for signatures. Once again this is not relevant to PostgREST since it is not supported. Curious readers can find more information in [this article](#). Recommendations about high quality libraries for usage in API clients can be found on [jwt.io](#).

The last type of critique focuses on the misuse of JWT for maintaining web sessions. The basic recommendation is to [stop using JWT for sessions](#) because most, if not all, solutions to the problems that arise when you do, [do not work](#). The linked articles discuss the problems in depth but the essence of the problem is that JWT is not designed to be secure and stateful units for client-side storage and therefore not suited to session management.

PostgREST uses JWT mainly for authentication and authorization purposes and encourages users to do the same. For web sessions, using cookies over HTTPS is good enough and well catered for by standard web frameworks.

### 9.2.2 HTTPS

PostgREST aims to do one thing well: add an HTTP interface to a PostgreSQL database. To keep the code small and focused we do not implement HTTPS. Use a reverse proxy such as NGINX to add this, [here's how](#). Note that some Platforms as a Service like Heroku also add SSL automatically in their load balancer.

## 9.3 Schema Isolation

A PostgREST instance is configured to expose all the tables, views, and stored procedures of a single schema specified in a server configuration file. This means private data or implementation details can go inside a private schema and be invisible to HTTP clients. You can then expose views and stored procedures which insulate the internal details from the outside world. It keeps you code easier to refactor, and provides a natural way to do API versioning. For an example of wrapping a private table with a public view see the [Public User Interface](#) section below.

## 9.4 SQL User Management

### 9.4.1 Storing Users and Passwords

As mentioned, an external service can provide user management and coordinate with the PostgREST server using JWT. It's also possible to support logins entirely through SQL. It's a fair bit of work, so get ready.

The following table, functions, and triggers will live in a `basic_auth` schema that you shouldn't expose publicly in the API. The public views and functions will live in a different schema which internally references this internal information.

First we'll need a table to keep track of our users:

```
-- We put things inside the basic_auth schema to hide
-- them from public view. Certain public procs/views will
-- refer to helpers and tables inside.
create schema if not exists basic_auth;

create table if not exists
basic_auth.users (
  email    text primary key check ( email ~* '^.+@.+\.+$' ),
  pass     text not null check (length(pass) < 512),
  role     name not null check (length(role) < 512)
);
```

We would like the role to be a foreign key to actual database roles, however PostgreSQL does not support these constraints against the `pg_roles` table. We'll use a trigger to manually enforce it.

```

create or replace function
basic_auth.check_role_exists() returns trigger as $$
begin
  if not exists (select 1 from pg_roles as r where r.rolname = new.role) then
    raise foreign_key_violation using message =
      'unknown database role: ' || new.role;
    return null;
  end if;
  return new;
end
$$ language plpgsql;

drop trigger if exists ensure_user_role_exists on basic_auth.users;
create constraint trigger ensure_user_role_exists
after insert or update on basic_auth.users
for each row
execute procedure basic_auth.check_role_exists();

```

Next we'll use the pgcrypto extension and a trigger to keep passwords safe in the users table.

```

create extension if not exists pgcrypto;

create or replace function
basic_auth.encrypt_pass() returns trigger as $$
begin
  if tg_op = 'INSERT' or new.pass <> old.pass then
    new.pass = crypt(new.pass, gen_salt('bf'));
  end if;
  return new;
end
$$ language plpgsql;

drop trigger if exists encrypt_pass on basic_auth.users;
create trigger encrypt_pass
before insert or update on basic_auth.users
for each row
execute procedure basic_auth.encrypt_pass();

```

With the table in place we can make a helper to check a password against the encrypted column. It returns the database role for a user if the email and password are correct.

```

create or replace function
basic_auth.user_role(email text, pass text) returns name
  language plpgsql
  as $$
begin
  return (
    select role from basic_auth.users
    where users.email = user_role.email
      and users.pass = crypt(user_role.pass, users.pass)
  );
end;
$$;

```



anonymous roles. Below is an example of permissions that allow anonymous users to create accounts and attempt to log in.

```
-- the names "anon" and "authenticator" are configurable and not
-- sacred, we simply choose them for clarity
create role anon noinherit;
create role authenticator noinherit;
grant anon to authenticator;

grant execute on function login(text,text) to anon;
```

Since the above `login` function is defined as *security definer*, the anonymous user `anon` doesn't need permission to read the `basic_auth.users` table. It doesn't even need permission to access the `basic_auth` schema. `grant execute on function` is included for clarity but it might not be needed, see *Function privileges* for more details.

## 9.5 Binary Release

[ [Download from release page](#) ]

The release page has pre-compiled binaries for Mac OS X, Windows, and several Linux distributions. Extract the tarball and run the binary inside with the `--help` flag to see usage instructions:

```
# Untar the release (available at https://github.com/PostgREST/postgrest/releases/
→latest)

$ tar Jxf postgrest-[version]-[platform].tar.xz

# Try running it
$ ./postgrest --help

# You should see a usage help message
```

**Note:** If you see a dialog box like this on Windows, it may be that the `pg_config` program is not in your system path.



It usually lives in `C:\Program Files\PostgreSQL\<version>\bin`. See [this article](#) about how to modify the system path.

## 9.6 PostgreSQL dependency

To use PostgREST you will need an underlying database (PostgreSQL version 9.5 or greater is required). You can use something like Amazon RDS but installing your own locally is cheaper and more convenient for development.

- [Instructions for OS X](#)
- [Instructions for Ubuntu 14.04](#)
- [Installer for Windows](#)

On Windows, PostgREST will fail to run unless the PostgreSQL binaries are on the system path. To test whether this is the case, run `pg_config` from the command line. You should see it output a list of paths.

## 9.7 Configuration

The PostgREST server reads a configuration file as its only argument:

```
./postgrest /path/to/postgrest.conf
```

For a complete reference of the configuration file, see [Configuration](#).

## 9.8 Running the Server

PostgREST outputs basic request logging to stdout. When running it in an SSH session you must detach it from stdout or it will be terminated when the session closes. The easiest technique is redirecting the output to a log file or to the syslog:

```
ssh foo@example.com \  
  'postgrest foo.conf </dev/null >/var/log/postgrest.log 2>&1 &'  
  
# another option is to pipe the output into "logger -t postgrest"
```

## 9.9 Docker

You can get the official PostgREST Docker image with:

```
docker pull postgrest/postgrest
```

The image consults an internal `/etc/postgrest.conf` file. To customize this file you can either mount a replacement configuration file into the container, or use environment variables. The environment variables will be interpolated into the default config file.

These variables match the options shown in our [Configuration](#) section, except they are capitalized, have a `PGRST_` prefix, and use underscores. To get a list of the available environment variables, run this:

```
docker inspect -f "{{.Config.Env}}" postgrest/postgrest
```

There are two ways to run the PostgREST container: with an existing external database, or through docker-compose.

### 9.9.1 Containerized PostgREST with native PostgreSQL

The first way to run PostgREST in Docker is to connect it to an existing native database on the host.

```
# Run the server
docker run --rm --net=host -p 3000:3000 \
-e PGRST_DB_URI="postgres://postgres@localhost/postgres" \
-e PGRST_DB_ANON_ROLE="postgres" \
postgres/postgres
```

The database connection string above is just an example. Adjust the role and password as necessary. You may need to edit PostgreSQL's `pg_hba.conf` to grant the user local login access.

**Note:** Docker on Mac does not support the `--net=host` flag. Instead you'll need to create an IP address alias to the host. Requests for the IP address from inside the container are unable to resolve and fall back to resolution by the host.

```
sudo ifconfig lo0 10.0.0.10 alias
```

You should then use `10.0.0.10` as the host in your database connection string. Also remember to include the IP address in the `listen_addresses` within `postgresql.conf`. For instance:

```
listen_addresses = 'localhost,10.0.0.10'
```

You might also need to add a new IPv4 local connection within `pg_hba.conf`. For instance:

```
host      all             all             10.0.0.10/32      trust
```

## 9.9.2 Containerized PostgREST and db with docker-compose

To avoid having to install the database at all, you can run both it and the server in containers and link them together with docker-compose. Use this configuration:

```
# docker-compose.yml

version: '3'
services:
  server:
    image: postgres/postgres
    ports:
      - "3000:3000"
    links:
      - db:db
    environment:
      PGRST_DB_URI: postgres://app_user:password@db:5432/app_db
      PGRST_DB_SCHEMA: public
      PGRST_DB_ANON_ROLE: app_user #In production this role should not be the same as,
      ↪ the one used for the connection
    depends_on:
      - db
  db:
    image: postgres
    ports:
      - "5432:5432"
    environment:
      POSTGRES_DB: app_db
      POSTGRES_USER: app_user
```

(continues on next page)

(continued from previous page)

```
POSTGRES_PASSWORD: password
# Uncomment this if you want to persist the data.
# volumes:
#   - "pgdata:/var/lib/postgresql/data"
```

Go into the directory where you saved this file and run `docker-compose up`. You will see the logs of both the database and PostgREST, and be able to access the latter on port 3000.

If you want to have a visual overview of your API in your browser you can add `swagger-ui` to your `docker-compose.yml`:

```
swagger:
  image: swaggerapi/swagger-ui
  ports:
    - "8080:8080"
  expose:
    - "8080"
  environment:
    API_URL: http://localhost:3000/
```

With this you can see the `swagger-ui` in your browser on port 8080.

## 9.10 Deploying to Heroku

Assuming your making modifications locally and then pushing to GitHub, it's easy to deploy to Heroku.

1. Create a new app on Heroku
2. In Settings add the following buildpack <https://github.com/PostgREST/postgrest-heroku>
3. Add the require Config Vars in Heroku (see <https://github.com/PostgREST/postgrest/blob/master/app.json#L7-L57> for more details)
4. Modify your `postgrest.conf` file as required to match your Config Vars in Heroku
5. Create your Procfile and add `./env-to-config ./postgrest postgrest.conf`
6. Push your changes to GitHub
7. Set Heroku to automatically deploy from Master and then manually deploy the branch for the first build

## 9.11 Build from Source

---

**Note:** We discourage building and using PostgREST on **Alpine Linux** because of a reported GHC memory leak on that platform.

---

When a pre-built binary does not exist for your system you can build the project from source. You'll also need to do this if you want to help with development. `Stack` makes it easy. It will install any necessary Haskell dependencies on your system.

- Install `Stack` for your platform
- Install Library Dependencies

Operating System	Dependencies
Ubuntu/Debian	libpq-dev, libgmp-dev
CentOS/Fedora/Red Hat	postgresql-devel, zlib-devel, gmp-devel
BSD	postgresql95-client
OS X	libpq, gmp

- Build and install binary

```
git clone https://github.com/PostgREST/postgrest.git
cd postgrest

# adjust local-bin-path to taste
stack build --install-ghc --copy-bins --local-bin-path /usr/local/bin
```

**Note:** If building fails and your system has less than 1GB of memory, try adding a swap file.

- Check that the server is installed: `postgrest --help`.

## 9.11.1 PostgREST Test Suite

### Creating the Test Database

To properly run `postgrest` tests one needs to create a database. To do so, use the test creation script `create_test_database` in the `test/` folder.

The script expects the following parameters:

```
test/create_test_db connection_uri database_name [test_db_user] [test_db_user_
↪password]
```

Use the [connection URI](#) to specify the user, password, host, and port. Do not provide the database in the connection URI. The PostgreSQL role you are using to connect must be capable of creating new databases.

The `database_name` is the name of the database that `stack test` will connect to. If the database of the same name already exists on the server, the script will first drop it and then re-create it.

Optionally, specify the database user `stack test` will use. The user will be given necessary permissions to reset the database after every test run.

If the user is not specified, the script will generate the role name `postgrest_test_` suffixed by the chosen database name, and will generate a random password for it.

Optionally, if specifying an existing user to be used for the test connection, one can specify the password the user has.

The script will return the db uri to use in the tests—this uri corresponds to the `db-uri` parameter in the configuration file that one would use in production.

Generating the user and the password allows one to create the database and run the tests against any PostgreSQL server without any modifications to the server. (Such as allowing accounts without a password or setting up trust authentication, or requiring the server to be on the same localhost the tests are run from).

### Running the Tests

To run the tests, one must supply the database uri in the environment variable `POSTGREST_TEST_CONNECTION`.

Typically, one would create the database and run the test in the same command line, using the *postgres* superuser:

```
POSTGREST_TEST_CONNECTION=$(test/create_test_db "postgres://postgres:pwd@database-host
↪" test_db) stack test
```

For repeated runs on the same database, one should export the connection variable:

```
export POSTGREST_TEST_CONNECTION=$(test/create_test_db "postgres://
↪postgres:pwd@database-host" test_db)
stack test
stack test
...
```

If the environment variable is empty or not specified, then the test runner will default to connection uri

```
postgres://postgrest_test@localhost/postgrest_test
```

This connection assumes the test server on the `localhost:code:` with the user *postgrest\_test* without the password and the database of the same name.

### Destroying the Database

The test database will remain after the test, together with four new roles created on the PostgreSQL server. To permanently erase the created database and the roles, run the script `test/delete_test_database`, using the same superuser role used for creating the database:

```
test/destroy_test_db connection_uri database_name
```

### Testing with Docker

The ability to connect to non-local PostgreSQL simplifies the test setup. One elegant way of testing is to use a disposable PostgreSQL in docker.

For example, if local development is on a mac with Docker for Mac installed:

```
$ docker run --name db-scripting-test -e POSTGRES_PASSWORD=pwd -p 5434:5432 -d_
↪postgres
$ POSTGREST_TEST_CONNECTION=$(test/create_test_db "postgres://
↪postgres:pwd@localhost:5434" test_db) stack test
```

Additionally, if one creates a docker container to run `stack test` (this is necessary on Mac OS Sierra with GHC below 8.0.1, where `stack test` fails), one can run PostgreSQL in a separate linked container, or use the locally installed PostgreSQL app.

Build the test container with `test/Dockerfile.test`:

```
$ docker build -t pgst-test - < test/Dockerfile.test
$ mkdir .stack-work-docker ~/.stack-linux
```

The first run of the test container will take a long time while the dependencies get cached. Creating the `~/.stack-linux` folder and mapping it as a volume into the container ensures that we can run the container in disposable mode and not worry about subsequent runs being slow. `.stack-work-docker` is also mapped into the container and must be specified when using `stack` from Linux, not to interfere with the `.stack-work` for local development. (On Sierra, `stack build` works, while `stack test` fails with GHC 8.0.1).

Linked containers:

```
$ docker run --name pg -e POSTGRES_PASSWORD=pwd -d postgres
$ docker run --rm -it -v `pwd`:`pwd` -v ~/.stack-linux:/root/.stack --link pg:pg -w=
↪ "`pwd`" -v `pwd`/stack-work-docker:`pwd`/stack-work pgst-test bash -c "POSTGRES_
↪ TEST_CONNECTION=$(test/create_test_db "postgres://postgres:pwd@pg" test_db) stack_
↪ test"
```

Stack test in Docker for Mac, PostgreSQL app on mac:

```
$ host_ip=$(ifconfig en0 | grep 'inet ' | cut -f 2 -d' ')
$ export POSTGRES_TEST_CONNECTION=$(test/create_test_db "postgres://postgres@$HOST"
↪ test_db)
$ docker run --rm -it -v `pwd`:`pwd` -v ~/.stack-linux:/root/.stack -v `pwd`/stack-
↪ work-docker:`pwd`/stack-work -e "HOST=$host_ip" -e "POSTGRES_TEST_CONNECTION=
↪ $POSTGRES_TEST_CONNECTION" -w "`pwd`" pgst-test bash -c "stack test"
$ test/destroy_test_db "postgres://postgres@localhost" test_db
```

## 9.12 Hardening PostgREST

PostgREST is a fast way to construct a RESTful API. Its default behavior is great for scaffolding in development. When it's time to go to production it works great too, as long as you take precautions. PostgREST is a small sharp tool that focuses on performing the API-to-database mapping. We rely on a reverse proxy like Nginx for additional safeguards.

The first step is to create an Nginx configuration file that proxies requests to an underlying PostgREST server.

```
http {
    # ...
    # upstream configuration
    upstream postgrest {
        server localhost:3000;
        keepalive 64;
    }
    # ...
    server {
        # ...
        # expose to the outside world
        location /api/ {
            default_type application/json;
            proxy_hide_header Content-Location;
            add_header Content-Location /api/$upstream_http_content_location;
            proxy_set_header Connection "";
            proxy_http_version 1.1;
            proxy_pass http://postgrest/;
        }
        # ...
    }
}
```

**Note:** For ubuntu, if you already installed nginx through apt you can add this to the config file in `/etc/nginx/sites-enabled/default`.

### 9.12.1 Block Full-Table Operations

Each table in the admin-selected schema gets exposed as a top level route. Client requests are executed by certain database roles depending on their authentication. All HTTP verbs are supported that correspond to actions permitted to the role. For instance if the active role can drop rows of the table then the DELETE verb is allowed for clients. Here's an API request to delete old rows from a hypothetical logs table:

```
DELETE /logs?time=lt.1991-08-06 HTTP/1.1
```

However it's very easy to delete the **entire table** by omitting the query parameter!

```
DELETE /logs HTTP/1.1
```

This can happen accidentally such as by switching a request from a GET to a DELETE. To protect against accidental operations use the `pg-safeupdate` PostgreSQL extension. It raises an error if UPDATE or DELETE are executed without specifying conditions. To install it you can use the `PGXN` network:

```
sudo -E pgxn install safeupdate

# then add this to postgresql.conf:
# shared_preload_libraries='safeupdate';
```

This does not protect against malicious actions, since someone can add a url parameter that does not affect the result set. To prevent this you must turn to database permissions, forbidding the wrong people from deleting rows, and using row-level security if finer access control is required.

### 9.12.2 Count-Header DoS

For convenience to client-side pagination controls PostgREST supports counting and reporting total table size in its response. As described in *Limits and Pagination*, responses ordinarily include a range but leave the total unspecified like

```
HTTP/1.1 200 OK
Range-Unit: items
Content-Range: 0-14/*
```

However including the request header `Prefer: count=exact` calculates and includes the full count:

```
HTTP/1.1 206 Partial Content
Range-Unit: items
Content-Range: 0-14/3573458
```

This is fine in small tables, but count performance degrades in big tables due to the MVCC architecture of PostgreSQL. For very large tables it can take a very long time to retrieve the results which allows a denial of service attack. The solution is to strip this header from all requests:

```
-- Pending nginx config: Remove any prefer header which contains the word count
```

---

**Note:** In future versions we will support `Prefer: count=estimated` to leverage the PostgreSQL statistics tables for a fast (and fairly accurate) result.

---

### 9.12.3 HTTPS

See the *HTTPS* section of the authentication guide.

### 9.12.4 Rate Limiting

Ngix supports “leaky bucket” rate limiting (see [official docs](#)). Using standard Ngix configuration, routes can be grouped into *request zones* for rate limiting. For instance we can define a zone for login attempts:

```
limit_req_zone $binary_remote_addr zone=login:10m rate=1r/s;
```

This creates a shared memory zone called “login” to store a log of IP addresses that access the rate limited urls. The space reserved, 10 MB (10m) will give us enough space to store a history of 160k requests. We have chosen to allow only allow one request per second (1r/s).

Next we apply the zone to certain routes, like a hypothetical stored procedure called `login`.

```
location /rpc/login/ {
    # apply rate limiting
    limit_req zone=login burst=5;
}
```

The burst argument tells Ngix to start dropping requests if more than five queue up from a specific IP.

Ngix rate limiting is general and indiscriminate. To rate limit each authenticated request individually you will need to add logic in a *Custom Validation* function.

## 9.13 Debugging

### 9.13.1 Server Version

When debugging a problem it’s important to verify the PostgREST version. At any time you can make a request to the running server and determine exactly which version is deployed. Look for the *Server HTTP* response header, which contains the version number.

### 9.13.2 HTTP Requests

The PostgREST server logs basic request information to stdout, including the requesting IP address and user agent, the URL requested, and HTTP response status. However this provides limited information for debugging server errors. It’s helpful to get full information about both client requests and the corresponding SQL commands executed against the underlying database.

A great way to inspect incoming HTTP requests including headers and query params is to sniff the network traffic on the port where PostgREST is running. For instance on a development server bound to port 3000 on localhost, run this:

```
# sudo access is necessary for watching the network
sudo ngrep -d lo0 port 3000
```

The options to ngrep vary depending on the address and host on which you’ve bound the server. The binding is described in the *Configuration* section. The ngrep output isn’t particularly pretty, but it’s legible.

### 9.13.3 Database Logs

Once you've verified that requests are as you expect, you can get more information about the server operations by watching the database logs. By default PostgreSQL does not keep these logs, so you'll need to make the configuration changes below. Find `postgresql.conf` inside your PostgreSQL data directory (to find that, issue the command `show data_directory;`). Either find the settings scattered throughout the file and change them to the following values, or append this block of code to the end of the configuration file.

```
# send logs where the collector can access them
log_destination = "stderr"

# collect stderr output to log files
logging_collector = on

# save logs in pg_log/ under the pg data directory
log_directory = "pg_log"

# (optional) new log file per day
log_filename = "postgresql-%Y-%m-%d.log"

# log every kind of SQL statement
log_statement = "all"
```

Restart the database and watch the log file in real-time to understand how HTTP requests are being translated into SQL commands.

---

**Note:** On Docker you can enable the logs by using a custom `init.sh`:

```
#!/bin/sh
echo "log_statement = 'all'" >> /var/lib/postgresql/data/postgresql.conf
```

After that you can start the container and check the logs with `docker logs`.

```
docker run -v "$(pwd)/init.sh":/docker-entrypoint-initdb.d/init.sh -d postgres
docker logs -f <container-id>
```

---

### 9.13.4 Schema Reloading

Users are often confused by PostgREST's database schema cache. It is present because detecting foreign key relationships between tables (including how those relationships pass through views) is necessary, but costly. API requests consult the schema cache as part of *Resource Embedding*. However if the schema changes while the server is running it results in a stale cache and leads to errors claiming that no relations are detected between tables.

To refresh the cache without restarting the PostgREST server, send the server process a SIGUSR1 signal:

```
killall -SIGUSR1 postgrest
```

The above is the manual way to do it. To automate the schema reloads, use a database trigger like this:

```
CREATE OR REPLACE FUNCTION public.notify_ddl_postgrest()
RETURNS event_trigger
LANGUAGE plpgsql
AS $$
BEGIN
```

(continues on next page)

(continued from previous page)

```

NOTIFY ddl_command_end;
END;
$$;

CREATE EVENT TRIGGER ddl_postgrest ON ddl_command_end
EXECUTE PROCEDURE public.notify_ddl_postgrest();

```

Then run the `pg_listen` utility to monitor for that event and send a SIGUSR1 when it occurs:

```
pg_listen <db-uri> ddl_command_end "killall -SIGUSR1 postgrest"
```

Now, whenever the structure of the database schema changes, PostgreSQL will notify the `ddl_command_end` channel, which will cause `pg_listen` to send PostgREST the signal to reload its cache.

## 9.14 Daemonizing

For linux distros that use **systemd** (ubuntu, debian, archlinux) you can create a daemon in the following way.

First, create postgrest configuration in `/etc/postgrest/config`

```

db-uri = "postgres://<your_user>:<your_password>@localhost:5432/<your_db>"
db-schema = "<your_exposed_schema>"
db-anon-role = "<your_anon_role>"
db-pool = 10

server-host = "127.0.0.1"
server-port = 3000

jwt-secret = "<your_secret>"

```

Then create the systemd service file in `/etc/systemd/system/postgrest.service`

```

[Unit]
Description=REST API for any Postgres database
After=postgresql.service

[Service]
ExecStart=/bin/postgrest /etc/postgrest/config
ExecReload=/bin/kill -SIGUSR1 $MAINPID

[Install]
WantedBy=multi-user.target

```

After that, you can enable the service at boot time and start it with:

```

systemctl enable postgrest
systemctl start postgrest

## For reloading the service
## systemctl restart postgrest

```

## 9.15 Alternate URL Structure

As discussed in *Singular or Plural*, there are no special URL forms for singular resources in PostgreSQL, only operators for filtering. Thus there are no URLs like `/people/1`. It would be specified instead as

```
GET /people?id=eq.1 HTTP/1.1
Accept: application/vnd.pgrst.object+json
```

This allows compound primary keys and makes the intent for singular response independent of a URL convention.

Nginx rewrite rules allow you to simulate the familiar URL convention. The following example adds a rewrite rule for all table endpoints, but you'll want to restrict it to those tables that have a numeric simple primary key named "id."

```
# support /endpoint/:id url style
location ~ ^/([a-z_]+)/([0-9]+) {

    # make the response singular
    proxy_set_header Accept 'application/vnd.pgrst.object+json';

    # assuming an upstream named "postgrest"
    proxy_pass http://postgrest/$1?id=eq.$2;

}
```

- *Authentication*
- *Installation*
- *Administration*

PostgREST has a growing ecosystem of examples, libraries, and experiments. Here is a selection.

### 10.1 Example Apps

- [tatut/postgrest-ui](#) - ClojureScript UI components for PostgREST
- [priyank-purohit/PostGUI](#) - React Material UI admin panel
- [Qu4tro/pgrst-dev-setup](#) - docker-compose and tmuxp setup for experimentation.
- [subzerocloud/postgrest-starter-kit](#) - Boilerplate for new project
- [NikolayS/postgrest-google-translate](#) - Calling to external translation service
- [CodeforAustralia/heritage-near-me](#) - Elm and PostgREST with PostGIS
- [timwis/handsontable-postgrest](#) - An excel-like database table editor
- [Recmo/PostgrestSkeleton](#) - Docker Compose, PostgREST, Nginx and Auth0
- [benoror/ember-postgrest-dynamic-ui](#) - generating Ember forms to edit data
- [ruslantalpa/blogdemo](#) - blog api demo in a vagrant image
- [timwis/ext-postgrest-crud](#) - browser-based spreadsheet
- [srid/chronicle](#) - tracking a tree of personal memories
- [diogob/elm-workshop](#) - building a simple database query UI
- [myfreeweb/moneylog](#) - accounting web app in Polymer + PostgREST
- [tyrchen/goodfilm](#) - example film api
- [begriiffs/postgrest-example](#) - sqitch versioning for API
- [SMRxT/postgrest-demo](#) - multi-tenant logging system
- [PierreRochar/postgrest-boilerplate](#) - example auth back-end

- [marmelab/ng-admin-postgrest](#) - automatic database admin panel

## 10.2 External Notification

These are PostgreSQL bridges that propagate LISTEN/NOTIFY to external queues for further processing. This allows stored procedures to initiate actions outside the database such as sending emails.

- [diogob/postgres-websockets](#) - expose web sockets for PostgreSQL's LISTEN/NOTIFY
- [frafra/postgresql2websocket](#) - Websockets
- [matthewmueller/pg-bridge](#) - Amazon SNS
- [aweber/pgsql-listen-exchange](#) - RabbitMQ
- [SpiderOak/skeeter](#) - ZeroMQ
- [FGRibreau/postgresql-to-amqp](#) - AMQP
- [daurnimator/pg-kinesis-bridge](#) - Amazon Kinesis

## 10.3 Extensions

- [pg-safeupdate](#) - Prevent full-table updates or deletes
- [srid/spas](#) - allow file uploads and basic auth
- [svmnotn/postgrest-auth](#) - OAuth2-inspired external auth server
- [wildsurfer/postgrest-oauth-server](#) - OAuth2 server
- [nblumoe/postgrest-oauth](#) - OAuth2 WAI middleware
- [criles25/postgrest-auth](#) - email based auth/signup
- [ppKrauss/PostgREST-writeAPI](#) - generate Nginx rewrite rules to fit an OpenAPI spec

## 10.4 Client-Side Libraries

- [Kong/py-postgrest](#) - Python
- [datrium/postgrest-pyclient](#) - Python
- [tomberek/aor-postgrest-client](#) - JS, admin-on-rest
- [hugomrdias/postgrest-url](#) - JS, just for generating query URLs
- [john-kelly/elm-postgrest](#) - Elm
- [mithril.postgrest](#) - JS, Mithril
- [lewisjared/postgrest-request](#) - JS, SuperAgent
- [JarvusInnovations/jarvus-postgrest-apikit](#) - JS, Sencha framework
- [davidthewatson/postgrest\\_python\\_requests\\_client](#) - Python
- [calebmer/postgrest-client](#) - JS
- [clesiemo3/postgrestR](#) - R

- PierreRochard/postgrest-angular - TypeScript, generate UI from API description
- thejettdurham/postgrest-sharp-client (needs maintainer) - C#, RestSharp
- team142/ng-postgrest - Angular app for browsing, editing data exposed over Postgrest.

## 10.5 Commercial

- subZero - Automated GraphQL & REST API with built-in caching (powered in part by PostgREST)
- *Example Apps*
- *External Notification*
- *Extensions*
- *Client-Side Libraries*
- *Commercial*



# CHAPTER 11

---

## Release Notes

---

Here we'll include the most relevant changes so you can migrate to newer versions easily. You can see the full changelog of each release in the [PostgREST repository](#).

- *v6.0.2*
- *v5.2.0*



## CHAPTER 12

---

### In Production

---

Here are some companies that use PostgREST in production.

- Datrium
- Nimbus - See how Nimbus uses PostgREST in [Paul Coplestone's blog post](#).
- Catarse
- Moat
- Redmin
- Image-charts
- MotionDynamic - Fast highly dynamic video generation at scale
- Drip Depot
- Convene by Thomson-Reuters
- eGull
- Elyios
- Simply Connected Systems



# CHAPTER 13

---

## Testimonials

---

“It’s so fast to develop, it feels like cheating!”

—François-G. Ribreau

“I just have to say that, the CPU/Memory usage compared to our Node.js/Waterline ORM based API is ridiculous. It’s hard to even push it over 60/70 MB while our current API constantly hits 1GB running on 6 instances (dynos).”

—Louis Brauer

“I really enjoyed the fact that all of a sudden I was writing microservices in SQL DDL (and v8 javascript functions). I dodged so much boilerplate. The next thing I knew, we pulled out a full rewrite of a Spring+MySQL legacy app in 6 months. Literally 10x faster, and code was super concise. The old one took 3 years and a team of 4 people to develop.”

—Simone Scarduzio

“I like the fact that PostgREST does one thing, and one thing well. While PostgREST takes care of bridging the gap between our HTTP server and PostgreSQL database, we can focus on the development of our API in a single language: SQL. This puts the database in the center of our architecture, and pushed us to improve our skills in SQL programming and database design.”

—Eric Bréchemier, Data Engineer, eGull SAS

“PostgREST is performant, stable, and transparent. It allows us to bootstrap projects really fast, and to focus on our data and application instead of building out the ORM layer. In our k8s cluster, we run a few pods per schema we want exposed, and we scale up/down depending on demand. Couldn’t be happier.”

—Anupam Garg, Datrium, Inc.



# CHAPTER 14

---

## Translations

---

- [Chinese \(latest version v0.4.2.0\)](#)