

WenQuanYi Micro Hei [Scale=0.9]WenQuanYi Micro Hei Mono song-
WenQuanYi Micro Hei sfWenQuanYi Micro Hei "zh" = 0pt plus 1pt

PostgREST Documentation

åRŠåÿČ 4.1.0

Joe Nelson

8æIJL 06, 2017

Contents

1	áčřæŸŌæĀğcijÚçlŃ	2
2	Leak-proof æL;èsa	3
3	æÑěæLśâĚşçşzælađŃ	4
4	äÿÄäÿléĜmçĆz	5
5	æŤzè£ZâĚsāzn	6
6	çŤšæĀçşçzş	7
6.1	áčæLúçnrážŞ	7
6.2	éçlād' ŪéĀZçşě	8
6.3	çd'zäçNāzŤçŤl	8
6.4	Production	9
6.5	æŃŞásŤ	9
6.6	áz£âŞŁ	9
7	èŧđèl'	10
8	èŌuâçŪæŤræŃA	11
9	Tutorial 0 - èŏl'áoŃĈèùSèŧuæIě	12
9.1	Step 1. æŤçè;çæİçèĀAéŞA, æLŚsāznāijZāyŏä;áčZĎ	12
9.2	Step 2. áŏL'èçĚ PostgreSQL	13
9.3	Step 3. áŏL'èçĚ PostgREST	13
9.4	Step 4. äÿž API áLZâzæŤræmŏžŞ	14
9.5	Step 5. è£ŘèqŃ PostgREST	15
10	Tutorial 1 - éĜSéŞěâŃZ	17

10.1	Step 1. æúzáŁääÿÄäyĺáRŮäŁçŽDčŤĺæLŮ	17
10.2	Step 2. çŤšæLŘäÿÄäyĺáŕEçăA	18
10.3	Step 3. çŤšæLŘ token	18
10.4	Step 4. èŁZèaŇèŕŮæśĆ	18
10.5	Step 4. æúzáŁæèŁĜæIJšæŮúéŮŤ	20
10.6	éZĎáŁæécŸ: çŇŇá■şæŠd' éŤĂ	21
11	âŔŕæL'gèaŇæŮĜäzŮ	23
12	Homebrew	24
13	PostgreSQL ä;ìèŤŮ	25
14	æžŘäzčçăAçijŮèŕŚ	26
14.1	PostgREST æŤŇèŕŤæŮäzŮ	27
15	éĚ■ç;õ	30
15.1	âŔŕáŕLÍ Server	32
16	çañâŇŮ PostgREST	33
16.1	éŸzæ■çâĚlèaŁæŞ■ä;IJ	34
16.2	Count-Header DoS	34
16.3	HTTPS	35
16.4	éZŔéĂş	35
17	èŕČèŕŤ	37
17.1	Schema éĜ■è;ij	38
18	âĎĜçŤÍ URL çzŞæĎĎ	39
19	èaĺ & èĝEaZ;ç	40
19.1	æŕŤ' âzşèŁĜæzd' (Rows)	40
19.2	âĎČçŽŤ' èŁĜæzd' (Columns)	42
19.3	æŎŠâžŔ	42
19.4	Limit âŠŇâLÉéaŤ	43
19.5	çZÿâžŤæäijâijŔ	44
19.6	â■ŤæŤŕæLŮâĎ■æŤŕ	45
19.7	âžŇèŁZáLŮè;ŞâĜž	45
20	èŤĎæžŔaŤŇâèŮ	47
20.1	âŤŇâĚèèŁĜæzd' âŠŇæŎŠâžŔ	49
21	èĜlăõŽäzL'æşèèŕç	51
22	â■ŸâCíèŁĜçlŇ	52
22.1	èŎŮâŔŮèŕŮæśÇžĎ Headers/Cookies	53
22.2	âĎ■æiCâyČârŤéĂzè;Ś	54

22.3	æŁZãĜžéŤŽèřř	54
23	æŘŠãĚě/ãĚŃãŤž	56
23.1	æŁžéĜŘæŘŠãĚě	57
24	ãĹãéŽď	59
25	OpenAPI æŤřæŇA	60
26	HTTP çĹúæĀçãA	61
27	èĝŠèĹšçşçzşæçCèřř	63
27.1	éĪŇèřAžŘãĹŮ	63
27.2	çŤĹæĹúãŠŇçžĎ	65
27.3	èĜĹãŃžãžĹéĪŇèřA	67
28	ãŃçæĹúçńř Auth	68
28.1	JWT Generation	68
28.2	SSL	71
29	æďúæďĎéŽŤçž	72
30	SQL çŤĹæĹúçŃçŘE	73
30.1	ãŃŸãĀçĹŤĹæĹúãŠŇãřEçãA	73
30.2	Public çŤĹæĹúçŤŇéĹc	75



PostgREST je webový nástroj, který umožňuje poskytnout REST API z PostgreSQL databáze. Je to jednoduchý a rychlý způsob, jak získat REST API z jakékoli PostgreSQL databáze.

PostgREST umožňuje provádět CRUD operace na datích v PostgreSQL databázi. Je to jednoduchý a rychlý způsob, jak získat REST API z jakékoli PostgreSQL databáze. PostgREST je jednoduchý a rychlý způsob, jak získat REST API z jakékoli PostgreSQL databáze.

CHAPTER 1

ĀčřæŸŌæĂğcijŪçlŃ

æřTèřũárzçlĀæšèèrcçzšæđlJéĜ■ād'■āŁşāLlīijNèđl æšèèrcèđāāLšçRĒæyĔçzĒèŁĆä;£ Post-
greSQL äyžæĆlæũzāLăæŦřæ■ōæŸřäyĂäzũæZt' āóžæŸšçŽĐäžŃăĂĆäyžæŦřæ■ōžšărzèsāāLĒéĔ■ælČéZŘæřTā

CHAPTER 2

Leak-proof æŁjèsą

æšąæIJL ORM áRCäyÓãĂCáLZázž SQL æĂğèĈ;æRRçd'žçŽDèğEáZ;ãĂCæŦræ■óážŞçóaçŘEáŚŸijLDAi
APIijŇæŮáéIJAáijĂáRŚáRCäyÓãĂC

CHAPTER 3

æNěæŁśăĔşçşzæłąđŃ

1970ázt'rijŃE. F. Codd áIĴázŮçŽDæŮĜçná âAIĴád' ġáđŃăĔśăznæŤřæ■óázŞæŤřæ■óăĔşçşzæłąđŃăĀI
äy■æL' zèrĎăžEă;ŞæŮüäyzârijçŽDæŤřæ■óázŞăĴEăśCăĴăđŃăĀCăžŌéC;æŮĜçnáaijŽăŔŚçŌŕăśCăñăæŤřæ■óáz
http èŭŕçŤśăzŃéŮt' â■ŸăIĴăĀĴăžçŽyâijjæĀġăĀCèĀŃăĴĴ PostgREST
äy■iiĴŃăĴŚăzňăŕĴă;ĴçŤĴçAĴæt' zçŽĐèĴĜæzd' âŞŃăĴŃăĔĕiiĴŃèĀŃăy■æŸŕăĴŃăeŮèŭŕçŤśăĀC

CHAPTER 4

ÄyÄytleG■çZ

PostgREST æIJL' äyÄytleG■çZDçL' zçCzãÄCãóČéÄCçTlãžŎãçC Ng-
inx çŽDãũëãËũãÄCèŁZãRfãzëãijžëãNãrEãzëãTãrã■öãÿzãÿ■ãŁČçŽD CRUD
æS■ã;IJãÿŎãËũãzÛéUóécÿèŁZëãNãzšãGããLEççzãÄC

CHAPTER 5

æŤzèŁZăĚśăžń

äyŎäzza;ŤâijĂæžRéazçZóäyĂæăüijŊæŁSăžněČ;ăRfăzèăzŎăũěăEuăy■çZDăŁşèČ;ăŠŊăŁóăd'■ăy■èŎŭçZŁ

CHAPTER 6

PostgreSQL

PostgreSQL is a powerful, open source object-relational database system with over two decades of active development. It has many extensions for data types, routines, and indexes, and is highly extensible.

Client Libraries

- [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

Úvod

Úvod

Úvod

- [frafra/postgresql2websocket](#) - Websockets
- [matthewmueller/pg-bridge](#) - Amazon SNS
- [aweber/pgsql-listen-exchange](#) - RabbitMQ
- [SpiderOak/skeeter](#) - ZeroMQ
- [FGRibreau/postgresql-to-amqp](#) - AMQP

Úvod

- [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
- [marmelab/ng-admin-postgrest](#) - automatic database admin panel
- [myfreeweb/moneylog](#) - accounting web app in Polymer + PostgREST
- [tyrchen/goodfilm](#) - example film api
- [begriffs/postgrest-example](#) - sqitch versioning for API
- [SMRxt/postgrest-demo](#) - multi-tenant logging system
- [PierreRochard/postgrest-boilerplate](#) - example auth backend

Production

- [Catarse](#)
- [iAdvize](#)
- [Redsmin](#)
- [Image-charts](#)
- [Drip Depot](#)
- [OpenBooking](#)
- [Convene by Thomson-Reuters](#)
- [eGull](#)

æNŞåŞT

- [ppKrauss/PostgREST-writeAPI](#) - generate Nginx rewrite rules to fit an OpenAPI spec
- [diogob/postgrest-ws](#) - expose web sockets for PostgreSQL's LISTEN/NOTIFY
- [pg-safeupdate](#) - Prevent full-table updates or deletes
- [srid/spas](#) - allow file uploads and basic auth
- [svmnotn/postgrest-auth](#) - OAuth2-inspired external auth server
- [nblumoe/postgrest-oauth](#) - OAuth2 WAI middleware

åZ£åŚŁ

- [subZero](#) - Automated GraphQL & REST API with built-in caching (powered in part by PostgREST)

CHAPTER 7

ètdèl'

“ájĀāRSètuæIēād' ĩāñāžE, æDšèğL' āřsāČRāIJĭ;IJĭájL!”

—François-G. Ribreau

“æLSäy■ā;Uäy■èrt', äyÓ Node.js/Waterline ORM ædDāžžçŽD API āřzærT

CPU/Memory usage çóĀçŽt' æYréZ; äžç;ōäĕā. ā;šæLSäžñāIJĭ 6 äyłçd' žā;N ĩijLdynosiiijL' æNĀçž■æšCæĈĒ 1GB æTřæ■óæYřáóČçTŽèGšāRĭæIJL' 60/70 MB ād' ġārR.”

—Louis Brauer

“æLSéIdāyÿāŪIJæñcèŁZæüüyĀäyĭāžNāóđĭijNāAúçDūā;ŁçTĭ SQL
DDLĭijLāŠŃ V8 javascriptĭijL' āijĀāRSā;ōæIJ■āŁāāĀC æLSäžñāIJĭ 6
äyĭæIJLāEĒāóNāĒÉĜ■āEŽāžEäyĀäyĭ Spring + MySQL éAŪçTŽāžTçTĭçĭNāžRāĀC
éĀšāžēāĕñ 10 āĀ■ĭijNāžççāAā;ŁçóĀæť AāĀĈèĀNāžNāL'■çŽDāžžçTĭläžE 4
äyĭāžžèŁsäžE 3 āžť æŪúéŪť āĀĀC”

—Simone Scarduzio

CHAPTER 8


èÕùå; ÛæŤræÑÀ

èfèéazçZóæIJL' äyÄäyłáRNâe; äyŤäy■æŮ■æLRéŤçZDçd' ; åÑzãĂĆâŁääĚæŁSäzñçZD
èAŁád' r' áóđ' æİèèóÍèõžãŠÑæśĆâŁ' ãĂĆâRÑæŮüä; äžšâRřäzèâIJÍ Github çZD issues
äyŁæŘIJçt' ć bugs/featuresãĂĆ

Tutorial 0 - èó'áóÇèùSèùæläě

æñçè£Ōä;£çŦÍ PostgRESTijAèrèáL'■èl'ÄæYřäyÄäyŦ Quick
startiijNäyóáL'ä;ää£néÄšáL'ZázžçñnäyÄäyŦçóÄ■ŦçŽĐ APIäÄĆ

PostgREST æYřäyÄäyŦçNñçñŦçŽĐ Web æIJ■áLäáZlíijNäyž PostgreSQL
æŦřæ■óázšçŦšæL'Ŧ RESTful APIäÄĆ áóÇæRŦä;ZášžžÖázŦásÇæŦřæ■óázšçzšæđDáoZáLúçŽĐ
APIäÄĆ



æÇšèeAçŦšæL'Ŧ APIiijNæL'SázñäRléIJÄèeAáL'ZázžäyÄäyŦæŦřæ■óázšäÄĆæL'ÄæIJL'çnrçCzáŠNæIČéZŦé
áIJLæIJñæŦžçl'NçzšæIšçŽĐæÚúáZíijNæCíárEæNèæIJL'äyÄäyŦèC;çŦÍçŽĐæŦřæ■óázšŦiijNPostgREST
æIJ■áLäáZlíäŠNäyÄäyŦçóÄ■ŦçŽĐä■ŦçŦÍæLú todo list APIäÄĆ

Step 1. æŦç;è;žælä;èÄÄéŠA, æL'SázñäijŽäyóä;äçŽĐ

áIJLä;ääijÄägNè£ZäyŦæŦžçl'NæÚú, Ctrl+T äyÄäyNáIJLæŦřæäGç;ç;äy■æL'ŠäijÄéazçŽó
èAŁäd'Ŧ'áóđ' . æIJL'äyÄç;ç;ä;L nice çŽĐäžžáIJLèAŁäd'Ŧ'áóđ' äy■æŦ'žèüČiijNæçÄèIJä;ää■ä;RäžEæL'SázñäijŽä

Step 2. Installing PostgreSQL

For Linux/BSD/Windows, you can use Docker to run PostgreSQL 9.3. The Docker image is available on Docker Hub. The image is named postgres:9.3. The image is based on the Ubuntu 14.04 LTS image. The image is available on Docker Hub.

To run PostgreSQL in Docker, you need to have Docker installed on your system. You can find the installation instructions for Docker on the Docker website.

Run the following command to start a PostgreSQL container:

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

After running the command, you can access the PostgreSQL server from your host machine. The container will be named tutorial and will be listening on port 5432.

Step 3. Installing PostgREST

PostgREST is a REST API for PostgreSQL. It is written in Haskell and is available on Linux/BSD/Windows. You can find the installation instructions for PostgREST on the PostgREST website.

To install PostgREST, you need to have Haskell and Cabal installed on your system. You can find the installation instructions for Haskell and Cabal on the Haskell website.

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

After downloading the PostgREST binary, you can run it from the command line. The binary is named postgrest.exe on Windows and postgrest on Linux/BSD/Windows.

```
./postgrest
```

For Linux/BSD/Windows, you can use Docker to run PostgREST. The Docker image is available on Docker Hub. The image is named postgrest:latest. The image is based on the PostgreSQL 9.3 image.

PostgREST is a REST API for PostgreSQL. It is written in Haskell and is available on Linux/BSD/Windows. You can find the installation instructions for PostgREST on the PostgREST website.

ror while loading shared libraries: libpq.so.5”

Step 4. Creating the API Schema

Run the following SQL in the PostgreSQL shell:

```
sudo docker exec -it tutorial psq1 -U postgres
```

Run the following SQL in the PostgreSQL shell:

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

postgres=#
```

Run the following SQL in the PostgreSQL shell to create the API schema:

```
create schema api;
```

Run the following SQL in the PostgreSQL shell to create the todos 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');
```

Run the following SQL in the PostgreSQL shell to create the web_anon role and grant permissions:

```
create role web_anon nologin;
grant web_anon to postgres;

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

Run the following SQL in the PostgreSQL shell to create the web_anon role and grant permissions:

Článek psal jsem pomocí API aplikace

```
\q
```

Step 5. Nastavení PostgREST

PostgREST aplikaci nastavíme pomocí konfigurace `tutorial.conf` podle následujícího příkladu:

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

Pro spuštění aplikace použijeme příkaz `options --config tutorial.conf`:

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

Ukázka výstupu:

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

Pro otestování API použijeme příkaz `curl` v terminálu (pro spuštění PostgREST aplikace použijeme příkaz `options --config tutorial.conf`):

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

API vrátí:

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

```
}
]
```

`todos`
`todo`
`aijZècnaNSçziãÄ`

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

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!

CHAPTER 10

Tutorial 1 - éĜŚéŠěãŇŽ

âĪĪ *Tutorial 0 - èól'áoĈèùSèřuæĪě* äy■æĹSázňáĹZázžazĒäyÄäyĚèŌúâRŪ
todos æřræ■óâRĪèřzĉŽĎ APIãĀĈä;ĚçřĪĪ■řäyĪçñřĉĈžáĹŪâĜž todosãĀĈ
æĹSázňæĪĪĹ'â;Ĺâd'ŽâĹđæşřrâRázěä;ĚèĚŽäyĪ API æŽr' æĪĪĹ'èúĉĪĪjNä;ĒäyÄäyĪâč;ĉŽĎâĪjĀăĜNæŸrâĒĀèőyäyĀă

Step 1. æúžâĹăäyÄäyĪâRŪäĒçžĎĉřĪæĹú

äyĹäyÄèĹĈäy■ĪĪjNèĚZèqNâNĚâř■ Web èrúæśĈĉŽĎäžNâRŌâĪĪæřræ■óâžšäy■âĹZázžazĒäyÄäyĪ
web_anon èĝSèĹ'sãĀĈèól' æĹSázňâĪĪĹ'âĹZázžäyÄäyĪèĝSèĹ'sâRnáAž todo_user
ĉřĪläžŌâ;ĚçřĪĪ API èĚZèqNèžňáz;ĒĪNèřAçžĎĉřĪæĹúĪĪjNèĚŽäyĪèĝSèĹ'sâřĒæĪĪĹ'æĪĈârž todo
list âAžžäzä;řžňæĈĒãĀĈ

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

create role todo_user nologin;
grant todo_user to postgres;

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;
```

Step 2. çTšæLŘäYÄäyĽarEçäA

ãóæLúçnréÅžèĚĜ API ä;çTĪ JSON Web Token èĚZèaÑèznáz;éĪÑèrAāĀĆJTW æYřä;çTĪläzĒæIJL æLŠäznāŠÑæIJāLāāZĪçšééAšçŽDārEçäAèĚZèaÑāLāārEç;ĶāRçŽD JSON āřzèšāāĀĆ çTšäzŌãóæLúçnräyçšééAšçārEçäAĪijÑæL'ÄäzèäyēĈ;çřæTž token çŽDāEĒāóžāĀĆ PostgREST äijZæçÄæĪNäijléÄçŽD token äzúæNšçzĪāóČäznāĀĆ

æLŠäznæĪēāLZäzžäyÄäyĽarEçäAāžúæRŘä;ZçžZ PostgRESTāĀĆæIJāāē;æČšæČšäyÄäyĽād' ■æiCçŽDēĚĚ

æšĽèġč: OpenSSL toolkit æRŘäyĽäyÄäyĽçóĀā■ĚçŽDæŪžāijRæĪççTšæLŘāóL'āĒĪçŽDārEçäAāĀĆæČæĈæĪJä;äæĪ

```
openssl rand -base64 32
```

æL'ŠāijÄ tutorial.conf (āĪĪäyĽäyÄèĚČäyāLZäzžçŽD) äžúārEārEçäAæúžāLāāIJĪæŪřçŽDäyÄèāÑ:

```
# add this line to tutorial.conf
jwt-secret = "<the password you created>"
```

æČæĈæĪJ PostgREST server äž■æŪgāĪĪèĚĚrèāÑäy■ĪijÑéCčázĹéĪJĀèēAéĜ■āRřāóČäzèä;ĚāĽæĪ;æIJæŪřçž

Step 3. çTšæLŘ token

éŽäyÿä;äèĜĪāūsçŽDäzççāAāĪĪæĚřæ■óāžšæLŪāĒūāzŪæIJāLāāZĪläy■ārEāLZäzžāžúç;ç;šèznáz;éĪÑèrA tokenĪijÑä;EæYřāĪĪæIJĪæĚçĪNäy■ĪijÑæLŠäznārEāĀĪEĒĜĪāūsāĹæL'NāĀĪāĀĆèùšè;ñāĽř jwt.ioĪijÑāžúāāñāEŽæĈäyNā■ŪæōĪijž

èřúèōrā;RæČĪāāñāEŽçŽDārEçäAĪijÑèĀÑäy■æYřāZ;çĹ'ĜéĜNçŽD secretāĀĆāāñāEŽārEçäAāŠÑ payload äžNārŌĪijÑāuēä;ġçŽDçijŪçäAæĚřæ■ōāijžĀĹūæŪřĪijÑèræĚřæ■óāš token ād' ■āĹūāóČāĀĆ

æšĽèġč: èZ;çDŪäzd'çĹNārřèĈ;çIJNèĚūæĪēā;ĹæĪççĹĪijÑä;Eā;ĹāóžæYšéĀEārŠāĜžçŽD payloadāĀĆtoken äžĒäzĒæYřèçñç;ĶāR■ĪijÑæšāæIJL'āLāārEĪijÑæL'ÄäzèæĈæĪJä;äæIJL'äy■æČšèōĹ'ãóæLúçnrçĪJ

Step 4. èĚZèaÑèrúæšĆ

āžđāĽř terminalĪijÑæLŠäznæĪççTĪ curl æúžāLāāyÄäyĽ todoāĀĆèrèèrúæšĆārEāNĒæNñäyÄäyĽāNĒārNèznáz;éĪÑèrA token çŽD HTTP ād't'āĀĆ

Encoded PASTE A TOKEN HERE

```
eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJyb2x1IjoiaW9kbG9kb191c2VyIn0.IF9PKgLZ_XA70Uz5vm00xrqYTCQdXxBA00z4uK8lBqM
```

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), <input type="text" value="secret"/>) secret base64 encoded

1. Enter password →

2. Enter this JSON →

3. Copy the resulting token ↑

Figure 10.1: Generating a JWT Token via https://jwt.io

```
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"}'
```

The response from the PATCH endpoint will be:

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

The response from the GET endpoint will be:

```
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
}
]

```

Step 4. æúžáŁæèŁGæIJSæÚúéÚť

çŽóáL'■īījNæLSäzñçŽDèòd'èřA token áržäžŎæL'ÄæIJL'èřúæśCéČ;æÝřäyÄèGť æIJL'æŤLçŽDãĂCæIJ■áŁæJWT árEçăAīījNāřšäijŽéĂŽèŁGéI'NèřAãĂC

æŽť äè;çŽDç■ÚçŤæÝřèól' token ä;ŁçŤÍ exp äčřæÝŎäyÄäyłèŁGæIJSæÚúéÚť æŁšãĂCèŁŽæÝř PostgREST çL'žáL'náržâ;ĚçŽDäy'd'äył JWT äčřæÝŎäzNäyÄãĂC

Claim	Interpretation
role	The database role under which to execute SQL for API request
exp	Expiration timestamp for token, expressed in “Unix epoch time”

æšłèğč: Unix æÚúéÚťæŁš (Unix epoch time) ècńáòŽázL'äyžèGť 1970 ážť 1 æIJL 1 æÚě 00:00:00 á■RèřČäyÚçŤNæÚīījLUTCīījL'äžèæIéáLřçŎřáIJL'çŽDæĂžçğŠæŤřīījNäy■èĂČèŽŚéUřçğŠãĂC

äyžäzEáIJlèaŤNáLlāy■èğČáršèŁGæIJSīījNæLSäzñāřEæúžáŁæäyÄäyłáIJl 5min äžNāRŎèŁGæIJSçŽD exp äčřæÝŎăĂCéçÚăĚLæL'ġáLřäzŎā;ŠáL'■æÚúéÚť çóUèťúáLř 5min äžNāRŎçŽDæÚúéÚť æŁšãĂC áIJl psql äy■èŁRèaŤīījŽ

```

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

```

ážďáLř jwt.io ážúüäŁæŤž payload

```

{
  "role": "todo_user",
  "exp": "<computed epoch value>"
}

```

æNúèť IæŤřçŽD tokenīījNçDúăRŎăřEăĚüăŁIá■YäyžäyÄäyłæŤřçŽDçŎřácČăRÝéGRăĂC

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

Uvedeme si novy token a zkusime ho pouzít:

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

Uvedeme si novy token a zkusime ho pouzít: API vrati HTTP 401 Unauthorized:

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

Revokace tokenu: role a email

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>"
```

PostgREST 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 \  
  -H "Authorization: Bearer $TOKEN"  
  
# this one is rejected  
curl http://localhost:3000/todos \  
  -H "Authorization: Bearer $WAYWARD_TOKEN"
```

The server responds with 403 Forbidden:

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

CHAPTER 11

Installing PostgreSQL

[*ayNèj;éat*]

*ayNèj;éat*élcáĚúæIJL Mac OS XāĀWindows āŠNāĜäyġ Linux
āRŠèaŃçL'ŁçZĎécDçijŪerSæŪĜäzúāĀCèġcāŌNāzNāRŌāRfāzèèĚRèaŃāRfæL'ġèaŃæŪĜäzúāĻă
--help æāĜåĤŪæĪæšëçIJŃä;ĤçŦĪèŕ' æŸŌ:

```
# èġčāŌŃ tar āŃĚ (available at https://github.com/begriffs/postgrest/  
→releases/latest)  
  
$ tar Jxf postgrest-[version]-[platform].tar.xz  
  
# āřĪèŕŦèĚRèaŃ  
$ ./postgrest --help  
  
# You should see a usage help message
```

CHAPTER 12

Homebrew

âĪĪ Mac äÿLä;ããRřázěä;£çTĪ Homebrew æIěãóL'èčĚ PostgREST

```
# çąóăĹĪ brew æÿræIJĂæŮřçŽĎ  
brew update  
  
# æčĂæSě brew çŽĎ setup æIJL'æšæéŮóéćÿ  
brew doctor  
  
# áóL'èčĚ postgres  
brew install postgres
```

ěřěáŠ;äzd'äijŽeĜlâLlärE PostgreSQL ä;ŠâAŽä;IèťŮáóL'èčĚ.
ěřěè£ĜćlNâ;Ăâ;ĂéIJĂèçAéTfè; 15áLEéŠšæL■èČ;áóL'èčĚ;řázúãNěãRĹãĚüä;IèťŮãĀĆ
áóL'èčĚáóNæLRăRŌijNěřěãüèãĚüäijŽèćnæúžâLääLř \$PATH
äÿ■ijNä;ããRřázěãĪĪlázžæĎRä;■ç;óä;£çTĪijŽ

```
postgres --help
```

CHAPTER 13

PostgreSQL ä; lètÚ

9.3 [PostgREST](#) [æCÍlárEéIJÄðeAåL'ècĚæTřæ■óžšijLPostgreSQL](#)
[æLÚæZt'énYçL'LaIJñijL'ãĀĆ](#) [æCÍlárRäzëä;çTÍlâĀĀ](#) [Amazon RDS](#)
[èfZæüçZDäyIJeëfiijNä;EæYřaIJæIJñâIJrãL'ècĚæIJñèznærTè;Čä;£åóIJñijNæZt'ä;£äžŌaijĀãRŠãĀĆ](#)

- [OS X èrt'æYŌ](#)
- [Ubuntu 14.04 èrt'æYŌ](#)
- [Windows åL'ècĚãĀĀ](#)

CHAPTER 14

æžŘäzččäAçijŮèrŚ

æšlèġč: æĹŚäznäy■āzžèōōāIJĪ **Alpine Linux** äyŁæđDāzžāŠNä;ŁçTĪ Post-gRESTiijNāZāyžāIJlèrēāzšāRræIJL'èŁĠ GHC āĒĒā■ŸæšDæijRçŽDæŁēāŚĹāĀĆ

ā;ŠæĆĪçŽDçšzçzšæšæIJL'écDæđDāzžçŽDāRræL'gèāNæŮĠzūæŮiijNāRræžēāzŌæžŘäzččäAæđDāzžēāz Stack āĀĆāōCārĒāIJĪæĆĪçŽDçšzçzšäyŁāōL'èčĒäzä;ŤāĒĒèçAçŽD Haskell ä;ĪèŤŮāĀĆ

- āōL'èčĒ Stack
- āōL'èčĒä;ĪèŤŮāzš

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

- æđDāzžāzūāōL'èčĒ

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

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

æşlēċ: åĉĆæđIJă;ăæđĎăzžâd'sèt'ëiijÑèĂÑăyŦă;ăçŽĎçşzçzşâRlæIJL'ây■ăLř 1GB
 åĒĒă■ŸiijÑârİerŦæûzâLăăyĂăyŦ swap æŪĠăzŭăĂĆ

- æċĂæşēăđL'èċĒæŸrâRæLŘăLş: `postgrest --help`.

PostgREST æŦÑerŦăĕŪăzŭ

ăLŽăzžæŦÑerŦăžŞ

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 Postgres 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 postgres 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.

Deleting the test database

The test database will remain after the test, together with four new roles created on the *postgres* 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
```

Running the tests 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 MacOS 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 Postgres.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 "POSTGREST_TEST_CONNECTION=$(test/create_test_db
↳ "postgres://postgres:pwd@pg" test_db) stack test"
```

Stack test in Docker for Mac, Postgres.app on mac:

```
$ host_ip=$(ifconfig en0 | grep 'inet ' | cut -f 2 -d' ')
$ export POSTGREST_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=$host_ip" -e
↳ "POSTGREST_TEST_CONNECTION=$POSTGREST_TEST_CONNECTION" -w="`pwd`"
↳ pgst-test bash -c "stack test"
$ test/destroy_test_db "postgres://postgres@localhost" test_db
```

CHAPTER 15

éĚ■çjő

The PostgREST server 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 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"

# The database role to use when no client authentication is provided.
# Can (and probably 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 *èğŠèL'šçzçzšæęĆèłř*.

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	
server-host	String	*4	
server-port	Int	3000	
server-proxy-uri	String		
jwt-secret	String		
secret-is-base64	Bool	False	
max-rows	Int	∞	
pre-request	String		

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. 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". Also allows connections over Unix sockets for higher performance.

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

db-anon-role The database role to use when executing commands on behalf of unauthenticated clients.

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.

server-host Where to bind the PostgREST web server.

server-port The port to bind the web server.

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": "0.4.0.0",
    "title": "PostgREST API",
    "description": "This is a dynamic API generated by PostgREST"
  },
  "host": "postgrest.com:443",
  "basePath": "/",
  "schemes": [
```

```
"https"
]
}
```

jwt-secret The secret used to decode JWT tokens clients provide for authentication. 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.

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

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.

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.

Running 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 logfile or to the `syslog`:

```
ssh foo@example.com \
  'postgres foo.conf </dev/null >/var/log/postgres.log 2>&1 &'

# another option is to pipe the output into "logger -t postgres"
```

(Avoid `nohup postgres` because the HUP signal is used for manual *Schema refresh*.)

CHAPTER 16

çąńăŃŰ 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/;
        }
    }
}
```

```

    }
    ...
  }
}

```

Row-level security

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.

Count-Header DoS

For convenience to client-side pagination controls PostgREST supports counting and reporting total table size in its response. As described in [Limit aŠŇáŁÉéqŧ](#), 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:

```
Nginx stuff. Remove any prefer header which contains the word count
```

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

HTTPS

See the [SSL](#) section of the authentication guide.

ezReÅš

Nginx supports “leaky bucket” rate limiting (see [official docs](#)). Using standard Nginx 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 Nginx 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.

CHAPTER 17

èřČèřĚ

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. Note the `Server` response header as well which identifies the version of server. This is important when submitting bug reports.

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.

Schema Cache

Users are often confused by PostgreSQL'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 `EXPLAIN`. 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 PostgreSQL server, send the server process a SIGHUP signal:

```
killall -HUP postgres
```

In the future we're investigating ways to keep the cache updated without manual intervention.

CHAPTER 18

åd'ĜçŤÍ URL çžŞæđĎ

As discussed in *å■ŤæŤřæŤŤad'■æŤř*, there are no special URL forms for singular resources in PostgREST, only operators for filtering. Thus there are no URLs like `/people/1`. It would be specified instead as

```
GET /people?id=eq.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;

}
```

CHAPTER 19

èaí & èğEåŽ;

All views and tables in the active 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.

æşlèċ: 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 `èłDæžŘåłŇăěŮ`.

æřł'åžşèłGæzd' (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=lt.13 HTTP/1.1
```

Adding multiple parameters conjoins the conditions:

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

These operators are available:

abbreviation	meaning
eq	equals
gte	greater than or equal
gt	greater than
lte	less than or equal
lt	less than
neq	not equal
like	LIKE operator (use * in place of %)
ilike	ILIKE operator (use * in place of %)
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	checking for exact equality (null,true,false)
@@	full-text search using to_tsquery
@>	contains e.g. ?tags=@>.{example,new}
<@	contained in e.g. ?values=<@{1,2,3}
not	negates another operator, see below

To negate any operator, prefix it with `not` like `?a=not.eq.2`.

For more complicated filters (such as those involving disjunctions) 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
```

æſlëğç: We’re working to extend the PostgREST query grammar to allow more complicated boolean logic, while continuing to prevent performance problems from arbitrary client queries.

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=fname,age HTTP/1.1
```

The default is `*`, meaning all columns. This value will become more important below in [Vertical Filtering](#).

Filters

Filters may be applied to computed 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 $$
  SELECT $1.fname || ' ' || $1.lname;
$$ 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=@@.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
```

order

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


```
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
```

Content Negotiation

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.

Array of Objects

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 empty body and 404 status code rather than the usual empty array and 200 status.

⚠️: 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.

Raw Binary Data

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
```

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

CHAPTER 20

PostgreSQL

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:



As seen above in *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 Lumire 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(last_name) HTTP/1.1
```

Which would return

```
[
  { "title": "Workers Leaving The LumiÁre Factory In Lyon",
    "directors": {
      "last_name": "LumiÁre"
    }
  },
  { "title": "The Dickson Experimental Sound Film",
    "directors": {
      "last_name": "Dickson"
    }
  },
  { "title": "The Haunted Castle",
    "directors": {
      "last_name": "MÁl'liÁs"
    }
  }
]
```

æšlëč: As of PostgREST v4.1, parens () are used rather than brackets { } for the list of embedded columns. Brackets are still supported, but are deprecated and will be removed in v5.

PostgREST can also detect relations 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
```

æšlëč: Whenever foreign key relations change in the database schema you must refresh PostgREST's schema cache to allow resource embedding to work properly. See the section *Schema éšëjj*.

áŤNáĚëèŁGæzd'áŠNæÓŠážŘ

Embedded tables can be filtered and ordered similarly to their top-level counterparts. To do so, prefix the query parameters with the name of the embedded table. For instance, to order the actors in each film:

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

CHAPTER 21

èĜłáóŽázL'æšěèrc

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 and OR-conditions in the where clause
- More complicated joins than those provided by *èĬæžŘáŧŇáěŮ*
- Geospatial queries that require an argument, like “points near (lat,lon)”
- More sophisticated full-text search than a simple use of the @@ filter

CHAPTER 22

å■ÿåĆíè£ĞçÍÑ

Every stored procedure in the API-exposed database schema is accessible under the `/rpc` prefix. The API endpoint supports only POST which executes 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). However 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.

Procedures must be used with [named arguments](#). 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 $$
  SELECT $1 + $2;
$$ LANGUAGE SQL IMMUTABLE STRICT;
```

The client can call it by posting an object like

```
POST /rpc/add_them HTTP/1.1
{ "a": 1, "b": 2 }
```

The keys of the object match the parameter names. Note that PostgreSQL converts parameter names to lowercase unless you quote them like `CREATE FUNCTION foo("mixedCase"`

`text)` . . . 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.

æšlëġċ: We recommend using function arguments of type `json` to accept arrays from the client. To pass a PostgreSQL native array you'll need to quote it as a string:

```
POST /rpc/native_array_func HTTP/1.1
```

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

```
POST /rpc/json_array_func HTTP/1.1
```

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

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).

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.

æšlëġċ: 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.

We are considering allowing GET requests for functions that are marked non-volatile. Allowing GET is important for HTTP caching. However we still must decide how to pass function parameters since request bodies are not allowed. Also some query string arguments are already reserved for shaping/filtering the output.

èŎũåŔŮèŕũæšĆċŽĎ Headers/Cookies

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

XYZ. For example, to read the value of the Origin request header:

```
SELECT current_setting('request.header.origin', true);
```

22.2.2. Boolean Logic

For complex boolean logic you can use stored procedures, an example:

```
CREATE FUNCTION key_customers(country TEXT, company TEXT, salary_
→FLOAT) RETURNS SETOF customers AS $$
  SELECT * FROM customers WHERE (country = $1 AND company = $2) OR_
→salary = $3;
$$ LANGUAGE SQL;
```

Then you can query by doing:

```
POST /rpc/key_customers HTTP/1.1
{ "country": "Germany", "company": "Volkswagen", "salary": 120000.00 }
```

22.2.3. Error Handling

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

```
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"
}
```

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

CHAPTER 23

æRŠåĚě/ăŁóæŤž

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 *ăĎĈžŤ'èŁĜæzd' (Columns)* to these results.

æşlêċ: 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 [built for PostgREST](#).

To update a row or rows in a table, use the PATCH verb. Use [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=lt.13 HTTP/1.1
{ "category": "child" }
```

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

Warning: Beware of accidentally updating every row in a table. To learn to prevent that see [Prevent Accidental Updates](#).

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 backend 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 }  
]
```

CHAPTER 24

Řádění

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

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

Varování: Beware of accidentally deleting all rows in a table. To learn to prevent that see *Řádění*.

CHAPTER 25

OpenAPI æŦŦæŦŦ

Every API hosted by PostgREST automatically serves a full [OpenAPI](#) description on the root path. This provides a list of all endpoints, along with supported HTTP verbs and example payloads.

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, provides guidance with request headers and example request bodies.

æŦŦæŦŦ: 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 éŦŦèŦŦ](#).

CHAPTER 26

HTTP çŁúæĀçăĀ

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 prereq 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 authed 403, else 401	insufficient privileges
other	500	

CHAPTER 27

Authentication

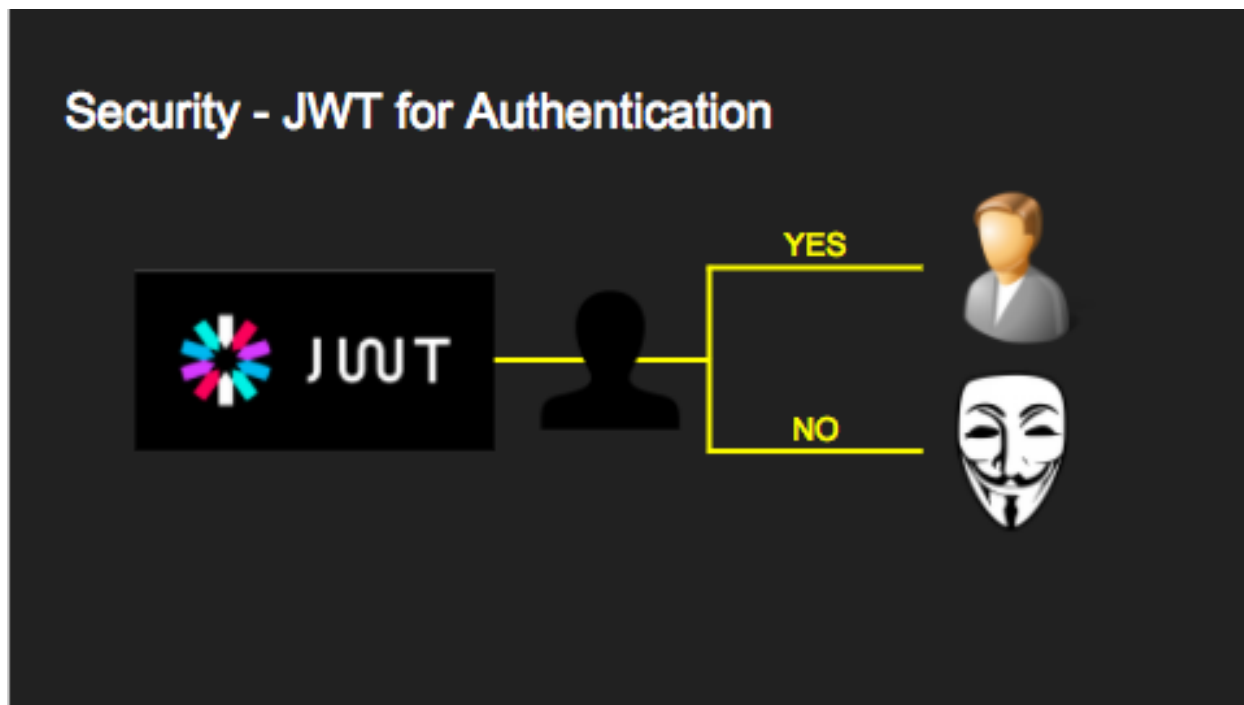
PostgREST is designed to keep the database at the center of API security. All authorization happens through database roles and permissions. It is PostgREST's job to **authenticate** requests – i.e. verify that a client is who they say they are – and then let the database **authorize** client actions.

Roles

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 PostgREST 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.

PostgreSQL

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.

Web Role Switching

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.

Web Authentication

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.

Database Roles

There is no performance penalty for having many database roles, although roles are namespaced per-cluster rather than per-database so may be prone to collision within the database. You

are free to assign a new role for every user in a web application if desired. 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)
```

èĜłáóŽázL'élÑérA

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
  LANGUAGE plpgsql
  AS $$
BEGIN
  IF current_user = 'evil_user' THEN
    RAISE EXCEPTION 'No, you are evil'
      USING HINT = 'Stop being so evil and maybe you can log in';
  END IF;
END
$$;
```


CHAPTER 28

JWT 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.
→GYDZV3yM0gqvuEtJmfpp1LBXSGYnke_Pvnl0tbKAjB4
```

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 passphrase – the signer and verifier share the secret. Hence any service that shares a passphrase with a PostgreSQL server can create valid JWT. (PostgreSQL currently supports only the HMAC-SHA256 signing algorithm.)

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` 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
  LANGUAGE sql
  AS $$
SELECT sign(
  row_to_json(r), 'mysecret'
) AS token
FROM (
  SELECT
    'my_role'::text as role,
    extract(epoch from now())::integer + 300 AS exp
  ) r;
$$;
```

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 '!!secret!!';

-- 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.

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'
  }
})
```

JWT aL'áĚÍ

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.

SSL

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 SSL. 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.

CHAPTER 29

æđúæđĐéŽŤçęž

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 çŤłæŁúçŤŇéíć* section below.

CHAPTER 30

SQL çŦíæŁũçóaçŦĒ

å■ŸåĆíčŦíæŁũåŠŦårĒçăA

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
  language plpgsql
  as $$
begin
  if not exists (select 1 from pg_roles as r where r.rolname = new.
→role) then
    raise foreign_key_violation using message =
      'unknown database role: ' || new.role;
    return null;
  end if;
  return new;
end
$$;

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
  language plpgsql
  as $$
begin
  if tg_op = 'INSERT' or new.pass <> old.pass then
    new.pass = crypt(new.pass, gen_salt('bf'));
  end if;
  return new;
end
$$;

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 $$
begin
  return (
    select role from basic_auth.users
    where users.email = user_role.email
      and users.pass = crypt(user_role.pass, users.pass)
  );
end;
$$;

```

Public autentizace

In the previous section we created an internal table to store user information. Here we create a login function which takes an email address and password and returns JWT if the credentials match a user in the internal table.

Uvod

As described in *JWT from SQL*, we'll create a JWT inside our login function. Note that you'll need to adjust the secret key which is hard-coded in this example to a secure secret of your choosing.

```

create or replace function
login(email text, pass text) returns basic_auth.jwt_token
  language plpgsql
  as $$
declare
  _role name;
  result basic_auth.jwt_token;
begin
  -- check email and password
  select basic_auth.user_role(email, pass) into _role;
  if _role is null then
    raise invalid_password using message = 'invalid user or password';
  end if;

  select sign(
    row_to_json(r), 'mysecret'
  ) as token
  from (

```



```

    select _role as role, login.email as email,
           extract(epoch from now())::integer + 60*60 as exp
    ) r
    into result;
    return result;
end;
$$;

```

An API request to call this function would look like:

```

POST /rpc/login HTTP/1.1

{ "email": "foo@bar.com", "pass": "foobar" }

```

The response would look like the snippet below. Try decoding the token at jwt.io. (It was encoded with a secret of `mysecret` as specified in the SQL code above. You'll want to change this secret in your app!)

```

{
  "token": "eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.
  →eyJlbWFpbCI6ImZvb0BiYXNpdG9sZSI6ImF1dG8uZm93IiwiaWF0Ijoi
  →ERi5qbW9E5NPzvaUjgvdmlm0zkIG9xSm2w5zmdw"
}

```

Permissions

Your database roles need access to the schema, tables, views and functions in order to service HTTP requests. Recall from the **'Overview of Role System'** that PostgREST uses special roles to process requests, namely the authenticator and 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;
create role authenticator noinherit;
grant anon to authenticator;

grant usage on schema public, basic_auth to anon;
grant select on table pg_authid, basic_auth.users to anon;
grant execute on function login(text,text) to anon;

```

You may be worried from the above that anonymous users can read everything from the `basic_auth.users` table. However this table is not available for direct queries because it lives in a separate schema. The anonymous role needs access because the public `users` view

reads the underlying table with the permissions of the calling user. But we have made sure the view properly restricts access to sensitive information.