

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

PostgREST Documentation

ãĖŠãĚĈ 4.1.0

Joe Nelson

2020 åžť 03 ælJĹ 11 æŮě

Contents

1	áčřæŸŌæĀgçijŮčlŇ	2
2	Leak-proof æL;èsa	3
3	æŇěæLsâĚšçszælađŇ	4
4	äÿÄäÿléĜmçĆz	5
5	æŤzè£ZâĚsāzn	6
6	çŤšæĀçszçzš	7
6.1	áčæLüçnrážš	7
6.2	éçlād' ŮéĀžçšě	8
6.3	çd'žäçNāzŤçŤl	8
6.4	Production	9
6.5	æŇšāsŤ	9
6.6	áz£āšŁ	9
7	èťđèl'	10
8	èŮuâçŮæŤræŇA	11
9	Tutorial 0 - èŋl'áočĚuŤæiě	12
9.1	Step 1. æŤçè;çæĭçèĀæéšA, æLšāznāijZāyŋäçčžĐ	12
9.2	Step 2. áŋl'èçĚ PostgreSQL	13
9.3	Step 3. áŋl'èçĚ PostgREST	13
9.4	Step 4. äÿž API áLžâzæŤræmŋōāžš	14
9.5	Step 5. è£ŘèqŇ PostgREST	15
10	Tutorial 1 - éĜšéšĚāŇž	17

10.1	Step 1. æúzáŁääyÄäyIáRŮäŁçŽDčŤlæLú	17
10.2	Step 2. çŤšæLŘäyÄäyIáRŮäŁçŽDčŤlæLú	18
10.3	Step 3. çŤšæLŘ token	18
10.4	Step 4. èŁŽèaŇèrŮæśĆ	18
10.5	Step 4. æúzáŁäèŁĜæIJšæŮúéŮt'	20
10.6	éŽDáŁæćŸ: çŇŇá■šæŠd' éŤĀ	21
11	âŔŕæL'gèaŇæŮĜäzú	23
12	Homebrew	24
13	PostgreSQL ä;IèŤŮ	25
14	æžŘäzčçäAçijŮèŕŚ	26
14.1	PostgREST æŤŇèŕŤæŮäzú	27
15	éĚ■ç;õ	30
15.1	âŔŕáŁÍ Server	32
16	çañáŇŮ PostgREST	33
16.1	éŸzæ■çáĚlèaŁæš■ä;IJ	34
16.2	Count-Header DoS	34
16.3	HTTPS	35
16.4	éŽŔéĀš	35
17	èŕČèŕŤ	37
17.1	Schema éĜ■è;ij	38
18	âd'ĜçŤÍ URL çzšædĎ	39
19	èaÍ & èĝEaŽ;	40
19.1	æŕt' ázšèŁĜæzd' (Rows)	40
19.2	âd'ČçŽt' èŁĜæzd' (Columns)	42
19.3	æŎŠázŔ	42
19.4	Limit âŠŇáĚÉaŤ	43
19.5	çŽyázŤæäijâijŔ	44
19.6	â■ŤæŤŕæLŮâd'■æŤŕ	45
19.7	äžŇèŁŽáLŮè;šâĜž	45
20	èŤDæžŔaŤŇæèŮ	47
20.1	âŤŇæĚèŁĜæzd' âŠŇæŎŠázŔ	49
21	èĜláoŽázL'æššèèŕć	51
22	â■ŸâĆlèŁĜçlŇ	52
22.1	èŎúâŔŮèŕŮæśĆçŽĎ Headers/Cookies	53
22.2	âd'■æiĆâyČârŤéĀzè;Ś	54

22.3	æŁZãĜžéŤZèřř	54
23	æRŠãĚě/ãĚóæŤz	56
23.1	æL'zéĜŘæRŠãĚě	57
24	ãĹãéZd'	59
25	OpenAPI æŤřæŇA	60
26	HTTP çĹúæĀçãA	61
27	èĝŠèL'sçşzçzšæçCèřř	63
27.1	éĪŇérAžŘãĹŮ	63
27.2	çŤĹæĹúãŠŇçzĎ	65
27.3	èĜĹãóŽãzĹ'éĪŇérA	67
28	ãóçæĹúçnr Auth	68
28.1	JWT Generation	68
28.2	SSL	71
29	æđúæđDěŽŤçez	72
30	SQL çŤĹæĹúçõçŘE	73
30.1	ãŸãĆĹçŤĹæĹúãŠŇãřEçãA	73
30.2	Public çŤĹæĹúçŤŇéĹc	75



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

PostgREST umožňuje provádět CRUD operace na datích v PostgreSQL pomocí REST API. Je to jednoduchý a rychlý způsob, jak získat REST API z PostgreSQL. PostgREST je webový nástroj, který umožňuje poskytnout REST API z jakýchkoli PostgreSQL databází. Je to jednoduchý a rychlý způsob, jak získat REST API z PostgreSQL.

CHAPTER 1

ĀčřæŸŌæĂğcijŪčlŃ

æřTèŧuárzçlĀæšèèrcçzšæđlJéĜ■ād'■āŁşāLlĳjNèđl æšèèrcèđāāLšçRĒæyĔçzĒèŁĆä;£ Post-
greSQL äyžæĆlæũzāLăæŦřæ■ōæŸřäyĂäzŭæZt' āóžæŸšçŽDăžŃăĂĆäyžæŦřæ■ōăžšărzèsāāLĒéĔ■ælČéZŘæřŦă

CHAPTER 2

Leak-proof æŁjèśą

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

CHAPTER 3

æNěæŁśăĔşçşzæłąđŃ

1970ázt'rijŃE. F. Codd áIĴázŮçŽDæŮĜçná âAIĴád' ġăđŃăĔśăznæŤřæ■óăžŞæŤřæ■óăĔşçşzæłąđŃăĀI
äy■æL' zèrĎăžEă;ŞæŮüäyzăriçŽDæŤřæ■óăžŞăĴEăśCăłąđŃăĀCăžŌéC;æŮĜçnáaijŽăRŚçŌřăśCăňăæŤřæ■óăž
http èurçŤsăzŃéŮt' â■ŸăIĴăĀŁăžçZyăijjæĂġăĀCèĂŃăĴĴ PostgREST
äy■iiĴŃăĴSăzňăriĕrŤă;ĕçŤĴçAĵæt' zçŽĐèĴĜæzd' âŞŃăĵŃăĔĕriĴŃèĂŃăy■æŸřăĵŃăeŮèurçŤsăĀC

CHAPTER 4

ÄyÄytleG■çZ

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

CHAPTER 5

æŤzèŁZăĚśăžń

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

CHAPTER 6

PostgreSQL

PostgreSQL is a powerful, open source object-relational database system with over two decades of active development. Its major design goals are simplicity, reliability, and performance.

6.1 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

6.2 Listening to PostgreSQL

Listening to PostgreSQL

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

6.3 Examples

- [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
- [begriiffs/postgrest-example](#) - sqitch versioning for API
- [SMRxT/postgrest-demo](#) - multi-tenant logging system
- [PierreRochard/postgrest-boilerplate](#) - example auth backend

6.4 Production

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

6.5 æŃŖåŖ

- [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

6.6 åŖåŖ

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

CHAPTER 7

ètdèl'

"ájĀāRSĕtūæIēād' ĩāĕńăžĒ, æDšĕğL' ārsāĈRāIJlä;IJāijL!"

—François-G. Ribreau

"æLSäy■ā; Ūäy■èrt', äyŎ Node.js/Waterline ORM ædDāzžçŽĎ API ārzæfT

CPU/Memory usage çōĀçŽt' æYréZ; äzĕç;ōäĕā. ā;šæLSäzñāIJĪ 6 äyĭçd' žä;N ĩijLdynosiiijL' æŃAçz■æšĈæĈĒ 1GB æTřæ■ōæYřāōĈçTŽĕGšāRĭæIJL' 60/70 MB ād' ġārĔ."

—Louis Brauer

"æLSéIdāyyāŪIJæñcèĒZæäüäyÄäyĭāzNāōđiijŃāAūçDūä;ĕçTĪ SQL
DDLĭijLāŠŃ V8 javascriptiijL' āijĀāRSā;ōæIJ■āŁāāĀĈ æLSäzñāIJĪ 6
äyĭæIJLāEĒāōŃāĒĪéG■āEŽāzĒäyÄäyĭ Spring + MySQL éAŪçTŽāžTçTĪçĪŃāžRāĀĈ
éĀšāžĕāĕń 10 āĀ■iijŃāzççāAā;ŁçōĀæf' AāĀĈĕĀŃāzNāL'■çŽDāzžçTĪläžĒ 4
äyĭāžžĕĒŁsäžĒ 3 āžt' æŪúéŪt' āĀĈ"

—Simone Scarduzio

CHAPTER 8


èÕùå;ÚæŤræÑÀ

èfèéazçZóæIJL'äyÄäyłáRNâe;äyŤäy■æÚ■æLRéŤçZDçd' ;ãÑzãĂĆâŁääĚæŁSázñçZD
èAŁád'Ťáóđ' æİèèóÍèõžãŠÑæśĆâŁŤãĂĆâRÑæÚüä;ääzšâRřäzěâIJÍ Github çZD issues
äyŁæŘIJçť'ć bugs/featuresãĂĆ

Tutorial 0 - èó'áóÇèùSètùæläě

æñcéfŌä;£çTÍ PostgRESTijAèrēāL'■ēlĀæYřäyÄäy† Quick
startiijNāyóāL'ä;ää£néĀšāLZāzzçñnäyÄäy†çóĀ■TçŽĐ APIāĀĆ

PostgREST æYřäyÄäy†çNñçñNçŽĐ Web æIJ■āLāāZlíijNäyž PostgreSQL
æTřæ■óāžšçTšæLŘ RESTful APIāĀĆ áóÇæRŘä;ZāšžžžŌāžTásÇæTřæ■óāžšçzšæđĐáóZāLúçŽĐ
APIāĀĆ



æČšèeAçTšæLŘ APIiijNæLŠäzñāRlēIJĀèeAāLZāžžyÄäy†æTřæ■óāžšāĀĆæL'ĀæIJL'çnrçCzāŠNæĪČéZŘē
āIJĀæIJñæTžçlNçzšæIšçŽĐæUúāZíijNæCíārEæNēæIJL'äyÄäy†èÇ;çTÍçŽĐæTřæ■óāžšíijNPostgREST
æIJ■āLāāZlíāŠNäyÄäy†çóĀ■TçŽĐā■TçTíæLú todo list APIāĀĆ

9.1 Step 1. æTçè;žæiçèĀĀéŠA, æLŠäzñāijŽāyóä;áčŽĐ

āIJĀ;āāijĀāgNē£Zāy†æTžçlNæUú, Ctrl+T äyÄäyNāIJĀæŪřæāĠç;äy■æL'ŠāijĀéazçŽó
èAŁād'P'áód' . æIJL'äyĀç;đ'ā;L nice çŽĐāžžāIJlèAŁād'P'áód' äy■æT'žèùČíijNāçCæđIJä;ää■āā;RāžEæLŠäzñāijŽā

9.2 Step 2. Installing PostgreSQL

For installing PostgreSQL on Docker, you can use the following command. This command will create a PostgreSQL container named 'tutorial' and start it. The container will have a PostgreSQL server running on port 5432. The password for the 'postgres' user is 'mysecretpassword'.

The following command will create a PostgreSQL container named 'tutorial' and start it. The container will have a PostgreSQL server running on port 5432. The password for the 'postgres' user is 'mysecretpassword'.

The following command will create a PostgreSQL container named 'tutorial' and start it. The container will have a PostgreSQL server running on port 5432. The password for the 'postgres' user is 'mysecretpassword'.

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

The following command will create a PostgreSQL container named 'tutorial' and start it. The container will have a PostgreSQL server running on port 5432. The password for the 'postgres' user is 'mysecretpassword'.

9.3 Step 3. Installing PostgREST

PostgREST is a REST API for PostgreSQL. It can be installed on Linux/BSD/Windows. The following command will create a PostgREST container named 'tutorial' and start it. The container will have a PostgREST server running on port 8080. The password for the 'postgres' user is 'mysecretpassword'.

The following command will create a PostgREST container named 'tutorial' and start it. The container will have a PostgREST server running on port 8080. The password for the 'postgres' user is 'mysecretpassword'.

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

The following command will create a PostgREST container named 'tutorial' and start it. The container will have a PostgREST server running on port 8080. The password for the 'postgres' user is 'mysecretpassword'.

```
./postgrest
```

The following command will create a PostgREST container named 'tutorial' and start it. The container will have a PostgREST server running on port 8080. The password for the 'postgres' user is 'mysecretpassword'.

Notes: PostgREST requires the 'libpq' library to be installed. The following command will create a PostgREST container named 'tutorial' and start it. The container will have a PostgREST server running on port 8080. The password for the 'postgres' user is 'mysecretpassword'.

while loading shared libraries: libpq.so.5: cannot open shared object file: No such file or directory

9.4 Step 4. Creating API Schema

Running the following SQL in psql (psql) will create the schema:

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

Running the following SQL in psql will create the schema:

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

postgres=#
```

The following SQL will create the API schema in the database:

```
create schema api;
```

The following SQL will create the API table in the database:

```
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');
```

The following SQL will create the API role in the database:

```
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;
```

The following SQL will create the API role in the database:

Chapter 9: PostgreSQL API

\q

9.5 Step 5. Running PostgreSQL

PostgreSQL API tutorial.conf:


```
tutorial.conf
```

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

Options for PostgreSQL API:

```
./postgres tutorial.conf
```

Output:

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

API curl todos:


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

API response:

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

(aynéatczgcz)

(çzäyŁéą)

```
}
]
```

éÅžèŁĜă;ŞăL'çŽĎèĝŞèL'sæiČéŽŘřijŇăŇŁăŘ■èrúæśCæIJL' todos
 èajčŽĎăRlèrzæiČéŽŘăĂCăęCăđIJæŁSăznèrŤăŽ;æûzâLăäyĂäyŁæŮřčŽĎ todo
 äijŽècŇæŇŞçziăĂĆ

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

âŞ■ăžŤæŸř 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!
 âIJlâyŇăyĂçřĜæŤžçlŇăy■rijŇæŁSăznârĚajjŽçIJŇăLŕăęCă;ŤæŇŞăśŤèŁŽăyŁă;Ňă■ŘřijŇă;ŁçŤlæŽt'ăđ'■æiČçŽĎ

CHAPTER 10

Tutorial 1 - éĜŚéŠěãŇŽ

âĪĪ *Tutorial 0 - èól'áoĈèùSèŭæĬě* äy■æĹSázňáĹZázžazĚäyÄäyĚèŌúâRŪ
todos æŦræ■óâRĬérzçŽĎ APIãĀĈä;£çŦĪã■ŦäyĬçñrçĈzãĹŪãĜž todosãĀĈ
æĹSázňæĪĴĹâ;Ĺâd'ŽãĹđæşŦâRrãžěä;£è£ŽäyĬ API æŽr' æĪĴĹèúċĪĪjNã;ĚäyÄäyĬæč;çŽĎâĪjĀğNæŸrãĒĀèőyäyĀã

10.1 Step 1. æúžãĹääyÄäyĬãRŪä£açŽĎçŦĪæĹú

äyĹäyÄèĹĈäy■ĪĪjNè£ŽèãNãŦãR■ Web èrúæśĈçŽĎäžNãRŌâĪĪæŦræ■óâžŞäy■ãĹZázžazĚäyÄäyĬ
web_anon èĝSèĹ'sãĀĈèól' æĹSázňâĪĴĹãĹZázžäyÄäyĬèĝSèĹ'sãRnáAŽ todo_user
çŦĪãžŌã;£çŦĪ API è£ŽèãNèžňãz;éĬNèrAçŽĎçŦĪæĹúĪĪjNè£ŽäyĬèĝSèĹ'sãrĚæĪĴĹæĪĈãrž todo
list âAŽãžžã;ŦãžNæĈĒãĀĈ

```
-- 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;
```

10.2 Step 2. çŤšæĹRäyÄäyĹárEçăA

áóçæĹúçnréÄŽèĜ API ä;çŤĪ JSON Web Token èĤZèaÑèžnáz;éĪÑèrAāĀĈJTW æŸřä;çŤĪläzĒæĪĴĹæĹSäzñāŠÑæĪĪāĹāZĪçšééAšçŽĎárEçăAæĤZèaÑāĹāārEç;ĹāR■çŽĎ JSON áržèšāāĀĈ çŤšäzŌáóçæĹúçnräy■çšééAšçĎárEçăAĪĪjÑæĹĀžžèäy■èĈ;çřæŤž token çŽĎāEĒāóžāĀĈ PostgREST äĪjŽæçĀæĪNāĪjĹéĀāçŽĎ token ážúæNšçžĪāóĈäžñāĀĈ

æĹSäzñæĹēāĹZāžžäyÄäyĹárEçăAāžúæRŘä;ŽçžZ PostgRESTāĀĈæĪĀāē;æĈšæĈšäyÄäyĹād' ■æĪĈçŽĎéŤĸæĪ

æšĹèĝč: OpenSSL toolkit æRŘäyĹäyÄäyĹçóĀā■ŤçŽĎæŪžāĪjRæĹèçŤšæĹRāóĹāĒĪçŽĎárEçăAāĀĈæĈæĎĪĶ;āæĪ

```
openssl rand -base64 32
```

æĹŠāĪjÄ tutorial.conf (āĪĪäyĹäyÄèĹĈäy■āĹZāžžçŽĎ) ážúārEārEçăAæūzāĹāĪĪæŪřçŽĎäyÄèaÑ:

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

æĈæĎĪ PostgREST server äž■æŪĝāĪĪèĤRèaÑäy■ĪĪÑéĈçäžĹéĪĪÄèēAéĜ■āRřāóĈäžžèä;ĸāĹæē;ĶæĪĪæŪřçŽĎ

10.3 Step 3. çŤšæĹR token

éŽžäyŸä;æĜĹāúšçŽĎäžççăAāĪĪæŤřæ■óāžšæĹŪāĒūāžŪæĪ■āĹāZĪäy■ārEāĹZāžžāžúç;ç;šèžnáz;éĪÑèrA tokenĪĪjÑä;EæŸřāĪĪæĪĪæŤžĈĪNäy■ĪĪjÑæĹSäzñārEāĪĪèĜĹāúšāĹĪæĹNāĀĪāĀĈèúšè;ñāĹř jwt.ioĪĪjÑāžúāāñāEŽæĈäyNā■ŪæóřĪjž

èřüèõřā;RæĈĪāñāEŽçŽĎárEçăAĪĪjÑèĀNäy■æŸřāž;çĹĜéĜNçŽĎ secretāĀĈāāñāEŽārEçăAāŠÑ payload äžNārŌĪĪjNāūēä;ĝçŽĎçĪjŪçăAæŤřæ■óāĪjžĀĹūæŪřĪjÑèrèæŤřæ■óāš token āđ'■āĹūāóĈāĀĈ

æšĹèĝč: èž;çĎŪāžđ'çĹNārřèĈ;çĪĪNèřüæĹēä;ĹæĹaçšĹĪĪjNä;Eä;ĹāóžæŸšéĀEārŠāĜžçŽĎ payloadāĀĈtoken äžĒäžĒæŸřèçñç;ĹāR■ĪĪjÑæšāæĪĴĹāĹāārEĪĪjÑæĹĀžžèæĈæĎĪĶ;āæĪĴĹäy■æĈšèóĹ' áóçæĹúçnrçĪĴ

10.4 Step 4. èĤZèaÑèrúæšĈ

āžđāĹř terminalĪĪjÑæĹSäzñæĹèçŤĪ curl æūzāĹäyŸäyĹ todoāĀĈèrèèrúæšĈārEāNĒæNñäyÄäyĹāNĒāRñèžnáz;éĪÑèrA token çŽĎ HTTP āđ't'āĀĈ

Encoded PASTE A TOKEN HERE

Decoded EDIT THE PAYLOAD AND SECRET (ONLY HS256 SUPPORTED)

1. Enter password

2. Enter this JSON

3. Copy the resulting token

obrázok 1: Jak vytvořit token pomocí <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"}'
```

POST /todos PATCH /todos

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

GET /todos

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

```
[
  {
    "id": 1,
    "done": true,
    "task": "finish tutorial 0",
```

(obrázky jsou zobrazeny v samostatném okně)

(çz■äÿŁéą)

```

    "due": null
  },
  {
    "id": 2,
    "done": true,
    "task": "pat self on back",
    "due": null
  },
  {
    "id": 3,
    "done": true,
    "task": "learn how to auth",
    "due": null
  }
]

```

10.5 Step 4. æúzáŁæèŁGæIJšæÚúéÚť

çZóáL■iijNæLSázñçZDèòd'èrA token árzážŎæL'ĂæIJL'èrúæśCéČ;æYřäÿĂèGť æIJL'æTŁçZDăĂCæIJ■áŁa JWT árEçăAiiijNársaijZéĂZèŁGéłNèřAãĂC

æZť äè;çZDç■ÚçTşæYřèòł' token ä;ŁçTł exp äçřæYŎäÿĂäÿłèŁGæIJšæÚúéÚť æŁšăĂCèŁZæYř PostgREST çL'záLńárzáŁEçZDäÿd'äÿł JWT äçřæYŎázNäÿĂãĂC

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

æşłèçč: Unix æÚúéÚťæŁş (Unix epoch time) èćnáòŽázL'äÿžèGł 1970 ázt' 1 æIJŁ 1 æÚě 00:00:00 á■RèřCäÿÚçTŇæÚúéÚťLUTCiiijL'ázèæİèáŁřçŎřáIJčZDăĂZçğšæTřiiijNäÿ■èĂCèZŚéUřçğšăĂC

äÿžăZÈáIJłèąNáLíäÿ■èğCárşèŁGæIJšiiijNæLSázñârEæúzáŁăäÿĂäÿłáIJł 5min äZNáRŎèŁGæIJšçZD exp äçřæYŎăĂCéçÚáĚŁæL'ŁáŁřăžŎá;şáL'■æÚúéÚť çóUèłúáŁř 5min äZNáRŎçZDæÚúéÚťæŁşăĂC áIJł psql äÿ■èŁRèąNiiijZ

```

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

```

áZďáŁř jwt.io áZúüæŁóæTž payload

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

Uložte si novou hodnotu tokenu do proměnné `NEW_TOKEN`:

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

Prozkoumejte, jak se chová API, když použijete expirující token:

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

API vrátí chybu `HTTP 401 Unauthorized`:

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

10.6 Revokace tokenů: `revoke_token`

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;
```

([vídeopříklad](#))

(çzäyŁeą)

```

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

[*ayNè;jeat*]

*ayNè;jeat*élcáEüæIJL Mac OS X Windows *áŠNáGäyI* Linux
áRŠèaÑçL'ŁÇZĐéçDçijUèrSæÚGäzúãÁCèğçáÓNázNáRÓáRfäzèèfRèaÑáRfæL'gèaÑæÚGäzúãLä
--help *æäGåfÚæIèæšèçIJNä;ŁçTlèrt'æYÓ:*

```
# èğçáÓN tar áNĚ (available at https://github.com/begriffs/postgrest/  
→releases/latest)  
  
$ tar Jxf postgrest-[version]-[platform].tar.xz  
  
# ářIèrTèfRèaÑ  
$ ./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'æšæéŮóécŸ  
brew doctor  
  
# âóL'èčĚ postgres  
brew install postgres
```

ěřěãŠ;äzd'äijŽeĜlãLlãřE PostgreSQL ä;ŠãAŽä;IetŮãóL'èčĚ.
ěřěè£ĜćlNã;Ăã;ĂéIJĂèçAéTfè; 15ãLEéŠšæL■èČ;ãóL'èčĚ;řázúãNěãRĹãĚüä;IetŮãĀĆ
ãóL'èčĚãóNãĹRãRŌijNěřěãüèãĚüäijŽècnæüzãĹããĹř \$PATH
äÿ■ijNã;ããRřázěãĪĪlãzzæĎRã;■ç;óä;£çTĪijŽ

```
postgres --help
```

CHAPTER 13

PostgreSQL a; lèù

9.3 [PostgREST](#) [æCíárÉéIJÀèçAáòL'èçĚæTřæ■óžšijLPostgreSQL](#)
[æLÚæZt'énYçL'áIJñijL'ãĀĆ](#) [æCíáRřázěä;çTíáČR](#) [Amazon](#) [RDS](#)
[èfZæüçŽDäyIèçfijNä;EæYřáIJæIJñáIJřáòL'èçĚæIJñèžnærTè;Čä;£áòIJñijNæZt'ä;£äžÓaijĀáRŠãĀĆ](#)

- [OS X èrt'æYŎ](#)
- [Ubuntu 14.04 èrt'æYŎ](#)
- [Windows áòL'èçĚãĚ](#)

CHAPTER 14

æžŘäzččäAçijÛèrŚ

æšlègč: æLŠäznäy■ázžèóóáIJÍ **Alpine Linux** äyŁæđDázžāŠNä;ŁçTÍ Post-gRESTiijNāZāäyžāIJlèrēāzšāRræIJL'èŁĠ GHC āĒĒā■ŸæšDæijRçŽDæŁēāSŁāĀĆ

ā;ŠæĆlçŽDçšzçzšæšææIJL'écDæđDázžçŽDāRræL'gèāNæŪĠäzūæŪiijNāRrāzēāzŌæžŘäzččäAæđDázžéāz Stack āĀĆāóĀrĒāIJlæĆlçŽDçšzçzšäyŁāóL'èĀĒzā;ŦāĒĒèĒAçŽD Haskell ä;ĪèŦŪāĀĆ

- āóL'èĀĒ Stack
- āóL'èĀĒä;ĪèŦŪāzŠ

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

- æđDázžāzūāóL'èĀĒ

```
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
```

æšlëğç: æĈæđIJä; äæđDäzžäd' sèt' ëiijNëÄNäyTä; äçZDçşçzçşâRlæIJL' äy■älĭ 1GB
 äĔËä■YriijNârIërTæûzâLääyÄäyĭ swap æŪGäzŭäĀĈ

- æĈÄæšëäöL' èĈËæYrâræLĖLâLš: `postgres --help`.

14.1 PostgreSQL æŧNërŧæÜäzŭ

14.1.1 äLZäzžæŧNërŧäzš

äyžäzEæ■ççäðèĔRëaŊpostgrestèĔZëaŊæŧNërŧriijNëëŪäĔLÉIJÄèèAäLZäzžäyÄäyĭæŧŧæ■öäzšâĀĈäyžæ■d' r
 èDŽæIJñéIJÄèèAäzëäyNârĈæŧriijŽ

```
test/create_test_db connection_uri database_name [test_db_user] [test_
  →db_user_password]
```

ä;Ĕçŧĭ'connection URI <<https://www.postgresql.org/docs/current/static/libpq-connect.html#AEN45347>> ' _ äŌzæŊGäöZæŧŧæ■öäzšçŧĭæLüäÄÄârEçäAäÄÄäyžæIJžäzžäRĔçnrârĈäĀĈäy■èèAäIJæŧŧæ■öä
 èDŽæIJñäy■çZD:code: 'database_name'ârĈæŧriijNæYrârEèèAèĔäèĔäLĔçZDæŧŧæ■öäzšâR■çğŕäĀĈæĈ
 äĈæđIJä;ĔçŧĭæŊGäöZçZDæŧŧæ■öäzšçŧĭæLüèĔZëaŊäæEæälæŧNërŧäĀĈæŕRæñæŧNërŧèĔRëaŊârŌiijŊ
 äĈæđIJæIJæŊGäöZçŧĭæLüriijNèDŽæIJñârEäijZçŧŧæLĔRèğšèLšâR■:code:postgres_test_riijNäzŭäzèæLÄ
 äĈæđIJä;ĔçŧĭäyÄäyĭäüšçzRâ■YäIJĔZDçŧĭæLüäIèèĔZëaŊæŧNërŧèĔäèŌëriijNéĈçäzLèĔYèIJÄèèAæŊGä
 èrèèDŽæIJñârEèĔTäZæŧNërŧèĔĔçĭNäy■ä;ĔçŧĭçZDæŧŧæ■öuri
 èrèuriäyŌârEäIJĔŧšäzğäy■ä;ĔçŧĭçZDèĔ■ç;öæŪGäzŭäRĈæŧŧ:code: 'db-uri' çZyâržäzŧäĀĈ
 çŧŧæLĔçŧĭæLüäšNârEçäAäĔÄèöyälZäzžæŧŧæ■öäzšäzŭâržäzžä;ŧpostgresæIJ■äLäZlèĔRëaŊæŧNërŧriijŊ

14.1.2 èĔRëaŊæŧNërŧ

äyžäzEèèRëaŊæŧNërŧriijNäĔĔëäzâIJĔŌŕäĈĈârYéGRäy■æRĖä;Zæŧŧæ■öäzšçZDuriäĔæAŕriijNâržäzŧçZD
 éÄZäyÿæĈĔäĔäyNriijNäLZäzžæŧŧæ■öäzšäyŌèĔRëaŊæŧNërŧäS;äzd' äijZâIJläRñäyÄäS;äzd' èaŊäy■æLğè

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

äIJläRñäyÄæŧŧæ■öäzšäy■éĔ■äd' ■èĔRëaŊæŪriijNäzŧèrèäriijäGzæŧŧæ■öäzšèĔäèŌäRŶéGRäĔæAŕriijZ


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

áĈæđIĈŎřáĈĈáRŸéĠRäyžçl' zæLŪæIJĥæŃĠáŏŽiijŃéĈčázLæŧNèřŧçŽDèĚŘèaŃçlŃážRârĚäijŽézŸèóđ'èĚ

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

äyŁèĚřèĚđæŎéåAĠáŏŽæŧNèřŧçŽDæIJ■åŁaáŽlâIJlæIJñâIJř:code:localhost:code:iijŃážüäyŧæŧřæ■ŏázŞçŧl

14.1.3 éŧĂæíAæŧřæ■ŏázŞ

æŧNèřŧáŏŃæLŘázŃâŔŎiijŃæŧNèřŧæŧřæ■ŏázŞârĚäijŽèĈnáĚİçŧŽiijŃâŔŃæŮüèĚŸäijŽâIJlpostgresæIJ■åŁ

```
test/destroy_test_db connection_uri database_name
```

14.1.4 äĬĚçŧÍ Docker æŧNèřŧ

äyžázĚçŏĂâŃŮèĚđæŎééİđæIJñâIJřçŎřáĈĈPostgreSQLçŽDæŧNèřŧçŎřáĈĈèŏç;ŏiijŃâŔřázèä;ĚçŧlâyĂçġ■

äĬŃæĈiijŃæĈæđIæŸřâIJlmacäyŁâAŽæIJñâIJřâijĂârŠiijLâyŧâũşçzŘâŏL'èĈĚäžĚDockeræIJ■åŁaŧiijL'iijŃ

```
$ 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
```

æ■d'äd'ŮiijŃæĈæđIĚĂžèĚĠáLŽázždockeráŏžázlĚĚŘèaŃæĬèĚŘèaŃâĚæăLæŧNèřŧiijLâržázŎGHĈä;Ŏâ
SierraæŸřâĚĚèĚçŽDiiijŃâIJl:code:'stack test'äijzæŘŔçđ'žâijĈâyŷiijL'iijŃä;âârřázèâIJlâ■ŧçŃñçŽDáŏžázlây■

äĬĚçŧlâyžèäyŃèĎŽæIJñæđDázžæŧNèřŧä;ĚçŧlĚŽDáŏžázl:code:test/Dockerfile.testiijŽ

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

âIJlæŧNèřŧáŏžázlĚéŮæñæĚŘèaŃæŮüiijŃârĚäijŽèĚĚset'zèçĈĚŧçŽDæŮüéŮt'iijŃâŔžäyžéIJĚèĚAçijŞâ■Ÿçž
linux'æŮĠäžüäd'žä;IJäyžáŏžázlĚçŽDæŃĈè;Ĭâ■üiijŃäžèçqăŏĚĬæĚSâžñâIJlâyĂæñqæĂġæĬâijRâyŃèĚŘèaŃâŏžáz
work-docker'ârŃæăüéIJĚèĚAæŸârĎèĠşáŏžázlây■iijŃâIJlâ;ĚçŧlLinuxäy■çŽDstackæŮüâĚĚæžæŃĠáŏžázlây■
workăĈiijLâIJlSierraäyŁ:code:'stack build'ârřázèæ■čâyŷä;ĚçŧliijŃèĎŃ:code:'stack
test'âIJlĠĠĠ 8.0.1äy■äy■äijŽèĚüä;IJçŧl'iijL

æŮĠäžüäd'zæŸârĎèĠşdockeráŏžázlây■iijŽ

```
$ 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"
```

14.1.1. PostgREST 4.1.0

```
$ 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'šçzçzšæçĈèłř*.

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"
  },
}
```

(äyNéatçzğçz■)

(çz■äyŁéat)

```

"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.

15.1 ăŘřăĹÍ 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 \
  'postgrest foo.conf </dev/null >/var/log/postgrest.log 2>&1 &'

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

```

(Avoid `nohup postgrest` because the HUP signal is used for manual *Schema éĜ■èjj*.)

CHAPTER 16

çãñăŃŮ 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;
        }
    }
}
```

(äÿNéatçzğçz■)

(çzäyŁéat)

```

    proxy_pass http://postgrest/;
  }
  ...
}
}

```

16.1 éYzæ■cáĚléaÍæŞ■ä;IJ

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.

16.2 Count-Header DoS

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

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

16.3 HTTPS

See the `ssl` section of the authentication guide.

16.4 rate limiting

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
```

([more info](#))

(ĉz■ăĽŁéą)

```
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.

Nginx 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

ěřČèřĚ

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"
```

(äÿÑéąçžğçż■)

```
# 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.

17.1 Schema éĜ■èjĭ

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 *èŕĎæžŘáŕŇăěŮ*. 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 & èġEåŽ;

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.

æšlèċ: 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 `èġDæžŘåġÑåĚŮ`.

19.1 æřt'åžšèŁGæzd' (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
```

Adding multiple parameters conjoins the conditions:

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

These operators are available:

abbrev- viation	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
```

æşĳèċ: 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.

19.2 19.2 Columns (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 19.2.1.

19.2.1 19.2.1 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
```

19.3 19.3 Order

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` to order the results, even though the computed columns will not appear in the output.

19.4 Limit `items`

PostgREST 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
```

19.5 19.5 19.5

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.

19.6 19.6

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.

æšlëċ: 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.

19.7 19.7

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`.

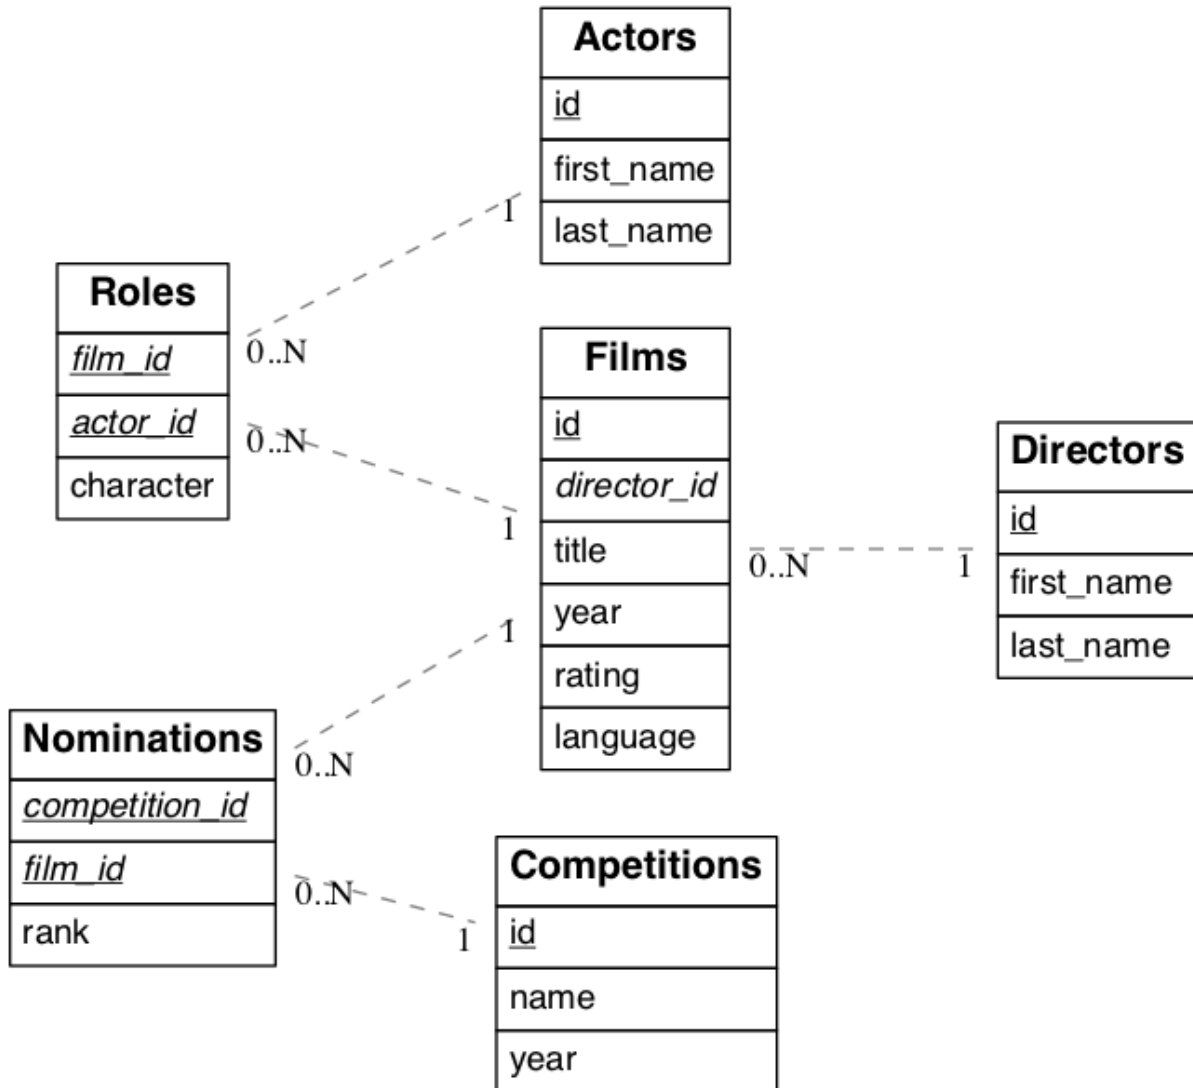
```
GET /items?select=bin_data&id=eq.1 HTTP/1.1
Accept: application/octet-stream
```

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 Lumière Factory In Lyon",
    "directors": {
      "last_name": "Lumière"
    }
  },
  { "title": "The Dickson Experimental Sound Film",
    "directors": {
      "last_name": "Dickson"
    }
  },
  { "title": "The Haunted Castle",
    "directors": {
      "last_name": "Méliès"
    }
  }
]
```

poznámka: 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
```

poznámka: 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 refresh*.

20.1 Embedded Tables

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'æŞěèřć

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 *èĬăžŘáŧŇăěŮ*
- 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.

zvlãc: 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.

zvlãc: 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.

22.1 èŔãŔŪèrŭæsŒŒ 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 22.2

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 }
```

22.3 22.3

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/ăŁóæŤź

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

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

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

(ăÿŦéatçžġčz■)

(çzäyŁeał)

```
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 *ãócaŁuçnrâžš* built for PostgREST.

To update a row or rows in a table, use the PATCH verb. Use *ært'ázšèŁĞæzd' (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 *ãđĆžT'èŁĞæzd' (Columns)*.

æšlēč: Beware of accidentally updating every row in a table. To learn to prevent that see *éŸæ■cãĚlèqáæš■ä;IJ*.

23.1 æL'zéĞRæRŠăĚě

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 *řádky (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 æŦræŦA

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.

æŦlêċ: 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 éĜèĵĵ](#).

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

èġŠèL'sçşzçz\$æęĆèĚř

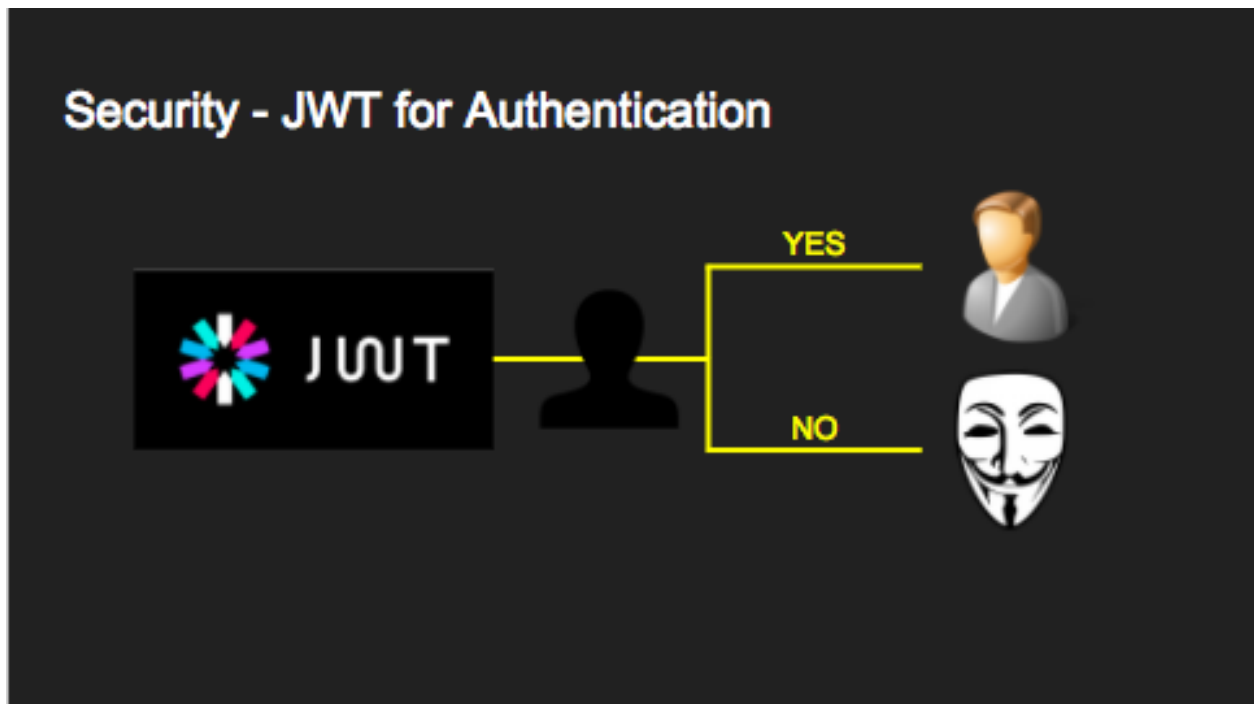
PostgRESTæÚlâIJlârEæTřæ■óâžŠä£IæÑAâIJlAPIâóL'âÉIæĂġçŽDäy■â£ČãĀĆ
æL'ĂæIJL'æÓLæiČéČ;éĂŽè£GæTřæ■óâžŠèġŠèL'sâŠÑæiČéŽRè£ZèqÑãĀĆ Post-
gRESTçŽDâüëä;IJæYř**éIÑèrA**erúæśĆ - â■şéIÑèrAâóçæLúçnræYřâRęæYřázÚázñæL'Ăèrt'çŽD
- çDúâRÓèól æTřæ■óâžŠ**æÓLæiČ**âóçæLúçnræŞ■ä;IJãĀĆ

27.1 éIÑèrAâžRâLÚ

the **authenticator**, **anonymous** and **user** roles. Post-
gRESTä;£çTlâyLçġ■çşzâdNçŽDèġŠèL'siijÑèžñâz;éIÑèrAâžRâLúijN**âÑ£âR■**âŠÑ**çTlæLú**èġŠèL'sãĀĆ
æTřæ■óâžŠçóaçRĚâSÝâlZâžžè£ZâžZèġŠèL'sâžúéĚ■ç;óPostgRESTäzèä;£çTlâóČazñãĀĆ



The authenticator is a role that is used to authorize requests. It is not a user and does not have any permissions. It is used to authorize requests from the API. The authenticator is used to authorize requests from the API. The authenticator is used to authorize requests from the API.



Here are the technical details. We use JSON Web Tokens to authenticate API requests.

æL'ÄæIJL'çt' cèŧTéČ; æYřáĚAèöÿçŽĎiijŇä; EPostgRESTçL'žáLnáĚšæşlāÿÄäÿlāR■äÿžèĝSèL'sroleçŽĎáčřæYŎā

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

ā; ŠèřúæśCāŇĚāRnáĚüæIJL'èĝSèL'sáčřæYŎçŽĎæIJL'æTŁJWTæŮüiijŇPostgRESTārEāIJĪHTTPèřúæśCæIJ

```
SET LOCAL ROLE user123;
```

èřúæşlāĎRiijŇéĀŽèĚĜāĚLāL■çŽĎæŞ■ā; IiijŇæTřæ■ōāžŞçōaçŘEāŚYāĚĚéāžāĚAèöÿèžñāz; éĪŇèřAèĀĚèĝ

```
GRANT user123 TO authenticator;
```

æĈæĎIJāóçæLúçnrāÿ■āŇĚāRnJWTiijLæLŮæşæIJL'èĝSèL'sáčřæYŎçŽĎJWTiijL'iijŇāLŽPostgRESTārEā
æTřæ■ōāžŞçōaçŘEāŚYāĚĚéāžæ■ççāōèö; ç; ōāŇĚāR■èĝSèL'sæĪĈéŽRiijŇāžééYšæ■cāŇĚāR■çTlæLúæşççIJŇæL

27.2 çTlæLúāŠŇçžĎ

PostgreSQL; çTlèĝSèL'sçŽĎæĈāĚçōaçŘEæTřæ■ōāžŞèóĚéŮóæĪĈéŽRāĀĈ
āRřāžèārEèĝSèL'sroleèĝEäÿžæTřæ■ōāžŞçTlæLúæLŮäÿĀçžĎæTřæ■ōāžŞçTlæLúüiijŇāĚüā; ŞāRŮāEşāžŎèĝSèL's

27.2.1 æRřäÿĪ Web çTlæLúçŽĎèĝSèL's

PostgRESTārřāžééĀĈāžTāzzā; TāÿÄäÿlèĝĈçĈzāĀĈ æĈæĎIJæĈlārEèĝSèL'sèĝEäÿžā■TāÿlçTlæLúüiijŇéĈā
ā; ŞçzŘèĚĜèžñāz; éĪŇèřAçŽĎçTlæLúāRŚāĜžèřúæśCæŮüiijŇPostgRESTārEāLĜæ■cāLřèřçTlæLúçŽĎèĝSèL'siij

æĈlārřāžèä; ççTlæāŇçžĝāōL'āĚlæĀĝçAġæt' žāIJřéŽRāLūā; ŞāL'■çTlæLúçŽĎāRřèĝAæĀĝāŠŇèóĚéŮóæĪĈéŽ
āžèäÿŇæYřāĚèĝĜTomas VondraçŽĎçd' žā; Ň<<http://blog.2ndquadrant.com/application-users-vs-row-level-security/>>' _iijŇèĚZæYřäÿÄäÿlā■YāĈlçTlæLúāžŇéŮt' āRŚéĀAçŽĎæŮLæAřçŽĎ
çTlæLúāRřāžèāIJlāĚüāÿ■æRŚāĚèèāŇāžèāRŚāĚüāžŮçTlæLúāRŚéĀAæŮLæAřiijŇāžüæşèèçāōĈāžæşççIJŇāĚü

```
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
);
```

æL'SāžñāÿŇæIJZāóĎæŮ; äÿĀéāžæTřç■ŮiijŇçāōāĪçTlæLúāRlèĈ; çIJŇāLřāžŮāRŚéĀAæLŮæL'ŞçóŮāRŚçz
æ■d' ād' ŮiijŇæL'SāžñèĚYäÿŇæIJZéYžæ■ççTlæLúā; ççTlæLúāžŮāžžçŽĎāĝŞāR■āiijlæĀmessage_fromāLŮāĀĈ

PostgreSQLiijL9.5āRĚæZt' éŇYçL'LæIJŇiijL'āĚAèöÿæL'Sāžñā; ççTlæāŇçžĝāōL'āĚlæĀĝèö; ç; ōæ■d' ç■ŮçTřèiij

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

èøÉèÛöçŤšæLŔçŽĐèAŁád'P'èaĹAPIçnrçCzçŽĐäzza;ŤäzzéČ;ārEçIJNāLŔrazŪāznāžŤeréāĜEçāoçŽĐèāNĭijN

27.2.2 Web çŤĪæLŪāĒšāžnèğŠèL's

æLŪèĀĒĭijNæŤŕæ■ōāžšèğŠèL'sāŔŕāzèāzçèaĹçzĐèĀNāy■æŸŕāyĹāLŭçŤĪæLŪĭijLæLŪāyĹāLŭnéZd'ād'ŪĭijL'ā
æCĪāŔŕāzèēĀL'æNĪ'WebāžŤçŤĪçĪNāžŔçŽĐæL'ĀæIJL'āušçŽzā;ŤçŤĪæLŪāĒšāžnāŔNāyĀāyĹwebuserèğŠèL'sāĀC
æCĪāŔŕāzèēĀŽèĚĜāĪĪJWTāy■āNĒāŔnēciād'ŪāçŕæŸŌæĪçŤĐāLŭ/æŌŠéZd'āĒūā;šæšŔāyĹçŤĪæLŪĭijNā;NāçC

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

SQLāzççāĀāŔŕāzèēĀŽèĚĜPostgreSQLæNĪ'èŕŭāšCèøç;ōçŽĐGUCāŔŸèĜŔèøÉèÛöçŕæŸŌāĀC
āĹNāçCĭijNèçĀèŌūāŔŪçŤĪā■ŔéCōāžūāçŕæŸŌĭijNèŕŭèŕçŤĪæ■d'āĜ;æŤŕĭijŽ

```
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.

27.2.3 æūūāŔĹçŤĪæLŪçzĐèğŠèL's

æNèæIJL'èøyād'ŽæŤŕæ■ōāžšèğŠèL'sæšæIJL'æĀğèČ;æ■šād'sĭijNār;çōāèğŠèL'sæŸŕæNĪ'ç; d'ÉZEāš;āŔ■è
āçCādĪĒĪĀèçĀĭijNæCĪāŔŕāzèèĜçŤsāyžWebāžŤçŤĪçĪNāžŔāy■çŽĐæŕŔāyĹçŤĪæLŪāLĒēĒ■æŪŕèğŠèL'sāĀC
æCĪāŔŕāzèæūūāŔĹçzĐāŠNā■ŤāyĹèğŠèL'sç■ŪçŤĒāĀC āĹNāçCĭijNæLŠāžnāz■çĐūāŔŕāzèæNèæIJL'āyĀāyĹwebuser

```
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)
```

27.3. Security

PostgreSQL provides a mechanism for checking the user of a request before allowing it to execute. This is done by setting the `pre-request` parameter in the `pg_hba.conf` file. The parameter is a string that is evaluated as a SQL expression. If the expression returns a non-zero value, the request is rejected. If the expression returns zero, the request is allowed.

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

The `check_user` function is defined in the `pg_catalog` schema. It takes a single argument, the name of the user, and returns a non-zero value if the user is not allowed to connect.

```
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

Bearer Auth

HTTP Authorization Bearer <jwt>

```
GET /foo HTTP/1.1
Authorization: Bearer eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.
  eyJyY2xiOiJpamRvZSIsImV4cCI6MTQ3NTUxNjI1MH0.
  GYDZV3yM0gqvuEtJmfpp1LBXSGYnke_Pvn10tbKAjB4
```

28.1 JWT Generation

JWT generation process involving header, payload, and signature.

28.1.1 JWT from SQL

JWT generation from SQL database using pgjwt extension.

PostgREST JWT token generation example. The SQL below creates a type and a function to generate a JWT token. The function uses the `row_to_json` function to convert a row to JSON and signs it with a secret key. The function also sets the expiration time to 300 seconds.

```
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 endpoint: `POST /rpc/jwt_test` with headers: `Content-Type: application/json`

Example response: `{ "token": "eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJ1aW50ZW50IjoiYXNjaWkiLCJpYXQiOiIyMDIyMDUyMjE0MjE0IiwiaWF0IjoiYXNjaWkiLCJ0eXciOiJ0b2tlbiJ9" }`

```
-- 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 ...
```

28.1.2 JWT from Auth0

Auth0 is a cloud-based authentication service. It provides a REST API for generating JWT tokens. The example below shows how to generate a JWT token using the Auth0 REST API.

The example shows how to generate a JWT token using the Auth0 REST API. The `secret` is the secret key used for signing the token. The `is-base64` parameter is set to `true` to indicate that the token is base64 encoded. The `Auth0` is the name of the Auth0 application.

æLŠázñçŽDžčçãAéIJĀèçAJWTäy■çŽDæTřæ■óázŞèğŞèL'sãĀCèçAæúzáLáãóĈijNæCíéIJĀèçAārEæTřæ■ metadata <<https://auth0.com/docs/rules/metadata-in-rules>> '_äy■ãĀCçDúãRŎijNæCíārEéIJĀèçAçijŮãEžZäyÄäy param <<https://auth0.com/docs/libraries/lock/v10/sending-authentication-parameters#scope-string>> 'äy■ãNĚãRñijŽcodeijŽ'role'áčræYŎãĀC

```
// 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'
  }
})
```

28.1.3 JWT áŎL'áĚÍ

árzázŎã;ççTÍJWTijNèGšārSæIJL'äyL'çg■äyÿèçAçŽDæL'zèrDijŽIijL'éŠLārzáæãGãGEæIJñèžñijN2ijL'ár äŞşázŎãĀIJJWTæãGãGE<<https://tools.ietf.org/html/rfc7519>> '_çŽDæL'zèrDãIJç;SâyŁãĚŮázŮãIJræŮzèrçç //paragonie.com/blog/2017/03/jwt-JSON-WebçŽDæãGèõř - æYř - äIRæãGãGEæYřijNæfRäyŁäžžijNãžTèrèèçDèAç> '_ãĀC Post-gRESTæIJĀçZÿãĚşçŽDèCíãLEæYřæL'ĀèřŞçŽDijŽázççãAijijŽal = none'èŮóécYãĀCäyÄäžZãóđçŎřJWTçŽDæIJ■ãLqãZláĚĀèðÿãóçæLùçñréĀL'æNl'çTlãžŎç■;ç;şJWTçŽDçŎUçşT = noneãĀC

árzJWTázŞçŽDæL'zèrDžžĚéĀŽèŁGãóĈã;ççTlçŽDžžŞâyŎPostgRESTçZÿãĚşãĀCæçCäyŁæL'ĀèŁfrijNäy■ã //auth0.com/blog/critical-vulnerabilities-in-json-web-token-libraries/> '_äy■æL'çãLræZt'ãd'ZãŁæAřãĀCæIJL'ãĀ <<https://jwt.io/>> '_ãĀC

æIJĀãRŎäyĀçg■æL'zèrDçŽDèG■çZæYřæzèçTÍJWTæIèçzt'æLd'ç;ŞçzIJäijŽèrIãĀCãŞzæIJñãžžèðøæYřãĀ //crypto.net/~joepie91/blog/2016/06/13/stop-using-jwt-for-sessions/> '_ãZäyžãd'gãd'ZæTřijLãçCæđIJäy■æYřãĀ ijNã;Şä;ããZçŽDæŮũãĀZãGžçŎřçŽDèŮóécYçŽDègçãEşæŮzæãLijN'äy■ãũëã;IJ<<http://crypto.net/~joepie91/blog/2016/06/19/stop-using-jwt-for-sessions-part-2-why> - æCíçŽDæžũæűş- çŁrèçDãũëã;IJ/> '_ãĀCéŞ;æŎèçŽDæŮGçñãæũsãĚèèðlèðžãZèŁZãžZéŮóécYřijNã;EçŮóécYçŽDãóđèťlæYřfJW

PostgRESTäyžèèAä;ççTÍJWTèŁZèãNèžñãz;éIñèfAãŠNæŎLæIĈijNãžúëijŞãŁşçTlæLũãžşèŁZæũãAžãĀC .._ssl:

28.2 SSL

PostgREST can be configured to use SSL for connections to PostgreSQL. To do this, you need to install a reverse proxy like Nginx or Apache and configure it to use SSL. For more information, see the Nginx documentation: https://nginx.org/en/docs/http/configuring_https_servers.html.

CHAPTER 29

æđúæđĐéŽŤçęž

PostgREST aóđä; NéĚ■ç; óäyžaĚñâijĀæIJ■āŁāZléĚ■ç; óæŪĜäzúäy■æŃĜáóŽçŽDā■ŤäyŁæÍaâijRçŽDæL'Āæ
èŁZæĐRāSşçİĀçġAæIJL'æŤræ■óæLŪáóđçŎřçzEèŁCāRfäzèèŁZāĚĚçġAæIJL'æÍaâijRiiijNāzúäyŤárzHTTPáóçæL
çĐúāRŎiiijNæCÍāRfäzèāĚñâijĀèġEāZ; āŠNā■ŸāCÍeŁĜçÍNiiijNāzŎèĀNārEāEĚéCÍçzEèŁCāyŎād'ŪéCÍayŪçŤNé
áóČä; ŁäzççāAæZt' áóžæŸŞéĜ■æđDiiijNāzúæRŘä; ŽäžEäyĀçġ■èĜŤDŭçŽDæŪzâijRæĪèèŁZæāNAPIçL'ĹæIJñæŎ
æIJL'āĚşä; ŁçŤÍāĚñāĚšèġEāZ; āŃĚèçĚçġAæIJL'èaŁçŽDçd'žä; NiiijNèrúāRCéŸĚäyNéÍççŽDiiijŽrefiiijŽ'public_ui'ç

SQL ęTŁæŁuęóąęŘĚ

30.1 āYāĆłęTŁæŁuāŠŇārĚęăA

āęĆāyŁæL'ĀēřřijŇād' ŪēČłæIJ■āŁāāRřāzēæRŘā; ŽęTŁæŁuęóąęŘĚāzūā; ŁęTŁJWTāyŌPostgRESTæIJ■āŁāā
āzšāRřāzēāōŇāĒléĀŽēŁGSQŁæTřæŇĀęZzā; TāĀĆ ēŁZæYřāyĀēāzęZyā; Šād' ŽęZDāuēā; IJřijŇæL'ĀāzēāĠĚād' Č

āyŇēāłijŇāĠ; æTřāŠŇēęęāRŠāZłārĚā■YāIJlāžŌijŽcodeijZ'basic_auth'āłāāijRāy■ijŇæČłāy■āzTāIJĀPI
āĚŇāĒsēęĚāZ; āŠŇāĠ; æTřārĚā■YāIJlāžŌāy■āRŇęZDāłāāijRāy■ijŇēřēāłāāijRāIJlāĒĚēČłāijTęTŁæŁ'āĒĚēČłā

ēęŪāĒŁijŇæLŠāznēIJĀēęĀyĀāyłēāłāēĒēūšēyłæLŠāznęZDęTŁæŁūijŽ

```
-- æŁŠāznārĚāĒēĀōzę; ōāžŌbasic_authāłāāijRāy■ijŇ
-- āzēārĚāĒūēžŘēŪRāIJlāĒŇāĒsēęĚāz; āy■āĀĆ
-- æŠRāzžāĒŇāĒsēēŁĠčłŇ/ēęĚāz; āřĚāijTęTŁāĒĚēČłęZDāyōāŁl'čłŇāzřāŠŇēāłāĀĆ
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)
);
```

æŁŠāznāyŇæIJZēřēęŠēL'sroleYřāōđēZĚæTřæ■ōāžšēęŠēL'sęZDād' ŪēTŌijŇā; ĚæYřPostgreSQLāy■æTř
æŁŠāznārĚā; ŁęTŁēęęāRŠāZłæL'ŇāŁlāijžāŁūæL'ęēāŇāōČāĀĆ

```

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();

```

æÖçäyÑæIëijÑæLŠäznârEä;fcTlpgcryptoæLl'ásTãŠÑègæãRSãZlæIëäfIârEârEçãAijZcödeijZ'users'èalã

```

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();

```

ä;fcTlërèèalijÑæLŠäznârRräzèäyöãLl'æçÄæšèãLârEãLÛçZDârEçãAãÄÇ
æÇædIJçTlã■RèCöázúãŠÑârEçãAæ■ççãoijÑãóÇârEèfTãZçTlæLüçZDæTlæ■óázšègšèL'sãÄÇ

```

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

```

30.2 Public `user_role` Function

The `user_role` function is used to verify the email and password of a user. It returns the role name if the user is found and the password is correct, otherwise it returns `NULL`.

30.2.1 `login` Function

The `login` function is used to generate a JWT token for a user. It takes the email and password as input and returns the JWT token if the user is found and the password is correct, otherwise it returns an error message.

```

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,

```

(continued)

