

APPLICATION AUTHORIZATION WITH SET ROLE

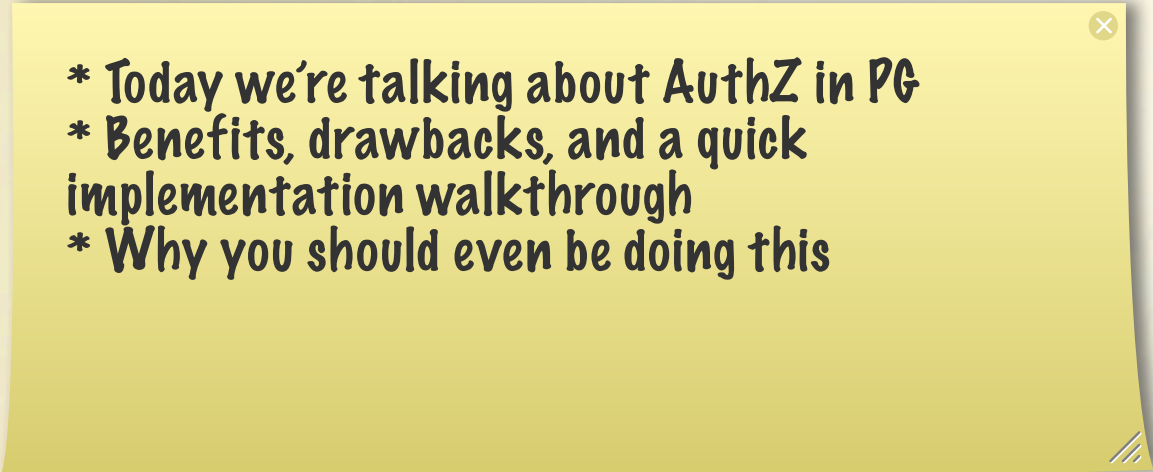
AURYNN SHAW, COMMAND PROMPT, INC.
PGCON 2010



Hi

Hi

- Aurynn Shaw
 - DBA/Lead Dev/PM/etc @ Command Prompt

- 
- * Today we're talking about AuthZ in PG
 - * Benefits, drawbacks, and a quick implementation walkthrough
 - * Why you should even be doing this

AND NOW I TALK MORE



PERMISSIONS SYSTEMS

PERMISSIONS SYSTEMS

- Data I can access

PERMISSIONS SYSTEMS

- Data I can access
- Data I can't access

*** Postgres handles this through standard GRANT and REVOKE statements.**

*** Most app fabrics handle this away from the data.**

APP-FOCUSED DESIGN

- * Easy enough to use the app to handle permissions
- * Few-no restrictions on application powers.
- * Permissions happen when the logic happens

DATA LAYER DISCONNECT

- * App-focus development treats the DB as a dumb store
- * Because the app embodies the AuthZ, the data fabric is at its most malleable.
- * Nothing stops a malicious, or badly-written app from unlimited data modification
- * Relying on a limited number of eyeballs to look for bugs

DB-FOCUSED DESIGN

- * Much tighter binding to the data layer
- * We can put permissions into the database, GRANT and REVOKE!
Not a problem!
- * Creating a user, not a problem. Everything can just work!

AUTHZ IS CLOSELY COUPLED TO YOUR LOGIN

- * Poolers, especially, have only a defined login
- * Forced into the broadest permissions set available
- * Can't attempt to restrict the data malleability - everything your app needs to do, your login has to be able to do, regardless of whether the user should.

WAIT, I LIED.

*** Your login DOES embody your core permissions, but, there's this great permissions-swapping feature in PG.**

SET ROLE TO STUN;

HEY, THIS IS IN THE TALK TITLE!

- * I've seen this before! It's like a Unix system!
- * So, SET ROLE is the funky mojo
- * Similar to SET AUTHORIZATION
- * Can be unwound - a very valuable aspect.

CAN ONLY SWITCH TO ROLES ALREADY IN YOUR TREE

only allows you to become roles you would have been able to be already - you can't just become a superuser, unless you already are one.

* By default, all the roles you have are already part of your user

WHY SET ROLE IS INTERESTING

- * Can swap permissions dynamically, without compromising the base connection
- * Vital in any pooled environment - long-lived connections don't need to be reset.
- * Trusted apps can easily set the data fabric to just the permissions they need
- * Can never exceed base fabric permissions

TRANSACTIONAL, TOO!

- * Single transactions can be in their own permissions space
- * Automatic, implicit `RESET ROLE` command on `ROLLBACK`

TRANSACTIONAL, TOO!

```
template1=# BEGIN;  
BEGIN  
template1=# SET ROLE test;  
SET  
template1=> ROLLBACK;  
ROLLBACK  
template1=#
```



A quick example.

TRANSACTIONAL, TOO!

```
test=> BEGIN;  
BEGIN  
test=> SET ROLE test;  
SET  
test=> SELECT * FROM test;  
--  
(0 rows)  
  
test=> ROLLBACK;  
ROLLBACK  
test=> SELECT * FROM test;  
ERROR: permission denied for relation test
```



And another

WELL, PARTLY.

```
template1=# BEGIN;  
BEGIN  
template1=# SET ROLE test;  
SET  
template1=> COMMIT;  
COMMIT  
template1=>
```

So it doesn't quite work like you'd expect for a committed transaction.

SO ALWAYS RESET ROLE

```
template1=# BEGIN;  
BEGIN  
template1=# SET ROLE test;  
SET  
template1=> COMMIT;  
COMMIT  
template1=> RESET ROLE;  
RESET  
template1=#
```

So it doesn't quite work like you'd expect for a committed transaction.

OUR WHY

- * Explored this to support a large Web application with very clear-cut access rules: A resource either is or isn't accessible.
- * In-app frameworks were insufficient - and not useful when we needed external software- Rewriting perms is a pain.

OTHER COOL WHYS

- * **Single definition of our permissions model, as close to the relevant data as possible.**
- * **Don't Repeat Yourself**
- * **Non-trusted clients can't manipulate your data fabric beyond your whim - you already have strong permissions on the data itself.**

BUT THERE'S ALL THOSE OTHER PERMISSIONS SYSTEMS...

- * Lots, in a variety of languages**
- * Including that one you're working on right now**
- * And that other one YOU LOVE.**
- * Should you use them? They work, to a point**
- * Valuable aspect of the permissions setup**
- * Exclusive use ends up looking like THIS**



THIS

- Data I can access
- Data I can't access
- **Data I shouldn't access, but can**

* Normal pooled application, single credentials relies on app to handle auth

* Never more than a strong warning about not using a resource, and some unfriendly language from your DBA.

PRINCIPLE OF LEAST PERMISSION

STOLEN FROM STEVEN FROST

- * You should never have more ability than you need.
- * Any time you do, Bad Things can happen.
- * In-app permissions systems tend to violate this

IMPLEMENTATION

(IT'S EASY)

So, let's look at how to go about implementing a SET ROLE-based system in your application framework.
It's surprisingly easy to do, too!
Let's begin.



GRANT AND REVOKE

First, a fairly core component is that you have to go through and **GRANT**, and **REVOKE** the various tables and views and suchly that make up your database.

REVOKE

```
test=# CREATE TABLE test ();
```

```
CREATE TABLE
```

```
test=# REVOKE ALL ON test FROM PUBLIC;
```

```
REVOKE
```

```
test=# SET ROLE TO test;
```

```
SET
```

```
test=> SELECT * FROM test;
```

```
ERROR: permission denied for
```

```
test
```

```
test=>
```

A simple REVOKE example.

GRANT

```
test=> SET ROLE TO aurynn;  
SET  
test=# GRANT ALL ON test TO test;  
GRANT  
test=# SET ROLE TO test;  
SET  
test=> SELECT * FROM test;  
--  
( 0 rows )  
test=>
```

And a GRANT

A PERMISSIONS TREE

Next, a permissions tree.

This aspect of a SET ROLE design is really, really, really dependent on your application structure.

To really get the most benefit from a SET ROLE environment, you should spend some time laying out every single last permission that you want to have - as fine-grained as you can. This ends up being very valuable later, when you need to add less trustworthy clients.

A PERMISSIONS TREE

```
CREATE ROLE content_read NOLOGIN;  
CREATE ROLE content_write NOLOGIN;  
CREATE ROLE content_delete NOLOGIN;
```

A PERMISSIONS TREE

```
CREATE ROLE user_base NOLOGIN;  
GRANT content_read TO user_base;  
GRANT content_write TO user_base;  
CREATE ROLE admin_base NOLOGIN;  
GRANT content_delete TO admin_base;  
GRANT user_base TO admin_base;
```


YOUR FINAL NODE POINTS

USER, ADMIN, MODERATOR, ETC.

Your final node points are the specific roles that a given user is going to be granted into - users, moderators, administrators, whatever. Your software would then issue `SET ROLE TO your_user_role` at the beginning of your transaction.

Caveat: Custom permissions are hard.

PERMISSIONS ENDPOINTS

```
CREATE USER user NOINHERIT;  
GRANT user_base TO user;  
CREATE USER admin NOINHERIT;  
GRANT admin_base TO admin;
```

NOINHERIT

The next piece is NOINHERIT. Right now, without this, you'd not exactly be restricting your permissions set - just granting the full set of useful permissions to a more limited, non-superuser user. Pretty much exactly the same as before.

With NOINHERIT, we mark that those endpoint roles that we just defined aren't applied to our login role - we have to explicitly SET ROLE to grab those permissions.

A FULLY REVOKE'D, LOGIN USER

- * The credentials that the application/pooler/whatever uses to connect.
- * This has pretty much every single possible permission, removed. All this role can do is SET ROLE to a different role, and pick up those permissions.
- * By default, no connections can actually do anything useful.

APPLICATION MODIFICATIONS

Lastly, modify your application. It's somewhat obvious, but it has to be said.

IT'S JUST THAT EASY!

**You've now successfully integrated a SET ROLE-based permissions system into your application.
It's just that easy.**



**I LIED AGAIN.
IT'S NOT QUITE THAT EASY**

Well, it's almost that easy. There are some bits that you do have to pay attention to, that you wouldn't otherwise



IT'S NOT QUITE THAT EASY

At least in Python's `psycopg2`, permissions errors aren't mapped to something useful - you have to handle it yourself.

- You have to catch permissions errors

PLUG, THE SHAMELESS KIND



- Exceptable, an exception-trapping library for Python

PLUG, THE SHAMELESS KIND



- Exceptable, an exception-trapping library for Python
- Turns PG exceptions into smarter Python exceptions.

PLUG, THE SHAMELESS KIND



- Exceptable, an exception-trapping library for Python
- Turns PG exceptions into smarter Python exceptions.
- We could use help with this - other language support &c.

IT'S NOT QUITE THAT EASY

- You have to catch permissions errors
- It's really coarsely grained

For one, this is fairly coarse-grained - you can restrict tables, but not individual rows in those tables. For that, there's nothing to be done but write a stored procedure, or a view that checks whether or not the user **can** read those roles.

The same applies for writes, obviously - but, that's a bit easier to solve with triggers to verify per-row permissions, as opposed to the per-table permissions.

IT'S NOT QUITE THAT EASY

*** You need to add the SET ROLE
mojo before you start running
queries**

- You have to catch permissions errors
- It's really coarsely grained
- Requires modifications to the DB interface

IT'S NOT QUITE THAT EASY

You actually have to send the **SET ROLE** and possibly **RESET ROLE** commands.

- You have to catch permissions errors
- It's really coarsely grained
- Requires modifications to the DB interface
- Adds additional wire traffic


IT'S NOT QUITE THAT EASY

- You have to catch permissions errors
- It's really coarsely grained
- Requires modifications to the DB
- Adds additional wire traffic
- Just as vulnerable to SQL injection as you were before

This doesn't give you any real additional protection against SQL injection attacks - it's pretty much security-by-obscurity at best, by requiring **SET ROLE** before your injection.

It does, however, grant you protection against random **DELETE** and **DROP** crap, which is good for something.

SO ALWAYS SANITIZE YOUR INPUTS.



It's just good data hygiene. Like brushing your teeth.

IT'S NOT QUITE THAT EASY

- You have to catch permissions errors
- It's really coarsely grained
- Requires modifications to the DB
- Adds additional wire traffic
- Just as vulnerable to SQL injection as you were before
- Not entirely transactional

As I showed you before, it's not really transactional - you have to pay pretty close attention to your **RESET ROLE** statements.

set session_authorization

- * The difference between SET ROLE and SET session_auth is a matter of semantics, mostly - both achieve the same effect.
- * set session_authorization changes what roles are available to SET to, though

set session_authorization

```
test=# SET SESSION_AUTHORIZATION TO pgcon;
```

```
SET
```

```
test=> SET ROLE TO aurynn;
```

```
ERROR: permission denied to set role  
"aurynn"
```


```
test=> SET SESSION_AUTHORIZATION TO  
aurynn;
```

```
SET
```

```
test=#
```

- * Remarkably similar
- * alters what roles are reachable from future SET ROLE requests.
- * Useful from perspective of additional layers of restriction over the connection

**SO THAT'S IT.
ANY QUESTIONS?**



THANK YOU!

SLIDES WILL BE AVAILABLE.