

# Writing Safe Postgres extensions with Rust





# Table des matières

Damien Clochard . . . . .	5
My journey . . . . .	5
My Story . . . . .	5
Menu . . . . .	5
A unique ecosystem . . . . .	7
So what is an Postgres extension ? . . . . .	7
Writing an extension in SQL or PL code . . . . .	7
20+ procedural languages . . . . .	7
PL/Rust ? . . . . .	8
Writing an extension C . . . . .	8
Can we have the best of both ? . . . . .	9
PGRX : a breath of fresh air . . . . .	9
PGRX : playing with crabs and elephants . . . . .	9
PGRX : principles . . . . .	10
PGRX : type conversions . . . . .	10
PGRX : A bridge between 2 worlds . . . . .	10
PGRX : Safety Fist . . . . .	10
PGRX : Modern development tooling . . . . .	11
PGRX : an open and active community . . . . .	11
A practical example . . . . .	12
Let's go ... . . . .	12
... to a new world ! . . . . .	12
src/lib.rs . . . . .	12
Let's try this ! . . . . .	12
Bonjour tout le monde ! . . . . .	13
Add a parameter . . . . .	13
Rebuild and try again . . . . .	13
Bonjour FOSDEM ! . . . . .	13
Same function with C ? . . . . .	14
Any Canadians in the room ? . . . . .	14
Do not reinvent the wheel . . . . .	15
Compute the checksum . . . . .	15
Rebuild the extension . . . . .	15
Manual testing . . . . .	15
Automatic testing . . . . .	16
Launch the tests . . . . .	16
Launch tests on Postgres 14 . . . . .	16
Create a User Defined Type (UDT) . . . . .	17
Implement basic I/O traits . . . . .	17

---

Input trait . . . . .	17
Output trait . . . . .	17
Let's try that new type . . . . .	18
Convert a TEXT into SIN . . . . .	18
Use this type in a column . . . . .	19
Nope . . . . .	19
Derive the default operators . . . . .	19
Now we can compare 2 SINS . . . . .	19
Now the column works . . . . .	20
There's so much more.... . . . . .	20
Feedback . . . . .	21
From C to Rust . . . . .	21
Rewriting PostgreSQL Anonymizer from scratch . . . . .	21
A feeling of « déjà vu » . . . . .	21
Immediate Gains . . . . .	21
Culture Shock . . . . .	22
It's not magic . . . . .	22
The Postgres Future is Rusty ! . . . . .	23
A new generation of extensions . . . . .	23
A great vector of innovation . . . . .	23
Join the revolution . . . . .	23
Links . . . . .	24
MERCI ! . . . . .	25

## Damien Clochard

- Co-founder of DALIBO, leading Postgres company in France since 2005
  - Active member of the French PostgreSQL community
  - Main developer of the PostgreSQL Anonymizer<sup>1</sup> extension
- 

## My journey

I discovered Postgres 25 years ago

I discovered Rust last year

---

## My Story

In 2018, I started a project called PostgreSQL Anonymizer<sup>2</sup>

Over the years, I wrote more and more C code...

Last year, I rewrote everything in Rust

This is my story :)

---

## Menu


- What are Postgres Extensions ?
  - The PGRX Framework
  - A practical example
  - Lessons learned from Postgres Anonymizer 2.0
  - Postgres future is rusty !
- 

<sup>1</sup>[https://gitlab.com/dalibo/postgresql\\_anonymizer/](https://gitlab.com/dalibo/postgresql_anonymizer/)

<sup>2</sup>[https://gitlab.com/dalibo/postgresql\\_anonymizer/](https://gitlab.com/dalibo/postgresql_anonymizer/)

## The Fallacy of the « Right Tool for the Job »

But there are so many possibilities on the modern web!.. Object stores! Key-value stores! Event sourcing! Time series! Graph databases! Blockchain! You can't just use the same tool for everything!



## Postgres is not a database, it's a platform

- **Graphs** ? Apache AGE, EdgeDB
- **Geo Data** ? PostGIS, pg\_pointcloud, pg\_routing
- **OLAP** ? pg\_DuckDB, Citus, Hydra, pg\_analytics
- **NoSQL** ? JSONB, FerretDB
- **AI** ? pgvector, PostgresML
- **ETL** ? 100+ Foreign Data Wrappers
- **Timeseries** ? TimescaleDB
- **Full Text Search** ? pgroonga, ParadeDB, zombodb
- **API** ? pg\_graphql, postgREST
- **Pluggable storage** ? OrioleDB, Neon
- etc....

## A unique ecosystem

More than 1000 known extensions<sup>3</sup>

... almost 250 are active and maintained

---

## So what is an Postgres extension ?

- Some SQL objects
  - and/or Procedural Language ( PL ) code
  - and/or a compiled library
- 

## Writing an extension in SQL or PL code

- Easy
  - A great way to share code between several databases
  - Very stable between major versions
  - Slow
  - 20 procedural languages<sup>4</sup> (PL)
- 

## 20+ procedural languages

---

PL/pgsql	PL/perl	PL/php	PL/Ruby	PL/Java
PL/Scheme	PL/tcl	PL/Lua	PL/python	PL/haskell
PL/Rust	PL/dotnet	PL/lolcode	PL/Julia	PL/sh
PL/XSLT	PL/R	PL/v8	PL/go	PL/brainfuck

---

<sup>3</sup><https://gist.github.com/joelonsql/e5aa27f8cc9bd22b8999b7de8aee9d47>

<sup>4</sup>[https://wiki.postgresql.org/wiki/PL\\_Matrix](https://wiki.postgresql.org/wiki/PL_Matrix)

```
CREATE OR REPLACE FUNCTION get_employee_name(emp_id INTEGER)
RETURNS VARCHAR AS $$
DECLARE
    emp_name VARCHAR;
BEGIN
    SELECT first_name || ' ' || last_name INTO emp_name
    FROM employees WHERE id = emp_id;
    RETURN emp_name;
END;
$$ LANGUAGE SQL;
```

```
SELECT get_employee_name(123);
```

---

## PL/Rust ?

- It's great
  - We're not going to talk about it today :)
  - <https://plrust.io/>
- 

## Writing an extension C

- Generally loaded when the instance starts
  - Fast and Direct access to the internal functions
  - Very low-level code / No Abstractions
  - Each new major version will probably break your extension
  - The dev/test framework (PGXS) is *very* limited
  - Absolutely no security barrier => segfaults fest
  - if the extension crashes, the entire Postgres instance will crash too
-



if the extension crashes,  
the entire Postgres instance will crash too


---

### Can we have the best of both ?

Safety of PL functions **AND** Performances of C  
High Level Abstractions **AND** Access to Postgres internals  
A modern language **AND** Stability

---

### PGRX : a breath of fresh air



### PGRX : playing with crabs and elephants

A framework that bridges the gap between rust and postgres  
You can write Rust extensions without it  
But it makes your life easier !

---

## PGRX : principles

- Expose your rust functions as user functions inside postgres
  - Automatic mapping postgres data types into rust types
  - ... and vice versa
- 

## PGRX : type conversions

---

postgres	rust
BYTEA	Vec<u8> or &u8
TEXT	String or &str
INTEGER	i32
DATE	pgrx::Date
DATERANGE	pgrx::Range<pgrx::Date>
NULL	Option::None

---

## PGRX : A bridge between 2 worlds

- Derive attributes and macros to export Rust functions in SQL
  - Rust abstractions over Postgres pointers (pgBox<T>)
  - Helpers to exchange memory with Postgres
  - Safe access to the Server Programming Interface (SPI)
  - Any Rust panic! is translated into a Postgres ERROR
- 

## PGRX : Safety Fist

If the extension crashes

An ERROR event is raised into Postgres

The transaction is cancelled (ROLLBACK)

The current session lives on

The Postgres instances survives

---

### **PGRX : Modern development tooling**

- A fully managed development environment ( cargo-pgrx )
  - All major versions are supported
  - An idiomatic Rust Test framework
  - A very nice development feedback loop
  - Easy commands to build and ship packages
- 

### **PGRX : an open and active community**

- Project launched by a single company (TCDI)
  - Transferred last year to the pgcentral foundation
  - A friendly Discord channel for beginners
-

## A PRACTICAL EXAMPLE

---

### Let's go ...

```
cargo install --locked cargo-pgrx
cargo pgrx init
```

---

### ... to a new world !

```
cargo pgrx new world
cd world
```

---

### src/lib.rs

```
#[pg_extern]
fn hello_world() -> &'static str {
    "Hello, world"
}
```

---

### Let's try this !

```
cargo pgrx run
```

---

## Bonjour tout le monde !

```
world=# CREATE EXTENSION world;

world=# SELECT hello_world();
 hello_world
-----
Hello, world
(1 row)
```

---

## Add a parameter

```
#[pg_extern]
fn hello(name: &str) -> String {
    format!("Hello, {name}")
}
```

---

## Rebuild and try again

```
cargo pgrx run
```

---

## Bonjour FOSDEM !

```
world=# DROP EXTENSION world;
world=# CREATE EXTENSION world;
```

```
world=# SELECT hello('FOSDEM');
      hello
-----
Hello, FOSDEM
(1 row)
```

---

## Same function with C ?

```
PG_FUNCTION_INFO_V1( hello );
Datum hello( PG_FUNCTION_ARGS ) {
    char hello[] = "Hello, ";
    text * name;
    int hellolen;
    int namelen;
    text * msg;
    name = PG_GETARG_TEXT_P(0);
    hellolen = strlen(hello);
    namelen = VARSIZE(name) - VARHDRSZ;
    msg = (text *)palloc( hellolen + namelen );
    SET_VARSIZE(hello, hellolen + namelen + VARHDRSZ );
    strncpy( VARDATA(msg), hello, hellolen );
    strncpy( VARDATA(msg) + hellolen, VARDATA(name), namelen );
    PG_RETURN_TEXT_P( msg );
}
```

---

## Any Canadians in the room ?

The Canadian Social Insurance Number<sup>5</sup> (SIN) is composed of  
8 digits + 1 control digit

046 454 286

The control digit is computed using the Luhn Formula<sup>6</sup>

$\text{Luhn}("046\ 454\ 28") = 6$

---

<sup>5</sup>[https://en.wikipedia.org/wiki/Social\\_insurance\\_number](https://en.wikipedia.org/wiki/Social_insurance_number)

<sup>6</sup>[https://en.wikipedia.org/wiki/Luhn\\_algorithm](https://en.wikipedia.org/wiki/Luhn_algorithm)

## Do not reinvent the wheel

```
cargo add luhn3
```

---

## Compute the checksum

```
#[pg_extern]
fn luhn_checksum(input: &str) -> char {
    use luhn3::decimal::checksum;
    checksum(input.as_bytes())
        .expect("Input should be decimal")
        as char;
}
```

---

## Rebuild the extension

```
cargo pgrx run
```

---

## Manual testing

```
SELECT luhn_checksum('04645428');
 luhn_checksum
-----
6
```

```
SELECT luhn_checksum('A');  
ERROR:  Input should be decimal
```

---

## Automatic testing

```
#[cfg(any(test, feature = "pg_test"))]  
#[pg_schema]  
mod tests {  
    use pgrx::prelude::*;  
  
    #[pg_test]  
    fn test_luhn_checksum() {  
        assert_eq!("8", crate::luhn_checksum("1"));  
    }  
}
```

---

## Launch the tests

```
cargo pgrx test  
[...]  
test tests::pg_test_luhn_checksum ... ok  
  
test result: ok. 1 passed; 0 failed; 0 ignored; 0 measured; 0 filtered out;  
finished in 5.26s
```

---

## Launch tests on Postgres 14

```
cargo pgrx test pg14
```

---



## Create a User Defined Type (UDT)

```
#[derive(PostgresType, Serialize, Deserialize, Debug)]  
#[inoutfuncs]  
pub struct SIN (i32);
```

---

## Implement basic I/O traits

```
impl InOutFuncs for SIN {  
  
    fn input(input: & core::ffi::CStr) -> Self { ... }  
  
    fn output(&self, buffer: &mut pgrx::StringInfo) { ... }  
  
}
```

---

## Input trait

```
fn input(input: & core::ffi::CStr) -> Self {  
    use luhn3::decimal::valid;  
    let val = input.to_str().expect("Invalid Input").replace(' ', "");  
    if ! valid(&val.clone().into_bytes()) {  
        error!("{}", "Not a valid SIN");  
    }  
    SIN(val.parse::<i32>()).expect("Value should be a number") / 10  
}
```

---

## Output trait

```
fn output(&self, buffer: &mut pgrx::StringInfo) {
    use luhn3::decimal::checksum;
    let part1 = self.0 / 100000 % 1000;
    let part2 = self.0 / 100 % 1000;
    let part3 = self.0 % 100;
    let part4 = checksum(&self.0.to_string().into_bytes())
        .expect("Checksum Failed")
        as char;
    let val = format!("{part1:03} {part2:03} {part3:02}{part4}");
    buffer.push_str(val.as_str());
}
```

---

## Let's try that new type

```
cargo pgrx run
```

---

## Convert a TEXT into SIN

```
SELECT CAST ( '046454286' AS SIN );
      SIN
-----
046 454 286
```

```
SELECT CAST ('999 999 999' AS SIN);
ERROR:  Not a valid SIN
```

## Use this type in a column

```
CREATE TABLE canadians (  
  id SIN PRIMARY KEY,  
  name TEXT  
);
```

---

## Nope

```
CREATE TABLE canadians (  
  id SIN PRIMARY KEY,  
  name TEXT  
);
```

ERROR: **data type sin** has **no default operator class for access method "btree"**

HINT: You must specify an **operator class for the index** or define a **default operator class for the data type**.

---

## Derive the default operators

```
#[derive(PostgresType, Serialize, Deserialize, Debug, PartialEq)]  
#[derive(Eq, PartialEq, Ord, Hash, PartialOrd)]  
#[derive(PostgresEq)]  
#[derive(PostgresOrd)]  
#[inoutfuncs]  
pub struct SIN(i32);
```

---

## Now we can compare 2 SINS

```
SELECT '483247862'::SIN > '483247870'::SIN ;
```

```
?column?
```

```
-----  
f  
(1 row)
```

---

## Now the column works

```
CREATE TABLE canadians (  
  id SIN PRIMARY KEY,  
  name TEXT  
);
```

```
INSERT INTO canadians VALUES ('483247862','James Howlett');
```

```
SELECT * FROM canadians;  
  id      | name  
-----+-----  
483 247 862 | James Howlett
```

---

## There's so much more....

You can also interact with the database engine itself!

- Foreign data wrappers
  - 30+ Hooks
  - WAL decoders for logical replication
  - Index Access Methods
  - Table Access Methods
-

## FEEDBACK

### From C to Rust

---

#### Rewriting PostgreSQL Anonymizer from scratch

- A data masking extension for PostgreSQL
  - I'll talk about it tomorrow at 15h00 in the Postgres devroom (UA2.220)
  - About 1000 lines of C code
  - Rewrote everything in a few weeks, without prior knowledge of Rust
- 

#### A feeling of « déjà vu »

There's some unspoken familiarity between Postgres and Rust

- The Rust compiler is dull and rough at the beginning
  - But once you climbed that learning curve, you're rewarded
  - .... pretty much like Postgres :)
- 

#### Immediate Gains

- Comfort of development
  - Dozens of unit tests => many bugs found along the way
  - Better performance by rewriting some PL/pgSQL in Rust
  - Using high level Rust crates ( `faker-rs`, `image` )
  - Stability ( « no more segfaults ! » )
-

## Culture Shock

- In Rust a variable is never NULL !
  - Some Postgres internal macros and some bindings are missing
  - Handling 2 memory contexts at once
- 

## It's not magic

- A lot of sections in the code are still unsafe
  - Building is very very slow (about 20x slower than C)
  - No support de Windows at the moment
  - For advanced features, I still need to read and understand the Postgres C code
-

## THE POSTGRES FUTURE IS RUSTY !

---

### A new generation of extensions

- supabase wrappers
  - PL/PRQL
  - Timescaledb-toolkit
  - pg\_graphql
  - pgvecto.rs
  - pg\_later
  - paradeDB
  - pgmq
  - neon
  - pgzx
- 

### A great vector of innovation

- New storage engines
  - New index methods
  - New connectors
  - New ideas that would never be accepted in the Postgres codebase
- 

### Join the revolution

- Bring back your code close to the data
  - Define your own types !
  - Use Postgres as a platform
  - Rust extensions are a great entrypoint to the Postgres community
-

## Links

PGRX

<https://github.com/pgcentralfoundation/pgrx>

A 4 hour tutorial

<https://daamien.gitlab.io/pgrx-tuto/>

Try out PostgreSQL Anonymizer !

[https://gitlab.com/dalibo/postgresql\\_anonymizer](https://gitlab.com/dalibo/postgresql_anonymizer)

---



**MERCI !**

