

# POSTGRESQL NETWORK FILTER FOR ENVOY PROXY

FABRÍZIO DE ROYES MELLO ÁLVARO HERNÁNDEZ

# **UNGRES**

#### `whoami`

- 25+ years on IT
- PostgreSQL Developer at OnGres
- PostgreSQL Contributor
- Brazilian Community Leader



Fabrízio de Royes Mello

<<u>fabrizio@ongres.com</u>>

สิfabriziomello



#### `whoami`

- Founder & CEO, OnGres
- 20+ years Postgres user and DBA
- Mostly doing R&D to create new, innovative software on Postgres
- Frequent speaker at Postgres, database conferences
- Founder and President of the NPO <u>Fundación PostgreSOL</u>
- AWS Data Hero



**Álvaro Hernández** <a href="mailto:aht@ongres.com">aht.es



# ENHANCING POSTGRES OBSERVABILITY

UNGRES

#### **Postgres monitoring**

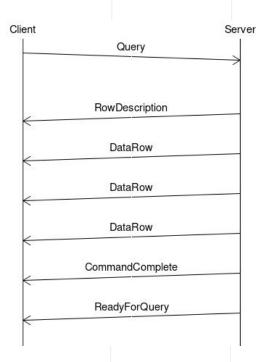
- Not in-core integrated solution.
- Postgres provides catalog views with rich monitoring information.
- But that means making queries to gather monitoring data.
- Eg. Prometheus exporter: dozens/hundreds/... of queries per monitoring cycle.
- Postgres monitoring extensions:
  - may require restart -> downtime.
  - may require configuration / external binaries-> complexity

#### Can we do better?



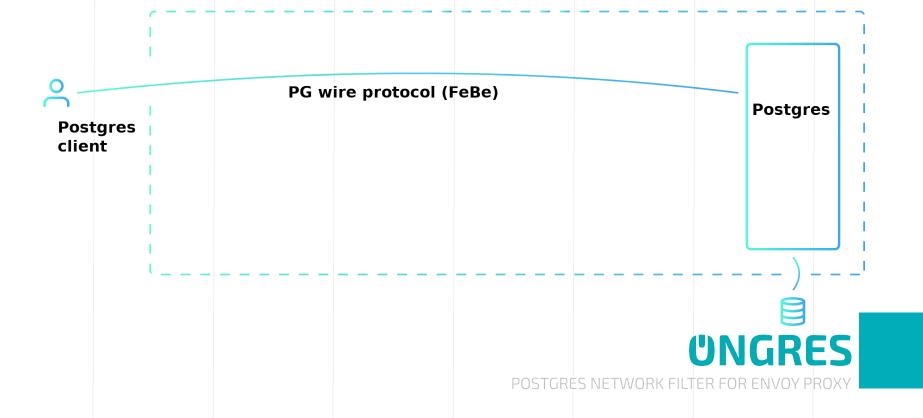
#### Postgres wire protocol (FeBe)

- Custom (layer 7) TCP protocol. Well documented.
- Well structured and defined messages (no "generic message" for many things).
- Very stable (current v3 was introduced in 2003 with PG 7.4).
- Implemented by countless tools and drivers.
- Used also for non-Postgres databases (Yugabyte, CockroachDB, Crate.io, NoisePage, ...).

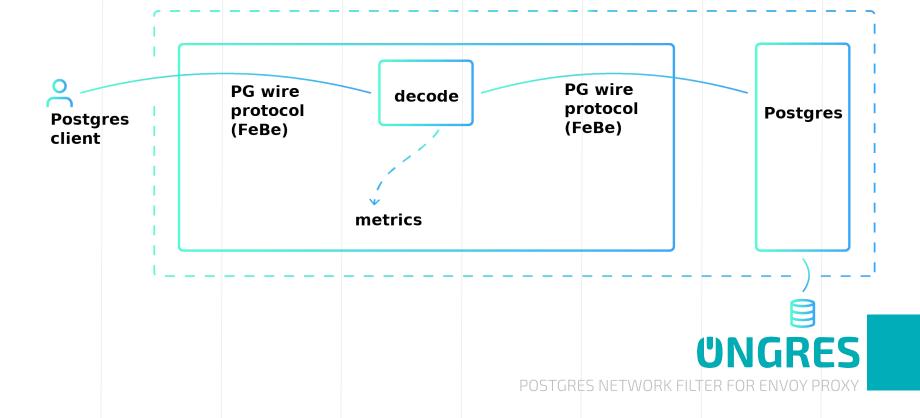




# FeBe protocol architecture



### Idea: proxy and decode the protocol to get metrics!



#### Advantages of metrics via proxying/decoding

- s/pull/push/
- **Zero impact on the database**. 100% transparent.
  - No performance impact.
  - No configuration required. No agents/tools to install.
- Can be deployed as a sidecar (eg. via injection in K8s).
- May significantly increase the volume of metrics obtained.
- Opens the door for other added functionality.



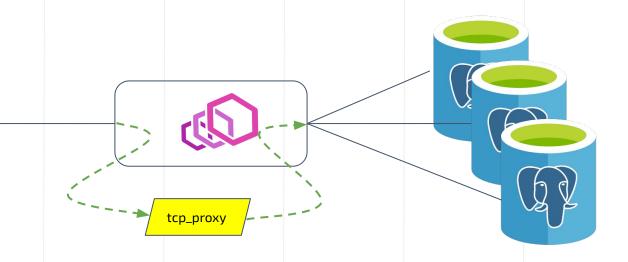
# ENVOY PROXY POSTGRES FILTER



# Envoy extensibility: simple connection model



Application (client)

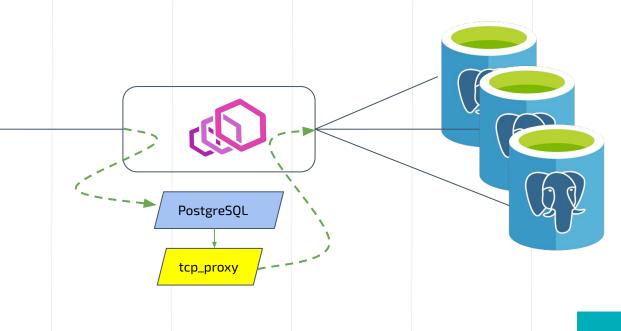




## Envoy extensibility: filter chain



Application (client)

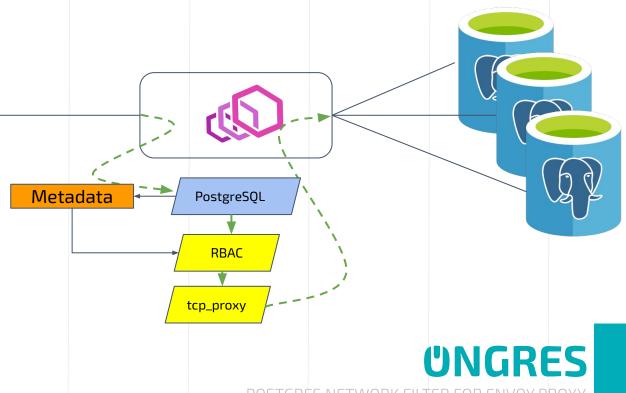




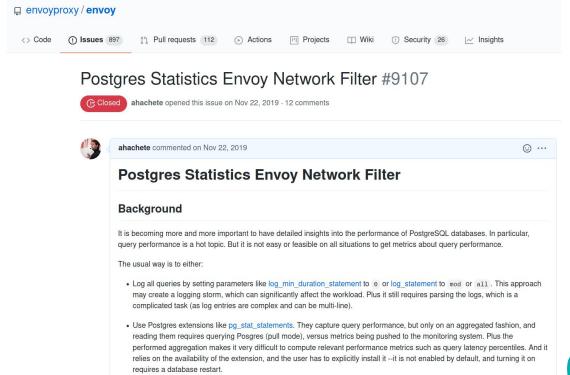
## PostgreSQL Filter Architecture: metadata



Application (client)



#### How it all started





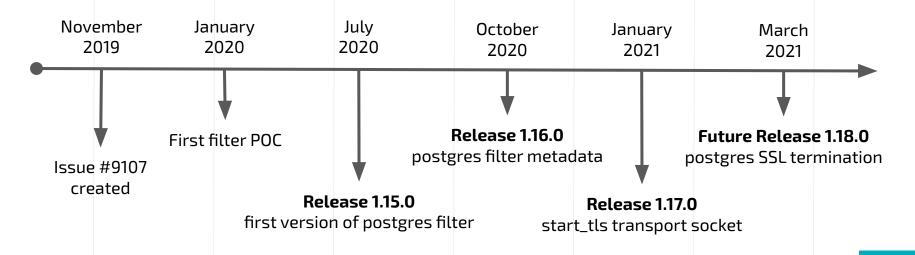
#### An effort developed by a great Community

- Contributions from Tetrate, Envoy Maintainers, OnGres and others.
- Merged and released with Envoy 1.15.
- Led to 10 issues and new functionality being implemented in several other areas.
- New features to come in new releases!

Help wanted!



### **Postgres filter Timeline**



## Metrics currently being exported

errors (error, fatal, panic,)	messages (frontend, backend)
sessions (encrypted, unencrypted)	statements (insert, update, delete,)
transactions (commit, rollback)	notices (notice, log, warning,)

Counters (metric / second)





# UNGRES

### **Expected result**



https://github.com/ongres/envoy-postgres-stats-example



# COMING TO ENVOY 1.18\*: POSTGRES SSL OFFLOADING AND MONITORING



#### **Postgres SSL**

- Doesn't operate at the TCP level 4, but rather application level (7).
- Initial connection is unencrypted, then a request to "upgrade the connection to SSL" is performed. Similar to STARTTLS in SMTPS.
- Database connection costs are high. SSL database connection costs are very high.
- Use a connection pooler! Like PgBouncer!
  - PgBouncer -> single threaded -> swamped under heavy SSL connection load.
- Turning on/off SSL or rotating certificates requires database restart -> downtime.



#### Offload Postgres SSL to Envoy!

- Avoids both Postgres and PgBouncer SSL performance problems.
- Allows monitoring of encrypted traffic!
- Turn on/off, rotate certificates without database impact.
- Programmatic management: use Envoy xDS APIs to manage certificates.

StartTls infrastructure already released on 1.17. Postgres specific filter implementation coming on 1.18\*.



#### **Architecture** Prometheus **Envoy Proxy** SSL (pg Wire) pg Wire pg Wire pg Wire **StartTIs** TCP Postgres **Postgres** Filter Extension Filter CRT **KEY** UNGRES

POSTGRES NETWORK FILTER FOR ENV

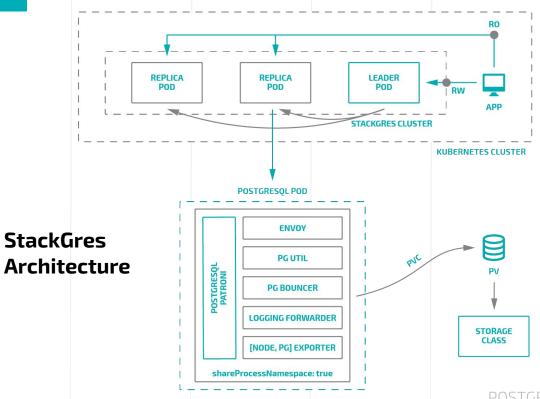
**SSL DEMO** 

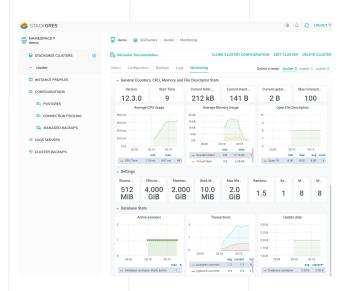
UNGRES

# USE CASES THE FUTURE OF ENVOY'S POSTGRES PLUGIN



#### Envoy's filter use-case: StackGres.io





UNGRES

#### **Future plans**

- Better SQL parsing
- Producing per-database statistics
- Routing based on Query type
- Traffic drain
- Opent Telemetry Integration

#### **Community:**

- Envoy slack: envoyproxy.slack.com
- PostgreSQL specific channel: #envoy-postgres
- PostgreSQL related issues:

https://github.com/envoyproxy/envoy/labels/area%2Fpostgres



#### References

#### Original Github Issue:

https://github.com/envoyproxy/envoy/issues/9107

First post about Envoy Postgres Filter:

https://www.cncf.io/blog/2020/08/13/envoy-1-15-introduces-a-new-postgres-extension-with-monitoring-support/

#### **Envoy Documentation:**

https://www.envoyproxy.io/docs/envoy/latest/intro/arch\_overview/o ther\_protocols/postgres

https://www.envoyproxy.io/docs/envoy/latest/configuration/listeners/network\_filters/postgres\_proxy\_filter

UNGRES

**Questions?** 

UNGRES