

# WHY ARE DBAS SO GRUMPY ALL THE TIME?

NORDIC PGDAY  
2018-03-12

Lætitia AVROT  
Loxodata

# WHO



Lætitia Avrot

---

- PostgreSQL consultant and teacher
- DBA PostgreSQL for more than 10 years (and also Oracle and SQL Server)
- @l\_avrot

# LOXODATA


Company built on 3 essential pillars



PostgreSQL



DevOps



Cloud



Why are DBAs so grumpy all the time ?

---

- Messing up from the beginning
- Worst developer's practices
- Ops can mess up too!



# THE DESIGN





## Remove potential interferences

---

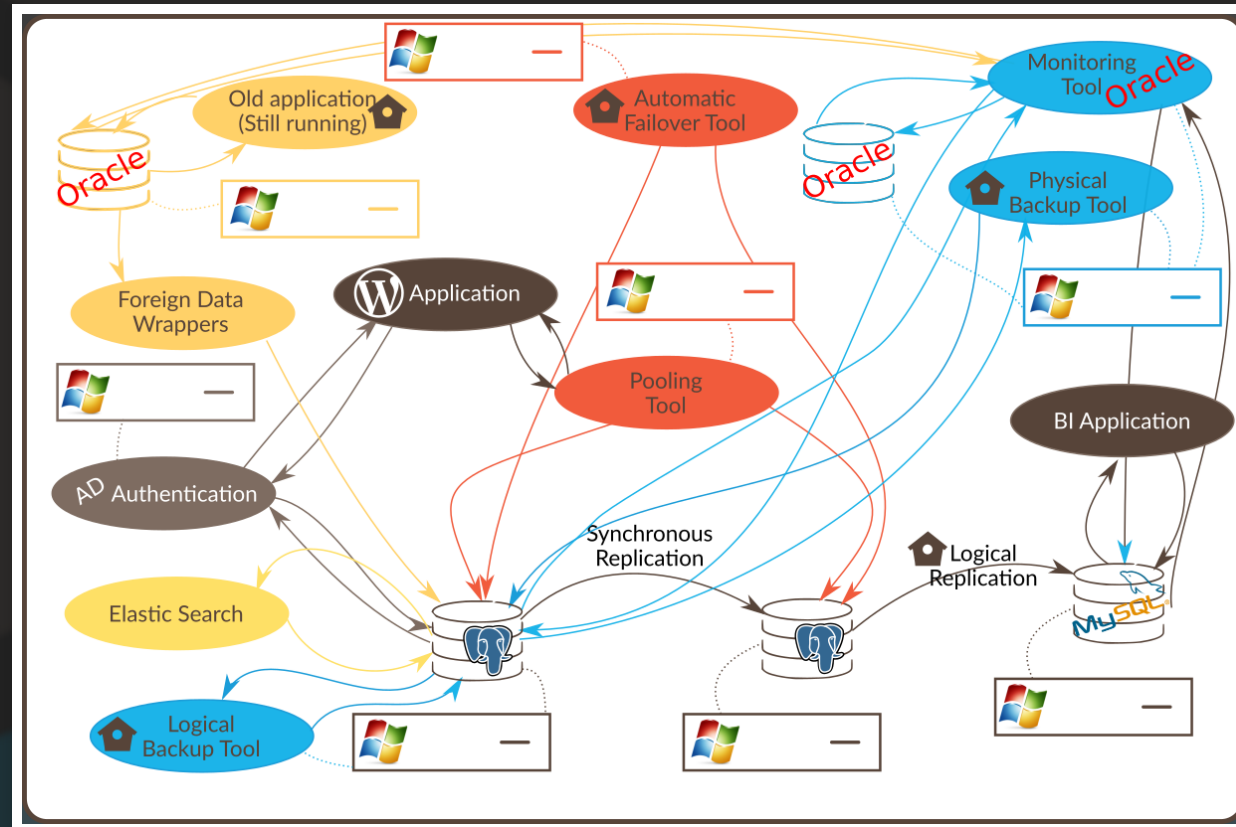
- Don't tell people about your project
- Don't tell "ops" about your project
- Keep your DBA busy (with production)





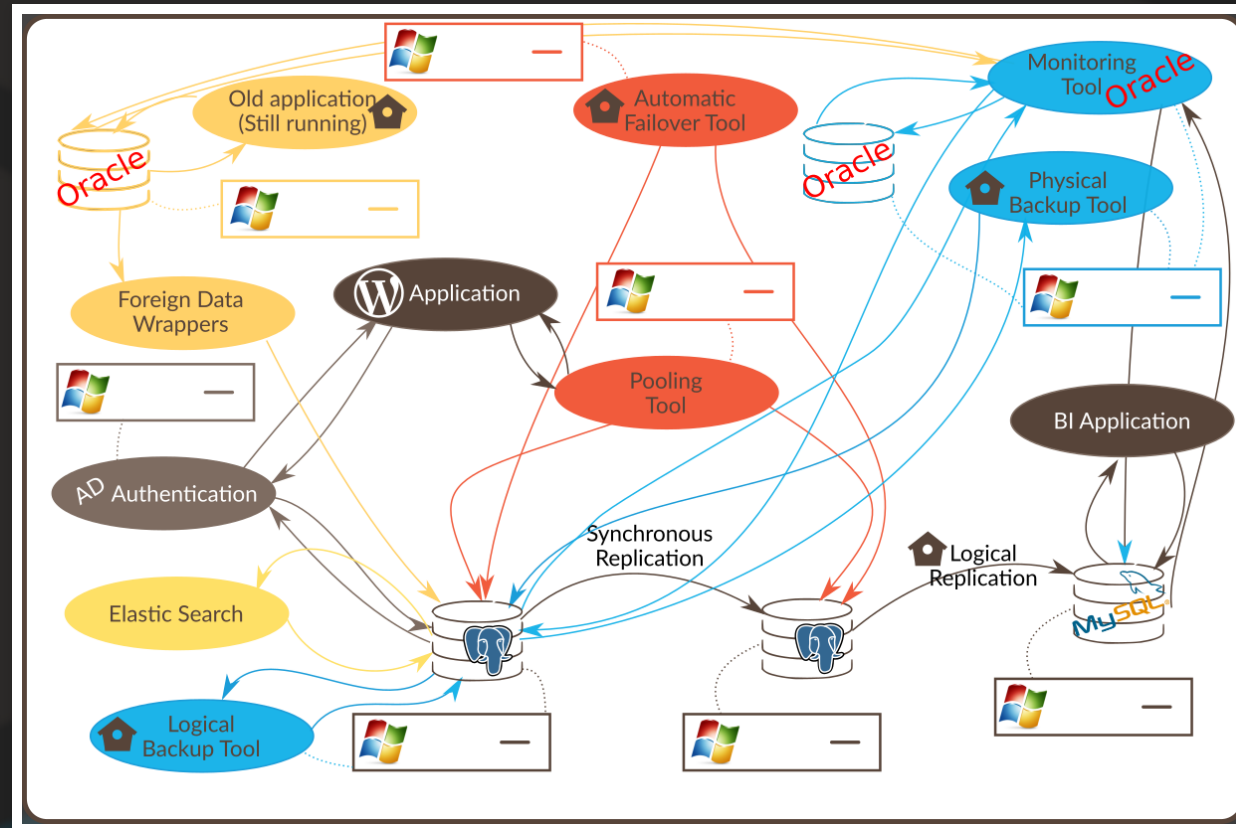
The more complex, the more fun!

- Replication (multimaster, synchronous, logical...)
- Authentication (AD, LDAP)
- External data sources (Datalinks, FDW, ETL...)



## Complexity by example

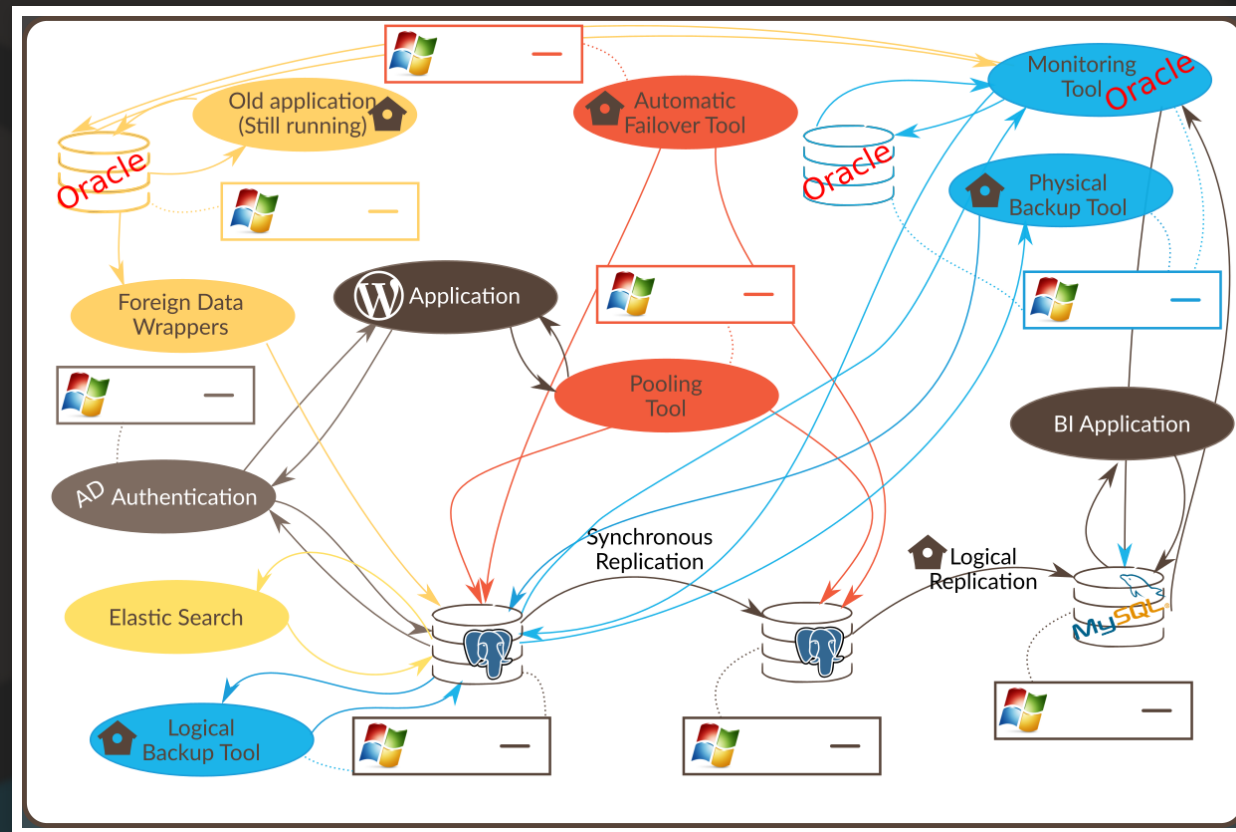
- A simple website
- Shows the company canteen menu each day



## Complexity by example


- Adding external authentication
- Adding synchronous replication
- Adding "0 administration architecture"
- Adding Backup and Monitoring tools





## Complexity by example

- Adding Elastic search for full text search
- Adding data from old application
- Adding logical replication for BI App



## Cost reduction

---

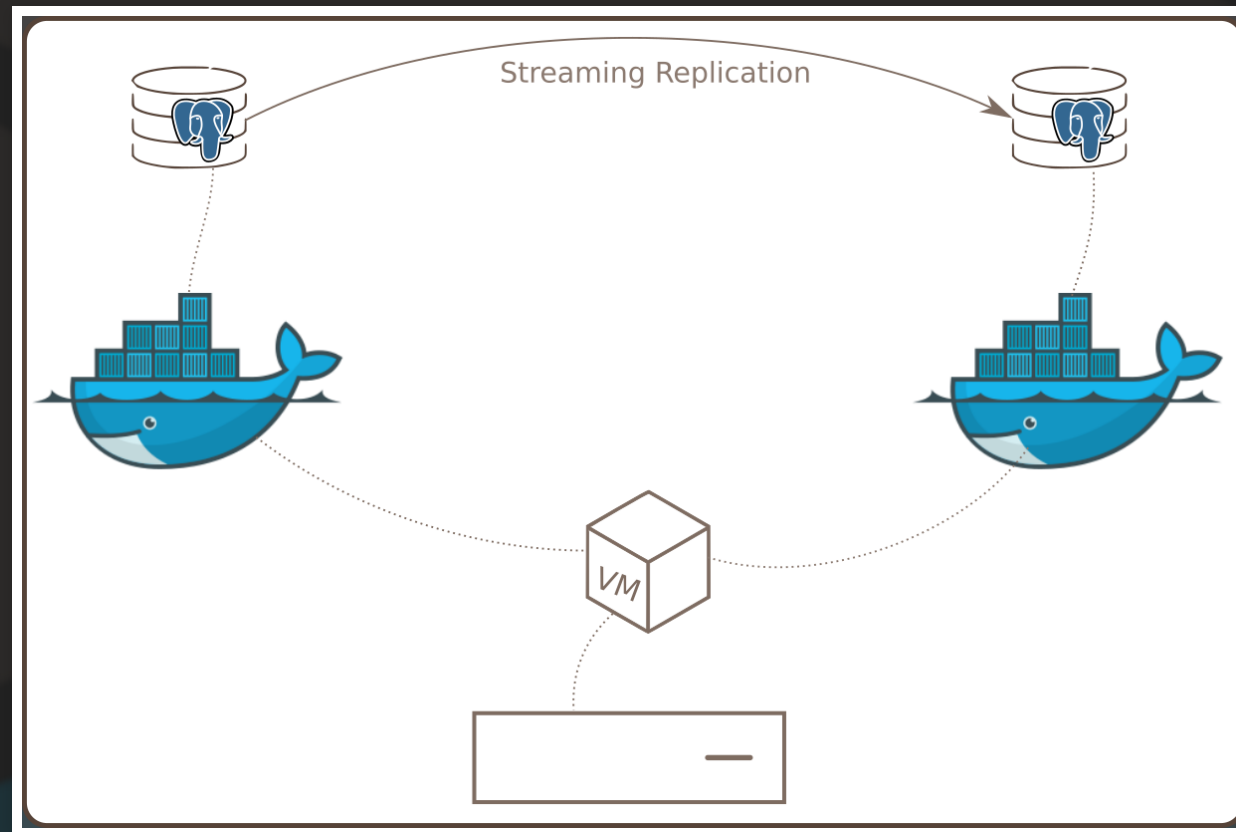
- Mutualization
- 50+ instances running on the same server
- One database per instance



## Docker is the new black

- Does everything!
- Composing ? What for ?
- Backing up ? What for ?





## "Bad" docker by example

- 2 dockers on same VM
- 1 PostgreSQL instance per container
- With streaming replication : call it high availability!

# APPLICATION DEVELOPMENT







## ORMs

---

- Let your ORM choose which model you need
- Activate indexing



## Model

---

- 1st Normal Form
- Drop constraints, your app can handle it!
- Don't normalize : call it denormalization!






Be trendy : try schemaless!

---

- Easier, faster, lighter!
- Mongo can do it!
- Store a lot : call it Big Data!
- You can fake a model too: Entity, Attribute, Value



## Time series

---

- Store current and archive data on same table
- No partitioning, of course!
- No purge ever : you need all your data!





## Joins

---

- Grab list of ids from first table
- Loop on it and query on the second table
- If it's slow, blame the database!





## Transactions, commits, locks and so on

---

- Open transactions and leave it open
- Ask for explicit commit and never commit
- Postgres can't manage correct locks : use explicit locks!

OPS





## Backups

---

- Don't need (replication, disk mirroring, VM snapshots)
- Only logical
- Home made!
- No testing required





## Monitoring

---

- Graphs are for weaklings!
- Provide graphs without unit or legend
- Take one measure per day, it's quite enough!



## "Cleaning" \$PGDATA

---

- pg\_xlog/pg\_wal
- pg\_clog
- use pg\_resetwal





RDBMS are magic!

---


- No DBA needed
- Nothing to understand: it's just magic!
- No maintenance needed!



## Autovacuum

---

- Setting it off
- `autovacuum*_scale_factor=0`
- `autovacuum*_threshold=100000`



# Hide behind ITIL!

- Strip the DBA off any permission
- One change per month, no more
- 6-months validation process for a minor update



# CONCLUSION



**KEEP  
CALM  
AND  
LISTEN TO THE  
CONCLUSION**



If you're a DBA...

---

- Improve your communication
- Smile
- Appease others



If you're not a DBA...

---

- Improve your communication
- Smile
- Don't let fear take control of you!






All we need is love!

---

- Sysadmin appreciation day : last friday of July
- Day of the Programmer : 256th day of the year
- DBA appreciation day : first friday of July



Questions ?

---

Thanks to Paul S Randall for the inspiration

Thanks to Freepik and Pixabay for illustrations and photos

Don't forget to leave feedback  
<https://2018.nordicpgday.org/feedback>