

Migrate to PostgreSQL

Reasons and Consequences

Jean-Paul Argudo

Dalibo - www.dalibo.com

PgConf.Russia 2016 - February, 4th

Thanks !

Great people makes great events

 **DALIBO**
L'expertise PostgreSQL



- ▶ Oleg Bartunov for the invitation
- ▶ Larisa Somova and Postgres Pro
- ▶ You, for being there in a non-tech talk

Jean-Paul Argudo

PostgreSQL geek

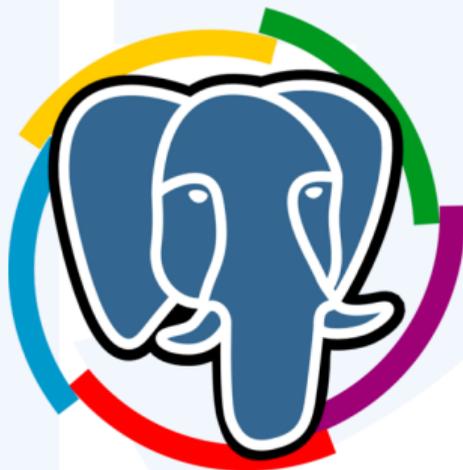
 **DALIBO**
L'expertise PostgreSQL

- ▶ PostgreSQL.fr website founder (2004 →)
- ▶ PostgreSQL Fr non-profit co-founder (2005 →)
- ▶ PostgreSQL Eu non-profit co-founder (2008 → 2013)
- ▶ Founder and CEO of Dalibo (2005 →)

PostgreSQLFr

the non-profit for french-speaking countries

- ▶ regroup users and companies
- ▶ booths on FLOSS events like ([Paris Open Source Summit](#))
- ▶ organize [pgDay.fr](#)
- ▶ sponsors [PUG Paris](#)
- ▶ legal matters (insurance, bank and other administrative things)



Francophony wheel

- ▶ PostgreSQL Expertise, Trainings and Support since 2005
- ▶ PostgreSQL "pure-player"
- ▶ Some tools: [OPM](#), [PoWA](#), [pgBadger](#) and more
- ▶ 20% of working time contributing to PostgreSQL
- ▶ Financial supporter of the PG community ([pgConf.eu](#), [pgDay.fr](#), etc)

Agenda

Topics covered by my talk

- ▶ Part 1: Reasons to migrate
- ▶ Part 2: How to migrate
- ▶ Part 3: Consequences
- ▶ Part 4: Usecases

PART 1

Reasons to migrate

Freedom... at all costs!

Fundamentals

PostgreSQL's pillars

- ▶ Stability
- ▶ Extensibility
- ▶ Security
- ▶ Respect of standards

State-of-the-art

since 2010

- ▶ built-in replication
- ▶ exclusion constraints, partitionning
- ▶ pg_upgrade
- ▶ extensive PL/pgSQL

Innovation

Searching perfection

- ▶ top-notch SQL (Window functions, CTEs, LATERAL, etc)
- ▶ Bi-Direction Replication (BDR)
- ▶ Geo Data

More Innovation

Fast pace of development

- ▶ Parallelism / External partitioning / Scale out ([pg_shard](#), [PL/Proxy](#), Foreign data wrappers, etc...)
- ▶ [Hypothetical Indexes support](#) inside [PostgreSQL Workload Analyzer](#)(POWA)
- ▶ more on [Dalibo's labs](#) (and more to come)

!!!Hot news!!! **PostgreSQL Automatic Failover, HA for Postgres, based on the industry references Pacemaker and Corosync** by Dalibo's labs, the release announcement should pop up really, really soon

PART 2

Conduct of change

How do I migrate ?

5 essential steps

(how to kick Oracle)

- ▶ Study : inventory, TCO and ROI
- ▶ Training : **convert** skills
- ▶ Industrialize : procedures, manuals and tools
- ▶ Migrate : schema, data and code
- ▶ Support : define your SLAs and choose a partner

(1) Studies

inventory

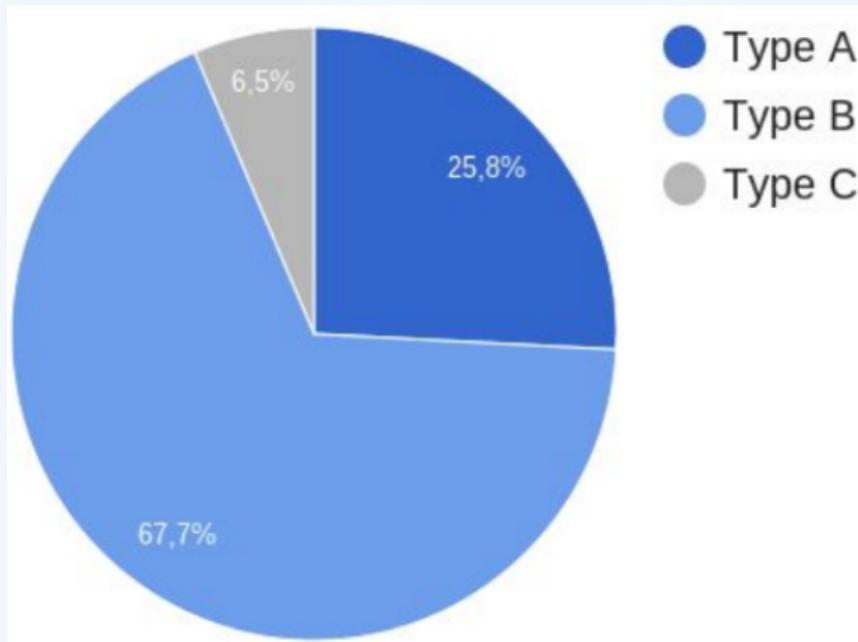
3 kinds of Oracle databases

- ▶ A : migrate in a few hours (no code)
- ▶ B : migrate in up to 10 man days (few code)
- ▶ C : migrate in more than 10 man days (code++)

Ora2Pg to estimate migration costs...

(1) Studies

ora2pg inventory sample



**Believe me or not,
type A+B is
usually 90% !**

Typical results are
like

- ▶ Type A : 25% → 40%
- ▶ Type B : 25% → 50%
- ▶ Type C : 0% → 20%

(1) Studies

TCO and ROI

- ▶ compare actual Oracle costs (TCO) and PostgreSQL one
- ▶ bear in mind investments (ROI?):
 - ▶ one shot costs: studies, trainings, industrialization and migration
 - ▶ recurring costs: trainings, support

(2) Trainings

This is the key

- ▶ discuss, propose POC and benchmark to convince
- ▶ **convert** skills : things are quite the same
- ▶ trainings once every 2 years or less, DBAs and DEVs

(3) Industrialize

Build solid things

- ▶ write a consistent set of documents to manage PostgreSQL in the company
- ▶ select a set of approved tools to cover daily needs:
 - ▶ administration
 - ▶ backup and recovery
 - ▶ tuning
 - ▶ monitoring
 - ▶ fleet management

(4) Migrate

- ▶ Type A : test and validation, easy ROI
- ▶ Type B : let the PostgreSQL expert guide your team
- ▶ Type C : PL/SQL rewriting to PL/pgSQL

Compare application's typology A, B, C with database's one...

(5) Support

- ▶ define your support perimeter
- ▶ best support unit so far: PostgreSQL instance (aka server)
- ▶ identify a PostgreSQL Pure Player
- ▶ you **need** a support contract

PART 3

Consequences

Okay I'm going full PostgreSQL now. What will be the consequences?

5 direct changes

Immediate wins

- ▶ no more licences to pay
- ▶ no more vendor lock-in
- ▶ no more tie-in sales
- ▶ no more weird licensing on virtualization
- ▶ no more intrusive audits

5 biggest threats

Immediate problems to solve

- ▶ lost in space (of Open Source)
- ▶ PostgreSQL is going so fast
- ▶ it's hard to find PostgreSQL DBAs
- ▶ almost 300 tools to do 5 things
- ▶ ROI and TCO with Open Source?

PostgreSQL TCO

Things to know

- ▶ beware hidden and indirect costs
 - ▶ each PostgreSQL server create costs
- ▶ costs to integrate
 - ▶ trainings: 0 → 3 k€/year/person
 - ▶ maintenance: 0 → 5 k€/year/person
 - ▶ support: 4 → 15 k€/year/server
- ▶ **you choose** the TCO

ROI on migration

with ora2pg's typology

- ▶ Type A : Automatic migration !
 - ▶ Immediate ROI !
- ▶ Type B : Migrate is <10 man/days
 - ▶ ROI < 1 year
- ▶ Type C : Complex cases (lots of PL/SQL code)
 - ▶ migration duration very variable
 - ▶ $1 < \text{ROI} < 3$ years

Some changes

PostgreSQL brings new methods

- ▶ be prepared to test and validate often
- ▶ 1 major version per year
- ▶ automatic testings
- ▶ include developers in changes
 - ▶ have last version in dev (9.5.0 / jan 2016)
 - ▶ have previous version in prod (from 9.4.0 / dec 2014 to 9.4.5 / oct 2015)
 - ▶ [Zalando's great example](#)

Stay in touch

Learn to find the information

- ▶ go to PostgreSQL events
- ▶ meet the people there
- ▶ depending the time you have
 - ▶ [Postgres Weekly](#)
 - ▶ [PostgreSQL Weekly News](#) by David Fetter
 - ▶ [planet.postgresql.org](#) blogs aggregator
 - ▶ and more than 80 [Mailing Lists](#) !
 - ▶ [pgsql-ru-general](#) since 2004

Find DBAs

How to have PostgreSQL skills in house

- ▶ You find one?! OMG lucky you. Just hire him/her!
- ▶ Otherwise
 - ▶ (as said before:) **TRAIN YOUR PEOPLE**
 - ▶ PostgreSQL companies will deliver great trainings
 - ▶ converting a DBA : 5 to 10 days of training
 - ▶ ...then everything is on the internet

- ▶ a standard DBA wants to manage PostgreSQL servers
- ▶ but also monitoring, statistics, backup tools and so on
- ▶ → need for a complete stack of tools
- ▶ → extremely active field in the last 2 years!

A field where an expert will help you

Support is obvious

How to identify a PostgreSQL "pure player" ?

- ▶ you **need** a support contract
- ▶ with a comprehensive and complete offer
- ▶ transforms a wide part of the income into contributing back to the community
- ▶ as a bonus, the governance is readable and assumed

HINT : sign contrat before they grow too much ;-)



PART 4

Usecases

Some stories on migrations

- ▶ French government example
 - ▶ sept. 2012: [Circulaire Ayrault](#) (prime minister) on Open Source
 - ▶ may 2013: [SILL 2013](#) (working station)
 - ▶ feb. 2014: [SILL 2014](#) (servers)

(SILL: the inter-ministerial free software base)



One result

First document

The very first document written by the Open Source group in the FR government is a "Guide to PostgreSQL Migration"

- ▶ [French version](#), 26th june 2015
- ▶ [Russian version](#) (courtesy of Postgres Pro)

- ▶ Migrations started **back in 2001** (french)
- ▶ 1st big turn in 2008: 70% PostgreSQL / still 30% Oracle
- ▶ biggest databases migrated in 2014
- ▶ 100% PostgreSQL in 2016
- ▶ **Migrations in big organisations takes many years**



Le Bon Coin

Top french ads website



leboncoin.fr

- ▶ 7th most visited in France, full PostgreSQL
 - ▶ 250M page views / day
 - ▶ 5M unique visitors / day
 - ▶ 18M UV / month (1/3 french people)
 - ▶ 600K+ new ads / day
 - ▶ 25M live ads
- ▶ 100+ PostgreSQL servers on 20 machines
- ▶ 8 Terabytes of data
- ▶ **read complete usecase:** [english](#)

More ?

If you need more references

- ▶ listing all usecases would take days
- ▶ you have about 100 references on [Dalibo's customer page](#)
- ▶ sorted in 19 categories: telecom, finances, space, medias, retail, public services, etc...

Conclusion

It's time to migrate, completely

- ▶ PostgreSQL **is great** and greater each year
- ▶ **Pure players** will help you
- ▶ to avoid problems you'll face if you go **alone**
- ▶ the best cards you have in hands are **your colleagues**
 - ▶ challenge your DBAs and DEVs
 - ▶ they'll add PostgreSQL in their skills and resume
 - ▶ they'll be tomorrow's highest salaries

Questions?

Jean-Paul Argudo <jean-paul.argudo@dalibo.com>

jpa @ irc.freenode.net

@jpargudo on twitter