

High Concurrency with Postgres

Bank of Brazil in real life

Fábio Telles Rodriguez

Timbira - The Brazilian Postgres Company

February 4, 2016

PGConf Russia

timbira

Fábio Telles Rodrigues

- ▶ Open software evangelist +15 years
- ▶ PostgreSQL DBA +10 year
- ▶ Contributor of the brazilian PostgreSQL community
- ▶ Blog: <http://savepoint.blog.br> (Portuguese only)
- ▶ @telles

- ▶ <http://www.timbira.com.br>
- ▶ The Brazilian Postgres Company
- ▶ Consulting / Development / Support
- ▶ Partnership with software development companies
- ▶ In-Company and On-Line training

Outline

- ▶ About Bank of Brasil
- ▶ Multi Document Processing
- ▶ Challenges
- ▶ Adopted solutions

- ▶ Founded at 1808, oldest bank in Brazil
- ▶ 59% of shares belong to the government
- ▶ The only bank that has branches in all Brazilian cities
- ▶ more then 110.000 direct employees

Multi Document Processing app

- ▶ Developed by Bull (now Atos)
- ▶ Development started as a client/server application 15 years ago
- ▶ Mostly written in C++
- ▶ Perform a decentralized dematerialization of documents combined with a complex central processing of them
- ▶ Each document have differentdifferent workflows

Multi Document Processing app

- ▶ Documents are captured using scanners at branches
- ▶ Local image recognition (reach about 85% of success)
- ▶ Shipping images and meta data to a central point (images are not stored inside database)
- ▶ Central image recognition using third-party tools like ABBYY and A2iA.
- ▶ Manual image recognition
- ▶ Other manual interactions like authorizations and deviations
- ▶ Thousands of business rules for each document
- ▶ Complex interaction with many legacy systems

- ▶ 6,364 branches
- ▶ 22,432 stations using application (will double by the end of the year)
- ▶ 25 servers making central recognition
- ▶ 12 application servers
- ▶ 200 stations making manual recognition
- ▶ Hundreds of leaders requesting reports on peak processing
- ▶ One PostgreSQL database

The database

- ▶ PostgreSQL 9.3
- ▶ 228 tables
- ▶ 23 triggers
- ▶ 17 sequences
- ▶ 402 functions

In a regular day

- ▶ check clearing: 600,000 (2 milion in the busiest day)
- ▶ escrow checks: 70,000
- ▶ signature cards: 30,000
- ▶ non financial documents: 50,000
- ▶ Critical window between 4pm to 7pm
- ▶ Growth of 10GB
- ▶ 80GB of archives generated

Challenges

- ▶ High number of connections
- ▶ Locks in processing queues
- ▶ High number of transactions
- ▶ Small processing window
- ▶ Many heavy queries for reports
- ▶ Need to keep the information for two years in the database
- ▶ New features being implemented constantly

Adopted solutions

- ▶ 3 PgBouncer instances
- ▶ Memcached + Listen / Notify to spread information across the stations
- ▶ Strict control of transactions
- ▶ The queue for image recognition was implemented in memory and integrated inside database with PL/Peru using sockets
- ▶ Use of advisory locks in other queues
- ▶ Memory adjustments for specific users
- ▶ Vacuum and fillfactor adjustments for specific tables
- ▶ Partition on 24 tables
- ▶ Use of temporary tables and unlogged tables
- ▶ Redesign critical process



Spasibo
Thank you
Obrigado
Fábio Telles Rodriguez
(telles@timbira.com.br)
<http://www.timbira.com.br>

timbira