

# The PostgreSQL Open Source Development Process

BRUCE MOMJIAN



This explores the Postgres open source development process and developer motivations.

*Creative Commons Attribution License*


*<http://momjian.us/presentations>*

*Last updated: November, 2018*

# Introduction

- ▶ How is open source development different from closed source?
- ▶ What is the pattern of PostgreSQL adoption?
- ▶ What does the future hold for PostgreSQL?

# Open Source World



# Revolutionary Change in Software Development: Challenging Totalitarianism

- ▶ Central Control
- ▶ Few Choices
- ▶ Limited feedback from population

# Bringing Democracy to Software Development

- ▶ Distributed Control
- ▶ Many Choices
- ▶ Driven by feedback from population

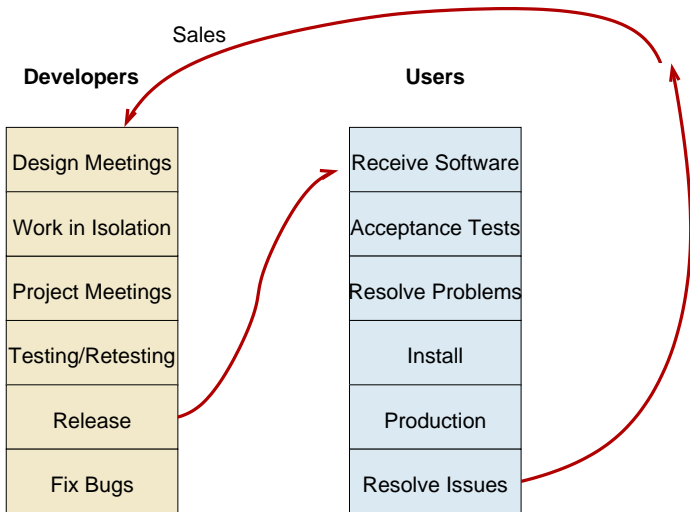
# Proprietary Software Development

- ▶ Put everyone in a room
- ▶ Discuss new features and how to implement them
- ▶ Everyone goes back to their office
- ▶ More meetings
- ▶ Testing
- ▶ Release
- ▶ Gather user bug reports
- ▶ Repeat

# Open Source Development

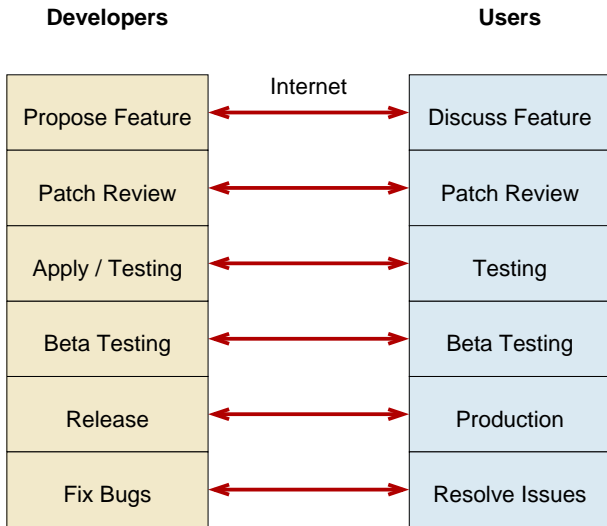
- ▶ Discuss via email, with access to a world-wide pool of talent
- ▶ New features discussed over several days by developers and users
- ▶ Coding
- ▶ Patch review
- ▶ Patch application and testing
- ▶ User testing
- ▶ Release

# Proprietary Software Development





# Open Source Development



# Typical Proprietary Database Product Life Cycle

- ▶ Create product
- ▶ Grow customer base
- ▶ Reduce development budget
- ▶ Decrease technical support costs
- ▶ Increase license fees
- ▶ Maximize revenue
- ▶ Customers leave
- ▶ End-of-life product

All possible because of high switching costs.

# Typical Open Source Life Cycle

- ▶ Create project
- ▶ Developers create initial software
- ▶ Users and developers improve the software
- ▶ No development budget
- ▶ No technical support costs
- ▶ No license fees
- ▶ No maximizing of revenue
- ▶ No customers leaving!
- ▶ No end-of-life

# PostgreSQL Is Completely Open Development

- ▶ No controlling company or individual
- ▶ Active multi-company and volunteer community
- ▶ Open development process
- ▶ Leadership encourages new members

### Users

General problem  
Other error xmin from before reifrozenxid  
Announce check\_pgactivity 2.4 released

### Developers

Hackers interval type additional option  
Commit Add --min-xid-age and --min-txid-age options to vacuumdb  
Versions Stable: 11.1+, 10.6+, 9.6.11+, 9.5.15+, 9.4.20+ | Development: 12 devel, in commitfest

### External

Blogs Keith Fiske: Managing Transaction ID Exhaustion (Wraparound) in PostgreSQL  
News PGConF APAC 2019 Schedule Announced  
Tweets It's #TalkTuesday! Rhys Eismore (@rhyselismore) says to "Just Use @PostgreSQL" at ...  
Media Microsoft Acquires Citus Data to Dive Deeper into Postgres - eWeek  
Events FOSSDEM PGDay 2019

### IRC (also Slack)

**pg\_dochot:** <https://explain.depesz.com/> :: <https://pasteboard.co/>  
**pg\_dochot:** <https://www.db-fiddle.com/> :: <https://paste.depesz.com/>  
**pg\_dochot:** <https://dpaste.de>

**szymon:** Hi, I have a question about FTS. A small example: I have "Lockheed Martin Corporation" in a column, when I search for 'lockhe' it's found, bug when I use 'lockeed' then it's not. This explains it a little bit <https://paste.depesz.com/s/3Ss>. I'm wondering how to fix it.

**reliability:** Am I correct in assuming that the `setspeed()` function only affects the current session?

**DarkUranium:** `peers()` (w.r.t. tables) yeah, that's how most of my tables have been.

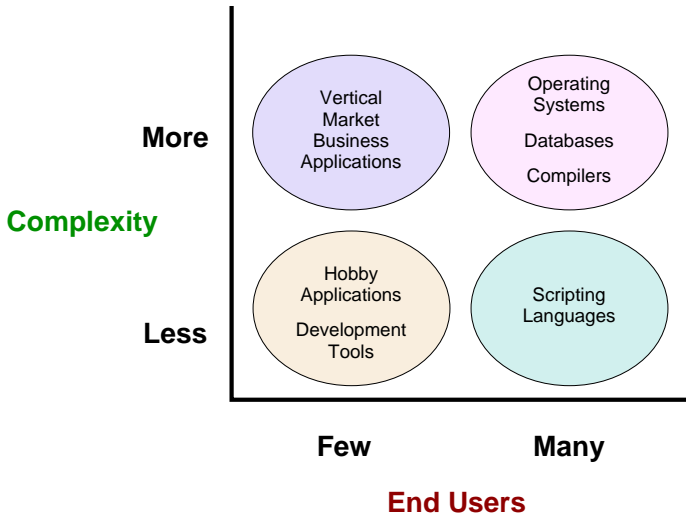
**DarkUranium:** I do typically need to fetch the primary key along with the rest of the data \*anyhow\* (either that, or join into somewhere where I can fetch it)

**DarkUranium:** Well, fetch or lookup by it.

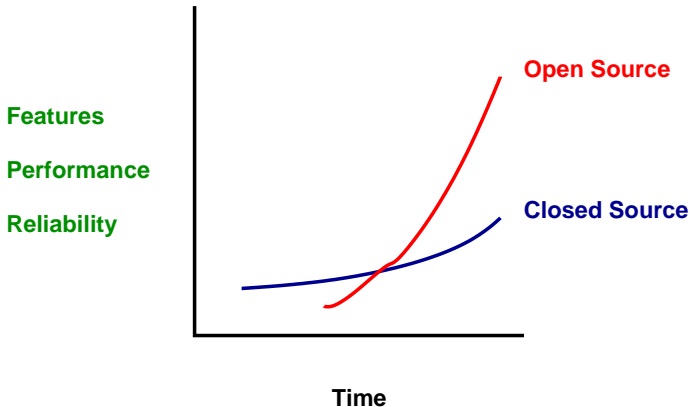
London 13:25 Berlin 14:25 Moscow 16:25 Mumbai 18:55 Beijing 21:25 Tokyo 22:25 Los Angeles 05:25 New York 08:25 São Paulo 11:25  
Content updates automatically | [About](#) | [Submit Feedback](#)

<http://pglife.momjian.us>

# Open Source Software Availability



# Is It As Good As Proprietary Software?



# Problems With Open Source (and Democracy)

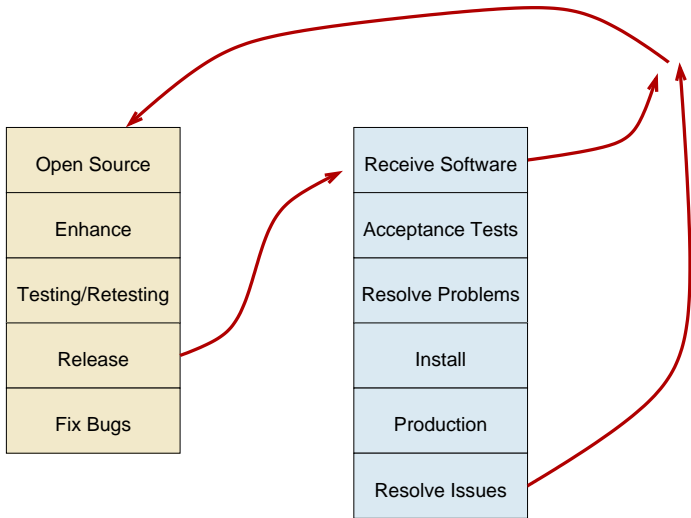
- ▶ Requires effort to stay informed
- ▶ Requires effort to receive assistance
- ▶ Database changes more rapidly
- ▶ More self-help



# Businesses Supporting Open Source Users

- ▶ Customization
- ▶ Technical support
- ▶ Packaged releases
- ▶ Add-on software
- ▶ Training
- ▶ Localized


# Hybrid Development



# Licensing

- ▶ BSD license, limitation of liability
- ▶ Not GPL, which requires unrestricted source distribution

# PostgreSQL Evolution



Flexibility includes:

- ▶ Application-specific data types, e.g., JSON, range types
- ▶ Advanced index types, e.g., GIN, SP-GiST
- ▶ Single and multi-node scalability

# Our Developers

Why is programming fun?

- ▶ Joy of making things
- ▶ Joy of making useful things
- ▶ Fun of a complex puzzle
- ▶ Joy of learning
- ▶ Tractable medium

Concepts derived from *The Mythical Man-Month* by Frederick P. Brooks, Jr.

# Open Source Motivations

- ▶ Professional advancement, learn new skills
- ▶ Practical need for the software
- ▶ Mental stimulation
- ▶ Belief in open source

Concepts derived from a survey conducted by OSDN and reported in “Open source a needed outlet for programming pros”, *The Register*, May 2, 2002, <http://www.theregister.co.uk/content/4/23935.html>.

# PostgreSQL Motivations

- ▶ Professional advancement, learn new skills
- ▶ Practical need for the software
  - ▶ *Business needs the enhancement*
  - ▶ *Business wants PostgreSQL to thrive*
- ▶ Mental Stimulation
- ▶ Belief in open source

## Extreme Case

*It was cheaper  
to pay to add missing features  
to PostgreSQL  
than to buy Oracle.*



# Questions



*<http://momjian.us/presentations>*

*<https://www.flickr.com/photos/andyarthur/>*