

# The PostgreSQL Open Source Development Process

BRUCE MOMJIAN



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

<https://momjian.us/presentations>




*Creative Commons Attribution License*

*Last updated: September, 2020*

# 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

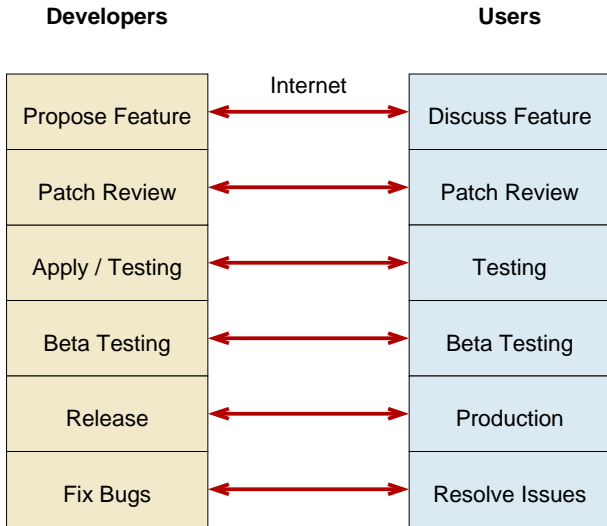
- Discuss via email, with access to a world-wide pool of talent
- New features discussed over several days by developers and users
- Coding
- Patch review
- Patch application and testing
- User testing
- Release

# 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

# PgLife

## PgLife Postgres Community Life

### Users

**General** Re: Obvious data mismatch in View2 which basically SELECT \* from View1  
**Other** Re: Policy on cross-posting to multiple lists  
**Announce** PostgreSQL 13 RC 1 Released!

### Developers

**Hackers** Re: Planner, check if can use consider HASH for groupings (src/backend/optimizer/plan/planner.c)  
**Commit** Re: Add support for building GiST index by sorting.  
**Versions** **Stable:** 12.4+, 11.9+, 10.14+, 9.6.19+, 9.5.23+ | **Development:** 13 rc1+, 14 devel, in commitfest

### External

**Blogs** Marco Slot: Talking about Citus & Postgres at any scale  
**News** PostgreSQL 13 RC 1 Released!  
**Tweets** RT by @PostgreSQL: The @PostgreSQL yum repository repository now includes ...  
**Media** Crunchy Data Announces Fully Managed Cloud Postgres Database Service - PRNewswire  
**Events** PGDay Austria 2020 - CANCELLED

### IRC (also Slack)

**Myon:** arrays are only good if you never want to search for a tag (which is unlikely to be the case)

**Myon:** make a table (noteid, tag)

**xocolatl:** ??tags

**pg\_dochot:** <http://www.databasesoup.com/2015/01/tag-all-things.html>

**dysfigured:** fwiw it's a personal note-taking system, and I doubt I'll ever make it a multi-user system, so it's unlikely there will be more than a few thousand entries

**xocolatl:** in that case, don't worry about anything

**dysfigured:** just trying to come up with the table structure for now, and debating whether it makes more sense if links between notes would be directed or undirected

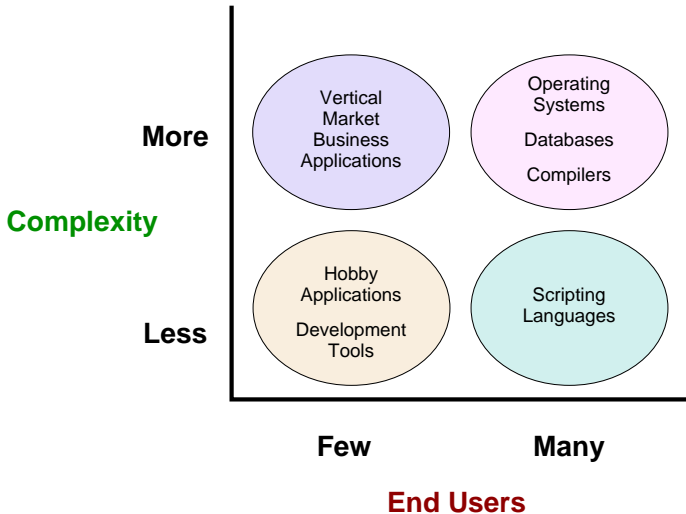
**RhodiumToad:** Myon: arrays are very good for searching on tags

London 21:15 Berlin 22:15 Moscow 23:15 Mumbai 01:45 Beijing 04:15 Tokyo 05:15 Los Angeles 13:15 New York 16:15 São Paulo 17:15


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>. For a developer survey, see <https://www.techrepublic.com/article/what-makes-developers-happy-contributing-to-open-source/>.

# 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



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

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