

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: February, 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

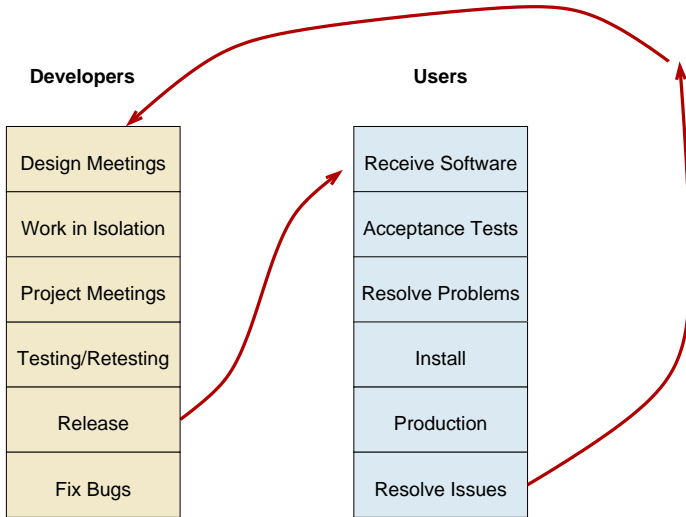
Proprietary Software Development

- ▶ 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

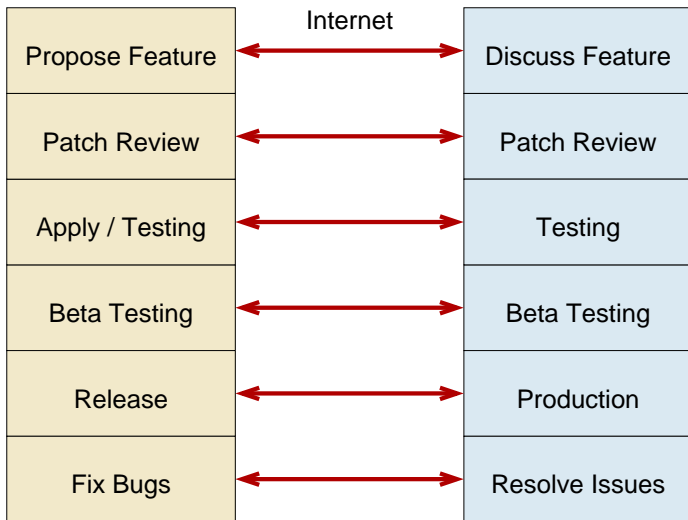
Proprietary Software Development



Open Source Development

Developers

Users



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

PG Life: <http://pglife.momjian.us/>

PgLife Postgres Community Life

Users

General

Re: How can I select rows by comparing an array data type column with multiple values ?

Other

Re: [BUGS] postgresql test failure with Python 3.4.0 in plpython_do

Announce

PG Partition Manager v1.7.1 Released

Developers

Hackers

Re: backup_label revisited

Commit

Improve the efficiency of certain jsonb get operations.

Versions

Stable: 9.3.4+, 9.2.8+, 9.1.13+, 9.0.17+, 8.4.21+ | Development: 9.4 beta, 9.4beta1

External

Blogs

Leo Hsu and Regina Obe: psql watch for batch processing

News

PostgreSQL Maestro 14.5 released

Media

Met Office selects 2ndQuadrant to help with data migration to open source - Computing

Events

PGDay France 2014

IRC

sz4man: plundra: i'm not the maintainer of this particular postgresql server, but i think it could be a good idea, can transaction id be manually set while doing begin? it should be always the same for the same script

plundra: sz4man: No, but in that case use application_name?

andres: Yea, I'd go for including the application_name in the log_line_prefix.

sz4man: so pgbadger or similar programs would distinguish these commits by application names?

plundra: pgbadger uses it to some extent, but not all that much I believe? But you can probably filter out based on it an generate a custom report.

sz4man: plundra: perhaps it would be not hard to find out, but could i ask quickly, how to set application name in psycopg2 ?

plundra: You have "Queries by application" in pgbadger.

plundra: sz4man: Throw application_name=foo on your DSN, maybe there it's possible with a (keyword)argument too, not sure.


sz4man: plundra: checking...

Zr40: can you only set application_name when connecting?


London 14:47 Berlin 15:47 Moscow 17:47 Mumbai 19:17 Beijing 21:47 Tokyo 22:47 Los Angeles 06:47 New York 09:47 São Paulo 10:47

Content updates automatically | [About](#) | [Submit Feedback](#)

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/>