

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: August, 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 Re: PostgreSQL System Views or Dictionary Tables
Other Re: [PG_UPGRADE] 9.6 to 10.5
Announce PgBouncer 1.9.0 released

Developers

Hackers Re: [HACKERS] Bug in to_timestamp().
Commit Remove unused configure test for ldpopen()
Versions **Stable:** 10.5+, 9.6.10+, 9.5.14+, 9.4.19+, 9.3.24+ | **Development:** 11 beta3+, 12 devel

External

Blogs Haroon .: PostgreSQL and IoT Data Localization, Integration, and Write Scalability
News PgBouncer 1.9.0 released
Tweets Have you used the @postgresql "jsonb_agg" function in your applications?
Media Connecting Apache Zeppelin and Apache Drill, PostgreSQL, etc.
Events PostgresOpen SV 2018

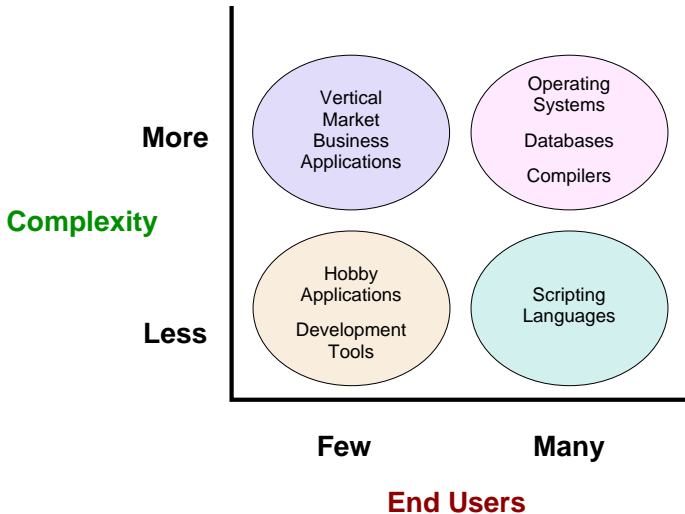
IRC

xocolati: on the plus side, barman can participate in synchronous replication
Snow-Man: apparently it also doesn't have the kind of verifications and safety checks that pgbackrest does either
Snow-Man: so it'll happily let you send WAL from two different clusters to the same archive, heh
Snow-Man: it's also got no facility for checking if a backup has latent corruption
Snow-Man: and doesn't check PG's page-level checksums.
Snow-Man: alk
ysch: rgn: You can make it into expressional index.
_rgn: perhaps something like this CREATE INDEX idx_name ON topic_snapshot(md5(guid || subject || message));
mlt-: Snow-Man: thanks on reconnect interval... just wanted to confirm I don't miss something. It is more like a yet another backup at the moment. Physical connection can be down, so there is no point to try to connect every 20 seconds.
jayjo: I've been researching and working on setting up a HA postgres cluster, I've diagramed what I have so far here - <https://t.imgur.com/nS4umjk.png> . I see some posts on dba.stack that have the pgbouncer before the haproxy. Does this architecture look sound?


London 22:31 Berlin 23:31 Moscow 00:31 Mumbai 03:01 Beijing 05:31 Tokyo 06:31 Los Angeles 14:31 New York 17:31 São Paulo 18:31
[Content updates automatically](#) | [About](#) | [Submit Feedback](#)

<http://pglife.momjian.us>

Open Source Software Availability



Is It As Good As Proprietary Software?



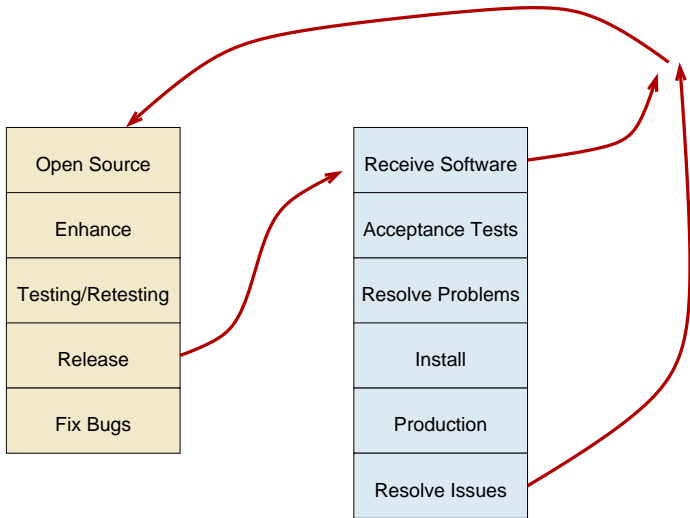
Problems With Open Source (and Democracy)

- ▶ 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/>