

The Democratization of Databases

BRUCE MOMJIAN



Democratic governments have a long history of success, with some setbacks. This talk explains how democratic principles underpin the success of the open source relational database Postgres.

<https://momjian.us/presentations>




Creative Commons Attribution License

Last updated: June 2024

Outline

1. History of governance structures
2. Strengths of each structure
3. Efficiency and messiness of democracy
4. Ultimate success of democracy
5. Software governance history
6. Where does this leave Postgres?
7. Democracy in action

1. History of Governance Structures



Pnyx Hill, Athens

<https://www.flickr.com/photos/wallyg/>

Autocracy

- Single person or small group in power
- Examples
 - dictatorship, e.g., Syria
 - absolute monarchy, e.g., Saudi Arabia
 - communist, e.g., North Korea
- First governing structure

Representative Democracy

- First established in ancient Athens
- Historically only a small group of the population could vote
- Elect representatives to vote on issues

2. Strengths of Each Structure



Chepstow Castle, Wales

<https://www.flickr.com/photos/damiavos/>

Autocracy Strengths

- Good for focusing a fixed amount of resources on a clear goal
- Examples
 - space exploration
 - military

Democracy Strengths

- Allows rapid adjustment for unclear goals
- Emboldens talent to act near the problem
- Expands the pool of talent
- Examples
 - consumer goods
 - software

3. Efficiency and Messiness of Democracy



Madison, Wisconsin

<https://www.flickr.com/photos/dennisdeery/>

Efficiency of Democracy

- Unleashes a flurry of activity based on agency, i.e., personal power
- Fluidity of solutions

Messiness of Democracy

- Hard to predict behavior
- Problems can get stuck
- Difficulty with large projects that span multiple elections
- Direct democracy rarely tried, yielded mixed results

4. Ultimate Success of Democracy



Moscow

<https://www.flickr.com/photos/varfolomeev/>

Democracy's Checkered History

- First democracy in Athens, suspended during wars
- Roman democracy ended in dictatorship
- Middle Ages, Magna Carta, Renaissance
- American Revolution, black and women's voting rights
- French Revolution leads to monarchy

Democracy Today


- Democracy continues to make steady progress
- There are still setbacks

Democracy in the Past 200 Years

<https://www.flickr.com/photos/dennisdeery/>

Number of world citizens living under different political regimes

The scale goes from -10 (full autocracy) to 10 (full democracy). Anocracies are those scoring between -5 and 5. "Colony" (coded as -20) includes not only colonies, but also countries that were not yet sovereign states (e.g. the Czech Republic and Slovakia in 1945–92).



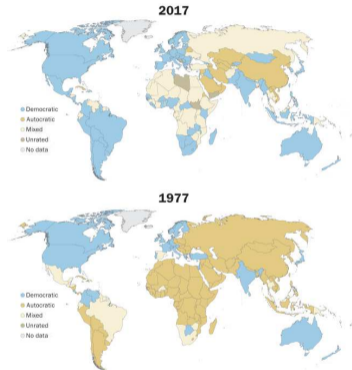
Source: World Population by Political Regime they live in (OWID (2016))
[OurWorldInData.org/a-history-of-global-living-conditions-in-5-charts/](https://ourworldindata.org/a-history-of-global-living-conditions-in-5-charts/) • CC BY

<https://ourworldindata.org/democracy>

Democracy in the Past 40 Years

Democracy has grown across the world over the past four decades

Regime types in each country, 1977 and 2017



Note: Map depicts data for the 167 countries included in the Polity IV database. Countries labeled "mixed" have a blend of democratic and autocratic regime characteristics. "Unrated" countries are those whose central government has or had completely collapsed, which are or were subject to foreign intervention or occupation, or which are or were in the midst of a regime transition. "No data" refers to states or entities that either have fewer than 500,000 people or are not internationally recognized as fully sovereign.

Source: Center for Systemic Peace's Polity IV Project.

PEW RESEARCH CENTER

<https://www.pewresearch.org/fact-tank/2019/05/14/more-than-half-of-countries-are-democratic/>

5. Software Governance History



Redwood City, California

<https://www.flickr.com/photos/nzdave/>

Proprietary Autocracy

- Executives make decisions
 - input from sales and marketing
 - indirectly from customers
- Decision matrix
 - “If we do this, what percentage of customers will we gain?”
 - “If we don’t, what percentage of customers will we lose?”
 - gains are more sensitive to the decision than losses
 - “Do gains or avoided losses justify the implementation costs?”

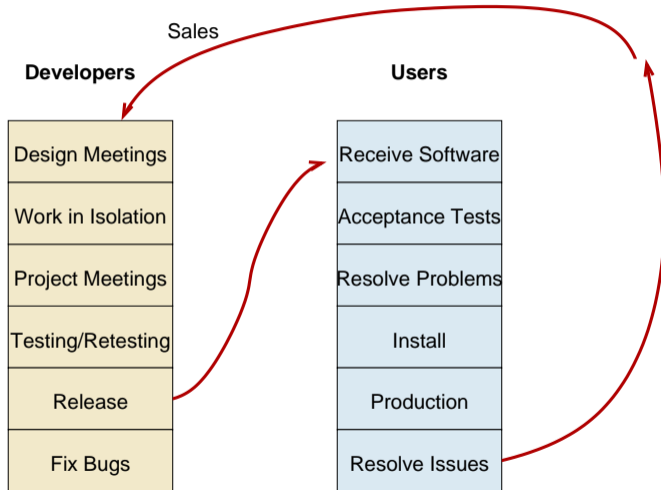
Open Source Democracy

- Mix of direct democracy and meritocracy
- Voting can be problematic
- Sometimes too much feedback
 - bike shedding
- Bad decisions can be quickly corrected, i.e., easy to revert
- No reliable road map


Internet Makes Direct Democracy Possible

- Easy to share ideas and iterate new solutions
- Rapid global communication

Autocratic Development Flow



Democratic Development Flow



Hybrid Development Models

- Single company controls development
- Open source distribution
- Uses autocratic decision matrix
- Lacks democratic feedback
- More comfortable for users transitioning from proprietary since there is a single company to contact
- Examples
 - MySQL, MariaDB
 - MongoDB

6. Where Does This Leave Postgres?




https://www.flickr.com/photos/tomas_vondra/

Postgres Wins Because Democracy Wins

- Uses democracy to attract talent
 - talent pool can easily compete with proprietary staff
- Superior feedback and decision matrix lead to success
- This is a challenge for niche software
- Setbacks still possible

Rise of Open Source

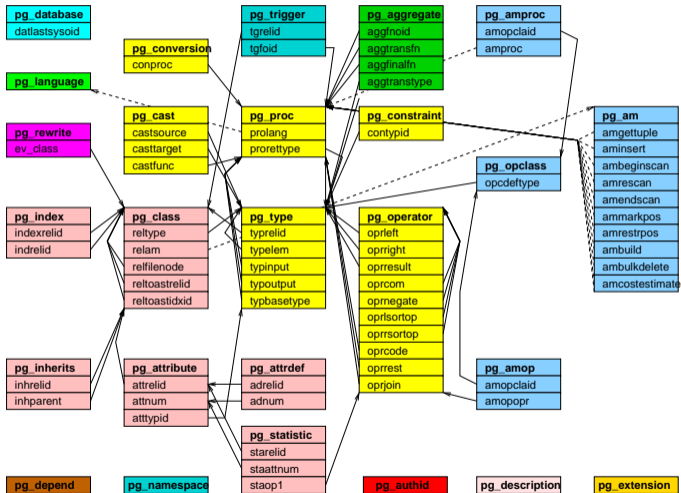


7. Democracy in Action

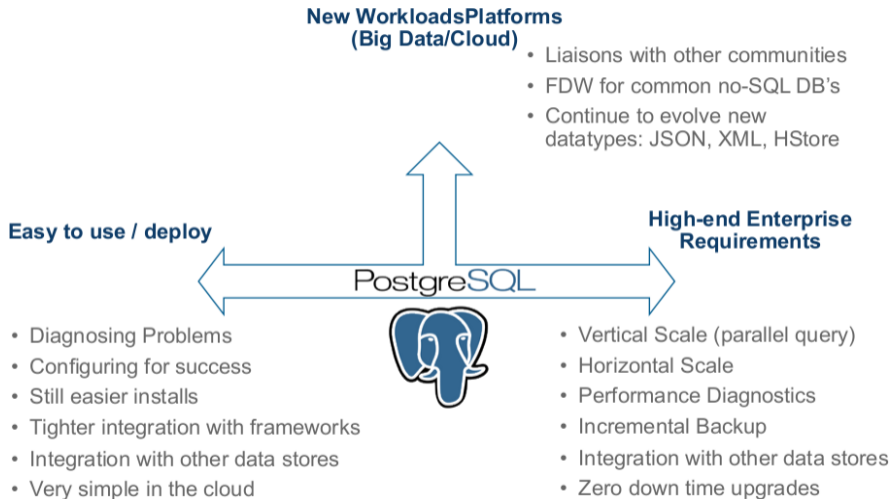


https://www.flickr.com/photos/tomas_vondra/


Postgres Extensibility



Many Focuses



Foreign Data Wrappers



Foreign Data Wrappers

- 100+ interfaces to foreign data
- Read/write
- Sophisticated push down of joins, sorts, and aggregates



users tags recent

in Distributions PGXN Search

acl administration aggregate
aggregate function analytics arithmetic array arrays audit
auditing count custom background worker data type
datatype dictionary distinct estimate example
extension external data fdw
foreign data wrapper
function functions hash hint index
json jsonb log logging maintenance
math md5 mean median meta mongodb
monitoring parallel partitioning
performance plan plpgsql postgresql
replication search sql sql med statistics table
template testing trigger type users

PGXN, the PostgreSQL Extension network, is a central distribution system for open-source PostgreSQL extension libraries.

Recent Releases

powa 5.0.1

An extension gathering pg_stat_statements and other plugins statistics

pg_partman 5.2.4

Extension to manage partitioned tables by time or ID

vasco 0.1.0

vasco: Maximal Information Coefficient (MIC) Extension for Postgres

ddsketch 1.0.1

ddsketch, a data structure for on-line accumulation of quantiles.

lower_quantile 1.0.3

aggregate calculating lower quantile

PostGIS

- PostGIS is a full-featured Geographical Information System (GIS)
- Implemented as a extension
- Independent development team and community



PostgreSQL Tenth Anniversary in 2006



https://www.postgresql.org/files/community/conference06/conference_group.html

My Postgres Activities



Voting

- Committers nominate new committers
- Core team nominates new core members
- Development is open to all, even occasional visitors
 - “Let the best idea win!”
 - “Where did that guy come from?”
 - Focus talent like a lens on every task

Roadmap

- Individuals and political parties have roadmaps
- Democratic governments don't
- Developers and companies have roadmaps
- Postgres doesn't

Users

- General Re: Will PQsetSingleRowMode get me results faster?
- Other Re: Aggressive vacuum
- Announce SQLPage v0.32: build web apps in SQL, now with more style !

Developers

- Hackers Moving the vacuum GUCs' docs out of the Client Connection Defaults section
- Commit Allow changing autovacuum_max_workers without restarting.
- Versions **Stable:** 17.2+, 16.6+, 15.10+, 14.15+, 13.18+ | **Development:** 18 devel, in commitfest

External

- Blogs Paul Ramsey: Running an Async Web Query Queue with Procedures and pg_cron
- News SQLPage v0.32: build web apps in SQL, now with more style !
- Media Amazon Aurora PostgreSQL Limitless Database is now generally available - AWS Blog
- Events CERN PGDay 2025

IRC (also Slack)

theseb: can you have a from clause with two tables separated by a space like that?

GrayShade: That's an alias, not a second table

GrayShade: If there's a comma between them, it's a cross join

theseb: GrayShade: thanks...never heard of alias

theseb: GrayShade: So the result is given the name B?

GrayShade: They're useful when you want to join a table by itself (or when the name is too long or chosen dynamically)

GrayShade: No, the result columns have the same name as before, but now you can also do where b.x = ...

Xgc: That SQL is actually invalid, because A is no longer available in that WHERE clause. Only table B is known in the result of the FROM clause. (re: question from theseb)

London 00:33 Berlin 01:33 Moscow 03:33 Mumbai 06:03 Beijing 08:33 Tokyo 09:33 Los Angeles 16:33 New York 19:33 São Paulo 21:33

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

Conclusion



https://www.flickr.com/photos/tomas_vondra/



<https://momjian.us/presentations>

<https://www.flickr.com/photos/atbaker/>