



Why to Choose Open Source Database Solutions in 2025

Takis Stathopoulos, PhD

Enterprise Architect, Percona



Takis Stathopoulos

Enterprise Architect

Tackling the more complex, challenging and interesting large scale cases

Solutions Architect

Delivering Complex Enterprise projects

IT Infrastructure Architect and Team Leader

Transforming from proprietary to Open Source

Research Engineer

Linux Kernel patch for testing a new network traffic control framework

Open Source has won!



Open Source is everywhere



CockroachDB

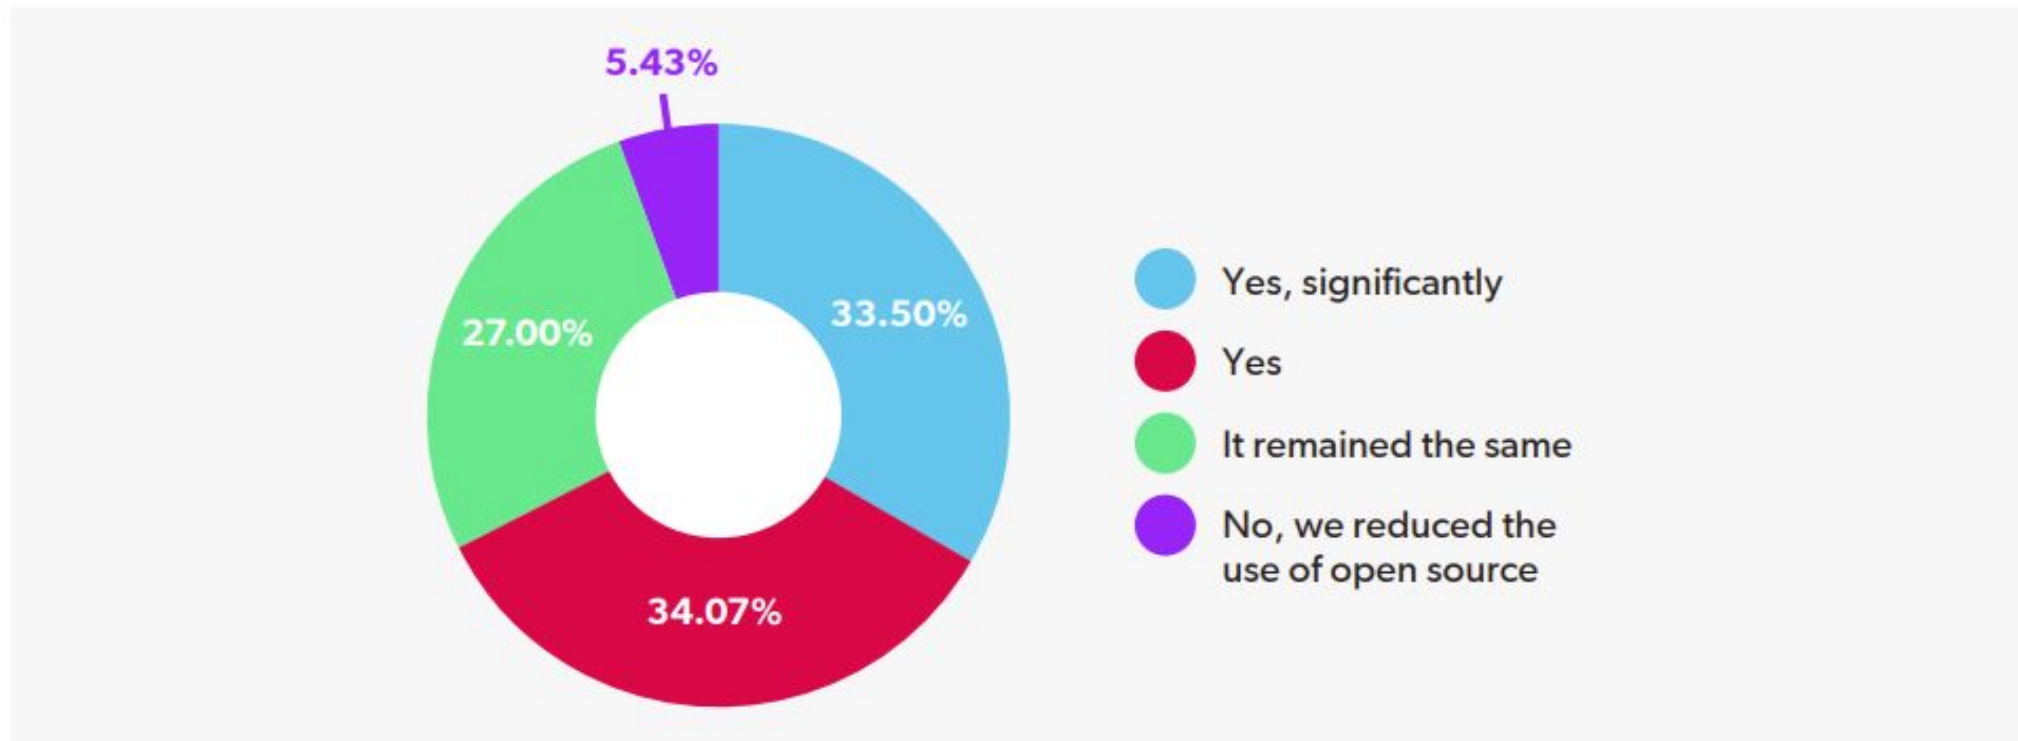




What will it be different in 2025?

- New workloads
- A more complex landscape
- Open Source corporate headwinds
- Cost-efficient growth

Has Your Organization Increased the Use of Open Source Software Over the Last Year?



<https://opensource.org/blog/announcing-the-2024-state-of-open-source-report>

Open Source vs Proprietary

Reason for Using OSS	Percentage
No License Cost, Overall Cost Reduction	36.64%
Functionality Available to Improve Development Velocity	30.71%
Stable Technology with Community Long-Term Support	27.64%
Access to Innovations and Latest Technologies	26.86%
To Reduce Vendor Lock-In	21.29%
Open Standards and Interoperability	20.93%
To Modernize Technology Stack	20.00%
Fast Moving / Constant Enhancements, Releases, and Patches	14.21%
Community-Oriented and Transparent	12.86%
Ability to Contribute to, and Influence Direction of, Open Source Projects	10.43%
Large Selection of Options for Similar Functionality	8.14%
Makes It Easier to Hire or Retain Employees	6.14%



No License Cost, Overall Cost Reduction



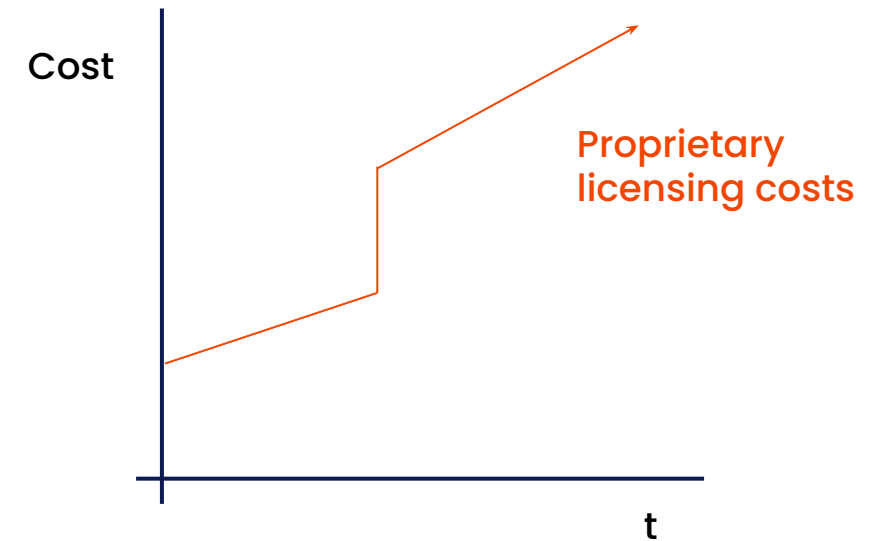
**Functionality to improve developer velocity
Access to Innovations and Latest Technologies**

51.5% of respondents working in government or public services said no license cost and overall cost reduction was their reason for using OSS.


<https://opensource.org/blog/announcing-the-2024-state-of-open-source-report>

Lower TCO

- **No license cost**
 - Reverse the trend
- **Overall cost reduction**
 - Support?
 - Training and Know-how
- **Support Optionality. Power to:**
 - Select Support Vendor
 - Select Support
 - Select Scope



Innovation Rate





Source: <https://momjian.us/main/writings/pgsql/forever.pdf>

New workloads

pgvector / pgvector Public Notifications Fork 597 Star 12.7k

[Code](#) [Issues 9](#) [Pull requests 7](#) [Actions](#) [Security](#) [Insights](#)

master [Go to file](#) [Code](#)

 ankane Added iterative index scans to troublesh... 5bc7937 · 3 days ago 
.github/workflows Added Postgres 18 to CI [...]
sql Added casts for arrays to ...
src Removed unneeded code
test Improved test [skip ci]
.dockerignore Added dockerignore [ski...]
.editorconfig Use consistent indentatio...
.gitignore Added log to .gitignore [...]

About

Open-source vector similarity search for Postgres

- [nearest-neighbor-search](#)
- [approximate-nearest-neighbor-search](#)

- [Readme](#)
- [View license](#)
- [Security policy](#)
- [Activity](#)
- [Custom properties](#)

-  12.7k stars
-  102 watching
-  597 forks

- AI
- Vectors
- Machine learning
- Time Series
- GIS
- ...

Freedom to innovate

- Deploy where you want, and how you want
- Modify, extend, optimize
- Test as long as you want
- Standardize only once successful
- Skip lengthy procurement and buying process





A complex landscape

Different kinds of Open

Source Code License: GPL, Apache, BSD, Public Domain, ...

Cost: Free as in **Beer** or Free as **Freedom**?

Ownership

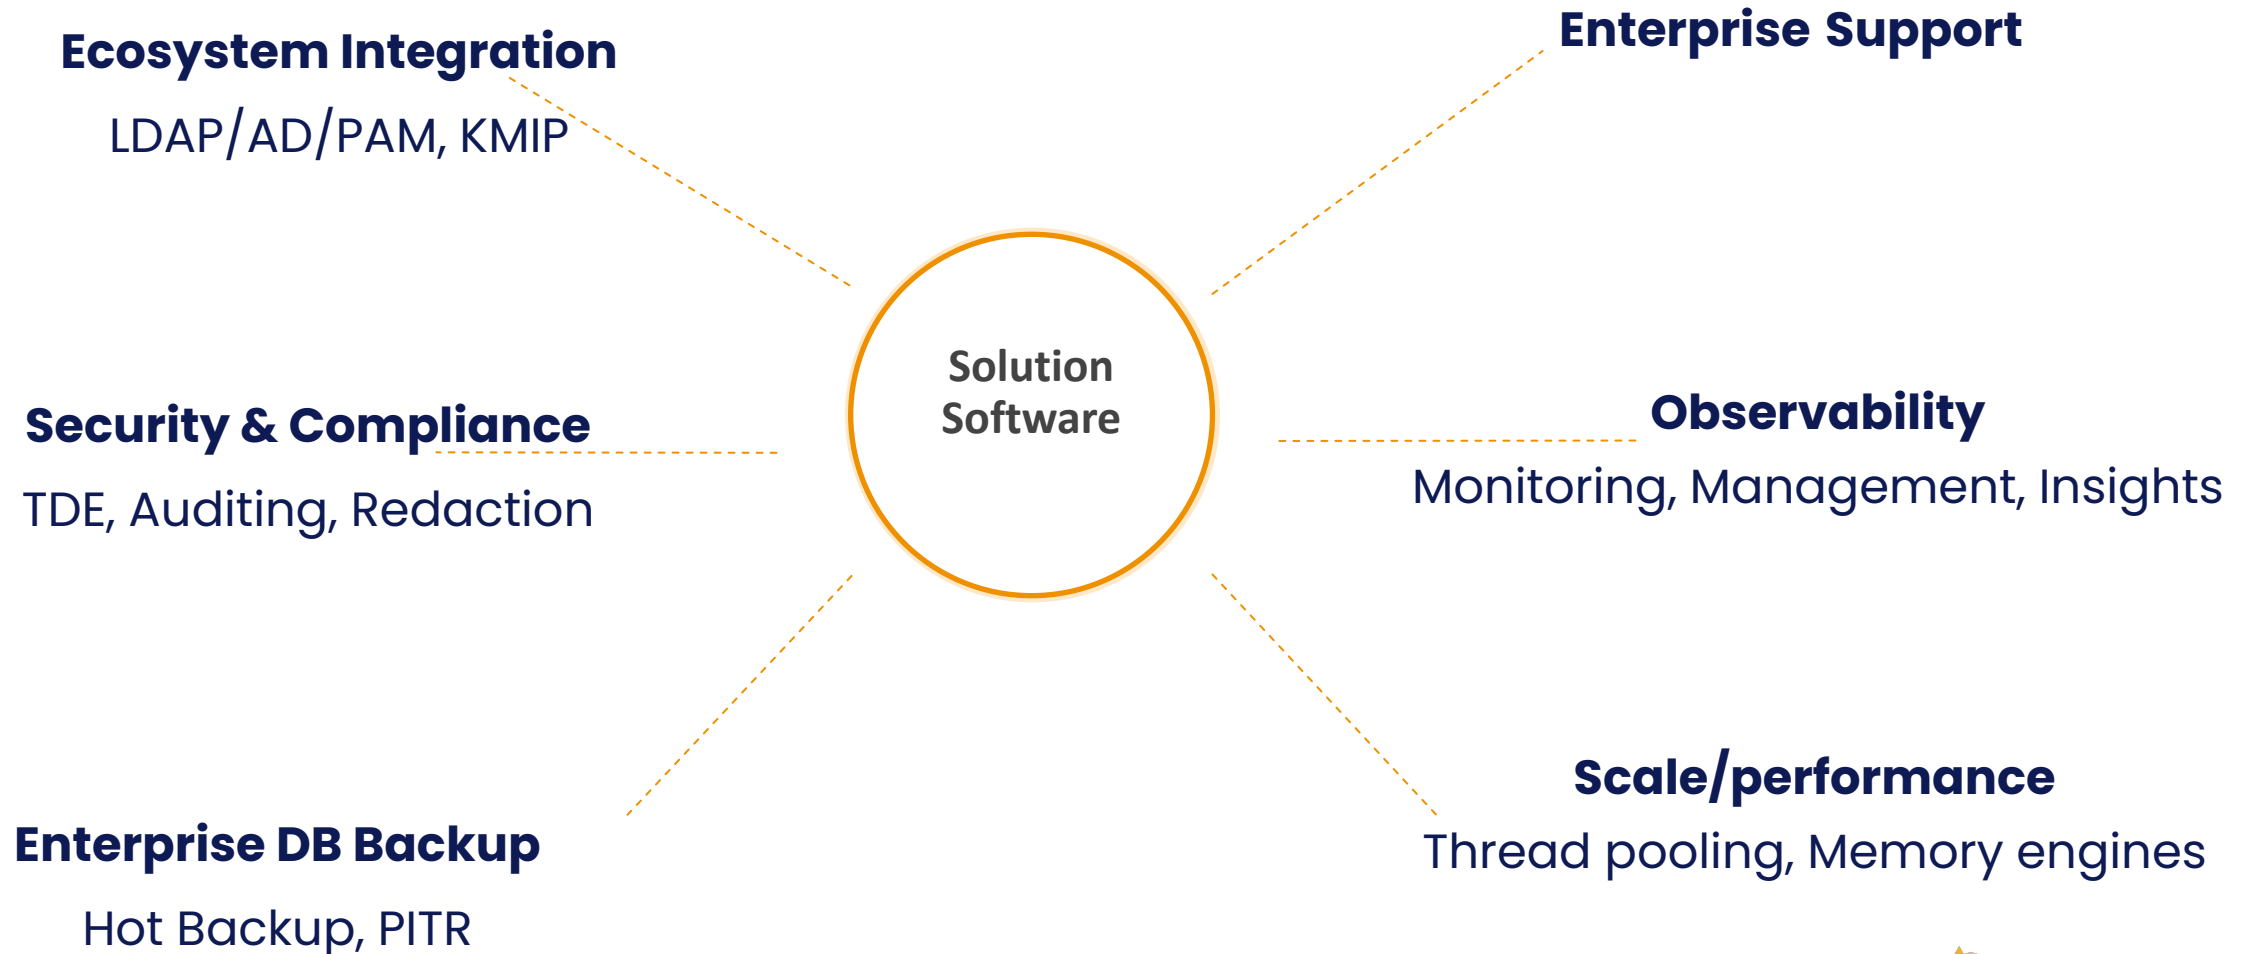
- Trademarks
- Source Code
- Intellectual Property



**open source
initiative**®

1. Free Redistribution
2. Source code available
3. Derived Works
4. Integrity of The Author's Source Code
5. No Discrimination Against Persons or Groups
6. No Discrimination Against Fields of Endeavor
7. Distribution of License
8. License Must Not Be Specific to a Product
9. License Must Not Restrict Other Software
10. License Must Be Technology-Neutral

Is your selected solution open?



Evaluation factors

Is the **core software** free to innovate?

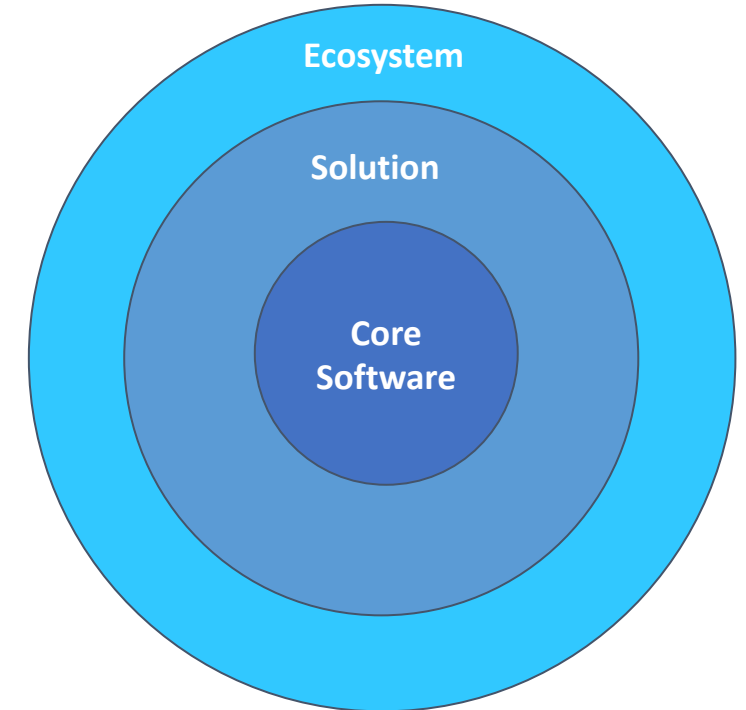
- Develop and expand?
- Use?

Is the **full solution software** open?

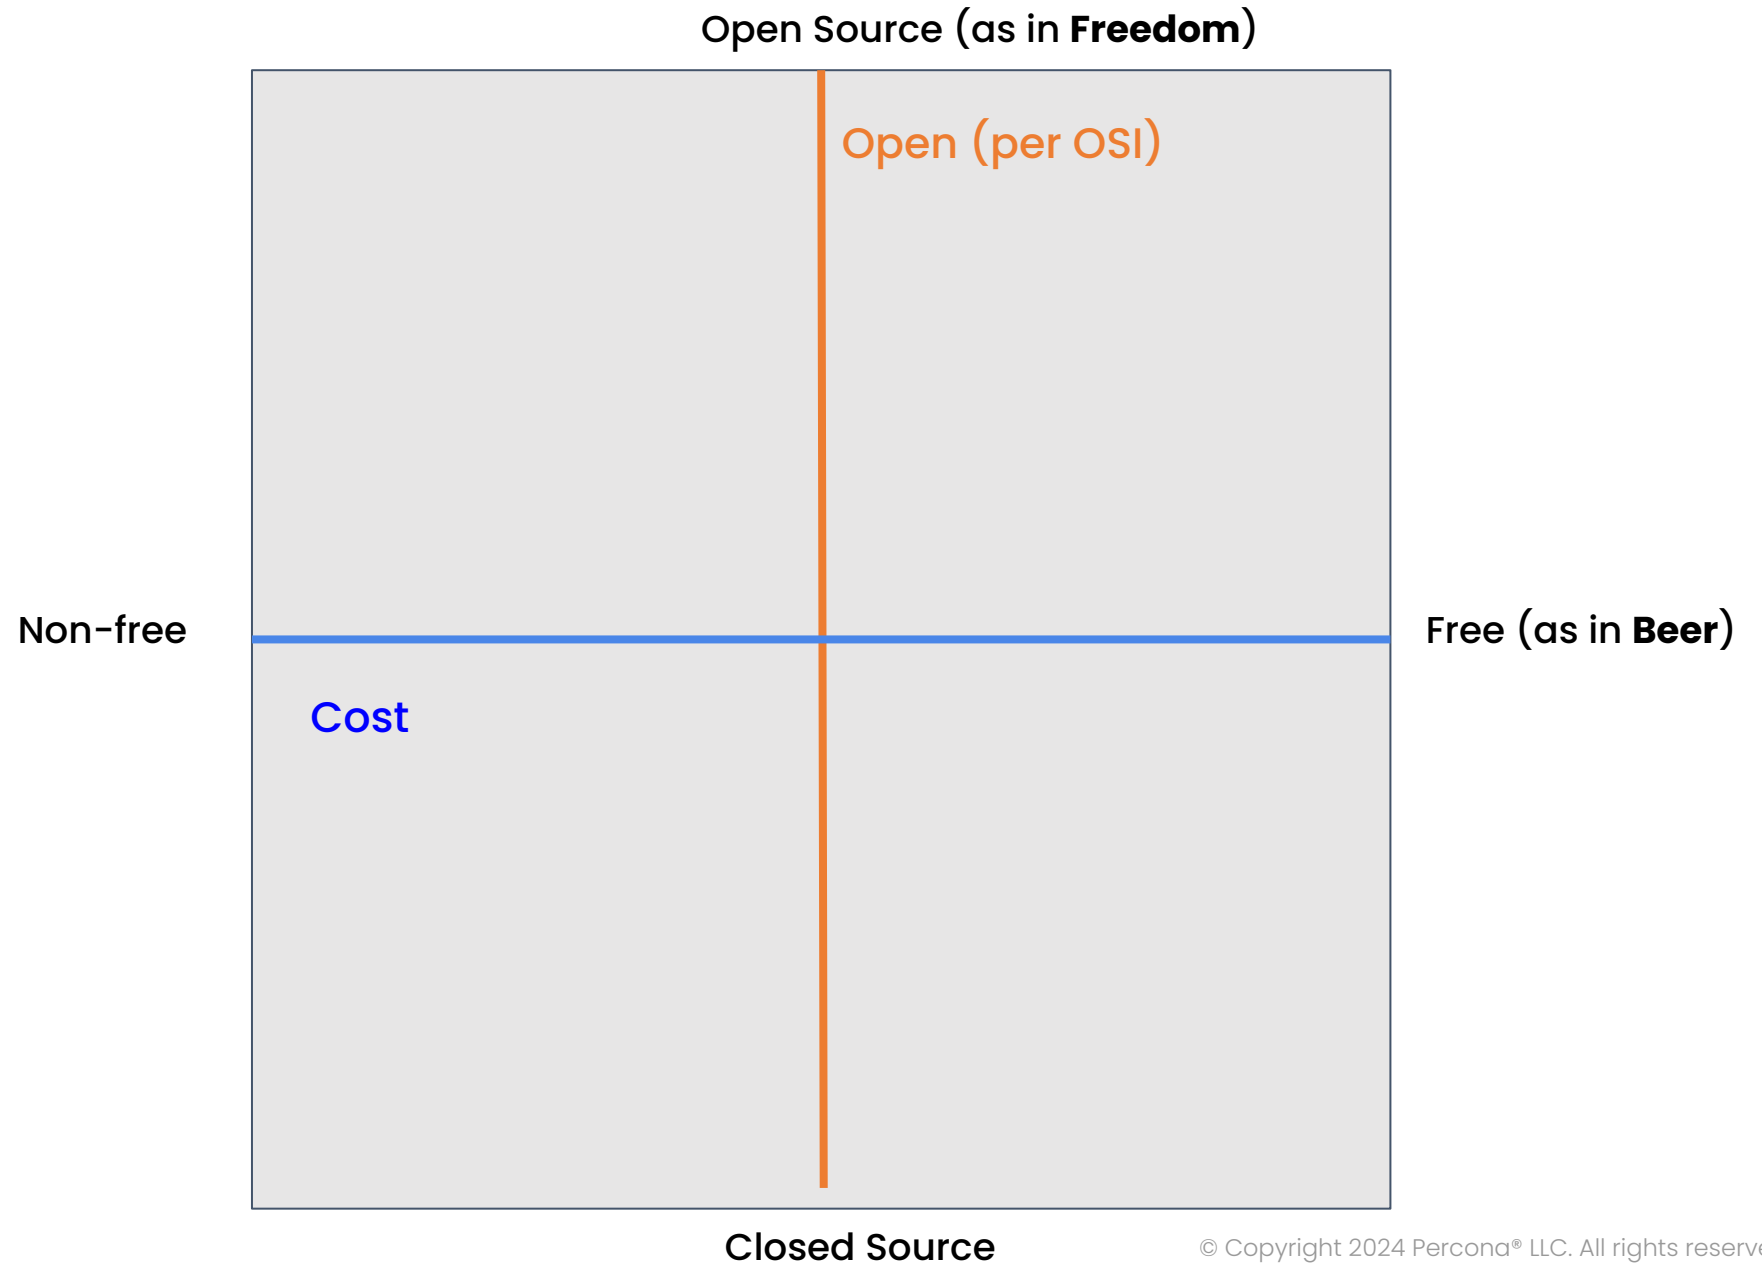
- Monitoring
- Compatibility
- Compliance

Is the **ecosystem** deployment open?

- Automations
- Extensions/plugins
- Know-how



An evaluation framework





Open Source (as in Freedom)

Non-free

Free (as in Beer)

<p>AWS RDS MySQL Google Cloud MySQL</p>	<p>MySQL CE PS for MySQL</p>
<p>AWS RDS Aurora for MySQL MySQL EE</p>	

MySQL CE
MySQL EE
Percona Server for MySQL
AWS RDS MySQL
AWS RDS Aurora MySQL
Google Cloud SQL MySQL

Closed Source



Open Source (as in Freedom)

PostgreSQL
EDB Postgres Advanced Server
Percona Distribution for PostgreSQL
EDB Postgres Advanced Server Non-free
Postgres Pro Enterprise
Fujitsu Enterprise Postgres
Google Cloud SQL PostgreSQL
AWS RDS Aurora PostgreSQL
AlloyDB for PostgreSQL

<p>AWS RDS PostgreSQL Google Cloud PostgreSQL</p>	<p>PostgreSQL</p> <p>Percona Distribution for PostgreSQL</p>
<p>AWS RDS Aurora AlloyDB for PostgreSQL</p> <p>EDB Postgres Advanced Server Postgres Pro Enterprise Fujitsu Enterprise Postgres</p>	

Free (as in Beer)

Closed Source



Open Source (as in Freedom)

Non-free

Free (as in Beer)

MongoDB CE
Percona Server for
MongoDB
MongoDB EA
Atlas
AWS DocumentDB

MongoDB CE
Percona Server for
MongoDB

AWS DocumentDB

MongoDB EA
Atlas

Closed Source

Solution partially built on open

Open Core business model

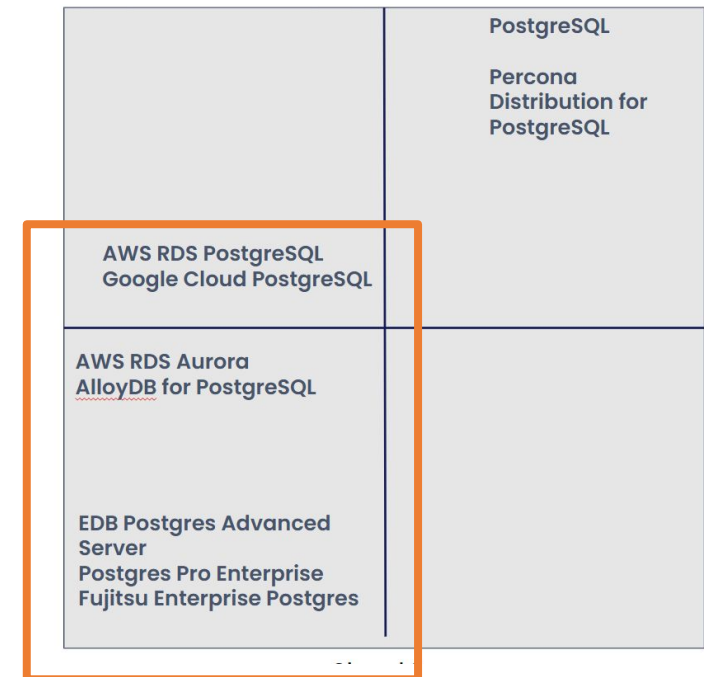
- Some functionality open and free

Other functionality closed and proprietary

- Potential for Vendor Lock-in

A different kind of proprietary?

- Potentially with greater data portability



Closed ecosystems built on open software

Cloud providers and packaged Solutions

Build on Vanilla Open Source or

- Cloud providers trying to differentiate
- Develop vendor specific proprietary editions

However ecosystem is proprietary:

know-how, free to use in any platform, licensing costs, vendor-lock in

Cost and Know-how implications

What is Cloud Computing?

An analogy: think of electricity services...

You simply plug into a vast electrical grid managed by experts to get a low cost, reliable power supply – available to you with much greater efficiency than you could generate on your own.



Power is a utility service - available to you on-demand and you pay only for what you use.



Cloud is not a utility service anymore



Open Source corporate headwinds



2023: is Corporate Open Source Viable?

Redis Adopts Dual Source-Available Licensing

MARCH 20, 2024



Rowan Trollope

Future Redis releases will continue to offer free and permissive use of the source code under dual RSALv2 and SSPLv1 licenses; these releases will combine advanced data types and processing engines previously only

MongoDB Issues New Server Side Public License for MongoDB Community Server



New License Leads the Way for Open Source in the Cloud Era

NEW YORK, NY - October 16, 2018 – Today [MongoDB, Inc.](#) (Nasdaq: MDB), the leading modern, general purpose database platform, issued a new software license, called [Server Side Public License](#) (SSPL), for MongoDB Community Server. The license clearly and explicitly states the conditions of deploying MongoDB - or any other open source project licensed under the SSPL - as a service. All versions of MongoDB's Community Server released after today, including patch fixes for prior versions, will be licensed under the SSPL.

 COMPANY



HashiCorp adopts Business Source License

HashiCorp adopts the Business Source License to ensure continued investment in its community and to continue providing open, freely available products.

AUG 10 2023 | [ARMON DADGAR](#)

03 JUNE 2021 NEWS

Elastic License Update

By [Steve Kearns](#)

Share



In January 2021, we [announced](#) that starting with version 7.11, we would be changing the Apache 2.0 portions of Elasticsearch and Kibana source code to be dual licensed under [Elastic License](#) and SSPL, at the users' discretion. As part of that change, we created Elastic License 2.0 (ELv2) as a permissive, [fair-code](#) license, which allows free use, redistribution, modification, and derivative works, with only three simple limitations, outlined in our original [announcement](#).

Redis history



vmware®

 **redislabs**
HOME OF REDIS

BSD 3-Clause	Redis Source Available License (RSAL)	Redis Enterprise License
<ul style="list-style-type: none"> Clients Tools Benchmarks 	<ul style="list-style-type: none"> RedisSearch RedisGraph RedisJSON RedisBloom RedisAI More... 	<ul style="list-style-type: none"> Shared-Nothing Architecture Multi-Tenant Proxy-Based Linear Scaling Pure In-Memory Replication Single Digit Failover Enhanced Storage Engine Multi-Layer Security Backups and DR Active-Active (CRDT-based) Geo-Distributed Active-Passive Geo-Distributed Redis on Flash (SSD, PMEM) Integrated Modules

Dual License: RSALv2 + SSPLv1 Integrated Modules	Redis Enterprise License
<ul style="list-style-type: none"> RedisSearch Redis Stack RedisJSON RedisGraph RedisTimeSeries RedisBloom RedisGears 	<ul style="list-style-type: none"> Shared-Nothing Architecture Multi-Tenant Proxy-Based Linear Scaling Single Digit Failover Multi-Layer Security Enhanced Storage Engine Data Integration Engine Backups and DR Active-Active Geo-Distributed Active-Passive Geo-Distributed Redis on Flash Integrated Modules

2009

2010

2015 - v4

2018 - v5

2024 - v7.4

REmote
DictionarY Server

VMware (later
Pivotal) sponsors
the project

Redislabs employs
Salvatore and
becomes the main
sponsor

SSPL for Redis
Modules

RSALv2

Salvatore
Sanfilippo

Valkey is created - 2024




Introducing the Valkey project



Fully compatible
with Redis OSS 7.2



Vendor Neutral
BSD-3 Licensed



Built by contributors in
the open source
community

Backed by industry leaders

Google Cloud

aws

ORACLE®



HUAWEI



ERICSSON



PERCONA



aiven



HEROKU



Alibaba Cloud



CHAIN
GUARD

verizon[✓]

The tide is turning

Valkey is a viable and robust solution

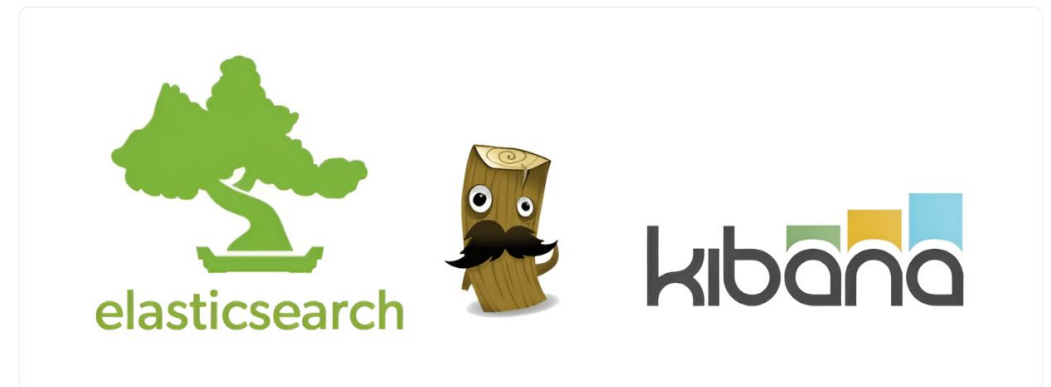
Open Source is not going anywhere

Has proved again the power of ecosystem

Elasticsearch Is Open Source. Again!

By Shay Banon

29 August 2024



[D.N.A] Elasticsearch and Kibana can be called Open Source again. It is hard to express how happy this statement makes me. Literally jumping up and down with excitement here. All of us at Elastic are. Open source is in my DNA. It is in Elastic DNA. Being able to call Elasticsearch Open Source again is pure joy.



Cost-efficient growth

Case studies and examples

From **Cloud-repatriation** to **Cloud Agnostic** and **Cloud Native**

Cost efficiency

- Longer runway for startups
- Increase Profits

Cloud Agnostic

- Heavy investment in MongoDB Atlas (400 clusters in prod)
- Now they must run in a location without Atlas (e.g. OCI)

K8s as the Cloud Native Cloud Vendor Agnostic layer

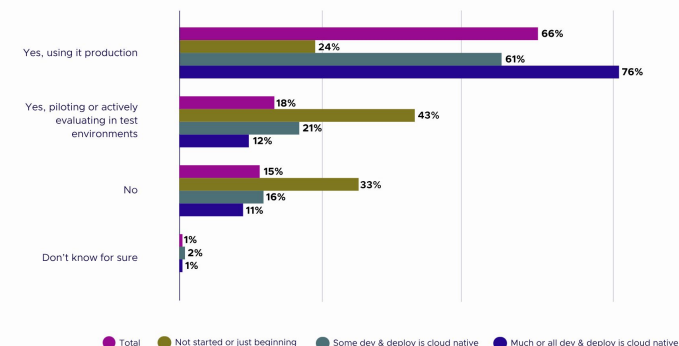
Moving from niche to the mainstream

- Cost, Complexity, Vendor Strategy



<https://arstechnica.com/information-technology/2024/10/basecamp-maker-37signals-says-its-cloud-exit-will-save-it-10m-over-5-years/>

To What Extent Has Your Organization Adopted Cloud Native Technologies?



<https://www.cncf.io/reports/cncf-annual-survey-2023/>



Conclusions

Key Takeaways

- Different kinds of Open exist
 - Open Software (OSI definition)
 - Open-based solutions (but not fully open)
 - Open Ecosystems (K8s as an example)
- Potential Lock-in comes in different ways
 - Data, Functionality, Automations, Convenience
- Open Source is again on the rise
 - Source Available and Open Core is still happening
 - But less successful as time goes by

Predictions for 2025

License changes and switching to Source Available will be less viable

- **The are still going to happen**

Open Source and Foundations on the rise

- E.g. CNCF
- vs vendor-provided Open Source

Permissive license mature projects e.g. PostgreSQL on the rise

- Rich, diverse, ecosystems
- **Need to read the fine lines!**

People making more conscious architecture decisions

- Cloud is great, K8s is great

Different use cases need different kinds of great:

- rapid
- cost-effective
- portable



PERCONA
.CONNECT



Gracias!