



Limassol, Cyprus 2024

PERCONA UNIVERSITY

Open Source Databases Meetup

Welcome



Peter Zaitsev,
Founder at Percona
April 20, 2024



Limassol, Cyprus 2024

PERCONA
UNIVERSITY

Veronica Urjan

Your friendly
Percona Sales Person

veronica.urjan@percona.com



WE ARE HIRING

We are **hiring**.
Check our openings.



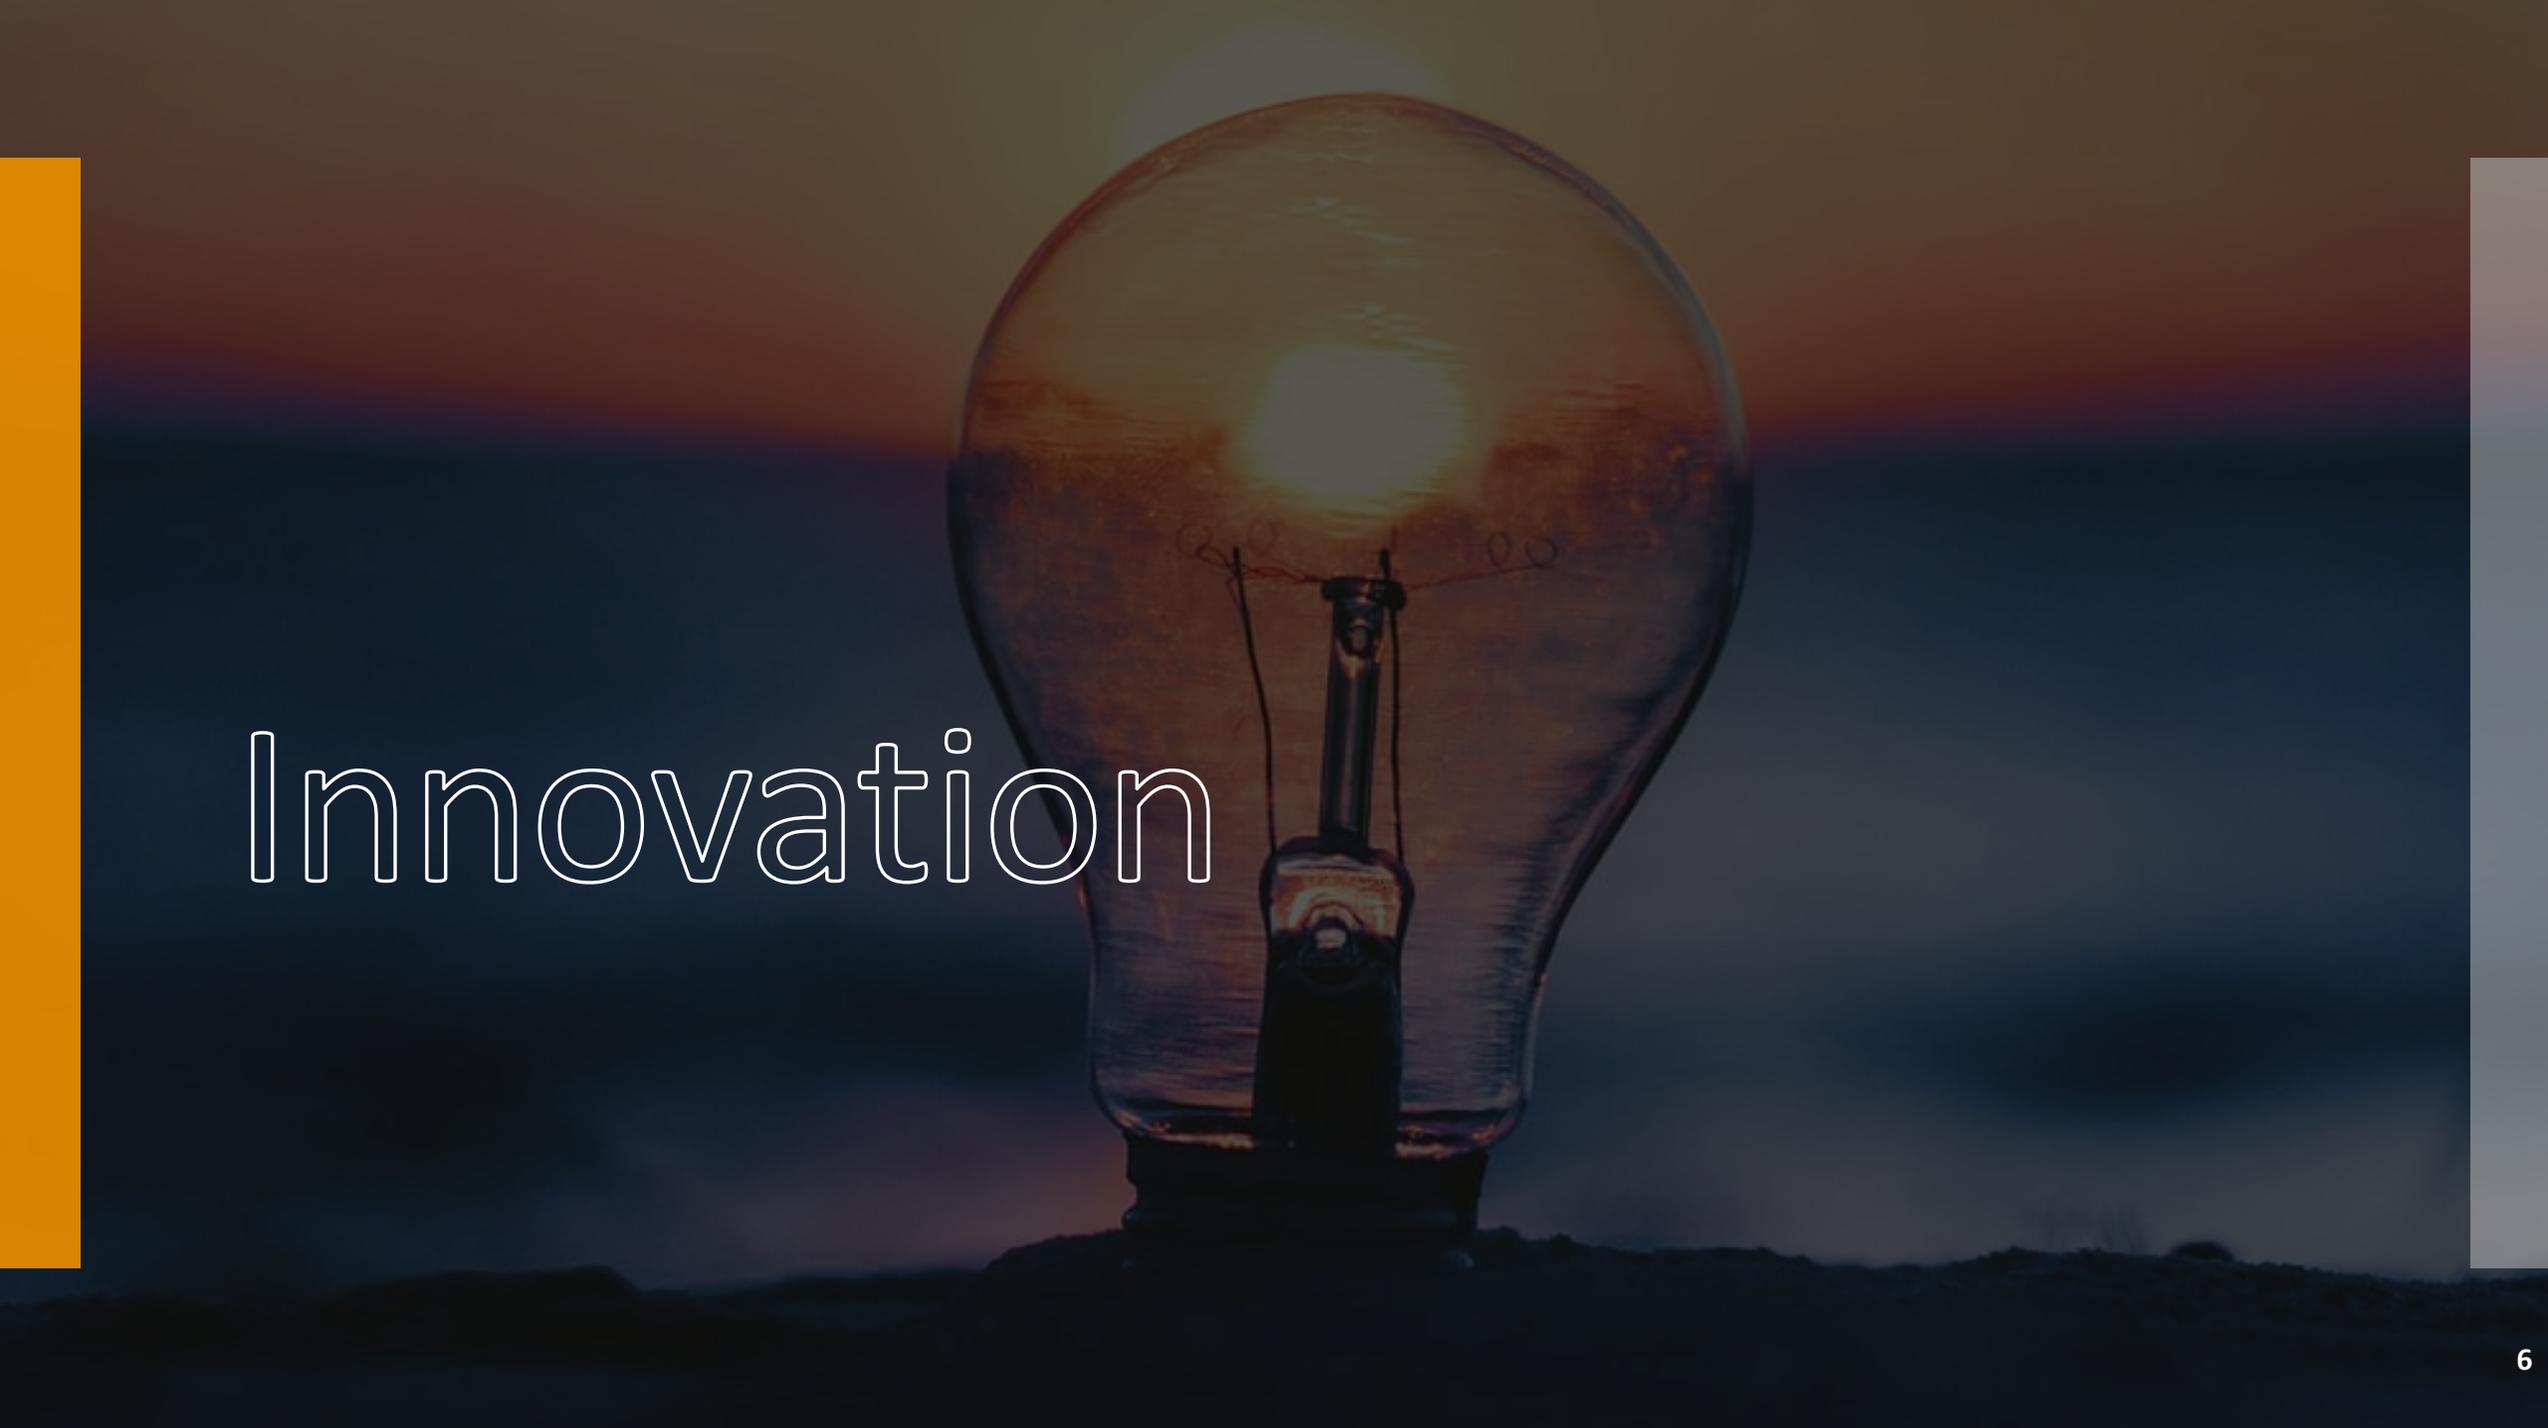
Current openings include:

- Senior Software Engineer (PostgreSQL)
- Support Engineer (PostgreSQL)
- PostgreSQL Evangelist
- ...and more!

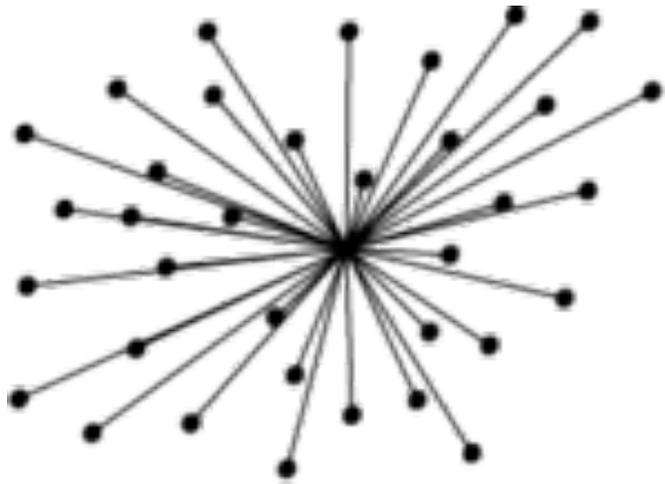


Bring Percona University to Your City!

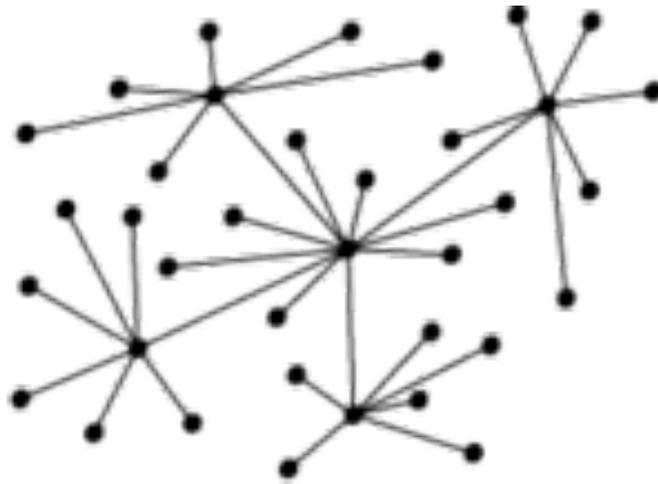


A glowing lightbulb is the central focus, set against a background of a sunset or sunrise. The lightbulb is illuminated from within, casting a warm, golden glow. The background shows a gradient of colors from dark blue at the bottom to orange and yellow at the top. The word "Innovation" is written in a white, outlined font across the middle of the image. On the left side, there is a solid orange vertical bar, and on the right side, there is a light gray vertical bar.

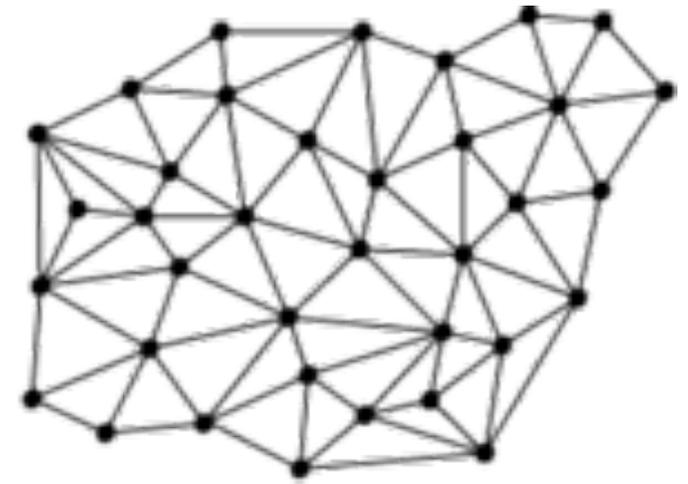
Innovation



centralised



decentralised



distributed

Distributed Databases



Separation of Storage and Compute

Serverless



A toy figure of a pink unicorn with large pink wings and a blonde fairy figure riding it, set against a background of a beach and ocean. The unicorn has a yellow mane with purple flowers, a purple tail, and blue floral patterns on its legs. The fairy has blonde hair, a gold crown, and a green and yellow dress. The background shows a sandy beach, blue water, and a dark rock formation on the right.

HTAP – Hybrid Transactional Analytical Processing



Data Pipelines

Multiple
Databases

Microservices

Developers Authority

Cloud



**Not just Relational
Any More**



Time Series



Graph



Data Structures

Different Models



Multi-Model

Some Databases Support Multiple Data Models, Some even Talk Different Languages/Protocols

Shapechangers

Clickhouse can speak PostgreSQL and MySQL Protocols

VictoriaMetrics has InfluxDB and Graphite API

FerretDB allows to use PostgreSQL as if it were MongoDB

Babelfish turns PostgreSQL in MS SQL Compatible Database



Vector Search

Technology Supporting building applications which use AI



AI Algorithms do not work with words and Images inside



They work with Vectors



Objects encoded to Vector form called "Embeddings"



Finding Similar Vectors is a key task



Traditional Database Algorithms do not handle it well

What's Up with Vectors

<https://manticoresearch.com/blog/vector-search-in-databases/>

Specific Tasks



Pattern Recognition



Semantic Search



RAG – “Retrieval Augmented Generation”

Vector search support in databases

Opensource vector dbs

Milvus	2019
Vespa	2020
Weaviate	2021
Qdrant	2022

Opensource dbs and search engines

PostgreSQL	2021
Lucene	2021
Opensearch	2022
Redis	2022
SOLR	2022
Cassandra	2023
Typesense	2023
Clickhouse	2023
Manticore Search	2023
Meilisearch	2023
MariaDB	In progress
MySQL	Not yet

Non-open source dbs

Elasticsearch	2019
Oracle	2023
MongoDB	2023

Clouds

Pinecone	2019
Amazon Elasticsearch / Opensearch	2020
Google Cloud Platform	2021
Alibaba Cloud AnalyticDB	2023
Azure	2023
Amazon DocumentDB	2023
Cloudflare Vectorize	2023

What Is the Biggest Factor Impacting Open Source Now?

Many will Say:
Cloud



**Maximize and
Simplify
Adoption**

**Change
Opportunities
for Monetization**

Cloud Impact

A close-up portrait of Marten Mickos, a man with short dark hair and blue eyes, wearing a light-colored suit jacket and a white shirt. He is looking directly at the camera with a slight smile. The background is a plain, light gray.

Remember...

**Marten Mickos: "Open Source
Is Not a Business Model"**

Open Source Ownership and Governance

- **Foundation Driven (Multiple Vendors)**
- **Single Vendor Driven**

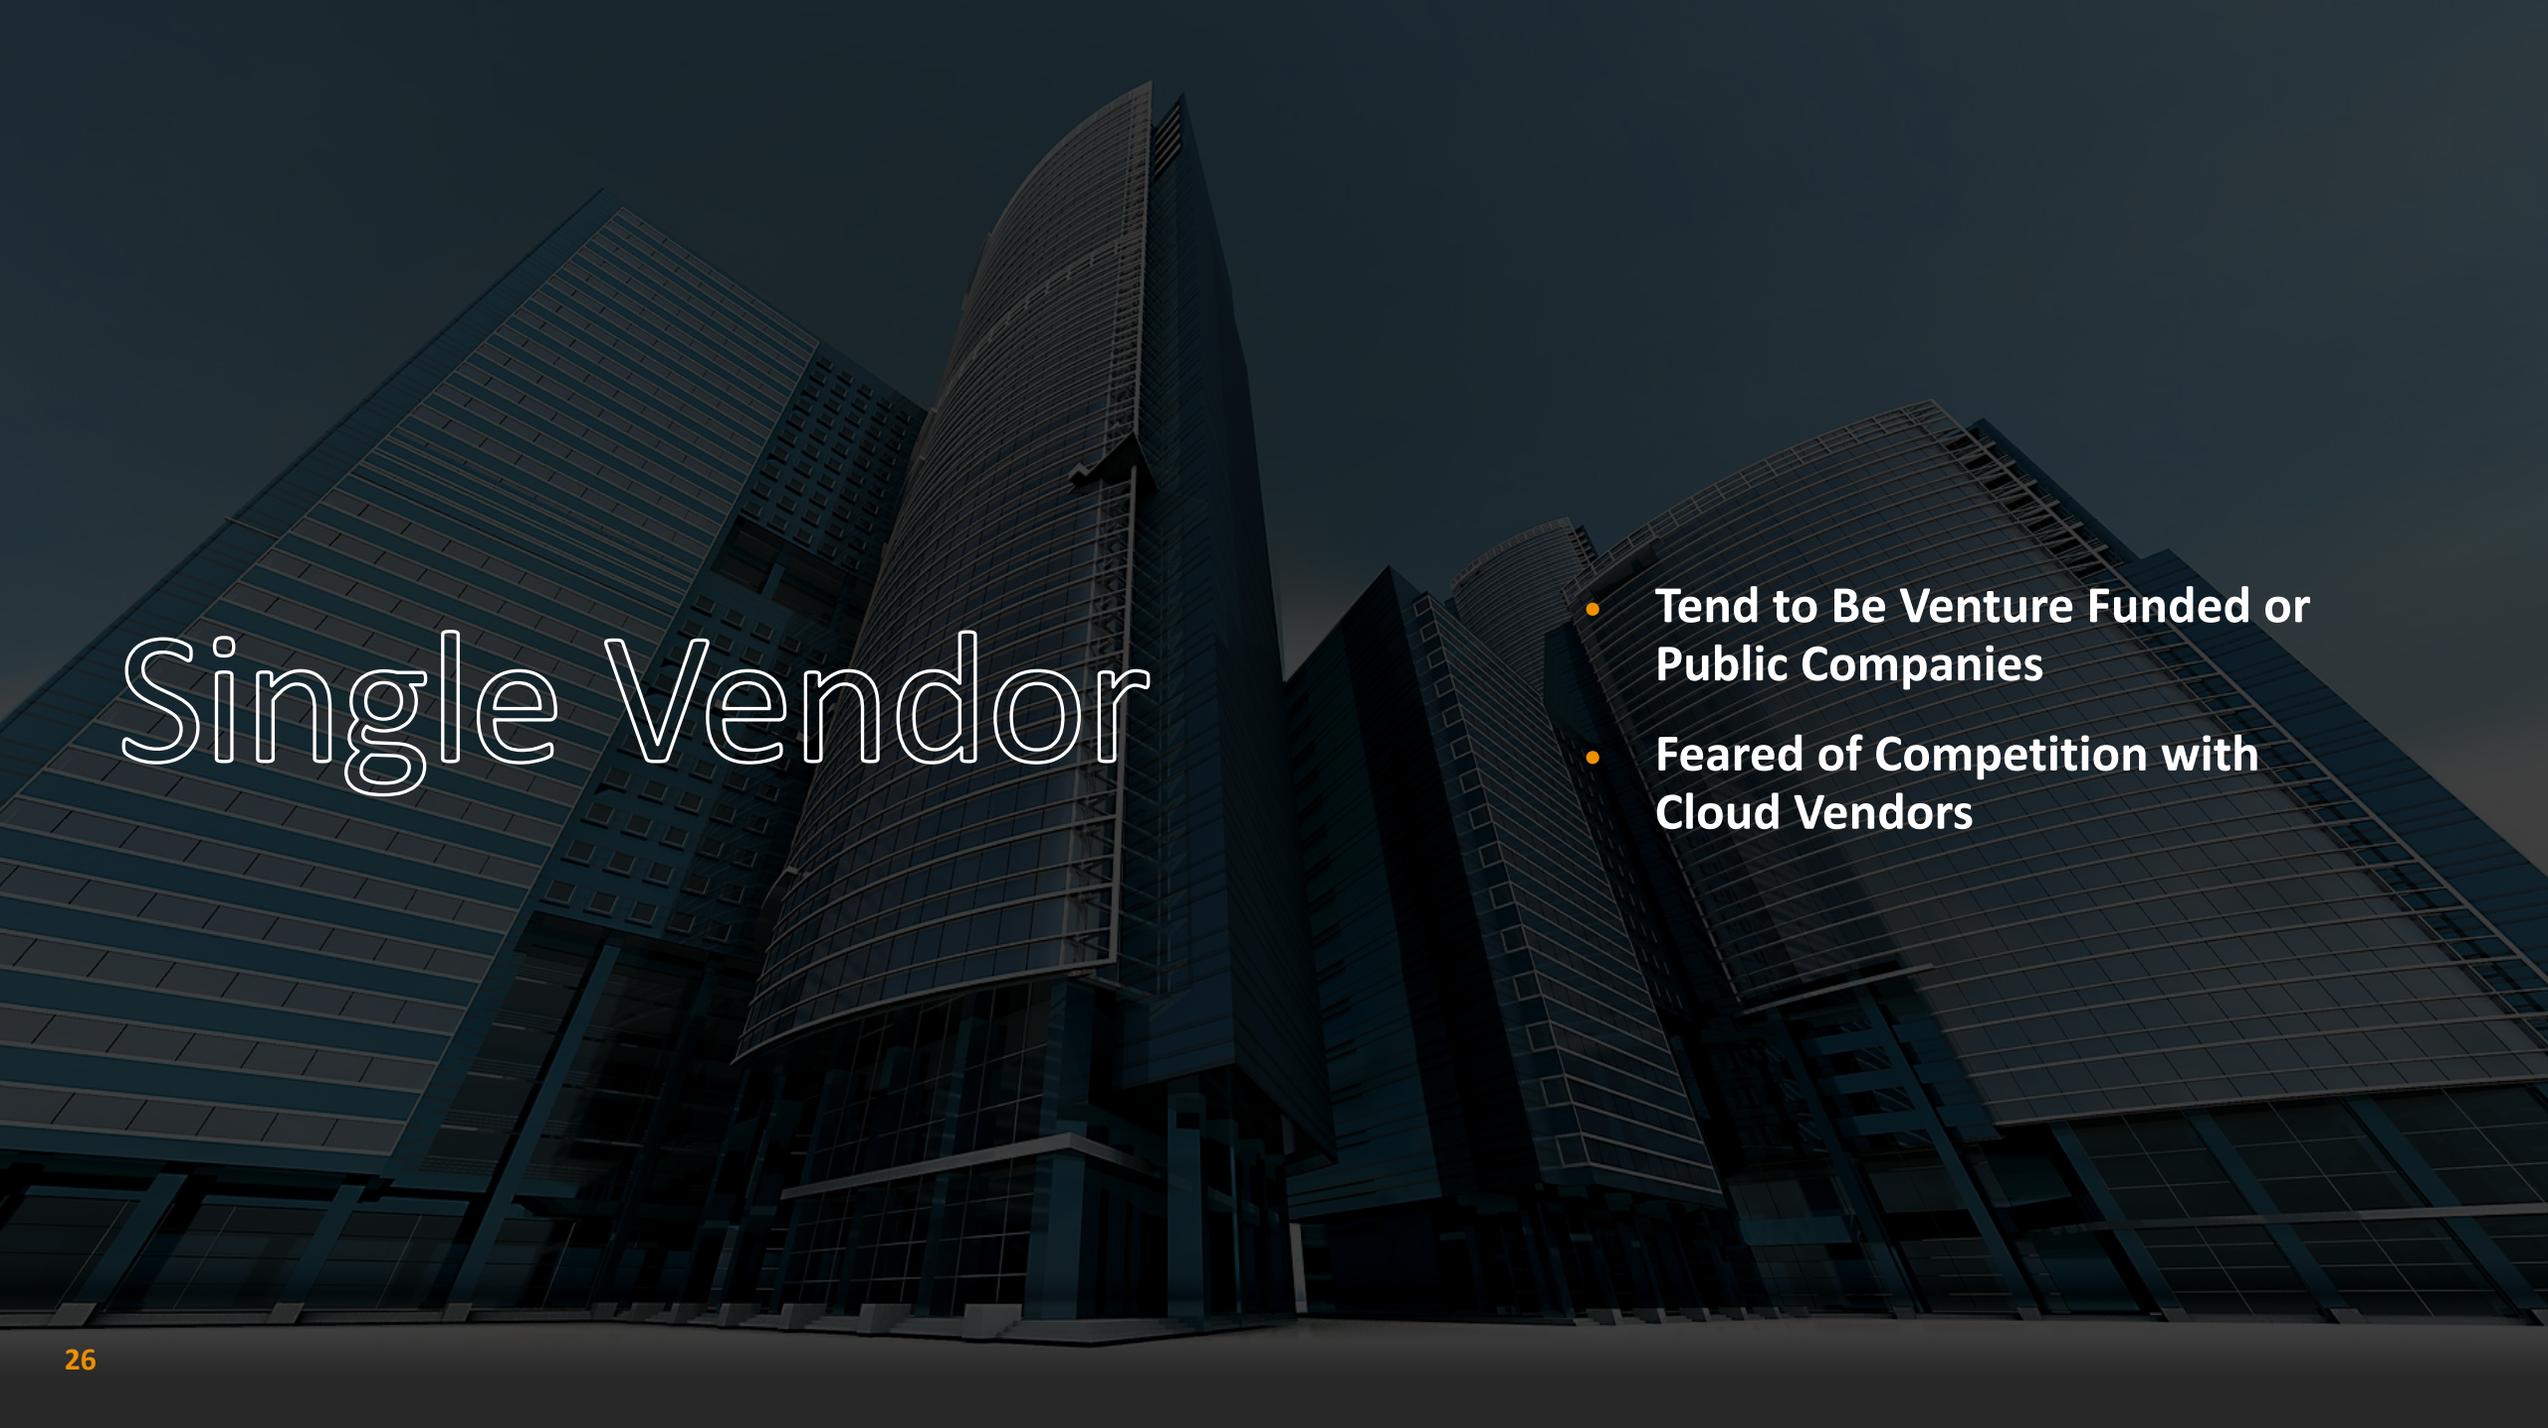


**CLOUD HELPS TO
ACCELERATE ADOPTION**



**CLOUD CHANGES WHO
CAPTURES THE VALUE**

Foundation Based Open Source



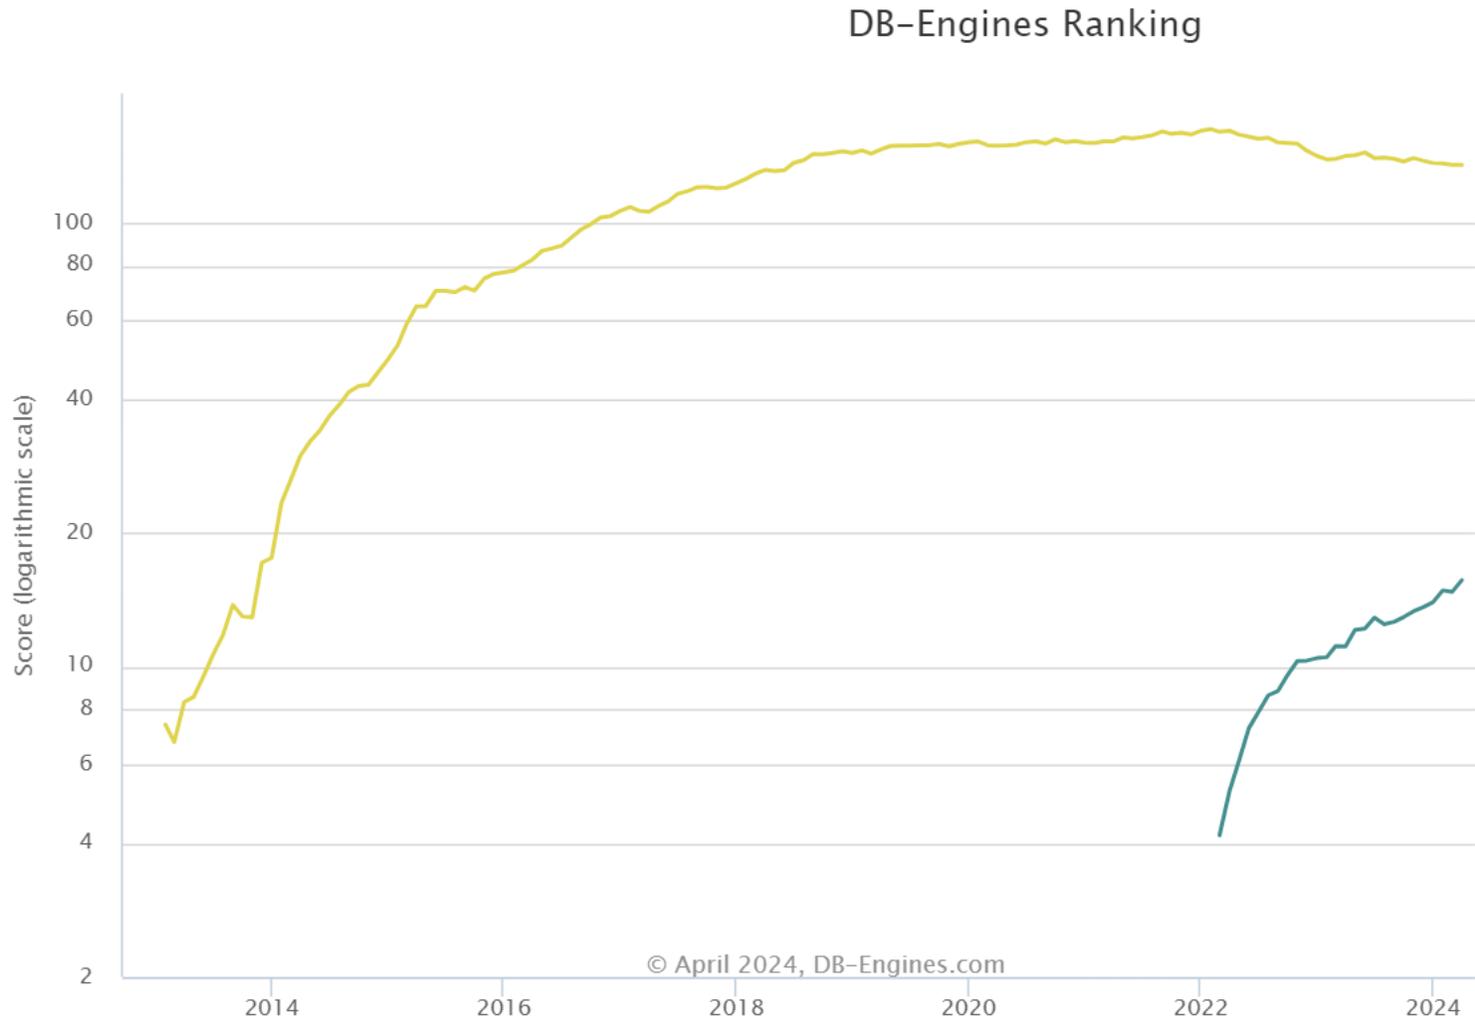
Single Vendor

- **Tend to Be Venture Funded or Public Companies**
- **Feared of Competition with Cloud Vendors**

Fully or Partially
Abandoning Open
Source Licenses

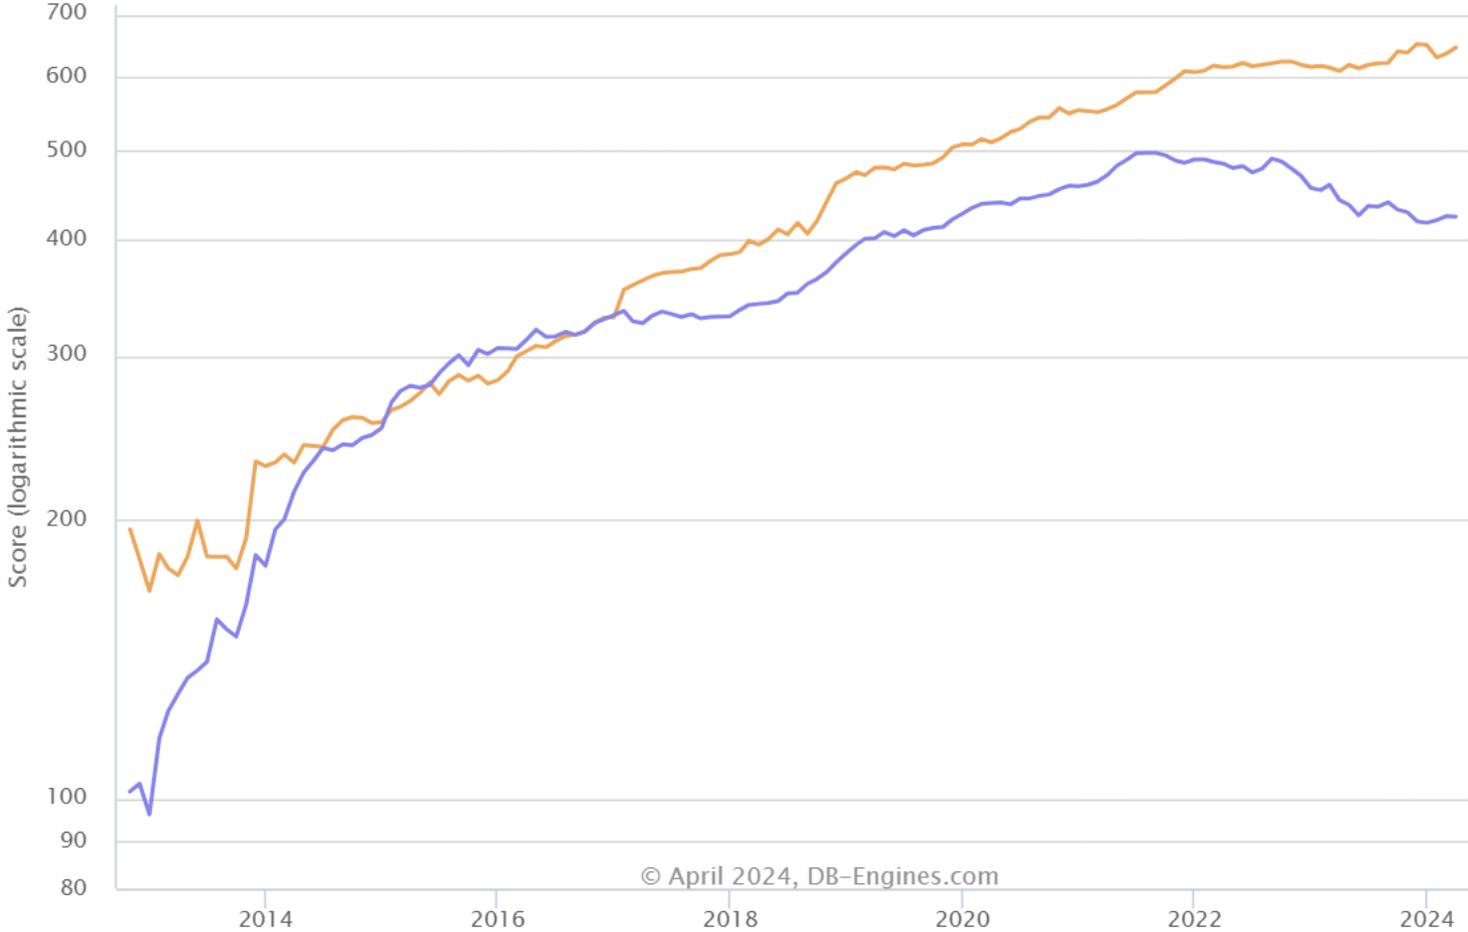


Elastic vs OpenSearch after Licence Change



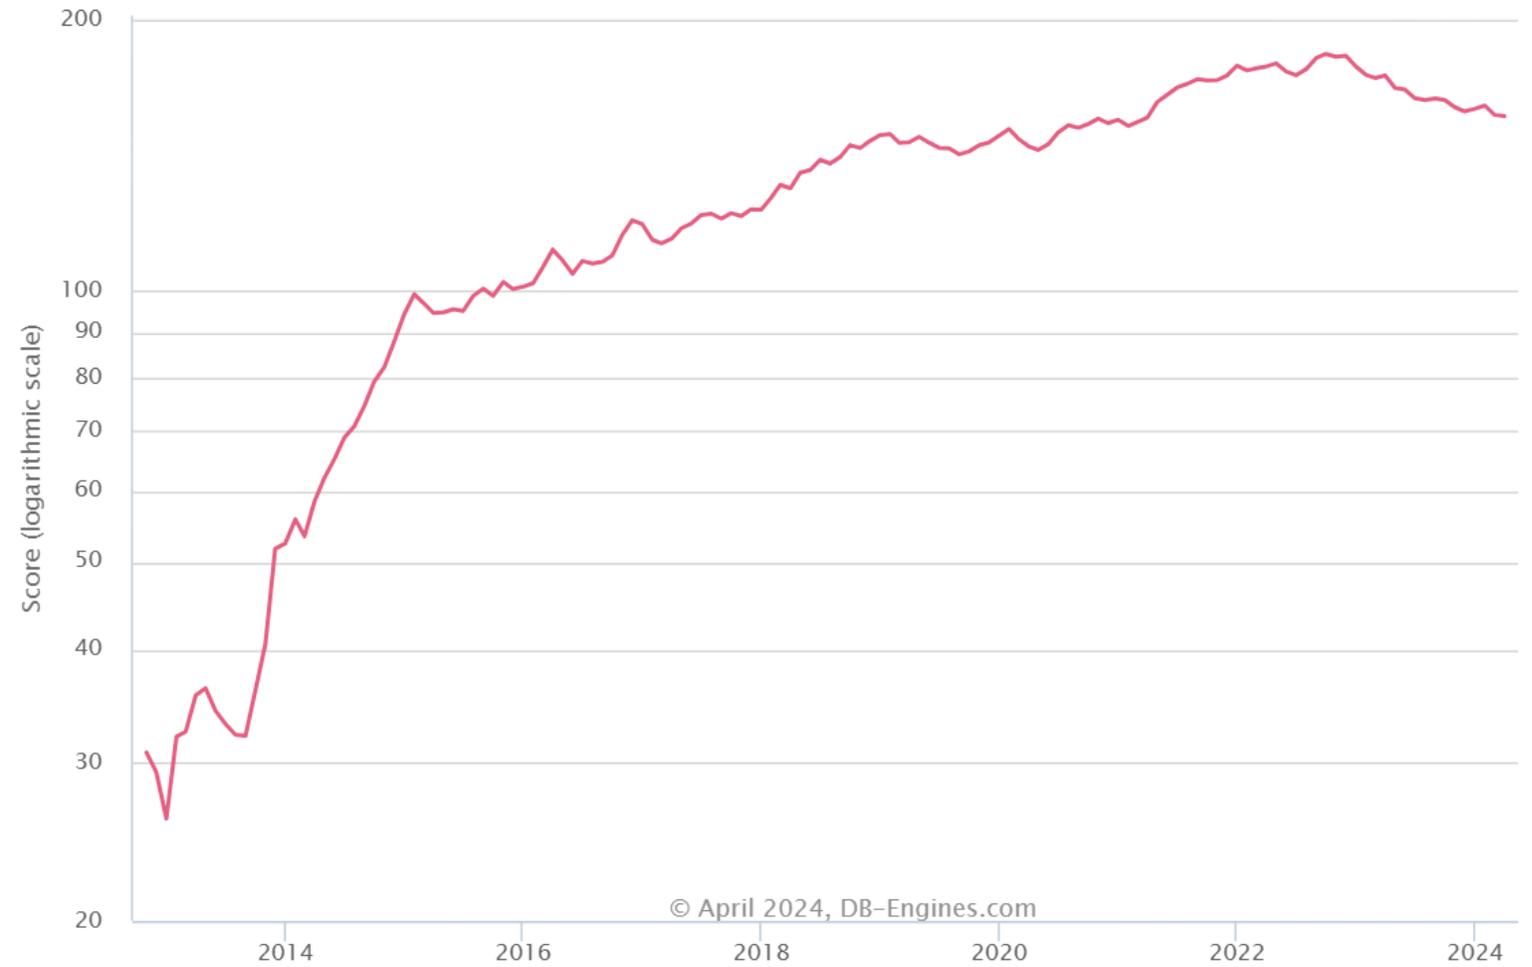
MongoDB vs PostgreSQL

DB-Engines Ranking



Redis

DB-Engines Ranking



Linux Foundation is Stepping Up!

 SIGN IN / UP

The Register®

DATABASES

Linux Foundation marshals support for open source alternative to Redis

7 

Follows the vendor's decision to overhaul licensing of the popular cache database

 [Lindsay Clark](#)

Wed 3 Apr 2024 // 11:15 UTC



Cloud giants AWS, Google, and Oracle have come out in support of a Linux Foundation open source fork of Redis, the popular in-memory database frequently used as a cache, following changes to its licensing.

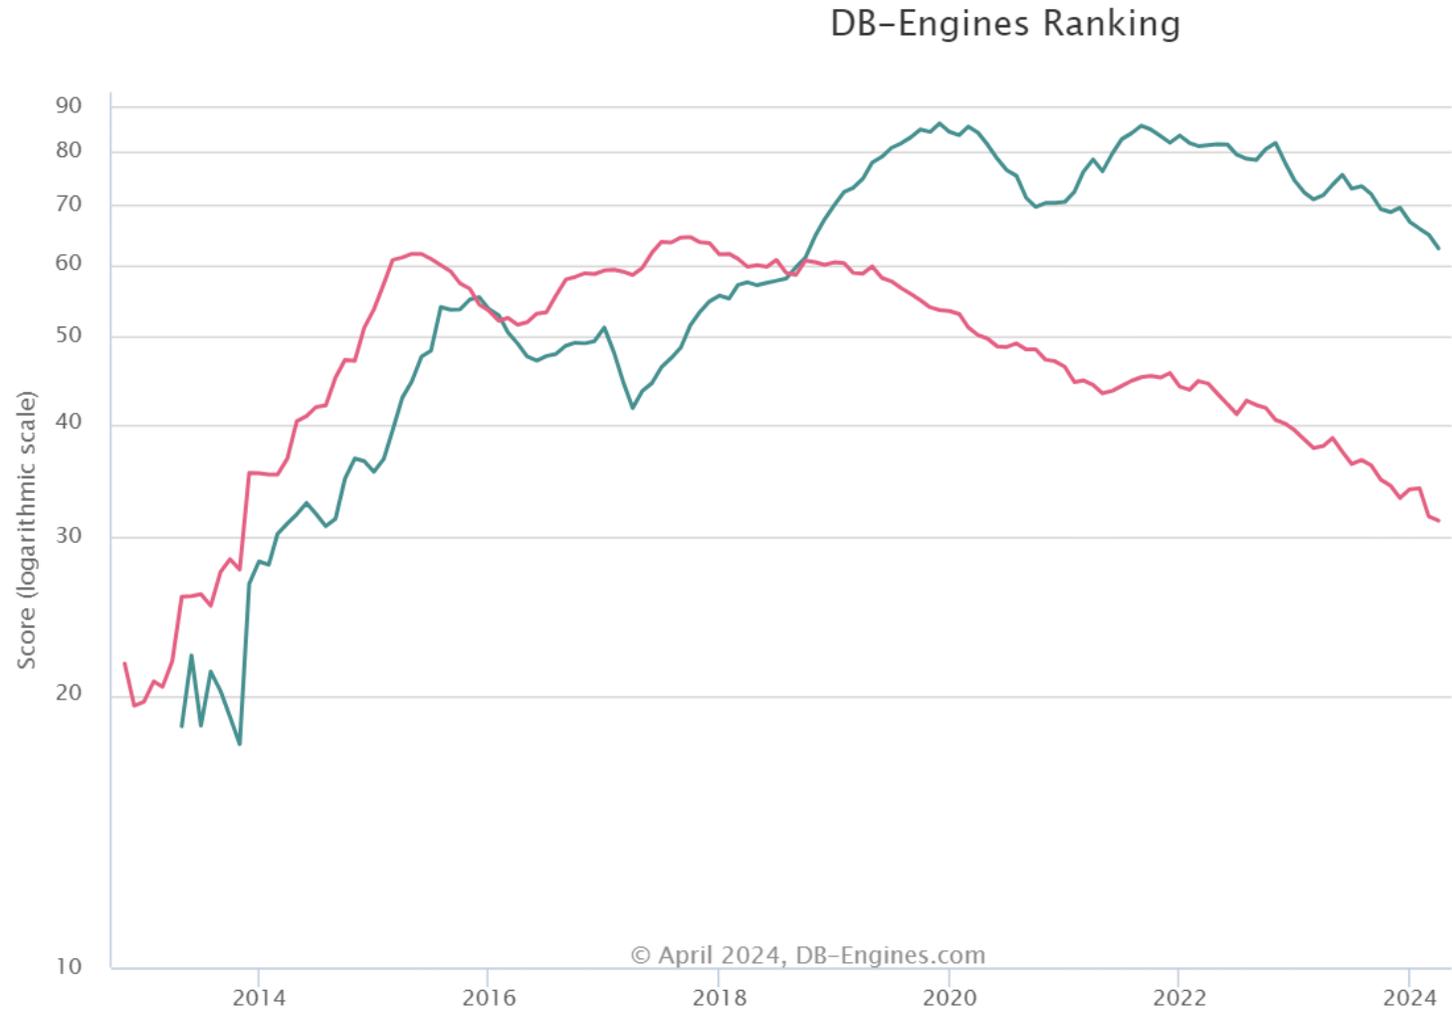
Last month, [Redis confirmed it was shifting its main key-value store system to a dual-license approach](#), imposing far more restrictive terms. Previously, the source code was available under the Berkeley Software Distribution (BSD) 3-clause license, which allows developers to make commercial use of the code without paying.

https://www.theregister.com/2024/04/03/open_source_redis_alternative/

Percona Joins the Effort!



Hive (Blue) and HBase





Primary Goal of the License Change?

- **Creating Monopoly on DBaaS Market**



STATE OF ART SIMPLICITY



**HIGH LEVEL OF
AUTOMATION**

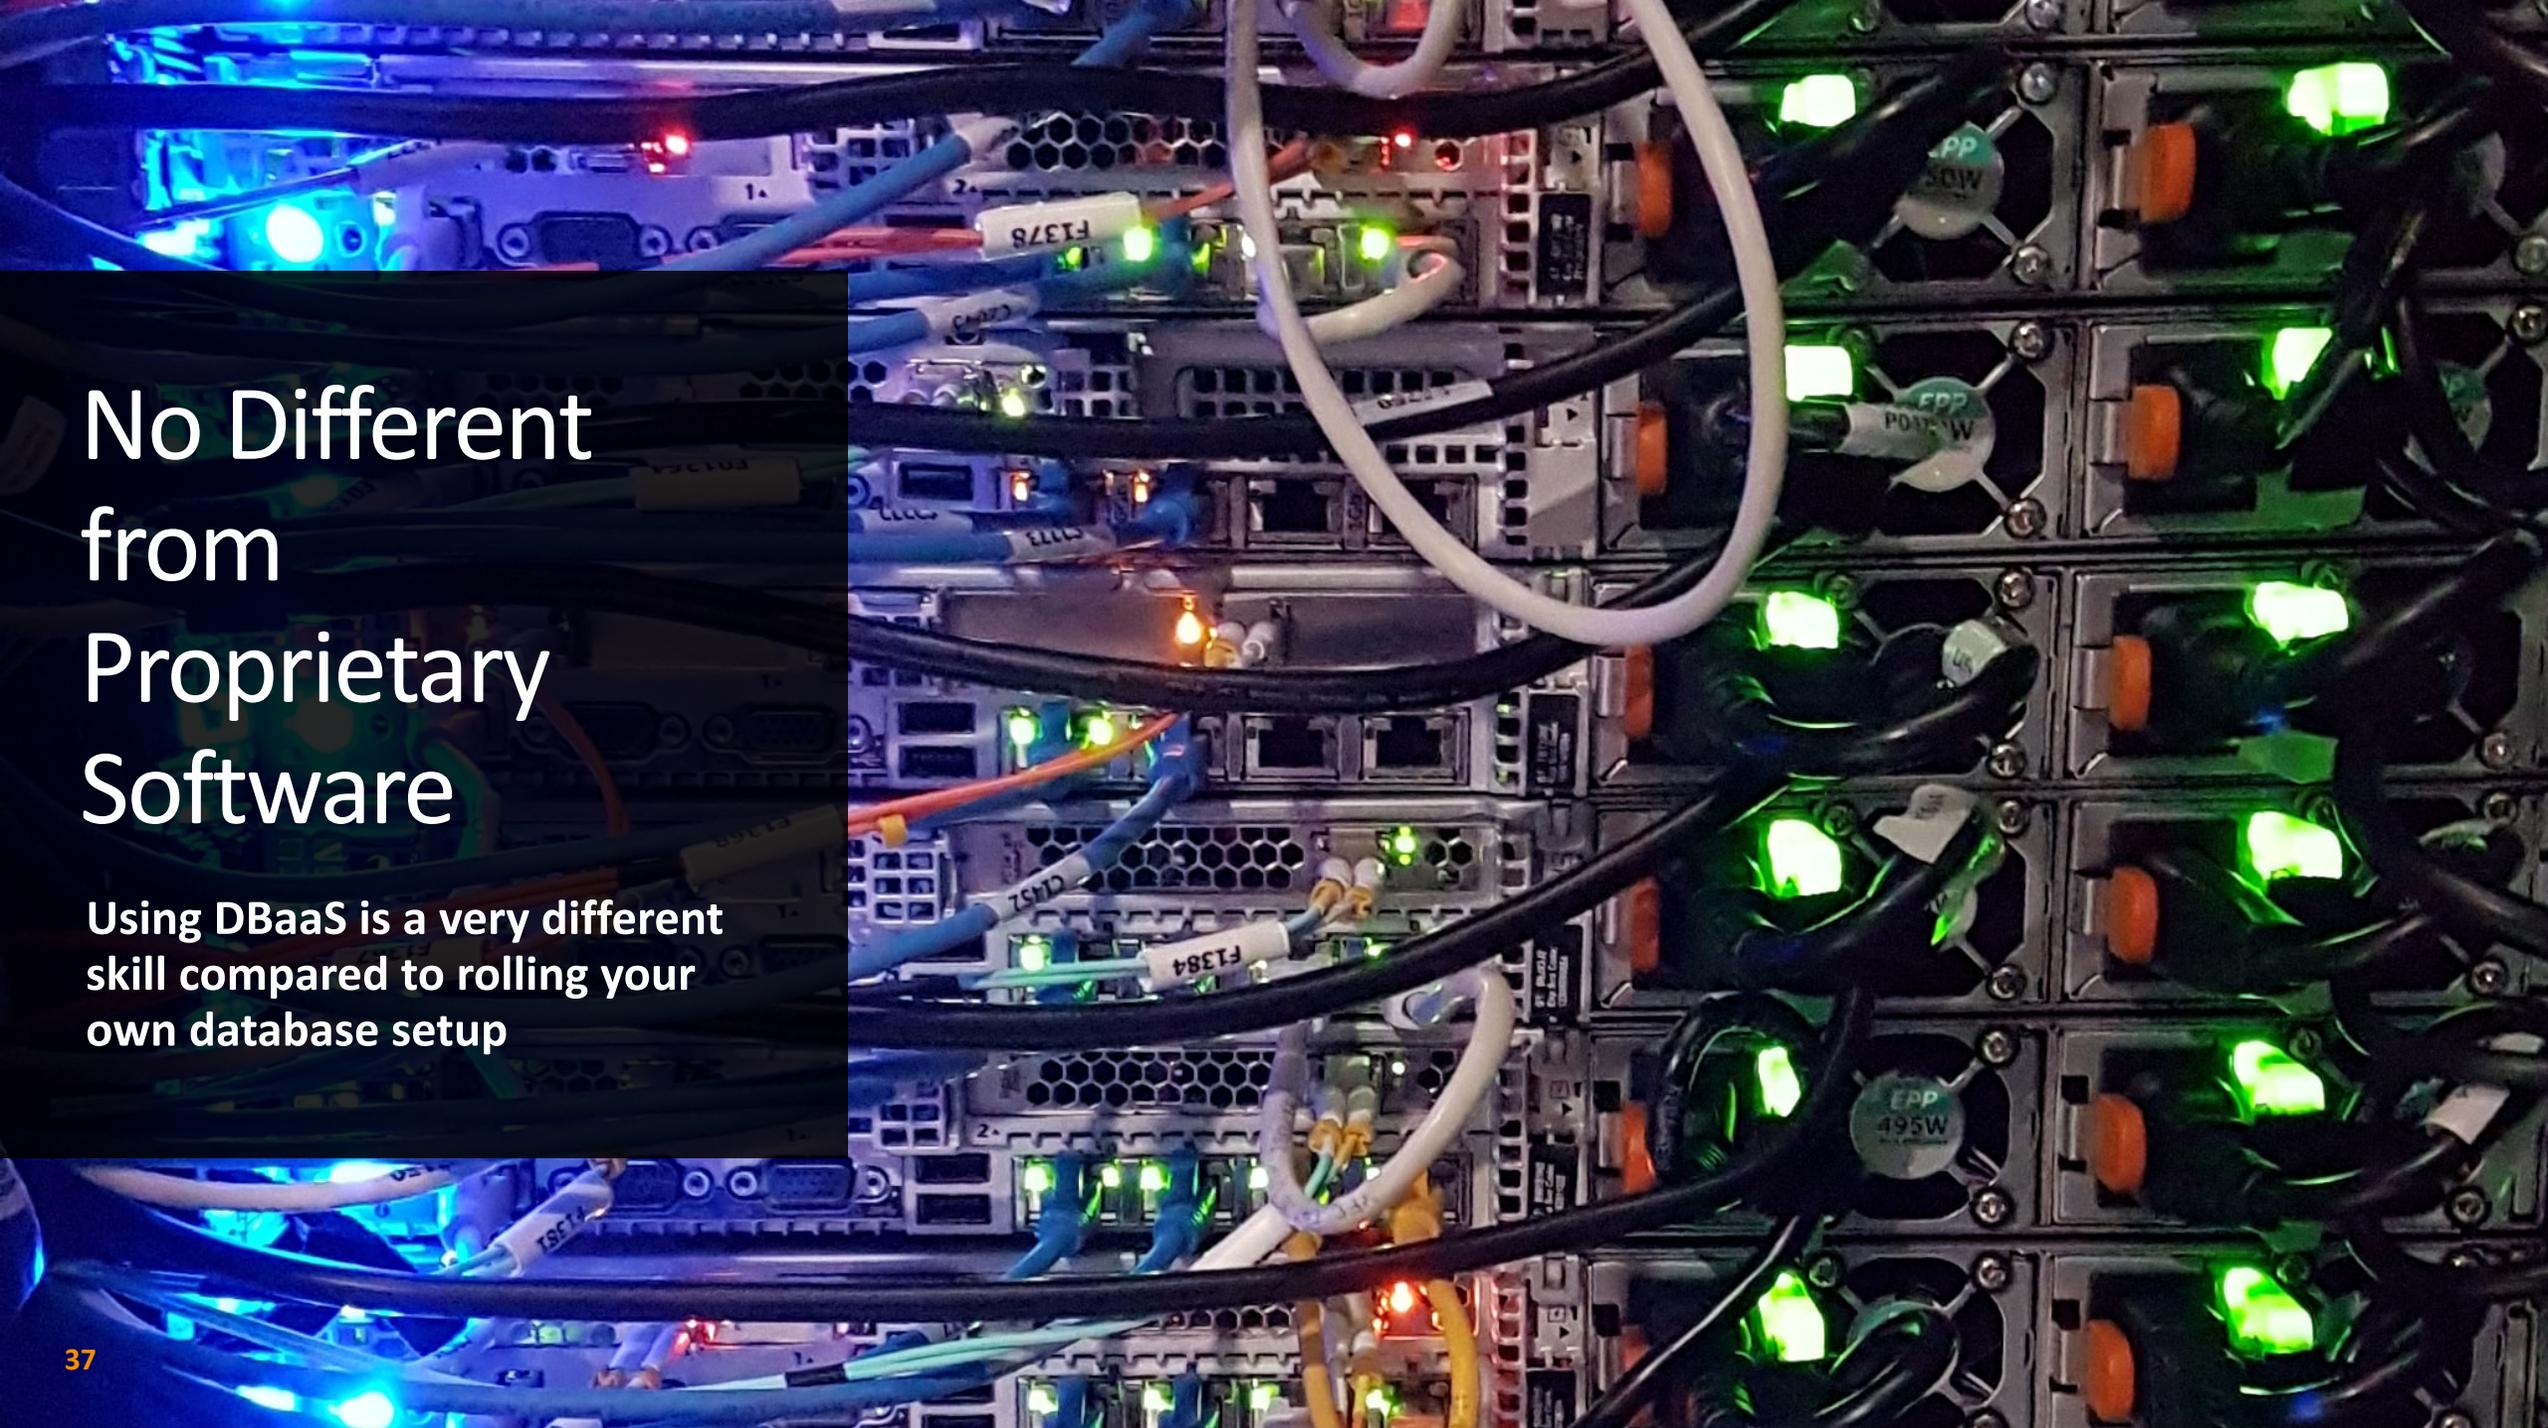


**MAXIMIZES DEVELOPERS
FOCUS ON APPLICATION**

Why DBaaS?

What Is the
Problem with
Monopoly on
DBaaS?

MONOPOLY



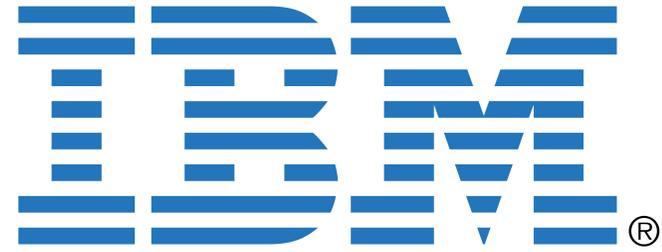
No Different from Proprietary Software

Using DBaaS is a very different skill compared to rolling your own database setup



DBaaS Vendor Lock-In

It may not be so painful **now**, but it is going to be painful in the **future**



History Lesson

Oracle used to Save its customer from IBM
Hardware Lock-in with Mainframe computers



Why Data is Special

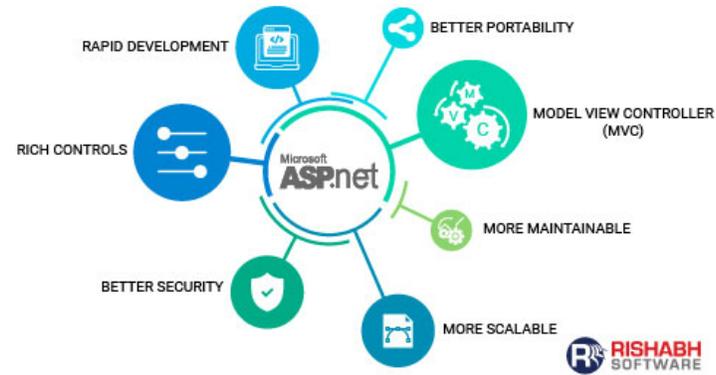
Moving Data is Expensive. Moving Lots of Data is Very Expensive – Principle of Data Gravity Applies

Not All Is Lost



Have We Been Here Before?

2000s



2020s



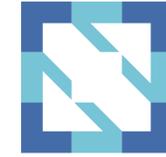
Operating Systems



Open Source Catches Up Again



- Lock-in with Cloud Vendor
- Use Proprietary Solutions
- Highly Differentiated Cloud
- Hostage
- No Vendor Choice

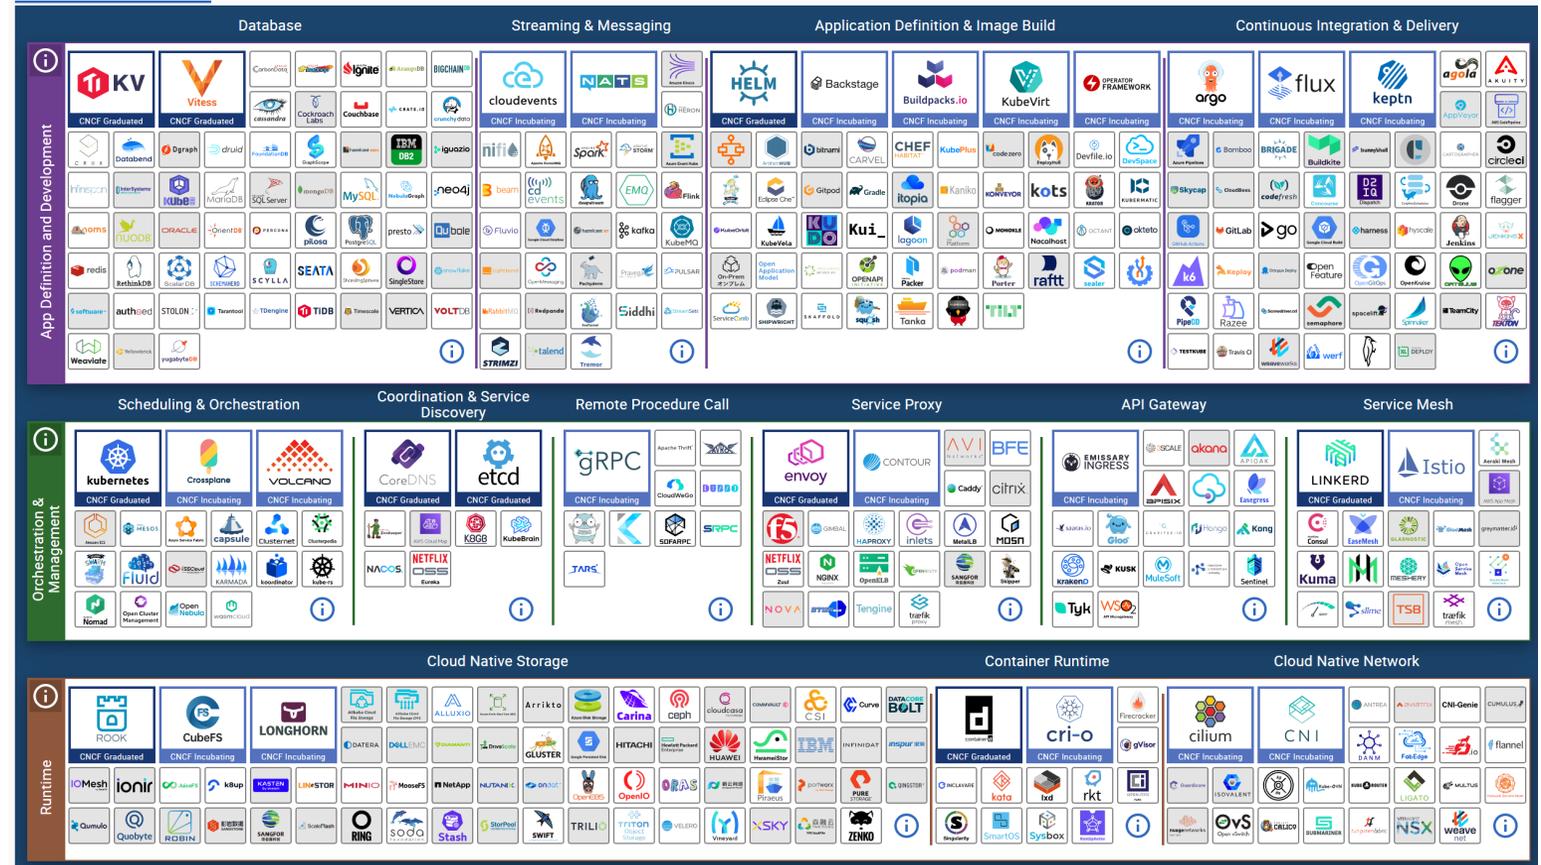


**CLOUD NATIVE
COMPUTING FOUNDATION**

- Freedom to Run Anywhere
- Use Open Source
- Cloud Is Commodity
- Customer
- Choice of Vendors



Cloud Native is Going Strong



<https://landscape.cncf.io/>

Giving Cloud Its Originally Intended Role of Commodity Infrastructure

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.



Kubernetes

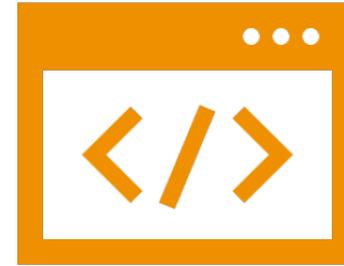


- **Kubernetes is universally available**
- **Kubernetes is getting better for stateful applications**
- **Kubernetes Operators are available for most popular Open Source Databases**

kubernetes



**Day 1 and Day 2 Automation,
Toil Reduction Similar to DBaaS**



**UX is Different, Requires
Kubernetes Expertise**

What's Up with Kubernetes Operators?

Can We Build DBaaS on Kubernetes?

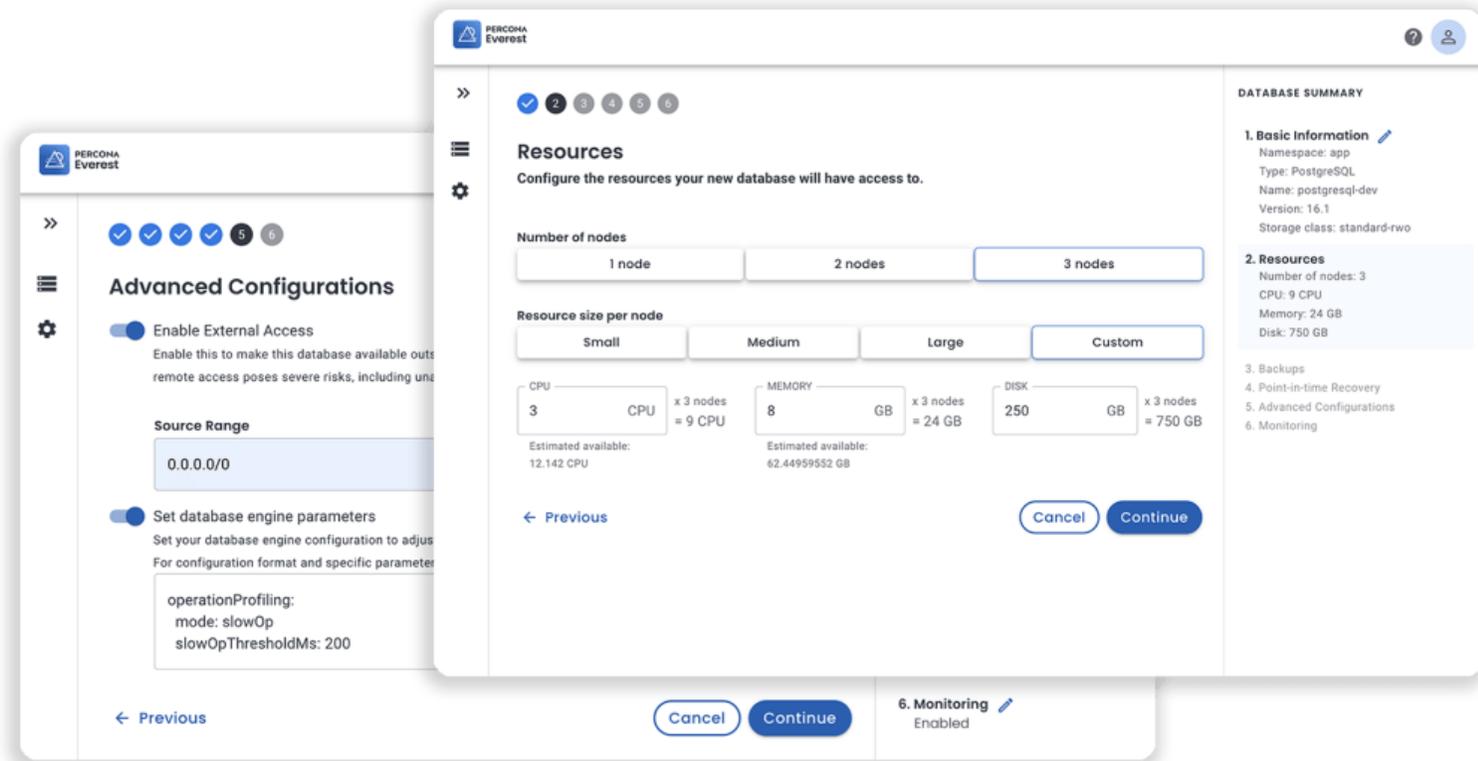


kubernetes

Many Modern
DBaaS Are Built on
Kubernetes
Operators



Open Source DBaaS Experience with Percona Everest



Summary



**Fantastic Pace
Innovation**



**Evolving Open
Source Models**



Great Future

Thank you, Let's Connect!

<https://www.linkedin.com/in/peterzaitsev/>

<https://twitter.com/PeterZaitsev>

<http://www.peterzaitsev.com>



Limassol, Cyprus 2024

PERCONA
UNIVERSITY