PERCONA

From Database in Container to DBaaS on Kubernetes

Peter Zaitsev,
Founder at Percona
June 19, 2024

stackconf

Early Days of Modern Open Source

```
Floppy drive(s): fd0 is 1.44M
floppy: FDC version 0x90
Swansea University Computer Society Net2Debugged [1.30]
IP Protocols: ICMP, UDP, TCP
NE*000 ethercard probe at 0x300: 52 54 00 12 34 57
eth0: NE2000 found at 0x300, using IRQ 9.
ne.c:v0.99-15k 3/3/94 Donald Becker (becker@super.org)
Checking 386/387 coupling... Ok, fpu using exception 16 error reporting.
Linux version 1.0 (root@softland) #1 Tue Apr 26 19:25:39 PDT 1994
Partition check:
  hda: hda1 hda2
EXT2-fs warning: mounting unchecked fs, running e2fsck is recommended
VFS: Mounted root (ext2 filesystem).
Cannot open *.o
[: argument expected
none on /proc type proc (rw)
/etc/rc.net: sls105(192.168.1.100), rpc.portmap, inetd, namedJul 22 13:57:32 nam
ed[45]: restarted

, rpc.nfsd, rpc.mountd


Welcome to Linux  0.99.15g

sls105 login:
```


PERCONA

Never-Ending Move towards Simplicity

PERCONA

Installation of Open Source Software

- **Download Sources, Patch and Compile**
- **Tar.gz binaries and INSTALL**
- **Packages with Dependencies – .deb .rpm**
- **Package repositories APT and YUM**
- **Docker, Snap etc.**

Database in Docker

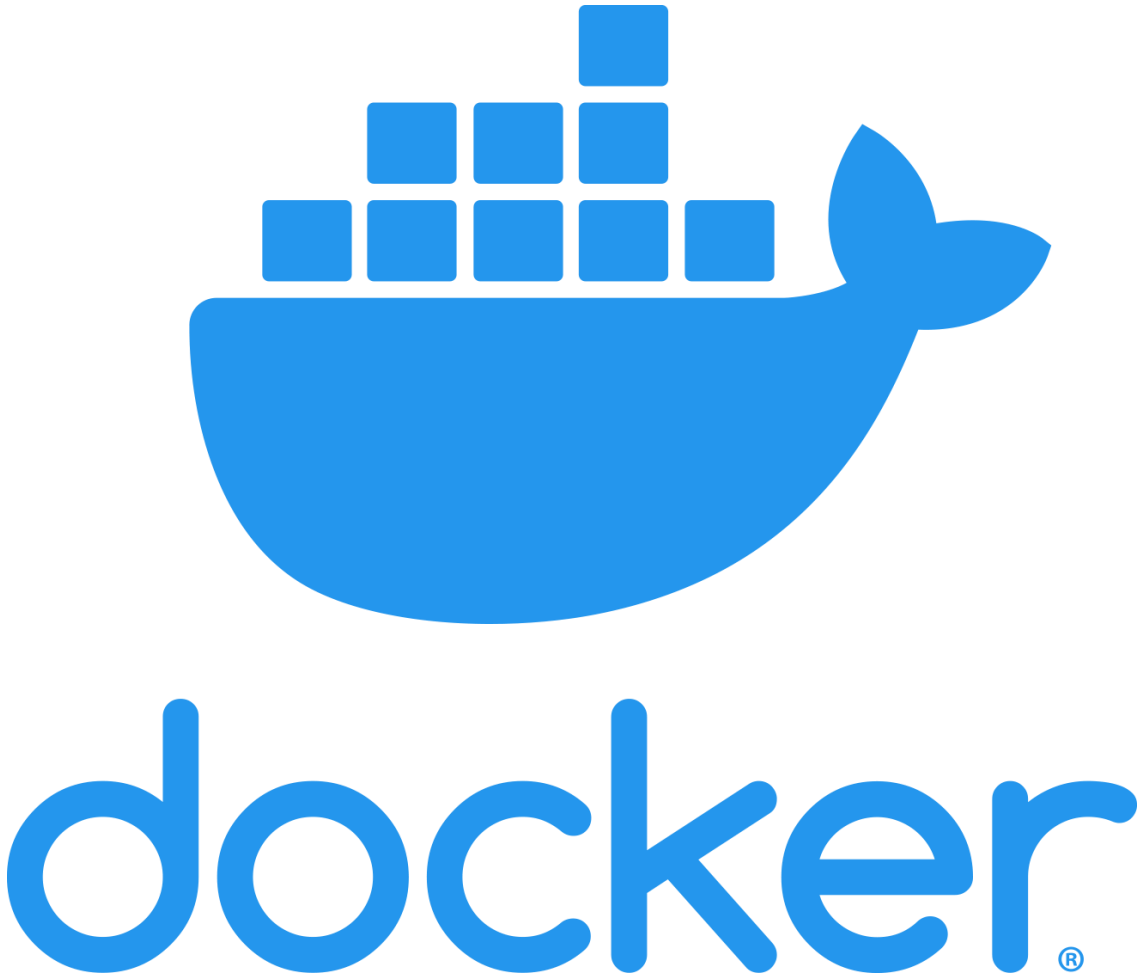

PERCONA

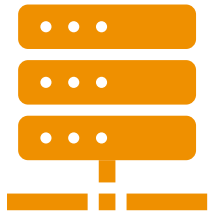

Clean. Isolated from Other Components in the Environment

Easy to have Multiple Environments

Simplify Deployment with Docker Compose

Test and Dev

Production

(Mostly Unfounded) Fears of overhead

Extra Complexity. Need to use Data Volumes for Best Result

Some monitoring tools initially lacked proper Docker Support

PERCONA

Most Open Source Databases have official Docker Images

Commonly Deployed for Test and Dev

Limited Deployment in Production

Open Source Databases with Docker

Percona Solution

We Provide Docker Packages for MySQL, MongoDB, PostgreSQL

Enhanced Enterprise Grade Distribution

100% Free and Open Source

PERCONA

Unsolved Problem

What about Day 2 Operations

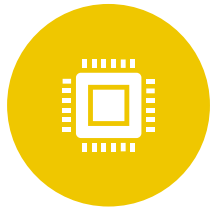

Upgrades, High Availability, Scaling Up and Down

Must be done in the context of "Cluster" rather than single node

Docker does not provide great Solution

Kubernetes

PERCONA

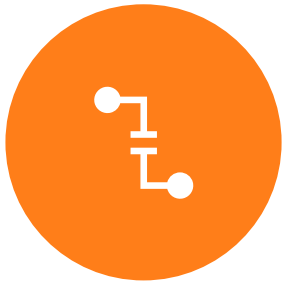

Complicated Relationship…

Kubernetes was designed for Stateless application first

… and Database is very opposite of Stateless

Improved to manage Stateful Applications in Recent Years

Kubernetes and Databases

Stateful Applications and Kubernetes

Kelsey Hightower
@kelseyhightower

I'm always going to recommend people exercise extreme caution when running stateful workloads on Kubernetes. Most people who are asking "can I run stateful workloads on Kubernetes" don't have much experience with Kubernetes and often times the workload they are asking about.

3:10 AM · Mar 24, 2019 · Twitter Web Client

317 Retweets **903** Likes

PERCONA

DOK Community

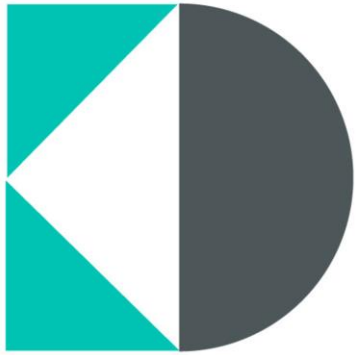

Dok Community

https://dok.community/

PERCONA

Data on Kubernetes

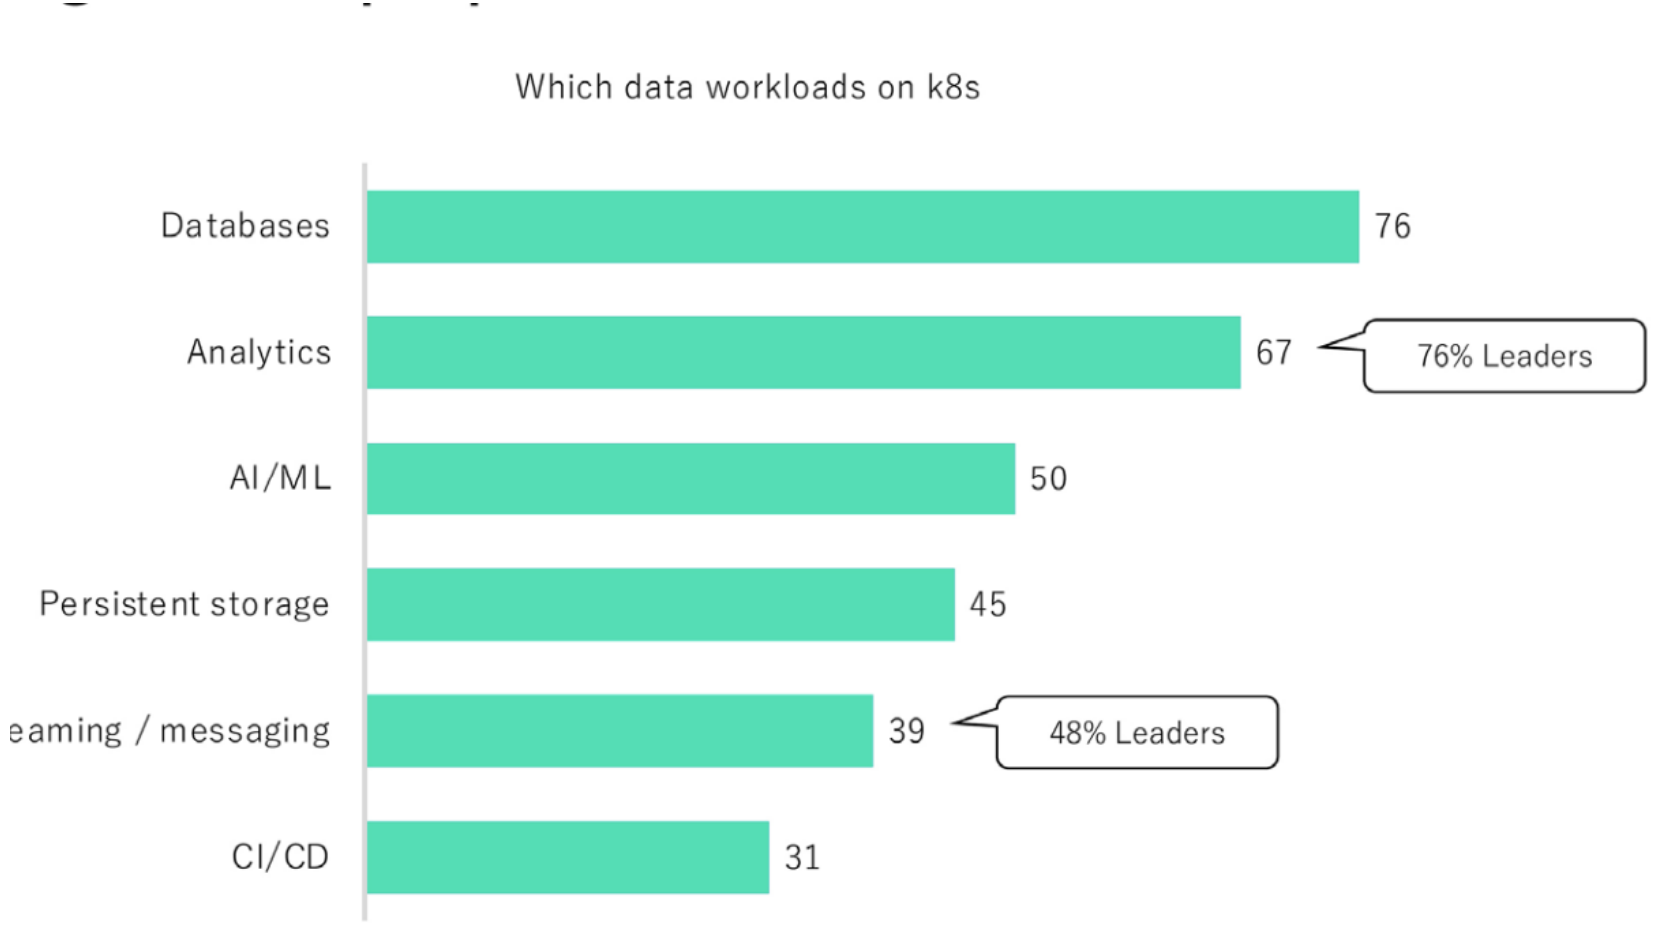

https://dok.community/wp-content/uploads/2022/10/DoK_Report_2022.pdf

PERCONA

Satisfied

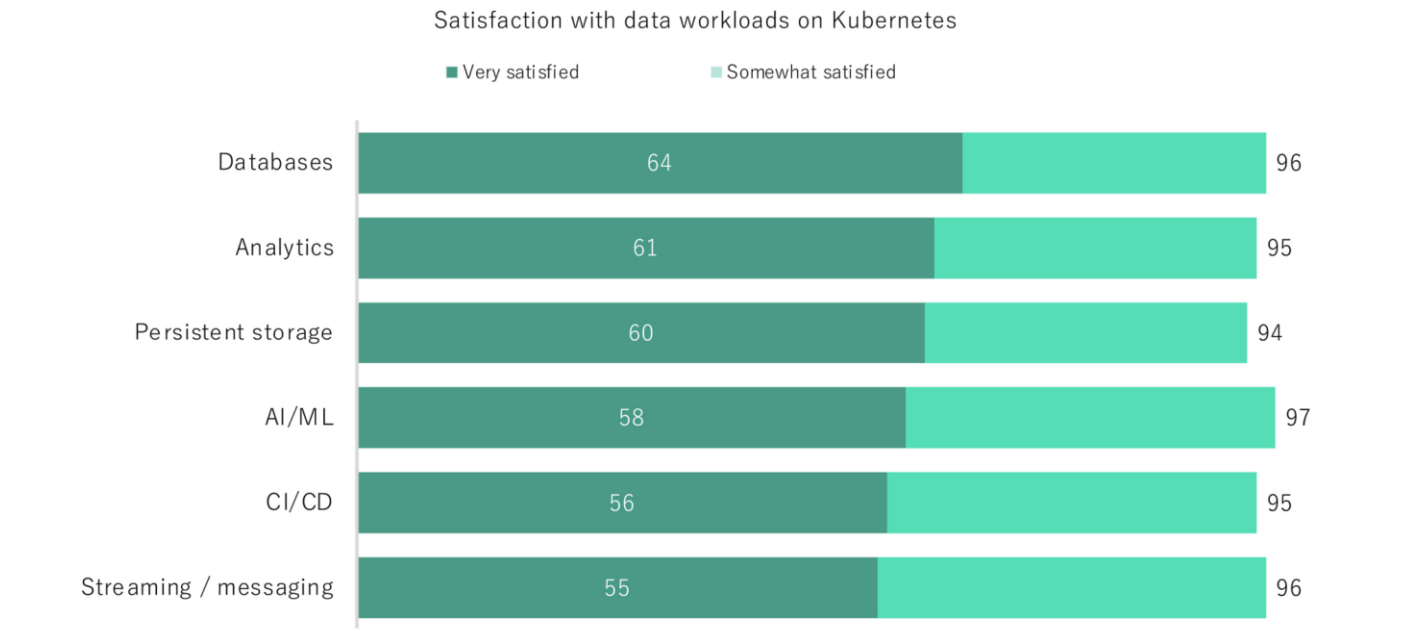

1. In general, how satisfied are you with using Kubernetes to run each of the following data workloads in your organization? Use a scale from 1 to 5 where 5 means "very satisfied" and 1 means "not at all

PERCONA

Adoption Areas

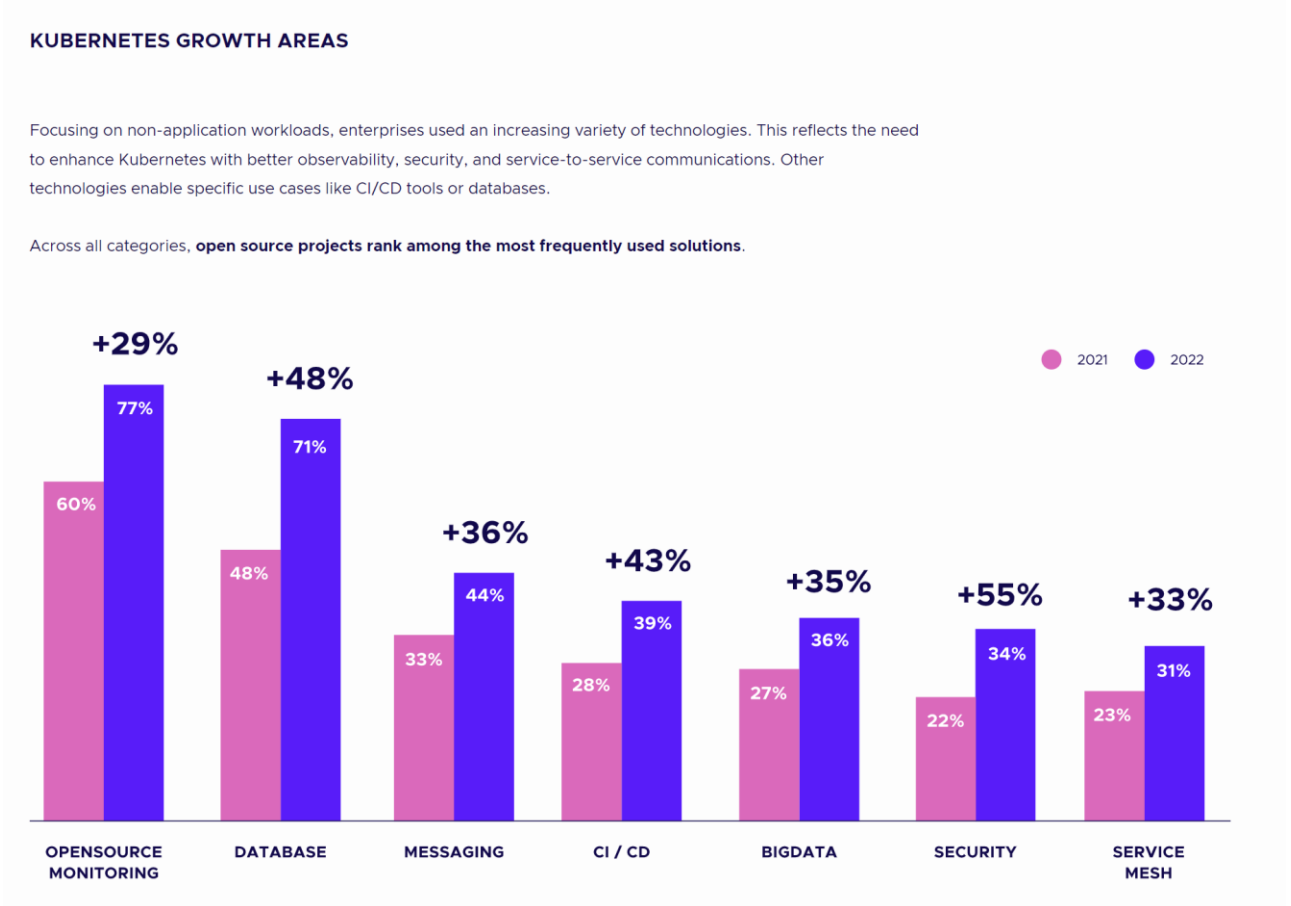

https://www.cncf.io/reports/cncf-annual-survey-2022/

PERCONA

Kubernetes Powers many DBaaS

Altinity

influxdata

MariaDB

crunchy data

planetscale

BIGANIMAL
Powered by EDB

PERCONA

Promise of Kubernetes

Operating System for your Data Center, like Linux was for a single server

Robust mechanics to deal with node failure

Operator Framework for Automating Complex Database Operations Tasks

PERCONA

Slower Pickup by vendors

Many Third Party Kubernetes Solutions Were Developed

“Helm Chart” and “Operator” packages

Open Source Databases on Kubernetes

Percona Solutions

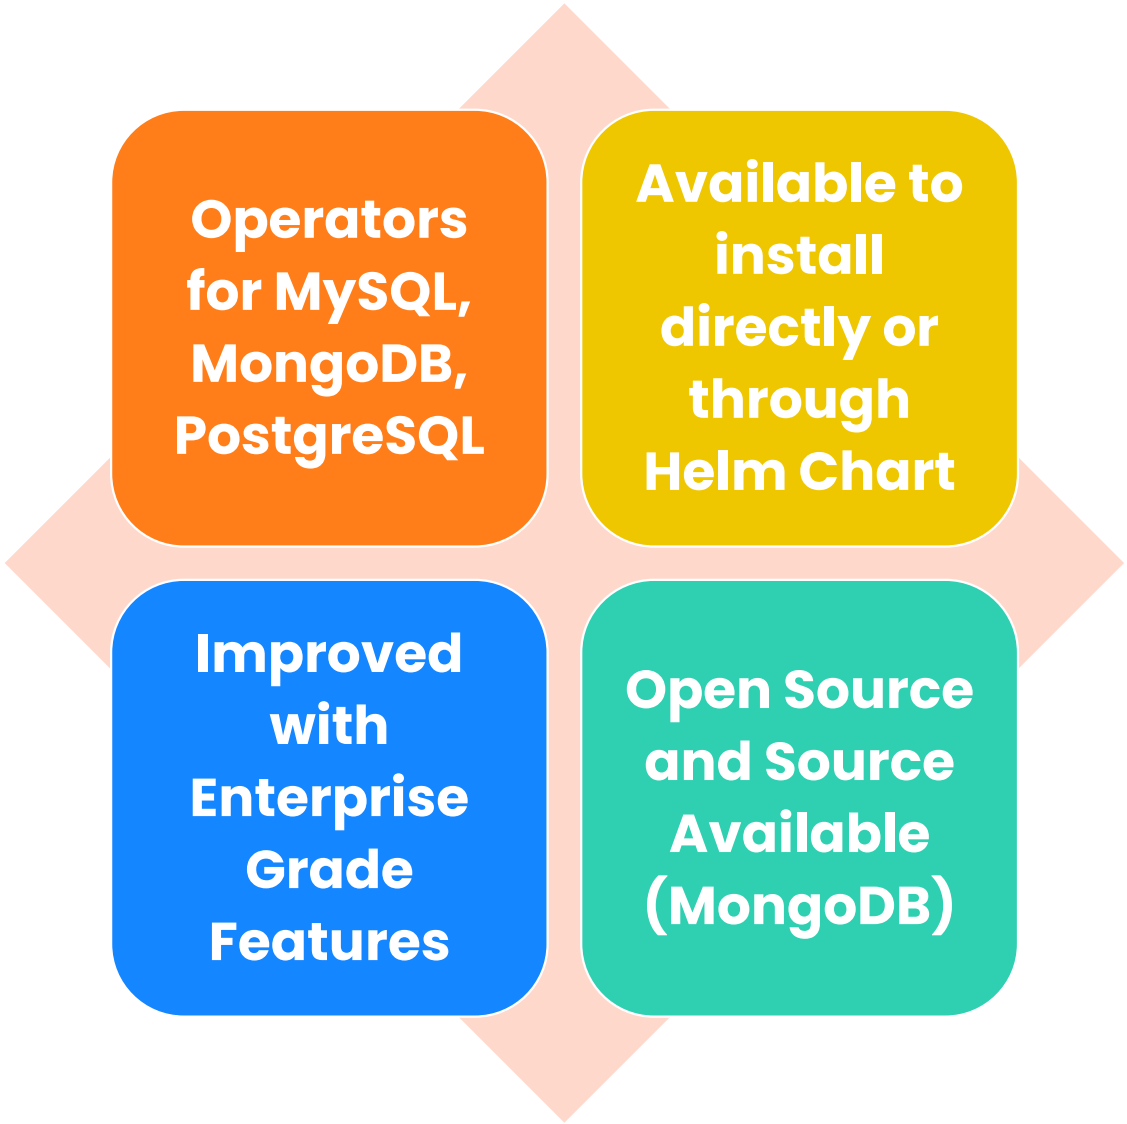

PERCONA

Mission Critical: Open Source Base Database Software

PERCONA Distribution for MySQL

PERCONA Distribution for PostgreSQL

PERCONA Distribution for MongoDB

PERCONA

Running Business Critical Statefull applications on Kubernetes is “Black Belt” level skill

Not an easy task for non-Kubernetes expert

Unsolved Problems

DBaaS “state of art” simplicity for databases

PERCONA

Cloud

Proprietary Clouds bring Great Usability at Great Cost

PERCONA

Current State

Major Clouds have Proprietary DBaaS Offerings

- Amazon RDS, Aurora, Google Cloud SQL etc

Database Vendor Own Proprietary Solutions

- MongoDB Atlas, SkySQL, CockroachCloud etc

Multi Database Proprietary Solutions

- Aiven, Instaclustr etc

Managing High Availability

Database Patching

Backups

(Some) Performance Tuning

Easy to Scale

DBaaS Removes a lot of “Toil”

Do not worry, DBaaS is Open Source Compatible

Limited, "Hotel California" Compatibility

DBaaS

PROMISE

REALITY

Fully Managed DBaaS tends to be Overmarketed

PERCONA

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

How about Going Back to Basics?

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.

amazon web services™

PERCONA

Kubernetes – Universal API for Public and Private Cloud

PERCONA

Cloud Choice

- **Lock-in with Cloud Vendor**
- **Use Proprietary Solutions**
- **Highly Differentiated Cloud**
-
-

CLOUD NATIVE
COMPUTING FOUNDATION

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

PERCONA

Cloud Native is Going Strong

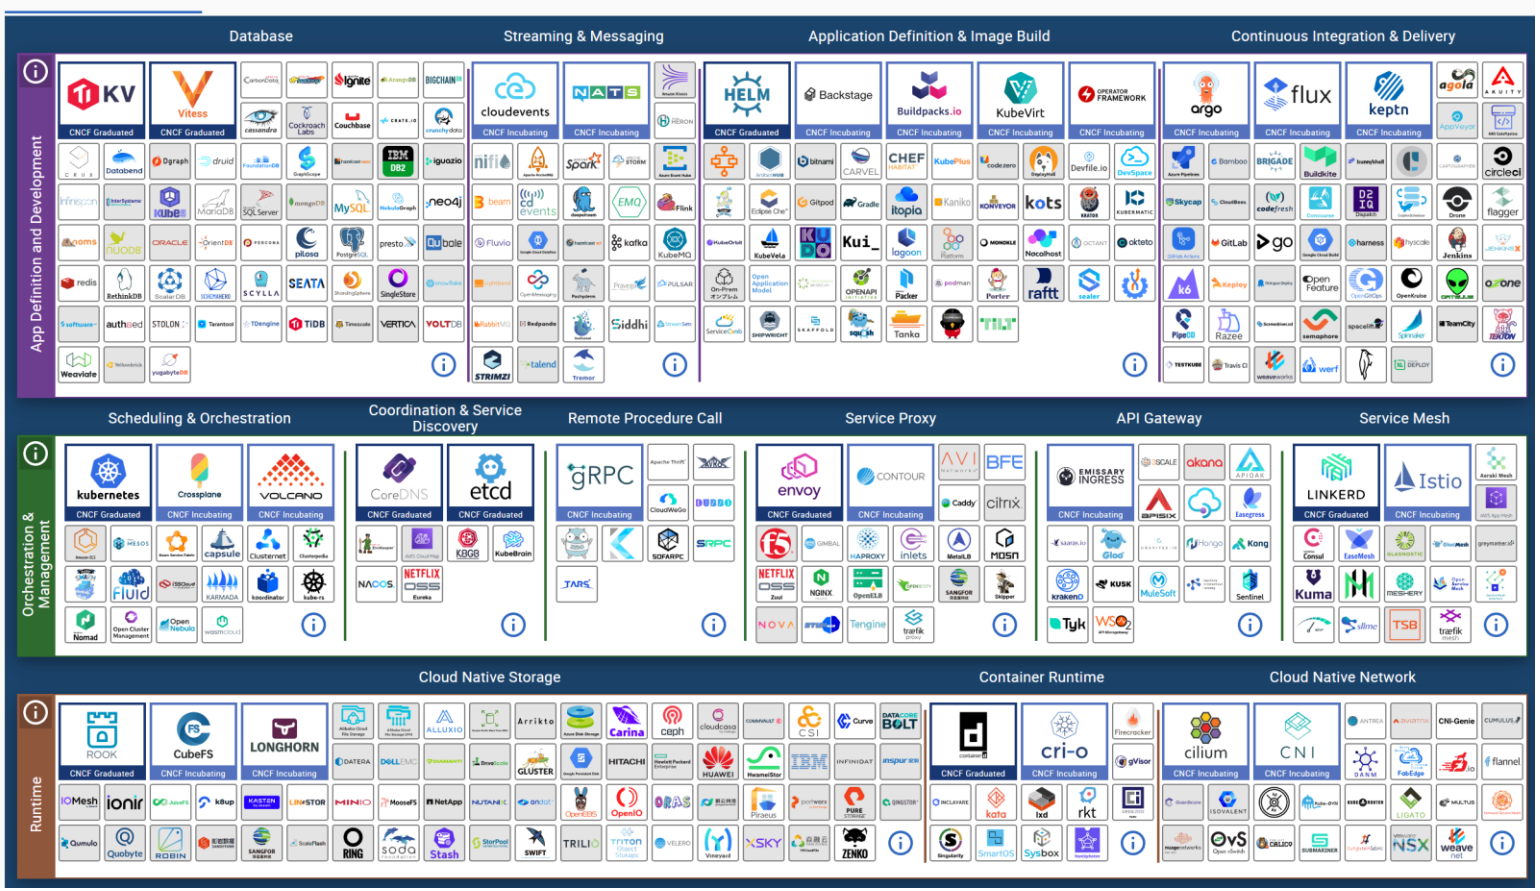

https://landscape.cncf.io/

PERCONA

Our Vision

Provide AWS RDS-Like DBaaS Experience in Public and Private with Open Source Software and No Vendor Lock-in

PERCONA

Get look and feel for Basics?

Kubernetes with Operators is Easy and Powerful

https://per.co.na/PXCMinkube

“Operators” to run them on Kubernetes

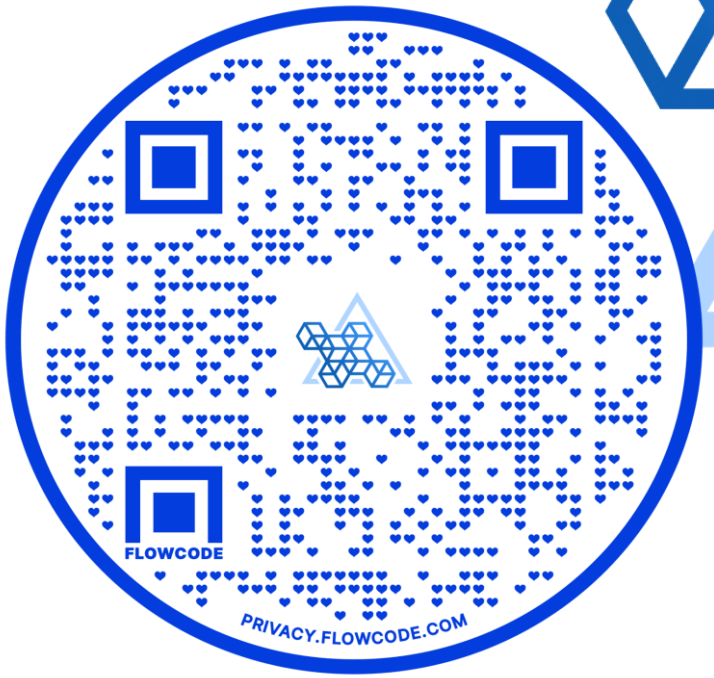

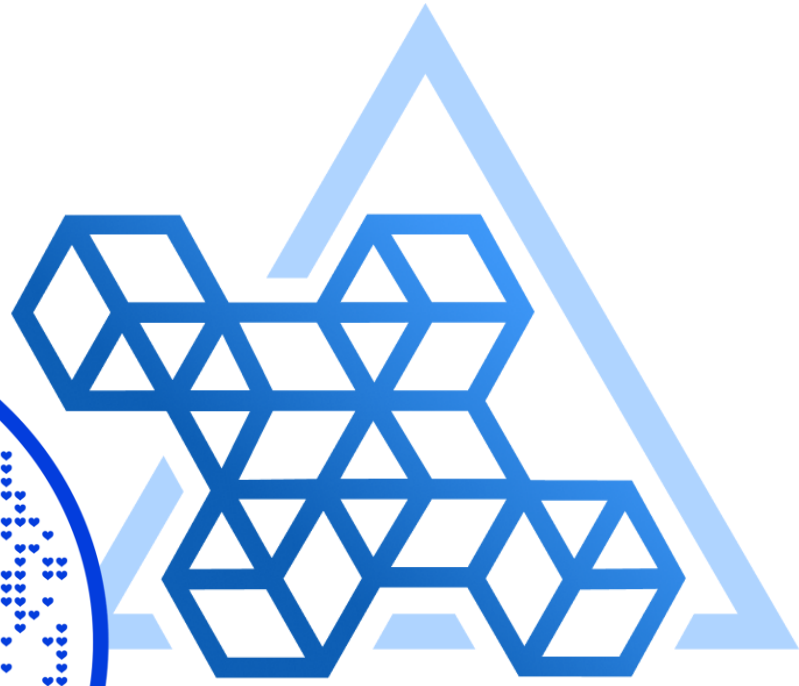

PERCONA

Kubernetes Operators

Work in Progress...

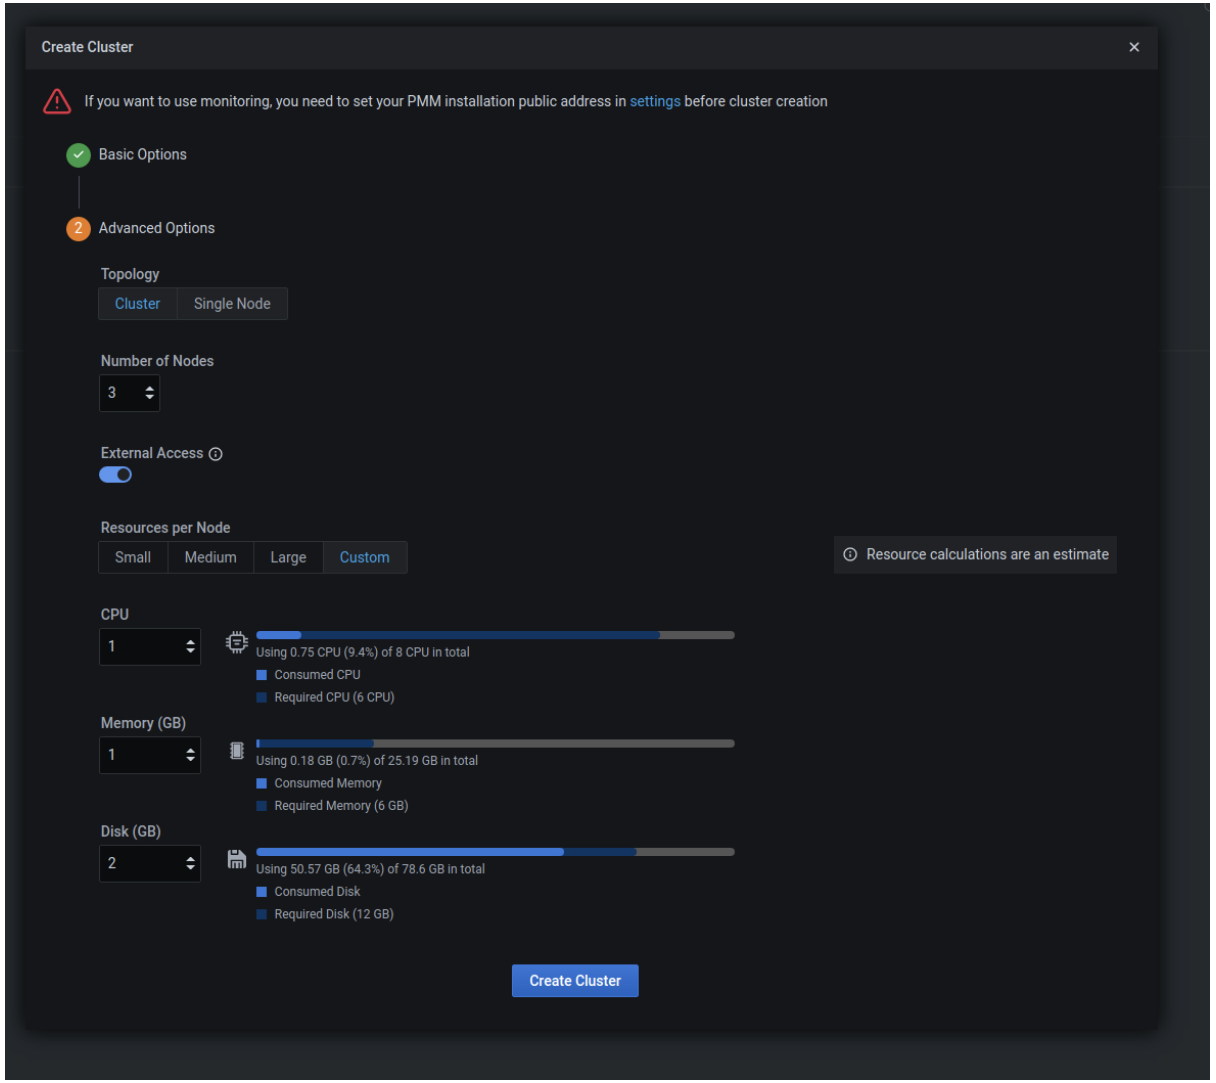

https://www.percona.com/doc/percona-monitoring-and-management/2.x/setting-up/server/dbaas.html

PERCONA

Open Source Databases are on the path from support of Containers to full DBaaS Experience

Docker Support Mature, Kubernetes Support is Getting Where and DBaaS Experience is still work in Progress

In Open Source you can always be part of solution – help build the projects you care about!

Summary

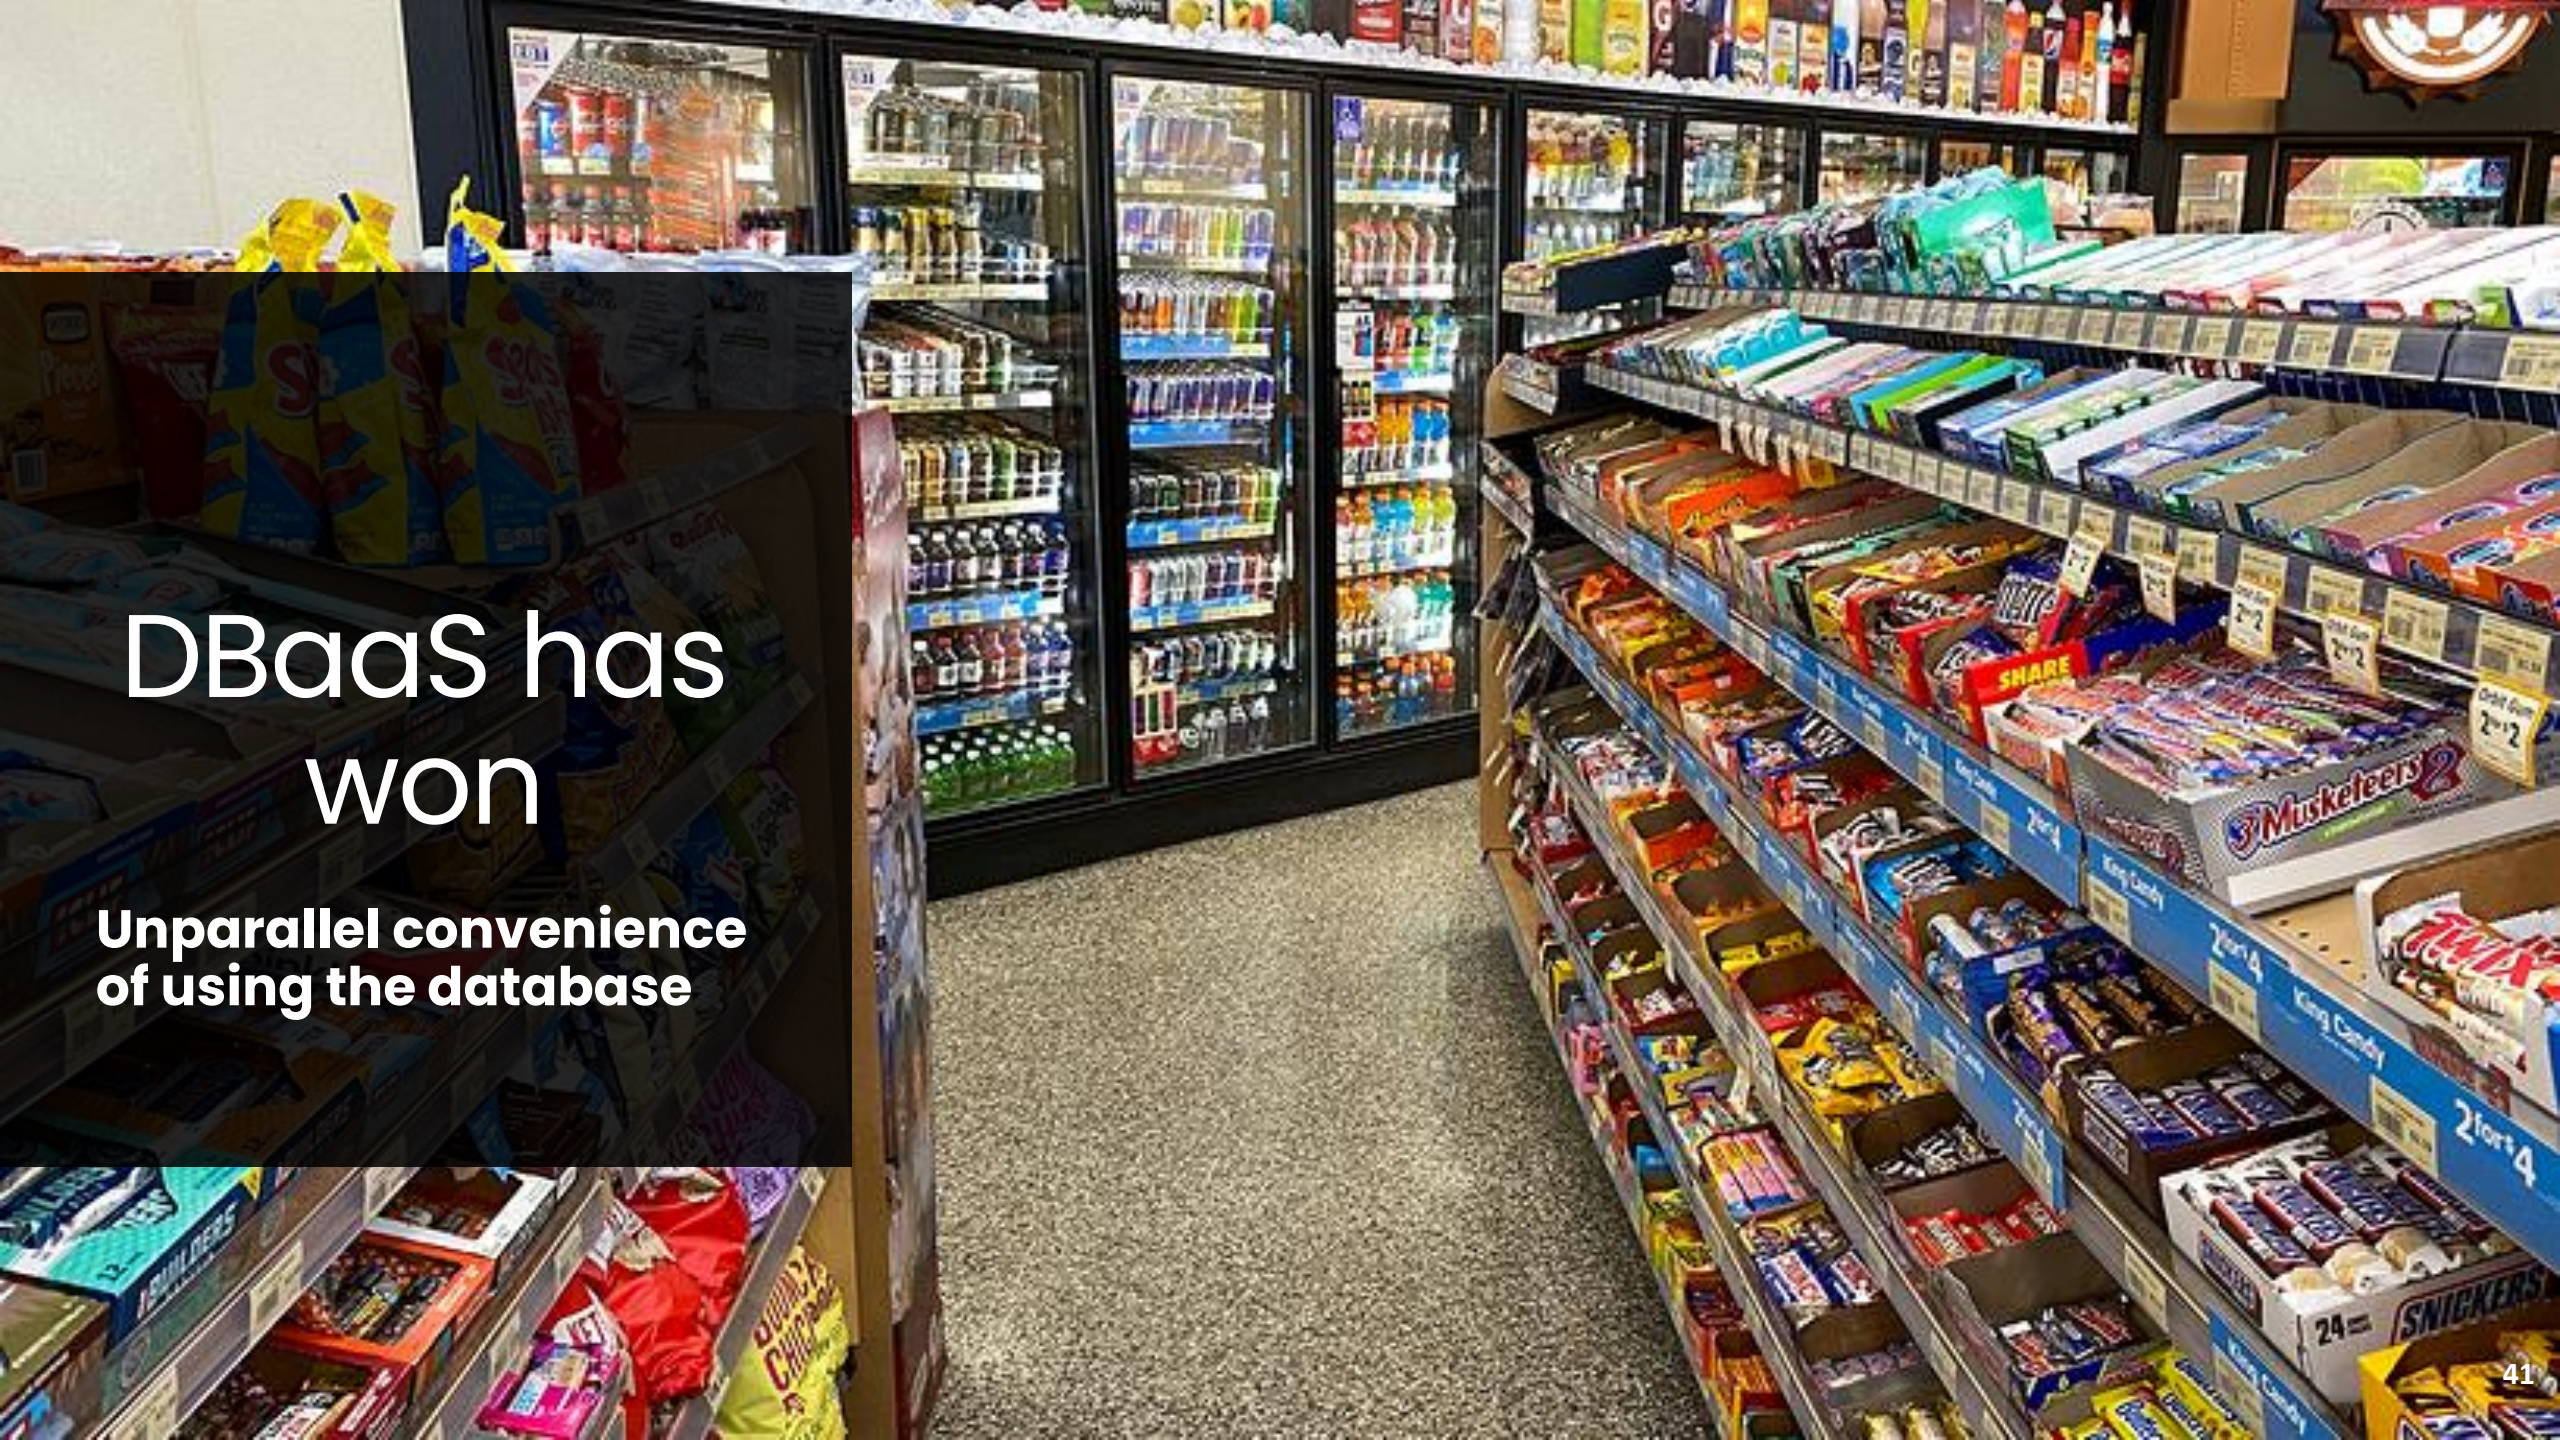

DBaaS has won

Unparallel convenience of using the database

Software Vendor Lock-in Sucks

Open Source to the Rescue

Thank you, Let's Connect!

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

https://twitter.com/PeterZaitsev

http://www.peterzaitsev.com