Cloud Native MongoDB: building a scalable infrastructure with Open Source components

Takis Stathopoulos, PhD

Enterprise Architect@Percona

Athens, Greece 2023

PERCONA UNIVERSITY

In partnership with FerretDB

Agenda

I. Cloud Native

II. Cloud Native MongoDB Options

    A. Proprietary

    B. Open

III. Comparing Proprietary vs. Open Source

IV. Conclusions

Why?

- Natively web-scale ready
- Built in resilience, High Availability
- Low entry barrier for development & installation
- NoSQL

Is everything perfect?

Relied on open source components

Initially open source itself

“Source-Available” now

SSPL means MongoDB being the Single Public DBaaS provider

Cloud Native MongoDB Options

Athens, Greece 2023

PERCONA UNIVERSITY

In partnership with FerretDB

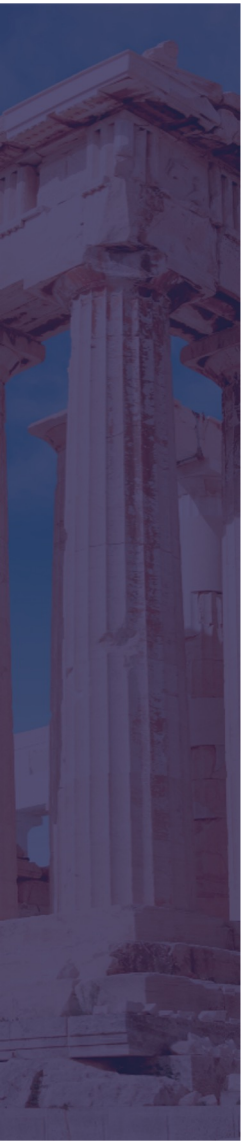

Cloud Native Definition

Source: https://github.com/cncf/toc/blob/main/DEFINITION.md

Cloud native technologies empower organizations to build and run scalable applications in modern, dynamic environments such as public, private, and hybrid clouds. Containers, service meshes, microservices, immutable infrastructure, and declarative APIs exemplify this approach.

These techniques enable loosely coupled systems that are resilient, manageable, and observable. Combined with robust automation, they allow engineers to make high-impact changes frequently and predictably with minimal toil.

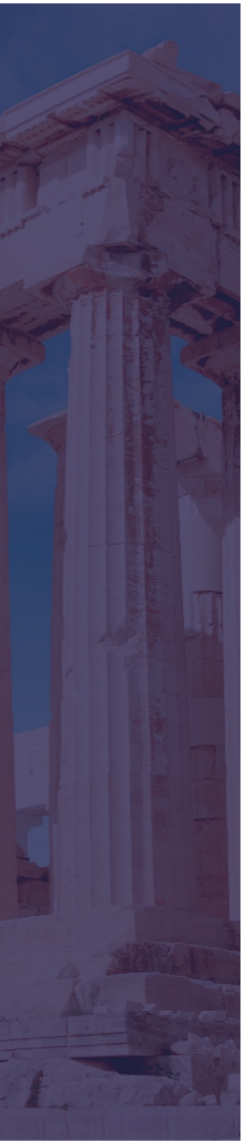

Cloud Native Definition

Source: https://github.com/cncf/toc/blob/main/DEFINITION.md

Cloud native technologies empower organizations to build and run **scalable** applications in modern, dynamic environments such as **public**, **private**, and **hybrid** clouds. **Containers**, service meshes, **microservices**, immutable infrastructure, and **declarative APIs** exemplify this approach.

These techniques enable loosely coupled systems that are **resilient**, **manageable**, and **observable**. Combined with robust **automation**, they allow engineers to make high-impact changes frequently and predictably with minimal toil.

PERCONA

The proprietary way

Athens, Greece 2023

PERCONA UNIVERSITY

In partnership with FerretDB

What?

Fully Managed

Database as a Service

(DBaaS)

MongoDB Enterprise

Features and more

Strengths

- Great UX!
- Pay-as-you go
- Cloud resources *seem* virtually unlimited
- Available in *major* Public Clouds
  - AWS, Azure, GCP

Traps?

Scaling Granularity

Couple of sizes fit all!

MongoDB Atlas, 3-node RS, AWS eu-west-1

M40	16 GB	80 GB	4 vCPUs	from **$1.15**/hr
M50	32 GB	160 GB	8 vCPUs	from **$2.20**/hr
M60	64 GB	320 GB	16 vCPUs	from **$4.36**/hr
M80	128 GB	750 GB	32 vCPUs	from **$8.06**/hr
M140	192 GB	1000 GB	48 vCPUs	from **$12.13**/hr
M200	256 GB	1500 GB	64 vCPUs	from **$16.10**/hr
M300	384 GB	2000 GB	96 vCPUs	from **$24.11**/hr
M400	512 GB	3000 GB	64 vCPUs	from **$24.73**/hr
M700	768 GB	4096 GB	96 vCPUs	from **$36.73**/hr

Scaling up:

- Only 10 production-grade tiers
- Any upgrade costs double
- Limited options available

M50 $2.20/h = **$19,272/year**

M60 $4.36/hr = **$38,193.6/year**

What?	How?
Autoscaling	**Scale Up** Avg. CPU Util > 75% for past 1h Avg. Mem > 75% for past 1h Next highest cluster tier **Scale down** Avg. Util. < 50% for the past 24h* **https://www.mongodb.com/docs/atlas/cluster-autoscaling/*
Topology	**No arbiter nodes** **No single node dev clusters**
Uniform offering	**AWS, Azure, GCP common denominator**

Overprovisioning

- Instead of Root Cause Analysis

Less careful planning

- Convenient to add nodes and shards
- Non-optimal architectures
- Building 20-shard 5-node each cluster? Easy!

Pay as you go quickly adds-up

Cloud DBaaS lock-in

You can check-out any time you like, but you can never leave

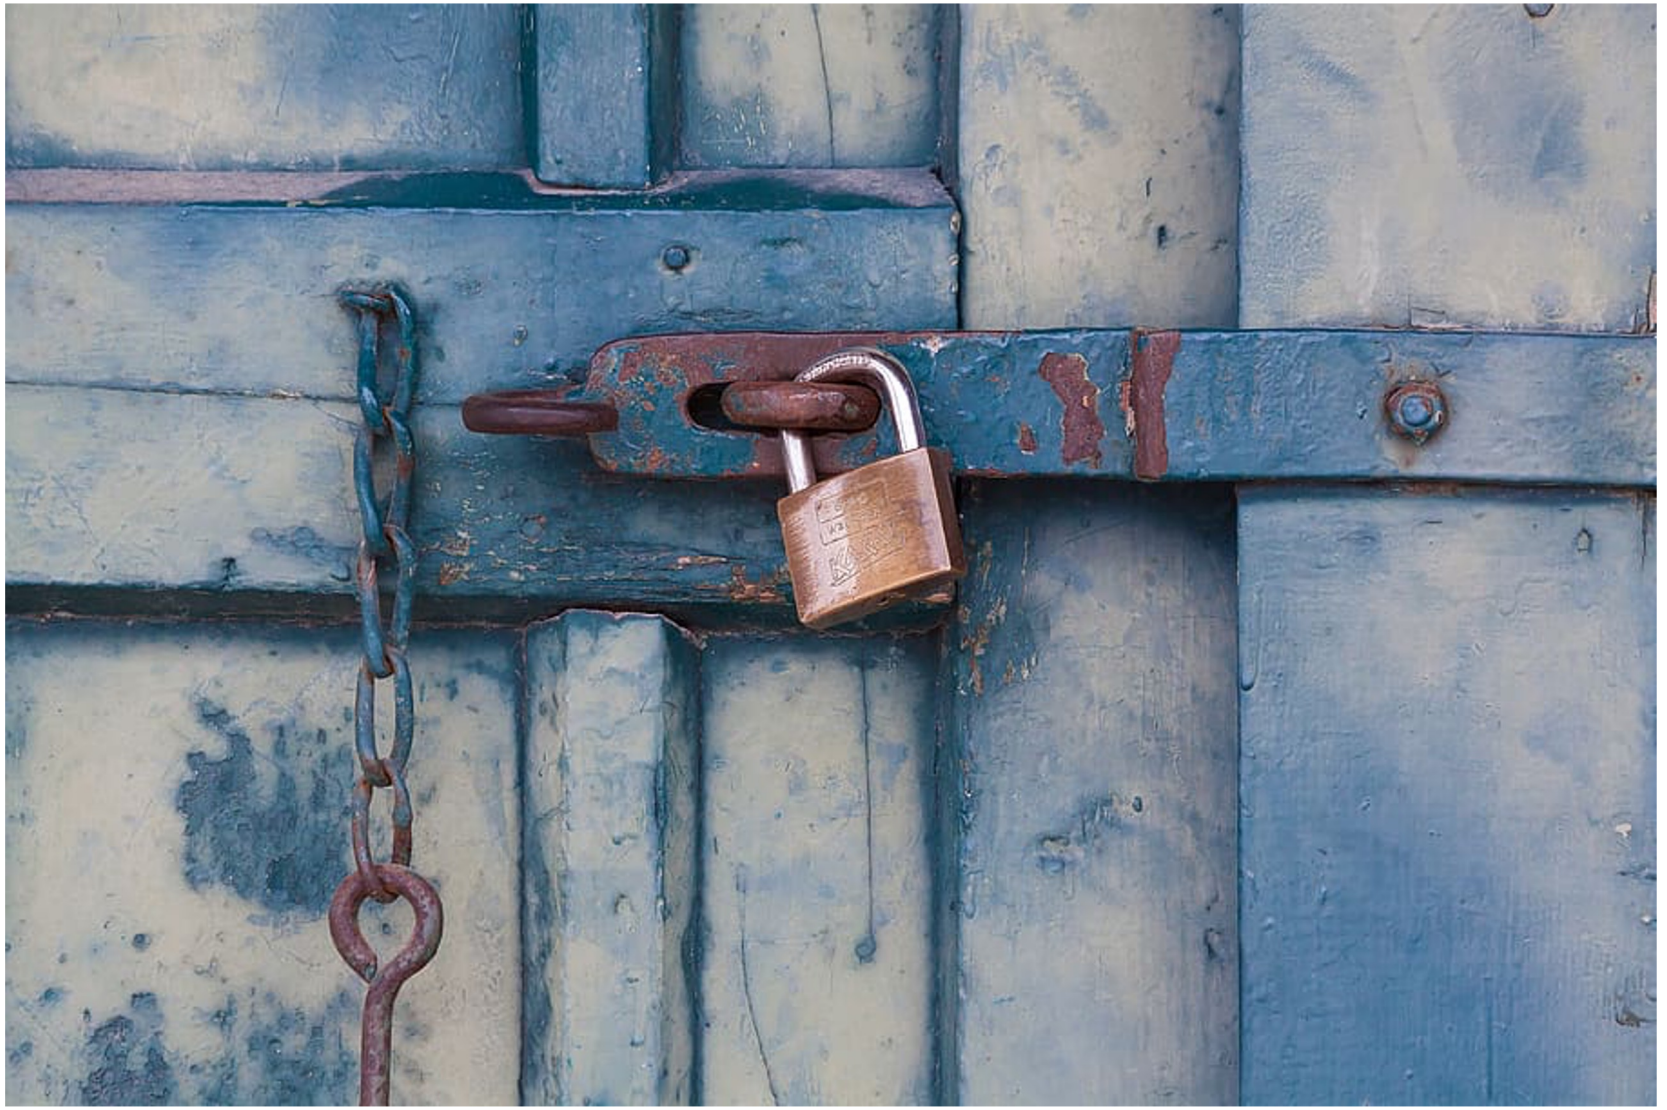

Athens, Greece 2023

PERCONA UNIVERSITY

In partnership with FerretDB

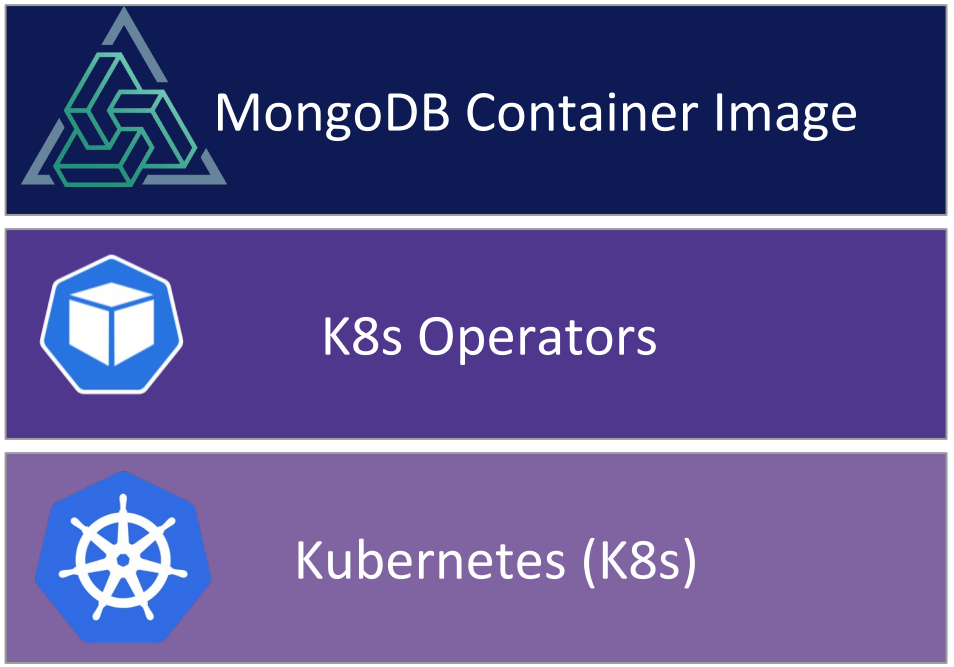

Observable, Containers, Microservices, …

Automated, Manageable, Easy to use, …

Scalable, Cloud Independent, Resilient, Declarative, …

Kubernetes in a single slide

Stateless vs.

Stateful

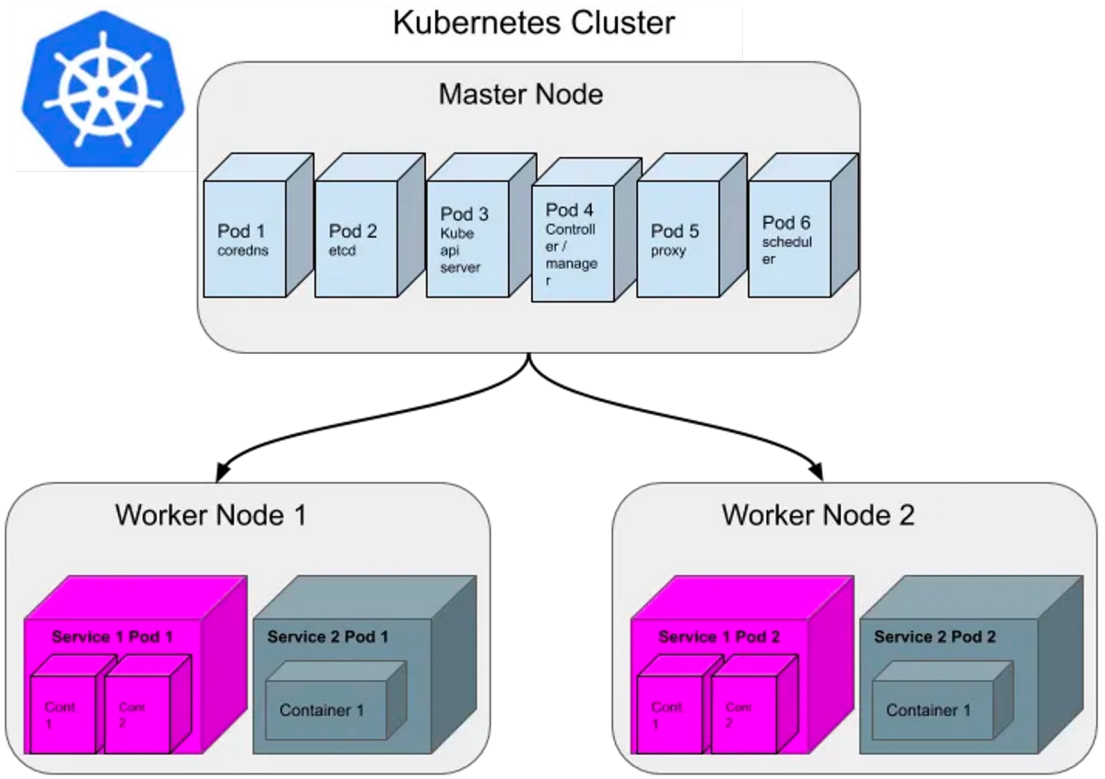

Basic objects

Cluster, Pods, Worker Nodes, Volumes, Secrets, Deployments, Services, ReplicaControllers, StatefulSets, Persistent Volume Claims …

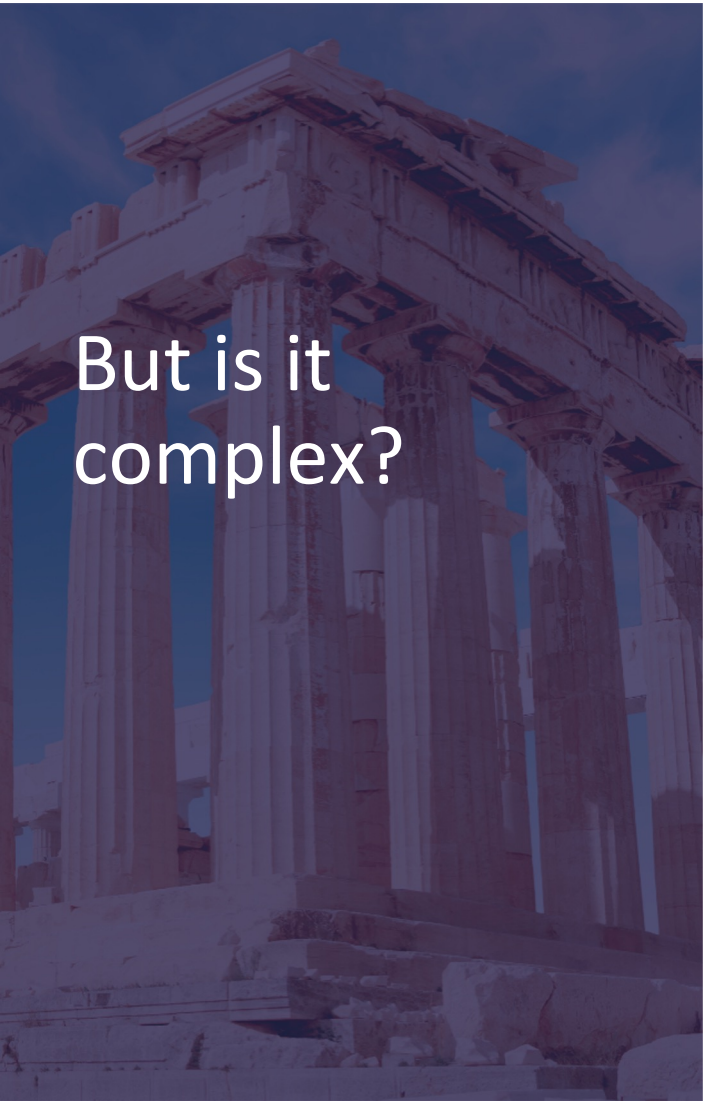

But is it complex?

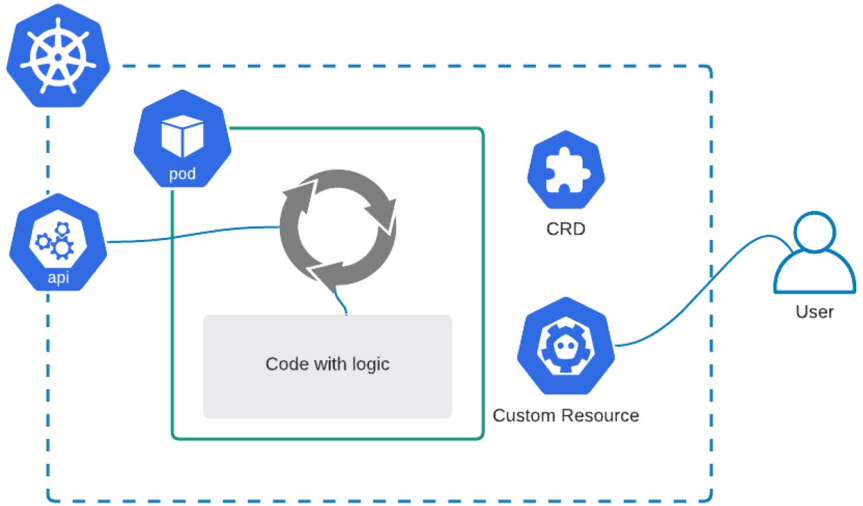

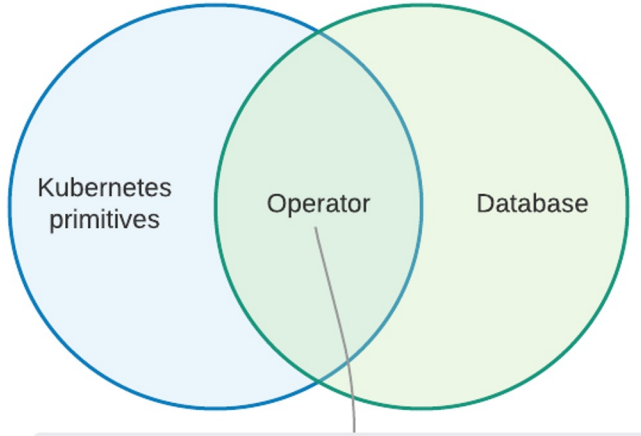


```
spec:
  image: percona/percona-server-mongodb:4.4.6-8
  replsets:
   - name: rs0
     size: 3
    sharding:
      mongos:
        size: 3
       configsvrReplSet:
         size: 3
     backups:
       ...
```

Operators are software extensions to Kubernetes that make use of Custom Resources Definitions (CRDs) to manage applications and their components.

Operators abstract and automate Database - level concepts to K8s primitive transparently for the end - user

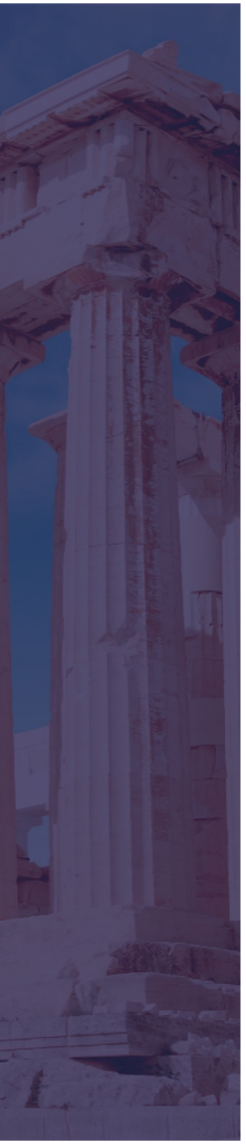

Deploying a MongoDB cluster via Percona Operator for MongoDB

```
apiVersion: psmdb.percona.com/v1
kind: PerconaServerMongoDB
metadata:
  name: percona-live-cluster
spec:
  crVersion: 1.15.0
  image: percona/percona-server-mongodb:6.0.4-
3
  secrets:
    users: minimal-cluster
  replsets:
  - name: shard1
    size: 3
    resources:
      limits:
        cpu: "4"
        memory: "8G"
      requests:
        cpu: "4"
        memory: "8G"
    volumeSpec:
      persistentVolumeClaim:
        resources:
          requests:
            storage: 30Gi
```

```
  sharding:
    enabled: true
    configsvrReplSet:
      size: 3
      resources:
        limits:
          cpu: "2"
          memory: "4G"
        requests:
          cpu: "2"
          memory: "4G"
      volumeSpec:
        persistentVolumeClaim:
          resources:
            requests:
              storage: 3Gi
    mongos:
      size: 3
```

```
$kubectl apply -f cr.yaml
```

https://docs.percona.com/percona-operator-for-mongodb/compare.html
Operators compared

PERCONA

Percona Operator for MongoDB

1. Deploy easily: replica sets, shards, (mongo/d/s/c)
2. Topology management (arbiters, node affinity, scaling)
3. Monitoring integration
4. Network exposure and load balancing
5. Backups management with Percona Backup for MongoDB
6. Self-healing
7. Upgrade automation (minor, manual major)
8. Configuration adjustments

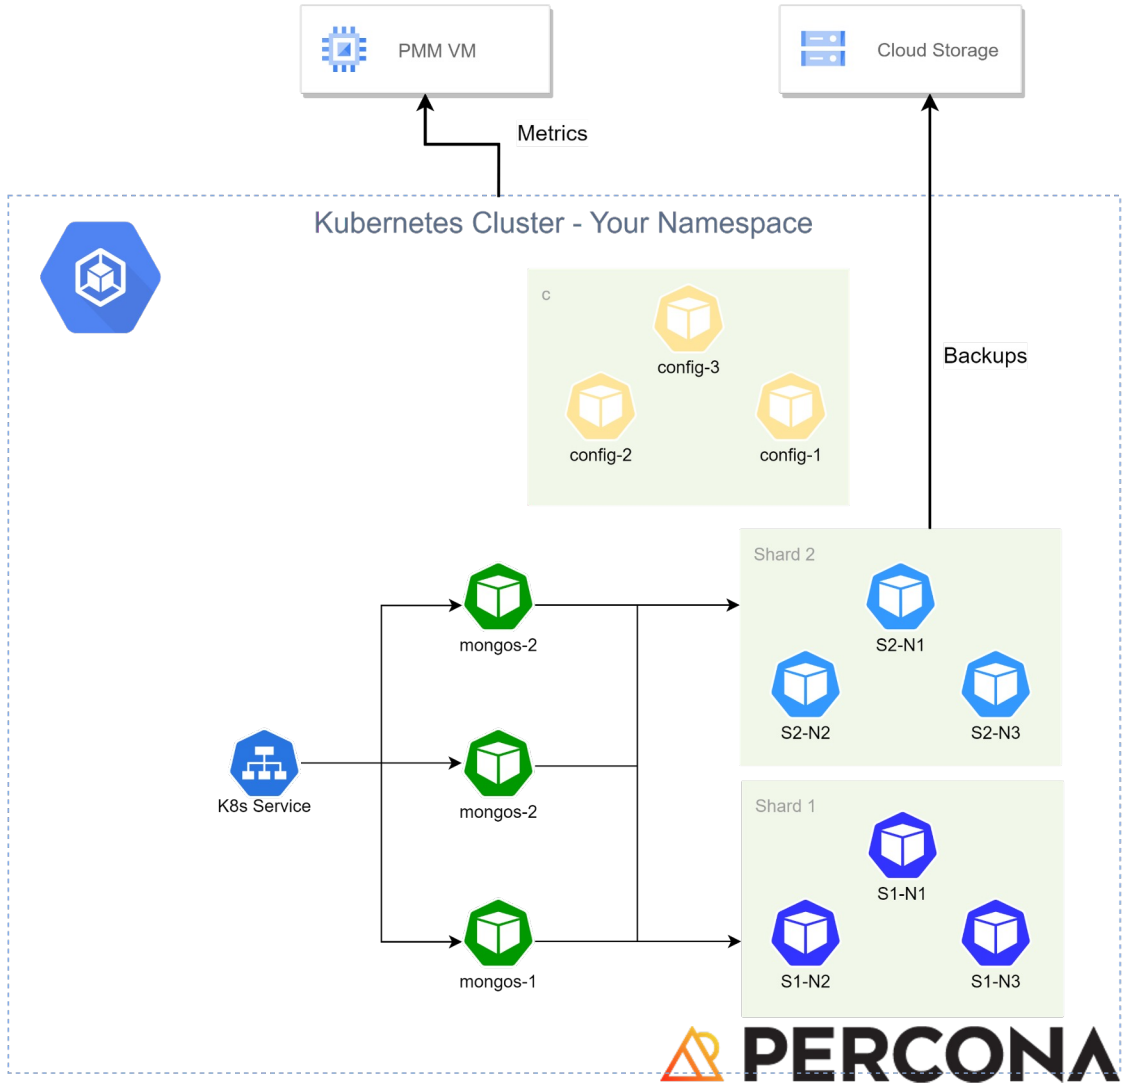

Percona Server for MongoDB

Binary compatible, drop-in replacement

for MongoDB CE . No license fees, free to use

Enterprise features, without the restrictions

1. Advanced backups (Physical, PITR)
2. LDAP Integration
3. Data-at-rest encryption
4. KMIP integration
5. Auditing
6. PMM Monitoring

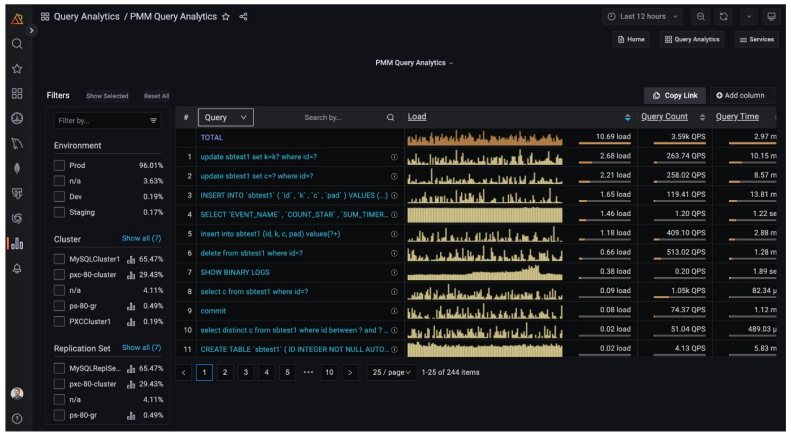

Enterprise Level QA

Test and package for everyone!

Enterprise Features

Bring in the enterprise features companies need.

PERCONA

It’s open.

Is it better?

Athens, Greece 2023

PERCONA UNIVERSITY

In partnership with FerretDB

K8s resources are cost-efficient

- **K8s runs everywhere!**
- **Managed K8s options**
  - ~$70/month at any major cloud provider
  - Runs on IaaS cost-efficient resources
  - Utilize “raw” resources at “raw” resource prices

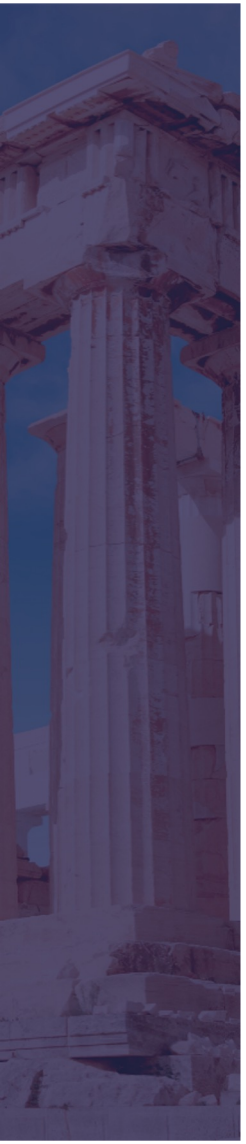

DBaaS convenience fee vs IaaS cost

Cloud and region	MongoDB Atlas	IaaS pay-as-you-go	IaaS reserved (3y)
M30 (3-node RS, 2vCPU, 8GB RAM, 40GB disk)			
AWS us-east1	$4,730	$2,666	$1,204
Azure centralus	$5,256	$2,450	$1,049
GCP us-east1	$3,854	$2,796	$1,393
M200 (4x3-node RS, 64vCPU, 256GB RAM, 2TB storage)			
AWS us-east1	$555,822	$355,086	$159,674
Azure centralus	$626,515	$366,672	$162,575
GCP us-east1	$525,755	$379,376	$198,303

No backup, no data transfer fees, yearly

PERCONA

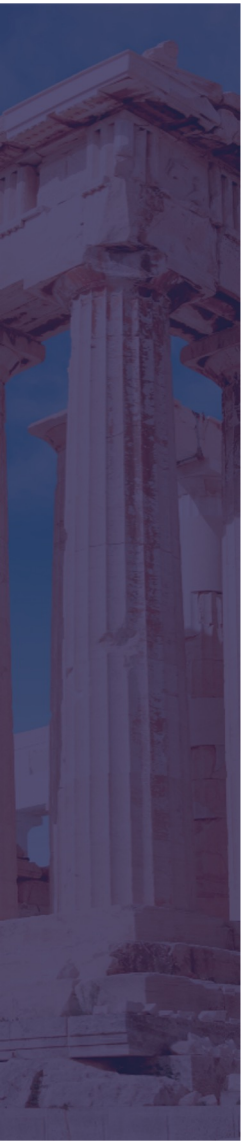

DBaaS convenience fee vs IaaS cost

Cloud and region	MongoDB Atlas	IaaS pay-as-you-go	IaaS reserved (3y)
M30 (3-node RS, 2vCPU, 8GB RAM, 40GB disk)			
AWS us-east1	$4,730	**$2,666**	$1,204
Azure centralus	$5,256	**$2,450**	$1,049
GCP us-east1	$3,854	**$2,796**	$1,393
M200 (4x3-node RS, 64vCPU, 256GB RAM, 2TB storage)			
AWS us-east1	$555,822	**$355,086**	$159,674
Azure centralus	$626,515	**$366,672**	$162,575
GCP us-east1	$525,755	**$379,376**	$198,303

No backup, no data transfer fees, yearly

PERCONA

Scaling Granularity

~~4 CPU -> 8 CPU -> 16 CPU -> 32 CPU~~
~~16GB -> 32GB -> 64GB -> 128 GB~~

1m (0.001 CPU) and 1 (CPU)
k, M, G, T...

12,5 CPU and 19GB RAM

- Multiple tiers
- Scale up earlier
- Scale up more gracefully

- **Optimize your K8s environment**
  - Define multiple node groups types
  - Choose CPU/memory ratio wisely
  - Utilize Local NVMe’s vs. EBS storage
    - Storage can be extremely expensive for larger workloads

What?	How?	
Autoscaling	**VPA**	
	HPA/KEDA	
	Cluster Autoscaler	
Topology	**Select arbiter nodes**	
	Any topology	
Efficiency	**Tap into local environment features (NVMe, …)**	
	Leverage workload understanding efficiencies	

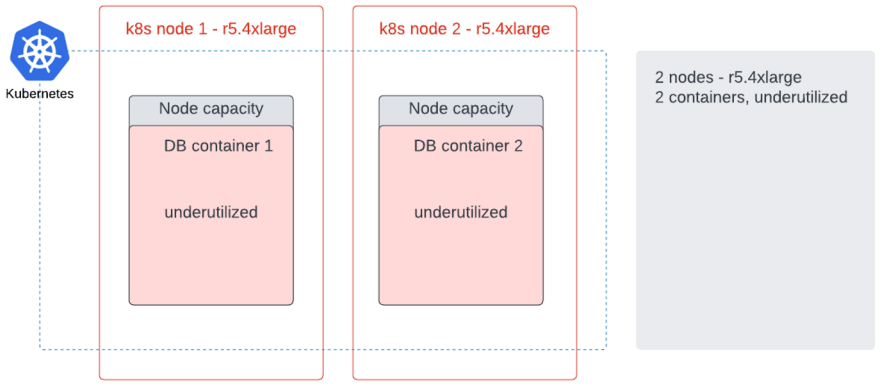

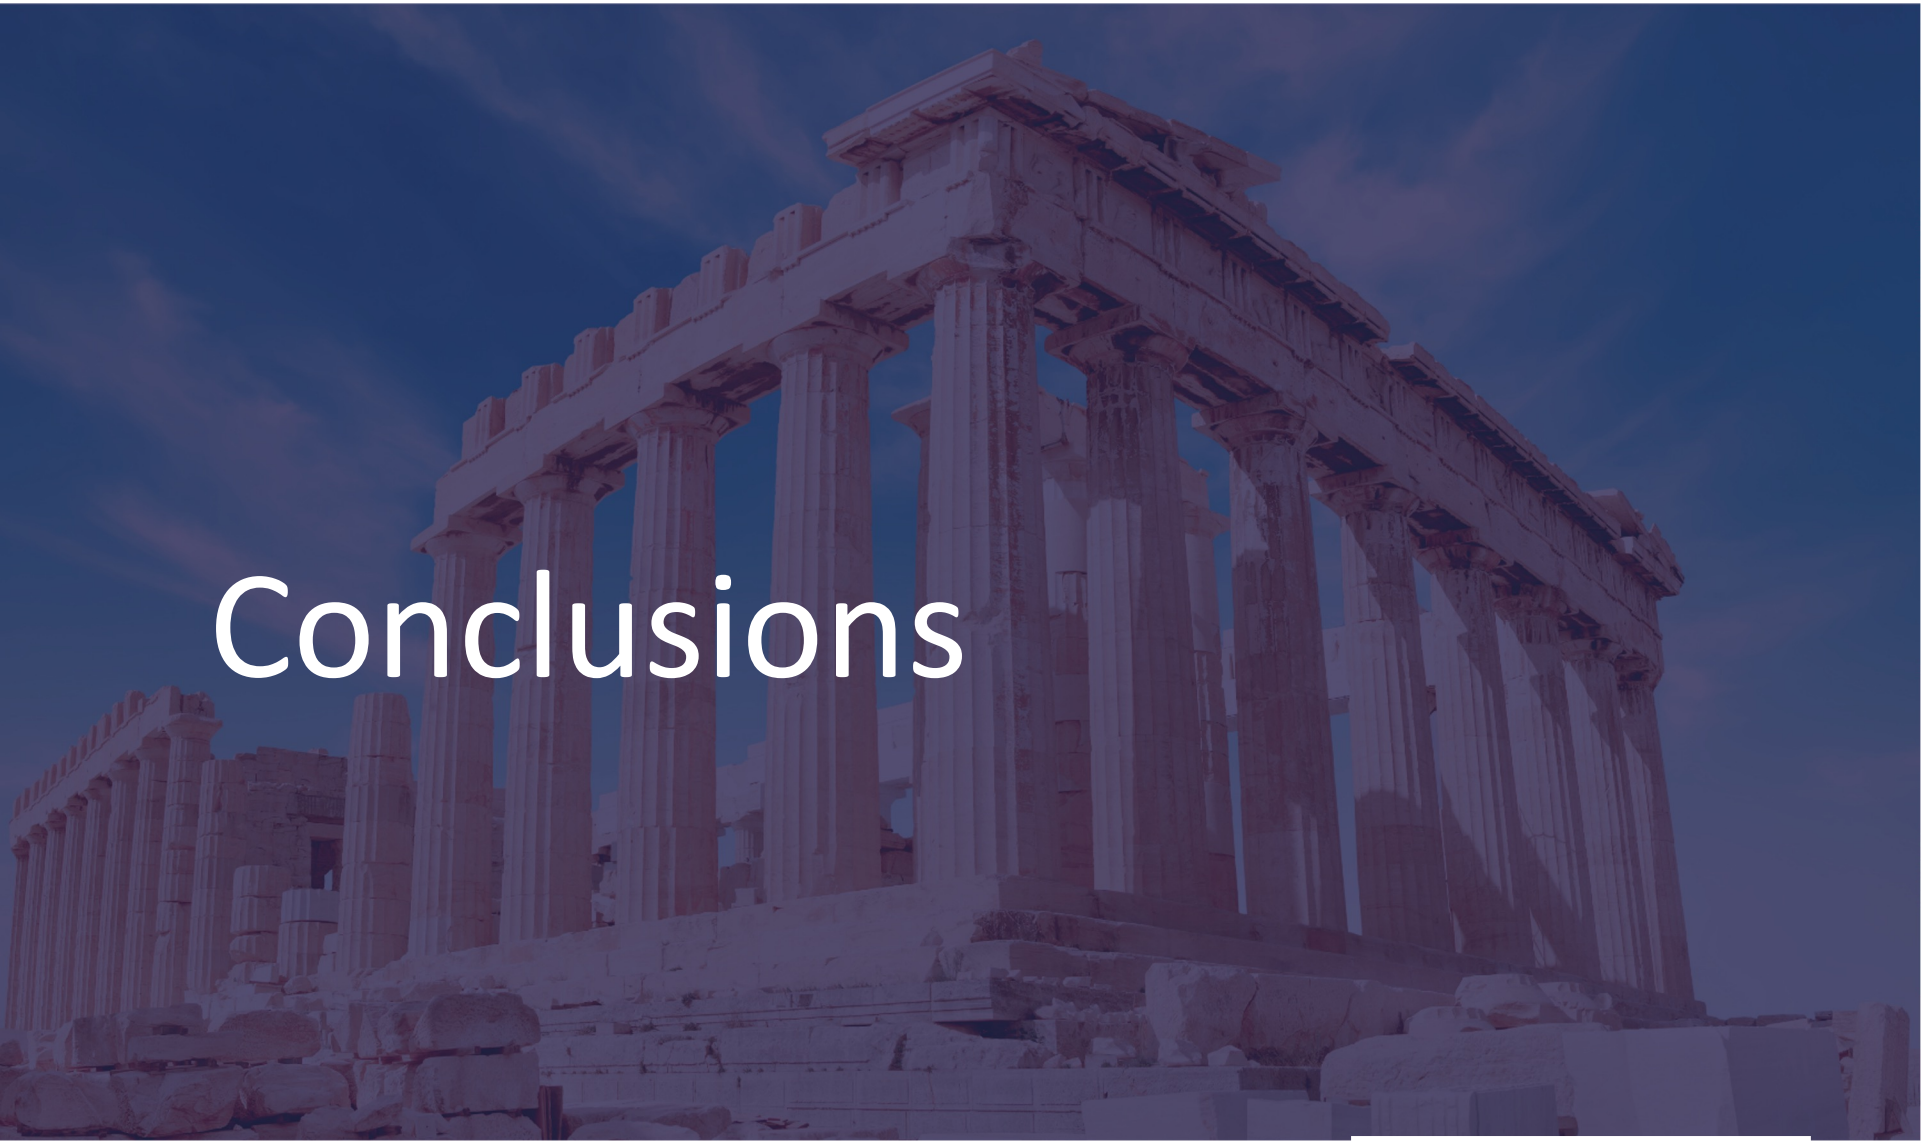

Conclusions

Athens, Greece 2023

PERCONA UNIVERSITY

In partnership with FerretDB

Key takeaways

1. No Silver bullet

    Understand your workload and growth projections

1. Cloud DBaaS such as MongoDB Atlas comes with a huge premium fee

    In some cases 5x cloud resources cost.

1. Hidden costs and suboptimal MongoDB configuration make the situation worse

    Cloud DBaaS fees are difficult to predict and understand.

1. Kubernetes and Percona MongoDB Operator can handle any MongoDB workload

    The solution has been proven in production for many years.

1. Running MongoDB on K8s has a significant cost saving potential

    It makes it possible to utilize cheap “raw” cloud resources without giving up convenience and automation

In partnership with FerretDB