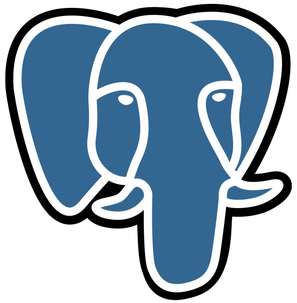

Databases on Kubernetes:

A Storage Story

Databases  Kubernetes

simplyblock

Chris Engelbert

Devrel @ simplyblock

Previous fun companies:

- Ubisoft / Blue Byte
- Hazelcast
- Instana
- clevabit
- Timescale

Interests:

- Developer Relations
- Anything Performance Engineering
- Backend Technologies
- Fairy Tales (AMD, Intel, Nvidia)

@noctarius2k

@noctarius2k@mastodon.online

@noctarius.com

Why Databases in Kubernetes?

Why not to run a database in Kubernetes?

Christoph Engelbert / Noctarius ツ / エンゲルベルト クリス
@noctarius2k

Why you ***SHOULD NOT*** run a database in Kubernetes?
What do you think? Please help me, I need you! 🙏💙

via: #postgresql #mysql #mariadb #kafka #kubernetes

Why not to run a database in Kubernetes?

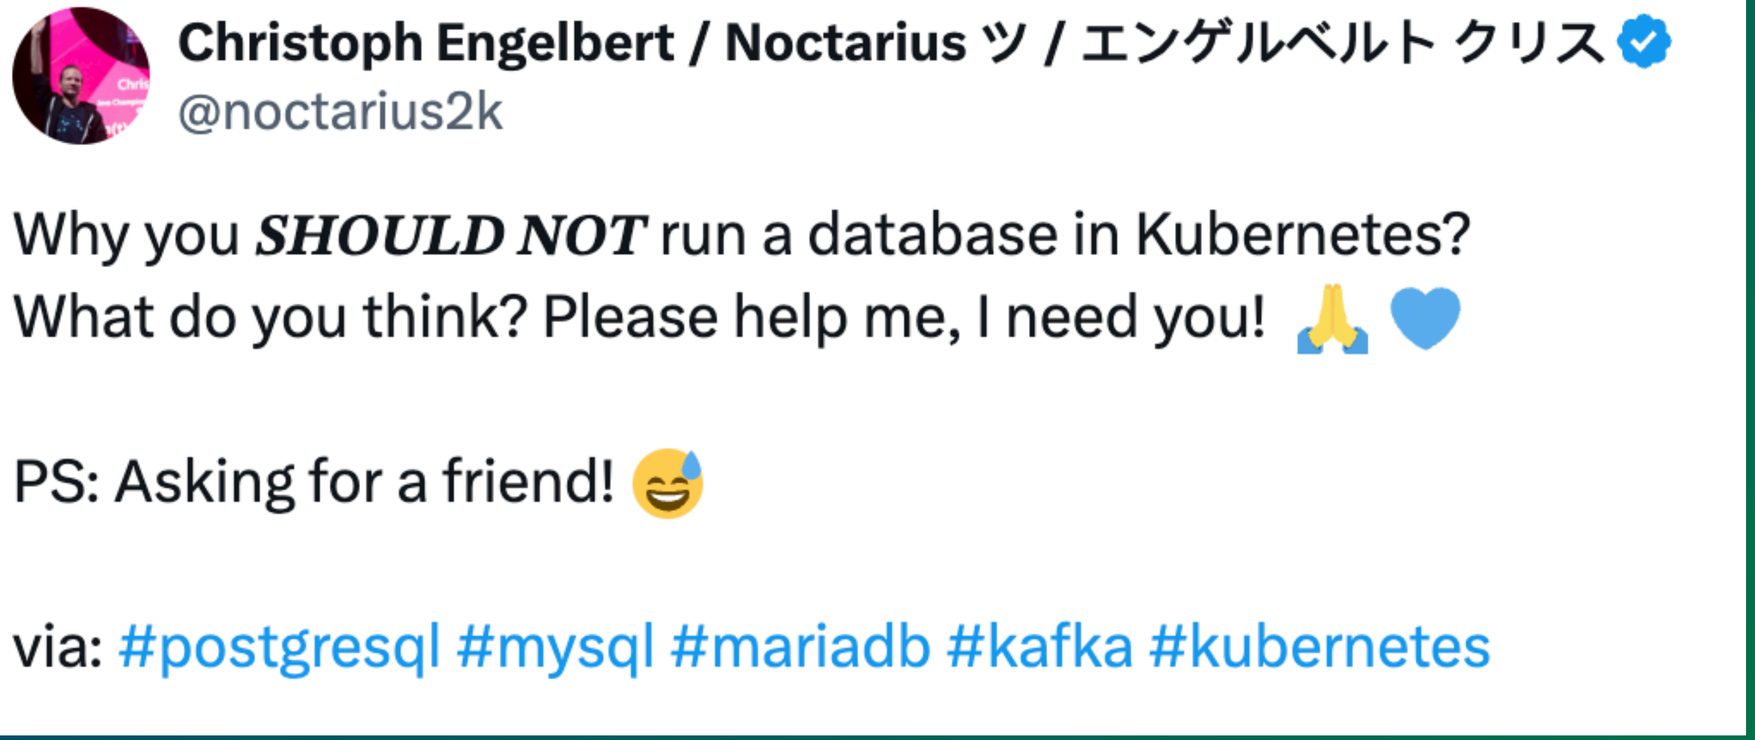

Why not to run a database in Kubernetes?

simplyblock

Why not to run a database in Kubernetes?

K8s is not designed with Databases in mind!

Why not to run a database in Kubernetes?

K8s is not designed with Databases in mind!

Never run Stateful Workloads in k8s!

Why not to run a database in Kubernetes?

K8s is not designed with Databases in mind!

Never run Stateful Workloads in k8s!

Persistent Data will kill you! Too slow!

Why not to run a database in Kubernetes?

K8s is not designed with Databases in mind!

Never run Stateful Workloads in k8s!

Persistent Data will kill you! Too slow!

Nobody understands Kubernetes!

Why not to run a database in Kubernetes?

K8s is not designed with Databases in mind!

Never run Stateful Workloads in k8s!

Persistent Data will kill you! Too slow!

Nobody understands Kubernetes!

What’s the benefit; databases don’t need autoscaling!

Why not to run a database in Kubernetes?

K8s is not designed with Databases in mind!

Never run Stateful Workloads in k8s!

Persistent Data will kill you! Too slow!

Nobody understands Kubernetes!

What’s the benefit; databases don’t need autoscaling!

Databases and applications should be separated!

Why not to run a database in Kubernetes?

K8s is not designed with Databases in mind!

Never run Stateful Workloads in k8s!

Persistent Data will kill you! Too slow!

Nobody understands Kubernetes!

What’s the benefit; databases don’t need autoscaling!

Databases and applications should be separated!

Not another layer of indirection / abstraction!

Why not to run a database in Kubernetes?

K8s is not designed with Databases in mind!

Never run Stateful Workloads in k8s!

Persistent Data will kill you! Too slow!

Nobody understands Kubernetes!

What’s the benefit; databases don’t need autoscalin

Databases and applications should be separat

Not another layer of indirection / abstractio

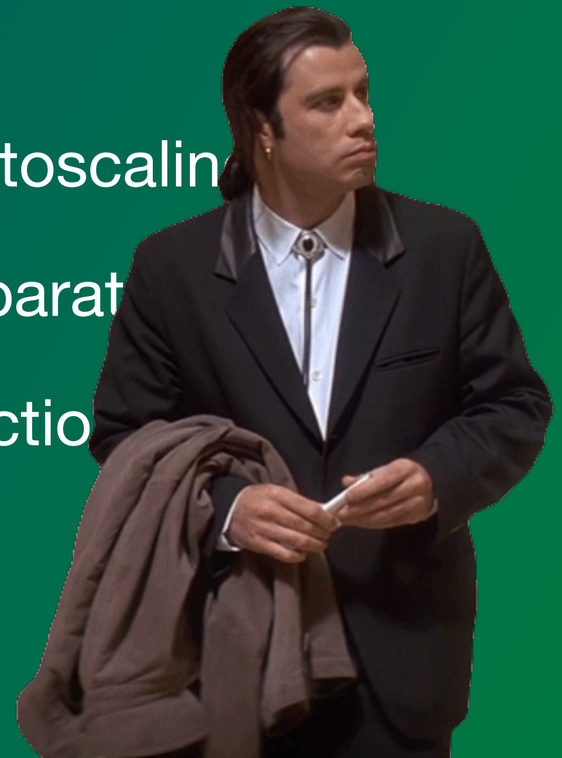

simplyblock

You Already Run A Database In Kubernetes

simplyblock

You already run a database in Kubernetes

simplyblock

Why?

simplyblock

Why?

No Cloud-Vendor Lock-In

simplyblock

Why?

No Cloud-Vendor Lock-In

Faster Time To Market

simplyblock

Why?

No Cloud-Vendor Lock-In

Faster Time To Market

Decreasing cost

simplyblock

Why?

No Cloud-Vendor Lock-In

Faster Time To Market

Decreasing cost

Automation

simplyblock

Why?

No Cloud-Vendor Lock-In

Faster Time To Market

Decreasing cost

Automation

Unified deployment architecture

simplyblock

Why?

No Cloud-Vendor Lock-In

Faster Time To Market

Decreasing cost

Automation

Unified deployment architecture

Need read-only replicas

simplyblock

Typical Challenges

Challenges with Databases on Kubernetes

simplyblock

Challenges with Databases on Kubernetes

Selecting The Right Database

simplyblock

Challenges with Databases on Kubernetes

Selecting The Right Database

Security Requirements (Data-At-Rest Encryption)

Challenges with Databases on Kubernetes

Selecting The Right Database

Security Requirements (Data-At-Rest Encryption)

Backups and Restores

Challenges with Databases on Kubernetes

Selecting The Right Database

Security Requirements (Data-At-Rest Encryption)

Backups and Restores

Access Control and Network Policies

Challenges with Databases on Kubernetes

Selecting The Right Database

Security Requirements (Data-At-Rest Encryption)

Backups and Restores

Access Control and Network Policies

Container Images and Extensions

Challenges with Databases on Kubernetes

Selecting The Right Database

Security Requirements (Data-At-Rest Encryption)

Backups and Restores

Access Control and Network Policies

Container Images and Extensions

Monitoring and Observability

Challenges with Databases on Kubernetes

Selecting The Right Database

Security Requirements (Data-At-Rest Encryption)

Backups and Restores

Access Control and Network Policies

Container Images and Extensions

Monitoring and Observability

Storage

Storage😩

simplyblock

The Biggest Issue: Storage

K8s is not designed with Databases in mind!

Never run Stateful Workloads in k8s!

Persistent Data will kill you! Too slow!

Nobody understands Kubernetes!

What’s the benefit; databases don’t need autoscaling!

Databases and applications should be separated!

Not another layer of indirection / abstraction!

The Biggest Issue: Storage

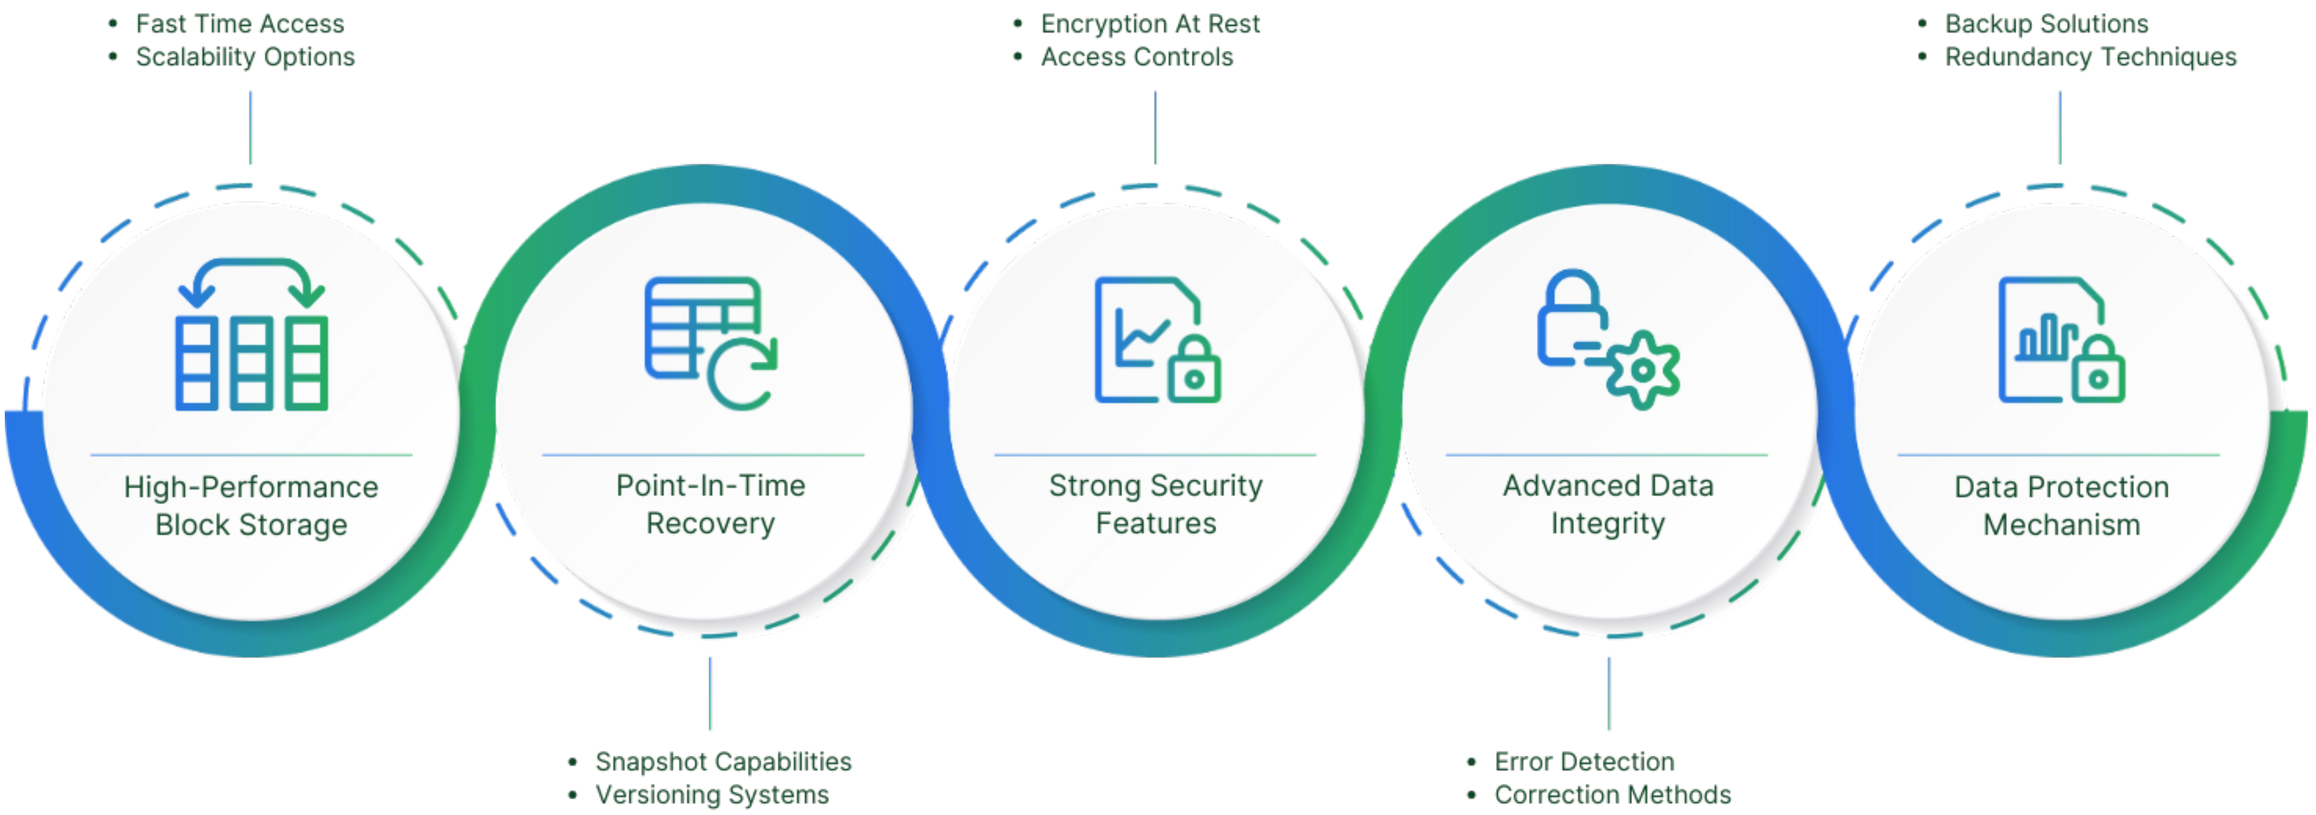

Database Storage Requirements

Storage

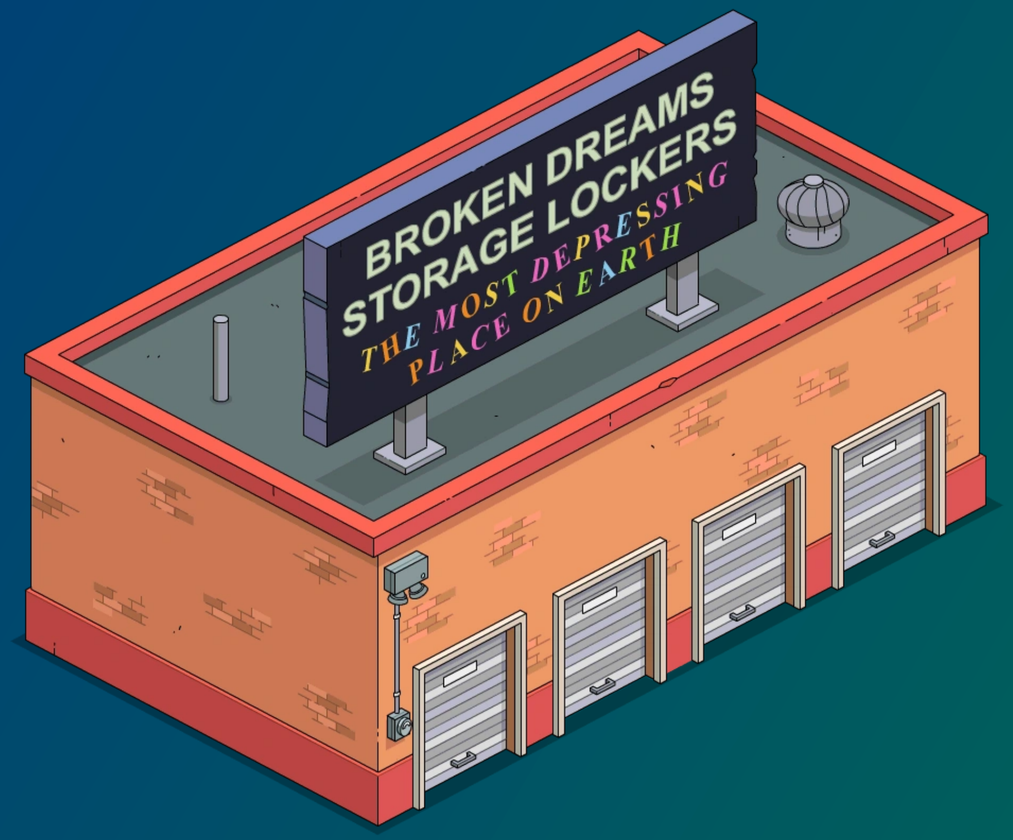

simplyblock

Storage

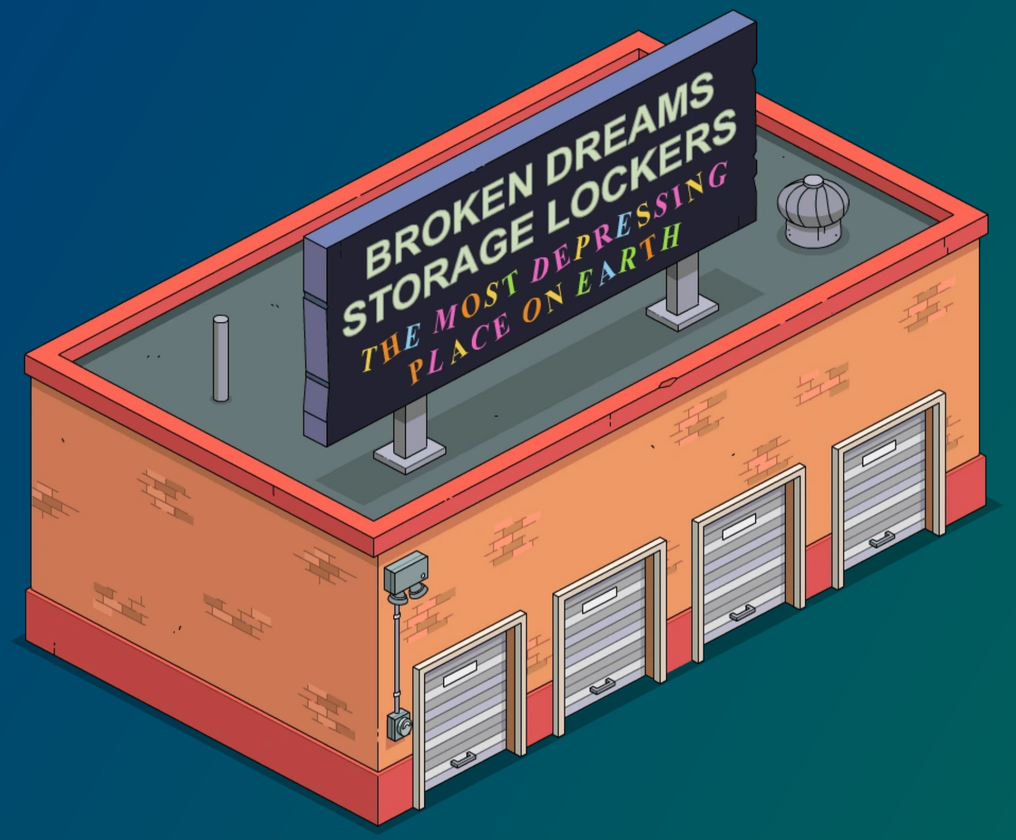

Should be dynamically provisioned

simplyblock

Storage

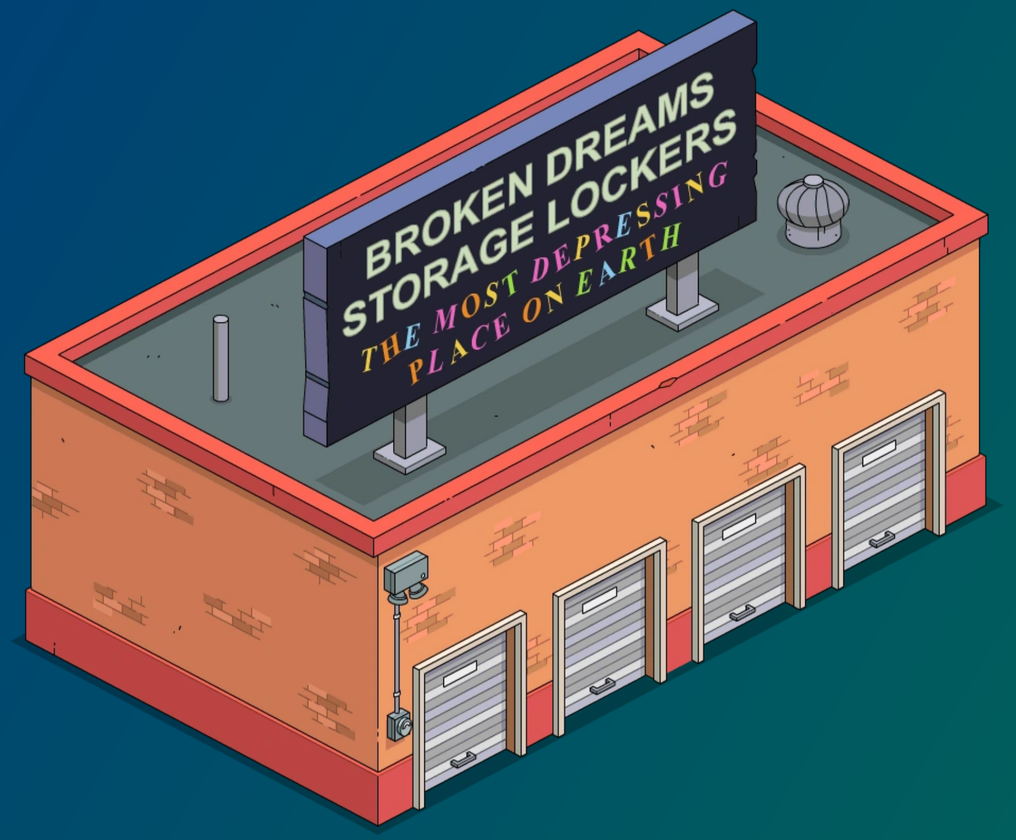

Should be dynamically provisioned

CSI provider enables encryption at rest

simplyblock

Storage

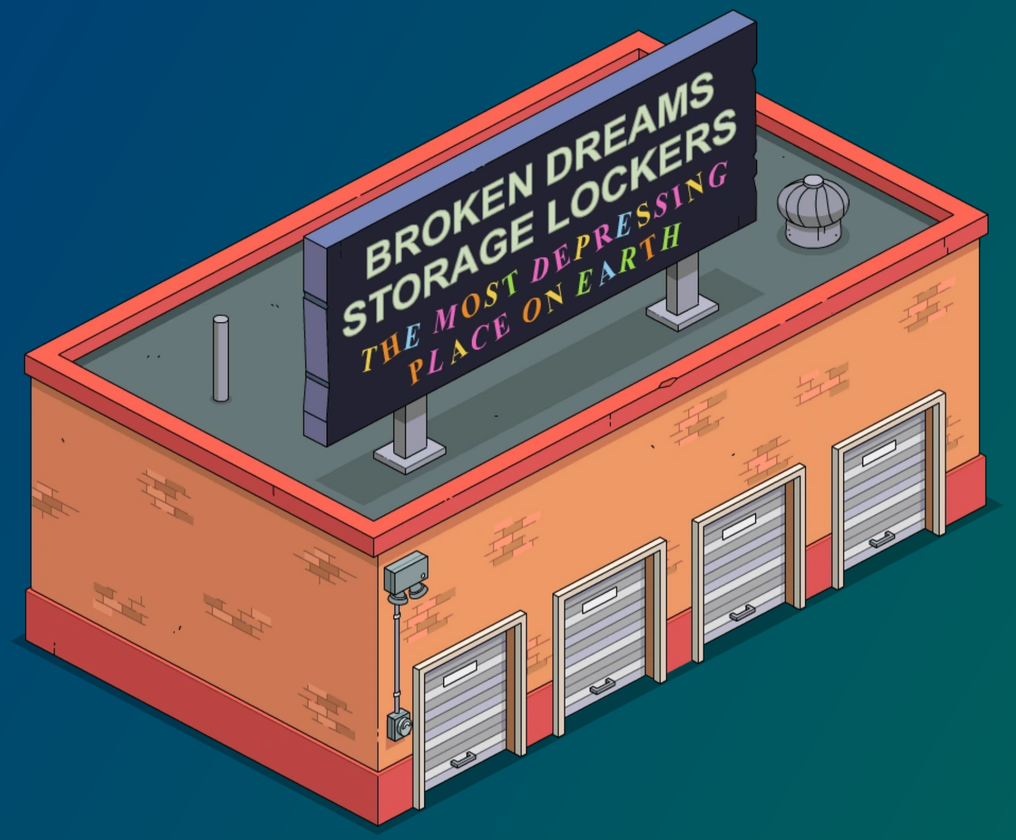

Should be dynamically provisioned

CSI provider enables encryption at rest

High IOPS (SSD or NVMe)

Storage

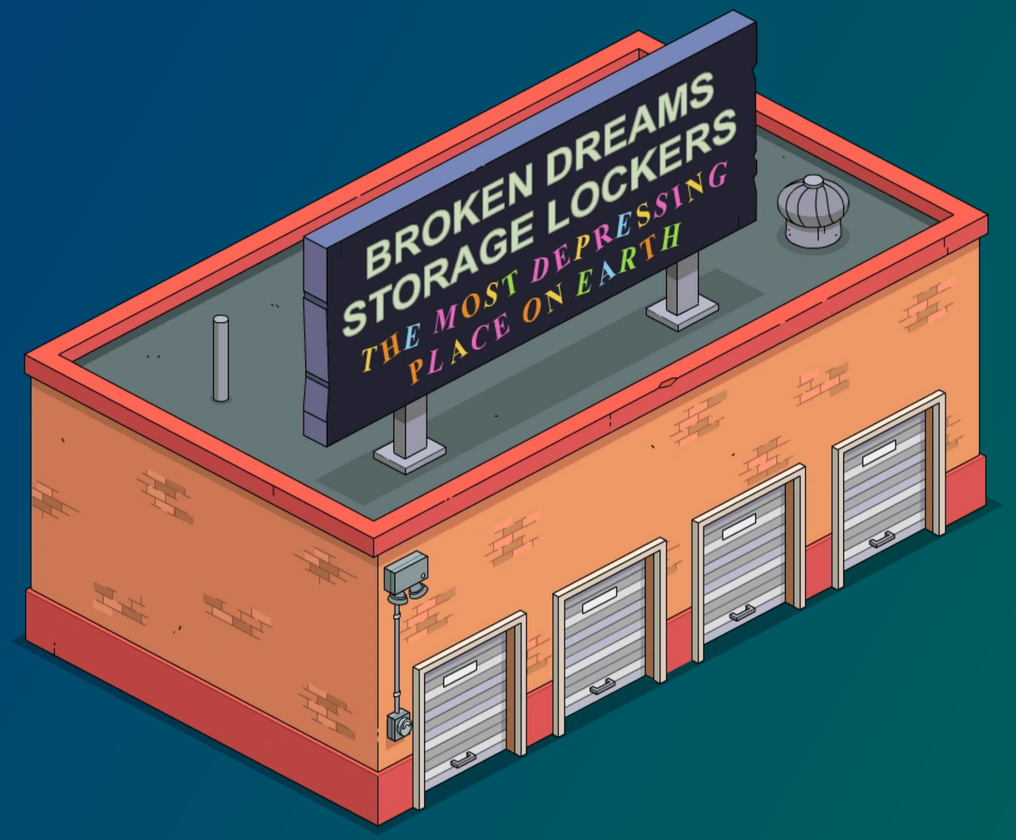

Should be dynamically provisioned

CSI provider enables encryption at rest

High IOPS (SSD or NVMe)

Low Latency

simplyblock

Storage

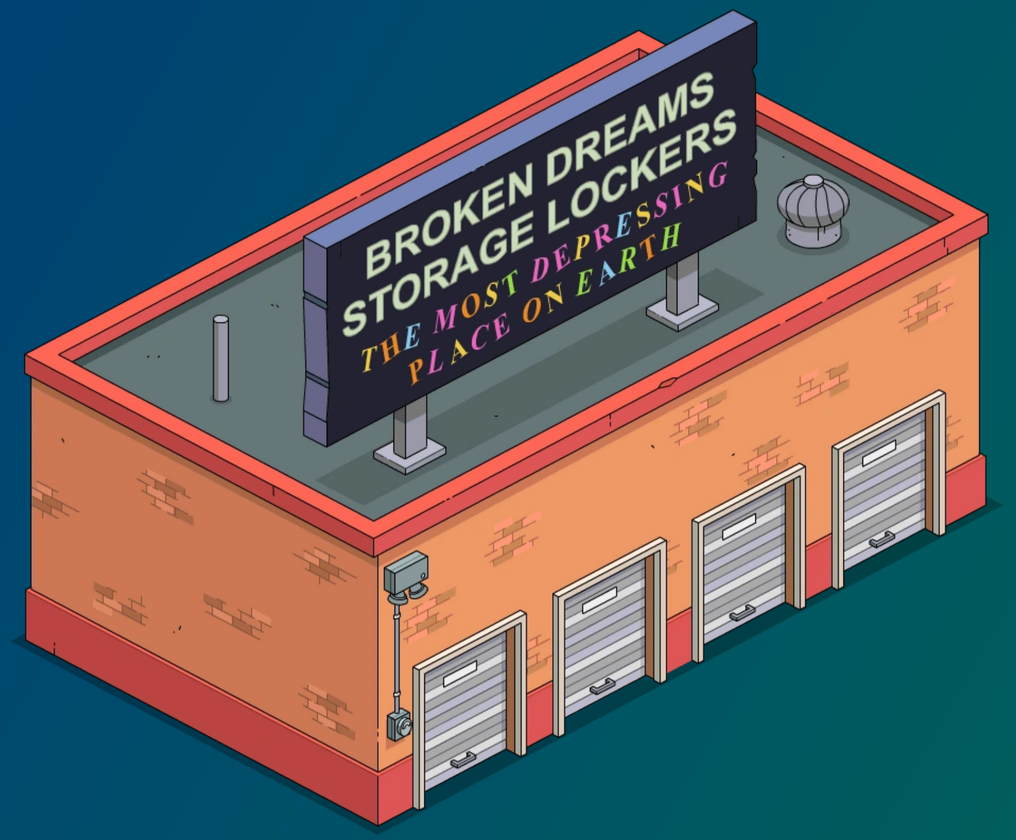

Should be dynamically provisioned

CSI provider enables encryption at rest

High IOPS (SSD or NVMe)

Low Latency

Database performance is as fast as your storage

simplyblock

Storage

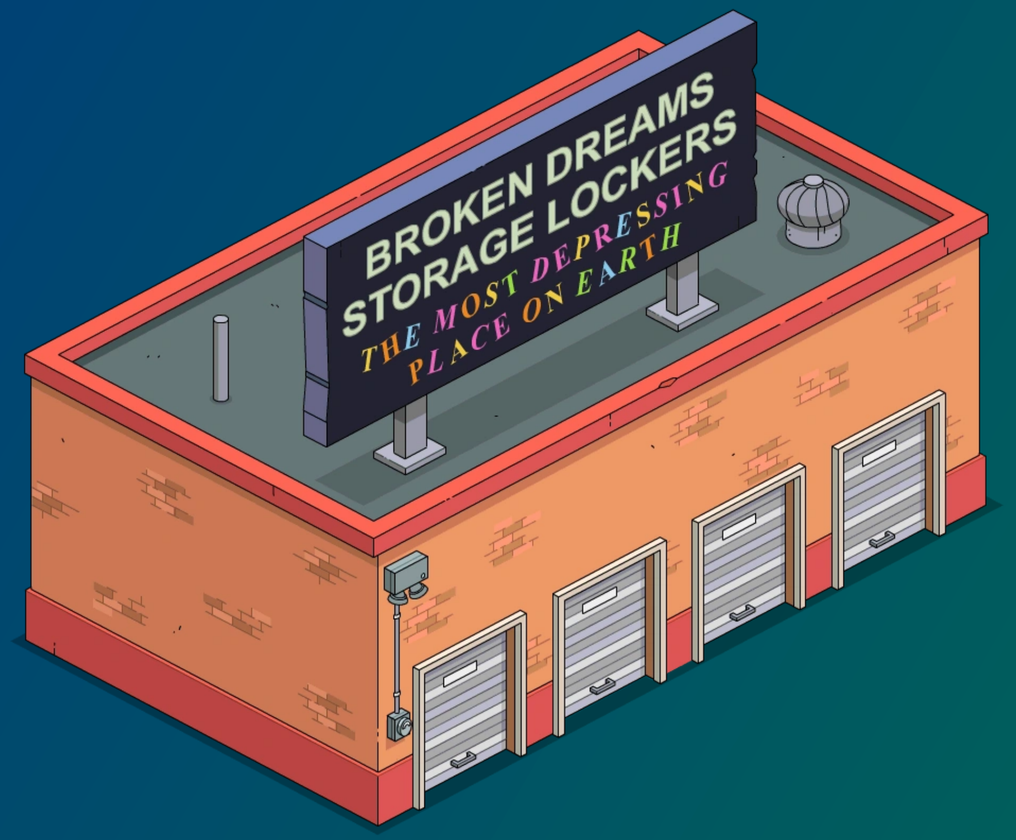

Should be dynamically provisioned

CSI provider enables encryption at rest

High IOPS (SSD or NVMe)

Low Latency

Database performance is as fast as your storage

I’d recommend a disaggregated storage!

simplyblock

Kubernetes CSI

simplyblock

CSI = Container Storage Interface

CSI = Container Storage Interface

```
interface CSIDriver {

    ListVolumes()
    ProvisionVolume()
    DeleteVolume()
    AttachVolume()
    ExpandVolume()
    SnapshotVolume()
    DeleteSnapshot()
    …

}
```

CSI = Container Storage Interface

CSI = Container Storage Interface

Avoids in-tree volume drivers

CSI = Container Storage Interface

Avoids in-tree volume drivers

Independent release cycles

CSI = Container Storage Interface

Avoids in-tree volume drivers

Independent release cycles

Works across Kubernetes versions

CSI = Container Storage Interface

Avoids in-tree volume drivers

Independent release cycles

Works across Kubernetes versions

Enables third-party volume drivers

Kubernetes CSI

simplyblock

Kubernetes CSI

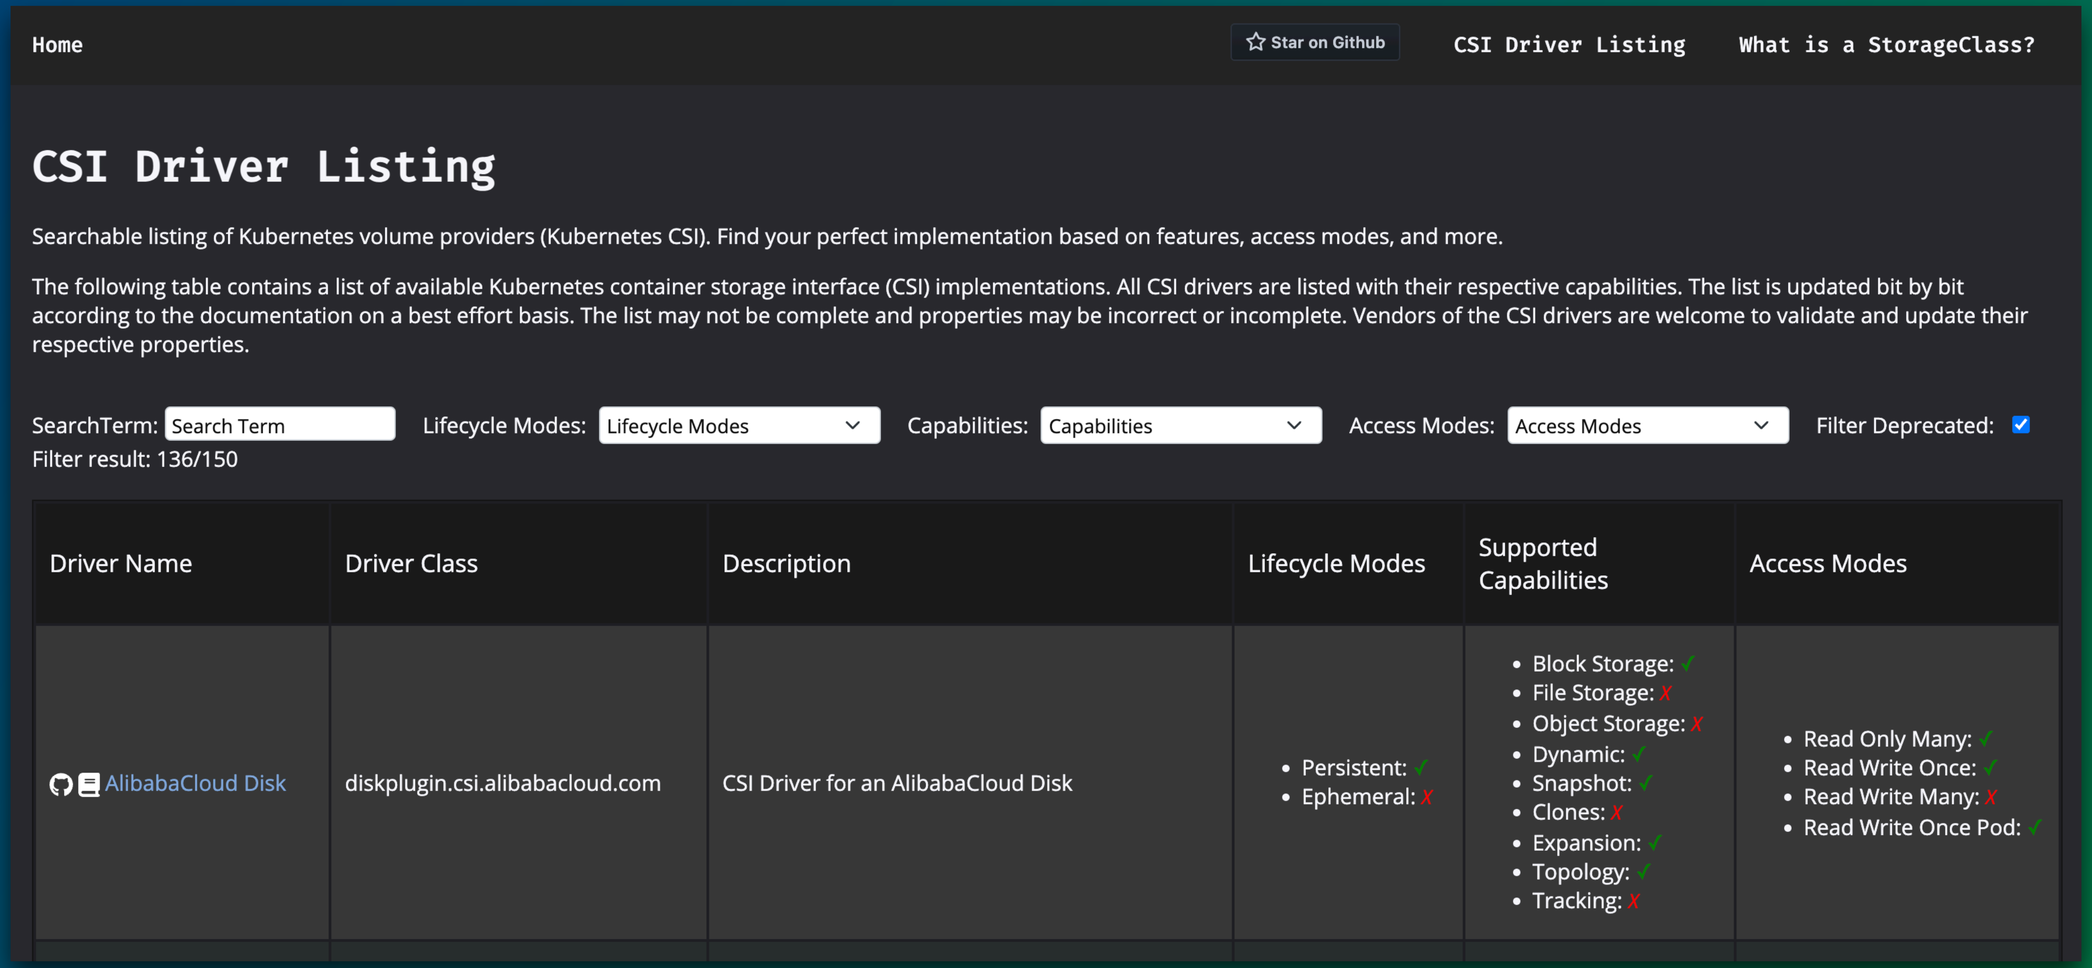

www.storageclass.info/csidrivers

simplyblock

Demo

Intelligent Cloud-Native Storage Platform

Intelligent Cloud-Native Storage Platform

- Copy-on-write
- Instant Snapshots
- Instant Clones / Branches / Forks (whatever you want to call it)
- NVMe over Fabrics (NVMe/TCP)
- Optimized for Database and I/O-heavy Workloads
- Disaster Recovery
- A Block Storage Device (basically a hard disk)

How to Choose?

How to Choose?

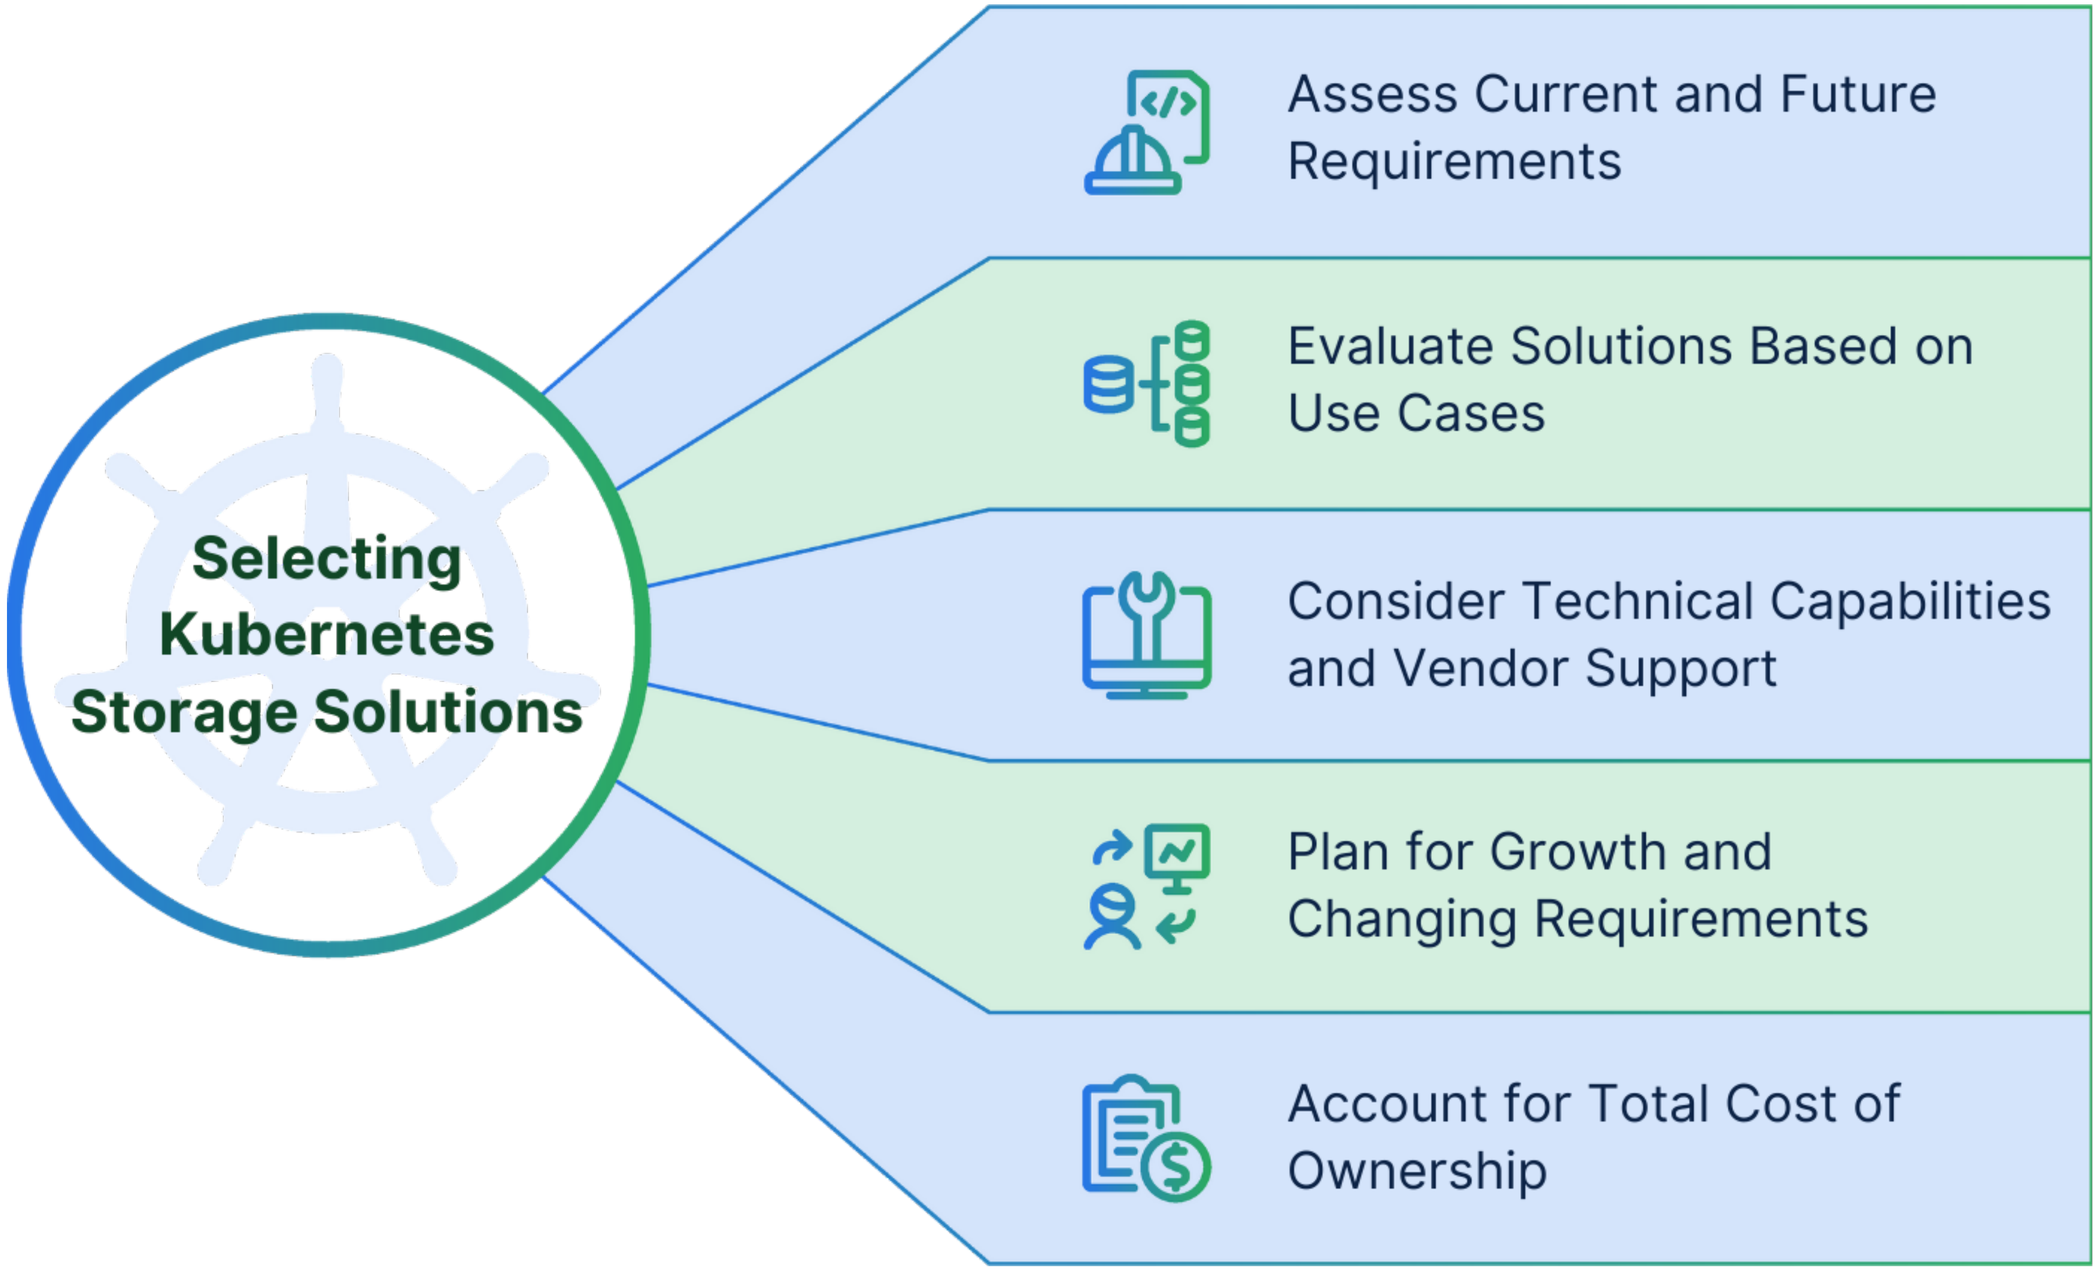

Copy-On-Write: The Hero of Storage

Nikita Shamgunov • Following

Entrepreneur (@neondatabase, @singlestore). Investor (@khoslaventures)

2d • Edited •

14.5 TB storage for the price of 2.9 TB. That's the power of copy-on-write.

(screenshot is of a logistics/IoT customer with a very common **Neon** Branching setup)

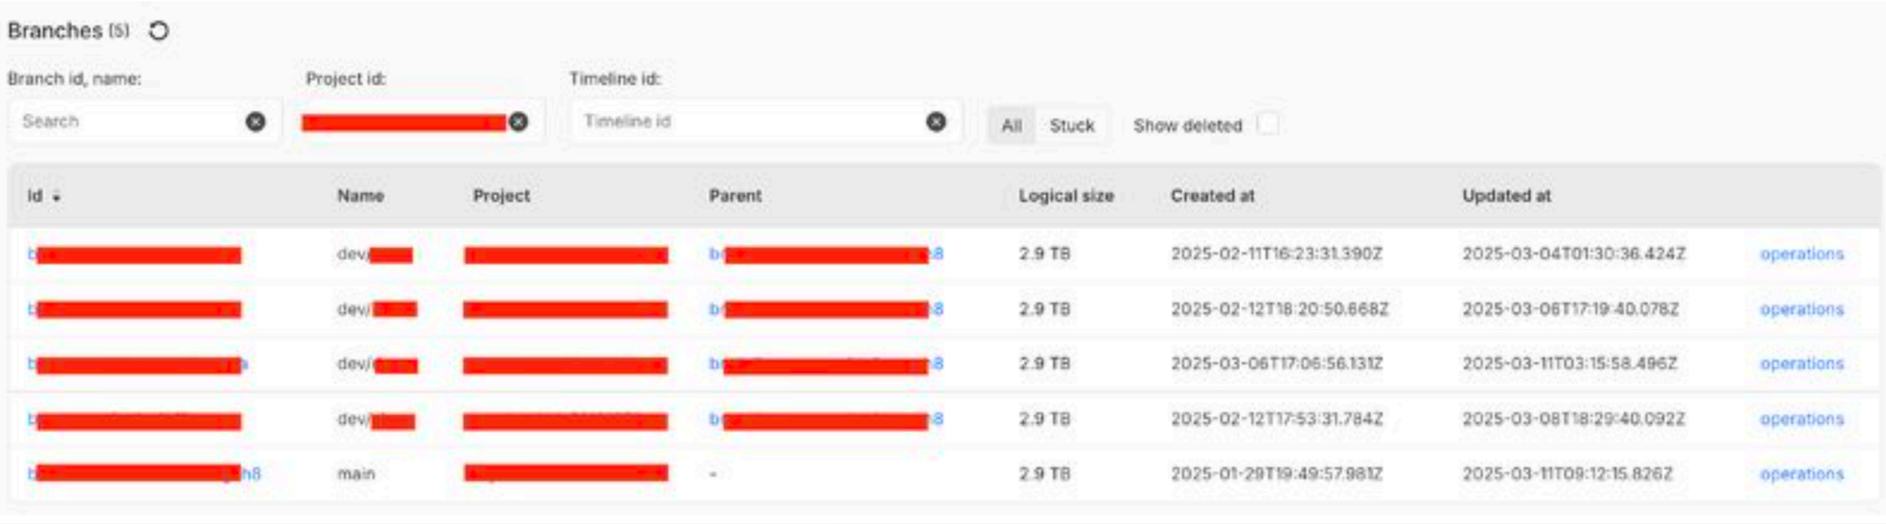

simplyblock

Trust me, I’m Kelsey!

Kelsey Hightower
@kelseyhightower

Soham Dasgupta @thesobercoder · Feb 10, 2023

@kelseyhightower Bust a myth for us please - running any sort of database on a Kubernetes instance is bad idea. I've heard this enough times to actually start believing it. #kubernetes #mythbuster

5:21 PM · Feb 10, 2023 · **326.7K** Views

43 149 839 224

https://x.com/kelseyhightower/status/1624081136073994240

simplyblock

Trust me, I'm Kelsey!

Kelsey Hightower
@kelseyhightower

You can run databases on Kubernetes because it's fundamentally the same as running a database on a VM. The biggest challenge is understanding that rubbing Kubernetes on Postgres won't turn it into Cloud SQL. 🧵

> **Soham Dasgupta** 🇮🇳 @thesobercoder · Feb 10, 2023
> @kelseyhightower Bust a myth for us please - running any sort of database on a Kubernetes instance is bad idea. I've heard this enough times to actually start believing it. #kubernetes #mythbuster

5:21 PM · Feb 10, 2023 · **326.7K** Views

43 149 839 224

https://x.com/kelseyhightower/status/1624081136073994240

simplyblock

More Resources

Data on Kubernetes Community: https://dok.community

Data on Kubernetes Whitepaper

@noctarius2k
@noctarius2k@mastodon.online
@noctarius.com

Thank you very much!

Questions?

simplyblock