

2022

1

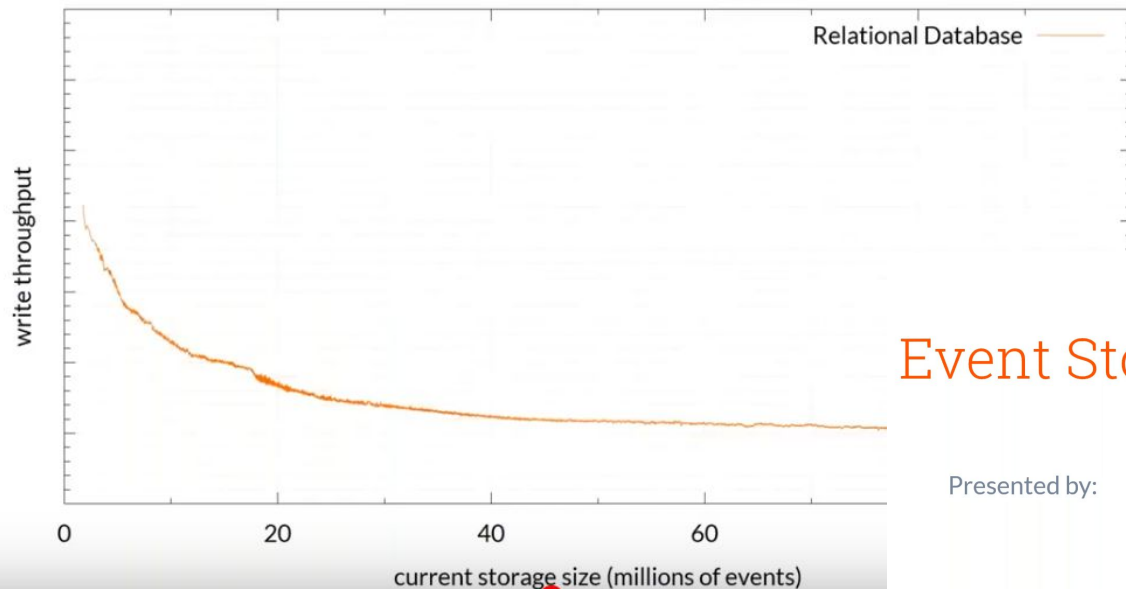
Axon Server Benchmarks

Marco Amann | digital frontiers



digital
frontiers

Motivation



- Talk by Allard
- Is it really this bad?

Webinar Event Storage in AxonServer

Presented by:



Allard Buijze
Founder & CTO, AxonIQ

✉ allard@axoniq.io
🐦 [allardbz](https://twitter.com/allardbz)

AxonIQ

Miniplay



digital
frontiers

Metrics



Throughput

Single Producer
Multiple Producers
Stability over time

Storage space

(Latency)

Contestants



Axon Server



Postgres



mongoDB

Mongodb

Machines

Small Virtual Server

14€/month

Shared

4 Cores
8 GB RAM
160 GB SSD

Beefy Virtual Server

59€/month

Shared

16 Cores
32 GB RAM
300 GB SSD

Bare Metal Server

> 1300€/month

Dedicated

80 Cores
256 GB RAM
2x 3.4 TB NVME

Dev Machine

N.A.

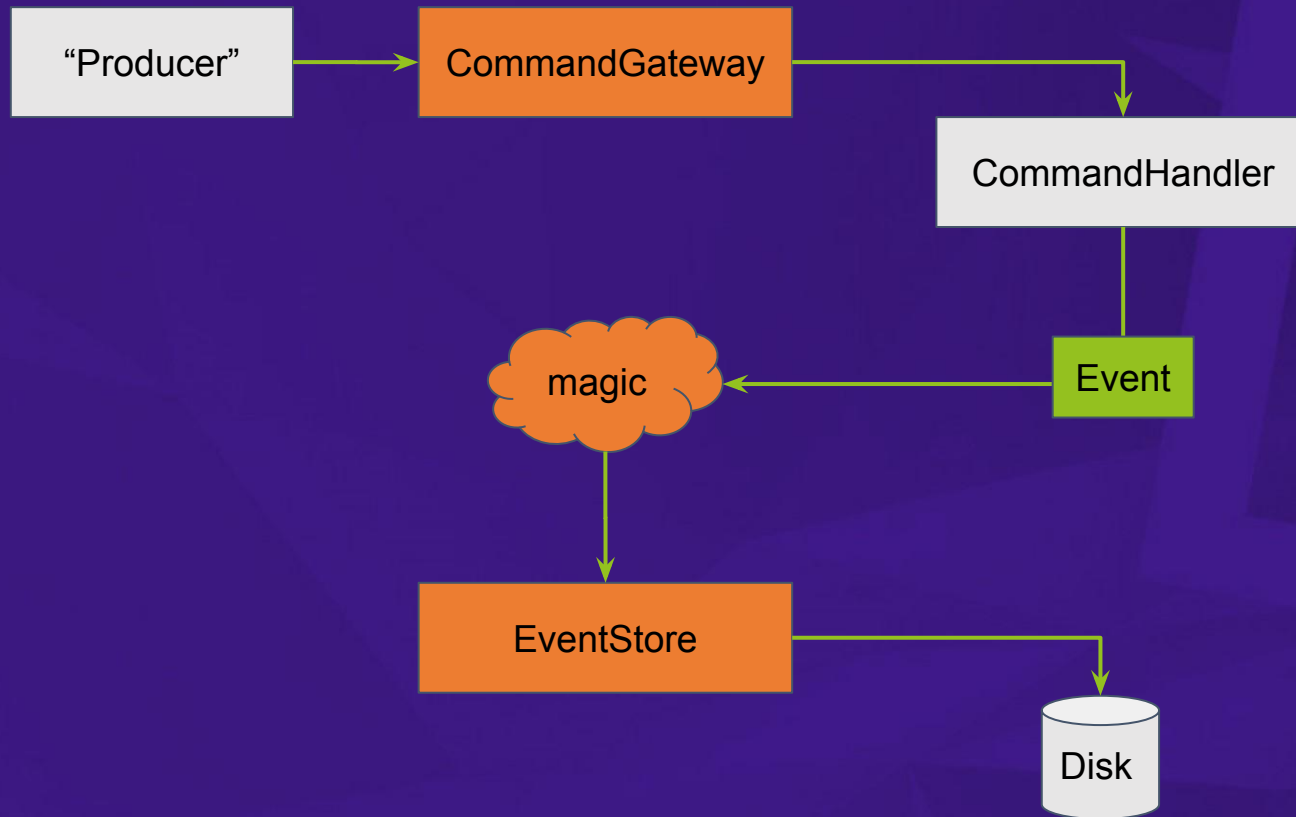
Dedicated

24 Cores
32 GB RAM
1x 256 GB M2 SSD

Completely different
system



Application



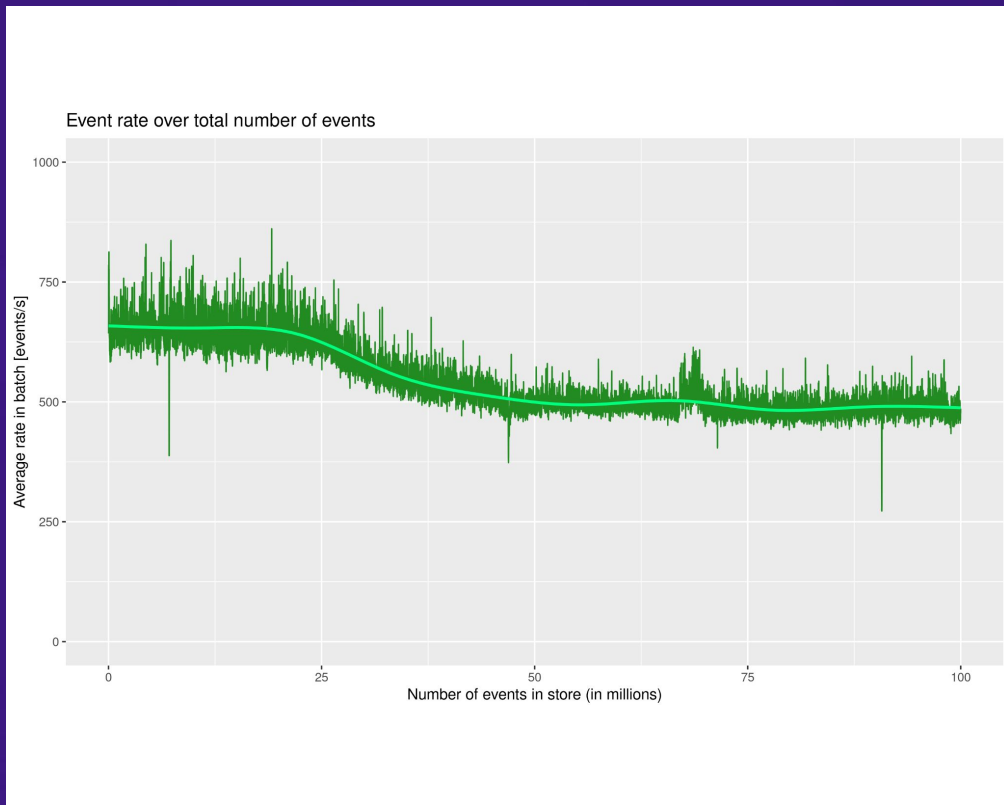
One or more producers?

Can we produce another command without knowing if the current one failed?

Stability

All these benchmarks ran on the small virtual server

Postgres



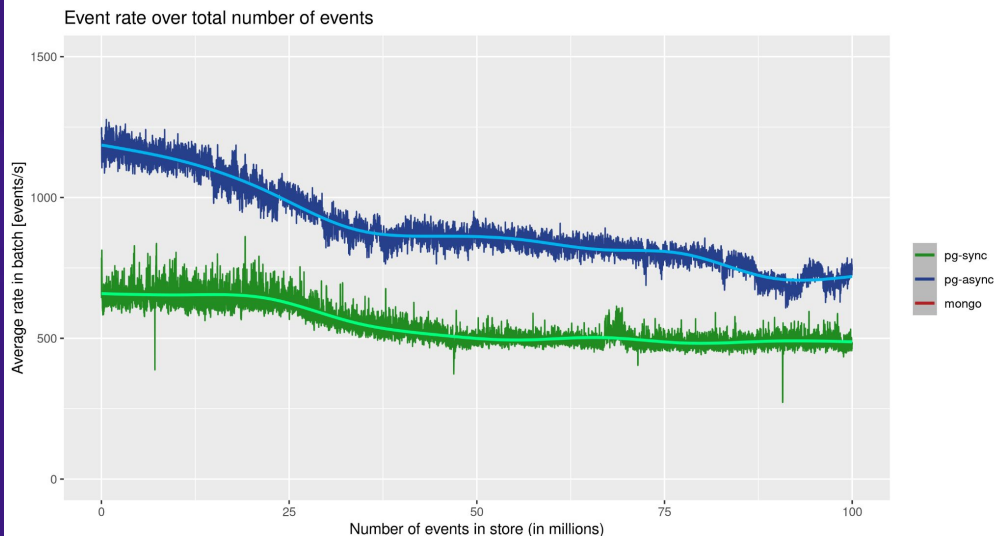
100 Million Events (about 2 days)

Averages of batches (1000 events each)

Executed on the small virtual server

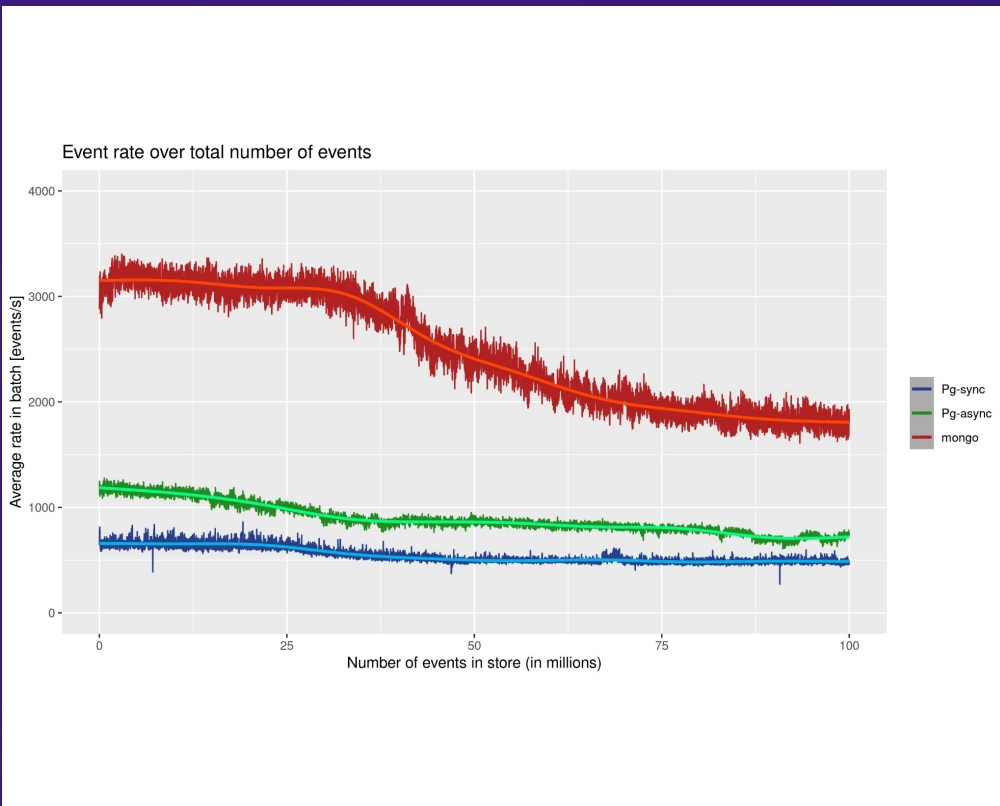
Postgres: Sync vs Async

Async: disabled fsync



Executed on the small virtual server

Postgres vs Mongo

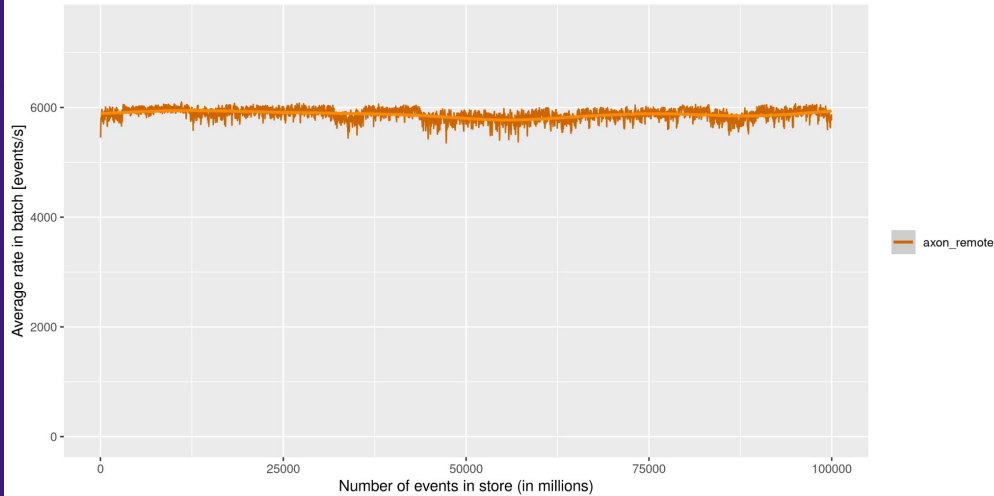


Mongo was much faster on remote machine.

On the dev machine, async-pg and mongo were equally fast.

Executed on the small virtual server

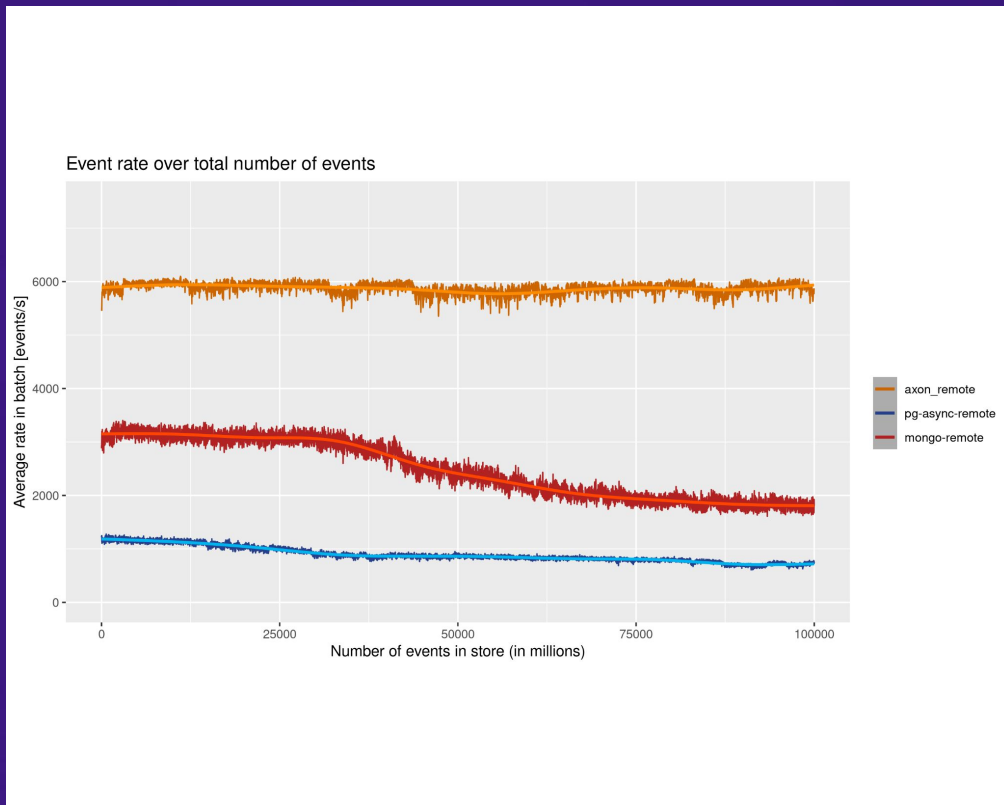
Event rate over total number of events



Nearly no decrease!

Executed on the small virtual server

Axon vs Mongo vs Postgres



Axon shows nearly no decrease over time.

Axon is much faster than postgres with a **single** producer (see slide 23).

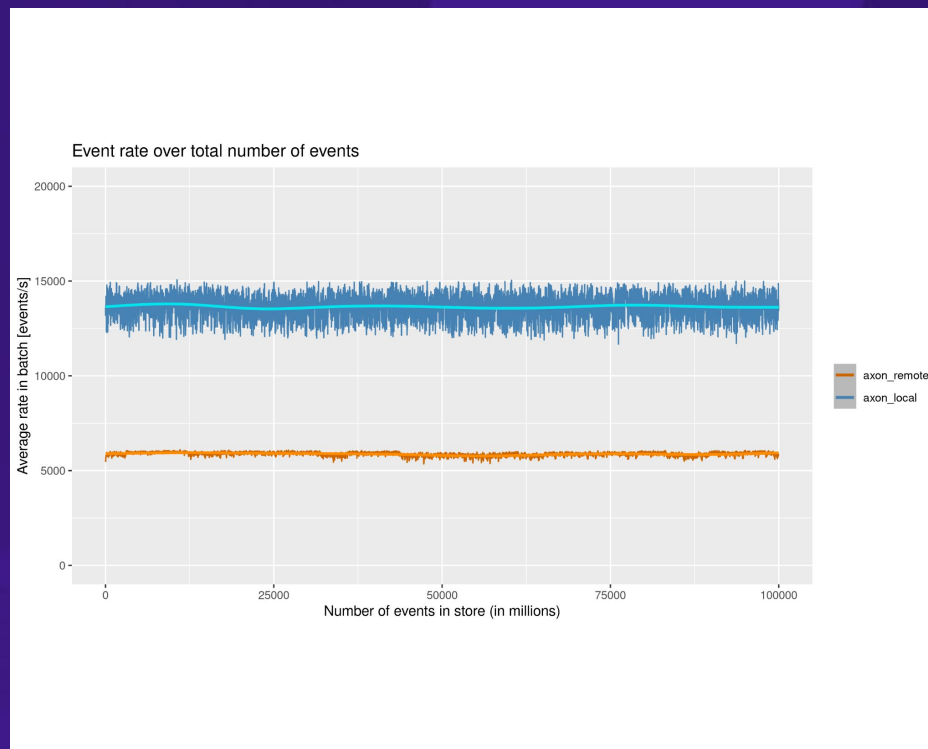
Executed on the small virtual server

Axon on better Hardware

Axon can make use of faster hardware.

Let's see how the others keep up!

But first: storage



Storage

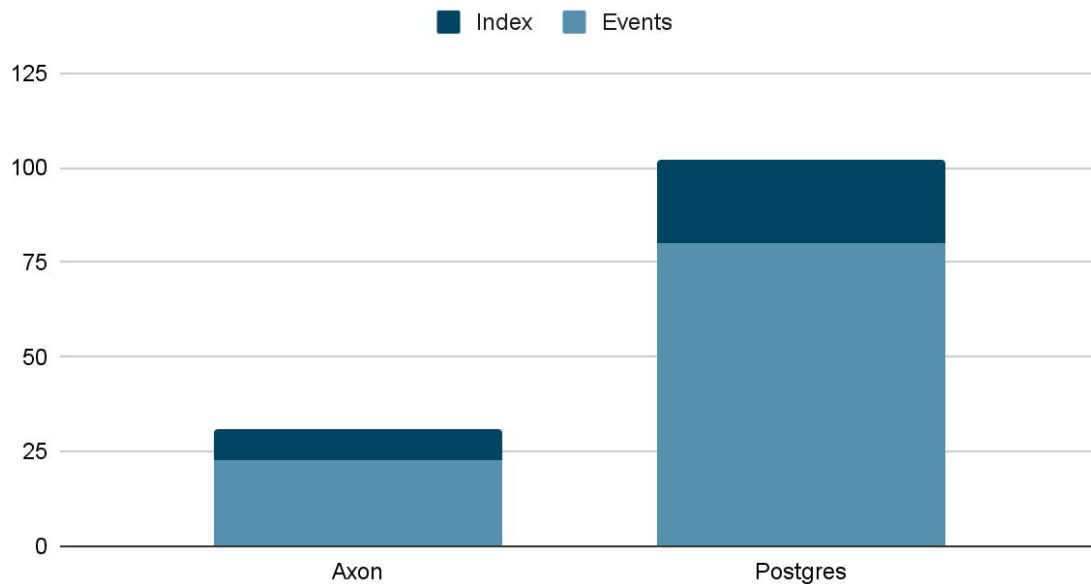
Storage size

Note:

- Small Events
- No tuning for space

Postgres storage cleanup (
VACUUM FULL) did never
complete

Storage (GB) for 100 million events



Storage

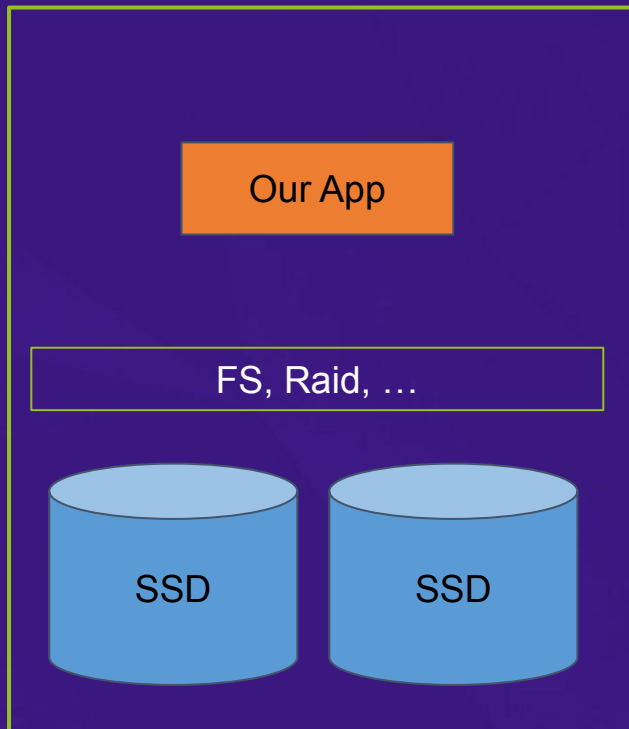
Latency and Throughput

For EventStores, storage throughput is no concern in most cases.

Latency:

- Read
- Write
- Sync

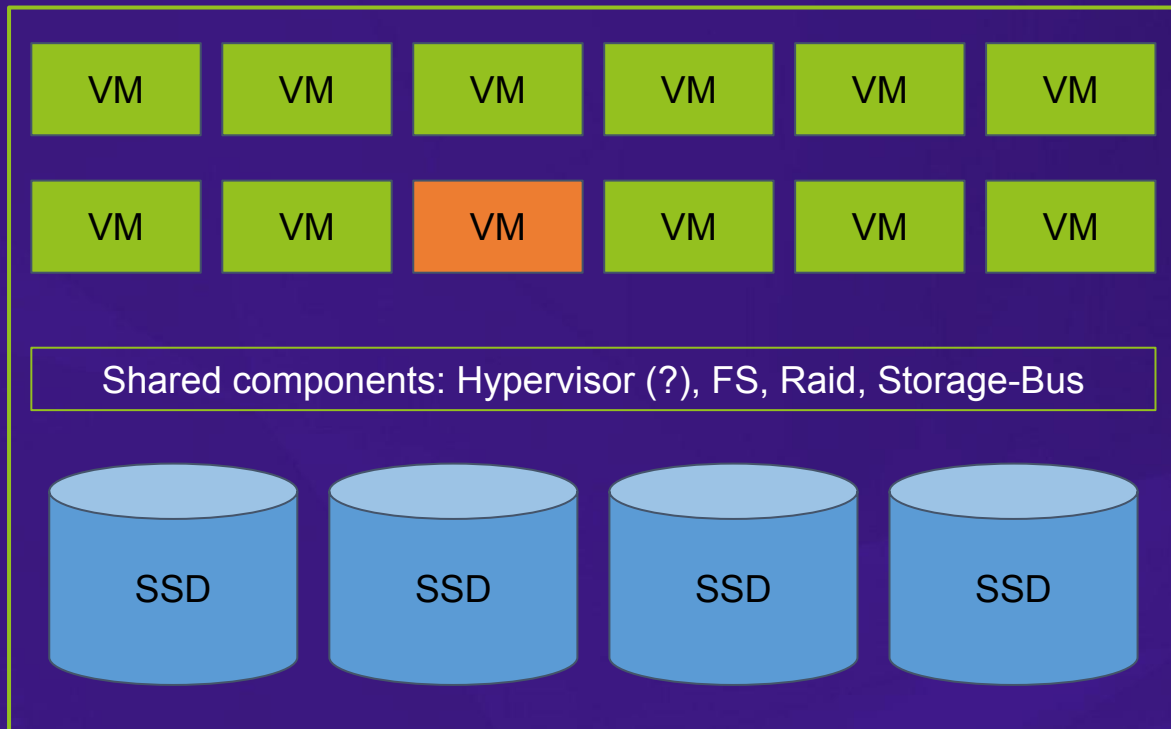
Storage



Dedicated Server

- Fast
- Stable performance
- Prohibitively expensive
 - Often better conditions with 1y-5y contracts

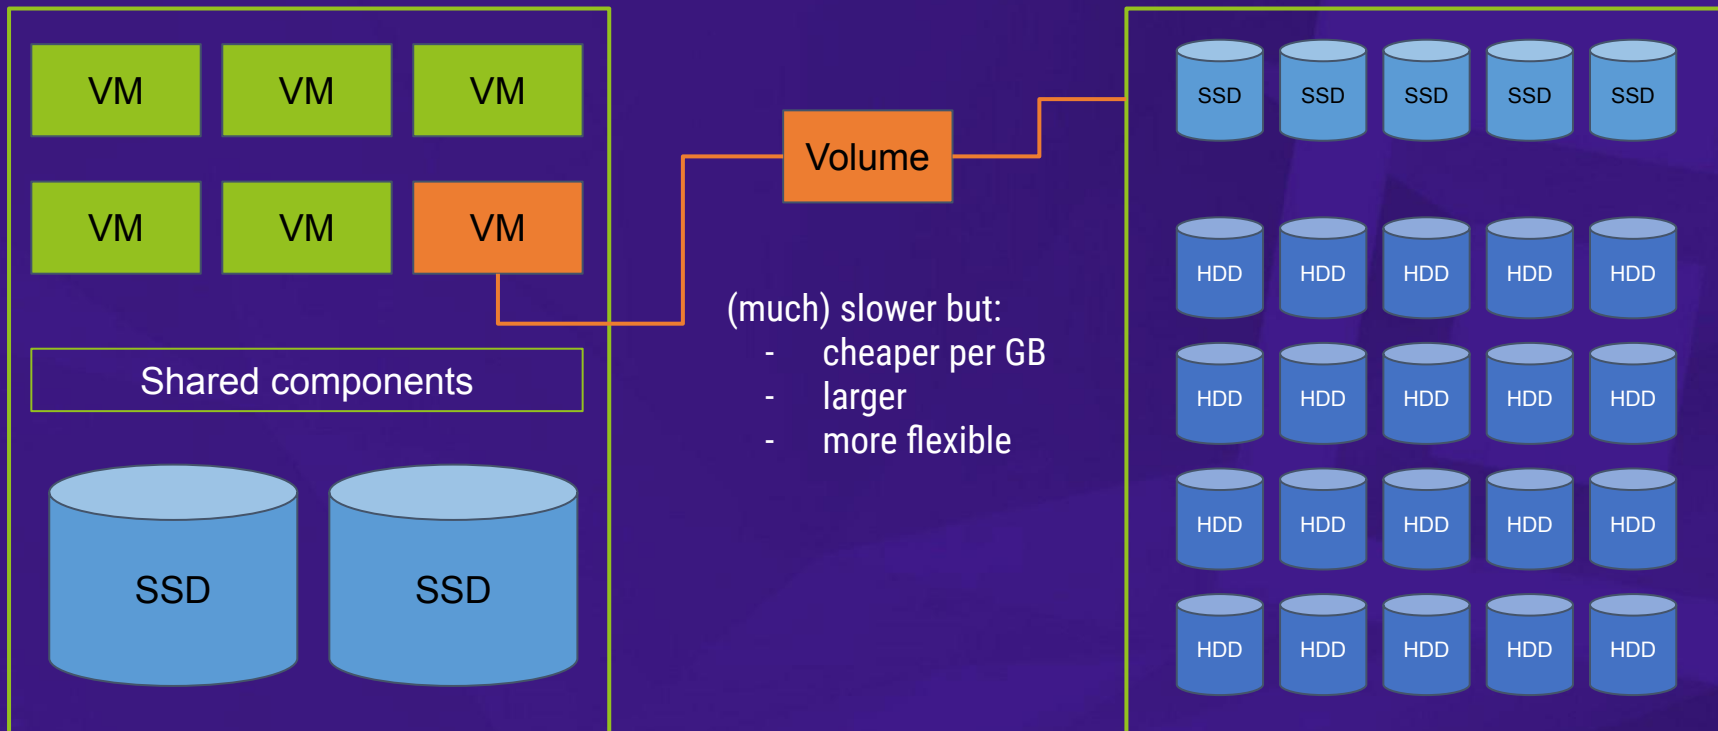
Storage



Shared Server

- Fast most of the time
- Possibly fluctuating performance
- Cheap
- Limited configurations

Storage



Fsync performance

Example: fsync performance

Dedicated:	47473.648 ops/sec	21 usecs/op
Volume:	368.807 ops/sec	2711 usecs/op

Factor: 129

Faster Hardware

All of the following benchmarks ran with 500K events

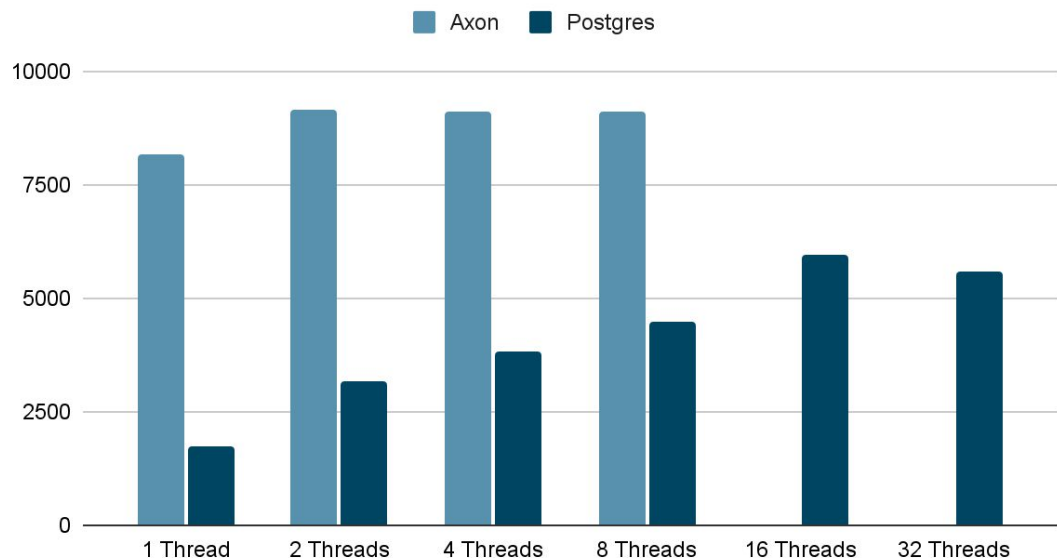
Multiple Threads

Increasing threads did not boost axon performance that much, since the CommandBus already allows for async processing.

Multiple threads with PG in Axon Framework has to be done by hand!

Can we allow for parallel/async processing?

Performance with different Threads



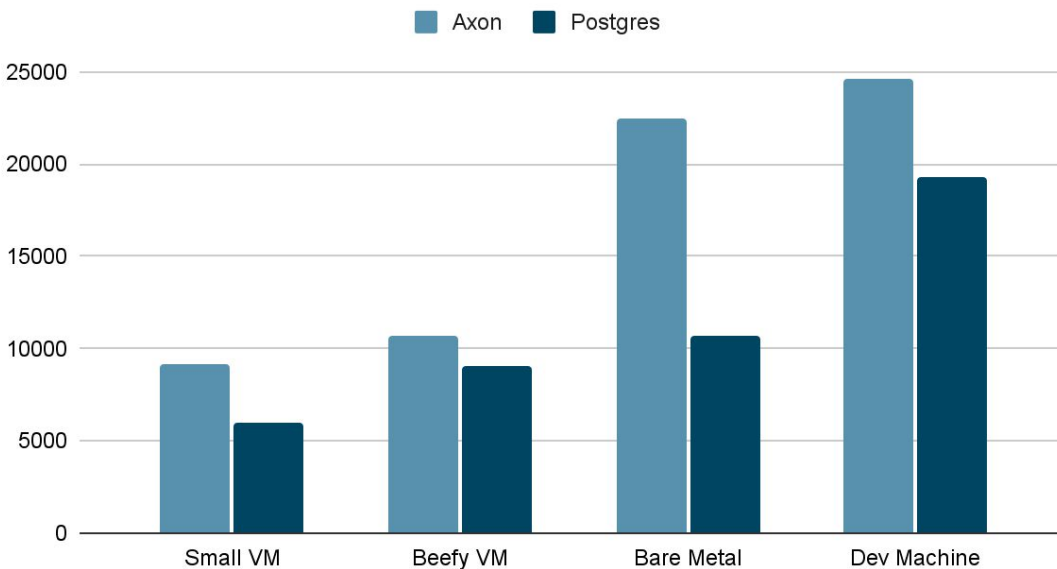
Executed on the large virtual server



digital
frontiers

Axon can make use of more/faster
RAM and CPU

Maximum Throughput by Server



Marco Amann

12 min read · Draft in Digital Frontiers—Das Blog

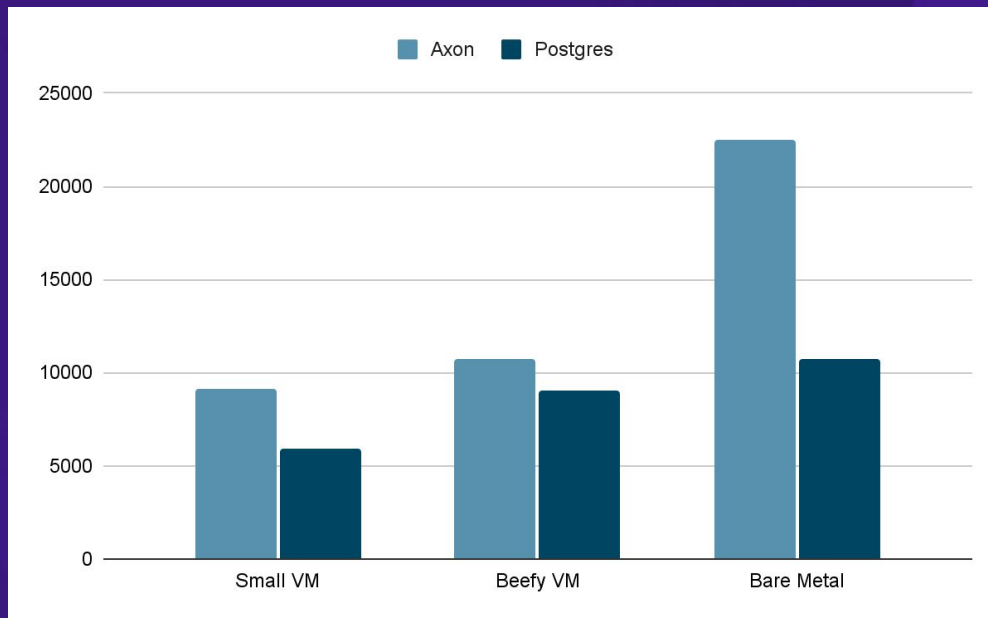
...

That one time I had a Postgres cluster on my veranda

This is the story of a performance problem with Postgres that required tearing apart every part of the tech stack, scratching the realms of absurdity.



Maximum throughput of different servers



14€/month

59€/month

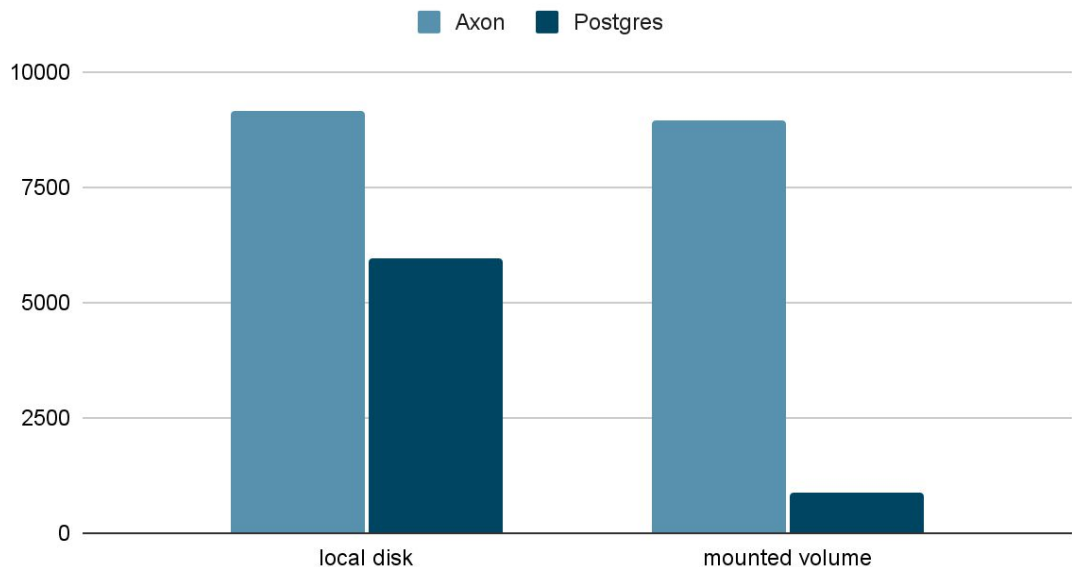
>1300€/month

Impact of slow disks

Axon showed nearly no effect when running on remote storage

Postgres performance was abysmal

Impact of Slow Disk



Executed on the large virtual server

Available Storage

**With parallel workloads, Axon Server
can run on cheap remote storage volumes**

VMs (at Hetzner): Max 960GB local storage
Volumes for 0.0476€ per GB/month, up to 10TB without contract

If you need low-latency (local)
storage, you need to buy larger VMs

VM with 10TB storage at AWS is at least 2100€/month

Summary


Axon

- Steady performance
- Less impacted by slow storage
- Can make use of more CPU
 - also: requires more CPU and RAM

Postgres

- Axon Framework connector requires many concurrent transactions
- With lots of tuning and on certain hardware, 75% of Axon performance is possible
- Less compact (default) storage representation: roughly x4



 @amann_dev

digitalfrontiers.de
blog.digitalfrontiers.de

What's the main driver of your server costs?