





Postgres on Metal

How **Postgres** typically runs and scales in the cloud



More compute doesn't solve **every** Postgres bottleneck.



Latency

Storage I/O


Availability



Convex reduced p99 query latency by over 50%


AWS Aurora

PlanetScale Postgres



AWS Aurora

PlanetScale Postgres



Query p99

10–15 ms → 5–7 ms

Batch-commit p99

75–200 ms → ~20 ms

"p50 is the new p99."

Jamie Turner | CEO of Context

Intercom: Faster Inbox, 60%+ Lower Cost



Slow Inbox at peak load

Intercom: Faster Inbox, 60%+ Lower Cost




More headroom → Higher cost



Intercom: Faster Inbox, 60%+ Lower Cost

Previous EBS io2 setup


PlanetScale Metal




Hourly database cost

What is Metal?

Network-attached storage (AWS)




PlanetScale Metal




More QPS & Lower p99 on Metal

Throughput (QPS)




p99 latency



PlanetScale Postgres Architecture


99.99%
single-region HA SLA



Neki: Horizontal Scaling for Postgres



118.5M queries / second



1.22 PiB
data

512
shards

6.06 ms
router p99

PlanetScale on your cloud



Google Cloud

PlanetScale-hosted



Your cloud account (BYOC)



Enterprise

Your database **operated by PlanetScale**

Migrate to PlanetScale



Self-service migration



 PlanetScale-led migration

Technical discovery + migration assessment