

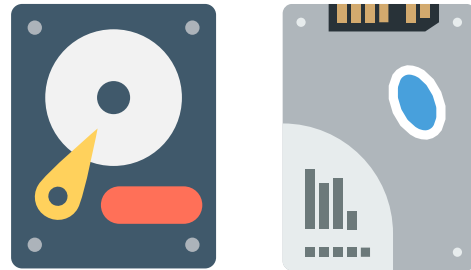
Make Your PostgreSQL 10x Faster on Cloud in Minutes

Sangwook Kim

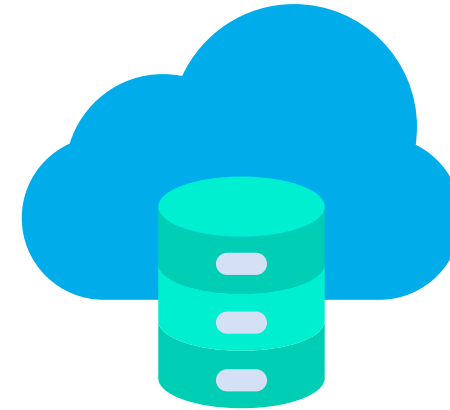
apposha

Physical vs. Cloud Storage

All storages are not created equal



≠



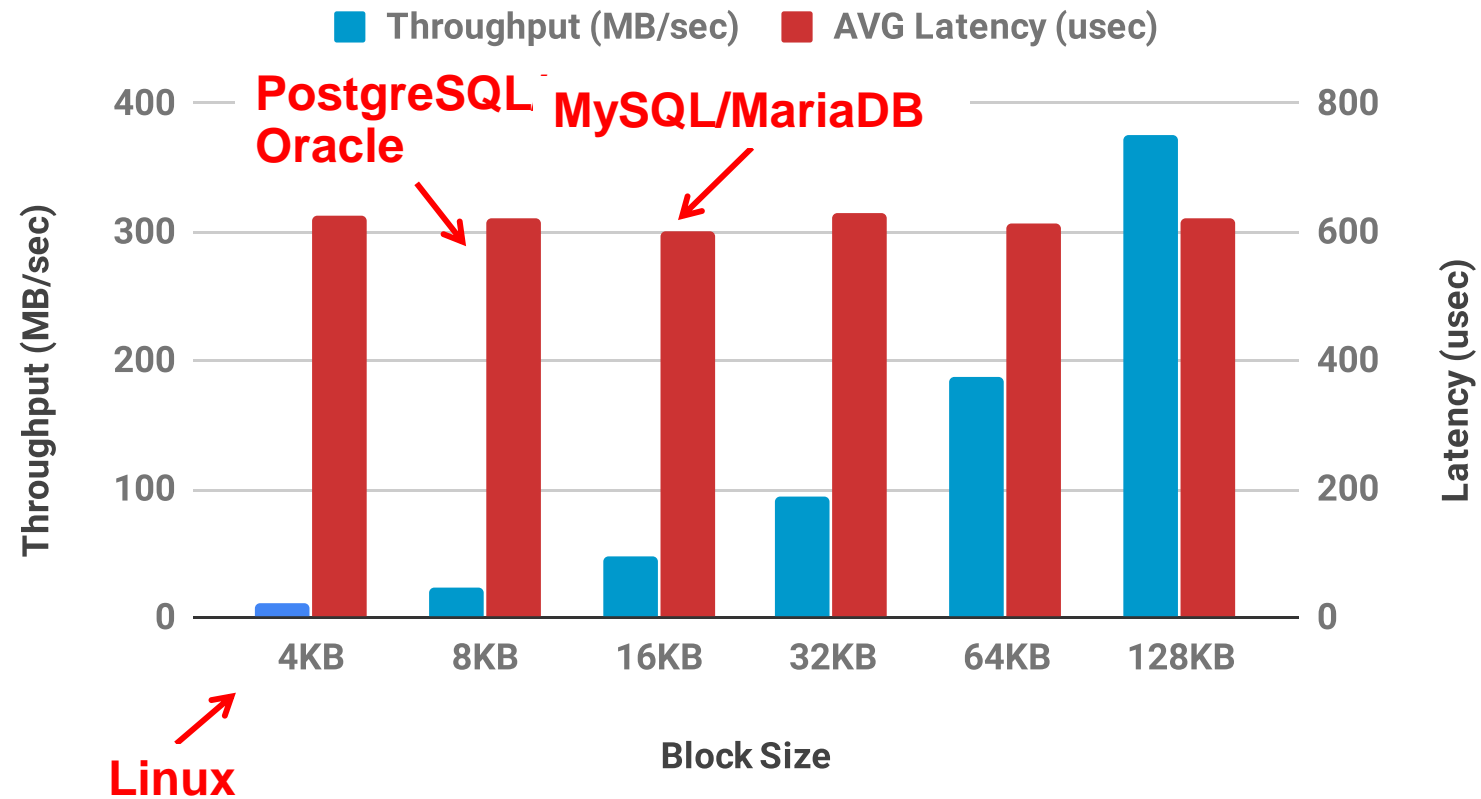
Purchase H/W
HW-controlled

Pay-as-you-go
SW-controlled

Example: AWS SSD Storage

I/O below 256 KiB consumes one IOPS

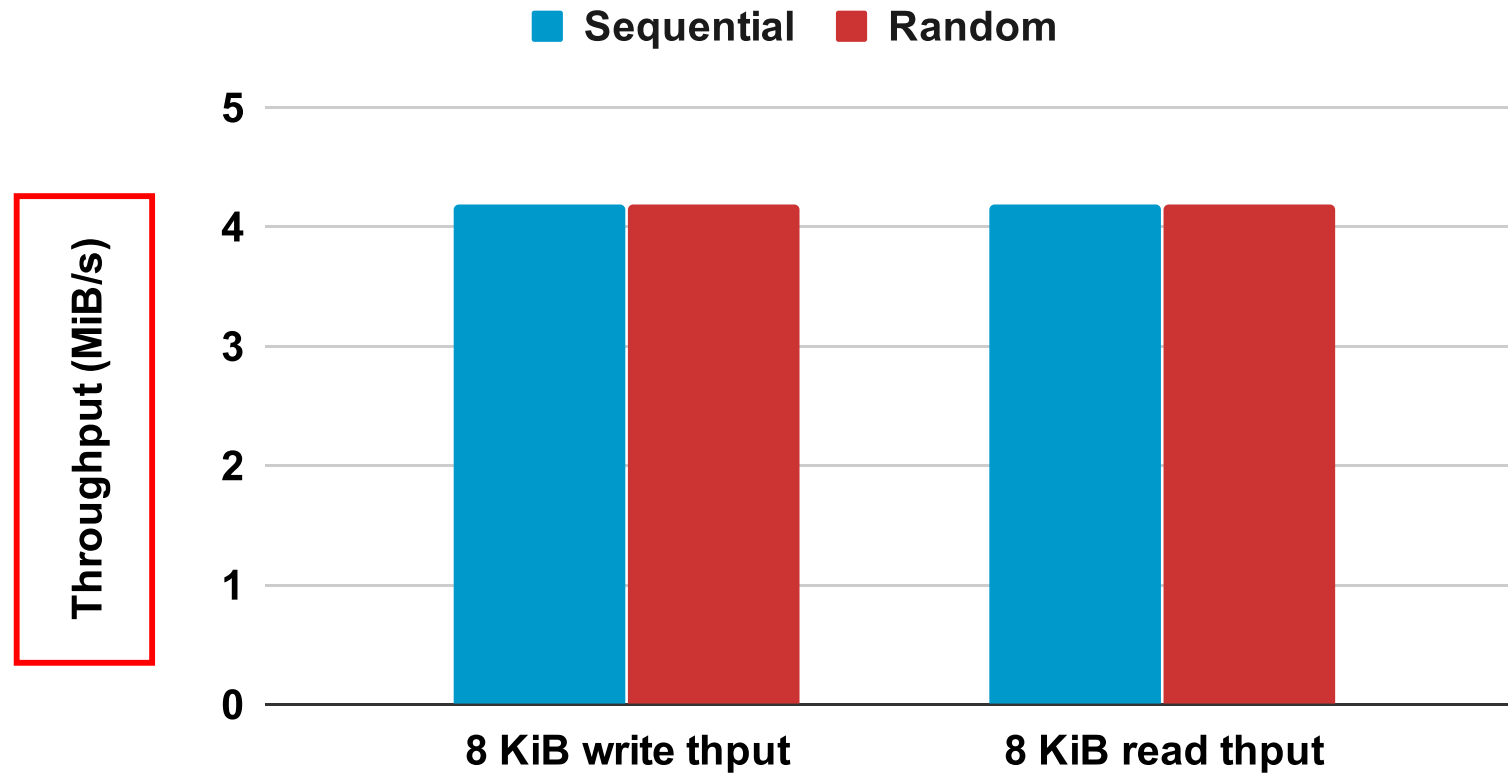
FIO random I/O test with 3000 IOPS AWS EBS



Example: AWS HDD Storage (1)

Sequential and random I/O show identical performance

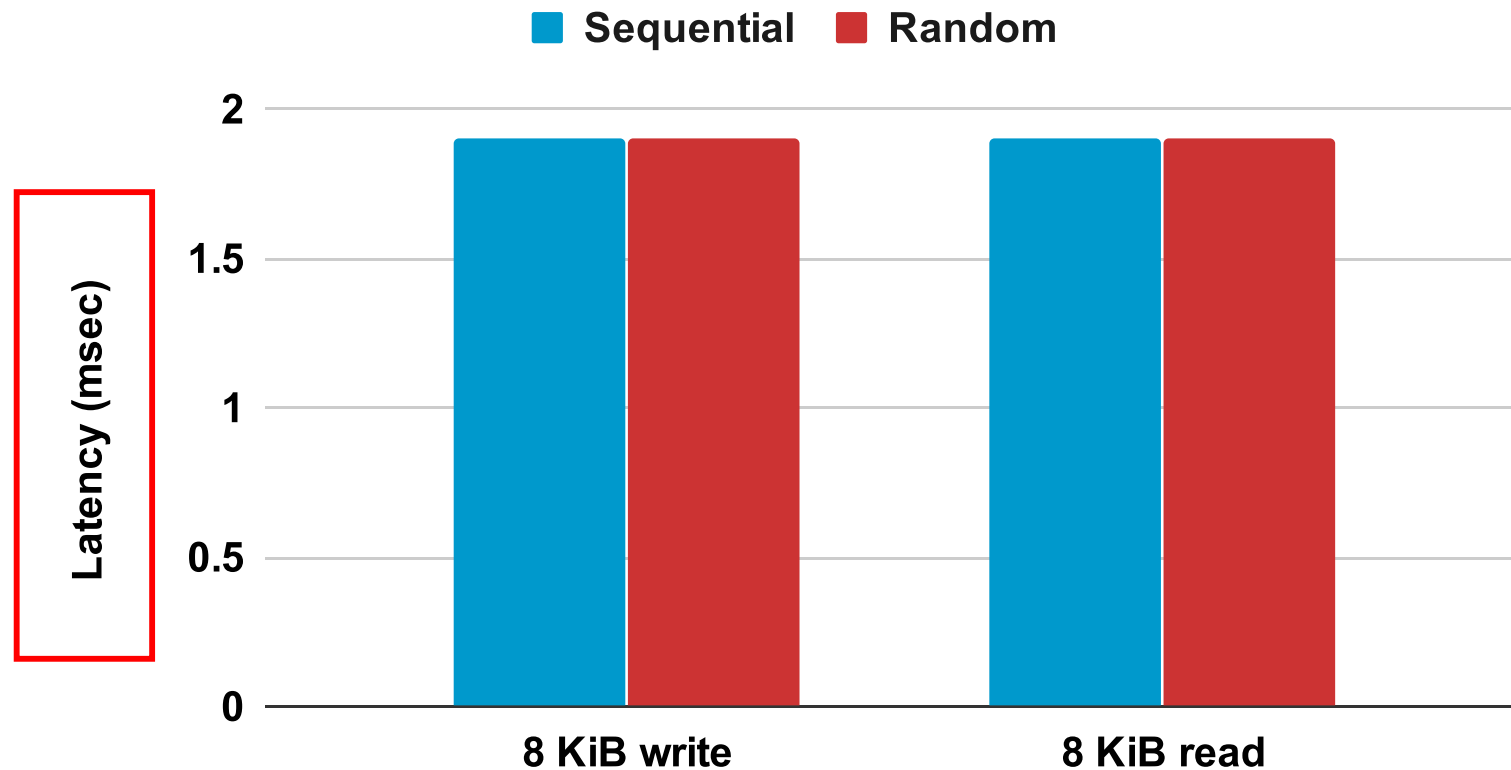
FIO throughput with 16 TiB AWS HDD EBS



Example: AWS HDD Storage (2)

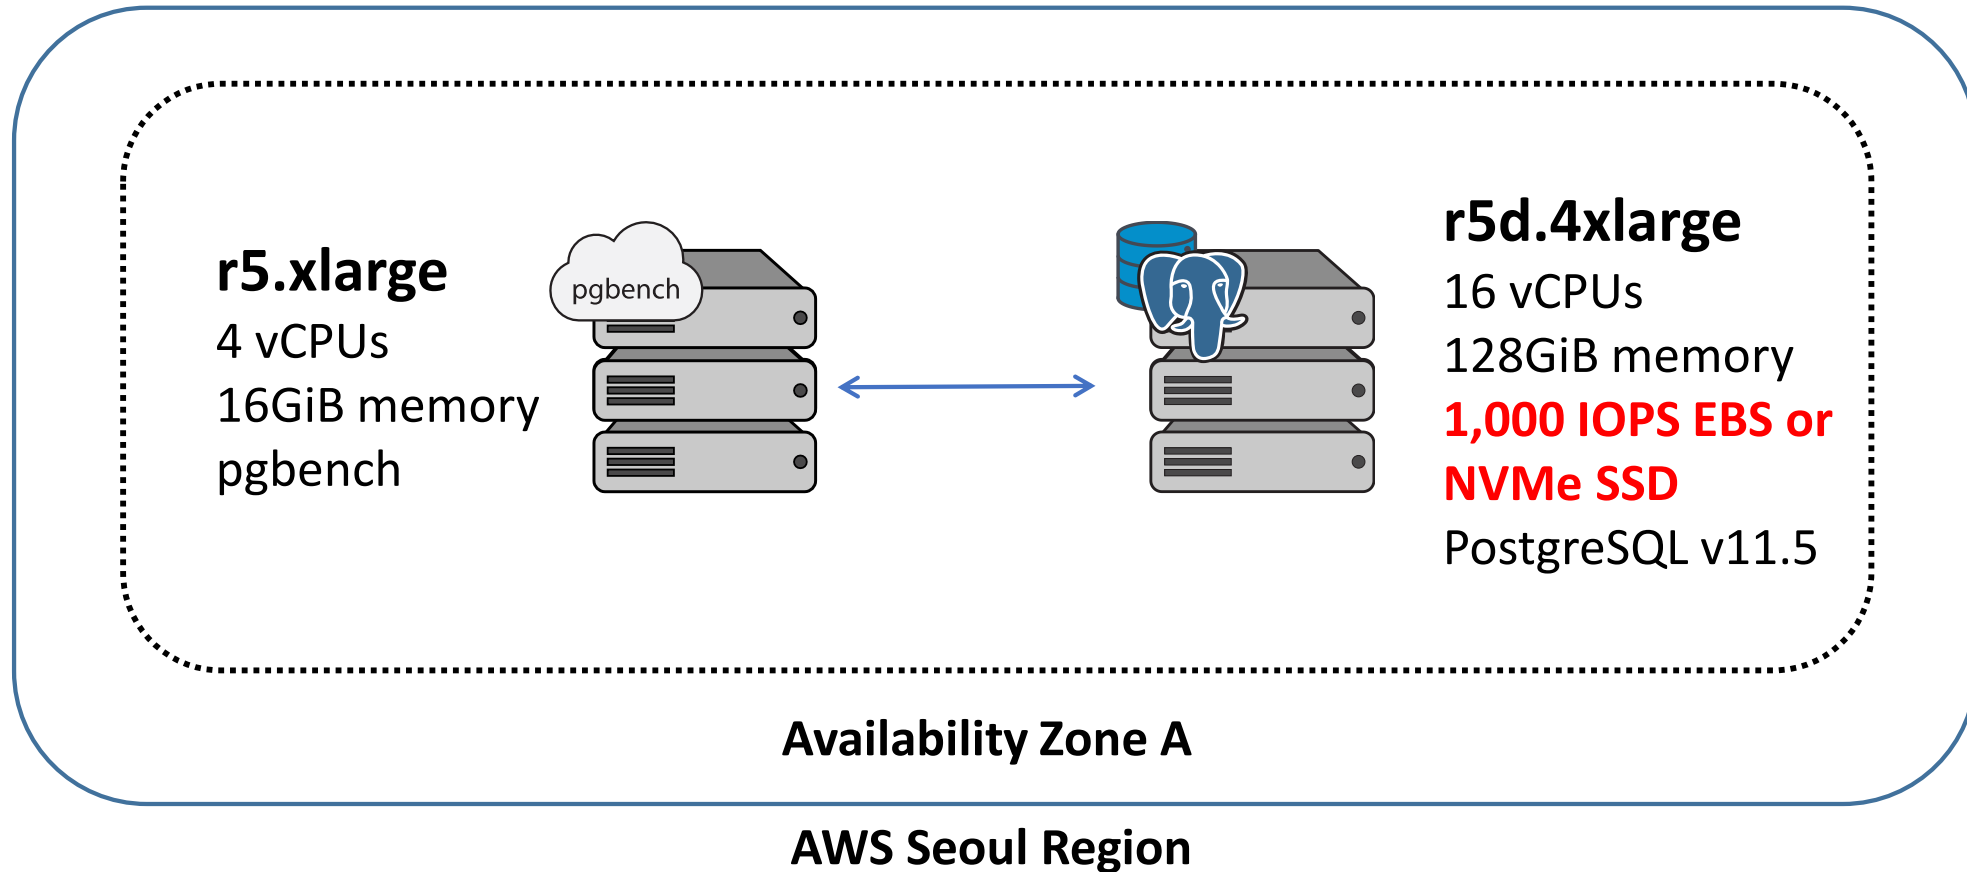
Sequential and random I/O show identical performance

FIO latency with 16 TiB AWS HDD EBS



Illustrative Experiment (1)

Test setup on AWS cloud



DB Type

[what is this?](#)

Total Memory (RAM)

[what is this?](#)

Number of CPUs

[what is this?](#)

Number of Connections

[what is this?](#)

Data Storage

[what is this?](#)

```
# this tool not being optimal
# for very high memory systems
```

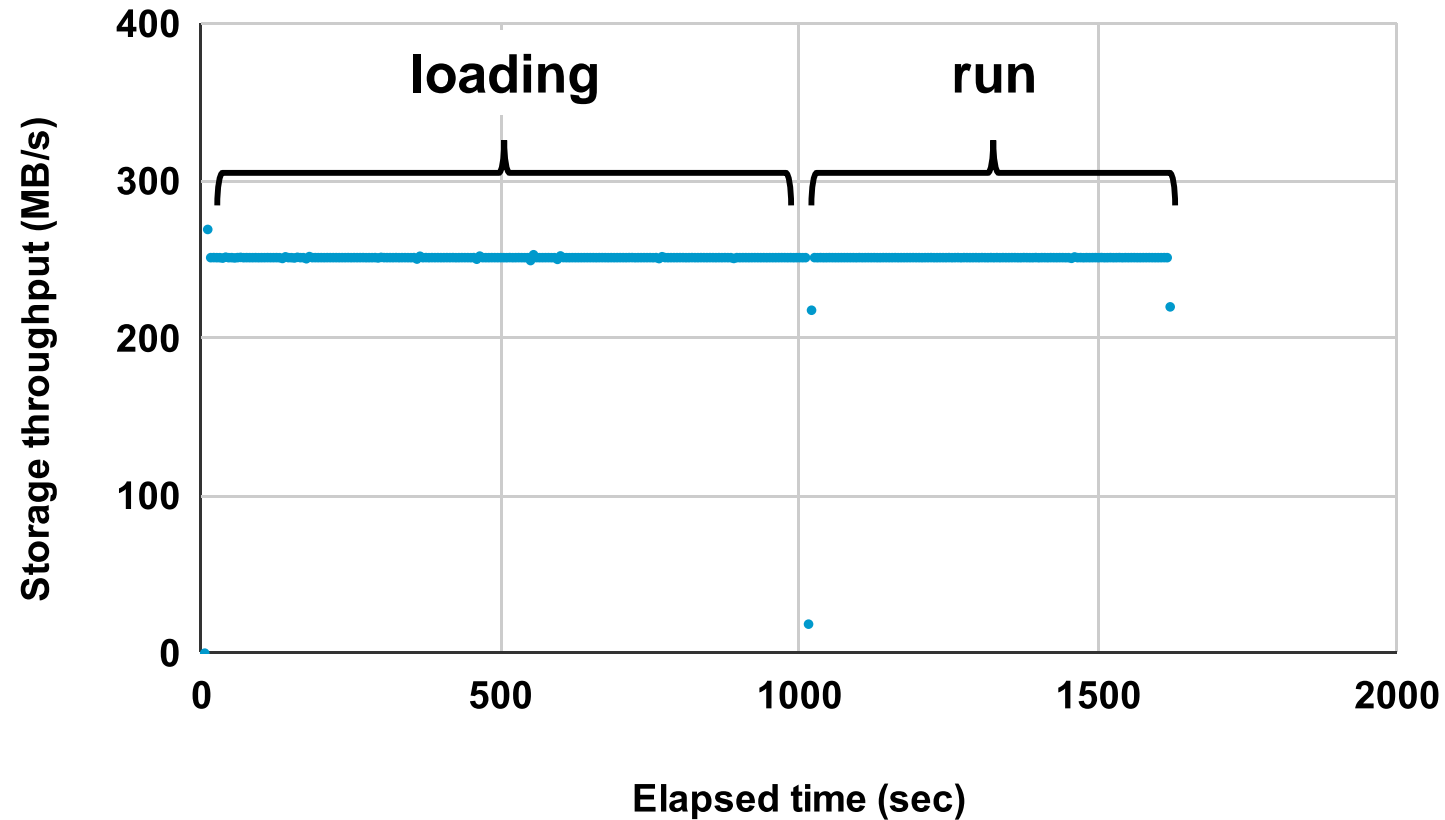
```
# DB Version: 11
# OS Type: linux
# DB Type: oltp
# Total Memory (RAM): 128 GB
# CPUs num: 16
# Connections num: 500
# Data Storage: ssd
```

```
max_connections = 500
shared_buffers = 32GB
effective_cache_size = 96GB
maintenance_work_mem = 2GB
checkpoint_completion_target = 0.9
wal_buffers = 16MB
default_statistics_target = 100
random_page_cost = 1.1
effective_io_concurrency = 200
work_mem = 8388kB
min_wal_size = 2GB
max_wal_size = 4GB
max_worker_processes = 16
max_parallel_workers_per_gather = 8
max_parallel_workers = 16
```

Illustrative Experiment (2)

PostgreSQL can fully utilize SSD

pgbench, 5K scale, 100 clients, 300 GiB NVMe SSD

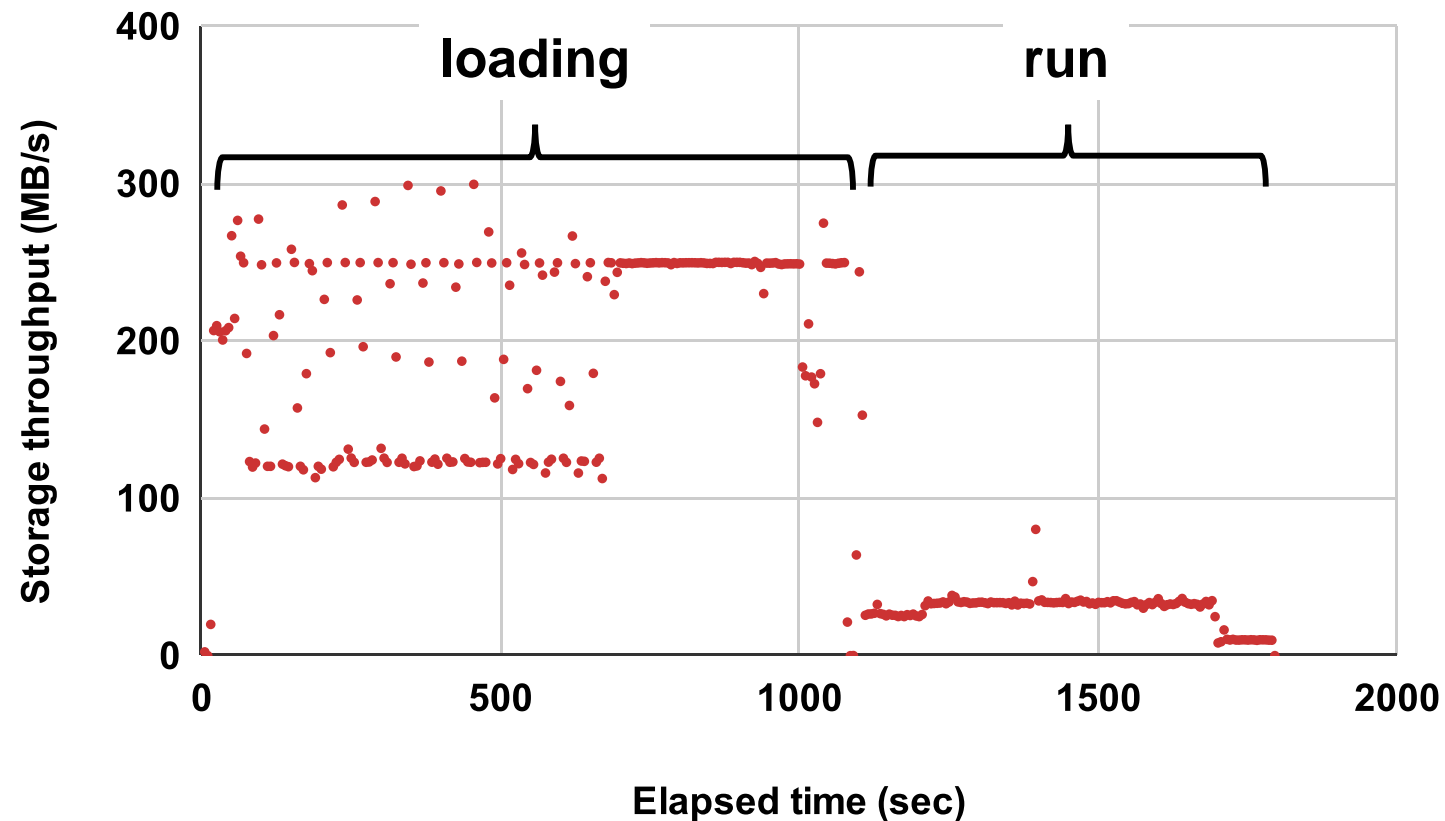


13,000 TPS

Illustrative Experiment (3)

PostgreSQL **cannot fully utilize cloud storage**

pgbench, 5K scale, 100 clients, 300 GiB EBS SSD (1K IOPS)



13,000 TPS

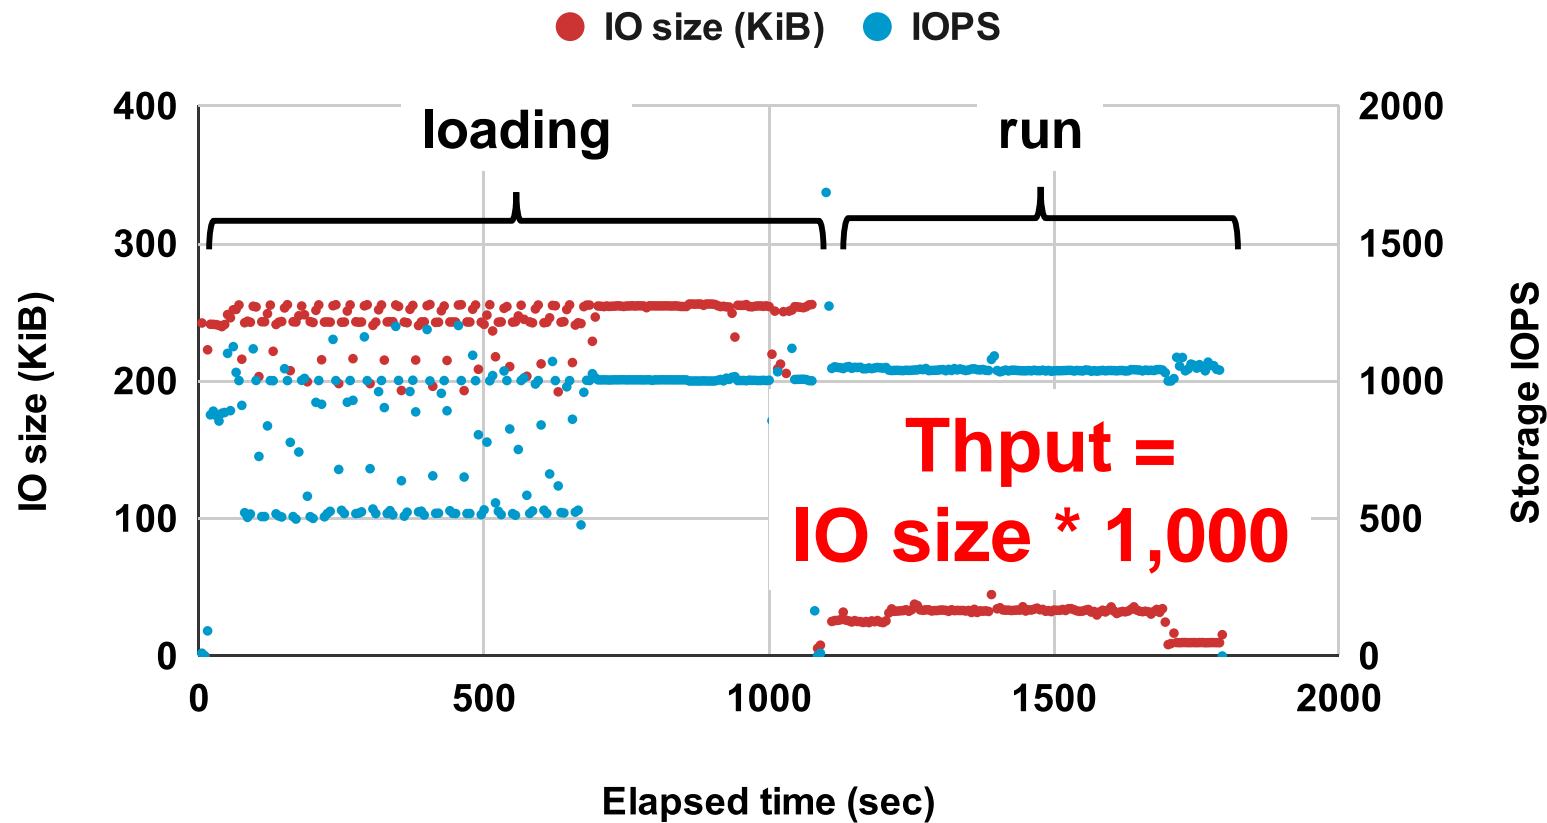


1,000 TPS

Illustrative Experiment (4)

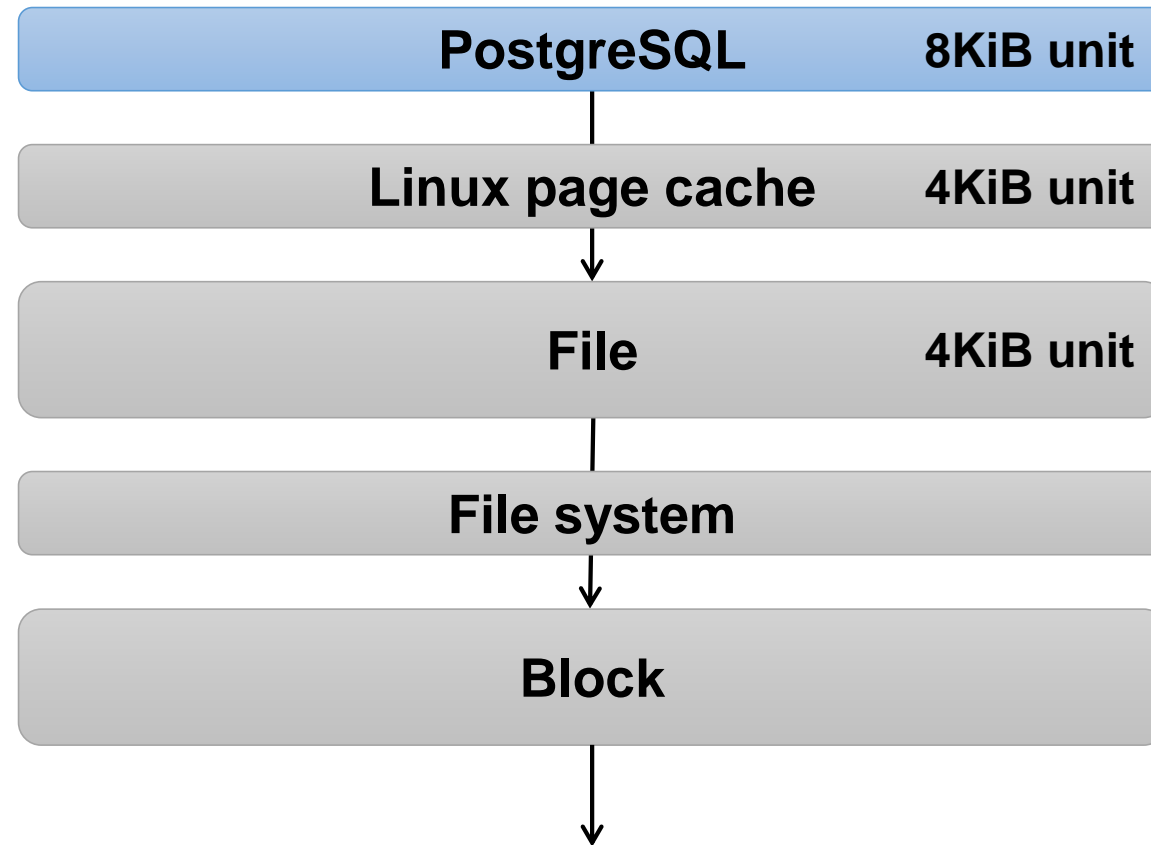
PostgreSQL cannot fully utilize cloud storage

pgbench, 5K scale, 100 clients, 300 GiB EBS SSD (1K IOPS)



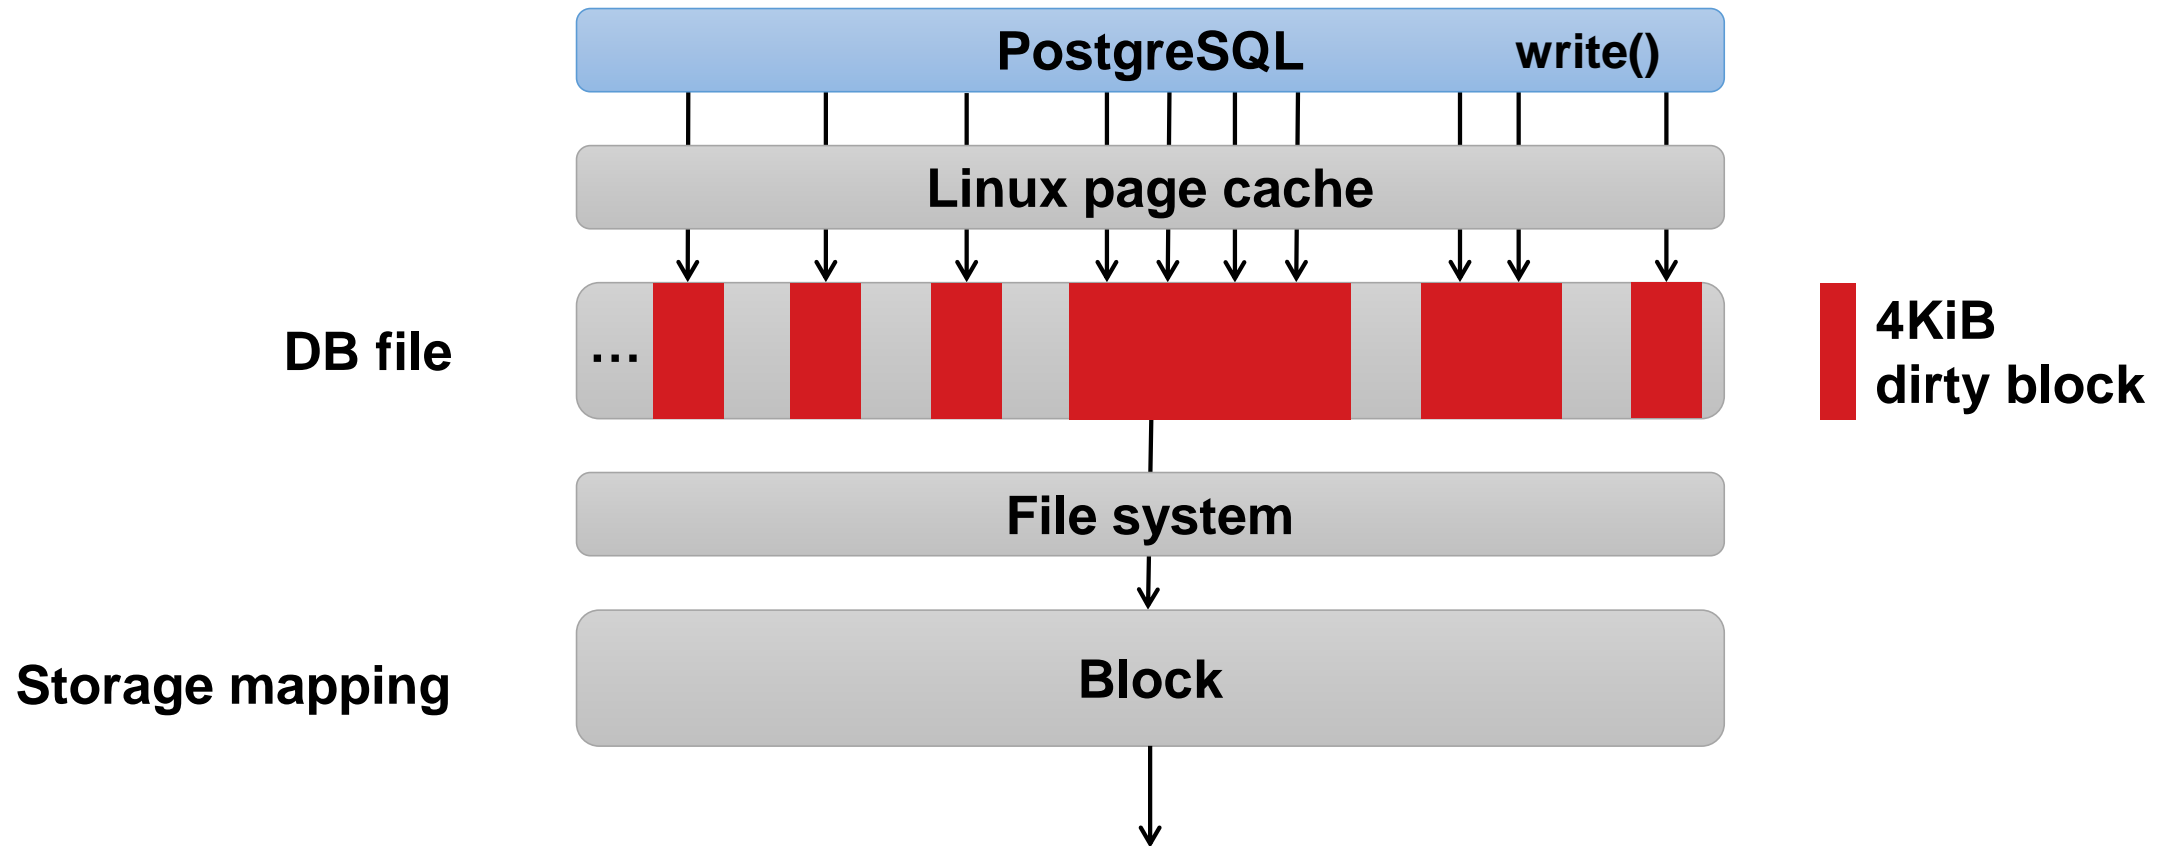
What's the Problem? (1)

Linux issues I/Os in 4 KiB unit



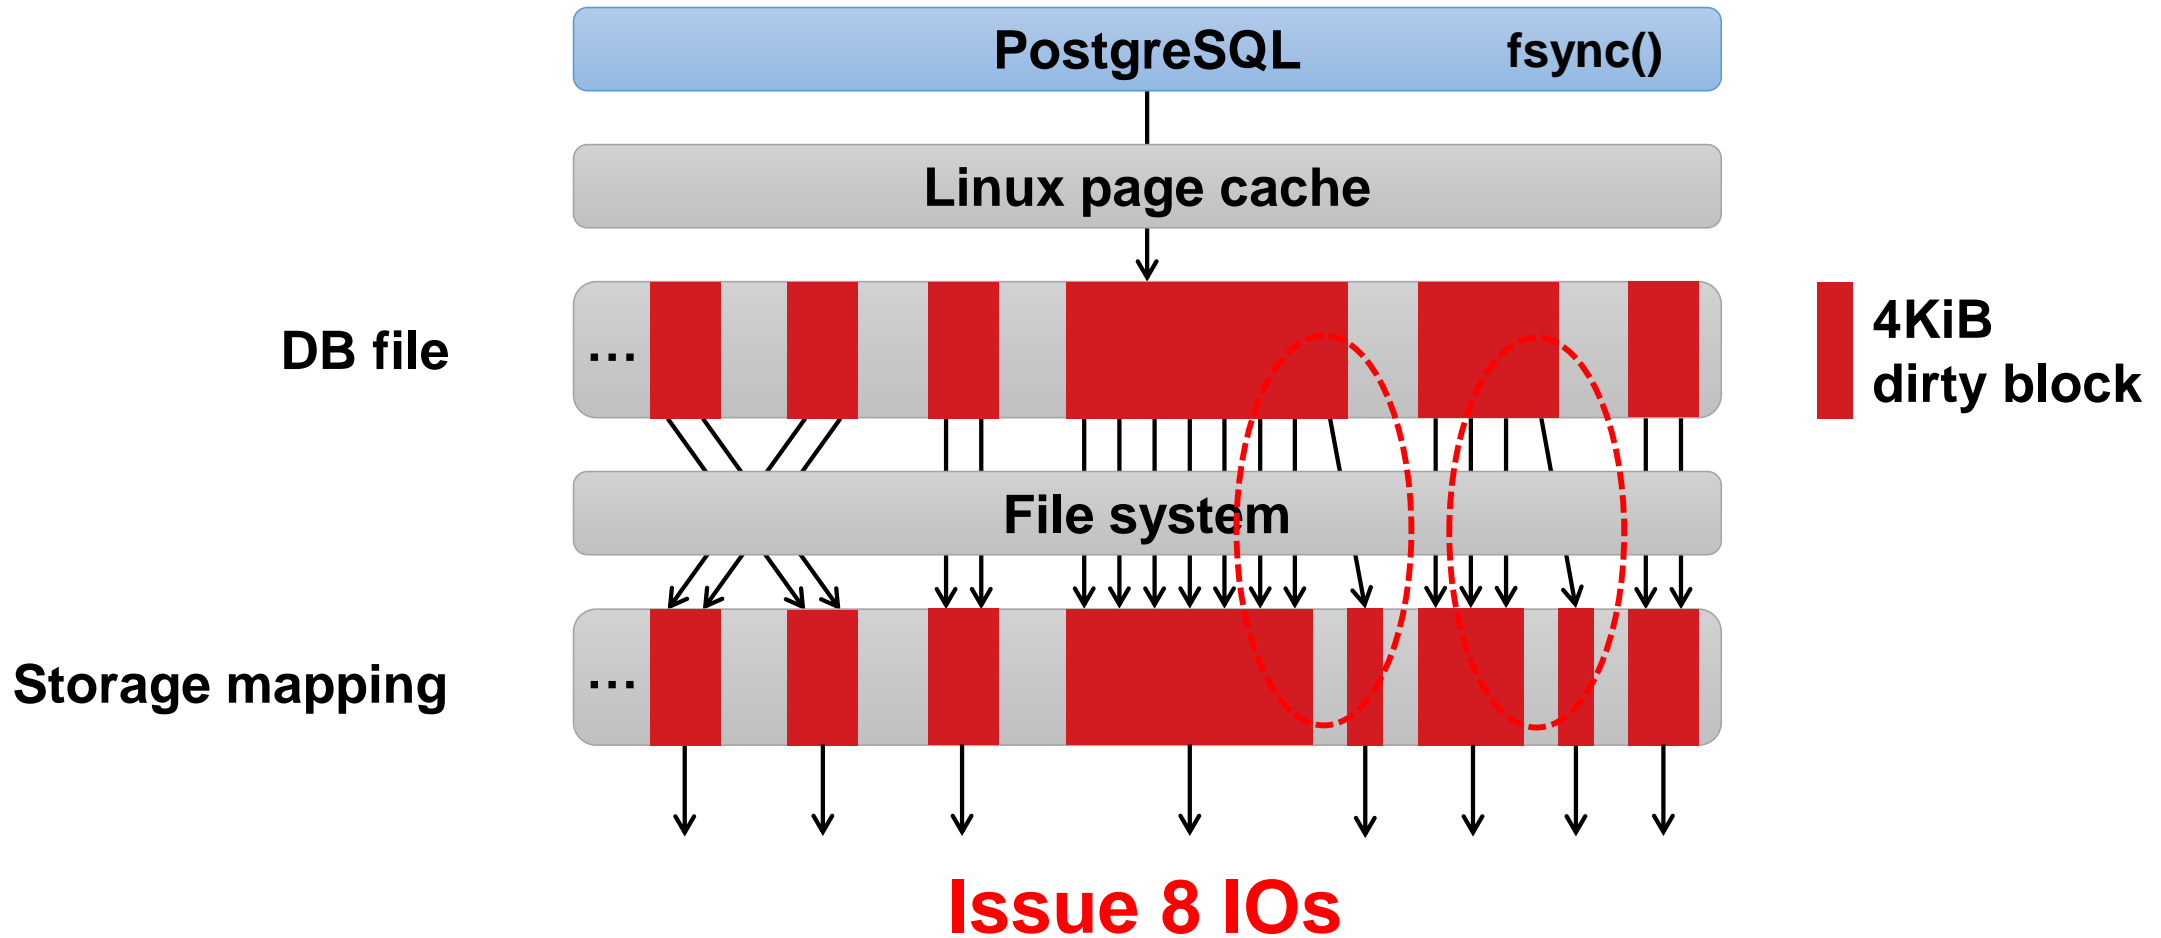
What's the Problem? (2)

Linux issues I/Os in 4 KiB unit



What's the Problem? (3)

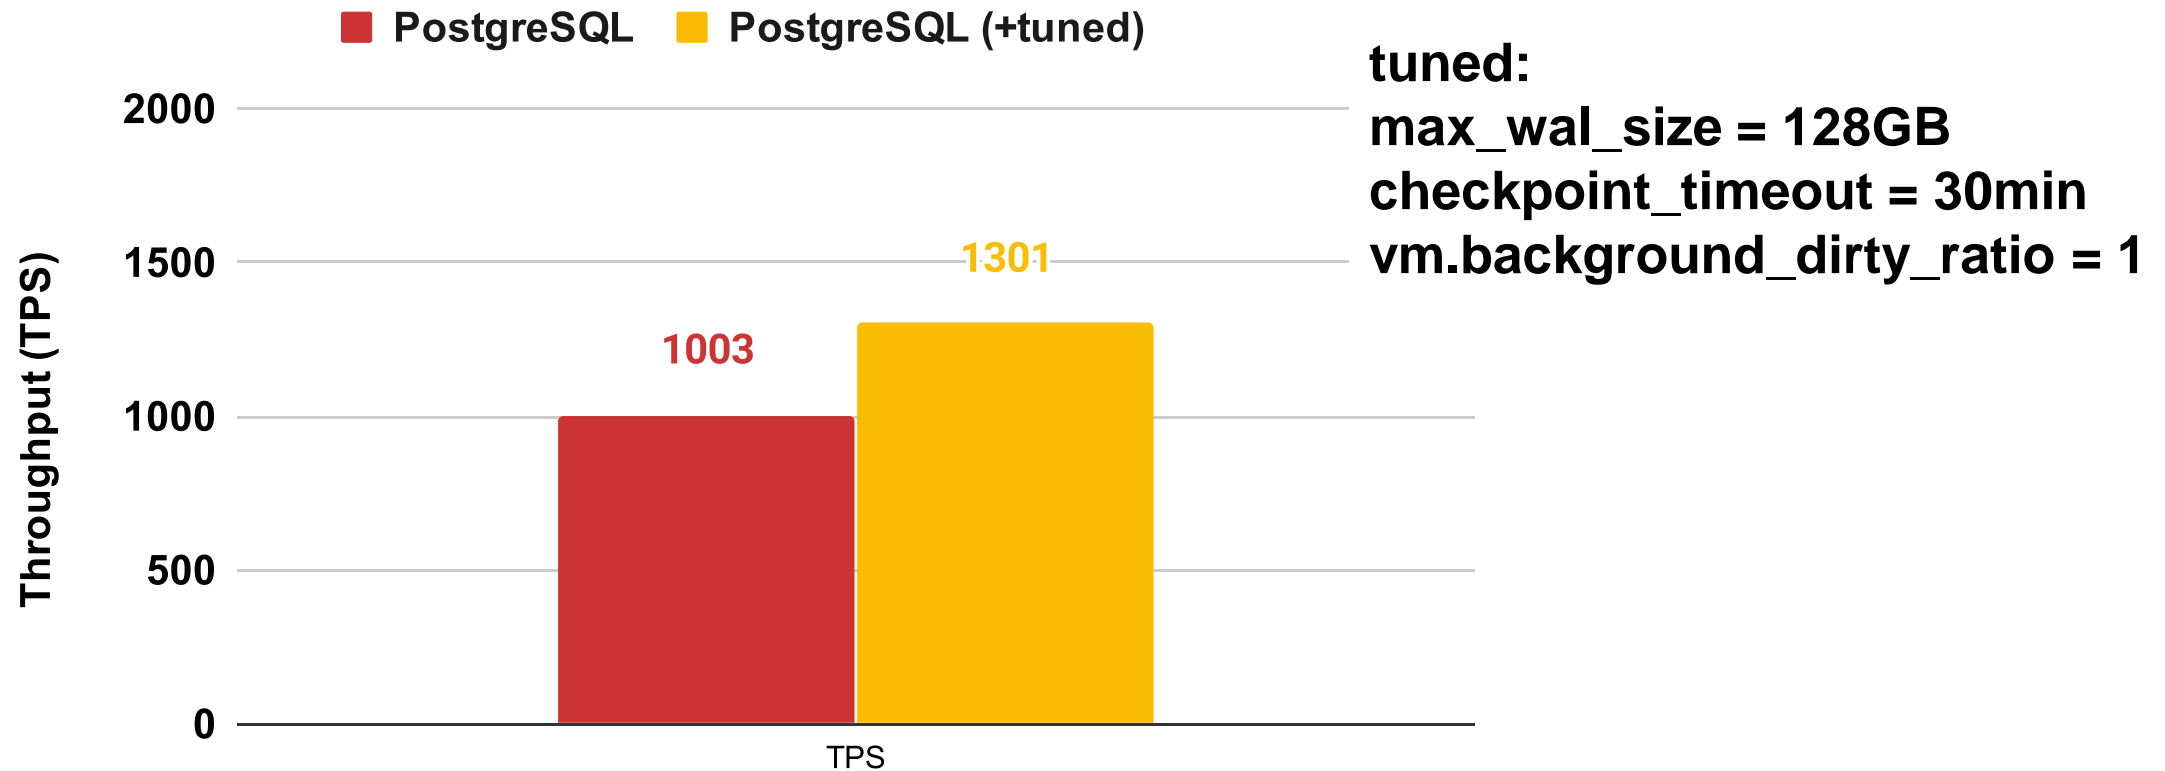
Linux issues I/Os in 4 KiB unit



Tuning PostgreSQL (1)

Tuning helps to get 30% better TPS

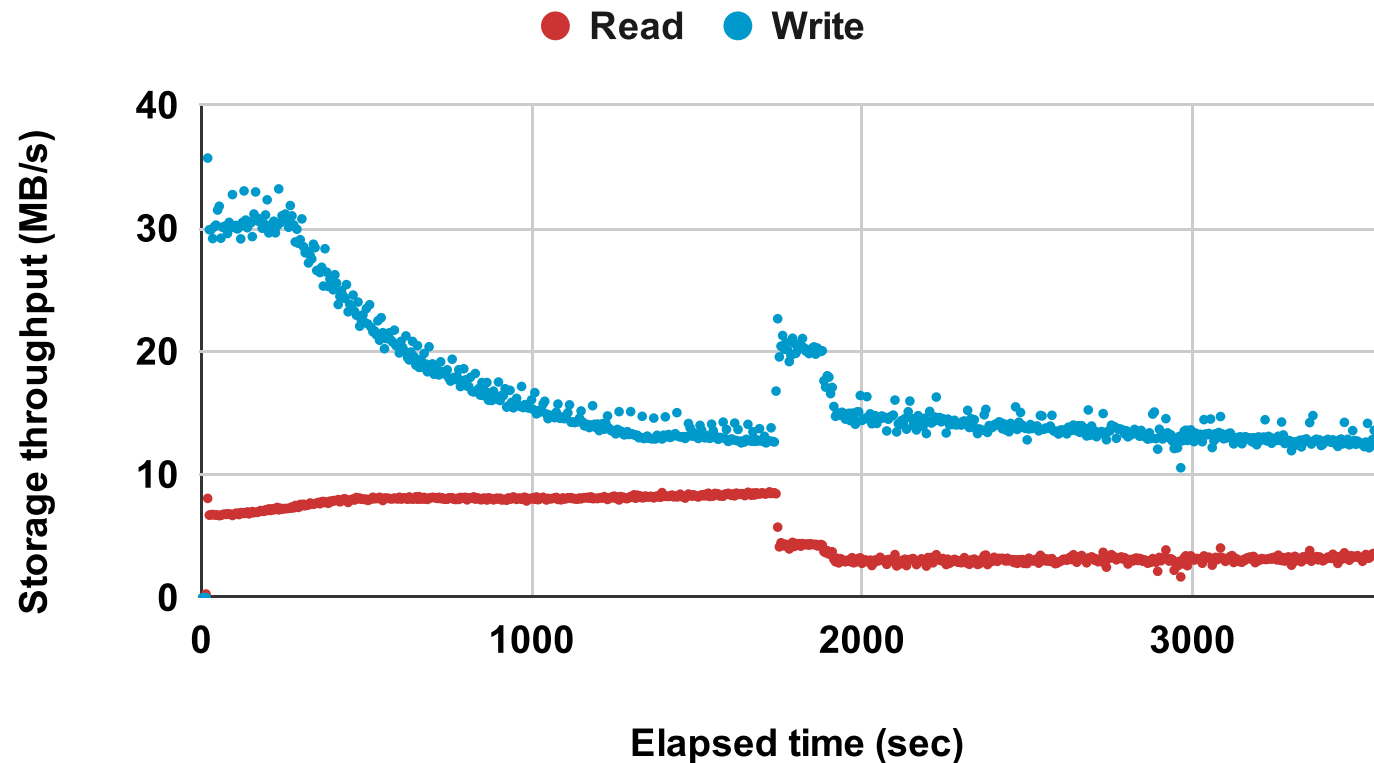
pgbench tpcb-like throughput, 5K scale, 100 clients



Tuning PostgreSQL (2)

Be aware of cache space contention inside OS

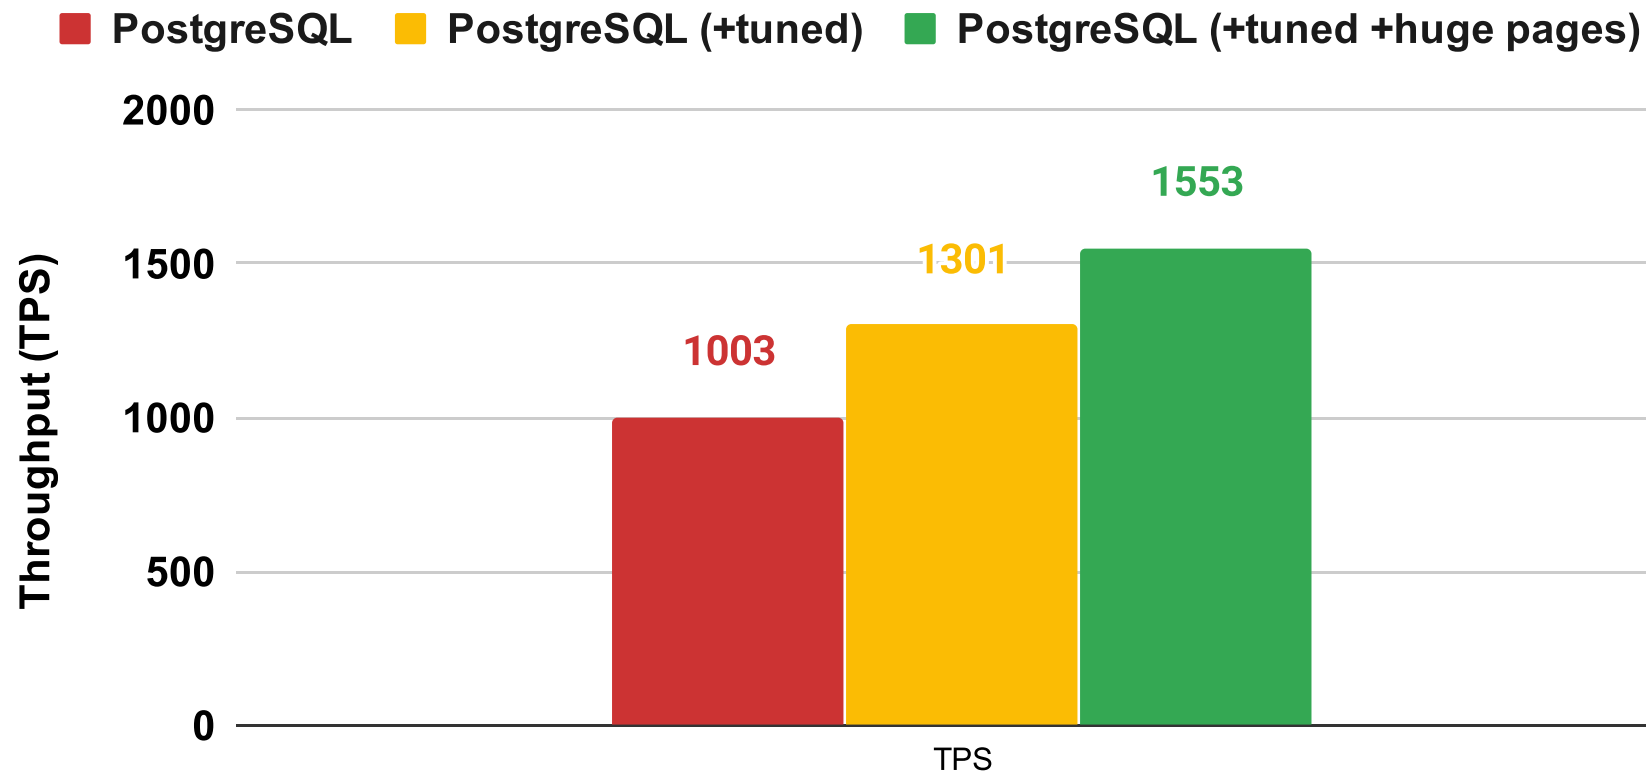
PostgreSQL (+tuned) storage throughput during pgbench run



Tuning PostgreSQL (3)

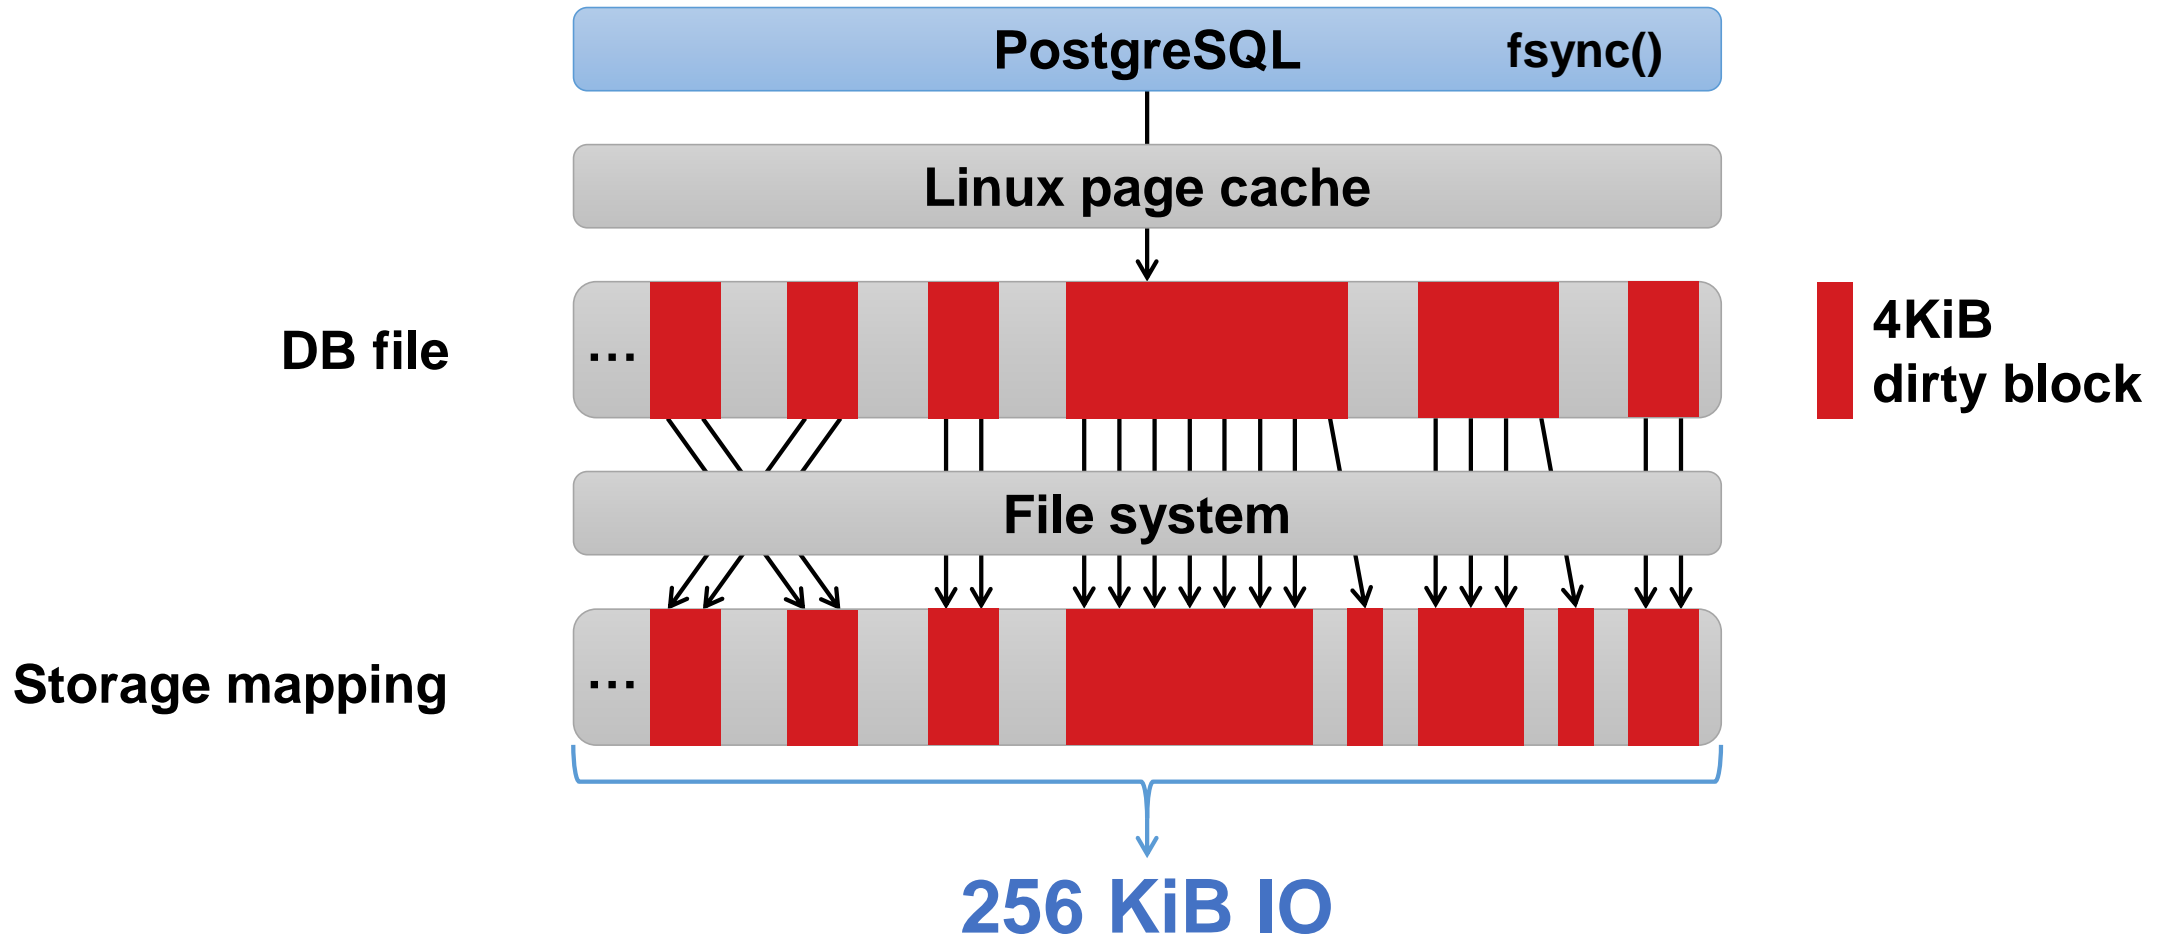
Using huge pages additionally improves 20%

pgbench tpcb-like throughput, 5K scale, 100 clients



Possible Solution for AWS EBS SSD

What if we can issue **1 IO** instead of 8?



Cloud Storage Diversity

Each cloud has different policy for its storage



	IOPS	Thput	Note
ebs-gp2	3/GiB	250	Burstable
ebs-io1	\$67/1K/M	500	
ebs-st1	500	40/TiB	Burstable
ebs-sc1	250	12/TiB	Burstable



	IOPS	Thput	Note
ultra	\$49/1K/M	\$1/MB/M	
pre-ssd	20,000	900	Burstable
std-ssd	6,000	750	Pay # IOs
std-hdd	2,000	500	Pay # IOs



Google Cloud

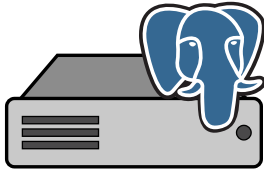
	IOPS	Thput	Note
pd-ssd	30/GiB	0.48/GiB	
pd-ssd-16kb	30/GiB	0.48/GiB	16 KiB
pd-std	1.5/GiB	0.12/GiB	
pd-std-16kb	1.5/GiB	0.12/GiB	16 KiB



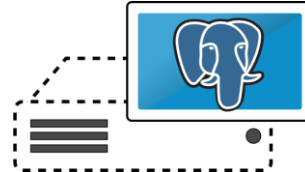
AppOS Extension

AppOS is a file engine built for diverse types of storage

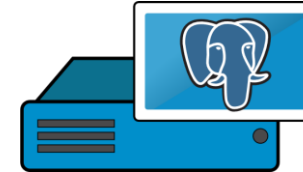
Bare metal



Virtual machine



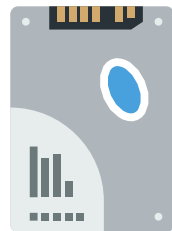
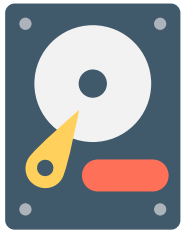
Container



Container



AppOS extension



aws



Google Cloud



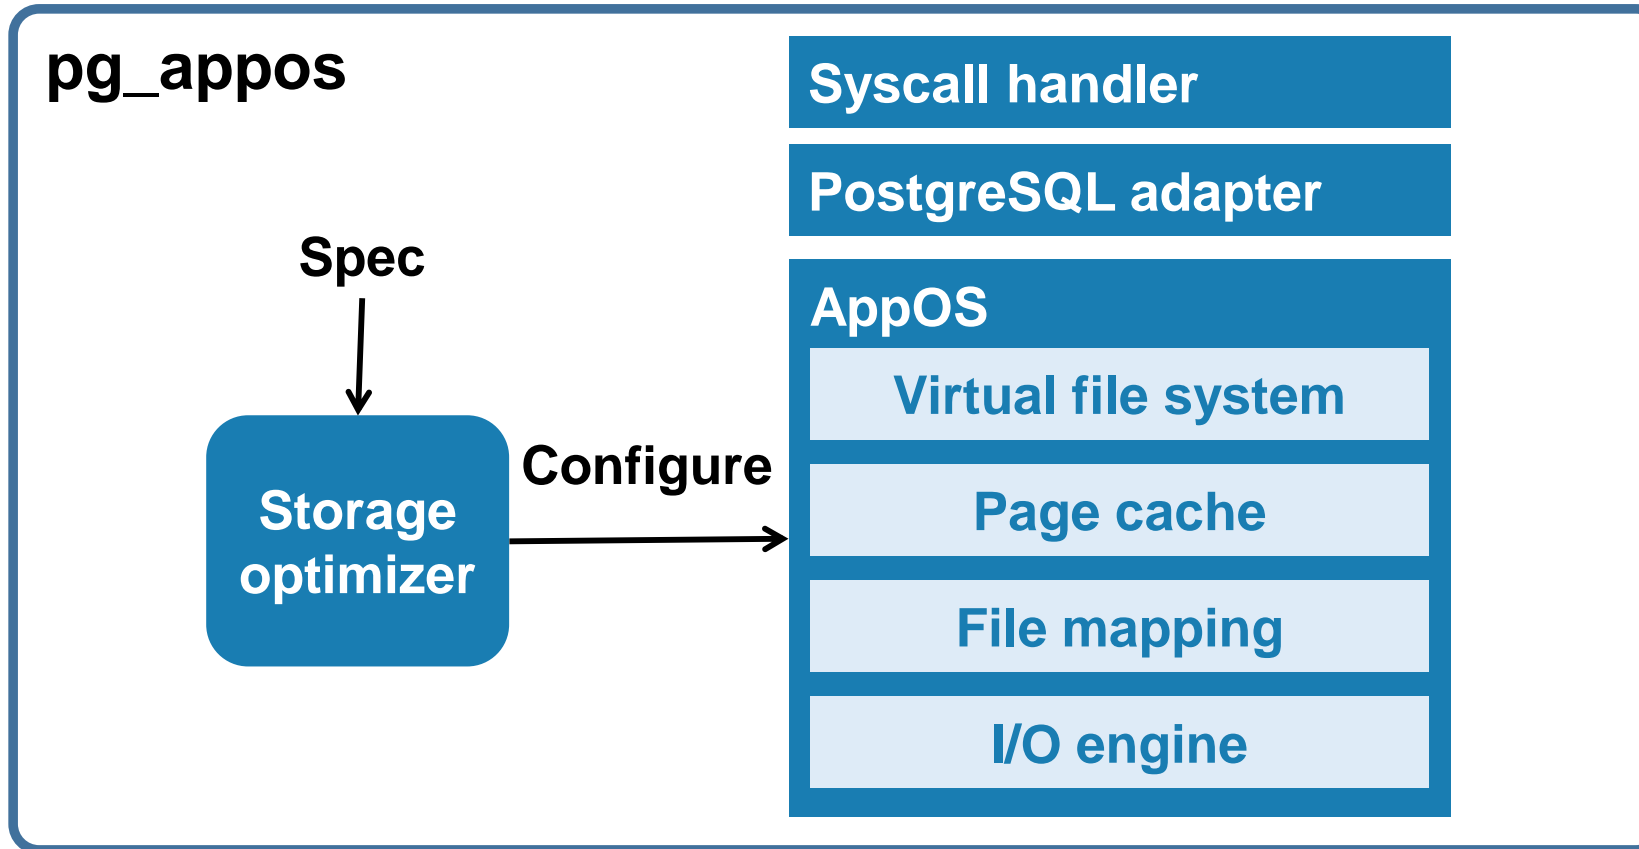
Microsoft
Azure

AppOS Design Goals

- No PostgreSQL code changes**
- PostgreSQL-aware I/O processing**
- Storage-aware I/O processing**

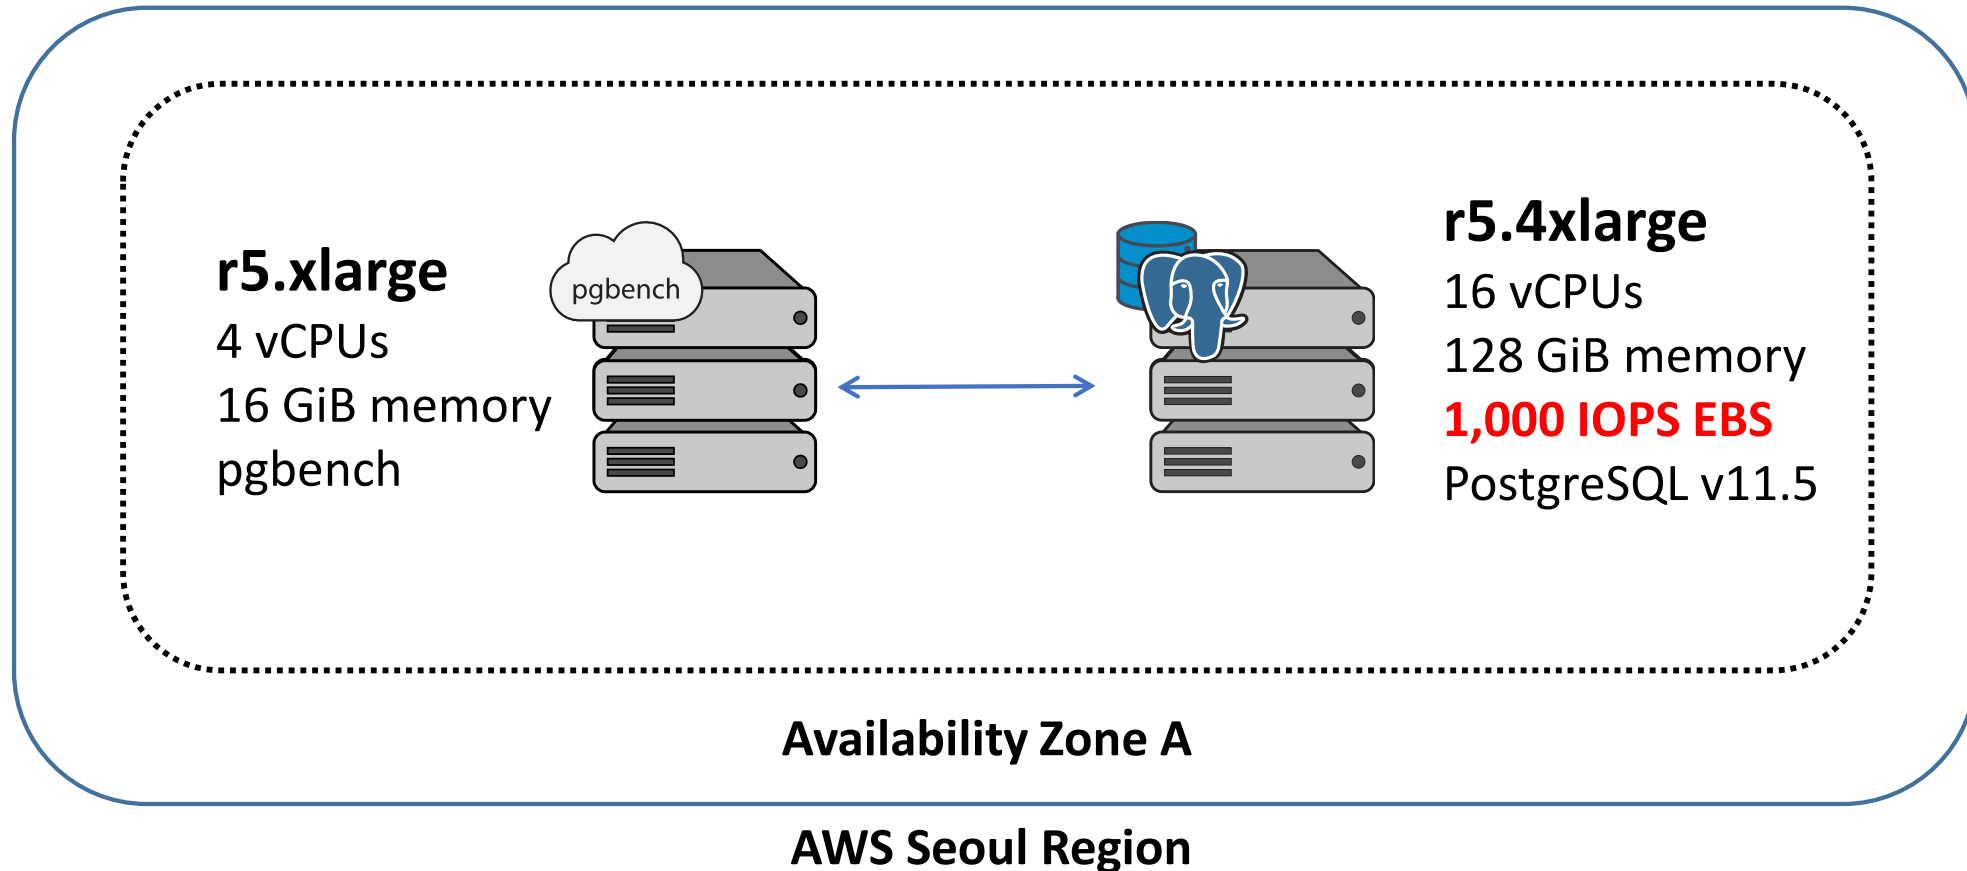
AppOS Architecture

AppOS extension is a collection of libraries



Performance on AWS Storage (1)

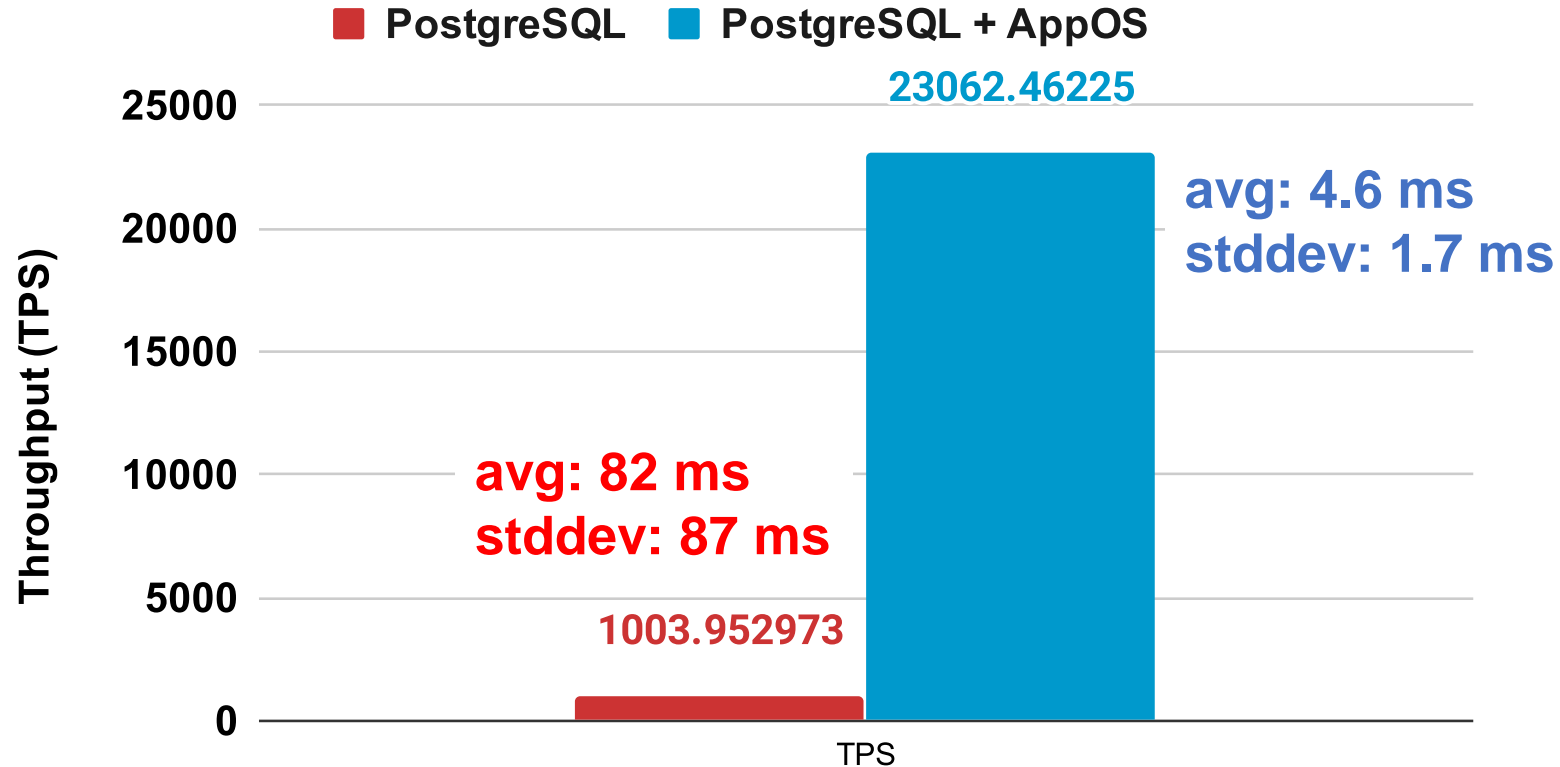
Test setup on AWS



Performance on AWS Storage (2)

PostgreSQL processes 23x more transactions with AppOS

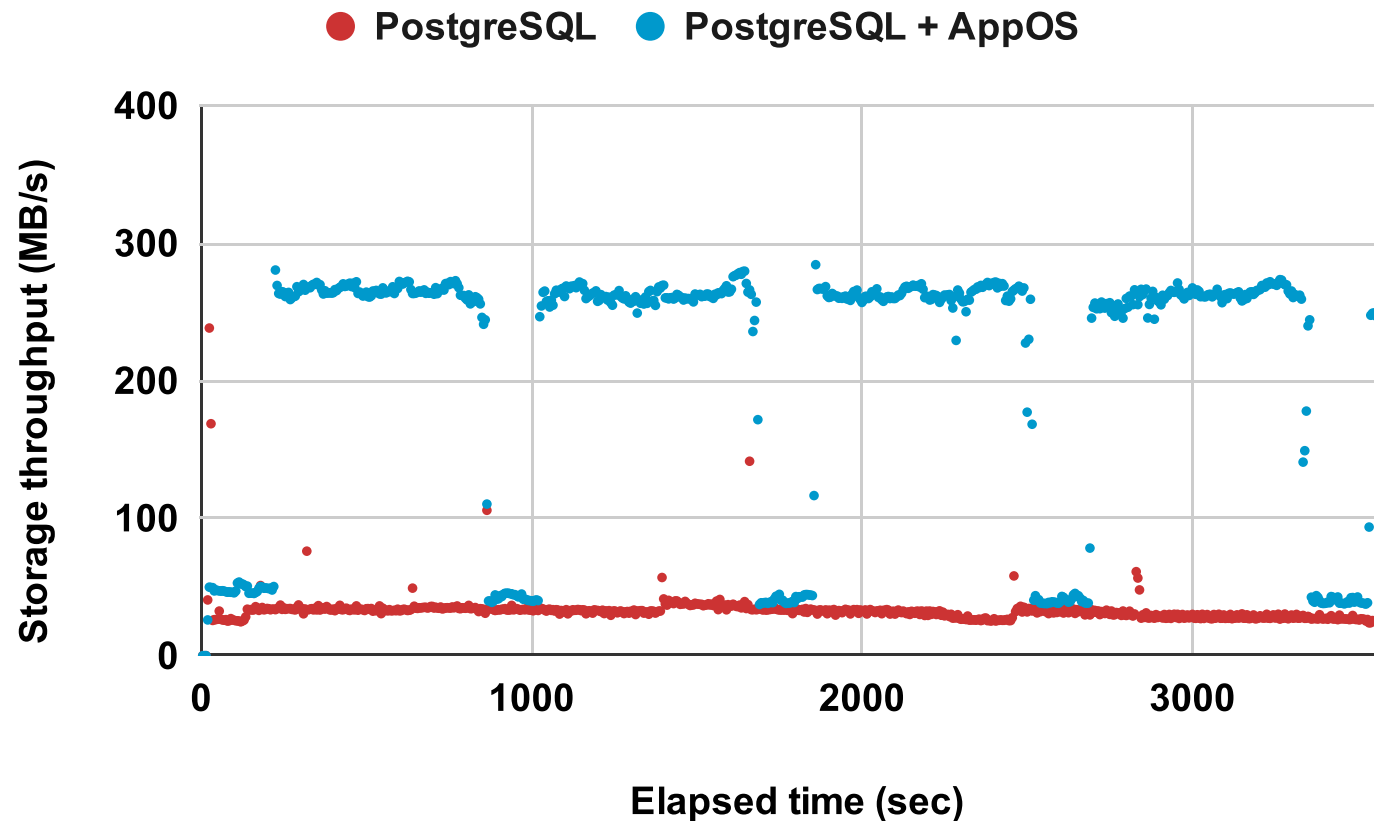
pgbench, 5K scale, 100 clients, 300 GiB EBS (1K IOPS)



Performance on AWS Storage (3)

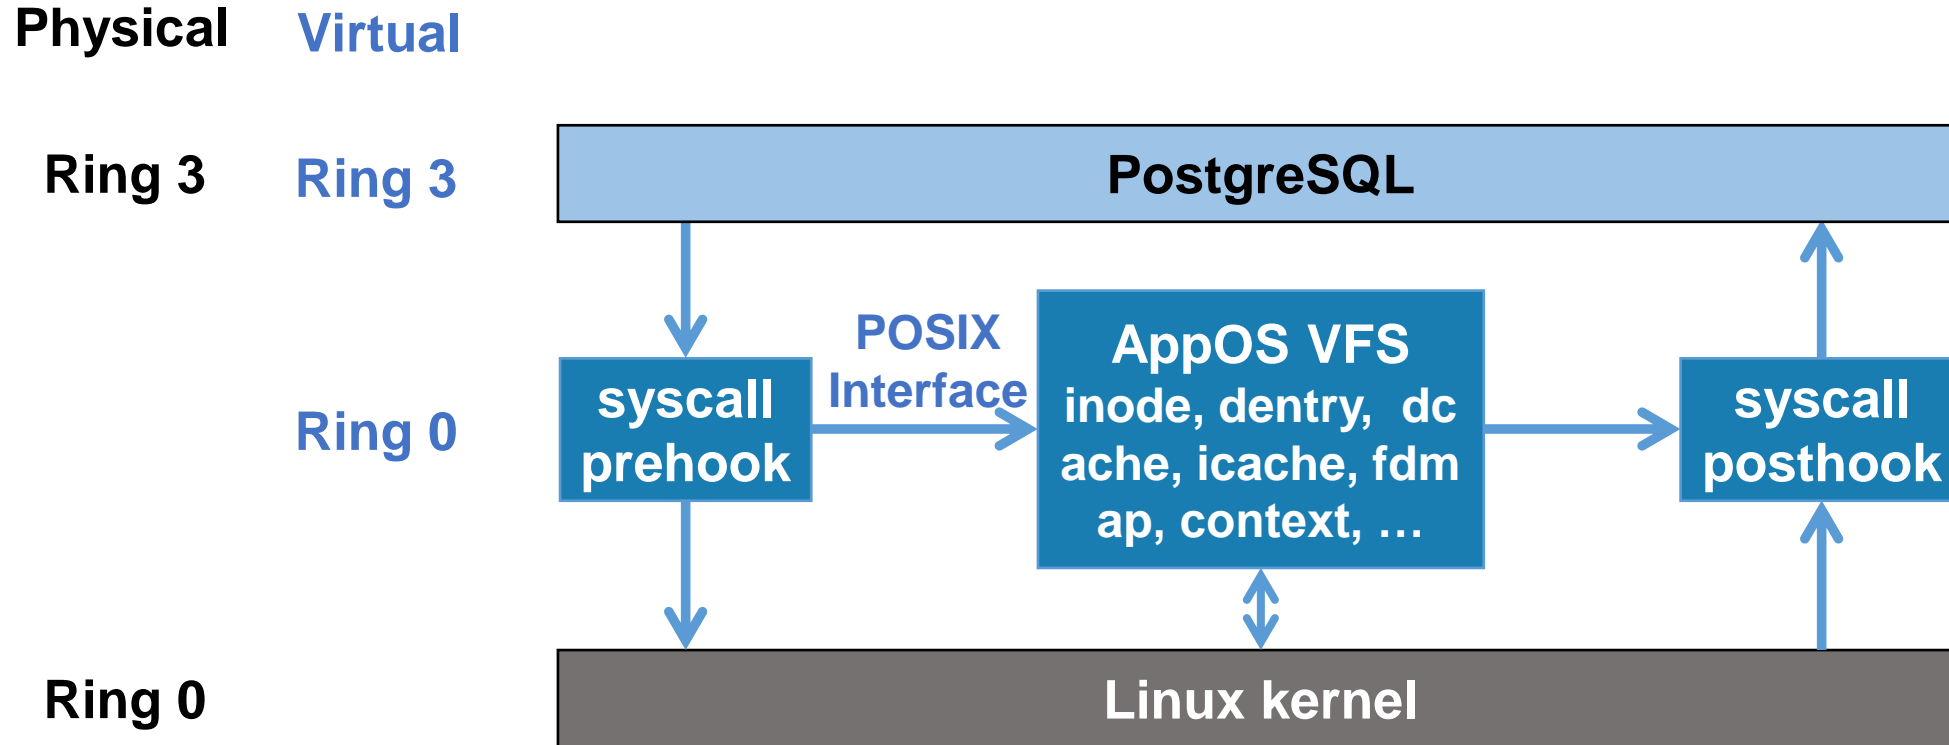
AppOS makes PostgreSQL to fully utilize storage bandwidth

Throughput of 1K IOPS EBS during pgbench run



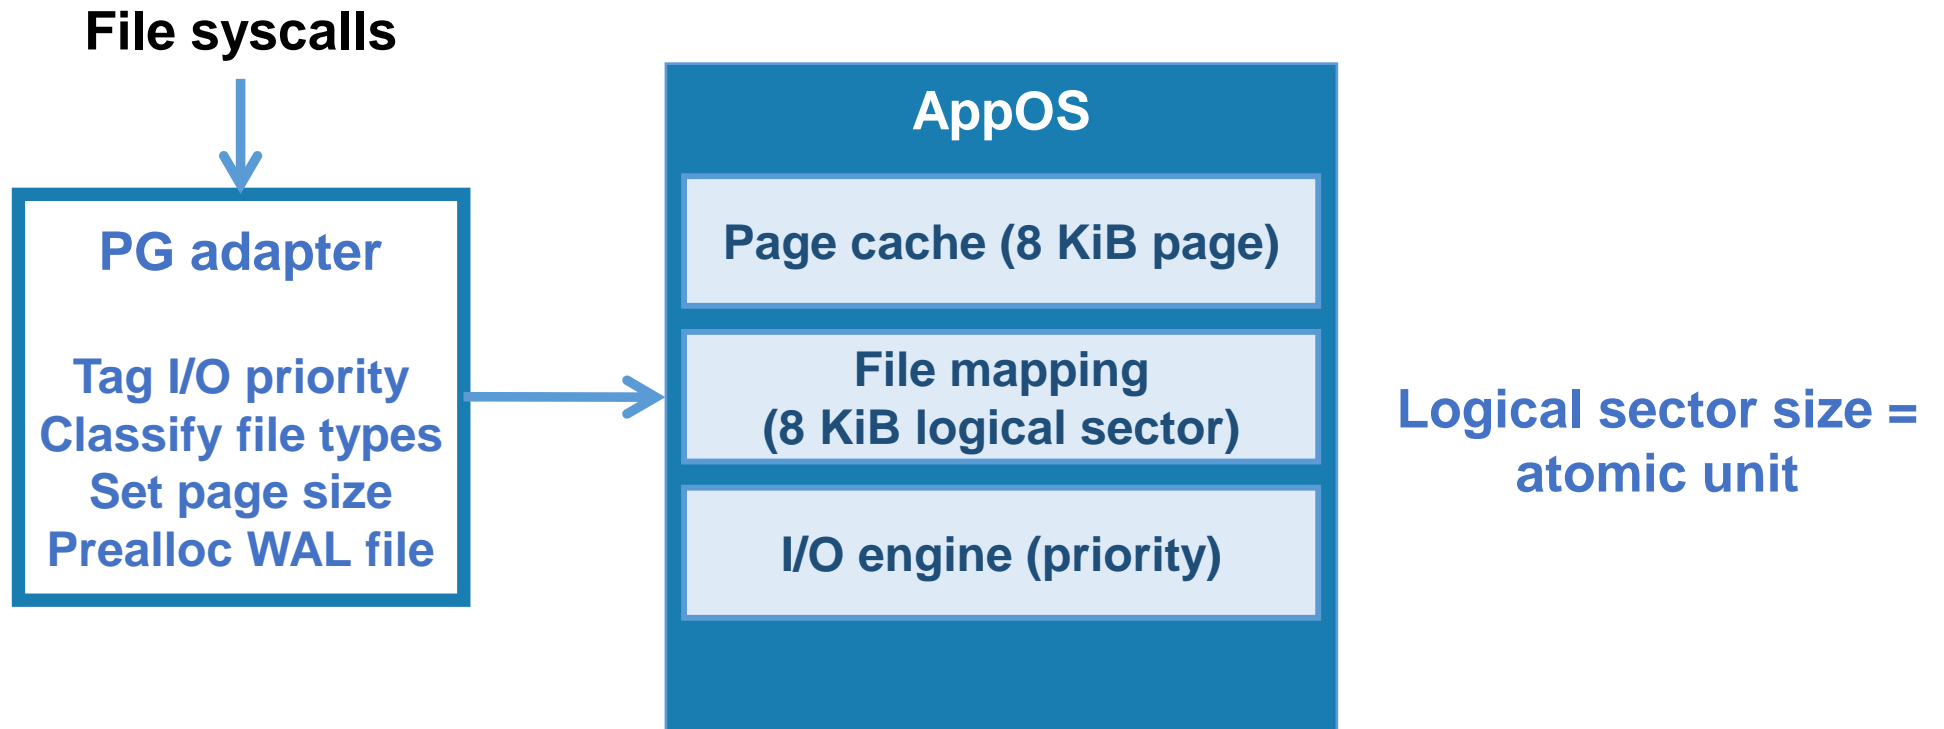
AppOS Internals (1)

AppOS does not require code changes



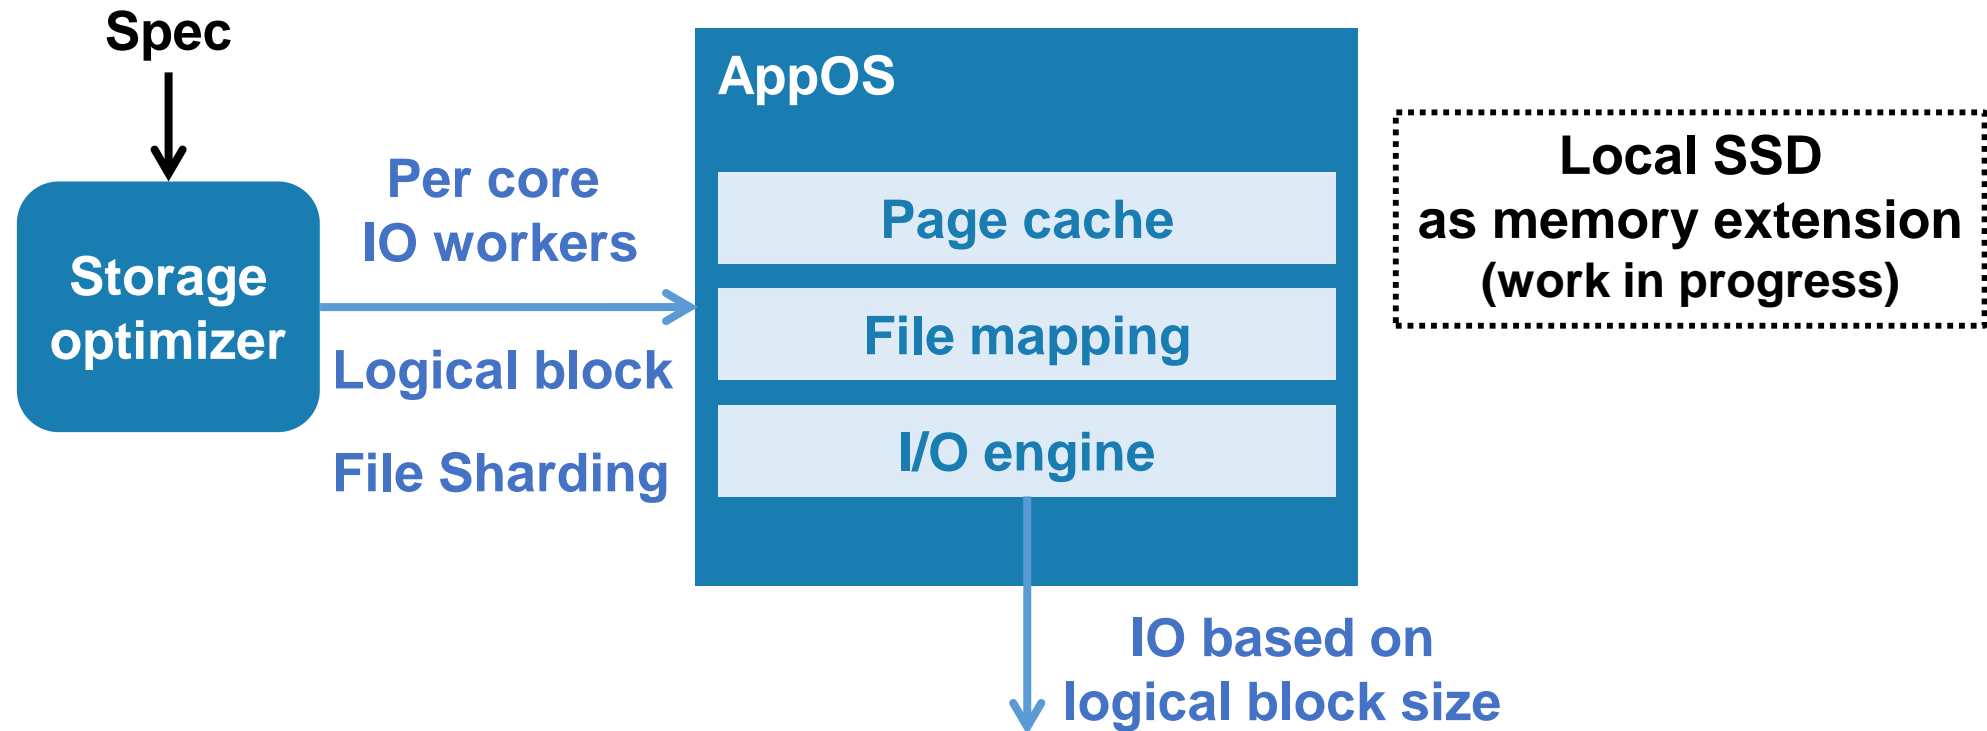
AppOS Internals (2)

AppOS provides PostgreSQL-aware I/O processing



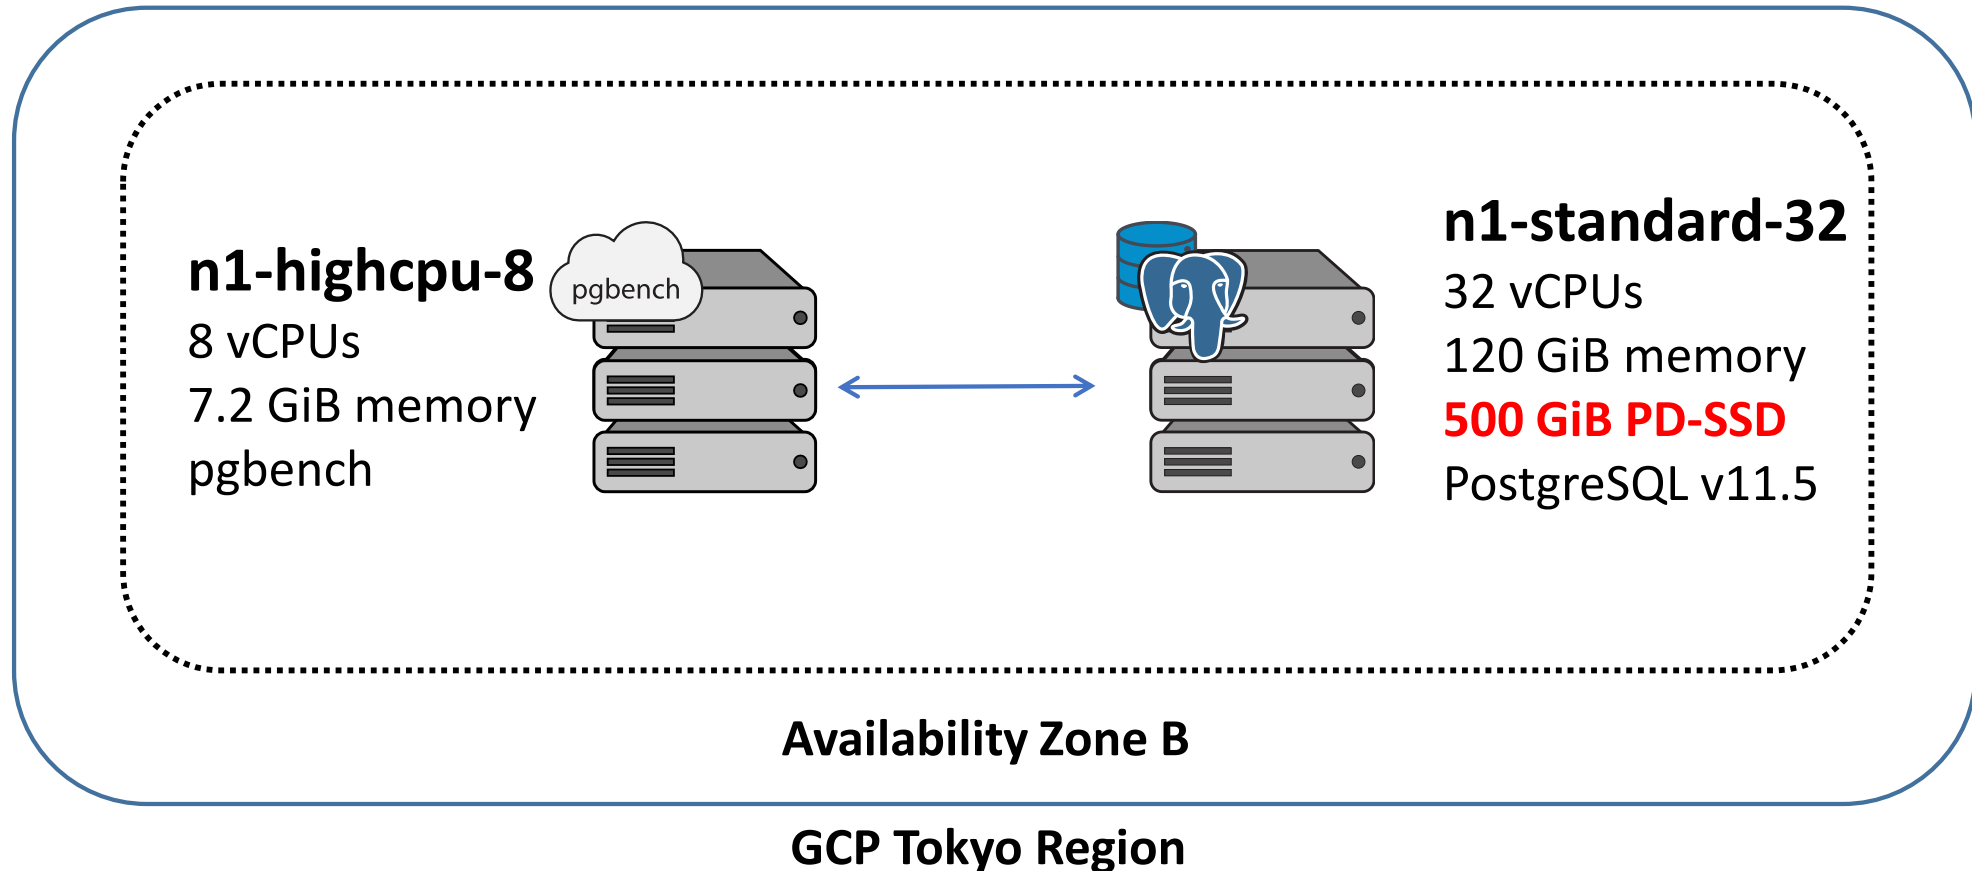
AppOS Internals (3)

AppOS provides storage-aware I/O processing



Performance on GCP Storage (1)

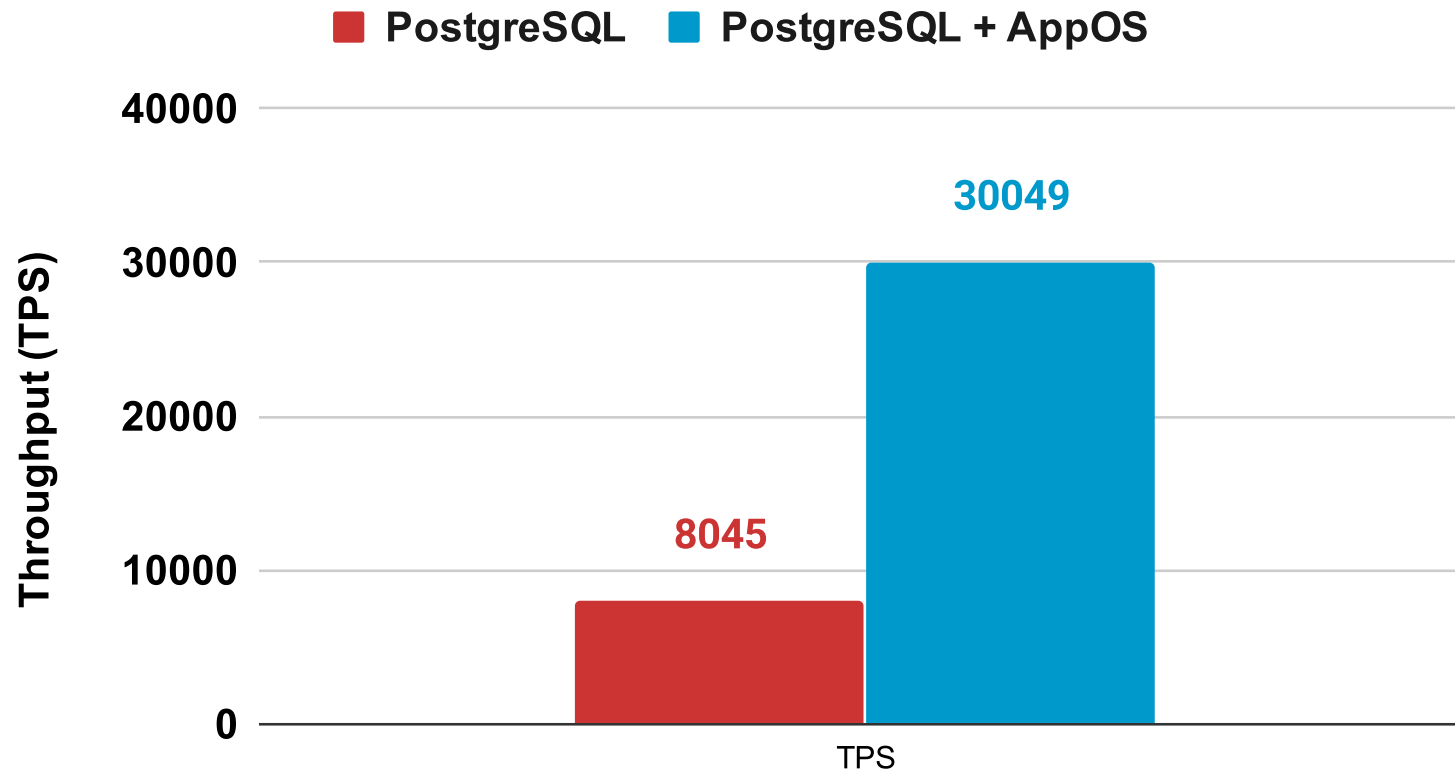
Test setup on GCP



Performance on GCP Storage (2)

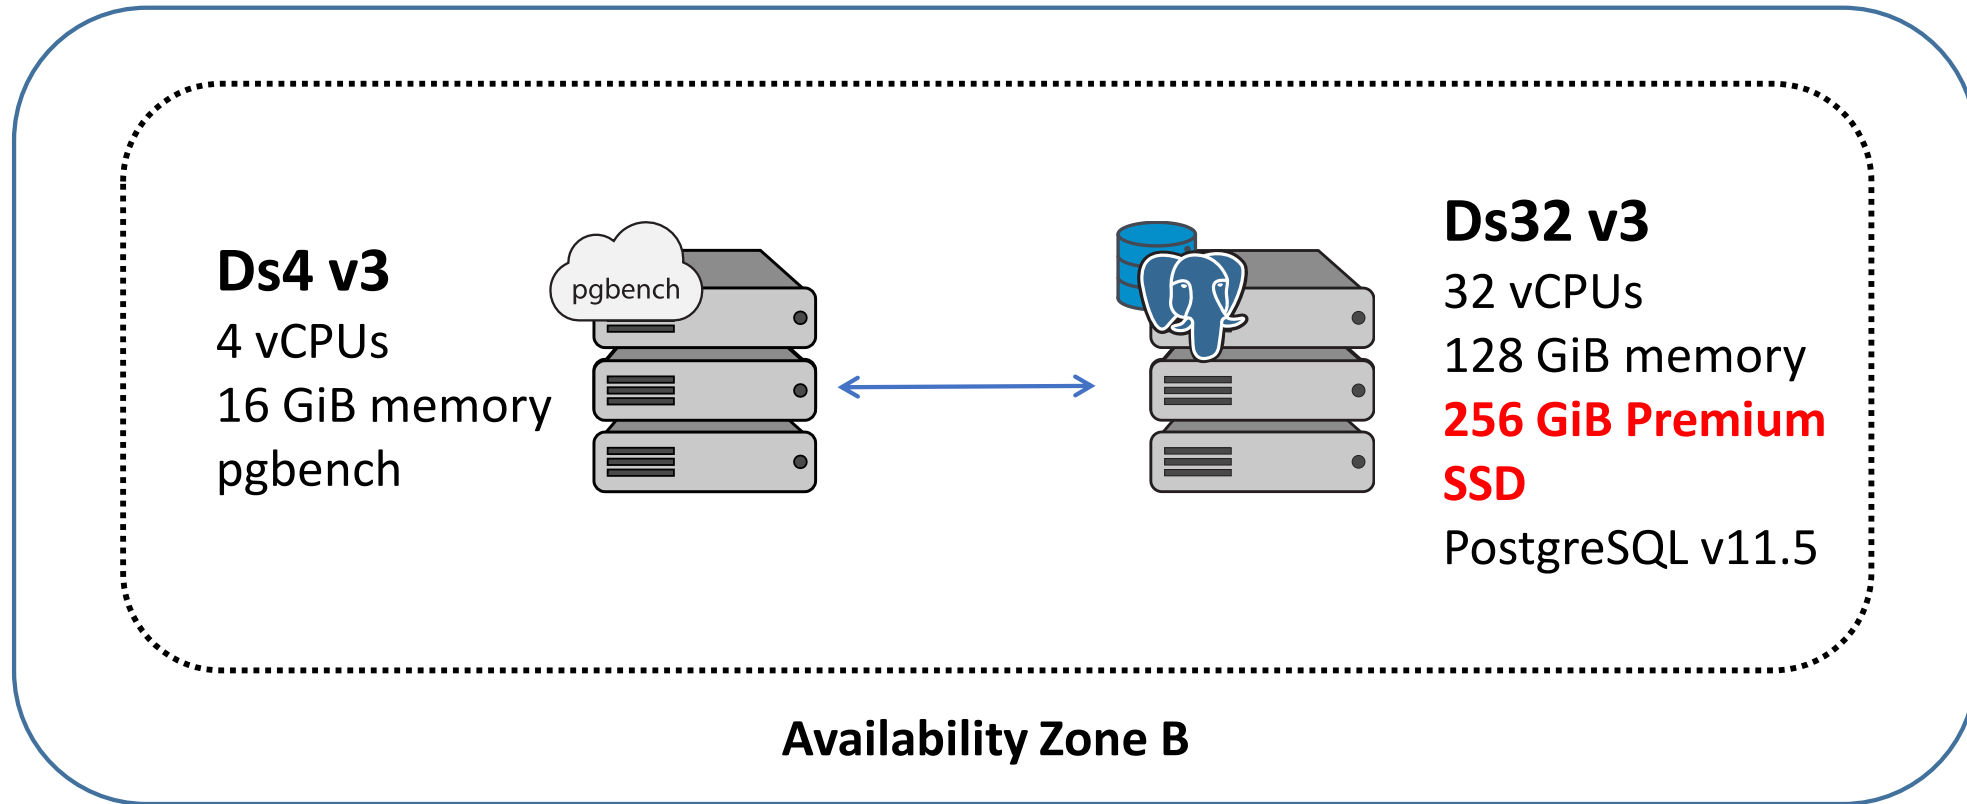
PostgreSQL processes 3.7 more transactions with AppOS

pgbench, 5K scale, 100 clients, 500 GiB PD SSD



Performance on Azure Storage (1)

Test setup on Azure



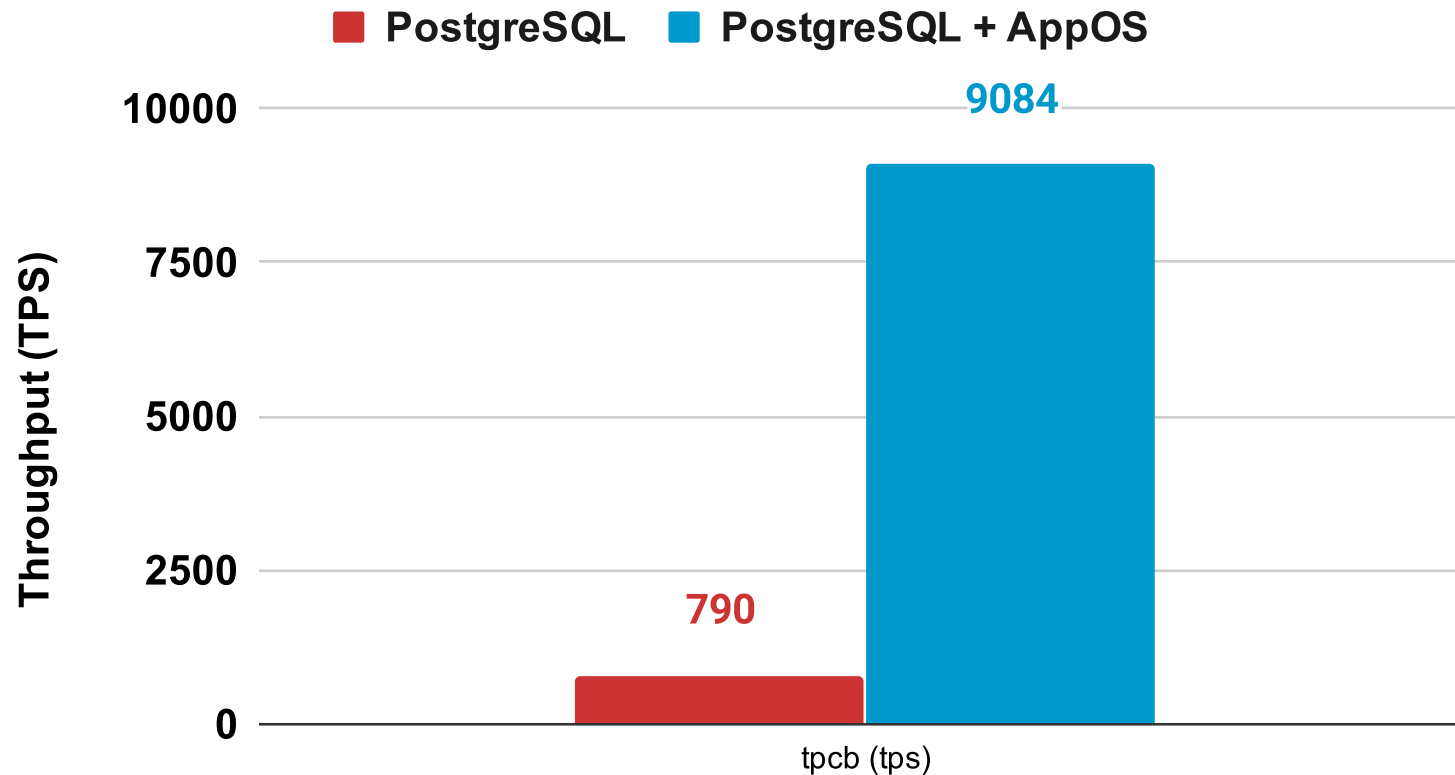
Availability Zone B

Azure Seoul Region

Performance on Azure Storage (2)

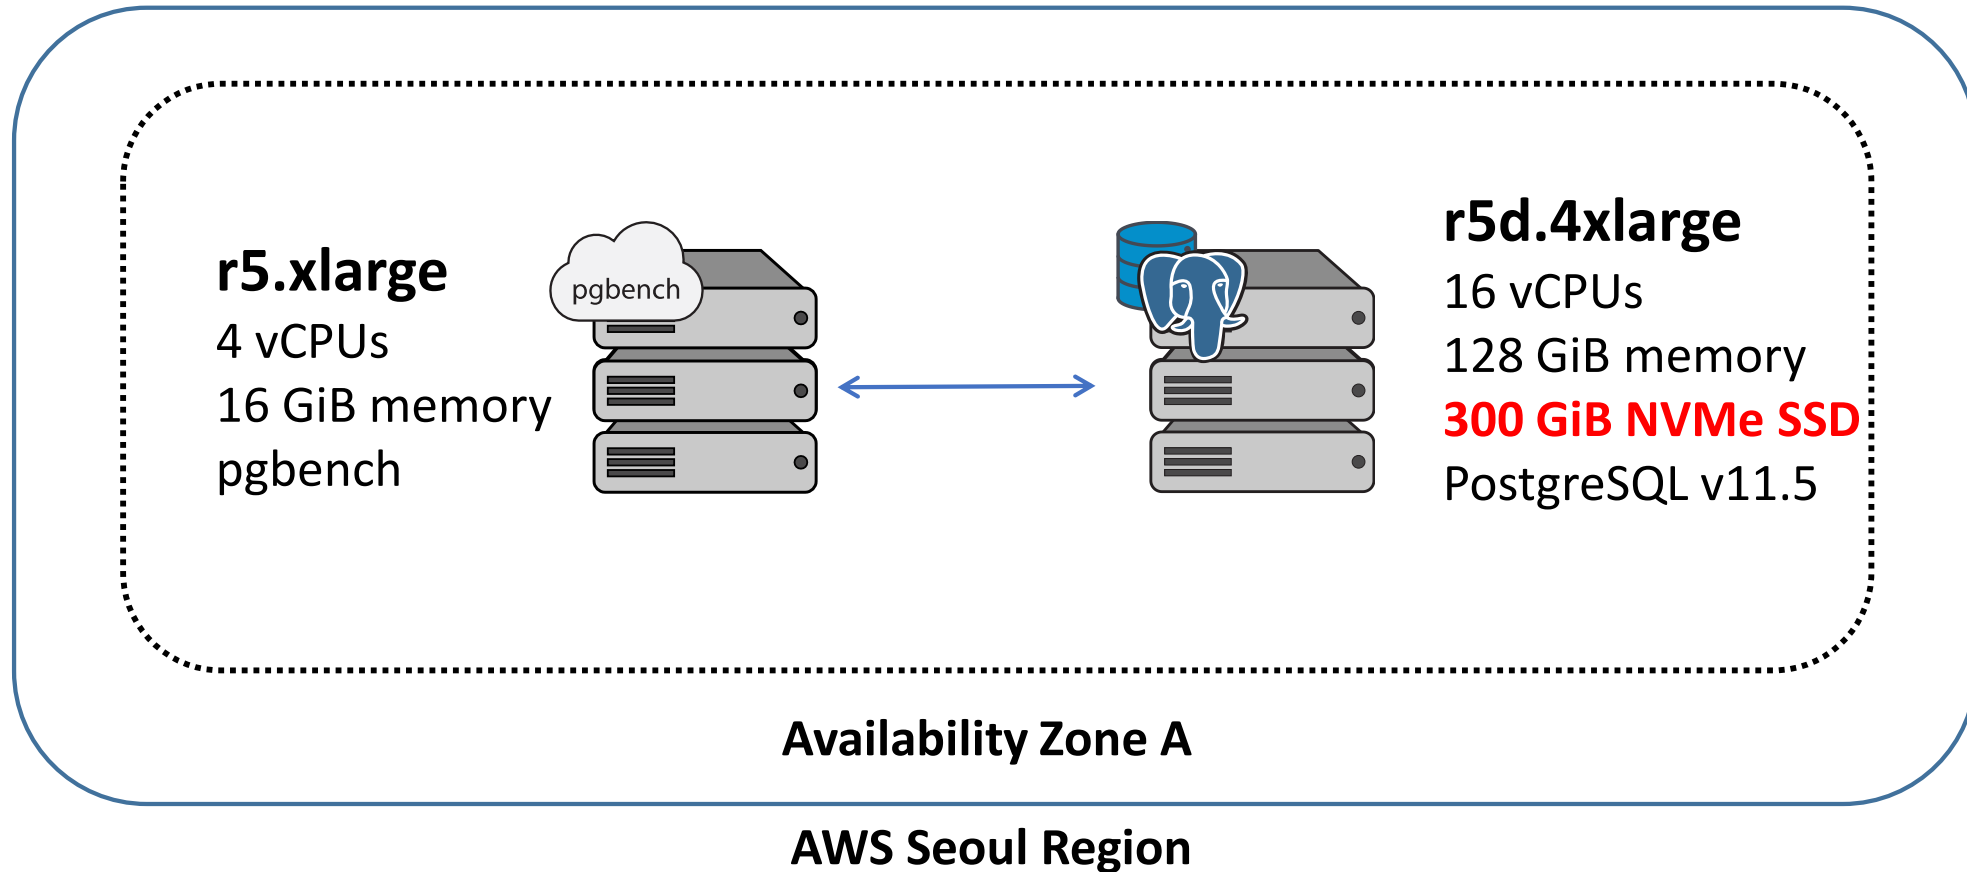
PostgreSQL processes 11.5x more transactions with AppOS

pgbench, 5K scale, 100 clients, 256 GiB Premium SSD



Performance on SSD (1)

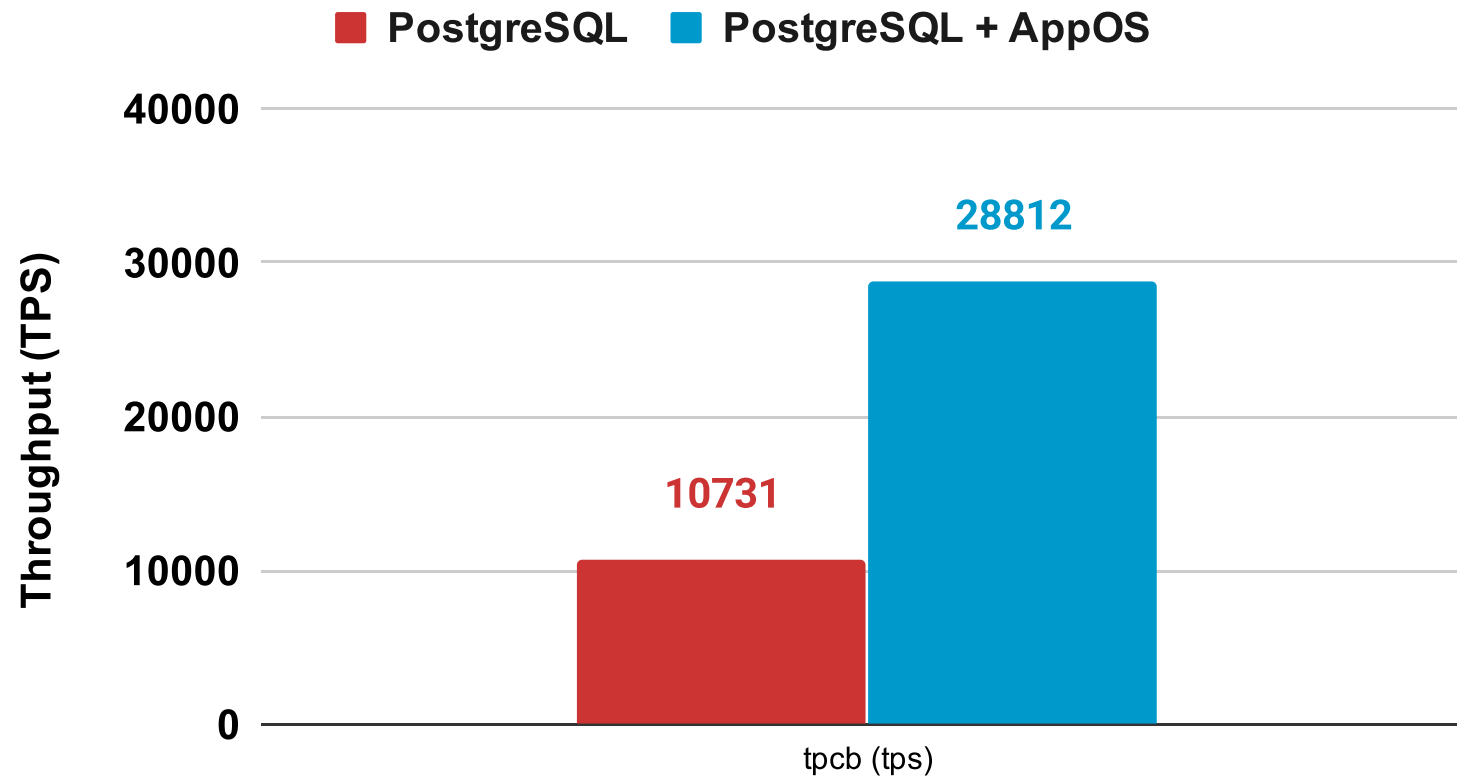
Test setup on AWS



Performance on SSD (2)

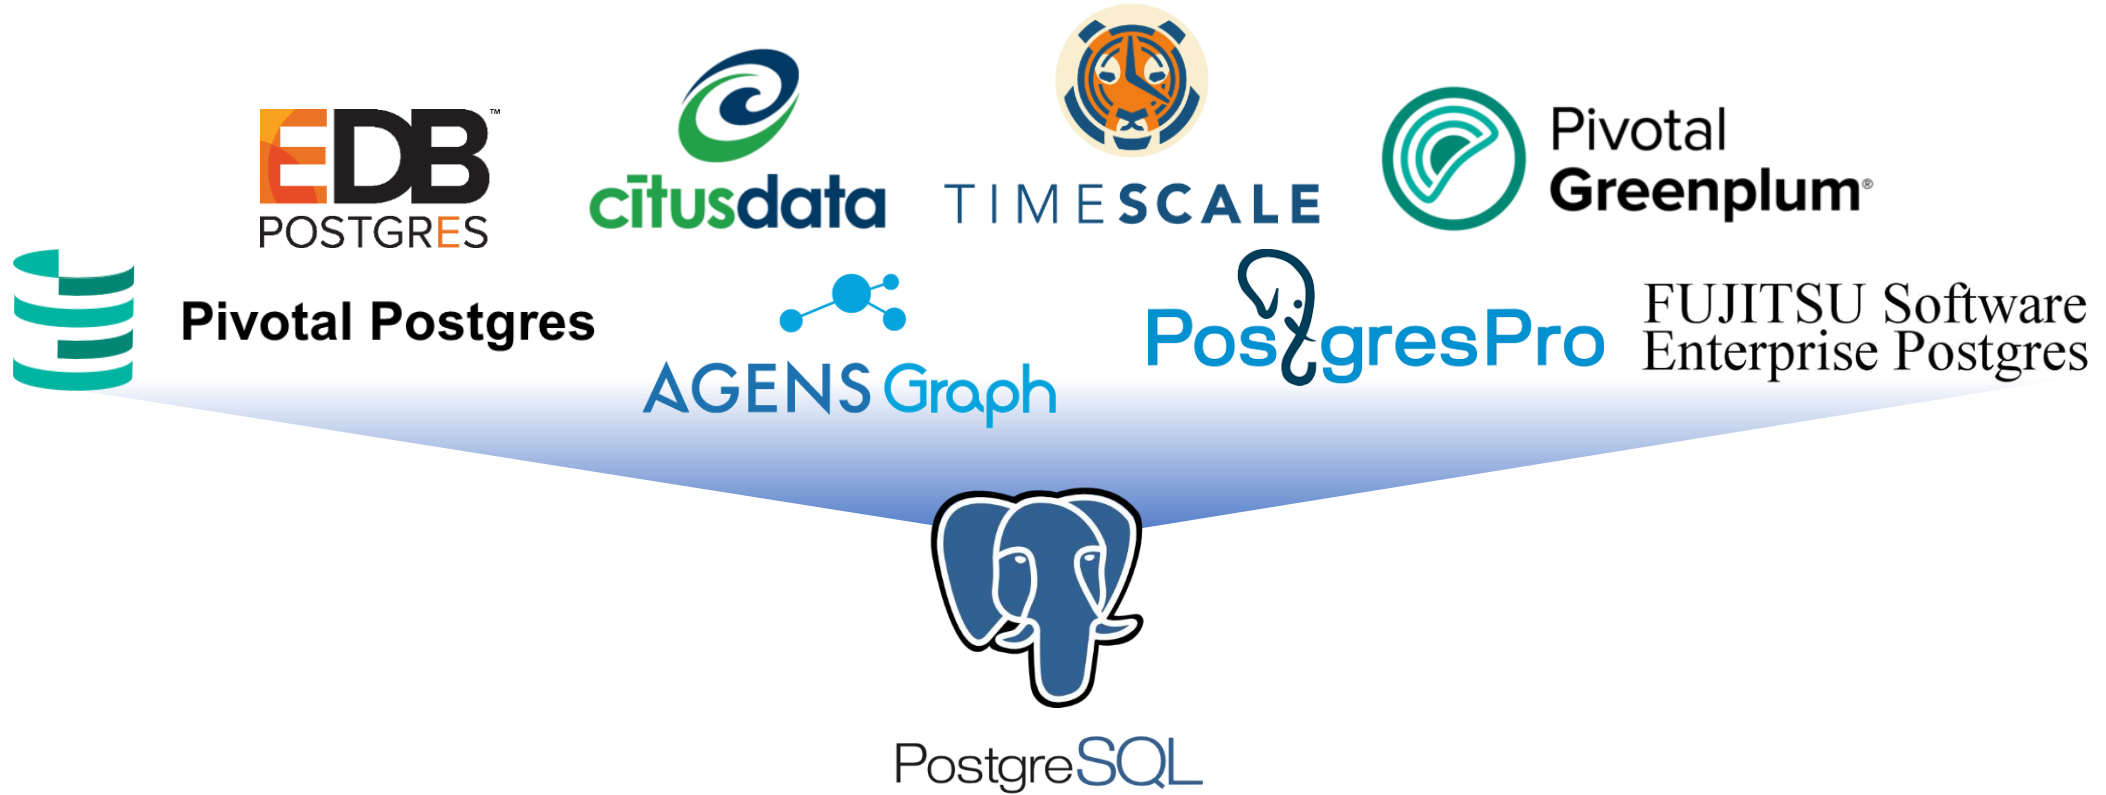
PostgreSQL processes 2.7x more transactions with AppOS

pgbench, 5K scale, 100 clients, 300 GiB NVMe SSD



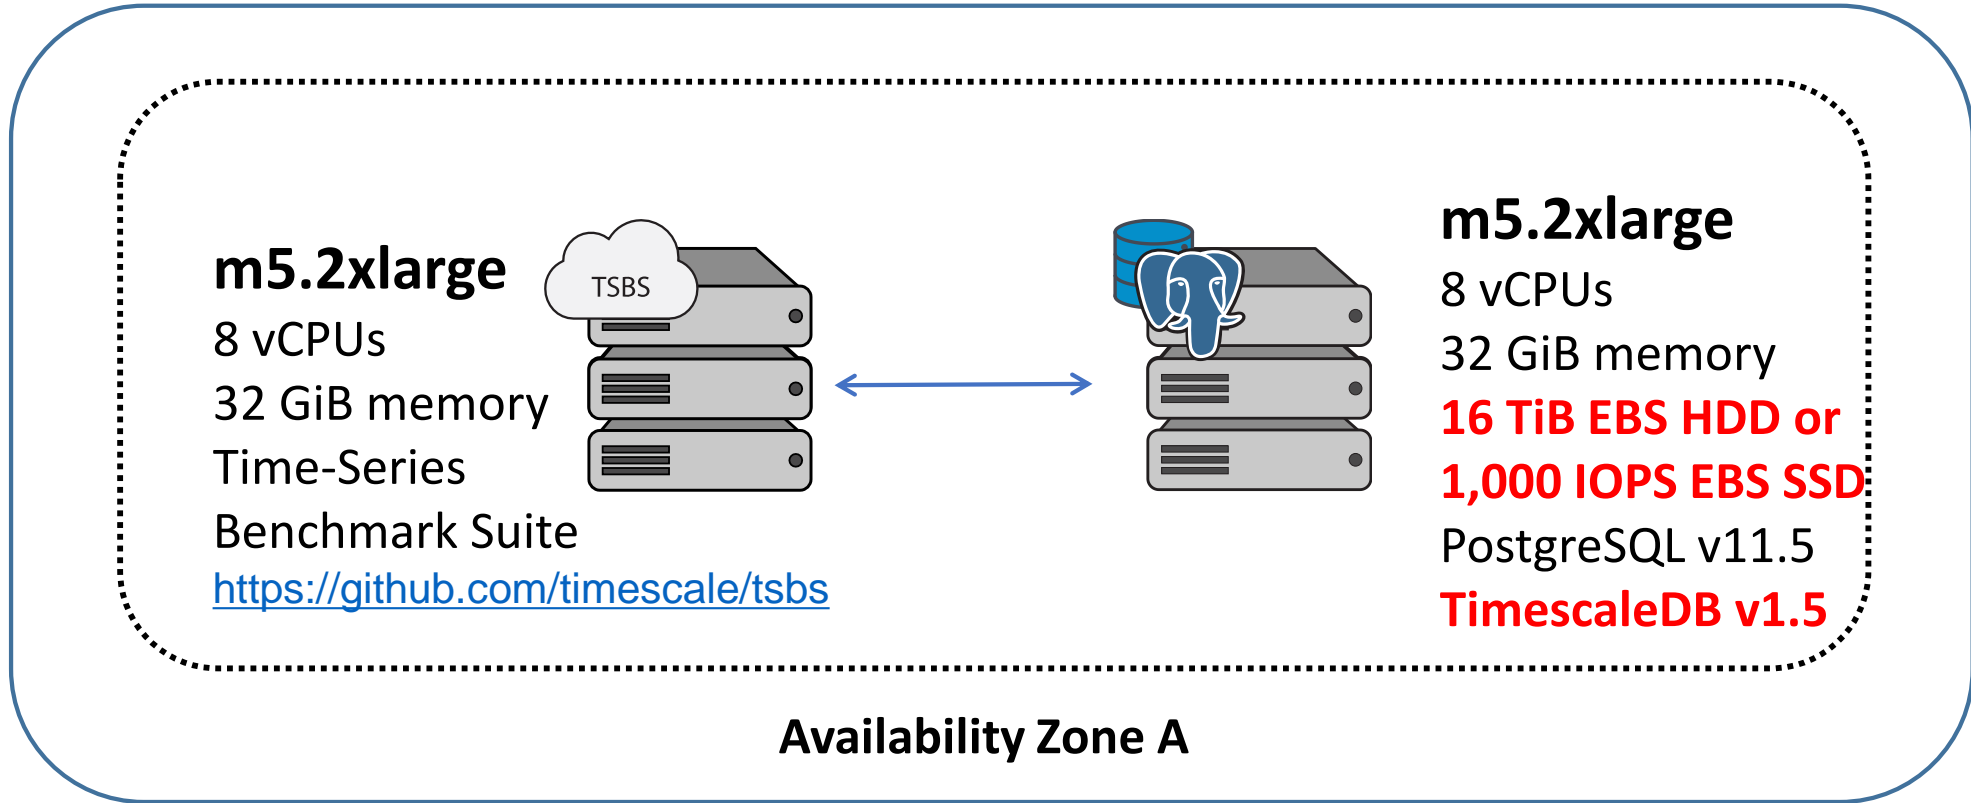
Extensibility

AppOS can seamlessly work with PostgreSQL-derived DBs



TimescaleDB Performance (1)

Test setup on AWS



m5.2xlarge

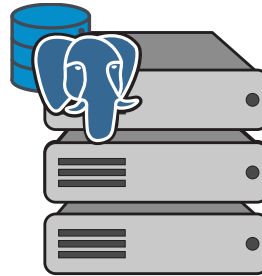
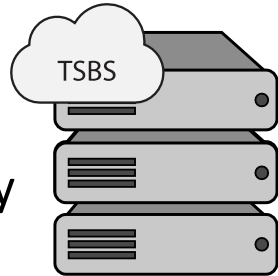
8 vCPUs

32 GiB memory

Time-Series

Benchmark Suite

<https://github.com/timescale/tsbs>



m5.2xlarge

8 vCPUs

32 GiB memory

**16 TiB EBS HDD or
1,000 IOPS EBS SSD**

PostgreSQL v11.5

TimescaleDB v1.5

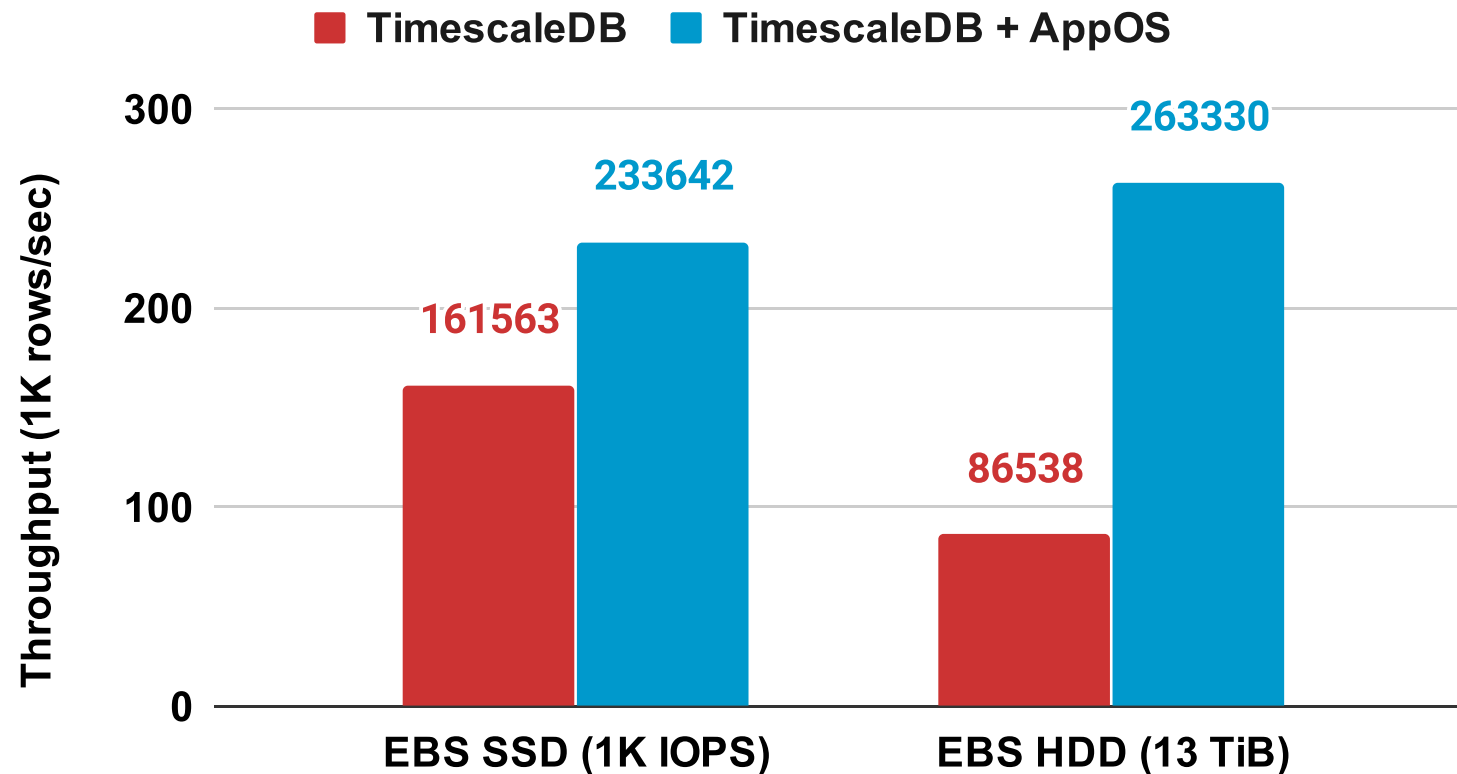
Availability Zone A

AWS Seoul Region

TimescaleDB Performance (2)

AppOS improves TimescaleDB inserts by 50%~3x

TimescaleDB inserts with 8 workers for 1B rows



We're partnering!

Try out AppOS via <https://apposha.io>

Join our Slack channel

<https://apposha.io/support>

<https://appos-postgres.slack.com/>



apposha



<https://apposha.io>
shawn@apposha.io