



StorPool
DISTRIBUTED STORAGE

**Latency: the #1 metric of your
cloud**

Boyan Krosnov



- Chief of Product
- Cloud architect
- "External technology"
- Previously service providers, networks, packet processing (SDN)
- Algorithms & data structures

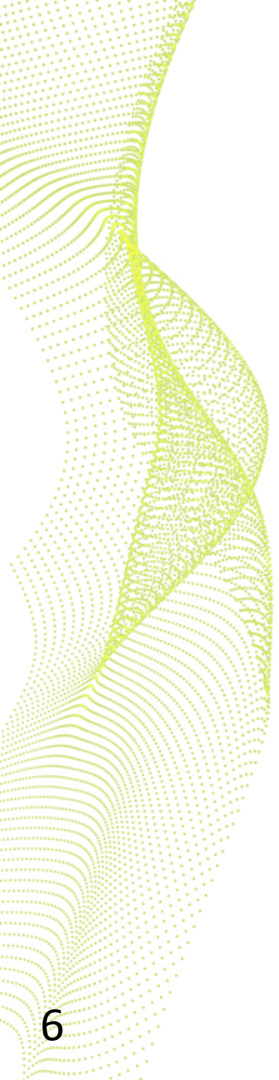
StorPool

- Fast and efficient software-defined storage system
- Used by Private clouds, MSPs, Cloud service providers, Mobile/Web Apps
- Started in 2011 to solve storage for cloud service providers
- Clean slate design - scale-out, API-controlled, end-to-end data integrity, CoW

StorPool

- Most deployments are with **KVM**
 - some Xen, some VMWare, some Hyper-V
 - some bare metal / dedicated server
- Deep integrations into **OpenStack**, **Kubernetes**, CloudStack, OpenNebula, OnApp






task


core


core



core

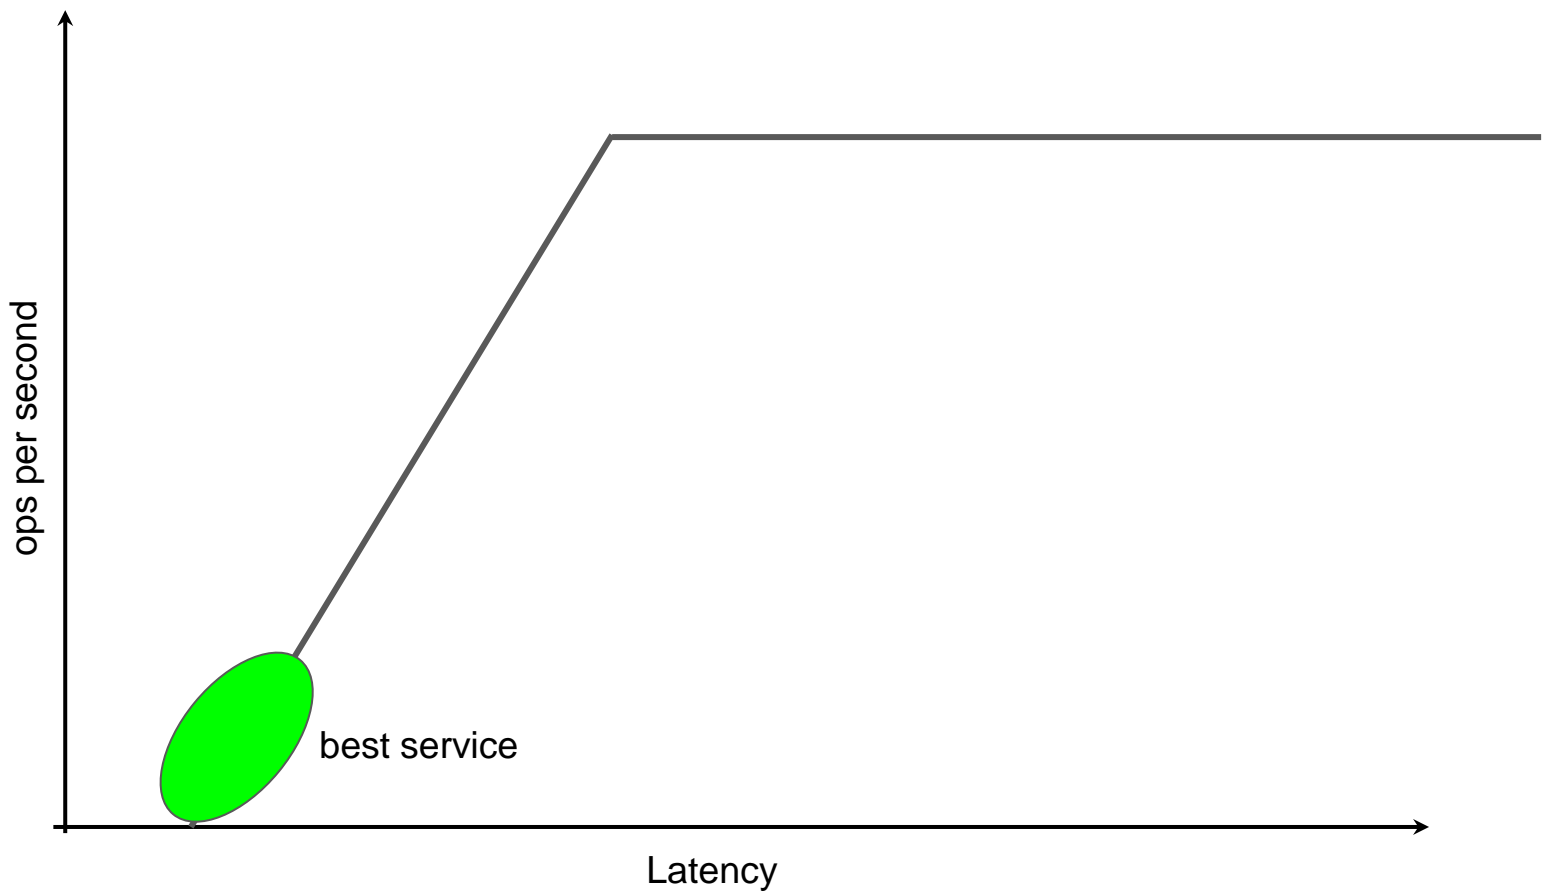

core

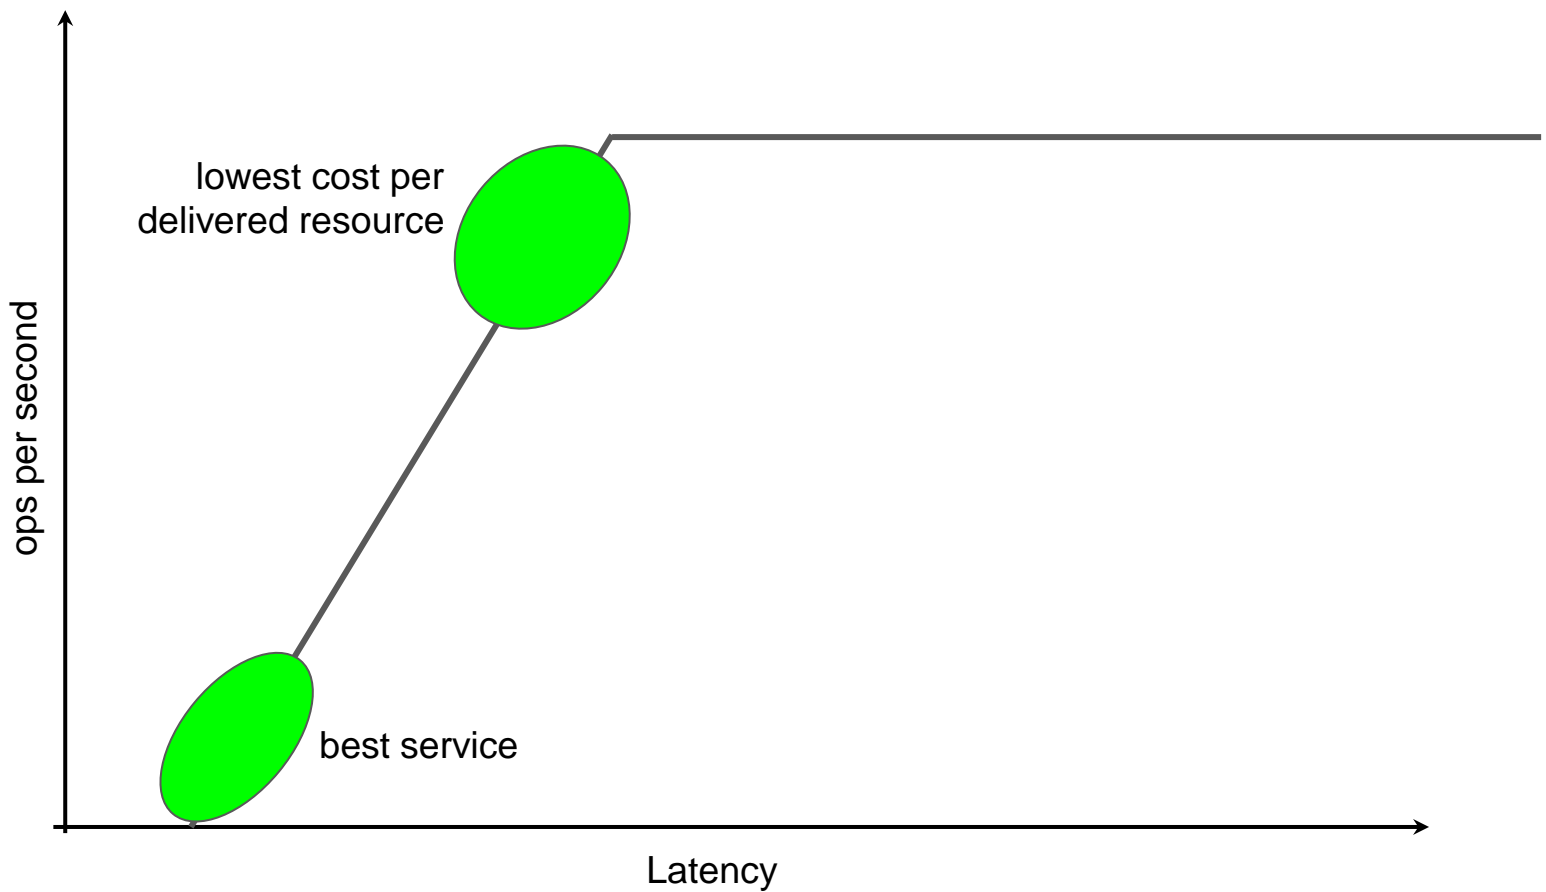



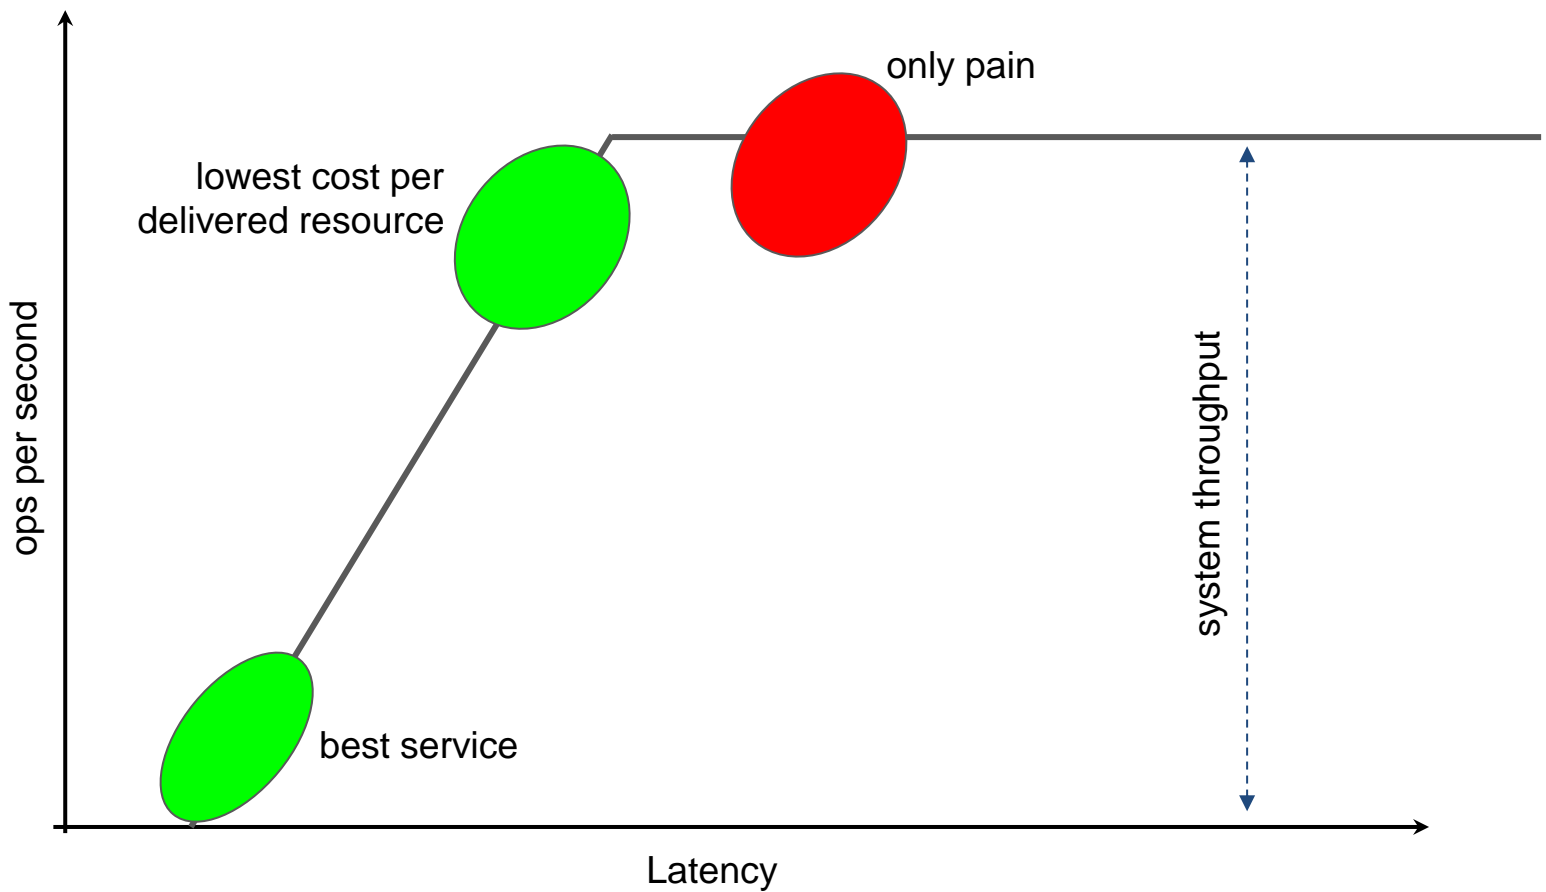



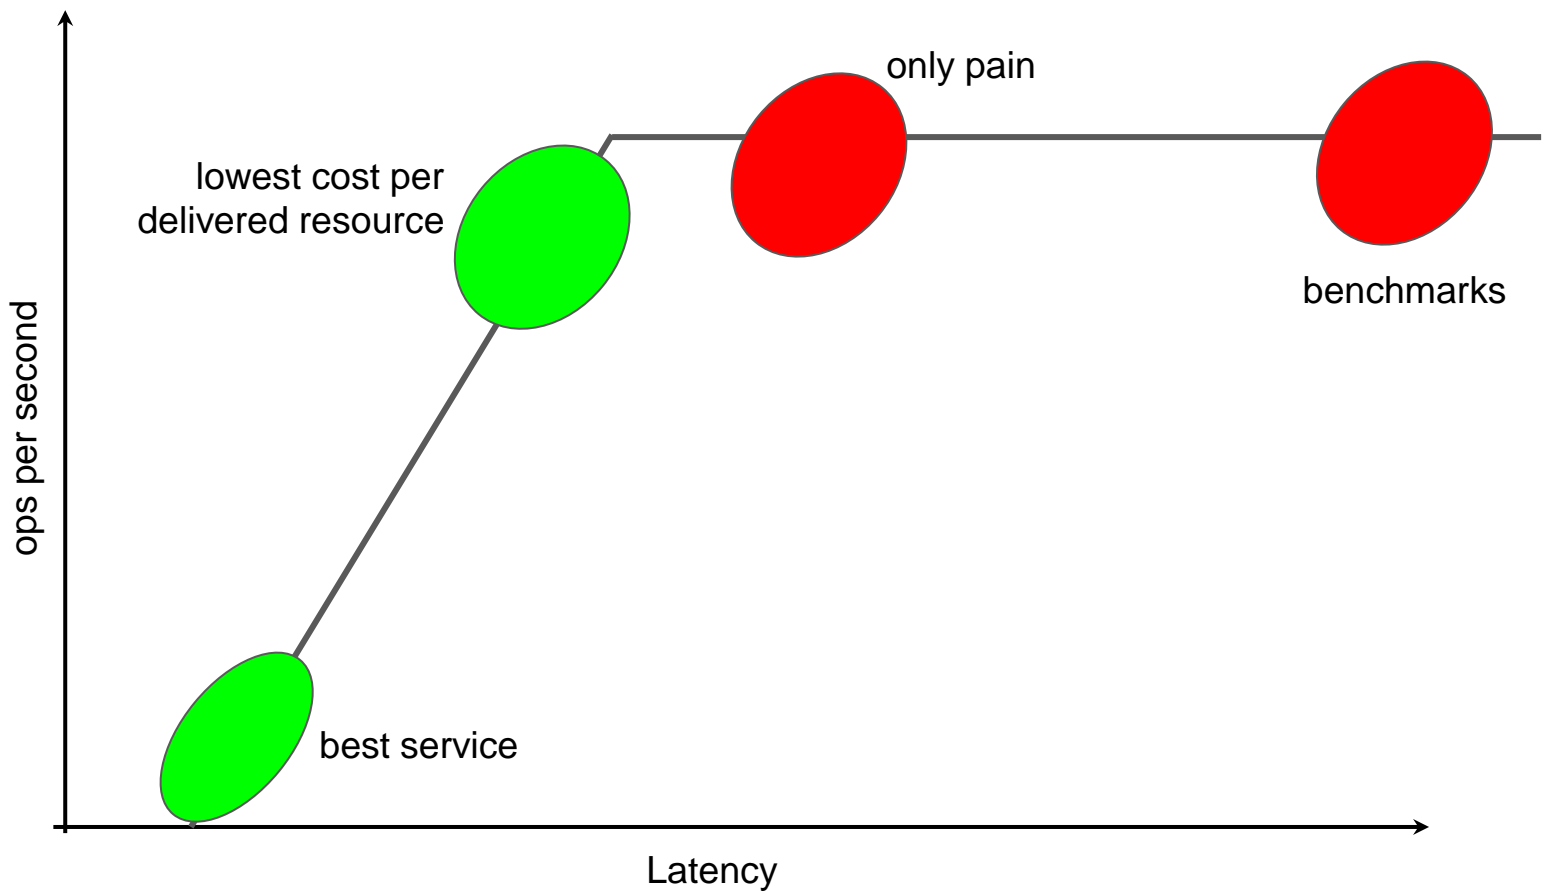


















8 vCPUs, 16 GB RAM



vdisk with 20k IOPS



"fast" volume with approx 0.15 ms latency (QD 1)

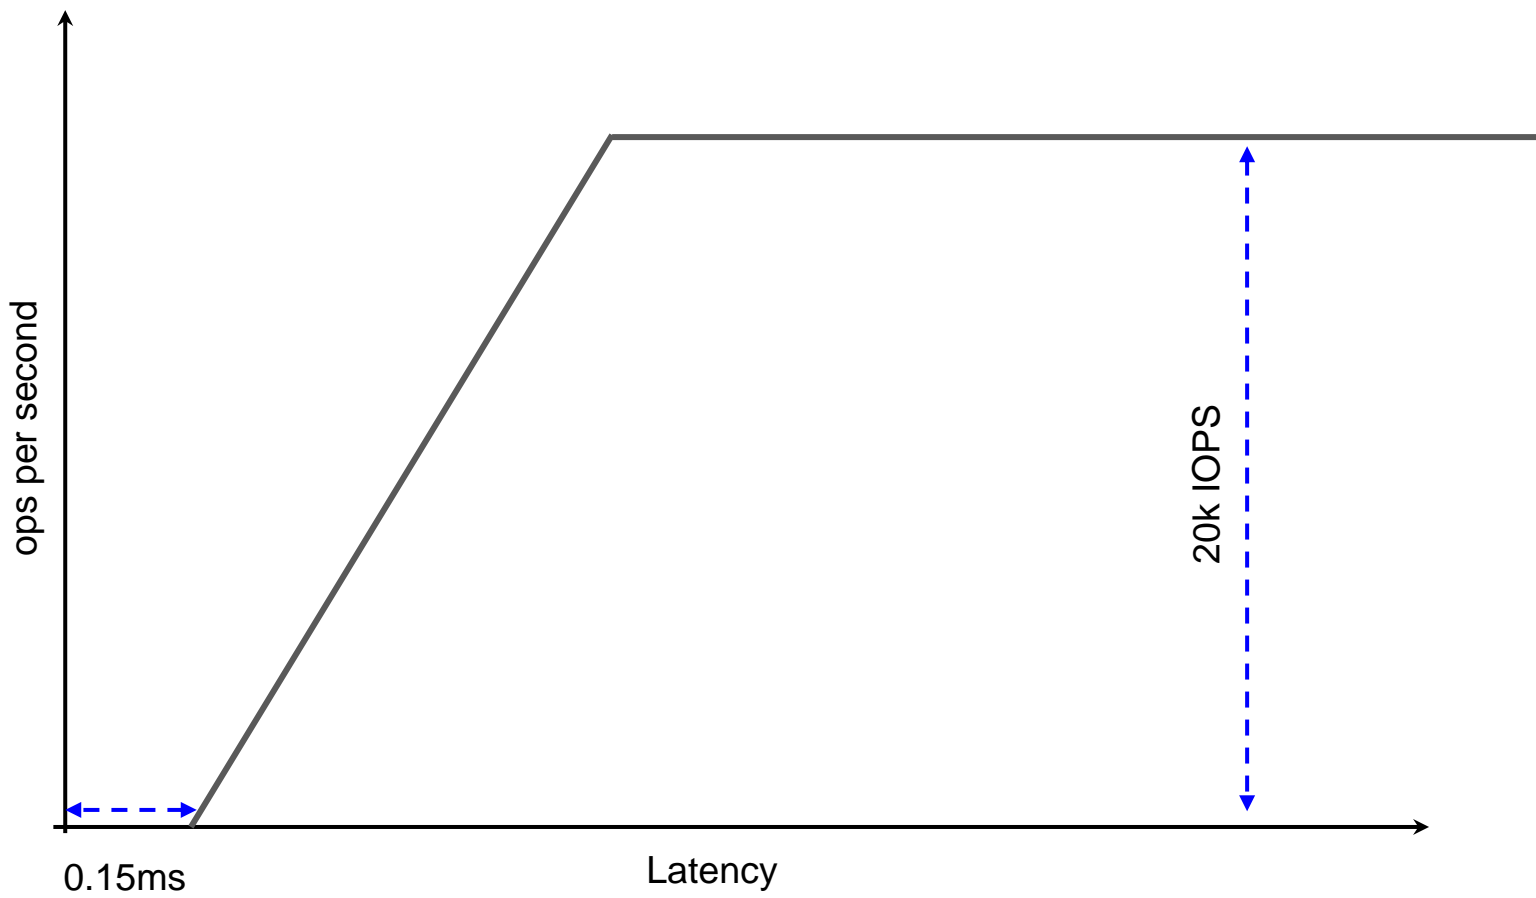

"slow" volume with 2 ms latency (QD 1)
added with dm-delay in host

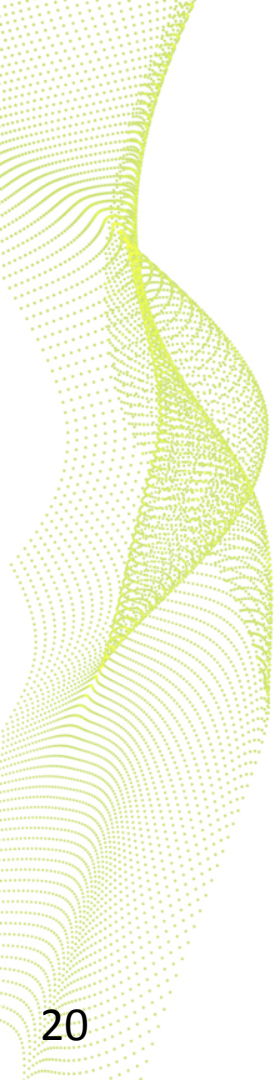
Both volumes are on the same pool

Both volumes measure 20k IOPS flat







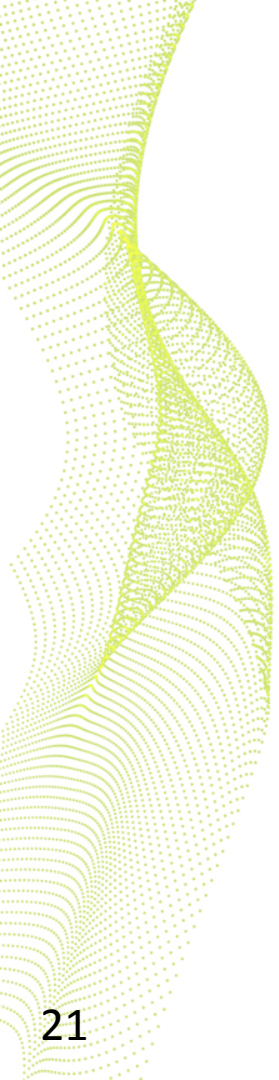


8 vCPUs, 16 GB RAM
vdisk with 20k IOPS

```
pgbench --client=16 --jobs=8 \  
--progress=1 --time=10 pgbench09
```

database size: 64 GB (4x RAM)

<https://wiki.postgresql.org/wiki/Pgbenchtesting>

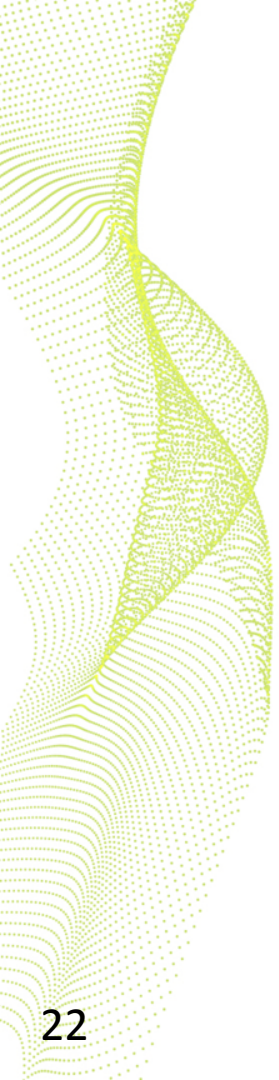


8 vCPUs, 16 GB RAM
vdisk with 20k IOPS

2ms storage latency

1600 TPS @ 8 ms in pgbench

-> if we ask for 2500 TPS -> pile up



8 vCPUs, 16 GB RAM
vdisk with 20k IOPS

0.15ms storage latency
3600 TPS @ 3 ms in pgbench
-> if we ask for 2500 TPS - no problem

random read/write 50/50 QD 1, avg. latency

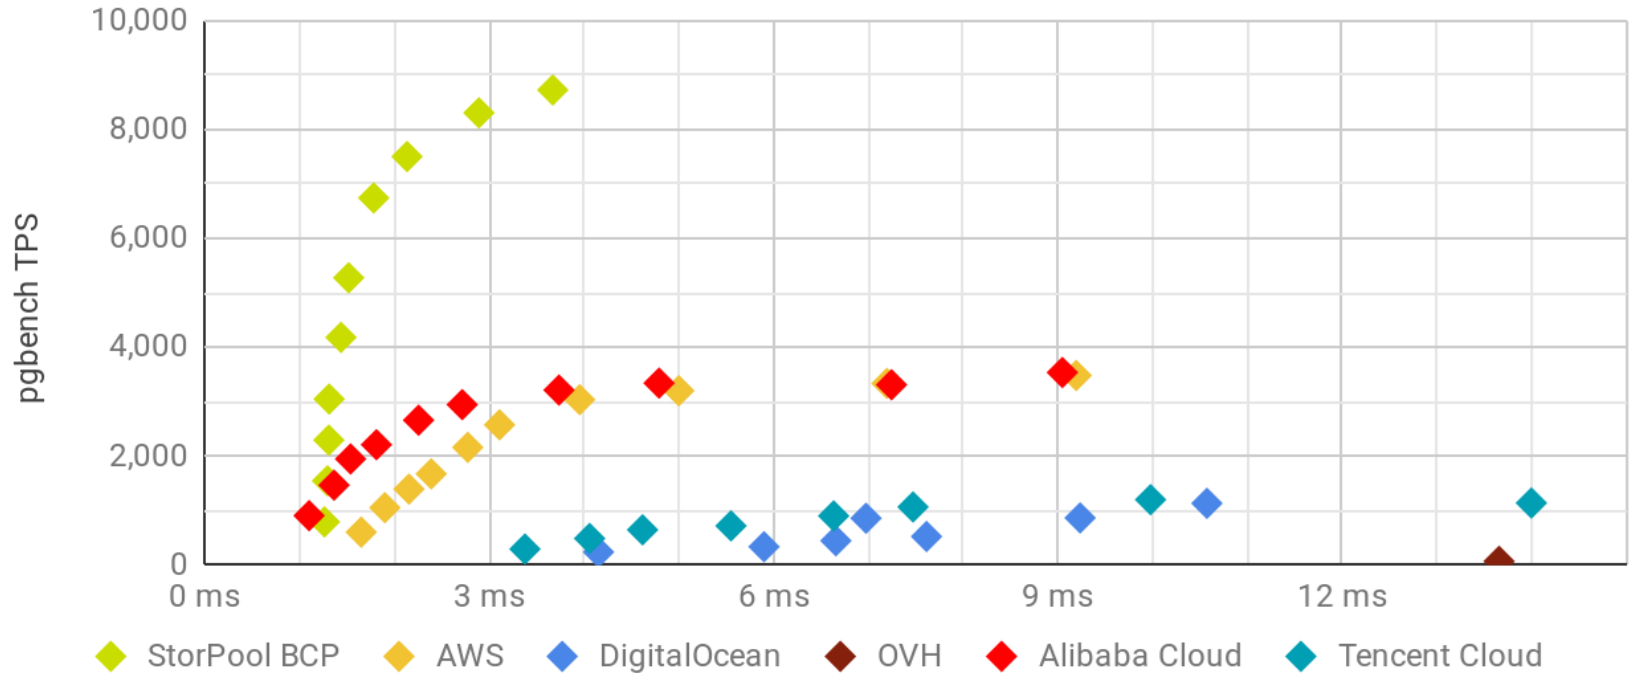
DigitalOcean (Ceph)	1.75 ms
OVH (Ceph)	1.53 ms
Tencent (?)	0.76 ms



random read/write 50/50 QD 1, avg. latency

DigitalOcean (Ceph)	1.75 ms
OVH (Ceph)	1.53 ms
Tencent (?)	0.76 ms
AWS gp2 10k	0.29 ms
Alibaba (?)	0.18 ms
StorPool BCP	0.17 ms

pgbench 4x RAM -- TPS vs latency [ms]



8 vCPUs, 16 GB RAM, dedicated if possible
database size = 4x RAM

<https://wiki.postgresql.org/wiki/Pgbenchtesting>



StorPool
DISTRIBUTED STORAGE

Q&A



StorPool
DISTRIBUTED STORAGE

Thank you!



StorPool
DISTRIBUTED STORAGE

Boyan Krosnov
StorPool Storage
bk@storpool.com

www.storpool.com
@storpool