# Public Cloud Performance

## MEASUREMENT REPORT

This report compares the block storage offerings of well-known public clouds (Amazon AWS, Alibaba, Google Cloud, and Microsoft Azure) with a StorPool-based public cloud offering.

StorPool
DISTRIBUTED STORAGE

# Introduction

In these public cloud performance tests, StorPool aimed to assess the block storage offerings of several public clouds – Amazon AWS, Alibaba, Google Cloud, and Microsoft Azure – and compare them against Katapult, a public cloud platform built on the StorPool block storage solution. To the best of our ability, we've selected VM instance types and their configurations to be identical. Therefore, we have an "apples-to-apples" comparison of only the underlying storage systems, adjusted for other aspects like memory where needed.

All storage systems tested for this report are in production clusters and part of generally available public cloud offerings, so our results are easily reproducible. The Katapult system is part of a production public cloud, so performance results on this service are directly comparable to performance results on the "big five" public clouds.

# About Katapult and the StorPool implementation in Krystal's infrastructure

Katapult is a virtual Infrastructure as a Service (IaaS) platform designed for extreme performance and simplicity. The solution is developed by Krystal, one of the largest independent UK web hosting companies. Katapult implements best-of-breed technologies and years of successful expertise in the cloud domain, backed up by an exceptional level of service.

Krystal selected StorPool because of its high performance, robust API, unique space-saving features, and extremely high level of data protection delivered by its triple data replication.

*"Storage forms the bedrock of any cloud platform, so whatever you use has to be bulletproof. At Krystal we've always had a standard rule; buy the very best solution available and sleep well at night! This is especially true when it comes to our clients' data, which is the most important thing we look after."*

Simon Blackler, CEO of Krystal

Explore our Customer Success Stories

## For our testing, we measured four benchmarks:

1. **PGBENCH** – a database benchmark, perhaps closest to "application performance"
2. **Sysbench/MySQL** – a second database benchmark for control of PGBENCH results
3. **fio** – a set synthetic benchmarks — random reads/writes, sequential reads/writes, latency measurements
4. **rsync** – copying files and syncing, simulating rapid deployment and backup workloads

Among these tests, we consider **PGBENCH** and **Sysbench/mysql** to be most representative of real-world transactional workloads (including most web applications). The remaining tests, using fio and rsync, are somewhat further away from being representative of real-world applications, but still serve as a synthetic measure of the ideal latency and maximum throughput possible with each storage technology.

In order to ensure a "level playing field", we used the same OS and software versions for each test. For completeness, we've shown the version of the StorPool software-defined storage (SDS) used as well.

| Software | Version |
|---|---|
| Ubuntu | 22.04 |
| postgresql | 14 |
| fio | 3.1 |
| rsync | 3.1.2 |
| StorPool | 21.0.606 |

## The Virtual Machines

We provisioned six virtual machines with identical parameters from all clouds we tested. We selected medium-sized VMs with 16 GB RAM and 8 dedicated vCPUs to represent a medium-sized database server, the same type of system at the heart of most web applications.

| Date of test | Provider | Instance name | Region | Montly cost (with 12 month commitment) | vCPUs | RAM |
|---|---|---|---|---|---|---|
| April 2024 | Katapult on StorPool Storage | k-8vcpu-24gb | London | $120 | 8 | 16GB |
| April 2024 | AWS | c7a.2xlarge | us-east-1f | $300 | 8 | 16GB |
| May 2024 | Alibaba Cloud | c7a.2xlarge | China (Hangzhou) | $131 | 8 | 16GB |
| May 2024 | Google Cloud | n2d-8vcpu-16gb | us-central1 | $214 | 8 | 16GB |
| May 2024 | Microsoft Azure | F8s_V2 | East US 2 | $265 | 8 | 16GB |

## The Storage Volumes

We provisioned block storage volumes associated with each VM and attached them to the VMs as virtual disks. The goal here was to give each provider a fair chance with their SSD-based block storage offering. We provisioned an over-sized 1TB virtual disk in each cloud, even though typical databases on a VM of this size would be 50-200GB. We did this to ensure that clouds that limit IOPS based on the number of GB of storage used would not be severely limiting in tests performed on those clouds or have an effect on overall performance. Interestingly, this actually skews the results in favour of the big four clouds.

With StorPool, our customers who build clouds either don't apply any IOPS limits or set a very high IOPS limit. This allows each VM to receive high performance for any bursts in their workloads. For example, Katapult provisions their customers' volumes with no IOPS limit. This is typical usage for StorPool systems. StorPool delivers a very high IOPS capability, so usually our customers only use IOPS limits as a policing action to combat abuse.

| Date of test | Provider | Service name | $/GiB/ month | IOPS limit | Size of volume [GiB] | Monthly cost |
|---|---|---|---|---|---|---|
| April 2024 | Katapult/StorPool | Shared disk NVMe | $0.15 | 100,000 | 980 | $147 |
| April 2024 | Amazon EBS io2 | EBS io2 | $2.52 | 40,000 | 1024 | $2,576 |
| May 2024 | Amazon EBS gp3 | EBS gp3 | $0.18 | 16,000 | 1024 | $182 |
| May 2024 | Alibaba Cloud | PL2 ESSD | $0.65 | 53,000 | 1024 | $668 |
| May 2024 | Google Cloud | SSD Persistant Disk | $0.17 | 15,000 | 1000 | $170 |
| May 2024 | Microsoft Azure | Premium SSD v2 | $0.18 | 16,000 | 1024 | $186 |

## PGBENCH Large DB test parameters

- Simulate a large OLTP (transactional) database — like most web and mobile applications use
- "On-Disk Test" settings from Postgresql Wiki
- Dataset is 64GB, i.e. 4x RAM
- Number of threads (-j parameter) is one-half of the number of clients
- The number of clients (-c parameter) was varied between 1 (to simulate a light load and lowest latency) and 32 (to simulate a workload that was more than the maximum recommended production workload for an - vCPU database)

Reference for test scenarios and parameters we tested:.
https://wiki.postgresql.org/wiki/Pgbenchtesting#Memory_vs._Disk_Performance

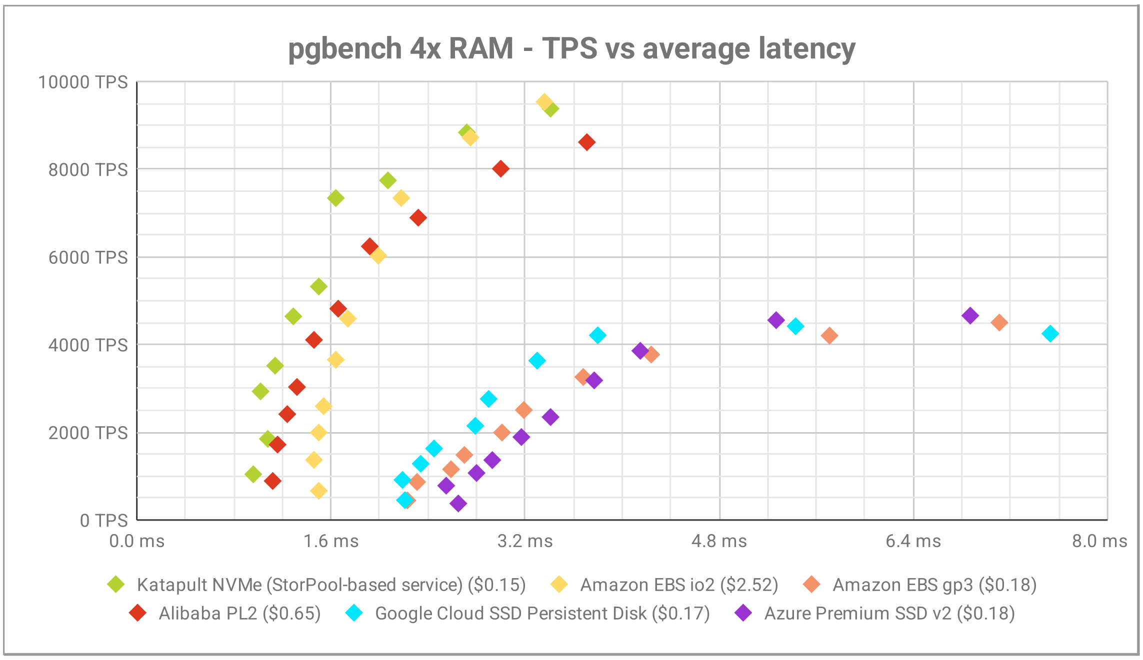


This graph shows extreme differences in the performance of the database on the VMs running on each of the 6 different storage systems (including 2 different services from Amazon) we tested. We investigated the cause of the large difference in results by running paired experiments (e.g., higher vs lower IOPS limits, higher-speed CPUs vs slower-speed CPUs). The different results can't be caused by differences in CPU or memory since we used the same exact configuration in all the tests. The fundamental difference in the test results observed can only be attributed to the differences in performance and behaviour of each storage system.

On the same storage system, an **IOPS limit** generally controls how many transactions per second (TPS) systems can support at large latency (to simulate batch processing) but **does not directly affect latency under a fixed transactional load**. For example, for a fixed workload of 2,000 TPS, the StorPool NVMe configuration used by Katapult delivered a per-transaction latency of 1ms while Google Cloud delivered a per-transaction latency of approximately 2.5ms. The difference between these latencies is almost wholly attributable to storage system latency.

## PGBENCH Small DB

- Simulate a small OLTP (transactional) database — like most web and mobile applications use
- "Mostly Cached" settings from Postgresql Wiki
- Dataset is 14.4GB, i.e., 0.9x RAM
- Number of threads (-j parameter) is one-half of number of clients
- The number of clients (-c parameter) is varied between 1 (to simulate a light workload and lowest latency) and 32 (to simulate a workload that was more than the maximum recommended production workload for a 8-vCPU database)

Reference for test scenarios and parameters we used:
https://wiki.postgresql.org/wiki/Pgbenchtesting#Memory_vs._Disk_Performance

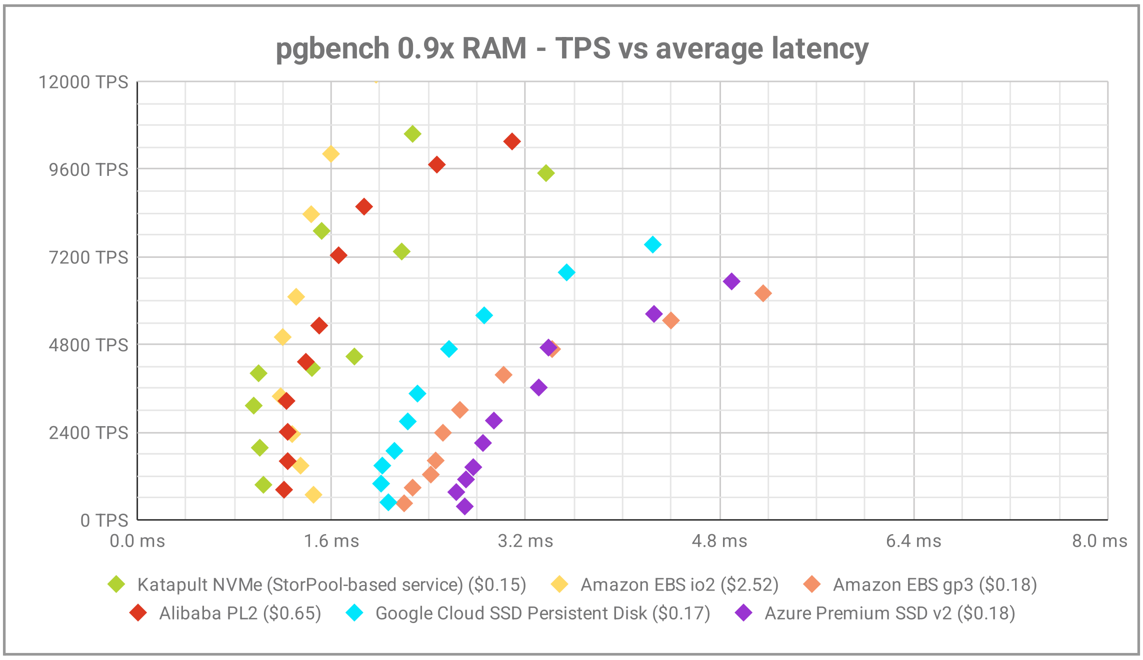


Even when a database is smaller than a server's RAM, the storage system's performance greatly influences the performance of the database.

## Sysbench/MySQL

We ran Sysbench with a MySQL database, for control over the PGBENCH results. This shows that the results and extreme differences are reproducible with a completely different benchmark and database stack.

The database sizes used in the testing were 50GB and 5GB, respectively.

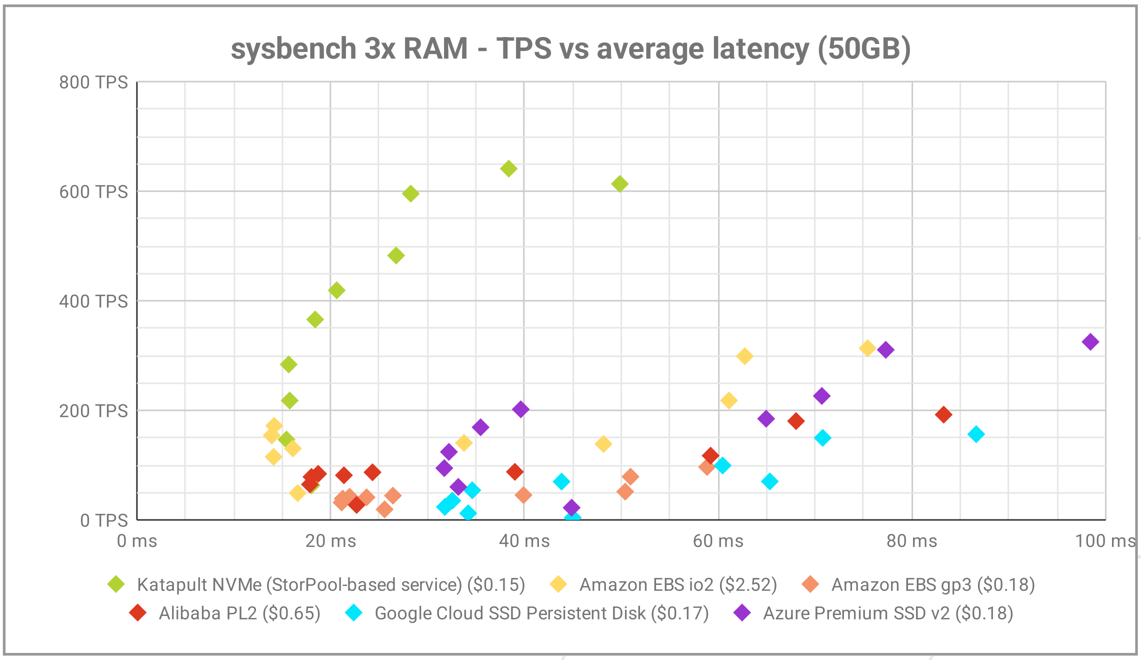

From the chart above, it's clear that, for large MySQL-like databases, Katapult's offering, built on StorPool Storage, stands out as the only solution that can serve many clients at consistently low latencies.

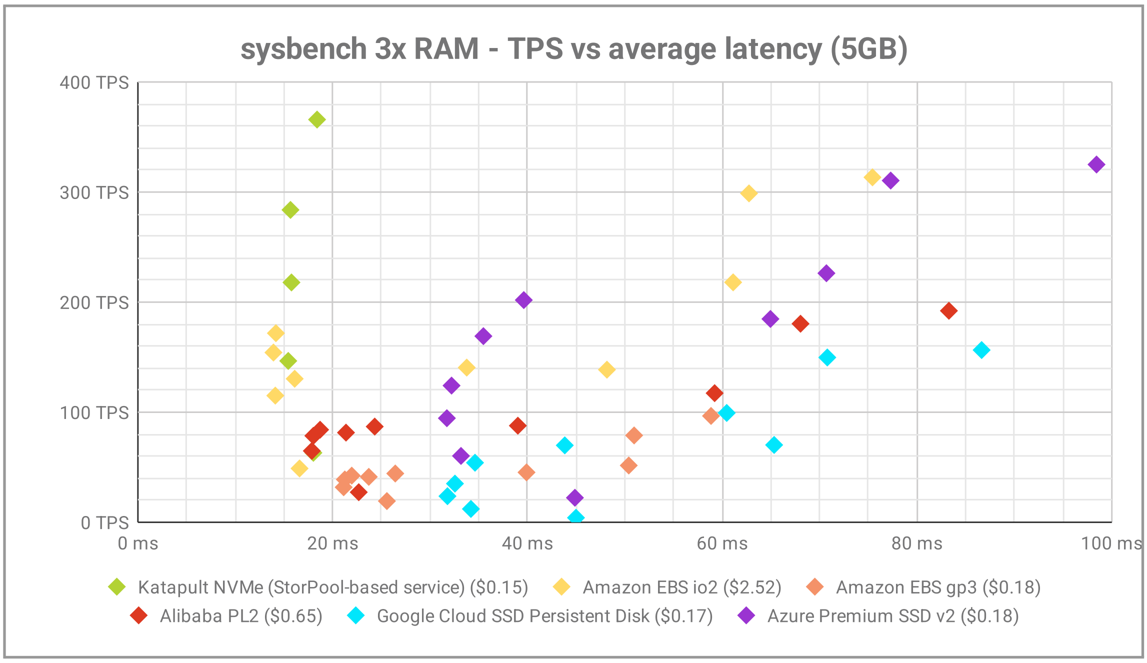


Meanwhile, the second chart shows that, for small databases, the only offerings capable of keeping up with Katapult's service levels are the Amazon EBS io2 and Alibaba However, at 100 or more TPS, Alibaba's latency increases dramatically, while Amazon's latency increases even more than that at 200 or more TPS.

It's worth noting that Alibaba's storage service comes at 4.5 times the cost of Katapult's all-NVMe storage offering, while Amazon's EBS io2 comes at more than 17 times the cost of Katapult.

## FIO Tests

We ran the usual suite of synthetic benchmarks, which show the performance envelope of the tested services. These are:

• Random 4k queue depth 1 – for latency under light load
• Random 4k queue large queue depth – for IOPS throughput
• Sequential workload with large queue depth – for MB/s throughput

# FIO Latency

- Transactional workload
- Random reads; random writes
- Block size 4k
- Queue depth 1
- Simulates small transactional workload

The poor results on these tests of Amazon, Alibaba, and Azure, especially for writes, are most likely the reason PGBENCH, and Sysbench/MySQL are showing such high latency on these tests.

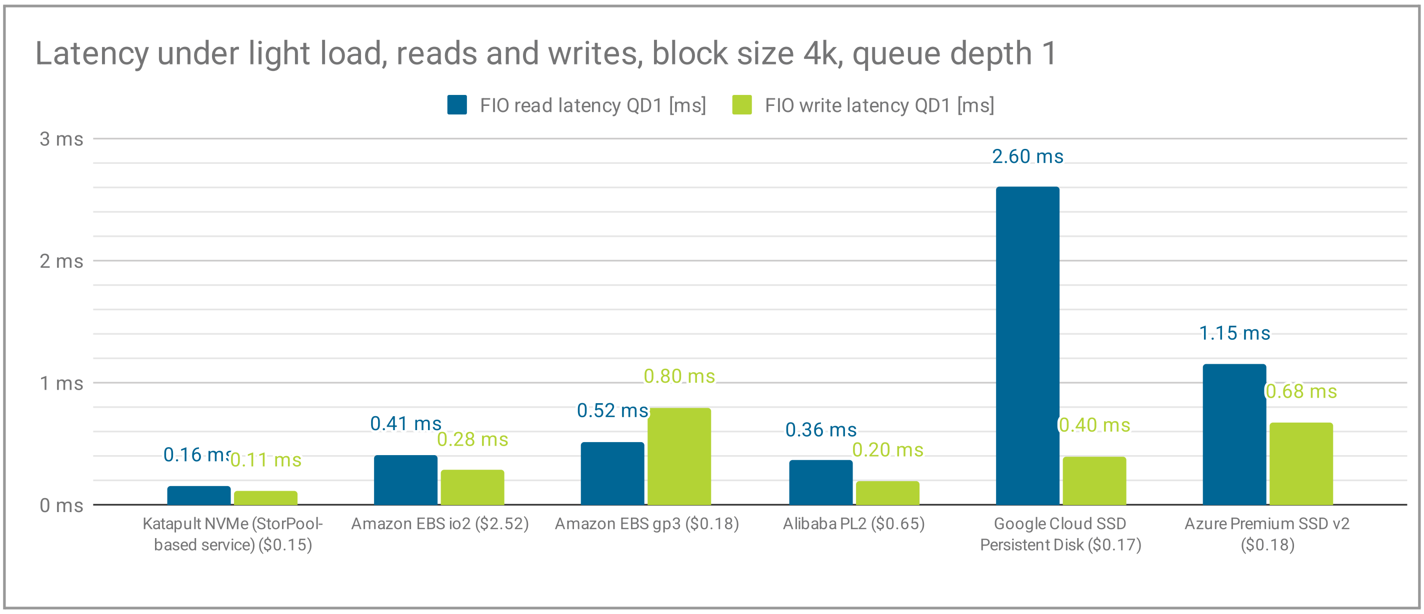


The chart above shows the latency results of the test (lower is better), while the chart below shows the test results as IOPS to make it easier to compare the top performers (higher is better):

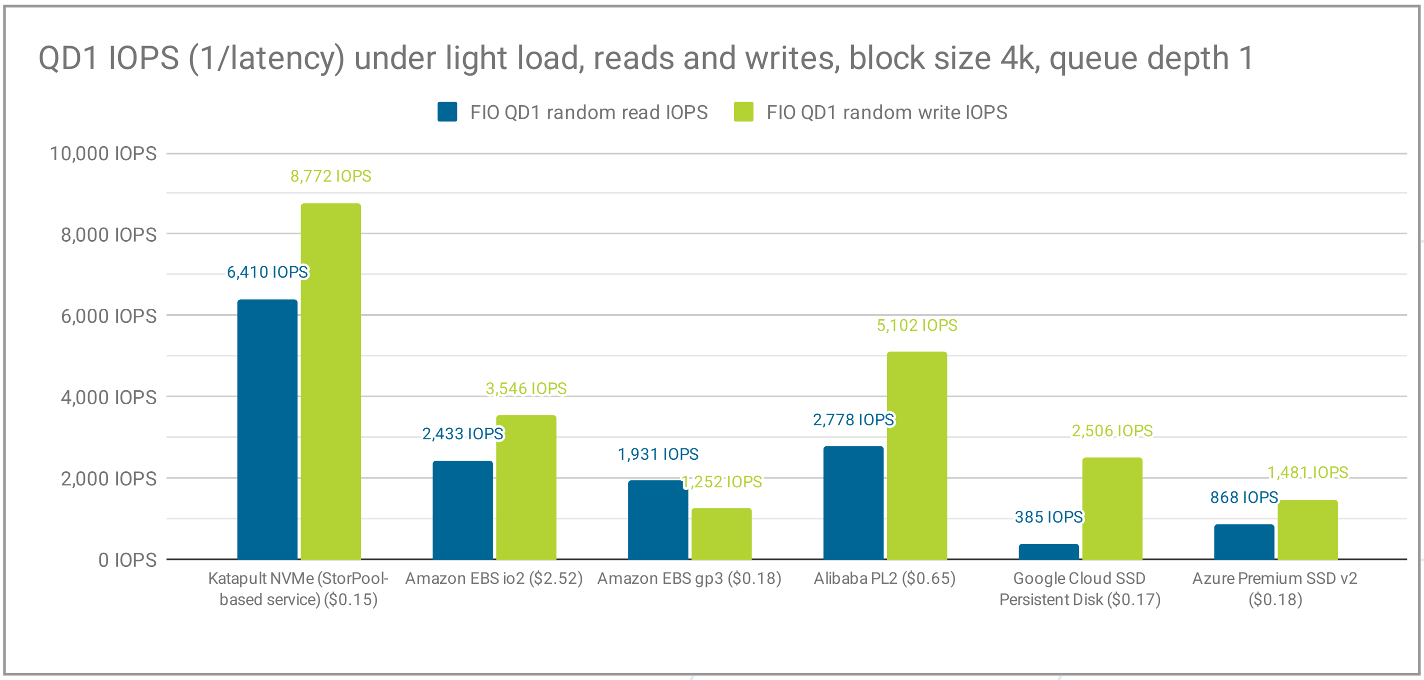

# FIO IOPS

- Parallel random workload
- Random reads and writes, 50/50 mix
- Block size 4k
- Queue depth 64

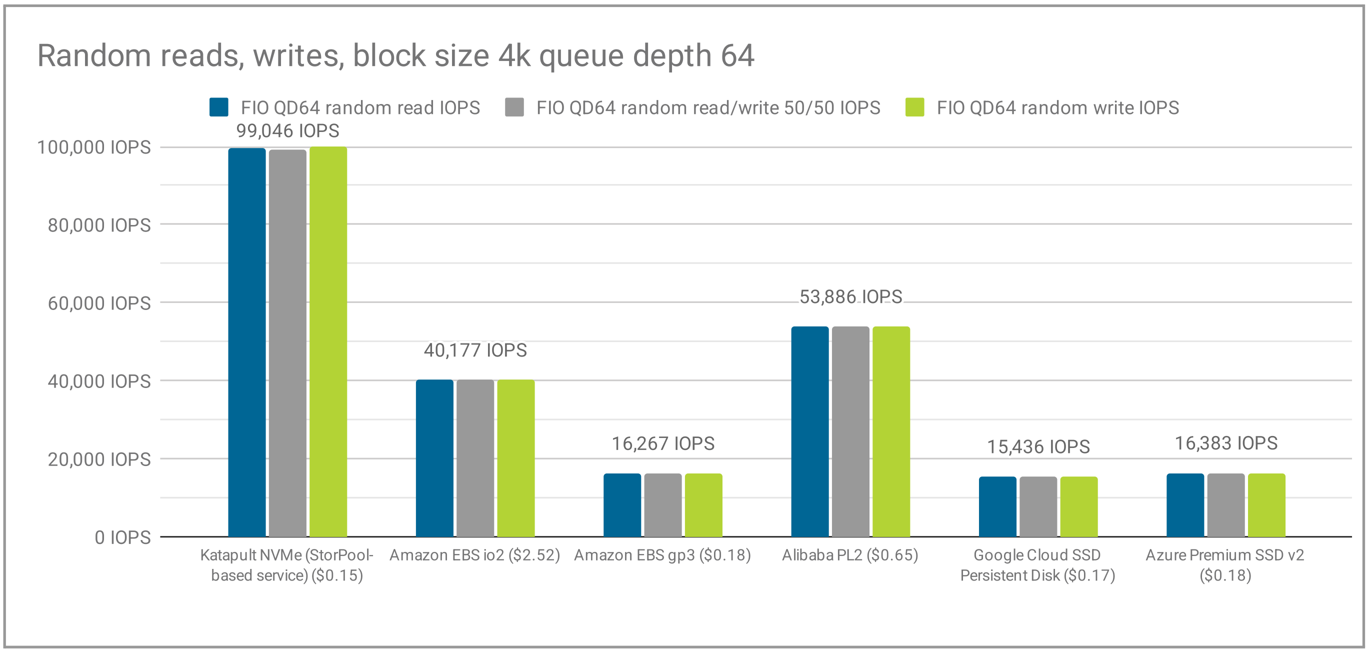


# FIO MB/s

- Streaming workload
- Sequential reads; sequential writes
- Block size 128k
- Queue depth 64

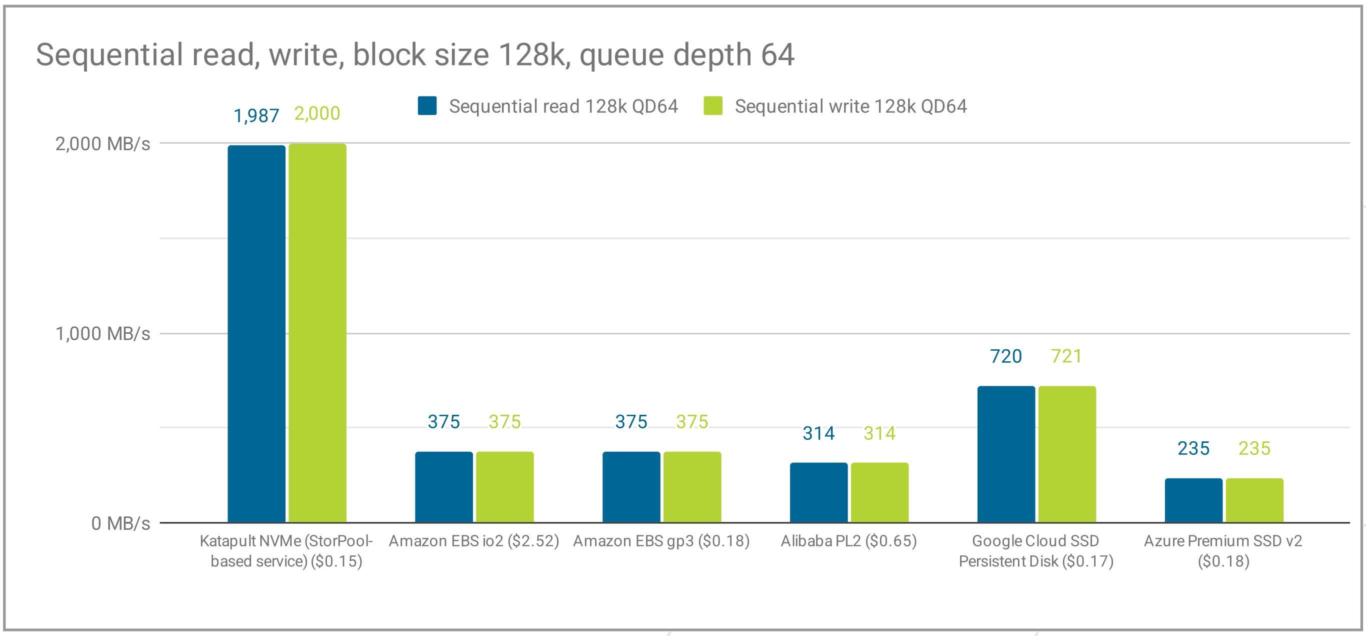

## RSYNC copy

- ext4 filesystem
- Linux kernel source (4.17.13) times 50
- 3.12M files
- 46 GB data
- Destination directory is empty
- Clear the cache, then copy files to destination directory
- This test approximates deployment workload – readdir(), sequential reads, sequential writes

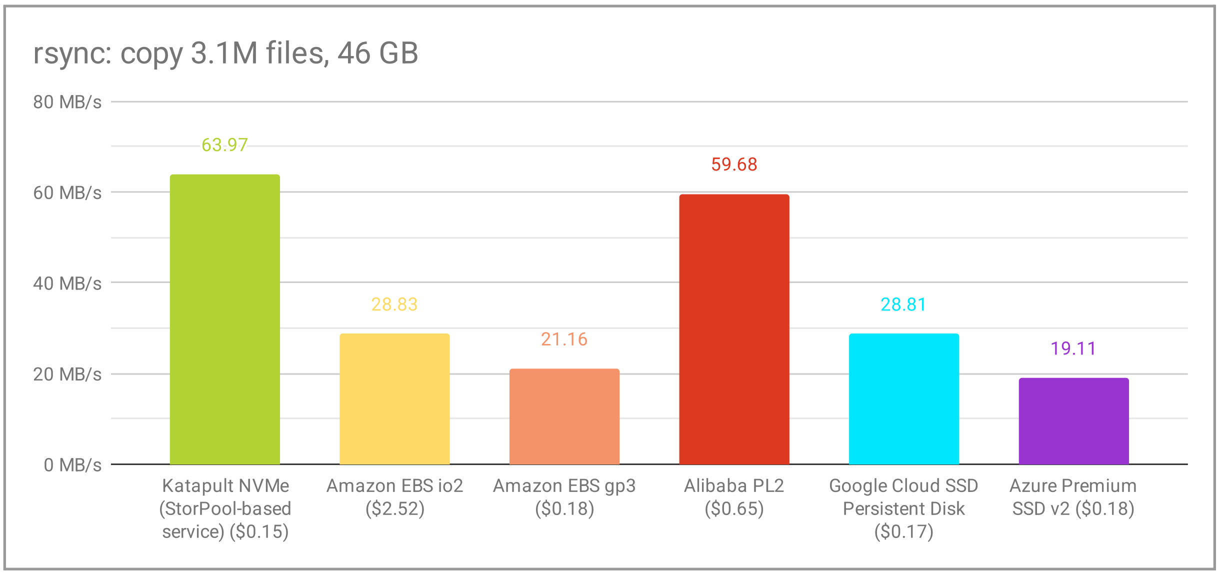


## RSYNC sync

- ext4 filesystem
- Linux kernel source (4.17.13) times 50
- 3.12M files
- 46 GB
- Source and destination directories are identical, each holding 3 million files
- Clear cache, then sync files to destination directory
- This test approximates backup workload – dominated by readdir() and stat()

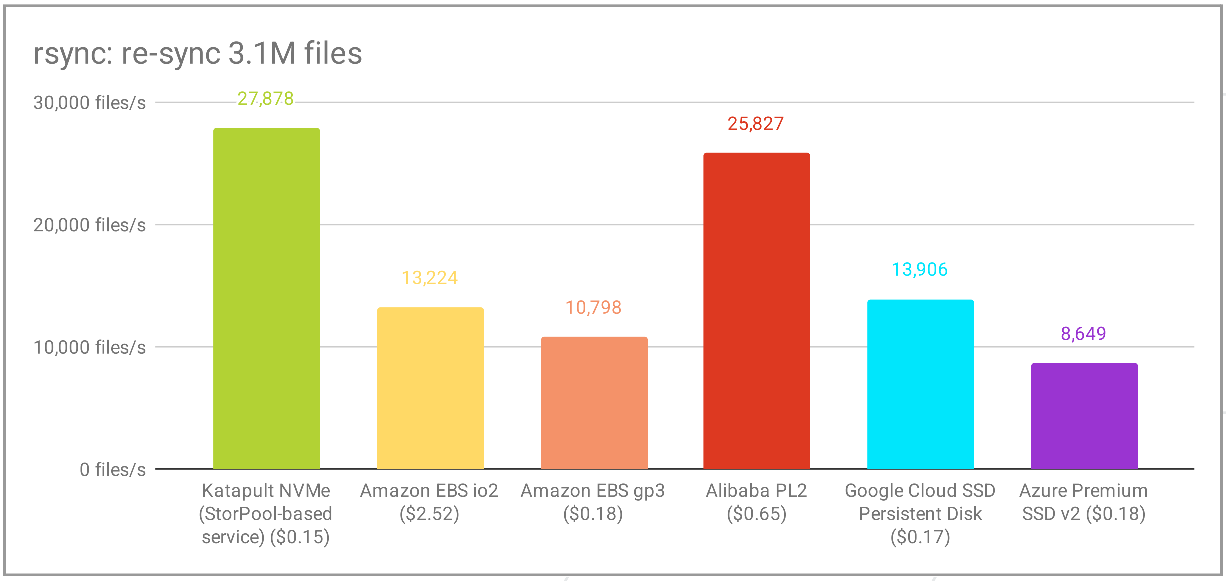

## Conclusions

The biggest contributor to slow application response times is the latency of the storage system. Configuration differences, like network bandwidth, number and speed of CPUs, and the amount of RAM can contribute to slow application response times, but still not as much as storage system latency. For these tests, we used VMs with identical parameters like CPU and amount of RAM, leaving the data storage as the only changing variable. Memory can have orders of magnitude differences in application performance. This demonstrates quite clearly that the choice of storage solution can make a great difference in overall performance.

While most people think of IOPS and throughput as the measurement for performance, these test results show that, while they're important, storage latency is the true measure of storage performance. The lower the latency of the storage system, the higher the application performance.

End-user applications running on StorPool-based public clouds perform demonstrably and measurably better (in some cases multiple times better) than even the second-best public cloud offering, while delivering storage at a significantly lower cost..

StorPool
DISTRIBUTED STORAGE

## About StorPool

StorPool is a software provider that develops the most reliable and speedy storage platform on the market. StorPool Storage is the easiest way to convert sets of commercial off-the-shelf servers into primary storage systems for cloud infrastructure. Public and private cloud builders - Managed Services Providers, Hosting Services Providers, Cloud Services Providers, enterprises, and SaaS vendors - use StorPool Storage as the foundation for their clouds.

StorPool Storage is designed for large-scale deployments, but it has efficient resource consumption and can start small. Each cluster scales seamlessly online - adding drives or servers expands both its capacity and performance. Adjusting StorPool volume capacity and performance also happens online without disrupting user workloads. Updates are also carried out online, without any interruptions to user-facing services.

The software comes as an utterly hands-off solution - the StorPool team architects, deploys, tunes, monitors, and maintains each StorPool Storage system so that end users experience speedy and reliable services.

StorPool Storage is a superior alternative to mid- and high-end SANs and All-Flash Arrays (AFA) for large-scale deployments (hundreds of terabytes to petabytes of storage).

StorPool
DISTRIBUTED STORAGE

## High Performance Linearly Scalable Primary Storage Platform

The ideal foundation for large-scale clouds running diverse, mission-critical workloads.

### Get in Touch


+ 1 415 539 0356

+359 (0) 88 834 6465

info@storpool.com

sales@storpool.com


Get Started | Talk to an Expert

www.storpool.com