

The Best Benchmark is Your Production Workload


Finding ways to measure the application performance differences between different NVM devices

Kevin O'Brien - StorageReview

Consistent Comparisons Are Key



- Use the same hardware as the final production environment
- Don't trust synthetic benchmarks in place of actual application performance
- Application workloads should be preconditioned and taken into steady-state
- Certain application workloads favor different NVM architectures
- Software that supports new flash-aware APIs can improve performance and endurance over legacy standards

Synthetic Benchmark Comparison




Synthetic benchmarks show SSD A is up to 37% faster at its peak than SSD B... how does that translate into application performance?

MarkLogic NoSQL Benchmark



- Comparing overall average latency, SSD B is 3% faster than SSD A in an actual application workload measuring NoSQL database performance.
- Large gains in synthetic benchmarks didn't pan out in NoSQL real-world conditions.

Sysbench OLTP Benchmark



- Measuring OLTP performance with Sysbench Percona in a MySQL database, SSD A offered 19% greater TPS than SSD B.

Synthetic vs Application Performance

Synthetic benchmarks only offer value up to a certain point and can exaggerate performance differences.
Application performance should be the primary consideration

- Synthetic benchmarks showed SSD A offered 37% greater I/O performance with an 8k 70/30 workload
- MySQL OLTP performance with Sysbench showed only 19% gains on the top-end from SSD A
- NoSQL performance with MarkLogic showed SSD B leading by 3%