



Performance Benchmarks

Reading Between the Lines

As we move into 2019, storage performance becomes more important than ever. At the media level, vendors are pushing the boundaries with fast solid-state technologies like NAND flash and 3D-Xpoint. These products provide performance capabilities we've never seen before. Storage system vendors then take components and build storage solutions, either as bespoke hardware or part of a broader technology solution.

These vendors like to talk about hero numbers. It's a subject we've covered before, both as a podcast and in blog posts. Pushing past the hype, how should we analyse storage performance figures? Does the top line number always represent the truth?

Measurement Rationale

There's an old adage that says, "If you can't measure it, you can't improve it". With enterprise storage, you could say that if you can't measure it, you can't differentiate between solutions. As a result, we see vendors creating their own performance figures as well as running industry accepted benchmarks. Let's think about vendor-provided numbers.

- [#58 - Storage Vendor Hero Numbers](#)
- [Why Deterministic Storage Performance is Important](#)

First of all, these figures will always show technology in the best light. There are lots of techniques to ensure that a vendor's solution looks the best. This might be achieved by quoting read-only I/O or showing the "at best" number, like the lowest expected latency rather than the average. We can treat vendor-generated data with a pinch of salt as they are unlikely to be independently verified or follow a standard process.

Industry Accepted

Then there are industry-accepted benchmarks. I step back from saying "industry standard" as there are no industry-standard benchmarks for data storage. We do have accepted units of measure (latency, bandwidth, throughput), but even here there are no accepted standards. What do we mean by this? Well, take bandwidth and throughput. Both measure the performance capability of a system - how much data can you push through the system in a measured time interval. Bandwidth typically measures IOPS or I/Os per second. Throughput measures data volume, usually in MB/s (megabytes) or GB/s (gigabytes). However, use small blocks of data and you get more IOPS. Use large blocks of data and you typically get more throughput. In both cases, there is no generally accepted figure that should be used as the block size in each instance.

Performance Profiling

Getting back to industry-accepted data, organisations such as the [Storage Performance Council](#) (SPC) provide a way to perform rigorous testing on storage hardware solutions using a consistently applied methodology. Storage vendors submit hardware configurations for testing that are then validated independently across a battery of test scenarios. The aim is to show how a storage system performs against its peers, taking into consideration common variabilities like cost and capacity. As we will go on to discuss, non-technical metrics like price can affect the way in which performance results are interpreted.

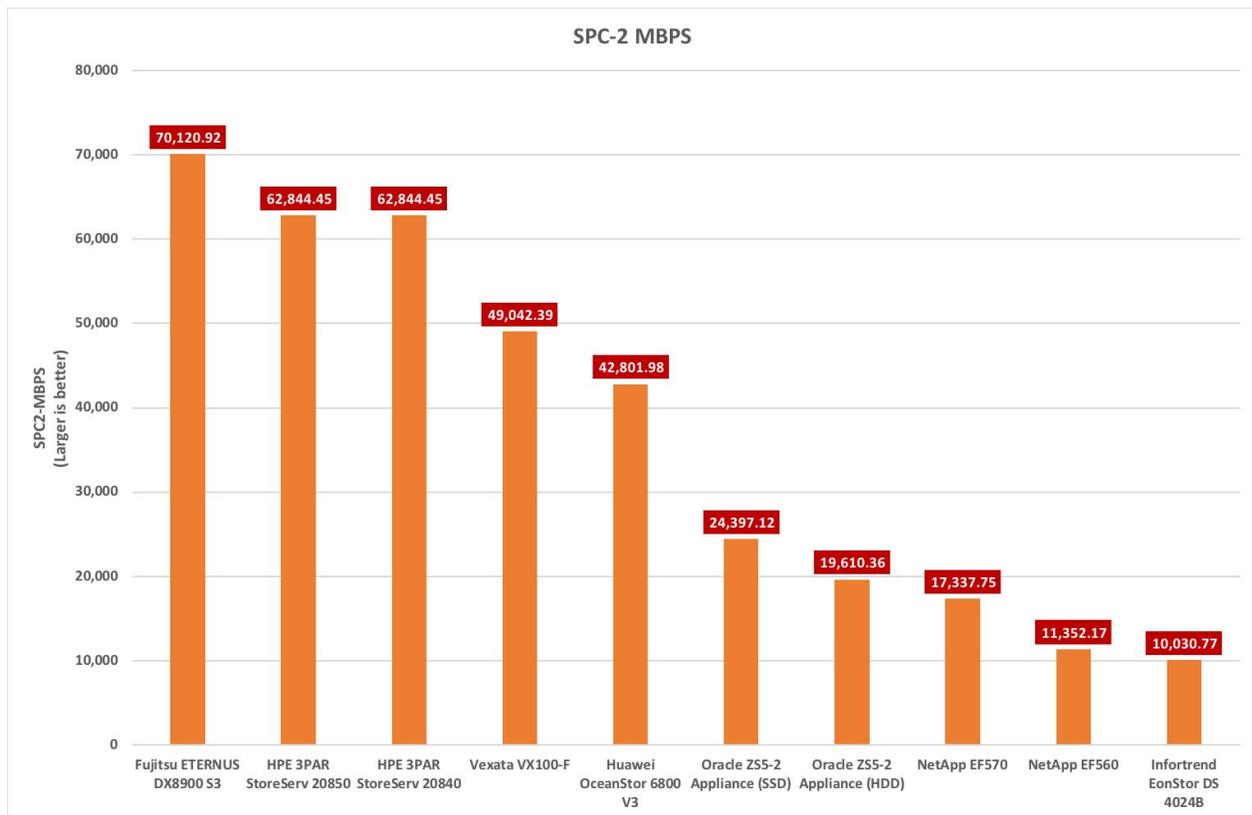
SPC Benchmarks

SPC provides two general benchmark categories. SPC-1 covers predominantly random I/O type workloads, such as those in a typical enterprise data centre running mixed, virtualised applications. SPC-2 covers workloads that are characterised by large-scale sequential I/O or movement of data. This can mean, for example, databases, streaming I/O (like Video on Demand) or other applications where throughput is important. In both cases, the benchmarks have variations for energy efficiency and measuring components rather than systems. Vendors pay to be part of SPC and testing is not a trivial (or cheap) exercise. As a result, we've generally seen the major storage vendors in the market taking part, although not everyone chooses to get involved. Let's dig a bit deeper into one of the tests - SPC-2.

High-Level Analysis

The headline ranking metric for SPC-2 is performance, otherwise described as SPC-2 MBPS. Checking the [latest version of the specifications](#), this value is calculated from the average of three individual tests that make up SPC-2 - LFP (Large File Processing), LDQ (Large Database Query) and VOD (Video on Demand). In turn, these three tests represent the average throughput measured during the audited test run. Looking purely at headline figures, then the Fujitsu ETERNUS DX8900 S3 comes out top with a score of 70,120.92. You can see the most recent test results summarised in figure 1 (data from the last 2 years).

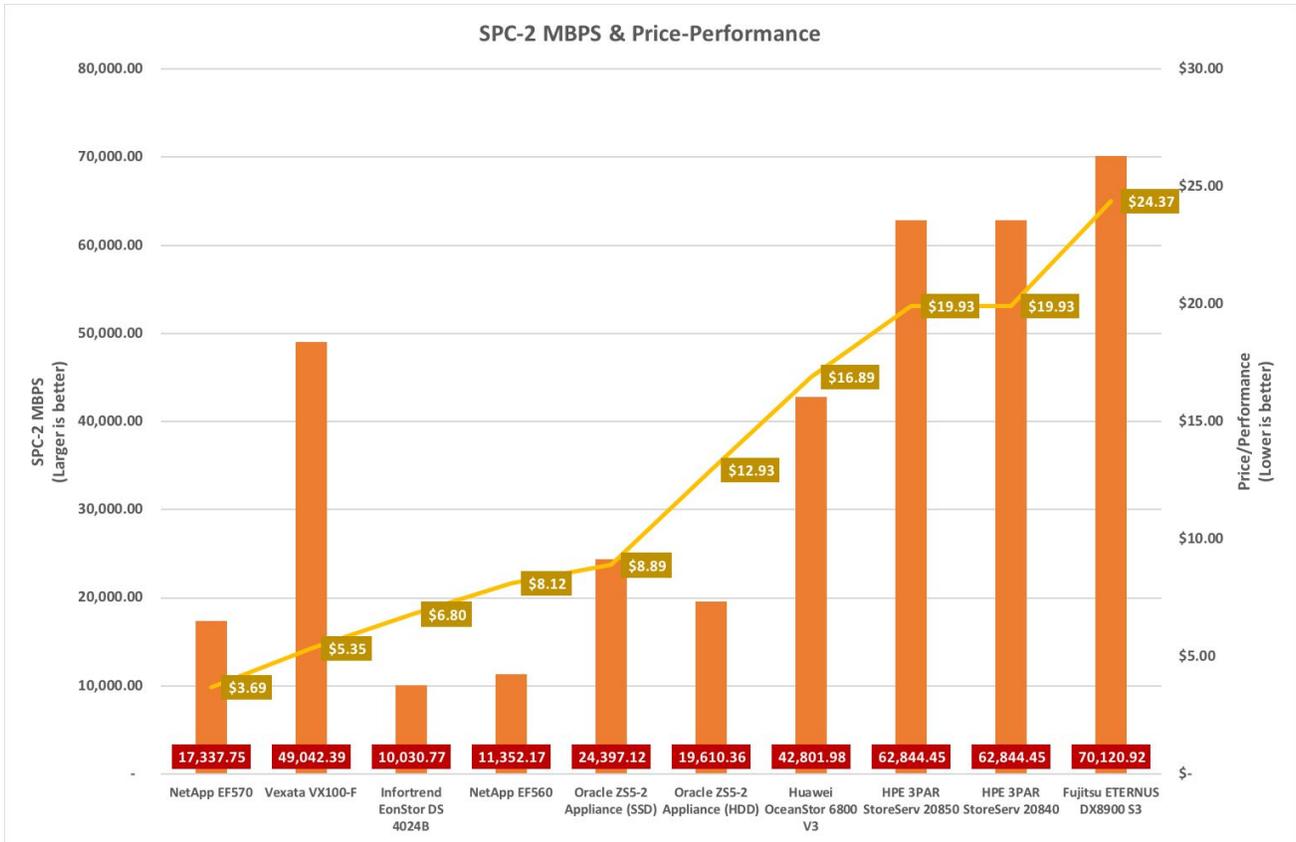
Figure 1 - SPC-2 MBPS



Digging Deeper

This is great, however, it's not the entire story. Looking at the cost of the system tested, the ETERNUS is the most expensive solution put forward (just over \$1.7m) and the worst on price/performance at \$24.37. The metric SPC-2 Price-Performance is officially defined as the ratio of the total system price to SPC-2 MBPS. Using this metric, we see a totally different view. So, if value for money in performance is an issue, then the best solutions that fits into an expected user performance profile of 20GB/s or higher is the Vexata VX100-F. See Figure 2.

Figure 2 - SPC-2 MBPS and Price/Performance

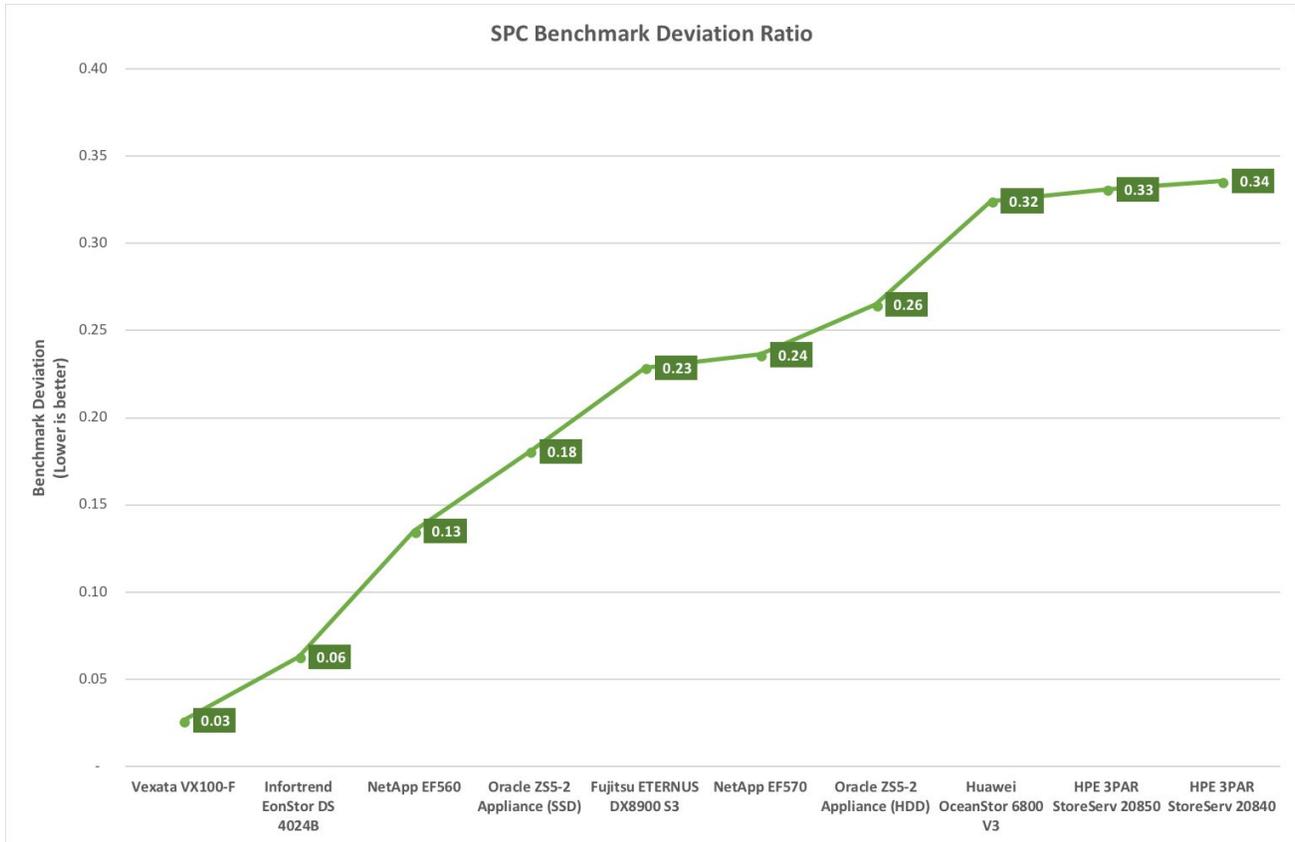


What about the test results themselves? The idea of having three separate tests is to demonstrate three similar but subtly different workloads. However, averaging these numbers out hides some of the inconsistencies between test results. Looking again at the ETERNUS figures and the three individual test results are:

- LFP Composite - 52,589.36
- LDQ Composite - 84,083.42
- VoD Composite - 73,689.99

The LDQ (Large Database Query) figures are 60% better than the LFP (Large File Processing). Clearly, the capability of this system is extremely dependent on the workload mix, making it difficult to predict how it would perform in the real world. One way to look at these figures is to calculate the variability of the data. Look at figure 3. Here we calculate the standard deviation and divide by SPC-2 MBPS (the mean). This gives us a ratio where the smaller the number, then the more consistent the results represent. Using this metric, we see that the Vexata VX100-F comes out on top, with a figure of 0.03.

Figure 3 - Deviation Ratio



Behind the Data

Of course, what we're showing here is data from tests of hardware put up by the vendor. The analysis doesn't show more abstract information, such as the depth of data services that are running on the platform - for example, deduplication and compression. Some vendors turn these features off to improve their results; some can't disable them and have to run them all the time. This can have a direct effect on the results.

What happens in failure scenarios? With enterprise-class storage, one important factor is understanding how a system performs when a component fails. It would be nice to see the result of failing a component like a drive, or controller and seeing how performance changes. Remember dual-controller storage array architectures with replicated DRAM? When those devices had a failure, the performance tanked because data in DRAM wasn't protected by mirroring to the failed controller. We shouldn't have this kind of scenario in 2019.

Then there's scalability. What happens when a solution is expanded? If more storage capacity is added to a system over time, generally the average performance declines, because the controller components become bottlenecked. This kind of behaviour isn't shown in these tests.

Finally, there's a question of future value. We've touched on scalability but tied to that is the ability to exploit an architecture effectively over the lifetime of the deployment. For example, can drives be added asymmetrically (that is, with differing capacities) or does the architecture expect an entire RAID set of drives to be added? Can individual components like controllers be easily upgraded (and in-place)? The reason this point is important is not just in the cost/performance calculation, but in measuring risk. If a system needs all drives to be replaced as part of an upgrade in capacity, then the risk exposure of a failure is significantly increased.

Evaluation Strategies

What's the right approach to picking a storage platform?

1. **Work out your sensitivities.** By this, we mean what's important to you. Is it cost, value for money, long-term scalability or reliability? Rank these in order because they will be important in choosing the right platform.
2. **Do a high-level analysis.** Create a shortlist of vendors using the data available. This is where vendor and SPC benchmarks become invaluable because vendor claims can be easily validated. What happens if a vendor is not on the SPC list? This is the time to ask the vendor why and what additional independent testing they can provide, in as much detail as possible.
3. **Proof of Concept.** The best way to test a system is via PoC. This is the opportunity to really put a system through its paces. Remember though that a vendor will want you to test on as close to production data as possible because if the results are good, you won't want to take the system out. If you can't do a PoC locally, ask your vendor what lab facilities they have to let you do testing.
4. **Choose.** Eventually, a decision has to be made. However, even at this point, it makes sense to agree on performance guarantees in a contract. If things don't work out, fixing the problem could prove very expensive without some kind of fallback or safety net.

The Architect's View

Benchmarks are a great baseline for understanding the performance of a storage system. But without digging into the figures in more detail, it's clear that headline numbers aren't always representative of the entire truth. Investigating the data behind the numbers shows a different set of metrics, ones which may prove more useful as a comparison.

Copyright (c) 2007-2019 Brookend Ltd, first published on <https://blog.architecting.it>, reprinted with permission. Originally posted at <https://blog.architecting.it/performance-benchmarks-reading-between-the-lines/>, and reproduced by Vexata Inc with permission.