What are IOPS and should you care?

Prepared by: George Crump, Lead Analyst
Prepared: February 2015

When evaluating a new storage system, especially an all-flash array, the number of IOPS (Inputs/Outputs per Second) that the storage system can sustain is often used to differentiate one storage system from another. But is this really a standard that has any merit given the demands of today's data center and the capabilities of today's storage systems?

There are three factors that when combined tell the full story of storage performance; bandwidth rate, latency and IOPS. Most storage vendors tend to focus on IOPS to brag about how fast their storage system is. But measuring storage system performance by IOPS only has value if the workloads using that storage system are IOPS demanding.

Transfer Rate vs. IOPS

There are many variables to consider when trying to determine the overall performance of a storage system. There are external factors like how the data is being read from or written to the storage system and the speed of the storage network fabric itself. There are also internal considerations like the CPU power of the storage compute engine (the storage controller), efficiency of the storage software and, of course, speed of the storage media installed in the storage system.

For the purposes of this article, we'll assume that all the external factors are equal. If that is the case then transfer rate is essentially the speed at which the storage controller can move a contiguous data block through the storage software to the storage media. It is typically measured in MB/s and a high transfer rate is important, especially for workloads that are sequential in nature.

IOPS however are different; they are measured as an integer number. It refers to the maximum number of reads and writes to non-contiguous storage locations. These operations are typically dominated by seek time, or the time it takes a disk drive to position its read/write heads to the correct location. Because this positioning of heads is so time consuming, the importance of storage controller CPU power and the efficiency of storage software are greatly minimized in a hard disk array. Flash arrays virtually eliminate seek time from consideration and as such they make the other variables like power of the storage controller and efficiency of the storage software far more important. The storage controller and storage software can no longer hide behind the bad performance of the hard drive. Flash exposes them for what they are.

By way of example let's compare how two workloads accessing the same amount of data require a significantly different amount of IOPS. The first workload requires reading ten 750MB files, 7.5GB and it takes 100 seconds for the transfer to occur. This means that the transfer rate is 75MB/s and consumes 10 IOPS, which is well within the capabilities of a single hard disk. The second workload requires reading ten thousand 750KB files, the same amount of data, 7.5GB, but it consumes 10,000 IOPS. Since the typical disk drive can't generate more than 200 IOPS this request won't get done in the same 100 seconds. This is an example of how different workloads can require significantly different performance, while using the same storage capacity.

Do IOPS Matter? - Simple Answer, No

With the definition of IOPS out of way, the next question is should an IT Professional be concerned about the potential IOPS performance of a storage system? IOPS were a far more important measurement in the hard disk array era because the potential number of IOPS was often less than what the data center needed. In the all-flash array era the opposite is true. Most all-flash arrays will deliver far more IOPS performance than the most data centers will need.

IOPS Measurements Can't Be Trusted

The other problem with using IOPS as a way to differentiate between flash storage systems is that there are too many ways to generate an IOPS number, as our illustration above indicates. IOPS can be significantly impacted by the size of the block used, the mix of read/write activity and amount of

randomness in that I/O stream. Even if vendors all standardized on how each of these variables were to be set, it would have little relevance to the data center. For example, if all vendors state they would report IOPS from tests using 4k block sizes and 50% random read/write mix, the resulting number would have little meaning to a data center whose workloads were generating 32k blocks with an 80% read to write ratio. Finally, most data centers are going to have multiple workloads running on their all-flash array. It will likely support a wide variety of workloads with varying read/write mixes.

The Right Measurement

The right way to measure the performance of an all-flash array or even a hybrid array is to develop performance statistics based on particular workloads or a mix of workloads. For example run a SQL performance test and a VDI performance test at the same time on the same storage system and instead of reporting on IOPS consumed, report on data that is more tangible and relevant to the data center. In this case it might be the number of simultaneous SQL users and VDI instances supported while still maintaining acceptable response times.

Upgradeable Performance

As stated above most all-flash arrays will deliver more performance than most data centers can take advantage of today. But today is the operative word here. As virtual server and virtual desktop density as well as users per database instance all continue to scale, data centers will need more and more performance. The flash media itself will become slightly faster but the key roadblock to expanding performance will be the storage controller and the efficiency of the storage software.

The features and capabilities of the storage software will add overhead to flash performance. The efficiency of that storage software in how it executes these various capabilities is critical to overall performance. Fortunately the all-flash vendor has access to ever increasing compute power that can mask most of the overhead of the storage software. It is critical though that the flash vendor be able to provide an upgrade path for their controller hardware so that their customers are able to take advantage of the increased power of each iteration of Intel CPUs.

Conclusion

Using IOPS as a way to differentiate between flash arrays is a risky practice. Most systems provide more IOPS than the typical data center needs. These data centers can better spend their time by looking for a flash array that provides the features they need at a price they can afford, as well as an upgrade path to continue to stay ahead of performance demand.

For those data centers that do need high enough performance that IOPS could be relevant, the vendor provided IOPS numbers have too much variability between them to provide a meaningful differentiation between them. These data centers are better off requesting specific results based on a mixture of workloads that closely match their environment.

Sponsored By Pure Storage

Pure Storage is an example of a company that understands the relative importance of IOPS and has instead shifted their performance reporting to more workload specific use cases. They provide a straightforward scale-up design but one that is upgradable. As a result, as faster storage controllers become available customers will see faster overall performance when they upgrade their controllers.