

The Value of an Independent Storage Performance Testing Platform



by George Crump, Lead Analyst

The more virtual machines per host, or users per application, the more cost effective an environment becomes. And the more environments that can be supported by a single storage system the more cost effective the storage infrastructure becomes. But increased density, scale and workload variety places greater performance demands on the storage system. A poor selection might mean that the cost effectiveness gained by increased density and scale will be consumed by an inefficient storage solution. Since IT planners seldom have time to test every solution, they have to count on vendor provided performance data to create their short list. The IT planner can only trust this data if the vendor is using an independent storage performance testing platform.

What is an Independent Storage Performance Testing Platform

When developing a short list of vendors the key is to look for vendors that used an independent testing platform in order to generate their performance data. An independent testing platform allows the IT evaluator to recreate the test criteria in the IT organization's lab or in a vendor's POC center. This requires a test that can be universally applied against off-the-shelf storage configurations. Ideally it should be one that is based on the organization's production I/O profile across a variety of application workloads – one that, in effect, replays those I/O profiles on the storage systems to be evaluated. This would allow the I/O generation engine to be simple, with an easily repeatable testing process. The only component that should change would be each vendor's storage system. Each of these storage systems should be configured in the proposed configuration that the vendor recommends to see each offering evaluated under its "production deployment guidelines".

Benchmarks Fall Short of Independence

When the word gets out that an organization is considering a performance focused storage system, vendors come from every direction with a variety of claims. To support these claims vendors will use a wide variety of benchmarking tools, often proprietary, to generate performance data.

Clear understanding of how your current performance-challenged production applications will work on potential new storage systems has always been a challenge. Vendors would make claims, typically based on using freeware tools or benchmarks, using unrealistic, best-case configurations that would never produce the same results once the storage system was put into production. This vendor supplied performance data could never be reproduced in the test lab of the IT organization and certainly never used to compare one vendor to another as the test set-up was never truly apples to apples.

The Benchmark Problem

The challenge with counting on vendor provided benchmarks is the variety of benchmarks that they can choose from in order to claim their system is the fastest. Sophisticated benchmarks like the SPC-1 and SPC-2 are largely dependent on the knowledge of the person executing the test and their ability to fine tune the test environment (servers and storage) to that test. Creating a consistent test environment and storage configuration across a variety of systems can be a challenge. The problem in using these types of benchmarks is compounded by the fact that some vendors refuse to support or test their systems against certain benchmarks.

Another challenge to these benchmark types of tests is they are very specific to a particular workload, often high transaction databases. Most modern performance focused systems will be better cost justified by their ability to support a wide variety of workloads, commonly found in virtualized environments.

The I/O Tester Problem

There is another type of performance data creating tool, known as I/O testers, which are made up of products like IOmeter, vdBench and Linux flexible I/O tester (fio). While somewhat simpler to execute, they provide little more than raw performance numbers with little translation to the real world. Also as we discuss in our article [“You can do better than ioMeter and Vdbench”](#) these solutions are often freeware and are subject to the limitations of those types of solutions. Another challenge with these solutions is easily creating a mixed workload benchmark that has any real world meaning for your data center.

Workload Specific Testing

There are a variety of workload specific tests available on the market, often put out by the supplier of the environment or application. Microsoft, VMware and Oracle as an example all provide various performance testing tools. Some of these tools are glorified I/O testers that are tuned for a specific workload but others do a very good job of simulating a real world environment. The challenge with these tests is that the more real world they become the more complex they become to set up. These solutions typically require extensive test labs with several servers and personnel dedicated to the performance test. They also are very specific to one application, so they are not very useful in mixed application workload or highly virtualized environments.

Like the benchmark tests some vendors will invest in these tests, others will not, so not all vendors can give you a number for the test you are most interested in. Finally, even more so than the benchmark tests they also can be impacted greatly by how well the storage system is configured for the specific task and are very workload specific. The only way to truly test a mixed workload environment would be to run the tools simultaneously on multiple test beds all connected to the same storage. This would obviously increase testing complexity and cost significantly and still create repeatability challenges across different vendors.

Establishing a Testing Platform

The key to making a wise decision based on performance benchmarks is to look for vendors that are using an industry recognized independent testing platform. It should be one whose setup and results are repeatable and constant. These solutions require a single server appliance that can generate a mixed and random I/O workload to very exacting specifications



that truly reflects your installed production workload I/O profiles. It's not a benchmark, but a production workload modeling and load generation appliance. It can also scale to massive levels to validate the largest flash and hybrid storage configurations.

A good example of this type of platform is the storage performance-testing platform from Virtual Instruments. It meets all of the above criteria but has an interesting advantage over its competitors; most storage vendors use WorkloadWisdom, formerly known as Load DynamiX, for their own internal performance validation of their storage systems. The good news for IT is that the test criteria and load assumptions used to create the performance data is easily transferred and the customer can then use the testing platform to perform the exact same test on the other systems on their short list. The tests and workloads can be easily re-run at any vendor's site or within your own testing lab.

More importantly, if the evaluator chooses, IT can bring the testing platform in-house on a permanent basis so that they could test future potential purchases. As we describe in our article ["Are You Planning for Storage Performance"](#) IT could also use the testing platform to test the impact of future application workloads on the current storage system. This would allow IT to know in advance if there will be any future performance problems and exactly when those performance problems may appear.

Conclusion

When reviewing vendor solutions, each will have numbers to support their performance claims, but if those claims cannot be independently repeated and compared in an apples to apples comparison then they have little value. IT planners should look for vendors that are using an independent performance testing platform that enables realistic workload profiles and provides the ability for testing criteria to be repeated consistently across a wide variety of storage systems.