



Analytico, Inc.
Analysis from all Angles™

Computing In The Clouds?



***Standardized Performance Tests Enable Users
to Discern Extremely Fast Storage Systems***

*Tom Trainer
President
Analytico, Inc.*

I. A Need For Standardized I/O Performance Tests

Storage system performance has long been a topic filled with broad vendor marketing claims. Performance claims by storage vendors have resulted in confusion, and at times great disappointment, on the part of application users, storage managers, and purchasing departments. Many users have realized too late that the storage system they purchased cannot live up to the claims made in the sales process. The negative results can include missed research opportunities; delayed projects; unhappy executive management; and negative impact to business operations ranging from increased costs to lost revenue.

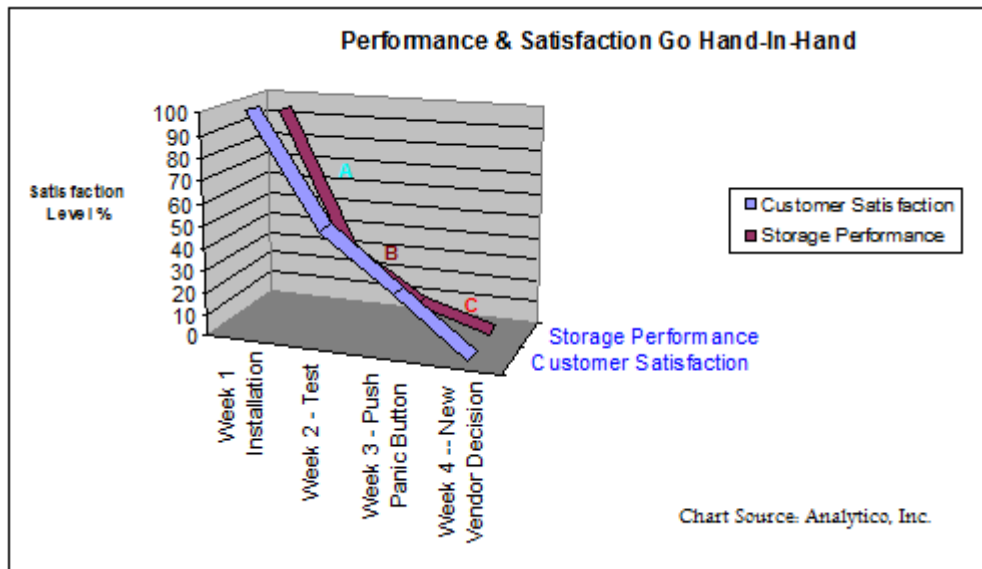


Figure 1: Performance & Satisfaction

Analytico has spoken to 10 users of high performance storage equipment. Each of these users had negative performance and customer satisfaction experiences in the past related to the purchase of storage systems. These experiences have typically resulted from a belief in marketing claims of stellar storage system performance—only to install, test, and realize that actual system performance paled in comparison to pre-sales claims.

Figure 1, above, illustrates the average user-vendor experience when storage performance does not meet user expectations. Plane “A” is representative of the first weeks of test time when results are not as expected. Plane “B” represents user faith and loyalty, even as performance still does not meet expectations. Place C, unfortunately, is “the slope of disappointment”™ and the point at which a user must make a rapid decision to replace the first vendor choice with an alternative. The weeks represented in the chart are an average indication of time.



High performance users and too many disappointed I.T. managers—who must build rapid-fire I/O infrastructures—have obviated the need for standardized benchmark testing in order to discern credibility from claims. Recently, a team of engineers from an internationally recognized supercomputing facility conducted benchmark testing of two well known storage vendors—who both supplied the site with what each term as storage designed for the supercomputing, or high performance computing, environment.

The benchmark tests, developed by established supercomputer engineers, are a standard set of tests used to measure the throughput of storage systems in extreme sequential read, sequential write and random write I/O conditions. In essence, these tests emulate the real world I/O conditions found in massively scalable cluster computing and grid environments. The workloads are characterized by the need for very high throughput levels to disk driven by large file sizes. These types of data patterns, considered unimportant just a few years ago, are becoming more common as the world shifts to creating and storing unstructured data, such as large media files, corporate documents, computer simulation files, photos, video, and Internet content. Unstructured data is driving the growth in demand for storage, leading storage systems vendors to position themselves at the forefront of serving extreme amounts of this type of information.

Analytico has examined these tests and can confirm that they leverage the compute environment under test, and exhaustively exercise the storage system under test. These developed tests move the high performance and cloud processing user much closer to being able to discern the short list of storage vendors who should be considered for testing and adding to their infrastructure.

II. Two Vendors Undergo Testing

The two systems under test included the DataDirect Networks S2A9900 and the LSI XBB-2. Data Direct Networks S2A9900 is powered by the Silicon Storage Architecture, an architecture that the company has evolved over 10 years of high performance storage development. The LSI XBB-2 storage subsystem, based on the XBB architecture within all mid-tier storage systems designed by LSI, is a mid-tier product initially shipping with Fibre Channel and InfiniBand front-end interfaces. The tested XBB-2 storage subsystem has been shipping for less than a year.

In the benchmark test, both vendors shared common functional attributes:

- Systems tested with:
 - Parity checking on reads
 - Cache failure protection enabled
 - RAID 6 (8+2) SATA drive configurations

Given the system capacity limitation of the LSI XBB-2 system, there was a difference in the number of drives configured between systems. In simple terms, this means that two LSI storage systems are required to test against one DataDirect Networks system. Given the additional cost



factor associated with two LSI systems, the testing was conducted with one system from each vendor. The storage configuration was as follows:

- LSI XBB-2; 24 RAID sets, 8+2 RAID 6 (240 drives)
- DataDirect Networks S2A9900; 28 RAID sets, 8+2 RAID 6 (280 drives)

Test Results

Why Throughput Matters In Storage Performance

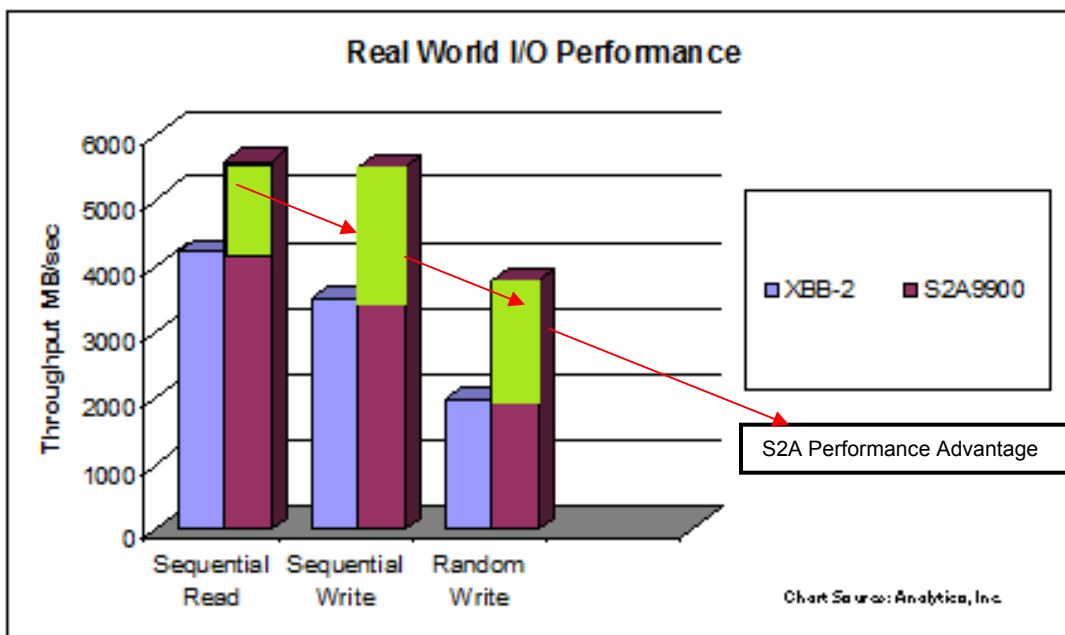


Figure 2: Throughput of LSI and DataDirect Networks Storage Systems

Figure 1, above, illustrates the results of benchmark testing for sequential read, sequential write and random write performance. In the test case of sequential read throughput, the DataDirect Networks S2A9900 performed at a 33% greater throughput vs. the LSI XBB-2. In the test case of sequential write performance, the DataDirect Networks S2A9900 performed at a 57% greater throughput vs. the LSI XBB-2. In the test case of random write throughput, the DataDirect Networks S2A9900 performed at 91% greater throughput vs. the LSIXBB-2. Taken together, the read/write performance benchmark results of the DataDirect Networks S2A9900 provided significantly greater overall throughput than the LSI XBB-2.



Note that there are technological differences between the XBB-2 and the S2A9900 architecture. The XBB-2 is designed for general-purpose computing environments (i.e. wide variety of commercial servers and applications). General purpose storage systems tend to display performance trail off at greater I/O rates due the nature of how they handle read and write I/O activity, check data integrity, and write data in a RAID 6 configuration. These systems tend to buffer large segments of I/O and this, in extreme I/O environments, results in reaching a throughput ceiling, as illustrated in figure 2, above.

Purpose-built storage systems are architected for environments that use the systems for specialized I/O requirements. With regard to the S2A9900, its purpose-built architecture leverages a patented DirectRAID technology which enables accelerated I/O performance. Within the S2A9900, all data is handled by dedicated, high-speed hardware instead of the more typical general purpose CPUs. This ensures that the system is not only fast, but also extremely predictable in its performance. The DirectRAID technology consists of a hardware-accelerated RAID engine that calculates Reed-Solomon error correcting code in real-time, combined with intelligent algorithms leveraging a high-speed parallel data path within the system. The high-speed parallel data path enables real-time reading and writing and on-the-fly RAID parity calculation and recording. The S2A9900 delivers 6GB/s of host throughput for both reads *and* writes; no other storage system in the industry can write as fast as it reads—let alone approach the performance levels of the S2A9900. Figure 3, below, illustrates the unique DirectRAID technology deployed within the S2A9900.

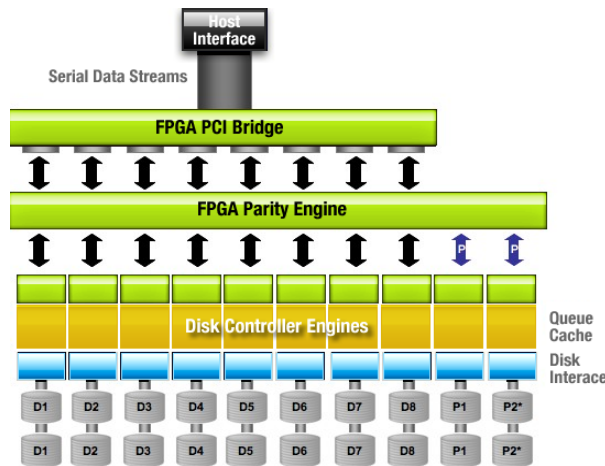


Figure 3: S2A9900 DirectRAID Engine & Parallel I/O Paths

III. Analytico’s Perspective

The numbers in Figure 1 tell a compelling story about performance benchmark capabilities and that of the two systems under test: the DataDirect Networks S2A9900 and the LSI XBB-2. Firstly, the fact that a set of quality benchmarks now exist and can be trusted in determining which vendors to put on the short list for building into the high performance infrastructure.



While these benchmark tests did not include testing under failure conditions, any storage administrator can tell you that if a disk drive fails you should expect severely degraded performance from your storage system as it allocates resources to rebuilding the failed drive, and fights contention issues communicating with the remaining healthy drives. This problem is typically exacerbated when using RAID 6 where the parity calculations are substantially more complex. The performance hit can be 50% or higher and can last for hours or days depending on the severity of the failure.

Analytico has seen a truly unique function of the DataDirect Networks' S2A design in that it completely avoids this problem, which means that the high levels of performance the system delivers can be counted on under drive failure conditions, not just when the system is in a completely healthy state. For supercomputing environments that can have tens of thousands of disk drives, this predictable, consistent performance is critical. While the independent tests conducted by the supercomputing site did not publish degraded performance numbers, Analytico believes that LSI's performance would diminish substantially given the RAID 6 implementation on that subsystem.

So, why is greater I/O throughput and benchmark testing a significant advantage for supercomputer, or high performance, application environments before production implementation? The answer is simple in nature; however, it varies depending upon the application environment.

Pharmacological research science environments complete complex calculations faster, where scratch pad memory is contained on external storage. Faster calculation completion enables scientific researchers to complete projects sooner, discover new—or improved—chemical formulations quicker, and produce results more rapidly.

Internet services providers take advantage of faster throughput by providing faster response time to users who use Slide applications on Facebook, an online social community, and XBOX LIVE for interactive game play on a real-time basis.

Faster throughput has positive effects in the video and CGI rendering environments, enabling news organizations to prep near real-time news stories for immediate broadcast. CGI environments increase production output, which results in getting movies to market faster and providing a competitive edge in release dates and increased revenue generation. It also helps video-on-demand providers serve more video streams concurrently.

Even enterprise environments benefit from faster throughput. Think about doing nightly backups to disk. These involve a series of high-throughput backup streams being written to the array at once. Having extreme write performance means more data can be backed up per unit time, helping organizations meet their backup windows. With performance like that seen in Figure 2, we can see that the S2A architecture would be ideal for disk-to-disk or Virtual Tape Libraries, as its write performance far exceeds anything witnessed on the market by Analytico.

Storage systems that deliver extremely fast throughput enable users to avoid direct costs by having to purchase 2X, or more, the amount of systems required to achieve desired throughput, and then dealing with how to aggregate the throughput through a common file system. Coupling the



original purchase cost with extended costs to power and cool additional hardware, it is easy to see the negative total cost of ownership (TCO) impact of slower storage systems.

IV. Summary

The old adage “Time is Money” is especially true in today’s high performance computing environments. Critical path processing environments, such as the number one ISP provider in the United States, AOL, simply cannot afford to take the financial risk of investing in poor performing data infrastructure products. Major motion picture companies, post production facilities and broadcasters understand the need for extreme storage performance and have deployed DataDirect Networks storage in their most demanding performance environments.

Analytico believes that over 10 years of experience in design and delivery of high performance storage systems has evolved into an ability by DataDirect Networks to deliver products that go beyond meeting the standard high performance requirements—they *exceed* extreme performance requirements. The performance testing discussed in this research report resulted in the customer choosing DataDirect Networks as their storage provider of choice and the facility is now working with the company to implement *the fastest file system in the world*.

With high performance storage platforms in leading companies such as the number one oil producer in the world, Saudi Aramco; to Lawrence Livermore Lab, which attains more than 100GB/s of usable bandwidth across three petabytes of scalable capacity; to China Central Television, which is using S2A platforms as the primary data infrastructure to air the Beijing Olympic games to over 1 billion people, DataDirect Networks maintains its leadership position in extreme storage environments with more major high performance systems installed than any vendor in this storage category.

Users should look to experienced storage providers who bring proven integrity and performance assurance to their environments; these storage vendors who submit to benchmark testing deserve respect as they strive to understand their products that much more and develop better products for the industry at large.

The Bottom Line: Standard benchmark suites now exist for high performance computing environments. Respected vendors are submitting to benchmark testing so that users can hit the ground running with the right storage system for their particular environment.

*For detailed specifics associated with this benchmark contact tom@analyticoinc.com
S2A9900 and DataDirect Networks are copyright and trade marks of DataDirect Networks
LSI XBB-2 is a model and trademark of LSI, Corporation.
“Slope of disappointment”™ is a trademark of Analytico, Incorporated, and is associated with Mt. Rainer.*

