



ISE 820 All Flash Array Performance Review

March 2015

Performance Review

March 2015

Table of Contents

Executive Summary	3
SPC-1 Benchmark Results.....	3
Virtual Desktop Benchmark Testing.....	3
Synthetic Workload Testing.....	4
The Storage Performance Council SPC-1 Benchmark	5
VMware View Planner 3.5 Benchmark Testing	6
Why is Group B so important?.....	6
Synthetic Workload Testing	8
Summary	10
Contact X-IO Technologies.....	10
Appendix A – Informational Links	11

Executive Summary

Storage performance requirements have been growing at an ever increasing rate, just as the demand for storage capacity did 10 years ago. Businesses are implementing new computing solutions that give their employees and customers better and faster access to information than ever before. While some of these solutions are not necessarily new, the speed at which the information must be accessed has increased tremendously. The ISE 700 Series Hybrid Storage Arrays have been utilizing flash technology for over four years, giving X-IO's customers the unmatched benefits of:

- High Reliability proven with a standard ISE hardware included for free with every system.
- High Performance with 60,000 to 90,000 IOPS proven in benchmarks, solution designs, and customer environments.
- High Capacity with virtually no performance degradation up to 100% capacity utilization

There are workloads however, that require more performance than even an ISE hybrid array can offer. For these workloads, "hundreds of thousands" of IOPS are the norm rather than the exception. The ISE 800 Series G3 All Flash Array was designed from the ground up to unlock the performance of flash devices, while integrating with the ISE Management infrastructure. Performance profiling for the ISE 800 Series is critical for X-IO partners and customers, and this paper will provide a brief summary of three performance evaluation tests that were run as part of the ISE 800 series development.

SPC-1 Benchmark Results

Databases are a hallmark of ever increasing performance requirements for data storage technologies. Getting access to data such as customer records, inventory levels, and sales pipelines gives business ever increasing visibility into how their organization functions. If the information can be mined effectively, data can be a transformational business asset. The Storage Performance Council's SPC-1™ benchmark was designed to give manufacturers and end users an "apples-to-apples" comparison for these transactional intensive solutions. With strict performance requirements and 3rd party auditing of submissions, the SPC-1 workload is the "Gold Standard" for independent, audited performance testing.

Virtual Desktop Benchmark Testing

The ISE 700 Series Hybrid Storage Arrays have proven that combining flash and HDD devices are an excellent method for servicing high numbers of virtual desktops. However, with the ever increasing server horsepower (CPU and RAM), wide adoption of GPU acceleration cards, and features such as VMware Linked Clones, desktop consolidation ratios are increasing and the functions they support are expanding. These desktops support operations such as high density "Kiosk Mode" desktops, AutoCAD workstations, High Definition Video editing, and DBA development workstations where concentrating these unique workloads together requires tremendous amounts of performance from all the data. User experience for these use cases is absolutely critical, and storage performance is usually the #1 cause of poor user experience. In addition to Login VSI, VMware View Planner is used by X-IO to evaluate the performance of ISEs in a sample Virtual Desktop configuration. This section will review the testing and latest "Group B" results for the ISE 820 G2 All Flash Array.

Synthetic Workload Testing

No two workloads are the same, just as no two environments are the same. No two SQL Server databases are ever the same, just as no two businesses do the same thing with their data. While synthetic load tools offer great degrees of flexibility in recreating a specific workload, performance characteristics change throughout the production work cycle. Testing with these tools to mimic a common function (OLTP – 25% Read, 50% Sequential, 4K) ignores performance demands of other database operations (TempDB, Logs, Indexing, etc). One method X-IO uses to evaluate the performance of the ISE systems is to run large numbers of tests varying multiple workload characteristics. This method allows for examination of over 1,000 specific tests, and allows X-IO to give our partners and customers guidance as to the performance potential of ISE systems.

The Storage Performance Council SPC-1 Benchmark

The Storage Performance Council is a non-profit organization that was founded in 1998 with the goal of establishing the first industry standard storage performance benchmark, and disseminating objective storage performance information to end users. SPC establishes a “level playing field” for comparing the performance between different storage vendors, while enforcing strict run and reporting rules to ensure the integrity of the benchmark results. With over 30 storage companies as active members of the SPC, the SPC-1 and SPC-2 benchmark workloads are considered the “gold standard” for audited 3rd party performance results.

The SPC-1 workload is designed to mimic the workload of OLTP/database applications, characterized by predominately random I/O operations, requiring both queries as well as update operations. While one of the main measures of this test is overall transactional performance (IOPS), submissions also have to include a cost of the solution as part of the audited documents. This allows end users to evaluate submissions not only on the performance the equipment was able to achieve, but also the cost of each transaction at that level. This comparison is expressed in terms of \$/IOPS, and is an excellent method for comparing solutions from different vendors.



Figure 1 - SPC-1 Top 10

X-IO has recently completed the audited performance testing process, and are proud to announce the new record setting value of just **\$0.32/IOPS**. This breaks the previous record by 20%, and represents one of the largest decreases in cost over the last several years. Furthermore, all other solutions in the top 10 **only include 3-yr warranties**, while all ISE systems come with a standard, no-cost, 5-year hardware warranty. Storage vendors typically charge ~20% of the original purchase price for years 4 and 5 warranties (in each year!), which would increase all of the other submission costs by ~40%. This cost savings through high reliability is a hallmark of the value that X-IO customers have been experiencing since the original ISE systems were launched (2008).

VMware View Planner 3.5 Benchmark Testing

The spikey nature of the VDI workload can vary throughout the day. Opening Microsoft Office applications, sorting in an Excel spreadsheet, running a PowerPoint slide show are all common tasks an end user will execute that generates I/O on the storage array. If most of the workforce comes in at 8:00am and starts opening applications, that daily storm of I/O becomes one more thing an engineer or architect must take into consideration before purchasing a storage array for a VDI solution. Overnight maintenance tasks (patches, recomposes) introduce heavy write bursts. On top of that, boot storms and anti-virus scans introduce huge read spikes. The daily VDI workload may average somewhere around 25 IOPS per desktop, but when the desktop boots it can be over 150. How do you know the storage array you are considering will handle the workload?

VMware View Planner 3.5 is designed give the engineers and architects of the world a level of assurance that their VDI infrastructure will work as designed. It will simulate a large-scale deployment of virtualized desktop systems and study its effects on an entire virtualized infrastructure. By running a set of application operations selected to be representative of real-world user applications, View Planner will report data on the latencies of those operations. The View Planner Quality of Service (QoS) methodology splits these operations into two main groups. Group A largely impacts the ESXi hosts infrastructure, Group B activities generate I/O that impacts the storage array where the desktops reside.

Group A	Interactive/fast-running operations that are CPU bound, like browsing through a PDF file, modifying a Word document, etc.
Group B	Long-running slow operations that are I/O bound, like opening a large document, saving a PowerPoint file etc.

Why is Group B so important?

View Planner 3.5 can be run as a benchmark that helps VMware partners and customers precisely characterize and compare both the software and hardware solutions in their VDI environments. It can be used to benchmark different CPU architectures, hosts, and storage architectures. Most VDI implementations fail due to storage solutions that cannot meet the demanding workloads of virtual desktops. The storage system can be benchmarked to see how many VMs can be concurrently supported without seeing a significant increase in the I/O latency for a given storage configuration. Measuring the quality of service in Group B directly translates to the impact on the end user experience.

The View Planner score is represented as VDI mark. This metric encapsulates the number of VDI users that can be run on a given system concurrently with application response times less than the set threshold. The score is determined separately for Group A user operations and Group B user operations by calculating the 95th percentile latency of all the operations in a group. The default thresholds are 1.0 seconds for Group A and 6.0 seconds for Group B. The scoring is based on several factors such as the response time of the operations, compliance of the setup and configurations, and so on.

View Planner has several different modes that benchmarks can run in, each having a different ratio of Target to Desktop virtual machines. The "Local Mode" option requires no extra VMs beyond the desktops being tested against, and increases the amount of server CPU and RAM that can be applied to the desktops under test. This allowed for the highest active VM density in the X-IO UCS based X-POD solution for the test. With 16x Cisco UCS M3 blade servers (320x CPU Cores and 3TB of RAM), the CPU

and RAM resources were exhausted at 1,000 desktops. The ISE 820 G3 All Flash Array, however, was able to easily accommodate the workload. Operations to setup and deploy the desktops required more performance than the View Planner workload itself, with the UCS blade servers being the main limiting factor to achieving higher workloads.

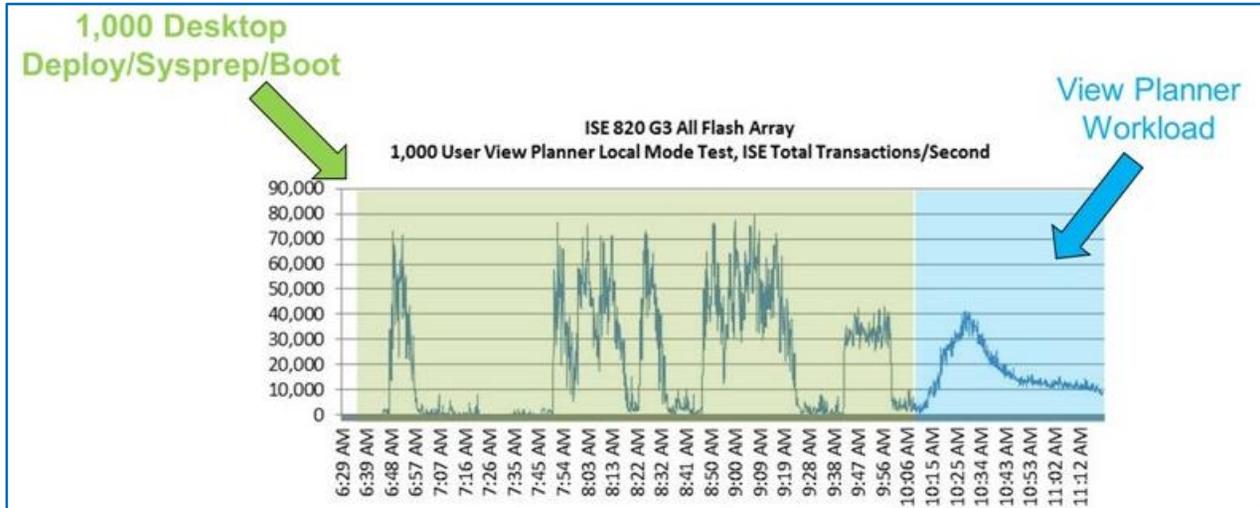


Figure 2 - View Planner Workload

Latency values for the entire run, from setup to the end of the View Planner workload, were under 3ms for 95% of the observed readings. This stress test shows that even with 16x servers (all running at saturation), the ISE 820 All Flash Array was easily able to accommodate not only the test operations, but the setup functions which were even more strenuous than the actual View Planner workload.

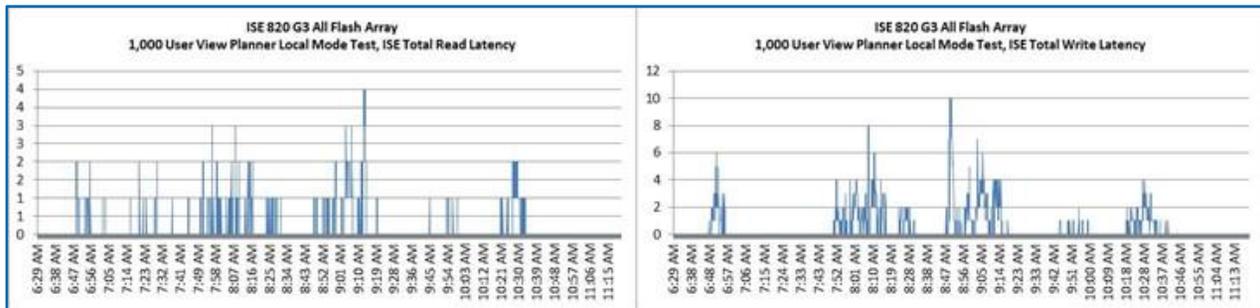


Figure 3 - View Planner Test Latency

While Group B scores tested with local mode are not suitable for submission to VMware as an “official benchmark”, local mode can be effective when the testing for the max storage IO that can be generated in a configuration. VMware reviews the results and certifies the scores as part of the Rapid Desktop Program. As there are many metrics (and benchmarks) that can be reviewed, some can be far more important than others. Group B is one of them.

Synthetic Workload Testing

Synthetic workload tools can tell a great deal about the response of storage systems with various workloads. These tools are used regularly in the ISE development lab to gauge the effect of firmware changes to the ISE system performance. The challenge with these tools is to generate something that is equivalent to workloads that are seen in production environments. While generalities can be made about any workload (OLTP – 25% Read, 50% Random), no workload stays constantly at these levels in production. Furthermore, applications often have several different IO streams with unique workload characteristics (Log Files vs TempDB vs DB).

To better understand the response of the ISE 800 series, X-IO has developed a “Spread Spectrum” testing methodology that performs over 1,000 iterations with various Read/Write, Random/Sequential, Block Size, and Queue Depths. This iterative testing technique allows X-IO engineering to examine the ISE performance at thousands of levels, not simply a single value as most performance testing provides.

Read %	Sequential %	Block Size	Queue Depth
100%	100%	4K	1
75%	75%	8K	2
50%	50%	16K	4
25%	25%	32K	8
0%	0%	64K	16
			32
			64
			128

Figure 4 - Spread Spectrum

Below is the graph of the total transactions per second the ISE 820 was servicing in the 100% Read workload iteration. High values approached **450,000 IOPS** (4K block size test) and throughput values of over **5GB/sec** (64K block size test). All Flash Arrays often have the problem of decreasing performance as the system is used due to “maintenance operations that SSDs must perform (Garbage Collection). Because of this, it is important to not only test the performance of the array, but the performance over time. Each major test series spans approximately 8 hours, with a total run time of 40 hours. No significant differences were seen in the individual test cases beyond what would be expected for the changing workloads, showing that the ISE eliminates this penalty that most other systems experience.

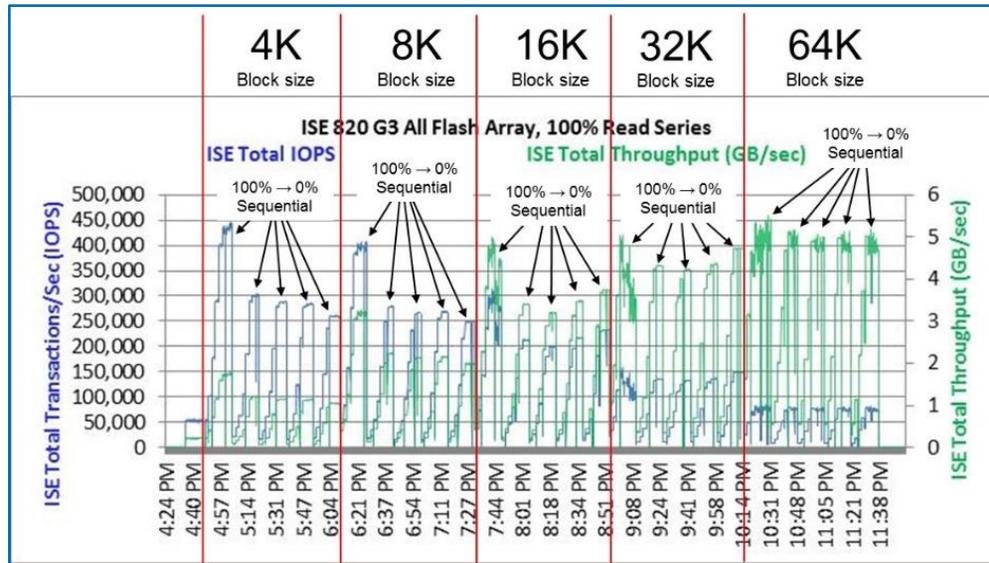


Figure 5 - Performance Test Profile

ISE systems have always been designed to provide the same performance regardless of how full the array is, and the ISE 820 All Flash Array is no different. These tests were performed over all of the system capacity (100% full) to ensure that X-IO customers and partners have accurate information for architecting these high performance solutions. For more information about the results of this testing, please contact us at the address at the end of this report.

Summary

Performance demands are different for every environment, and no single test can determine how a storage array will perform. X-IO executes extensive performance testing at the corporate headquarters in Colorado Springs, CO, giving our customers and partners the information they need to design solutions that have consistently high performance, high capacity utilization, and superior total cost of ownership.

The ISE 820 G3 All Flash Array is the next evolution in the ISE product line, extending the ISE value to the most demanding workloads in today's datacenters.

Contact X-IO Technologies

Website: <http://www.x-io.com>

Get in touch with us:

<http://x-io.com/contact/>

or

Visit our website and **chat** with us to get more information

Email: fastforever@XIOstorage.com



Talk with Us

United States »

866.472.6764

International »

+1.719.388.5500

Appendix A – Informational Links

- ISE System General Information
<http://xiostorage.com/products/ise-storage-systems/>
- ISE 800 Series G3 Hybrid Storage Array info
<http://xiostorage.com/ise-800-series/>
- ISE Performance Adapter for Windows Performance Monitor
<http://xiostorage.com/products/ise-manager/>
- The SPC-1 Full Disclosure Report (FDR) and Executive Summary for the ISE 820 are available at:
http://www.storageperformance.org/results/benchmark_results_spc1_active#a00155
- Iometer
<http://www.iometer.org>