

ESG Lab Review

EMC Isilon S200: Record-setting Scale-out Performance for Big Data Workloads


Date: June 2011 **Author:** Brian Garrett, Vice President, ESG Lab, and Terri McClure, Senior Analyst

Abstract: This ESG Lab report documents an ESG Lab audit of Isilon's S200 SPECsfs performance testing. Topics explored include the background and purpose of SPECsfs testing, the configuration of the Isilon S200 test bed, Isilon's CIFS and NFS SPECsfs test results with a plain language interpretation of the numbers, and the impact that attaining such results could have on an end-user's business.

Introduction

Keeping up with data growth continues to be a problem for IT organizations large and small. Unstructured data, such as audio, video, images, and office documents, is growing at an especially accelerated rate, generating much more data to be stored over time than either email or database applications (Figure 1). Challenges stem from the sheer amount of data generated, the growing number of endpoint devices that capture and generate that data, and file formats that continuously increase in size and density.

Figure 1. Projected Archive Data Growth By Type



Source: Enterprise Strategy Group, 2010.

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by EMC Isilon.

Background

As data continues to grow, management and efficiency problems plague storage administrators. Even with storage efficiency and capacity reduction technologies such as deduplication and thin provisioning available, users have not been able to stem growth entirely. According to ESG research, one-quarter of users are reporting unstructured capacity growth of more than 40% and 54% indicate that capacity is growing by at least 20% per year.¹ As shown in Figure 1, ESG research also indicates that unstructured data growth is exceeding that of other data types—estimating 226 exabytes of archived file data by 2015.²

This growth is likely to continue in the future. In part because the economy is reviving and in part because there is a need to accommodate the relentless increase in file data, organizations are increasing their NAS purchases. IT users tell ESG that they need a more scalable storage infrastructure to support rapid virtual machine growth as the physical scalability limits of traditional NAS storage systems become a real problem.

NAS vendors have responded by offering solutions targeted at specific use cases, depending on the performance, availability, and scalability required. This has resulted in increasing fragmentation of the NAS market with systems that specialize in certain performance characteristics. Some excel in throughput in terms of MB per second and do well supporting applications that have very large file sizes or that require streaming performance commonly found in line of business applications in specific industries: those that use geospatial imaging or in media and entertainment, oil and gas, or bioinformatics. Other solutions excel in file operations measured in operations per second and do well serving a large number of small files, a common enterprise IT workload.

ESG expects that over time, this specialized segmentation of the NAS market will shrink as next-generation scale-out systems are engineered for a wider range of applicable use cases. Indeed, that process is already underway. [Isilon](#) has long been known as a throughput performance leader, specializing in solutions for throughput-intensive applications like those found in media and entertainment, bioinformatics, and financial services. The new EMC Isilon S200 Scale-out NAS platform brings performance leadership for random-access file-based operations per second as well, making it an attractive solution for a growing number of enterprises.

Introducing the EMC Isilon S200


The EMC Isilon S200 is a high performance, rack-mounted 2U platform built on Isilon's proven unified scale-out storage technology. Tuned for highly transactional and random-access file-based applications, the Isilon S200 is simply deployed in a scale-out cluster accessed as a single file system. As shown in Figure 2, adding Isilon S200 nodes to a field-proven Isilon cluster can be used to meet a wide range of performance and capacity requirements.

The Isilon S200 combines the power of the latest multi-core processors, 10K RPM 2.5 inch SAS drives, quad 1 GbE or dual 1GbE and dual 10GbE front-end networking, a high performance Infiniband back-end, and high-speed SSD drives to accelerate namespace-intensive metadata operations, and, also provides the ability to place mission-critical, latency-sensitive data on SSDs in a SmartPools environment—all to deliver record-setting levels of performance. An S200 cluster of up to 144 nodes is rated at over 1.6 million CIFS operation per second and 1.1 million NFS operations per second with 85 GB of throughput per second.

¹ Source: ESG Research Report, [Scale-out Storage Market Trends](#), December 2010.

² Source: ESG Research Report, [Digital Archiving: End-User Survey & Market Forecast 2010-2015](#), July 2010.

Figure 2. Scaling with the EMC Isilon S200



ESG Lab Analysis

ESG Lab audited a series of record-setting, industry-standard S200 performance benchmarks. The balance of this report introduces the benchmark methodology, the configuration used during testing, the results achieved, and why they matter to the business.

Testing Network File System Performance with SPECsfs

SPECsfs is an industry standard benchmark that measures network attached file server performance and response times. It provides a standardized method for comparing the performance of NAS solutions. The benchmark is client-independent and vendor-neutral. Results must conform to a set of run and disclosure rules validated through a peer review process before being published. SPEC does not allow estimates or unofficial results to be publicized. This report is meant to complement the official results published on the SPEC web site at www.spec.org.

SPECsfs has become a de facto standard benchmark in the storage industry over the past 20 years. Nearly all vendors of NAS solutions publish results. A number of IT managers ESG has spoken with in recent years indicate that they routinely review SPECsfs results when making NAS purchasing decisions.

SPECsfs can be used to run tests and publish results using the CIFS or the NFS protocol. The CIFS protocol is used by computers running the Microsoft Windows operating system to access files on a network attached file system. The NFS protocol is used to emulate file serving activity in Linux, Unix, and virtual server environments.

SFS benchmark code running on an industry standard server generates remote procedure calls to simulate NAS clients. Each client generates a workload representing data traffic that is typical in a general-purpose file-serving environment. Each client generates an equal portion of the overall requested load. The benchmark run rules require the IO load be evenly spread across all network segments and disk storage capacity. The most important SPECsfs result is a measure of the maximum amount of operations per second (ops/sec) a system can sustain while maintaining acceptable response

times. This number can be used as a “score” to compare performance capabilities of two NAS solutions. An overall response time (ORT), a measure of how quickly the server services NAS operations over a range of ten increasing levels of traffic, is also reported.


The New and Improved SPECsfs2008 Workload

The SPECsfs workload was designed to emulate general-purpose file serving traffic. The profile used for the last decade, known as SPECsfs97, can be generally characterized as random file operations using relatively small files (27 KB or less). The workload was updated recently to reflect the changing nature of general purpose file system IO patterns. The updated workload is based on recent data collected by SPECsfs committee members from thousands of file servers operating at customer sites. Just as the files we use in our everyday business activities have grown in recent years, the file and IO request sizes used by SPECsfs2008 have increased. The total file set size for each load point has increased from 10 to 120 MB per SFS operation and the maximum logical transfer size has increased from 135 to 256 KB.

What was Tested

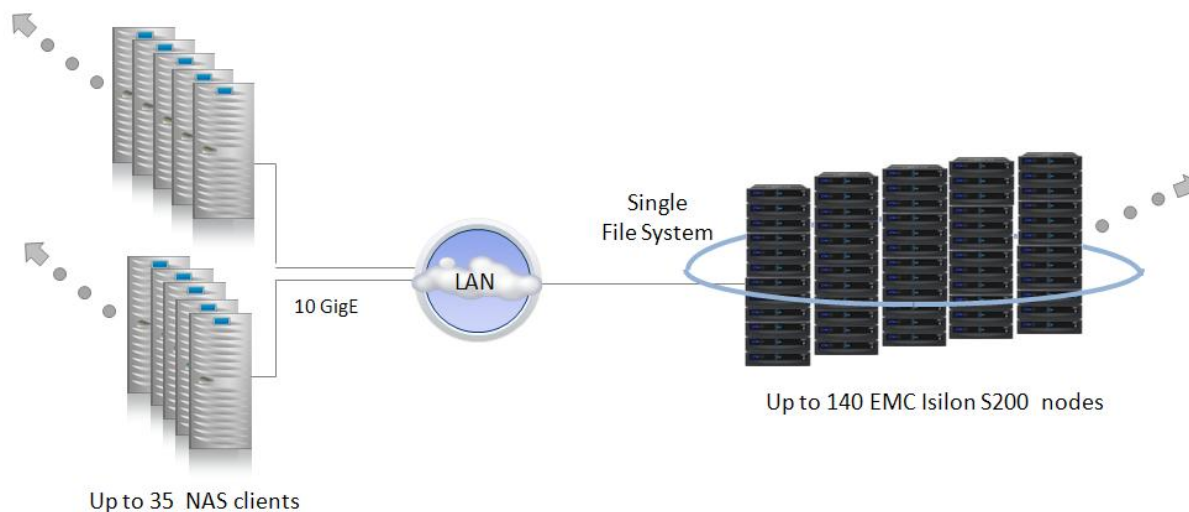
EMC Isilon S200 SPECsfs2008 testing was performed using the hardware shown in Figure 3. Up to 140 Isilon S200 nodes were configured in a cluster and accessed as a single NAS file system. Up to 35 1U servers were configured as NAS clients and connected to the Isilon S200 cluster through a 10 Gigabit Ethernet network.

Figure 3. A Physical View of the Test Bed



A logical view of the test bed is shown in Figure 4. Tests were performed with a scalable number of NAS clients and Isilon S200 nodes with a goal of demonstrating cost-effective, near-linear performance scalability. SPECsfs2008 workloads were tested with Linux and Windows clients using the CIFS and NFS version 3 protocols, respectively.

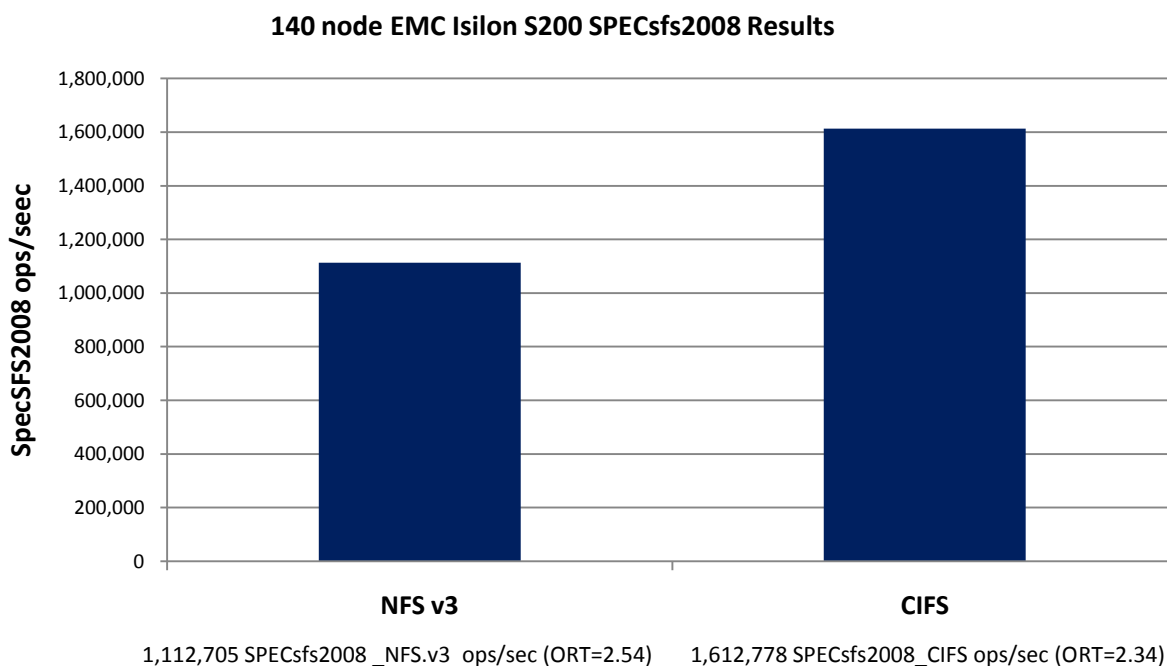
Figure 4. A Logical View of the Test Bed



Record-Setting Performance

The fully configured 140 node S200 solution was used to achieve the record-setting results summarized in Figure 5.

Figure 5. Record-Setting EMC Isilon S200 Performance



What the Numbers Mean

- The 140-node Isilon S200 cluster which was accessed as a single file system sustained more than 1.6 million operations per second with the CIFS protocol and 1.1 million file operations per second using the NFS version 3 protocol.
- Like other previously published SPECsfs2008 results, the CIFS results are greater than the NFS results. This is due to differences in the SPECsfs2008 workload, not due to differences in the NAS system being tested.

- Both results were achieved with a relatively low overall response time (ORT) of fewer than 3 milliseconds. A response time of 3 milliseconds feels instantaneous from an end-user perspective.
- These record-setting results are over twice the performance of the highest published CIFS SPECsfs2008 result and nearly twice the performance of the highest published NFS SPECsfs2008 result when this report was written.
- While the performance is impressive, what’s more impressive in ESG’s opinion is the affordability of the solution from a price/performance perspective. More than a million file operations per second was achieved with a cluster of generally available Isilon S200 nodes using industry standard SAS drives for cost-effective capacity and a reasonably small amount of expensive SSDs (one SSD per node) for metadata acceleration.


Why This Matters

A NAS solution that delivers predictably scalable performance for responsive time sensitive file sharing workflows reduces the cost and complexity associated with supporting a growing number of applications, users, and files in enterprise-class IT environments. ESG Lab confirmed that a cluster of EMC Isilon S200 nodes accessed as a single file system delivered record-setting, industry-audited performance of more than 1.6 million SPECsfs2008_cifs operations per second—more than twice the highest published result when this report was published.

Predictable Performance Scalability

SPECsfs2008 workloads were tested with a scalable number of NAS clients and Isilon S200 nodes with a goal of demonstrating cost-effective, near-linear performance scalability. The results are summarized in Figure 6.

Figure 6. EMC Isilon S200 SPECsfs Performance Linear Scalability



What the Numbers Mean

- A cluster accessed as a single file system was tested with an increasing number of Isilon S200 nodes (7, 14, 28, 56, and 140).
- ESG Lab hands-on testing has confirmed that adding a node to an existing Isilon cluster is extremely simple. Using the front panel menu interface, ESG Lab found that a node can be added to an existing cluster in a matter of minutes—Isilon claims that the time it takes to add a node has been further reduced since ESG conducted the initial test, and it is now under a minute. Existing data is automatically balanced in the background over all of the nodes in the cluster for optimal load balancing and performance.
- Industry-audited SPECsfs2008 results have proven that Isilon performance scales in near-linear fashion as Isilon S200 nodes are added to the cluster.

Why This Matters

Meeting the performance needs of ever-growing random-access file-based application workloads often requires the deployment of multiple traditional NAS solutions. Solving a performance problem with multiple NAS solutions increases the cost and complexity of meeting the needs of the business.

Industry-audited SPECsfs2008_cifs and SPECsfs2008_nfs.v3 results have proven that EMC Isilon S200 delivers cost-optimized general purpose NAS performance scalability which scales in near-linear fashion.

The Bigger Truth

Isilon was the very first solution tested by ESG Lab in 2003—we were extremely impressed by the massive throughput and amazingly simple scalability of its scale-out NAS solution. Back then, it was the perfect system for early adopters that needed massive amounts of throughput for large files. Since then, the requirements of early adopters have turned into the big data requirements of every-day organizations. Recently acquired by EMC, Isilon is in the envious position of already having the kind of scale-out architecture required for the most demanding storage performance environments. The advent of “big data” has driven application performance requirements right into EMC’s wheelhouse.

ESG Lab has confirmed that EMC has set another world record with its latest SPECsfs performance benchmark results. With new Isilon S200 platform nodes tuned for IO-intensive applications, EMC recorded more than 1.6 million file operations per second. While a million file operations per second is a really big and impressive number, what’s even more impressive is that it was achieved using a modular architecture that can be used to cost-effectively meet the performance needs of a wide variety of workloads from high performance big data workloads, through primary and secondary storage, to near-line archives—all from a single file system architecture.