

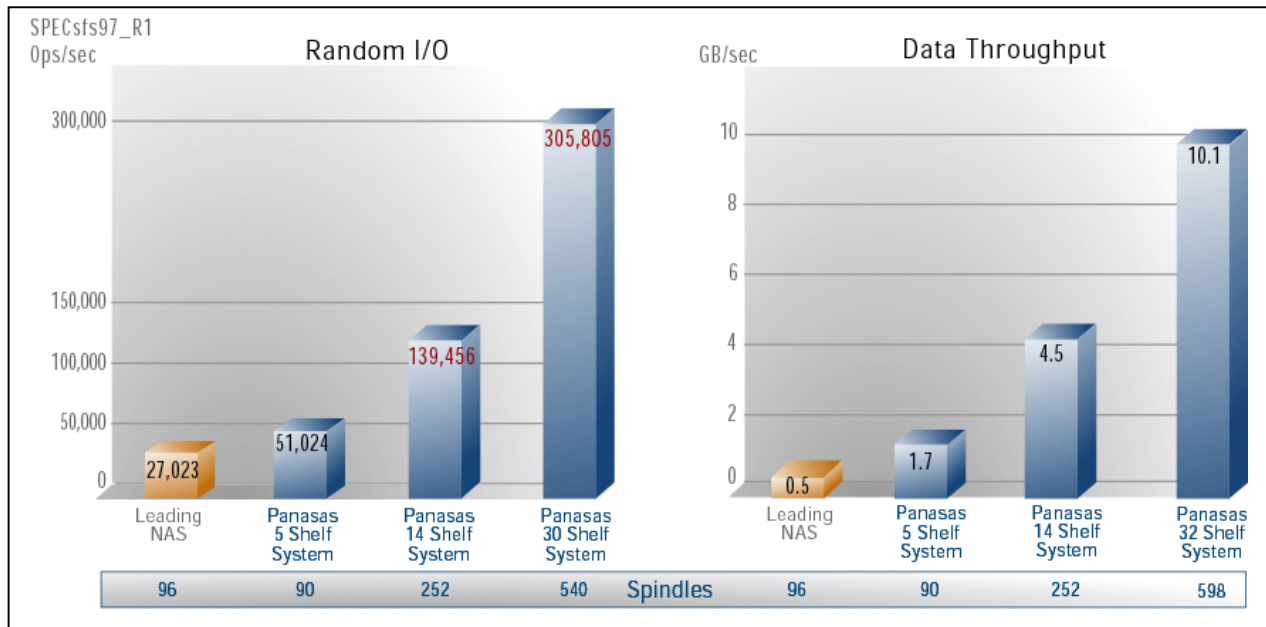


# Panasas ActiveScale Storage Cluster™ Performance Backgrounder

13 November, 2003

## Summary

The Panasas ActiveScale Storage Cluster™ delivers industry-leading performance and scalability in both random I/O and throughput benchmarks. The results demonstrate Panasas' unrivaled performance: **305,805 NFS ops/sec** and **10.1 GBytes/sec sustained sequential read performance**.



## Benchmark Results

Figure 1 presents the results of Panasas' tests, along with published performance numbers of the leading NAS vendor.

### Random I/O Results<sup>1</sup>

In its largest tested configuration, a 30-shelf system, the Panasas Storage Cluster delivered **305,805 NFS ops/sec** with an overall **response time (ORT) of 1.76 ms**. These results are more than **ten times** the performance of the leading NAS vendor's comparable configuration, and more than double the largest SPEC™ SFS results posted to date.<sup>1</sup>

### Data Throughput Results

A similar 32-shelf configuration was used to demonstrate unequaled sequential I/O throughput for Linux clusters. This configuration delivered **10.1 GBytes/sec sustained** performance on sequential I/O read tests.

### Superior Performance and Linear Scaling

The Panasas system demonstrates superior performance at comparable configurations and linear scaling to record-setting performance levels for both random I/O and sequential workloads.

At comparable configurations, the system delivers **89% greater random I/O performance** and **3.4 times greater sequential I/O performance** than the leading NAS vendor's configuration.

Panasas achieves an unprecedented **99% efficiency** in scaling performance to record-setting levels in a 72-TB storage cluster, the largest configuration tested to date. The unique storage architecture is capable of extending this linear scaling to single managed systems of more than a Petabyte.

<sup>1</sup> NFS ops/sec as measured by the SPEC™ benchmark. SPEC™ and the benchmark name SPECsfs97\_R1™ are registered trademarks of the Standard Performance Evaluation Corporation. Competitive benchmark results stated above reflect results published on [www.spec.org](http://www.spec.org) as of Nov 13, 2003. The comparison presented above is based on results of a *comparably configured* product from the industry leading NAS. For the latest SPECsfs97\_R1 benchmark results visit [www.spec.org/sfs97r1](http://www.spec.org/sfs97r1).

## About the Performance Benchmarks

### Testing Methodology

Panasas conducted both sequential and random I/O throughput testing in two benchmarking scenarios. The first is a *comparative* scenario, designed to demonstrate performance and price-performance for a configuration similar to those of key competitor's comparably provisioned systems. The second is a *scale* scenario, designed to demonstrate system scalability and assert world-record results for general purpose I/O in a verified, reproducible configuration.

For both of these scenarios, general purpose I/O testing was performed using the SPECsfs97\_R1 (v3.0) benchmark (<http://www.spec.org/sfs97r1/>). The SPECsfs benchmark is an industry accepted standard benchmark that characterizes a storage system's performance under a representative real-world NFS workload. The workload includes a mix of file and directory operations, including lookup, read, write, create, and others. The average file size is 27KB, with 85% of the files being 16KB or less in size.

Sequential throughput (bandwidth) testing was conducted using a number of Linux clients executing parallel read and write commands between the clients and the storage system, with each client reading and writing its own files. This benchmark measures the ability of the Panasas Storage Cluster to deliver high bandwidth I/O simultaneously to many clients in a typical cluster computing workload.

### System Configurations

The *comparative* scenario was run on a system provisioned comparably to the configuration of the leading NAS vendors as determined from their SPECsfs results submissions. More specifically, the Panasas system was configured to match, as closely as possible in a full-shelf configuration, the *spindle count* of this system. Spindle count – which refers to the number of disks in the system – is an often-used “equalizer” to explore the performance of similarly equipped storage systems. The NAS vendor's non-failover configuration includes 96 spindles. Panasas' nearest full-shelf configuration employs 90 spindles (45 StorageBlades). The Panasas Storage Cluster configuration specifications for the *comparative* scenario are presented in table 1, below.

The *scale* scenario was run on two different configurations to highlight the linear scaling of the Panasas ActiveScale Storage Cluster architecture, and to demonstrate record-setting performance numbers. These configurations deliver world-record performance for the SPECsfs benchmark and sequential I/O tests; the configurations are presented in tables 2 and 3.

**Table 1. Comparative Scenario Configuration**

<b>Storage</b>	5 shelf system <ul style="list-style-type: none"><li>• 45 240GB Panasas StorageBlades</li><li>• 10 Panasas DirectorBlades</li></ul>
<b>Clients</b>	35 2.4GHz P4 Xeon / GbE; RedHat 7.3 (2.4.21)
<b>Network</b>	Extreme Networks® BlackDiamond 6816 5 G8Ti 100/1000BASE-T 8-port line cards

**Table 2. Scale SFS Configuration**

<b>Storage</b>	30 shelf system <ul style="list-style-type: none"><li>• 270 240GB Panasas StorageBlades</li><li>• 60 Panasas DirectorBlades</li></ul>
<b>Clients</b>	35 2.4GHz P4 Xeon / GbE; RedHat 7.3 (2.4.21)
<b>Network</b>	Extreme Networks® BlackDiamond 6816 16 G8Ti 100/1000BASE-T 8-port line cards

**Table 3. Scale Bandwidth Configuration**

<b>Storage</b>	32 shelf system <ul style="list-style-type: none"><li>• 299 240GB Panasas StorageBlades</li><li>• 1 Panasas DirectorBlade</li></ul>
<b>Clients</b>	150 IA32 / GbE nodes (various makes and models); RedHat 7.3 (2.4.21)
<b>Network</b>	Extreme Networks® BlackDiamond 6816 16 G8Ti 100/1000BASE-T 8-port line cards 32 Dell 2400 GbE switches