

This marks our first analysis of the new SPECsfs®2008* benchmark results. One will no doubt recall that SPEC® SFS97 was closed out earlier this year and a new SPECsfs2008 benchmark was released. Unfortunately there are not a lot of SPECsfs2008 results posted at this time but what exists is reported here.

Latest SPECsfs2008 results

One advantage of the new SPECsfs2008 benchmark is the support for both CIFS and NFS performance data. Although as stated on the SPECsfs2008 website, “SPECsfs2008 results may only be compared to other SPECsfs2008 results for the same protocol. SPECsfs2008_cifs and SPECsfs2008_nfs.v3 are not comparable because they are generated using completely different workloads.”# SCI has provided below a hypothetical comparison of the two protocol

SPECsfs2008* NFSv3 vs. CIFS with linear trend line, as of 21 December 2008

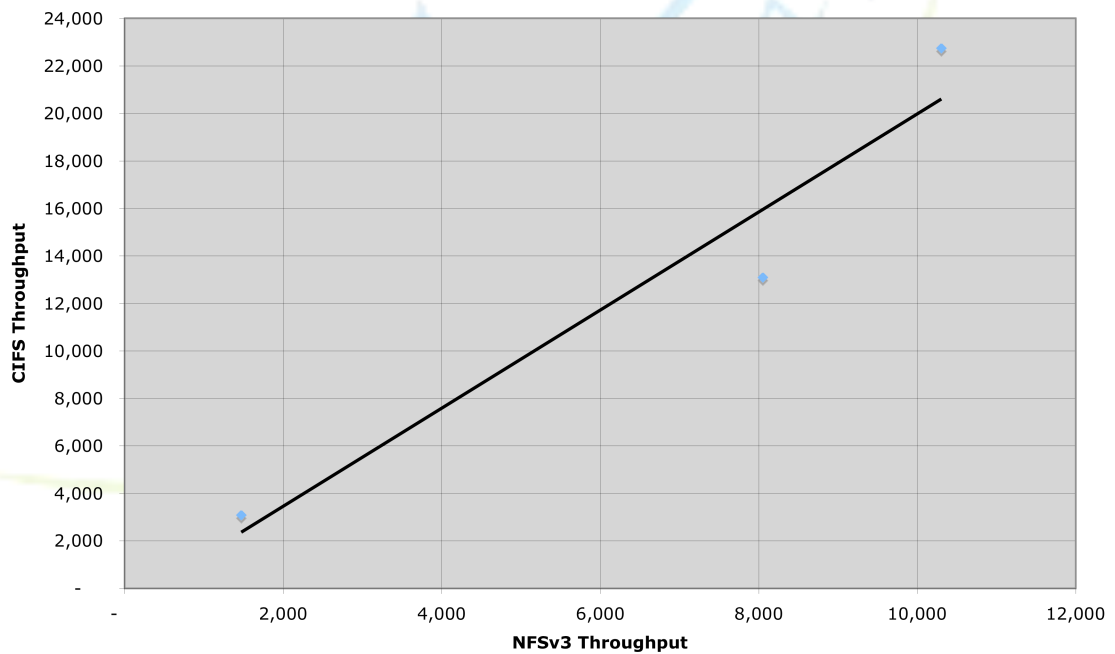


Figure 1 SPECsfs2008 NFS V3 vs. CIFS throughput

benchmark results. My friends at Microsoft will be pleased to see that when comparing the same subsystems running CIFS and NFSv3 for SPECsfs2008 workloads CIFS throughput is on average twice the NFS throughput. In figure 1 we have a scatter plot of the three subsystems that reported both NFS and CIFS results.

* SPECsfs2008 results from <http://www.spec.org/sfs2008/results/>

SPECsfs2008 statements from <http://www.spec.org/sfs2008/results/> as of

1/22/09

As can be seen by the chart the linear regression is quite good (R^2 at 0.93) and the regression equation is

$$\text{CIFS_throughput} = \text{NFS_throughput} * 2.1 - 679$$

A couple of provisos to watch for here:

- Only three subsystems recorded results for both interfaces but the data looks pretty consistent at the moment. As more data is reported for both CIFS and NFSv3 we will update the regression but for now it looks pretty solid.
- NFS workloads are not readily comparable to CIFS in a number of dimensions not the least of which that NFS is stateless and CIFS is state-full. SPECsfs benchmarks for the two are implemented to reflect those differences. Also, the relative proportions of the actual workload counterparts don't exactly matchup, e.g. workload percentages for NFS read and write operations versus CIFS read_andx and write_andx operations are slightly different (NFS read@18% vs. CIFS read_andx@20.5% and NFS write@10% vs. CIFS write_andx@8.6%)^{\$}, file sizes used are different, and all the remaining operations, which in all fairness represent a significant majority of their respective workloads, are by definition relatively impossible to compare.

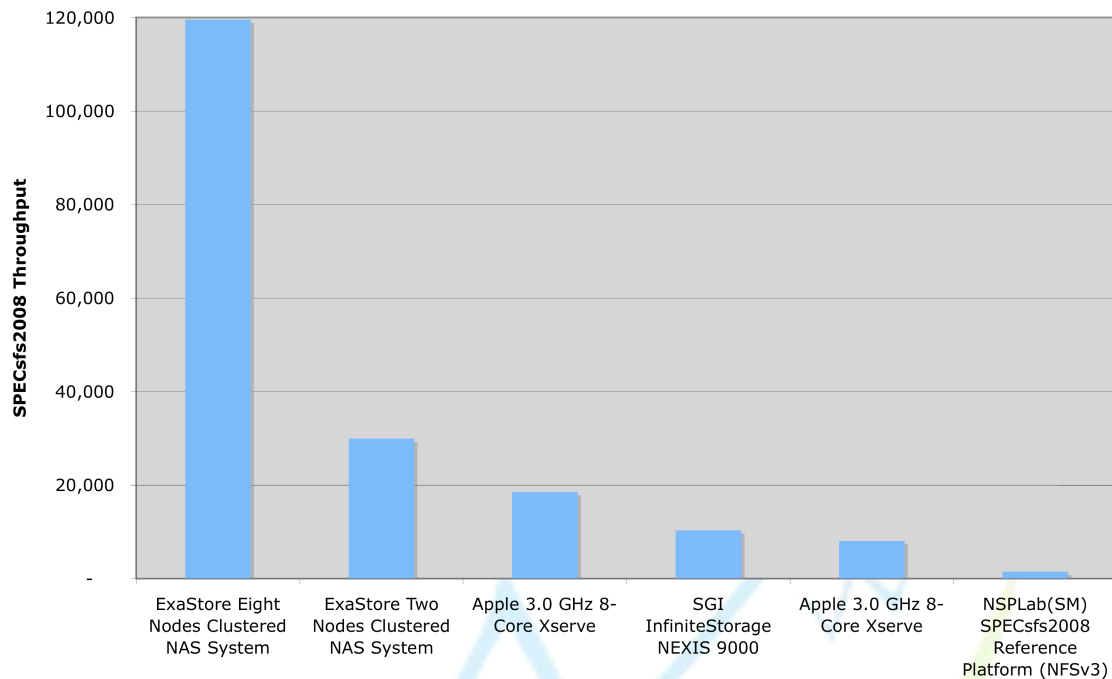
Nonetheless, taking a significant leap of faith that at the user level all of these protocol details result in emulating comparable end-user workloads, the results do show a significant difference in throughput results between the two protocols.

Both SPEC SFS97 and SPECsfs2008 support any networking hardware supplying TCP or UDP but SPECsfs2008 is the first one with published results for Infiniband, GigE, and 10GbE as system interfaces. One would think that Infiniband would have the edge in any comparison among these three but as figure 2 shows it's not apparent that networking hardware has any advantage. The SGI product is using Infiniband and both ExaStore benchmarks use 10GbE for these results. In all fairness the networking connection may not be a limiting factor in SPECsfs2008 results.



^{\$} *SPECsfs2008 User's Guide*, available from <http://www.spec.org/sfs2008/> as of 1/22/09

SPECsfs2008* NFS v3 Results as of 21 December 2008

**Figure 2 SPECsfs2008* NFSv3 results**

The *SPECsfs2008 User's Guide* also states the workloads have changed between the old SPEC SFS97 and the latest benchmark. Some of these changes include:

- Operation percentages have been adjusted to reflect recent usage
- Maximum file sizes have been increased
- Total file set size for a given workload has been increased
- Percentage of files accessed during a run has been increased
- Logical file transfer size has been increased
- NFS block (physical) transfer size is now negotiated with the server

All of this would make comparing SPEC SFS97 against SPECsfs2008 workloads a difficult endeavor. Also, no vendor has currently released a SPECsfs2008 and a SPECsfs97 benchmark for the same system. However, if more than a few vendors did this someday we might be able to calibrate results between the two. By not showing how results between the old and new benchmarks can differ may hold vendors back from releasing new SPECsfs2008 results. However, if one were able to compare results under both the old and new benchmarks, one could more easily show that current results were due to the changes in the benchmark and not an artifact of the system being benchmarked.

SPECsfs2008* NFSv3 ORT results as of 21 December 2008

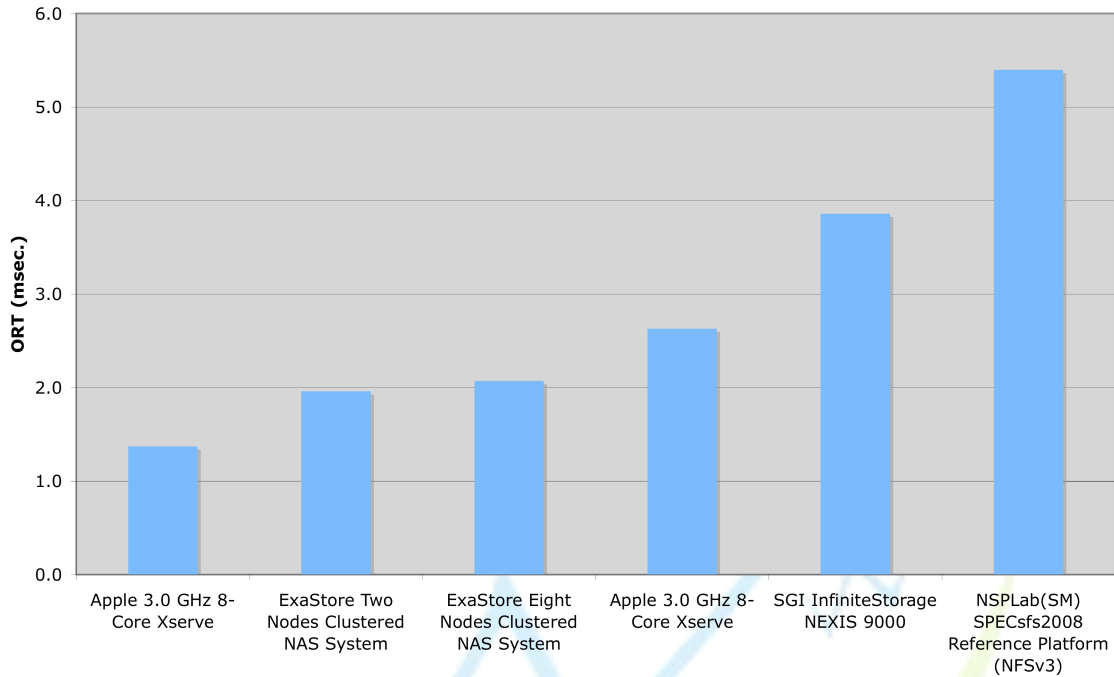


Figure 3 SPECsfs2008* ORT results

Another nail in the Infiniband coffin is delivered by ORT results. Again one result does not make a statistical relationship and physical network connections may not be a limiting factor in ORT measurements. However, as you may recall this is the average overall response time delivered by a storage system and SGI is the only Infiniband user present. Of course this doesn't say much about 10GbE either as ExaStore was the only vendor reporting 10GbE results. Also of interest is the great showing of the Apple Xserve device (with 49 disk drives), not bad for this class system. If only they could grow the top end they could have something here.

Next we turn to CIFS results. So far only three results have been released but what is shown below. Recall the SGI is using Infiniband while the other two are using GigE hardware interfaces.

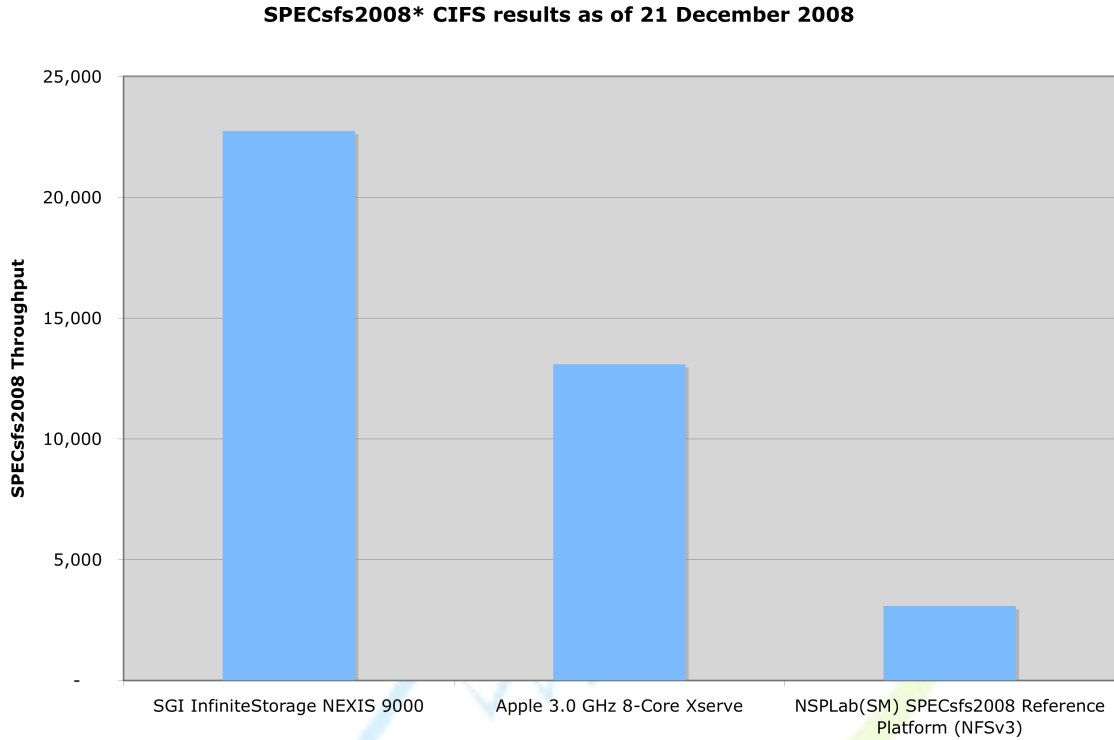


Figure 4 SPECsfs2008* CIFS results

Significance

SPEC SFS97 has served the IT industry well but its time has come. SPECsfs2008 has the potential to be a worthy substitute but only wider adoption and more result submissions can prove this out.

Notwithstanding all the provisos against comparing NFS and CIFS SPECsfs2008 results, we especially like that we can at least try to compare CIFS and NFS results albeit with "... completely different workloads". Whether these comparisons are real or not only time will tell. More data here will help refine the apparent advantage CIFS enjoys. Perhaps this was obvious to most vendors inasmuch as NFS was a stateless protocol and CIFS was state-full protocol, but it had never been shown quite so publicly before.

Intermixing Infiniband, 10GbE and GigE with the same results should be informative. Most HPC users would swear by the advantage inherent in using Infiniband for I/O operations. However, current results show that this doesn't seem to offer much of an advantage in SPECsfs2008 results. Results are skimpy as of yet and we may live to see these words thrown back at us but for now they stand as is. It would certainly be better if the same system released three results with GigE, 10GbE and Infiniband hardware but its unclear why any vendor would see an advantage doing this.

In retrospect, it's probably the right time to move onto a benchmark that can show CIFS and NFSv3 results and doesn't care what hardware is used. Hopefully, the end-user community will require their storage vendors to publish more SPECsfs2008 results and thereby help to make this a worthy successor to SPEC SFS97.

Silverton Consulting, Inc. is a Storage, Strategy & Systems consulting services company, based in the USA offering products and services to the data storage community