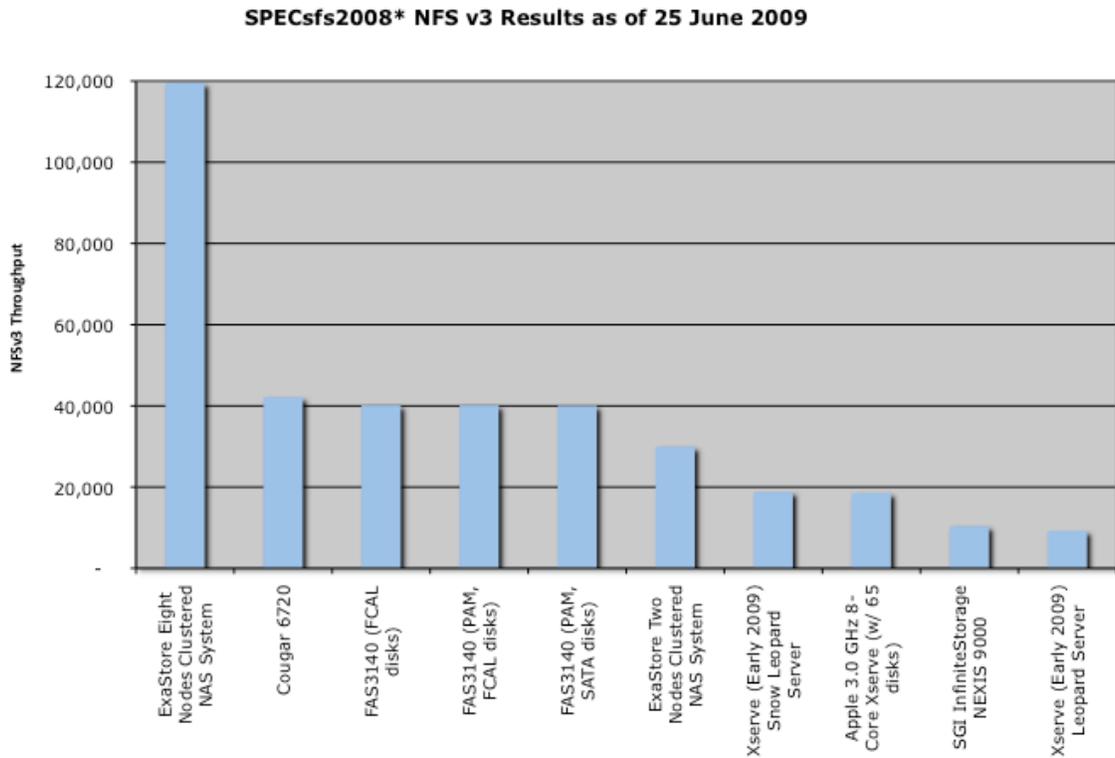


We now turn to analysis of the new SPECsfs®2008\* benchmark results. Unfortunately there were not a lot of highend SPECsfs2008 results, most notably ONStor and Apple for NFS, and Apple and Fujitsu Siemens for CIFS.

## Latest SPECsfs2008 results



**Figure 1 SPECsfs2008\* NFSv3 throughput results**

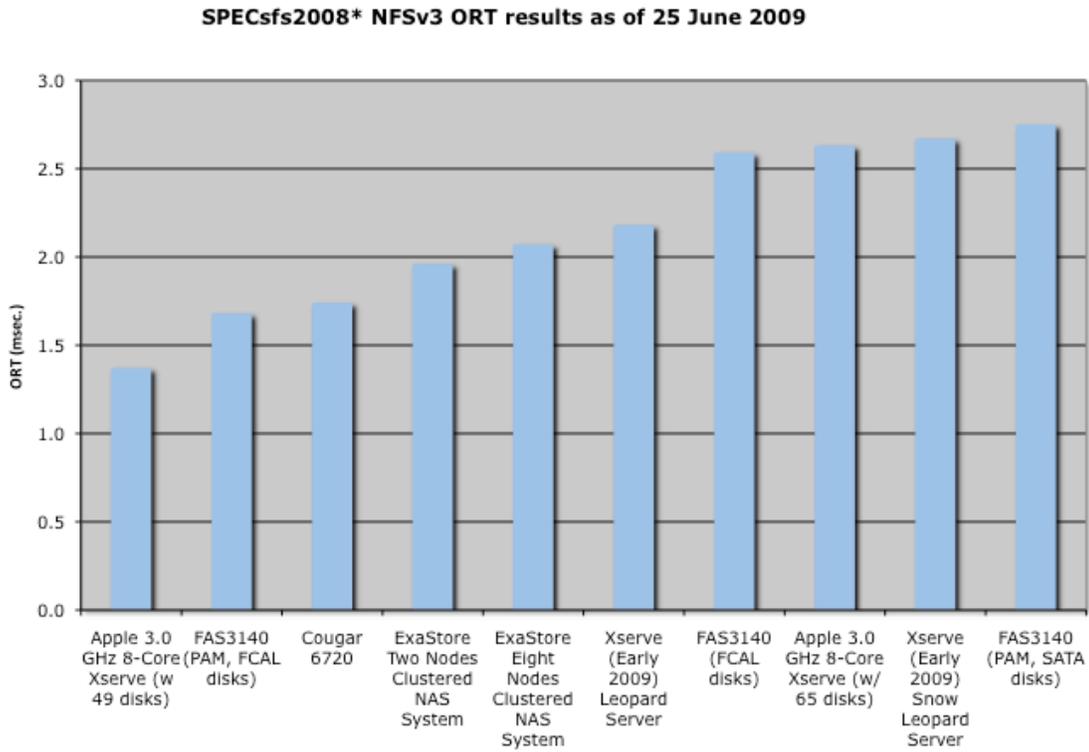
ONStor Cougar now takes second spot in the top 10 throughput results. The Cougar system had ~½ the disks of the ExaStore box and ~7 times less memory (cache). Given all that, its results standup pretty well. The two new Apple NFS benchmark results (Snow Leopard and Leopard server) round out the rest of the new members to the top 10 list at numbers 7 and 10 respectively.

Recall from our last report# that some NetApp results utilized their PAM card. Also, the SGI product result used Infiniband, both ExaStore benchmarks used 10GbE and all the rest use GigE. In all fairness the networking connection may not be a limiting factor in SPECsfs2008 results.

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

# Available at

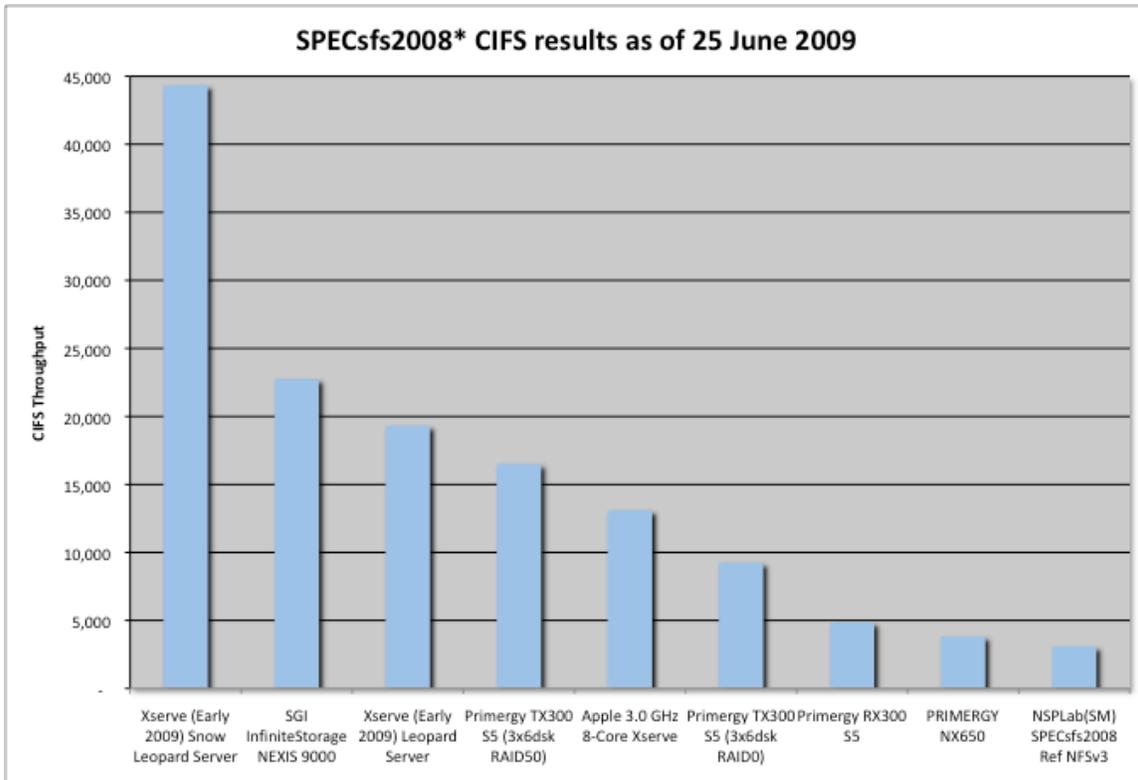
[http://www.silvertonconsulting.com/page2/page2d/storage\\_int\\_dispatch.html](http://www.silvertonconsulting.com/page2/page2d/storage_int_dispatch.html)



**Figure 2 SPECsfs2008\* NFSv3 ORT results**

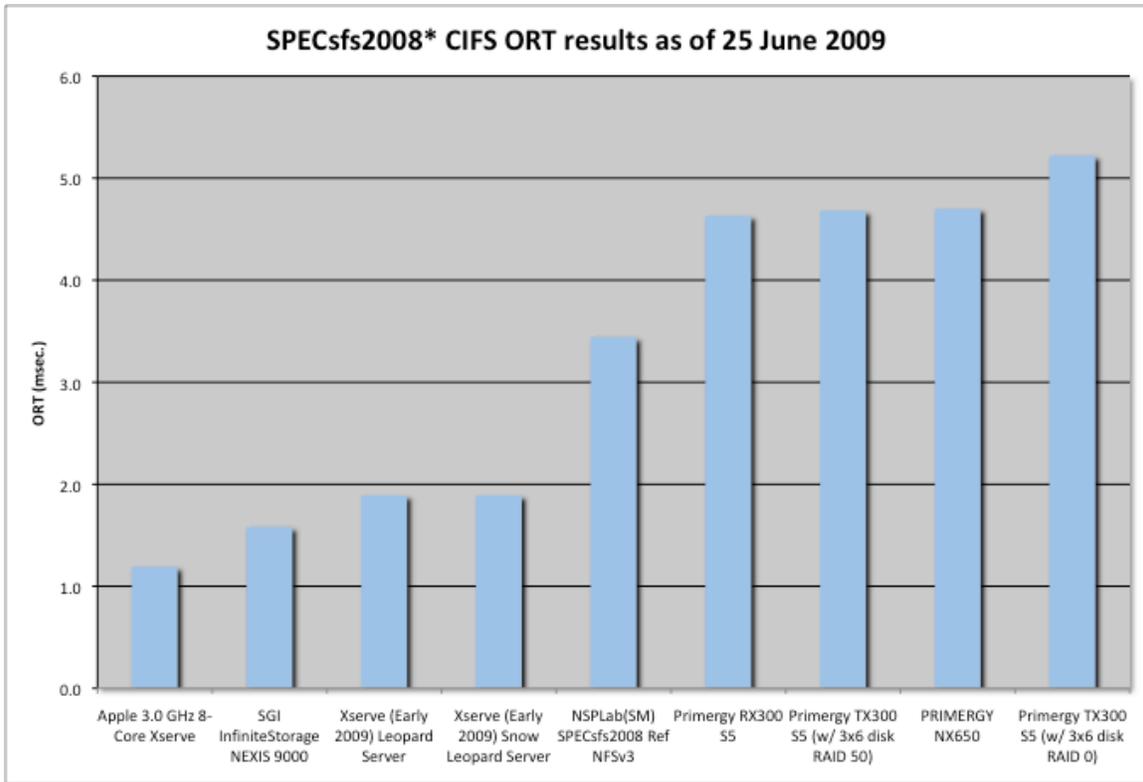
As discussed last time for NFS ORT results, one can clearly see the advantage of NetApp’s PAM with FC disks and yet, the new ONStor Cougar benchmark shows up at number 3, only ~60 microsec behind the NetApp/PAM result. The only other new showing was Apple’s Snow Leopard server coming in at number 9.

Next we turn to CIFS results, the five new results have more than doubled SPECsfs2008 CIFS benchmarks. Recall the SGI is using Infiniband while all the others use GigE interfaces.



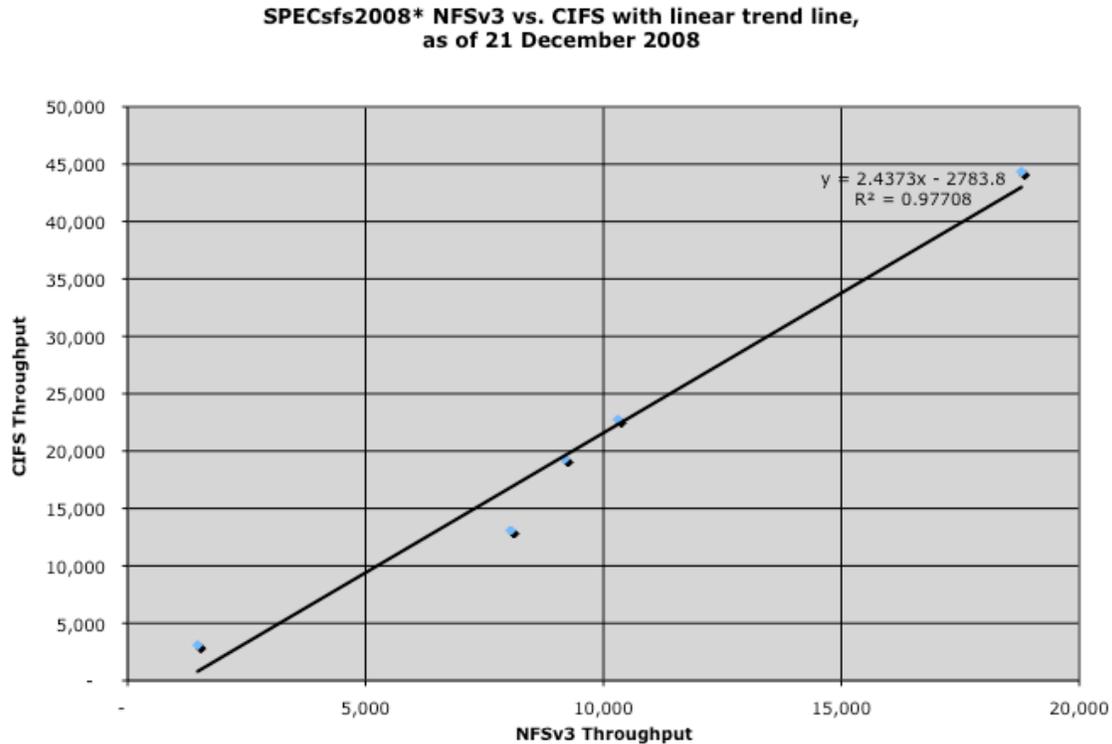
**Figure 3 SPECsfs2008\* CIFS throughput results**

We suppose it's not surprising to see Apple's Snow Leopard leading the pack, coming in at the new #1 in CIFS throughput considering its market place but one would think others could do better. More impressive is that the Snow Leopard result used only 65 disks whereas the SGI result sported 242 disks (~4X). It's unclear to us whether this is the new Apple OEM of Sun ZFS file system at work here, but clearly Apple CIFS performance has improved significantly.



**Figure 4 SPECsfs2008\* CIFS ORT results**

Once again, Apple shows up well in CIFS ORT results. Although, as best we can determine this #1 result was an early Leopard version (Mac OSX10.5.1) whereas the #3 result (using Mac OSX10.5.7) had a 2.93Ghz Nehalem processor. The other major difference was a dual port GigE card for the #1 result vs. a 6-port GigE card in the slower version.



**Figure 5 SPECsfs2008\* CIFS vs. NFS throughput correlation**

We have discussed this in earlier reports but once again the results would support our contention that the CIFS protocol results in better throughput than NFSv3. As pointed out to me, a couple of provisos are warranted here, namely:

- NFS workloads are not readily comparable to CIFS in a number of dimensions not the least of which is that NFS is stateless and CIFS is state-full. Also, the relative proportions of the actual workloads don't exactly matchup, e.g. 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 are different, and all the remaining operations, which, to be fair, represent a significant majority of their respective workloads, are by definition, nigh impossible to compare. SPECsfs benchmarks for the two are implemented to reflect all of these differences.
- A majority of these results (3 of 5) come from the same vendor (Apple) and their great CIFS and/or poor NFS implementations may be skewing results.
- Only five subsystems have recorded results for both interfaces but the correlation looks pretty good for now.
- Normally, host operating system affects could skew these results but the SPECsfs2008 benchmarks emulate their own client side stacks for both protocols, thus negating any operating system affects.

Nonetheless, once again, considering that at the user level all specific protocol details result in emulating comparable end-user workloads, the results do show a significant advantage for CIFS (~2.4X) throughput over NFS.

## Significance

Our earlier discussion on CIFS vs. NFSv3 throughput differences resulted in quite a lot of discussion. It was early then, and still is now, but we continue to stand by our claim, given benchmark results, CIFS seems to perform better than NFSv3. More dual protocol results should help clarify this relationship.

Slowly, more SPECsfs2008 results are being released. But, where are the major NAS systems. It's been 10 months since the old SPECsfs benchmark was retired and we still lack benchmark results for all the major NAS systems. In the mean time, smaller players continue to release results; just happy to get any visibility, validity and traction they can muster.

---

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