

We once again return to discuss the latest Storage Performance Council (SPC) results*. Also, SPC has recently introduced new versions of its benchmarks to measure storage component performance. We report on these new results below.

SPC-1* results

There have been only two new SPC-1 results these past three months one for Pillar Data Axiom 600 and the other for Sun Storage 6780. Neither of these two made the top 10 in IOPS™, LRT™, or \$/IOPS™ so these charts can be found in prior dispatches#. However, both of these products did make the top 10 for IOPS/\$/GB (see Figure 1 below).

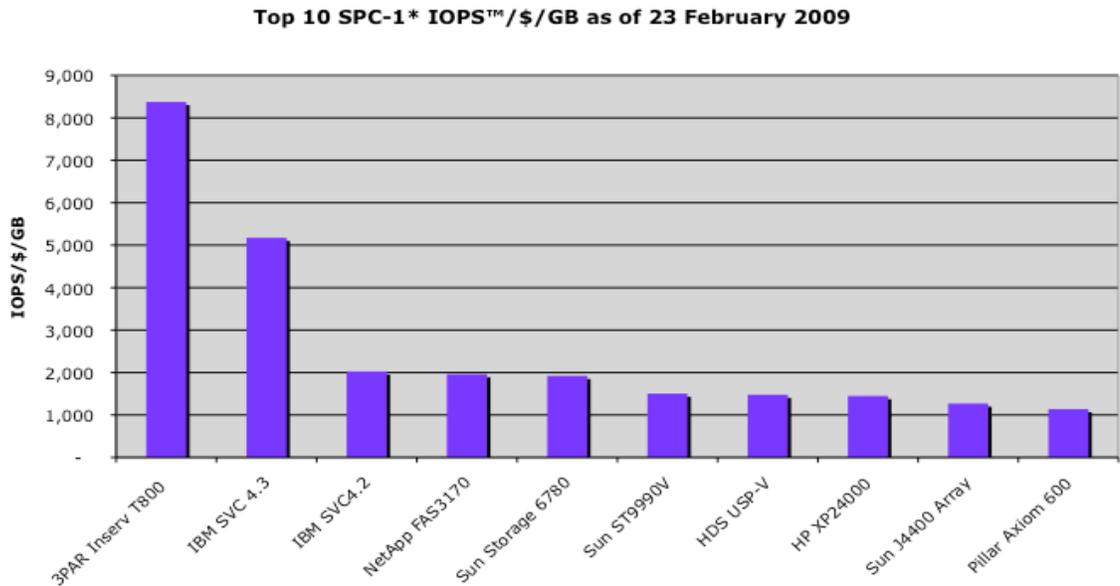


Figure 1 Top 10 SPC-1* IOPS/\$/GB

The big two players here from last report are 3PAR and SVC 4.3. Again, we calculate this metric as = IOPS / (\$/GB) and created this metric as another way to factor in performance against cost and capacity.

On the other hand, as has been pointed out to SCI, this metric may unfairly advantage big, monolithic subsystems at the expense of smaller subsystems. The monolithic subsystems generate such high IOPS counts that they're relatively expensive \$/GB doesn't impact their ranking on this chart. In contrast, smaller

* All results from www.storageperformance.org as of 23 February 2009

Prior SPC performance dispatches can be found at

http://www.silvertonconsulting.com/page2/page2d/storage_int_dispatch.html

subsystems, such as Xiotech’s Emprise may be capable of putting up high IOPS rates by aggregating a number of smaller subsystems but in their current instantiation, their relatively modest \$/GB doesn’t compensate for the resultant IOPS and hence they cannot compete on this chart. In such a configuration, even when taking additional switch port costs into account, may still be significantly less costly than the systems shown here.

Again it would seem the FAS3170, Sun Storage 6780 and Pillar Axiom 600 look out of place here with these multi-million \$ subsystems (3PAR, IBM SVC 4.3&4.2, Sun T9990V, HDS USP-V, and HP XP2400) but seems to provide relatively good performance for its price and capacity. Also, once again the Sun J4400 has no RAID protection whatsoever and probably does not belong here.

SPC-2 results

There were also two new SPC-2 benchmarks recorded for this update, both from Sun, one the SUN 6780 with RAID5 and the other the same system with RAID6. Both these systems are new entries in the top 10 MBPS™

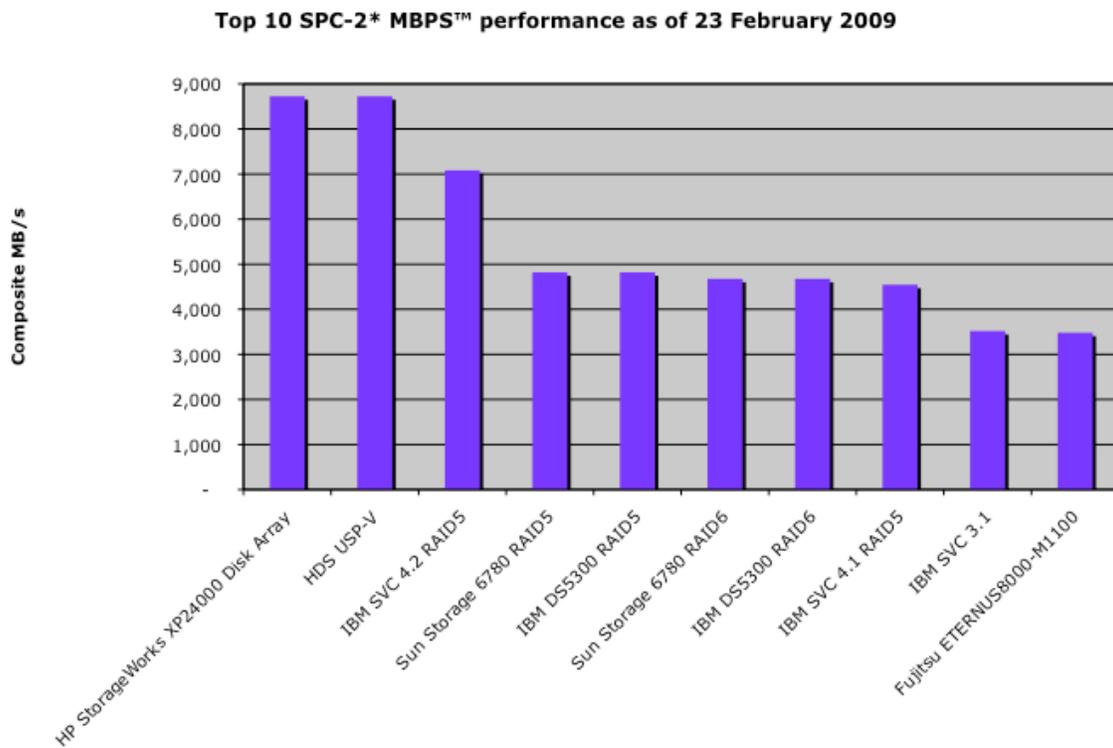


Figure 4 Top 10 MBPS™

There was a slight degradation in performance for the RAID6 over the RAID5 version of the Sun Storage 6780 of roughly 3%. All the other systems were reported on in prior SPC performance dispatches.

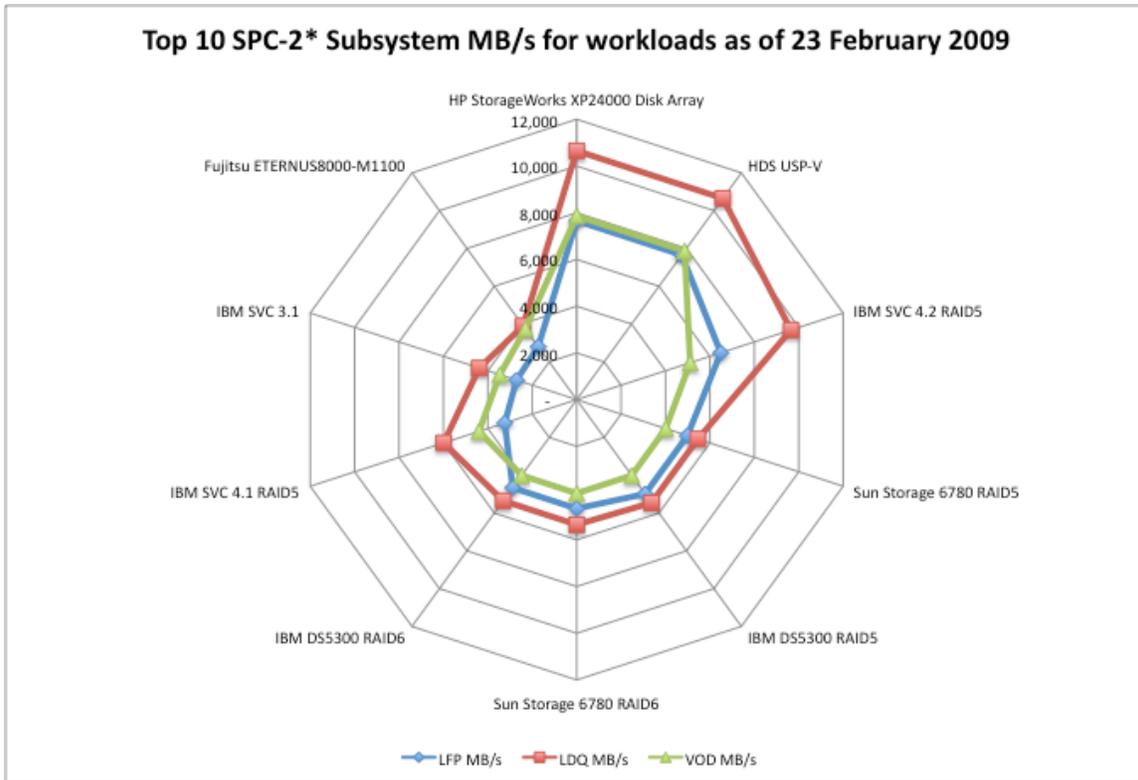


Figure 5 Top 10 SPC-2* Workload Results

With this report SCI is introducing a new way to show subsystem performance for each of the three workloads that constitute the composite MPBS™ score. Here one can see some variability in the scale of workloads each vendor’s product can attain. There seems to be some interesting difference between the LDQ (large database query) workloads and the other two LFP (large file processing) and VOD (video on demand).

It’s unclear why the HP XP24000, HDS USP-V and the SVC 4.2 do so well on the LDQ workload as compared to the other two workloads. SCI believes the secret lies somewhere in their respective caching algorithms being optimized for LDQ and not as well optimized for the other two workloads. This may say that the other products such as the Fujitsu ETERNUS8000, Sun Storage 6780 and IBM DS5300 could improve their LDQ performance if they tweaked their caching algorithms somewhat.

New SPC-1C and -2C Results

SPC has been busy and have created new versions of their SPC-1 and SPC-2 benchmarks specifically to measure storage component performance. Currently there are not many released results. Also, SCI is having some difficulty trying to compare a 24 drive RAID subsystem against a single drive SATA system but other than that we believe the results are worthy of elaboration.

SPC-1C Results

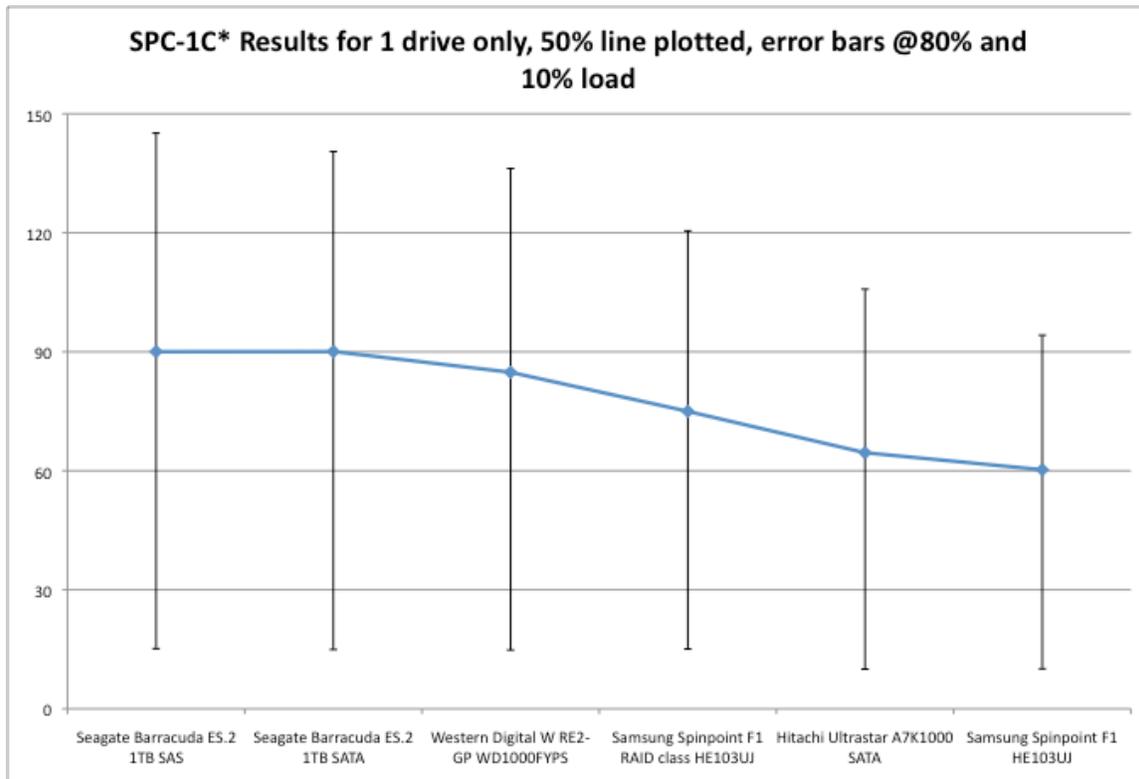


Figure 2 SPC-1C Results

SPC-1 results have always reported details for 10%, 50%, 80%, 90%, 95%, and 100% load levels with IOPS counts and response times in msec. 100% is based on exceeding a pre-determined response time cap, which SCI believes to be greater than 30msec.

For the single drive results instead of reporting IOPS™ which is at the 100% level SCI decided to use the 50% load level and plot error bars up to 80% load and down to 10%. At the 100% load point many of these single drives are producing response times just shy of the 30msec cut off. Unclear whether anyone would run these drives at that slow a response time. Thus, SCI decided to show the 50% load level with error bars. Somewhere around 80% load, the drive busy-ness takes over and how it manages its seek queue drives performance rather than seek speed.

All of these devices are running 7200RPM and have 1TB of storage. The first drive listed uses a SAS interface, all the others use SATA. LSI SAS3041X-R HBAs were used for all these runs. A couple of caveats here

- Seagate sponsored all of these results. But the results as reported were from benchmark runs executed at the SPC lab with sponsored supplied hardware.

- There are three other SPC-1C results, for 12, 24, and 15 drive RAID systems that don't seem comparable to these results so they have been left out of this analysis.

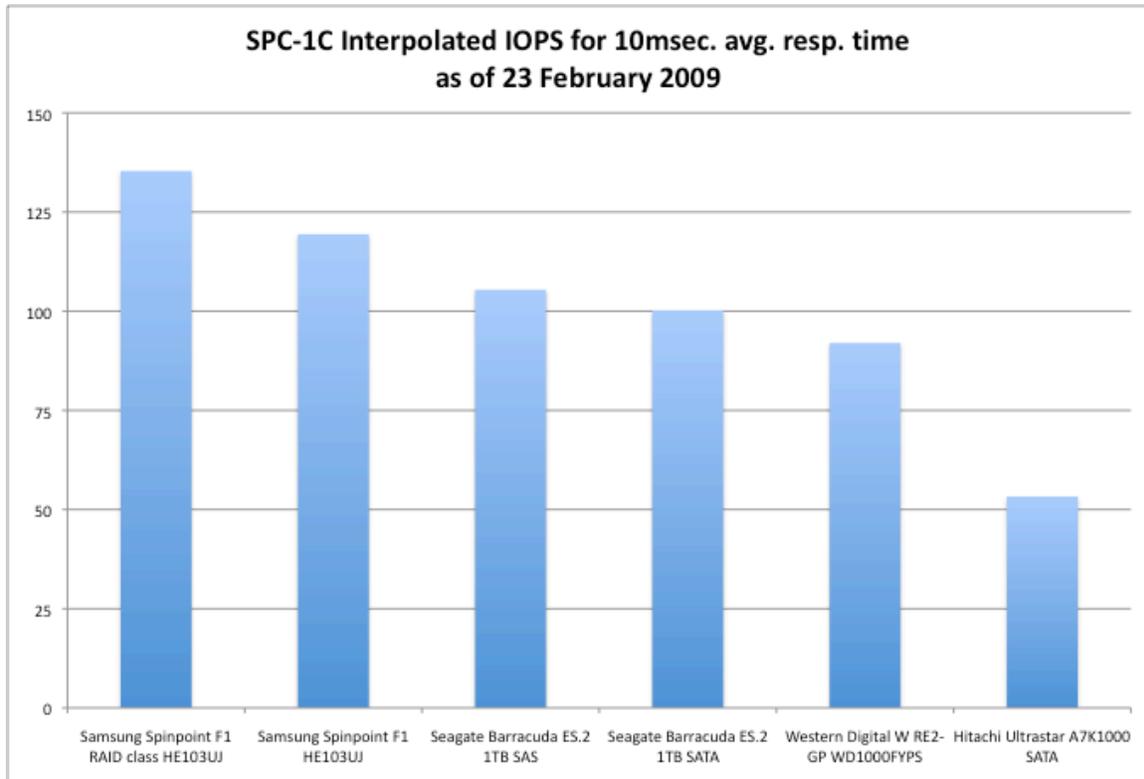


Figure 3 SPC-1C Interpolated IOPS at 10msec average response time

Many end users are more interested in what a storage subsystem could do at a constant 10msec response time. As such, we introduce another new chart for the SPC-1C that depicts the expected IOPS at an average response time to 10msec. For most drives, the IOPS shown in the chart is an interpolated value. However for the Samsung Spinpoint F1 HE103UJ the value shown represents what they actually achieved at the 100% workload value.

For some reason both the Samsung devices seem to do well at 10msec response time while the others fall off. This could possibly be due to just different seek optimization profiles targeted at varying workload levels. The 50% workload level is typically closer to the 10msec mark for all the other drives. For the other Samsung drive, the 10msec response time is closer to its 95% load.

SPC-2C results

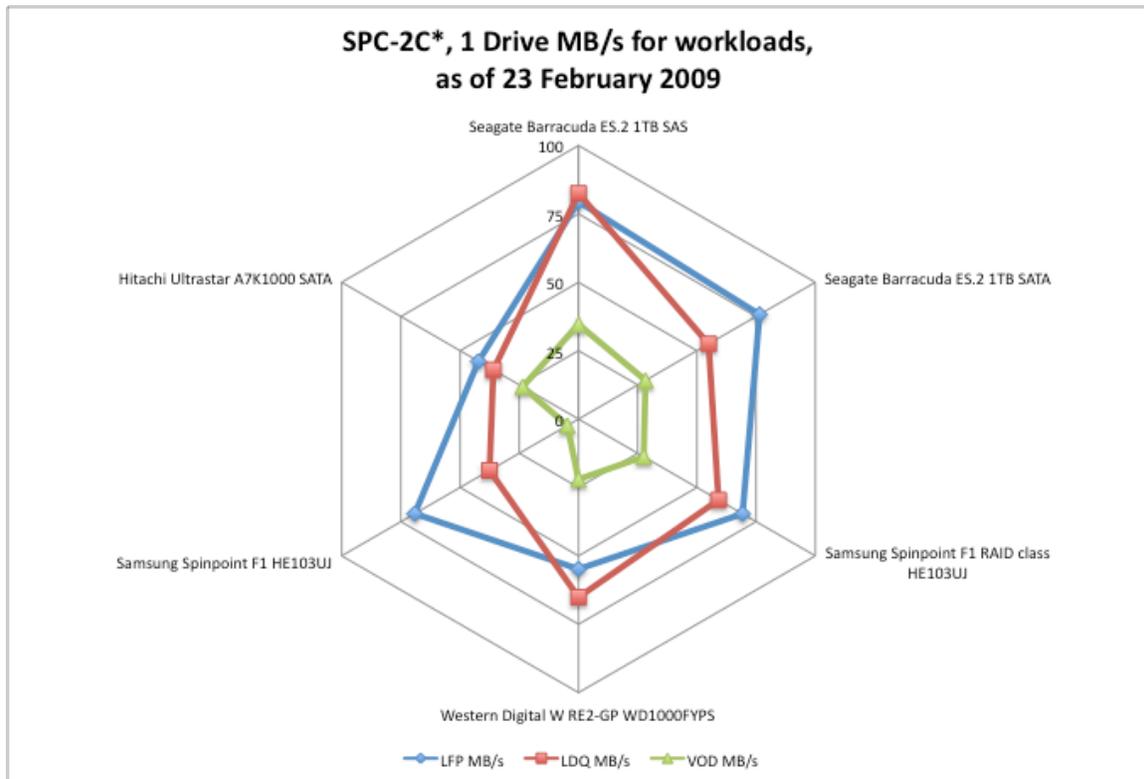


Figure 4 SPC-2C results

Finally, for the new SPC-2C results we show the various workloads in a radar chart format. Here we can obviously see the slowness of the VOD workload vis-à-vis the other workloads. Also there appears to be something amiss with the HDS Ultrastar drive during the LFP workload and the Samsung non-RAID class drive during the VOD workload. Possibly another artifact of varying seek optimization profiles.

Aside from the fact that these are 7200RPM drives, SCI would have assumed this chart would have looked closer to the subsystem SPC-2 chart shown earlier. Obviously, subsystem caching and data striping has an impact and can radically alter subsystem performance in comparison to drive performance on the same workloads. However, the difference between all these drives for the LDQ vs. VOD workloads are significantly different than what was shown on the earlier chart. For example, a maximum of 30% difference in performance between LDQ and VOD workloads on the subsystem chart vs. a minimum of almost 50% difference in performance for the current chart seems disquieting and worthy of more research.

The same caveats apply to SPC-2C that applied to SPC-1C results (see above).

Significance

SCI is always interested in understanding subsystem performance. The new SPC-1C and -2C benchmark results open up a new way to judge storage components rather than storage subsystems. Seagate has taken a shot at 7200RPM drives. It would be even more interesting to show 10K and 15KRPM drives, as well as HBA's.

As much as we like the new SPC-1C and -2C benchmarks, some method needs to be implemented to insure comparable benchmarks are run. It's unclear to SCI how to compare a 12, 15, and 24 drive storage subsystem against a single drive subsystem other than on a per drive basis. Of course this problem also occurs with the SPC-1 and SPC-2 benchmarks as well but at least there we are talking about complete storage subsystems rather than storage components. Perhaps as more multi-drive SPC-1C and -2C results are released we can use the per drive basis as a comparison.

Finally, SPC has once again allowed a sponsor to supply results for other vendor products (Seagate for SPC-1C and -2C results). SCI believes that while this may be expedient, it may hurt the long run value of SPC benchmarks. Because some non-sponsoring vendors feel slighted when not getting an opportunity to review, submit their own benchmark results, and/or reject the results entirely. Unclear how to solve this dilemma, given SPC's stated policy but perhaps allowing the non-SPC member vendors to review and dispute benchmark results ahead of publication may be a necessary first step.

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