SPEC SFS®2014 Performance Report

Silverton Consulting, Inc. StorInt™ Dispatch

This Storage Intelligence (StorInt™) dispatch covers SPEC sfs2014 benchmarks[1]. There were many new submissions this past quarter, mostly from WekaIO Matrix 3.1.9 (VDA, SWBUILD, DATABASE, VDI, & EDA) and NetApp AFF A800 (SWBUILD). As mentioned in our last report, we now have more than 10 submissions for VDA and SWBUILD, and as a result, now only show top 10 rankings for these workloads.

SPEC sfs2014 _vda (video data acquisition) results

SPEC sfs2014_vda workload simulates video data streams. One stream equates to a single camera being recorded. Figure 1 shows the VDA concurrent stream counts. The first metric of interest is number of concurrent (camera) streams supported.

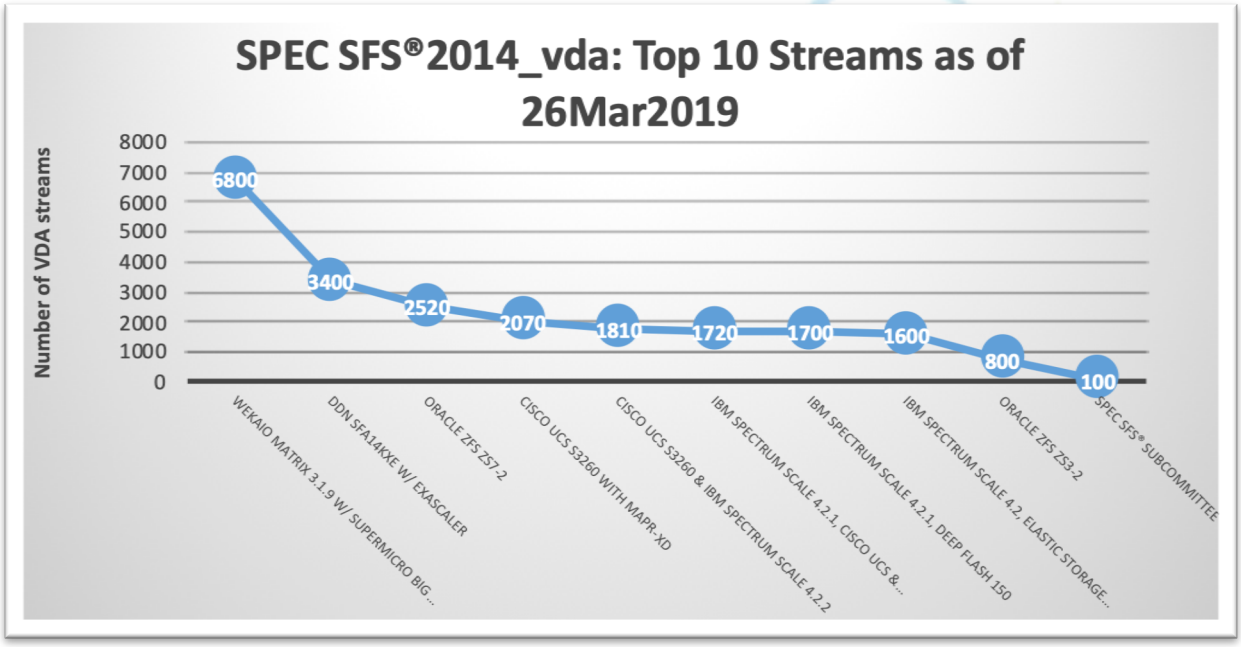

Figure 1 SPECsfs2014_VDA top 10 number of concurrent streams supported

In Figure 1, the new WekaIO Matrix 3.1.9 did extremely well, doubling the previous best from DDN SFA14KXE, with support for 6800 ongoing streams. This WekaIO solution used 138 3.84TB NVMe SSDs, across 23 Supermicro Big Twin storage nodes. Aside from the IBM Spectrum Scale with FlashSystem 900, this is the first time we've seen NVMe or NVMe like SSDs being used on this workload.

[1] All SPEC sfs2014 information is available at https://www.spec.org/sfs2014/ as of 26Mar2019

The WekaIO also used 161 100 GbE ports to support their storage access across their storage cluster and in aggregate had 11.7TB of system memory.

Next, we show the Overall Response Time (ORT) for VDA submissions in Figure 2.

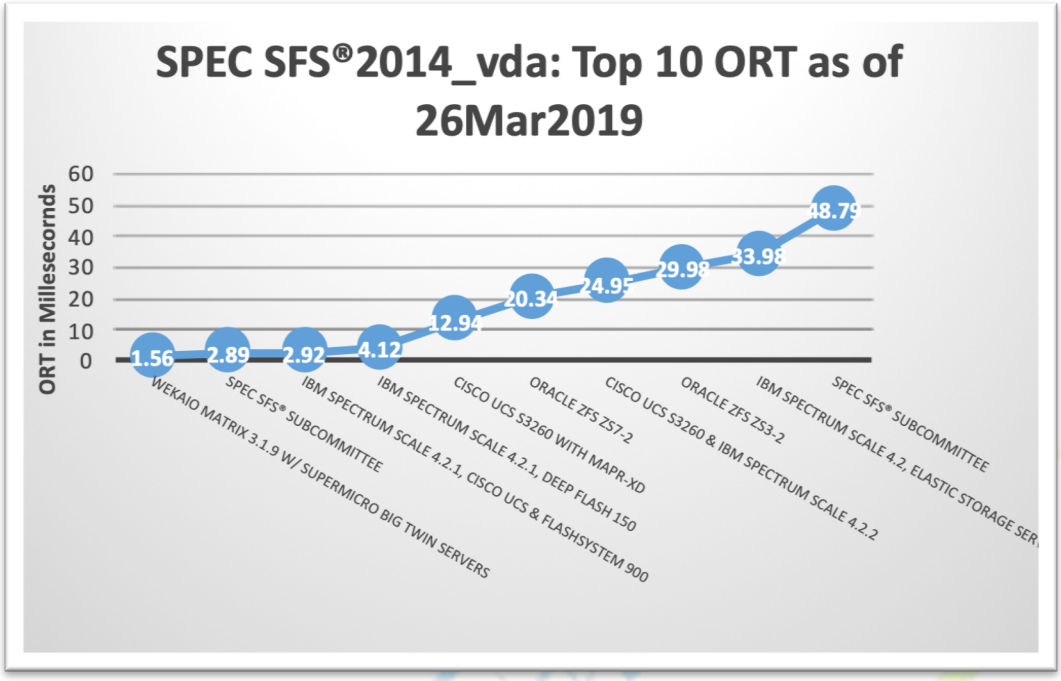

Figure 2 SPEC sfs2014_SWBUILD top 10 minimum (Min RT) and overall response times (ORT)

In Figure 2, the WekaIO submission came in at #1 with an ORT of 1.56 msec. Again, their use of NVMe SSDs probably helped them place well here.

As ORT is an average of response times across all of a submission's workloads. It can depend on how hard the storage system is being driven and has a lot to do with storage media. Disk-only systems can't match the RT of SATA, SAS or NVMe SSD media and systems that are being driven hard tend to have worse ORTs.

SPEC sfs2014_swbuild (software build) results

In Figure 3, we show the top 10 SWBUILD concurrent builds metric

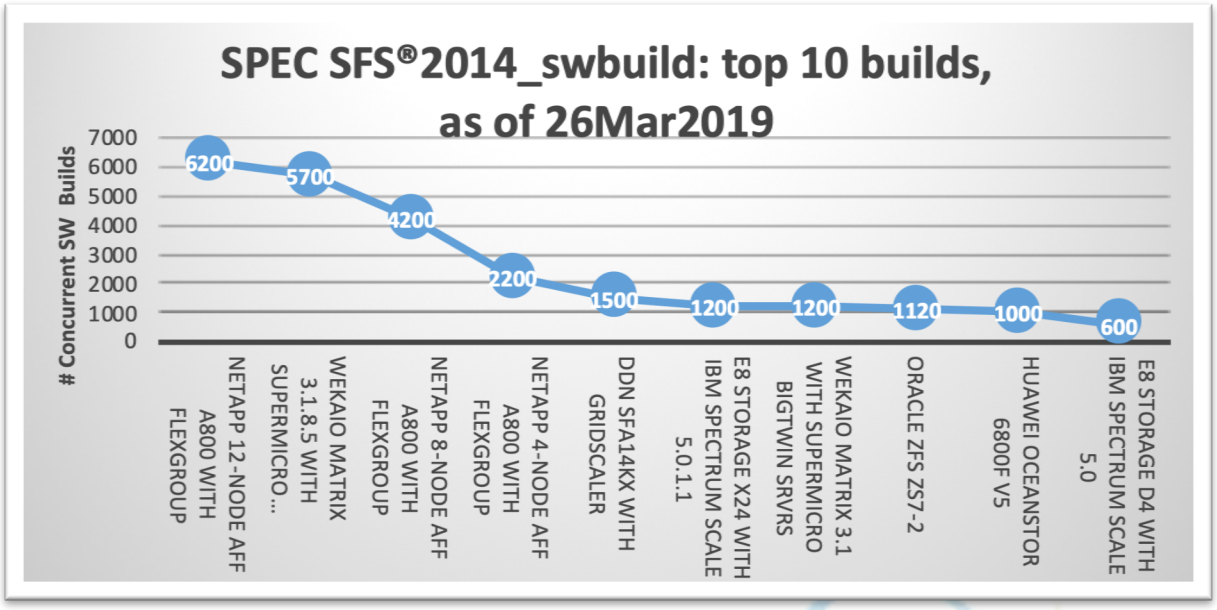

Figure 3 SPEC sfs2014_SWBUILD top 10 concurrent builds

In Figure 3, the two new submissions from NetApp and WekaIO performed well. The new 12 HA (24 controller) node, NetApp AFF A800 submission, used 288 3.84TB NVMe SSDs and came in as our new #1, supporting 6200 builds. The new WekaIO Matrix 3.1.8.5, with 138 3.84TB NVMe SSDS came in at #2 at 5700 builds.

For the NetApp AFF A800, 12 HA (24 controller) nodes cluster is at their max cluster size. There doesn't appear to be any known limit to the WekaIO cluster size, but at 23 nodes, it's approximately the same number as the AFF A800 controller nodes configuration. The NetApp AFF A800 also had 13.1TB of system memory as compared to WekaIO's 12.1TB of system memory.

Next, in Figure 4 we show the Top 10 ORT (ranked and sorted) for the SWBUILD workload.

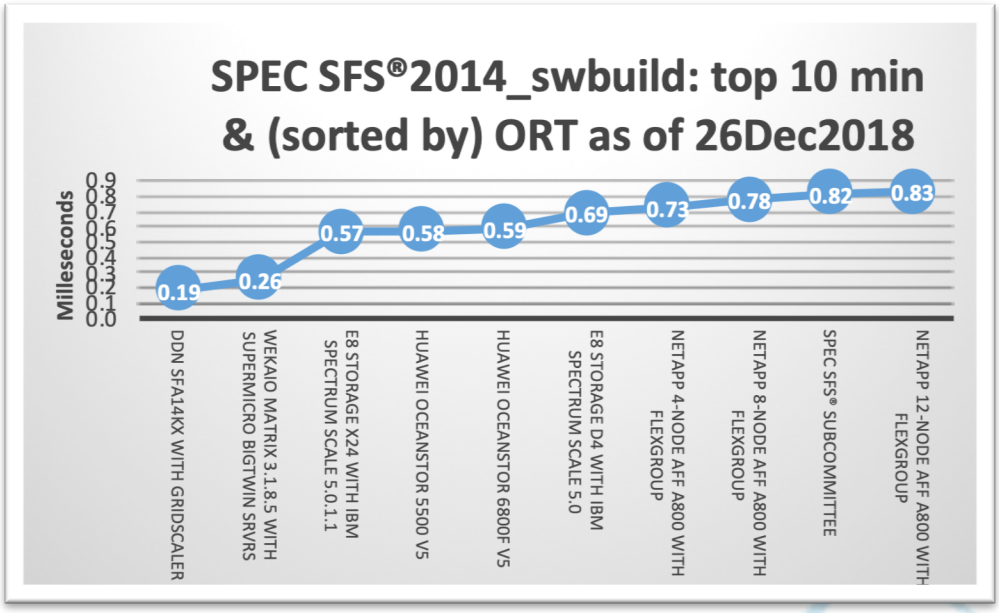

Figure 4 Top 10 SPEC sfs2014_SWBUILD ORT

In Figure 4, the new WekaIO Matrix 3.1.8.5 came in at 2nd place with an ORT of 0.26 msec. The new NetApp AFF A800 came in as our new #1, with an ORT of 0.83 msec. It turns out that the NeApp AFF A800 maximum RT was 2.8 msec., which hurt them as ORT is an average across all benchmark workloads. Both systems came close to one another on minimum RT, at 0.2 and 0.3 msec., for WekaIO and NetApp AFF, respectively. We can only conclude that the NetApp system was being driven quite hard as compared to the WekaIO submission.

In addition, we have never discussed the number of workload generating clients in previous swbuild reports but this number varies considerably between submissions. For example, the NetApp AFF #1 concurrent swbuilds submission used 72 clients (one prime, 71 generating workloads) while the Weka IO #2 concurrent swbuilds submission used 20 (one prime and 19 generating workloads). Thus, the Weka IO solution performed 300 swbuilds/workload client/sec., while the NetApp AFF solution performed ~87 swbuilds/workload client/sec. There are many factors used to determine client server counts but one may well be ORT. Indeed, the number one ORT above, DDN's submission, used 5 workload client nodes and as such, also performed 300 builds/workload client nodes/sec.

SPEC sfs2014_database results

In Figure 5, we present our chart for the SPEC sfs2014_database workload. The DATABASE workload represents a typical transactional workload for a SQL DATABASE. For more information, we described this workload in our previous SPEC sfs2014 performance report, one quarter ago, which should be up on our website (https://silvertonconsulting.com) next month.

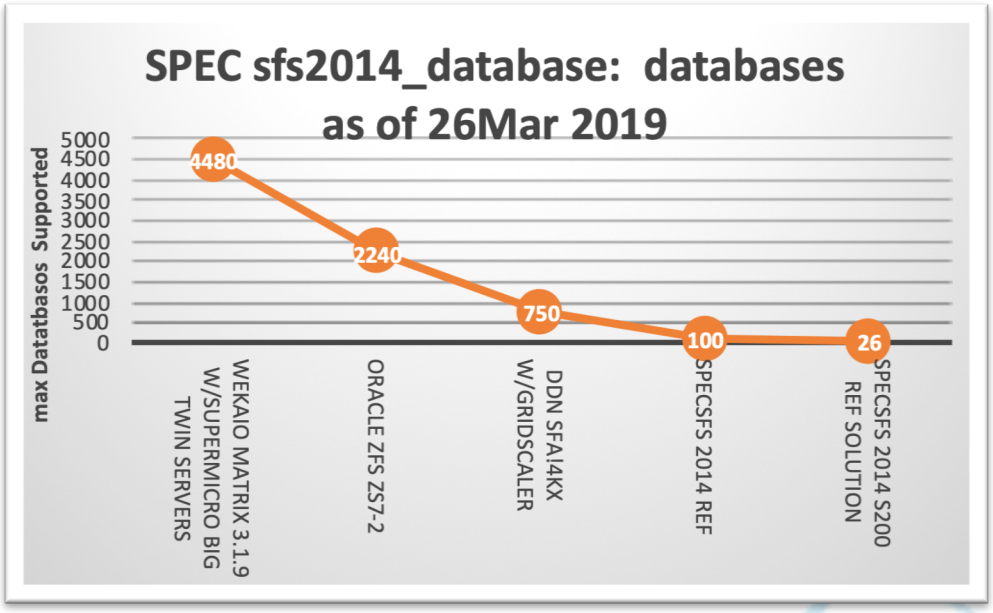

Figure 5 SPEC sfs2014_database concurrent databases supported

In Figure 5, the WekaIO Matrix 3.1.9 solution came in at as our new #1 with 4480 databases supported. All DATABASE submissions but the last used SSDs and again, the only one using NVMe SSDs, was the WekaIO solution with 138 3.84TB NVMe SSDs.

In Figure 6, we show the minimum and maximum response time (RT) on the DATABASE workload.

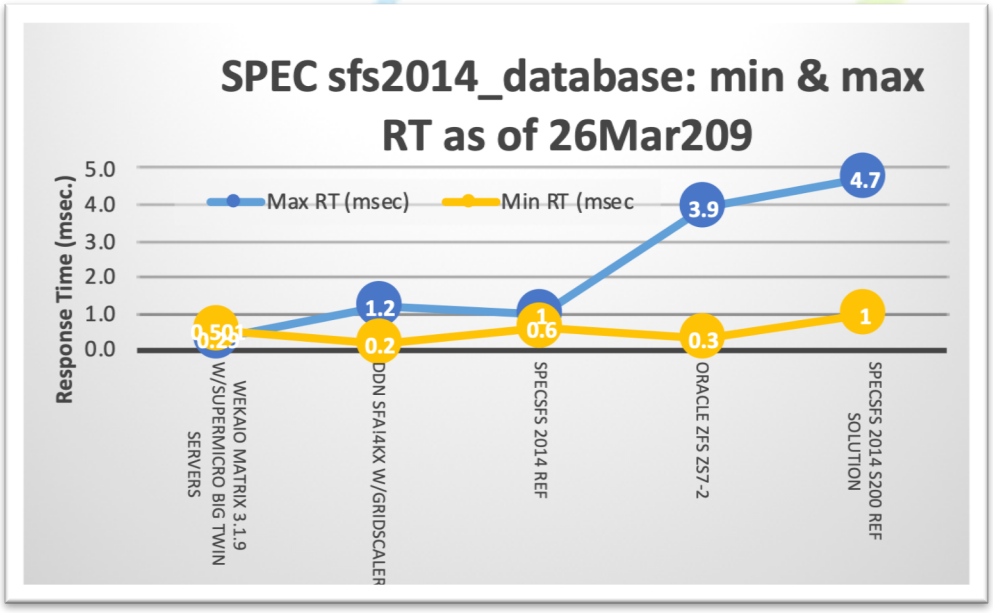

Figure 6 SPEC sfs2014_database Min and Max RT

It's a shame you can barely make out the Min and Max RT for the WekaIO submission, but it had a faster RT at their highest (Max) workload level than at the lowest (Min) workload level.

That is, 0.201 msec. and 0.501 msec. for Max (workload) RT and Min (workload) RT, respectively.

Min RT and Max RT are determined for SPECsfs2014 at the lowest and highest workloads for a submission under test. Normally, the Min RT is faster than the Max RT. This is not the case for WekaIO's submission. In all our years of reporting on IO performance, this is the first time we've seen a reported RT be faster for more work than less.

We asked Liran Zvibel, CEO, WekaIO and others at WekaIO to comment on this. They said there were a couple of factors at work here.

1. WekaIO's (IO) queues get more efficient with more work in them.
2. WekaIO storage supplies more parallelism when there's more IO to be done across a cluster.
3. For all the new WekaIO SPECsfs2014 submissions, they targeted to double the then current highest benchmark metric. So, for some of these benchmarks, the WekaIO submissions weren't breaking a sweat.

WekaIO's reversal of RT speeds was seen on DATABASE, VDI and EDA workloads as well but not for the VDA or SWBUILD workloads. This probably says there's less competition on the SPECsfs2014 DATABASE, VDI and EDA workloads

Significance

We are always eager to see new SPEC sfs2014 submissions, especially when rankings change. And we enjoy learning anything new about IO and storage performance. The quandary presented by WekaIO's RT reversal reminds us of a discussion we once had with IBM about Z/OS work queues which took some effort to startup when there was little work but performed better with more work.

When we see more submissions for EDA and VDI workloads, we will present them as well. At the moment, suffice it to say that the WekaIO submissions are #1 in desktops for VDI and #1 in jobsets for EDA.

It's nice to see NetApp AFF A800 and WekaIO compete on SWBUILD. At the moment, NetApp seems to have an edge, but something tells me this won't stand for long.

We are still trying to determine the best way to report SPEC sfs2014 results. We may experiment with a few variants of the above charts. Any ideas on other metrics of interest to report, please do let us know

Furthermore, suggestions on how to improve any of our performance analyses are always welcomed. Additionally, if you are interested in more file performance details (Top 20 SPEC sfs2008 results) and our NFSv3 and CIFS/SMB (SPEC SFS) ChampionsCharts™, please consider

purchasing our recently updated (March 2019) NAS Buying Guide, available on SCI's website (see QR code below left).

This report was sent out to subscribers as part our **free, monthly Storage Intelligence e-newsletter**. If you are interested in receiving future storage performance analyses along with recent product announcement summaries, please use the QR code (below right) to sign up for your own copy.

Silverton Consulting, Inc., is a U.S.-based Storage, Strategy & Systems consulting firm offering products and services to the data storage community

NAS Buying Guide QRcode

Newsletter signup QRcode