




# DELIVERING WORLD RECORD PERFORMANCE ON DATA ANALYTICS WORKLOADS

*With Supermicro AS -1114S-WN10RT Powered by AMD EPYC™ 7003 Series Processors*

## TABLE OF CONTENTS

- Introduction..... 1
- Test Submissions ..... 2
- System Under Test (SUT)..... 2
- Cluster Configuration ..... 3
- Benchmark Workload ..... 3
- Scale Factors ..... 4
- Measurements..... 4
- Interpreting Results ..... 5
- Results ..... 5
- Conclusions ..... 6
- Footnotes..... 7
- Additional Resources ..... 8



## Introduction

Big Data technologies such as Apache® Hadoop® and Spark are an essential part of the enterprise IT ecosystem. The TPC Express Benchmark™ HS (TPCx-HS) provides an objective measure of hardware, operating system, and commercial Apache Hadoop Filesystem API compatible software distributions. It also delivers the industry's verifiable performance, price-performance, and availability metrics. The benchmark models a continuous system availability of 24 hours a day, seven days a week. TPCx-HS can assess a broad range of system topologies and implementation methodologies in a technically rigorous and directly comparable, and vendor-neutral manner.



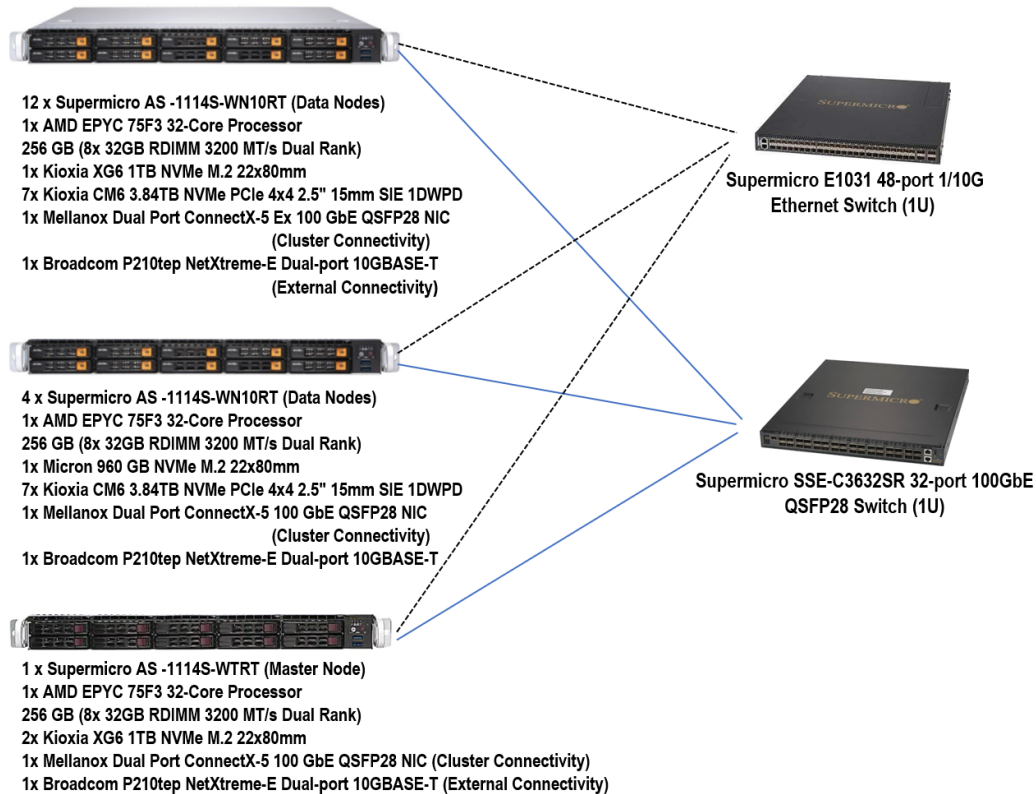
## Test Submissions

### SUPERMICRO

Supermicro (Nasdaq: SMCI), the leading innovator in high-performance, high-efficiency server and storage technology is a premier provider of advanced server Building Block Solutions® for Enterprise Data Center, Cloud Computing, Artificial Intelligence, and Edge Computing Systems worldwide. Supermicro is committed to protecting the environment through its “We Keep IT Green®” initiative and provides customers with the most energy-efficient, environmentally-friendly solutions available on the market..

Supermicro's ongoing benchmarking efforts yielded five TPCx-HS results that set 10 world records<sup>2</sup> with data sizes ranging from 1TB to 100TB on Supermicro WIO servers powered by 3<sup>rd</sup> Gen AMD EPYC™ processors. TPC Express (TPCx-HS) Big Data benchmark categorizes results by data sizes called Scale Factors (SF). This paper showcases Big Data performance results at both the low-end 1 TB SF and the high-end 100 TB SF. In addition, it includes detailed TPCx-HS results and describes their meaning and relevance to Big Data uses from aficionados to engineering, marketing, and sales professionals.

### System Under Test (SUT)



**Figure 1: Supermicro Hadoop cluster (SUT)**

## Cluster Configuration

Hardware					
# of Servers	16x AS-1114S-WN10RT (1P) 1x AS-1114S-WTRT (1P)				
# Cores/Threads	544/ 1,088				
Configuration	<table border="1"> <thead> <tr> <th>DATA NODES</th> <th>MASTER NODE</th> </tr> </thead> <tbody> <tr> <td>           16x Supermicro AS-1114S-WN10RT            1x AMD EPYC™ 75F3 (32 cores)            256 GiB (8x 32GB RDIMM 3200 MT/s dual rank)            1x NVMe (OS):           <ul style="list-style-type: none"> <li>Kioxia XG6 1 TB M.2 22x80mm - used in 12 nodes</li> <li>Micron 960 GB M.2 22x80mm - used in 4 nodes</li> </ul>           7x 3.84 TB NVMe (data):           <ul style="list-style-type: none"> <li>Kioxia CM6 PCIe® 4x4 2.5" 15mm SIE 1DWPD)</li> </ul>           1x 100 GbE (Mellanox ConnectX-5)            1x 10 GbE (Broadcom NIC)         </td> <td>           1x Supermicro AS-1114S-WTRT            1x AMD EPYC 75F3 (32 cores)            256 GiB (8x 32GB RDIMM 3200 MT/s dual rank)            2x 1 TB NVMe:           <ul style="list-style-type: none"> <li>Kioxia XG6 NVMe M.2 22x80mm 1x 100 GbE (Mellanox ConnectX-5)</li> </ul>           1x 100 GbE (Mellanox ConnectX-5)            1x 10 GbE (Broadcom NIC)         </td> </tr> </tbody> </table>	DATA NODES	MASTER NODE	16x Supermicro AS-1114S-WN10RT 1x AMD EPYC™ 75F3 (32 cores) 256 GiB (8x 32GB RDIMM 3200 MT/s dual rank) 1x NVMe (OS): <ul style="list-style-type: none"> <li>Kioxia XG6 1 TB M.2 22x80mm - used in 12 nodes</li> <li>Micron 960 GB M.2 22x80mm - used in 4 nodes</li> </ul> 7x 3.84 TB NVMe (data): <ul style="list-style-type: none"> <li>Kioxia CM6 PCIe® 4x4 2.5" 15mm SIE 1DWPD)</li> </ul> 1x 100 GbE (Mellanox ConnectX-5) 1x 10 GbE (Broadcom NIC)	1x Supermicro AS-1114S-WTRT 1x AMD EPYC 75F3 (32 cores) 256 GiB (8x 32GB RDIMM 3200 MT/s dual rank) 2x 1 TB NVMe: <ul style="list-style-type: none"> <li>Kioxia XG6 NVMe M.2 22x80mm 1x 100 GbE (Mellanox ConnectX-5)</li> </ul> 1x 100 GbE (Mellanox ConnectX-5) 1x 10 GbE (Broadcom NIC)
DATA NODES	MASTER NODE				
16x Supermicro AS-1114S-WN10RT 1x AMD EPYC™ 75F3 (32 cores) 256 GiB (8x 32GB RDIMM 3200 MT/s dual rank) 1x NVMe (OS): <ul style="list-style-type: none"> <li>Kioxia XG6 1 TB M.2 22x80mm - used in 12 nodes</li> <li>Micron 960 GB M.2 22x80mm - used in 4 nodes</li> </ul> 7x 3.84 TB NVMe (data): <ul style="list-style-type: none"> <li>Kioxia CM6 PCIe® 4x4 2.5" 15mm SIE 1DWPD)</li> </ul> 1x 100 GbE (Mellanox ConnectX-5) 1x 10 GbE (Broadcom NIC)	1x Supermicro AS-1114S-WTRT 1x AMD EPYC 75F3 (32 cores) 256 GiB (8x 32GB RDIMM 3200 MT/s dual rank) 2x 1 TB NVMe: <ul style="list-style-type: none"> <li>Kioxia XG6 NVMe M.2 22x80mm 1x 100 GbE (Mellanox ConnectX-5)</li> </ul> 1x 100 GbE (Mellanox ConnectX-5) 1x 10 GbE (Broadcom NIC)				
Connectivity	1x SSE-C3632R 32-port 100 GbE switch 1x E1031 48-port 1/10 GbE switch				
Software					
Operating System	SUSE Linux Enterprise Server 12 SP5				
Apache Hadoop	CDP Private Cloud Base Edition 7.1.6				
Java	OpenJDK 1.8.0_232-cloudera				

**Table 1: Hardware and Software Stack**

## Benchmark Workload

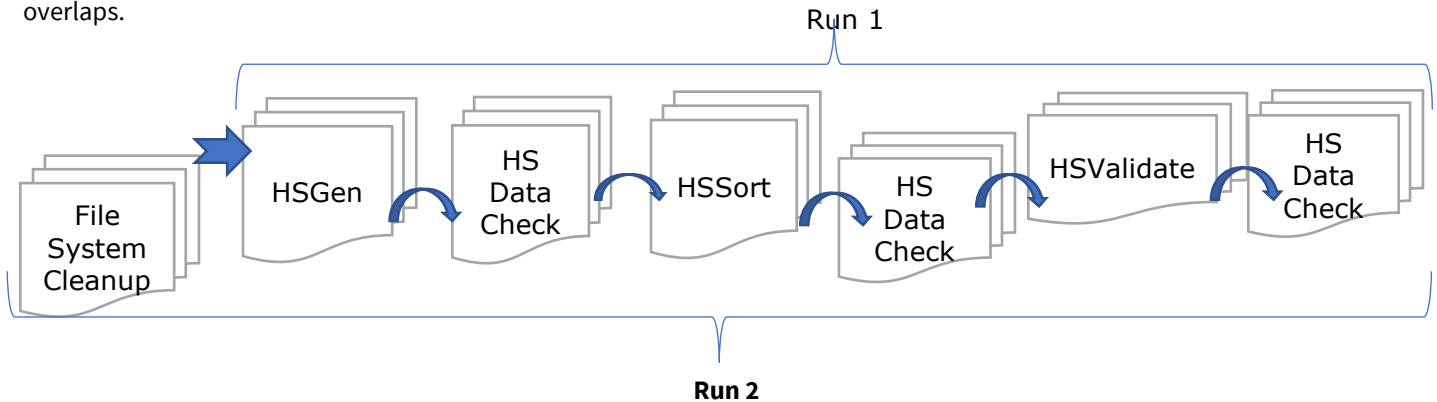
TPCx-HS is the first Big Data industry-standard benchmark based on Apache HDFS API-compatible distributions that stress both hardware and software. TPCx-HS extends the TeraSuite workload definitions (TeraGen, TeraSort, TeraValidate) with formal implementation, execution, metrics, result verification, publication, and pricing rules. As a result, it can assess a broad range of Big Data Hadoop system topologies, implementation methodologies, and systems in a technically rigorous, directly comparable, and vendor-neutral manner. Please visit the TPC Documentation webpage to view the current TPCx-HS specification.

You can download the latest TPCx-HS kit from [www.tpc.org/tpcx-hs](http://www.tpc.org/tpcx-hs). However, if you want to publish a compliant TPC Express result, you must use the TPC-provided kit for your selected benchmark, TPCx-HS. The TPCx-HS benchmark workload consists of the following modules:

- **HSGen:** Generates data at a particular Scale Factor; based on TeraGen.
- **HSDataCheck:** Checks the validity of the dataset and replication.

- **HSSort**: Sorts and orders the data; based on TeraSort.
- **HSValidate**: Validates the sorted output; based on TeraValidate.

The benchmark test occurs in five phases run from a TPCx-HS-master script. These phases must run sequentially without any overlaps.



**Figure 2: TPCx-HS sequence and runs**

The benchmark test consists of Run 1 and Run 2 that follow the run phases shown in Figure 2. No activities are allowed between Run 1 and Run 2 except filesystem cleanup. The TPCx-HS Performance Metric calculation uses the total elapsed runtime  $T$  in seconds. The performance run is defined as the run (either Run 1 or Run 2) with the lower TPCx-HS Performance Metric. The repeatability run is defined as the run (either Run 1 or Run 2) with the higher TPCx-HS Performance Metric. The reported Performance Metric is the TPCx-HS Performance Metric for the performance run.

## Scale Factors

The SF is the dataset size relative to the minimum required size of a test dataset. TPCx-HS requires selecting a test dataset size from the following set of fixed SFs:

1 TB	3 TB	10 TB	30 TB	100 TB	300 TB	1000 TB
------	------	-------	-------	--------	--------	---------

## Measurements

All TPC-published results disclose a Primary Metric that consists of a Performance metric, Price/Performance metric, and an availability date. For TPCx-HS:

1. Performance Metric ( $HSph@SF$ ) reflects the throughput of a run (Run 1 or Run 2) at Scale Factor  $SF$ . This metric is the elapsed time  $T$  for a performance run to perform all five phases shown in Figure 2.

$$HSph@SF = \frac{SF}{T/3600}$$

2. Price/Performance Metric ( $\$/HSph@SF$ ) indicates the Total Cost of Ownership  $P$  needed to own and sustain the SUT that scored the Performance Metric.

$$\$/HSph@SF = \frac{P}{HSph@SF}$$

- System availability date is the day all components used in the Performance test will be available to customers, as defined in the TPC Pricing specification.

## Interpreting Results

In general, faster performance run completion translates to a higher performance score. The score is obtained by normalizing the run times using the above formulas. For Price/Performance, the lower the metric score, the better. In this case, a higher Performance score achieved on a SUT with a lower Total Cost of Ownership P will show a better Price/Performance metric.

## Results

Supermicro published five TPCx-HS results on September 16th, 2021, across multiple scale factors from 1 TB to 100 TB. These results used Supermicro WIO servers powered by AMD EPYC™ 75F3 processors and configured as described in Table 1. Table 2 shows these published results, which set 10 new world records<sup>2</sup> for the Performance and Price/Performance metrics.

Data Size (Scale Factor)	SF 1 TB		SF 3 TB		SF 10 TB		SF 30 TB		SF 100 TB	
	Run1	Run2	Run1	Run2	Run1	Run2	Run1	Run2	Run1	Run2
Number of nodes	17									
Framework	Spark				MapReduce					
HSGen (s)	25.1	24.8	62.365	63.3	174.2	173.0	515.92	517.9	1,624.5	1,637.5
HSSort (s)	79.0	78.6	187.5	187.6	552.9	552.5	1,547.6	1,572.9	5,820.7	5,799.4
HSValidate (s)	15.8	16.4	36.50	37.5	87.1	88.4	256.5	261.4	751.0	747.3
Elapsed run time (s)	131.0	130.0	294.0	296.0	828.0	828.0	2,353.0	2,381.0	8,218.0	8,206.0
<b>TPCx-HS Performance Metric (HSph@SF)</b>	<b>27.54</b>	27.70	36.76	<b>36.49</b>	<b>43.47</b>	<b>43.47</b>	45.89	<b>45.36</b>	<b>43.80</b>	43.87
<b>TPCx-HS Price/Performance (\$/HSph)</b>	<b>\$32,166.53</b>		<b>\$24,276.96</b>		<b>\$20,378.79</b>		<b>\$19,529.68</b>		<b>\$20,225.26</b>	

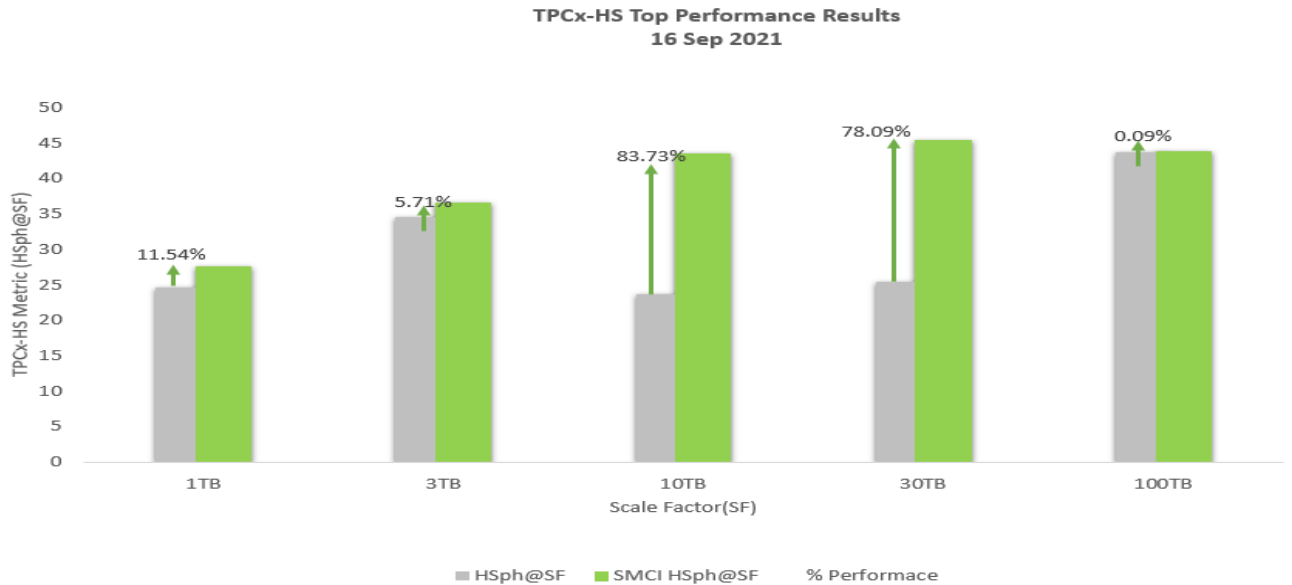
**Table 2: Supermicro TPCx-HS results published on September 16<sup>th</sup>, 2021**

\*\* The lowest of the Run1 & Run2 (shown in bold) is the reported Performance Run, and the other is Repeatability Run

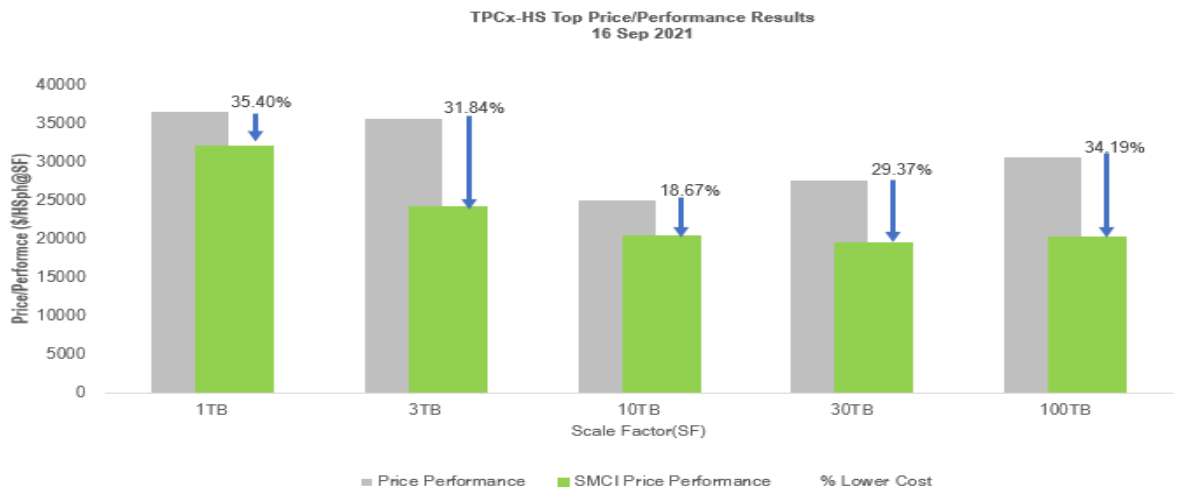
## Conclusions

Single-socket 1U and 2U Supermicro WIO servers powered by 3<sup>rd</sup> Gen AMD EPYC processors support up to 10 PCIe 4.0 NVMe drives to accommodate ever-growing Big Data deployment storage requirements. Supermicro WIO servers deliver the following benefits compared to both other OEMs and systems powered by previous-generation AMD EPYC processors:

- Consistently superior performance across multiple scale factors
- Up to 83% Performance metric uplift
- Up to 35% lower Total Cost of Ownership (TCO) based on Price/Performance metrics.



**Figure 3: Performance gain over a previous generation or competition<sup>3,4,5,6,7</sup>**



**Figure 4: Price/Performance (TCO) savings over the previous generation or competition<sup>3,4,5,6,7</sup>**

Supermicro WIO servers powered by AMD EPYC 7003 Series Processors help leading enterprises reduce time-to-solution across a wide range of applications, provide enhanced security features, and allow running all workloads either on-premises or in a public or private cloud.

Supermicro offers many certified solutions and reference architectures that empower organizations to create deployments that deliver data insights faster than ever before. These solutions include servers optimized for:

- AI/ML/DL training inference
- Hyperconverged infrastructure (HCI)
- Software-defined infrastructure (SDI)
- Software-defined storage, such as CEPH, VMWare vSAN, and Weka.IO.
- Data management, such as Oracle 19c, Apache Hadoop, and Cassandra.
- HPC application optimization, such as Ansys® Fluent®, OpenFOAM®, and WRF.

## Footnotes

1. *These results held world record performance prior to new results from Supermicro: Dell R6515 (AMD EPYC 75F3 processor) cluster HSph@1TB 24.69: <http://tpc.org/5551> HSph@3TB 34.52: <http://tpc.org/5548> HSph@100TB 43.76: <http://tpc.org/5552>; HPE DL325 (AMD EPYC 7502P) cluster: HSph@10TB 23.66: <http://tpc.org/5532>; HSph@30TB 25.47 <http://tpc.org/5533> ;*
2. *Supermicro WIO cluster 3<sup>rd</sup> Gen AMD EPYC 75F3 17 Nodes, 17 processors, 544 cores 27.54 HSph@1TB, 32,166.53 USD per HSph@1TB, 2 world records in price and price-performance for 1TB SF <http://tpc.org/5553> ; 3rd Gen AMD EPYC 75F3 17 Nodes, 17 processors, 544 cores 36.49 HSph@3TB, 24,276.96 USD per HSph@3TB, 2 world records in price and price-performance for 3TB SF <http://tpc.org/5554>; AMD EPYC 75F3 17 Nodes, 17 processors, 544 cores 43.47 HSph@10TB, 20,378.79 USD per HSph@10TB, 2 world records in price and price-performance for 1-TB SF <http://tpc.org/5555>; AMD EPYC 75F3 17 Nodes, 17 processors, 544 cores 45.36 HSph@30TB, 19,529.68 USD per HSph@30TB, 2 world records in price and price-performance for 30TB SF <http://tpc.org/5556> ; AMD EPYC 75F3 17 Nodes, 17 processors, 544 cores 43.80 HSph@100TB, 20,225.26 USD per HSph@100TB, 2 world records in price and price-performance for 100TB SF <http://tpc.org/5557> ; see also: <https://www.amd.com/en/processors/epyc-world-records>*
3. *Supermicro WIO cluster 3<sup>rd</sup> Gen AMD EPYC 75F3 17 Nodes, 17 processors, 544 cores 27.54 HSph@1TB, 32,166.53 USD per HSph@1TB, <http://tpc.org/5553> ; Dell R6515 (AMD EPYC 75F3 processor) cluster HSph@1TB 24.69 49,795.35 USD per HSph@1TB <http://tpc.org/5551> ;  $27.54/24.69=11.54\%$ ,  $\$32166.53/\$49795.35=-35.40\%$*
4. *Supermicro WIO cluster 3<sup>rd</sup> Gen AMD EPYC 75F3 17 Nodes, 17 processors, 544 cores 36.49 HSph@3TB, 24,276.96 USD per HSph@3TB, <http://tpc.org/5554> ; Dell R6515 (AMD EPYC 75F3 processor) cluster HSph@3TB 34.52, 35,615.50 USD per HSph@3TB <http://tpc.org/5548> ;  $36.49/34.52=5.71\%$ ,  $\$24276.96/\$35615.5=-31.84\%$*
5. *Supermicro WIO cluster 3<sup>rd</sup> Gen AMD EPYC 75F3 17 Nodes, 17 processors, 544 cores 43.47 HSph@10TB, 20,378.79 USD per HSph@10TB, <http://tpc.org/5555> . HPE DL325 (AMD EPYC 7502P) HSph@10TB 23.66 USD per HSph@10TB 25,057.91 <http://tpc.org/5532>;  $43.47/23.66=83.73\%$ ,  $\$20378.79/\$25057.91=-18.67\%$*

6. Supermicro WIO cluster 3rd Gen AMD EPYC 75F3 17 Nodes, 17 processors, 544 cores 45.36 HSph@30TB, 19,529.68 USD per HSph@30TB, <http://tpc.org/5556> . HPE DL325 (AMD EPYC 7502P) HSph@30TB 25.47 USD per HSph@30TB 27,649.40 <http://tpc.org/5533>;  $45.36/25.47=78.09\%$ ,  $19529.68/27649.4=-29.37\%$
7. Supermicro WIO cluster 3rd Gen AMD EPYC 75F3 17 Nodes, 17 processors, 544 cores 43.80 HSph@100TB, 20,225.26 USD per HSph@100TB, <http://tpc.org/5557> [tpc.org]; Dell R6515 (AMD EPYC 75F3 processor) cluster HSph@100TB 43.76 USD per HSph@100TB 30,732.52 <http://tpc.org/5552> [tpc.org];  $43.8/43.76=0.09\%$  ;  $20225.26/30732.52=-34.19\%$

## Additional Resources

[TPCx-HS is a Big Data System Benchmark](#)

[TPCx-HS Top Performance Results](#)

[TPCx-HS Top Price/Performance Results](#)

AMD, the AMD logo, EPYC, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.