



Super Micro Computer, Inc.

TPC Express Benchmark™ HS Full Disclosure Report

Supermicro Cluster

(with 16x AS-1114S-WN10RT Servers; 1x AS-1114S-WTRT Servers)

Running

CDP Private Cloud Base Edition 7.1.6

on

SUSE Linux Enterprise Server 12 SP5

TPCx-HS Version
Report Edition
Report Submitted

2.0.3
First
September 16, 2021

First Edition - September 2021

Super Micro Computer, Inc. (Supermicro), the Sponsor of this benchmark test, believes that the information in this document is accurate as of the publication date. The information in this document is subject to change without notice. The Sponsor assumes no responsibility for any errors that may appear in this document.

The pricing information in this document is believed to accurately reflect the current prices as of the publication date. However, the Sponsor provides no warranty of the pricing information in this document.

Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary because of these and other factors. Therefore, the TPC Express Benchmark™ V should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or product evaluation decisions are contemplated.

All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly. No warranty of system performance or price/performance is expressed or implied in this report.

Supermicro and the Supermicro Logo are trademarks of Super Micro Computer, Inc. and/or its affiliates in the U.S. and other countries. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Supermicro and any other company

TPC Express Benchmark™ HS, TPCx-HS, and HSph, are registered certification marks of the Transaction Processing Performance Council.

The Supermicro products, services or features identified in this document may not yet be available or may not be available in all areas and may be subject to change without notice. Consult your local Supermicro business contact for information on the products or services available in your area. You can find additional information via Supermicro's web site at <https://www.supermicro.com/en/>. Actual performance and environmental costs of Supermicro products will vary depending on individual customer configurations and conditions.

Copyright© 2021 Super Micro Computer, Inc.

All rights reserved. Permission is hereby granted to reproduce this document in whole or in part provided the copyright notice printed above is set forth in full text or on the title page of each item reproduced.

Abstract

This document contains the methodology and results of the TPC Express Benchmark™ HS (TPCx-HS) test conducted in conformance with the requirements of the TPCx-HS Standard Specification, Revision 2.0.3.


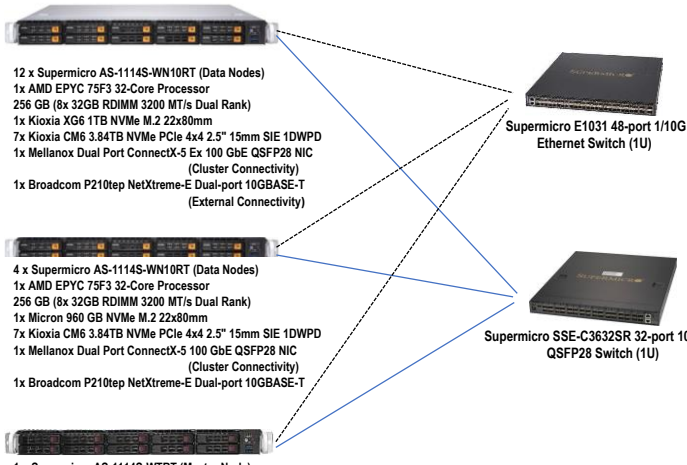
The benchmark results are summarized below.


Measured Configuration			
Company Name	Cluster Node	Hadoop Software	Operating System
Supermicro	AS-1114S-WN10RT	CDP Private Cloud Base Edition 7.1.6	SUSE Linux Enterprise Server 12 SP5

TPC Express Benchmark™ HS Metrics			
Total System Cost	HSph@30TB	Price/Performance	Availability Date
\$885,866	45.36	\$19,529.68	Currently Available


Executive Summary


The [Executive Summary](#) follows on the next several pages.

		<h1>Supermicro Cluster</h1>		TPCx-HS 2.0.3 TPC Pricing 2.7.0 Report Date Sep. 16, 2021
Availability Date Currently Available	TPCx-HS Performance 45.36 HSph@30TB	Price/Performance \$19,529.68 \$ / HSph@30TB	Total System Cost \$885,866 USD	
System Under Test Configuration Overview				
Scale Factor 30	Hadoop Software CDP Private Cloud Base Edition 7.1.6	Operating System SUSE Linux Enterprise Server 12 SP5	Other Software None	
<div style="display: flex; justify-content: space-around;"> <div style="width: 45%;"> <p>12 x Supermicro AS-1114S-WN10RT (Data Nodes) 1x AMD EPYC 75F3 32-Core Processor 256 GB (8x 32GB RDIMM 3200 MT/s Dual Rank) 1x Kioxia XG6 1TB NVMe M.2 22x80mm 7x Kioxia CM6 3.84TB NVMe PCIe 4x4 2.5" 15mm SIE 1DWPD 1x Mellanox Dual Port ConnectX-5 Ex 100 GbE QSFP28 NIC (Cluster Connectivity) 1x Broadcom P210tep NetXtreme-E Dual-port 10GBASE-T (External Connectivity)</p> <p>4 x Supermicro AS-1114S-WN10RT (Data Nodes) 1x AMD EPYC 75F3 32-Core Processor 256 GB (8x 32GB RDIMM 3200 MT/s Dual Rank) 1x Micron 960 GB NVMe M.2 22x80mm 7x Kioxia CM6 3.84TB NVMe PCIe 4x4 2.5" 15mm SIE 1DWPD 1x Mellanox Dual Port ConnectX-5 100 GbE QSFP28 NIC (Cluster Connectivity) 1x Broadcom P210tep NetXtreme-E Dual-port 10GBASE-T</p> <p>1 x Supermicro AS-1114S-WTRT (Master Node) 1x AMD EPYC 75F3 32-Core Processor 256 GB (8x 32GB RDIMM 3200 MT/s Dual Rank) 2x Kioxia XG6 1TB NVMe M.2 22x80mm 1x Mellanox Dual Port ConnectX-5 100 GbE QSFP28 NIC (Cluster Connectivity) 1x Broadcom P210tep NetXtreme-E Dual-port 10GBASE-T (External Connectivity)</p> </div> <div style="width: 45%;"> <p>Supermicro E1031 48-port 1/10G Ethernet Switch (1U)</p> <p>Supermicro SSE-C3632SR 32-port 100GbE QSFP28 Switch (1U)</p> </div> </div> 				
Physical Storage/Scale Factor: 14.93		Scale Factor/Physical Memory: 7.06		
Total Number of Servers: Total Processors/Cores/Threads:		17 (16x AS-1114S-WN10RT; 1x AS-1114S-WTRT) 17/544/1,088		
Server Configuration: Processors Memory Storage Device Network	Per AS-1114S-WN10RT 1x AMD EPYC 75F3 256 GiB 1x 1 TB NVMe (12 nodes) 1x 960 GB NVMe (4 nodes) 7x 3.84 TB NVMe (all nodes) 1x Mellanox Dual-port ConnectX-5 Ex 100 GbE (12 nodes) 1x Mellanox Dual-port ConnectX-5 100 GbE (4 nodes) 1x Broadcom Dual-port 10 GbE (all nodes)	Per AS-1114S-WTRT 1x AMD EPYC 75F3 256 GiB 2x 1 TB NVMe 1x Mellanox Dual-port 100 GbE 1x Broadcom Dual-port 10 GbE		
Connectivity: Total Rack Units:	1x SSE-C3632R 32-port 100 GbE; 1x E1031 48-port 1/10 GbE 16x(1U)+1x(1U)+1x(1U)+1x(1U) = 16U+1U+1U+1U = 19U			

		<h1>Supermicro Cluster</h1>		TPCx-HS	2.0.3	
				TPC Pricing	2.7.0	
				Report Date	Sep. 16, 2021	
Description	Part Number	Source	Qty	Unit Price	Extended Price	3 Yr. Maint. Price
HARDWARE						
Data Nodes						
H12SSW-NTR, CSE-116TS-R706WBPS-N10,RoHS	AS-1114S-WN10RT		1	16	\$1,477.00	\$23,632.00
32GB DDR4-3200 2Rx4 ECC REG DIMM	MEM-DR432L-HL01-ER32		1	128	\$184.60	\$23,628.80
Kioxia CM6 3.84TB NVMe PCIe 4x4 2.5" 15mm SIE 1DWPD	HDS-TUN-KCM6XRUL3T84		1	112	\$859.00	\$96,208.00
Mellanox ConnectX-5 EN network card 100GbE dual-port	AOC-MCX516A-CDAT		1	12	\$1,060.00	\$12,720.00
MCX516A-CCAT PCIe 2-port 100GbE QSFP28 Gen3.0 x16 CX-5	AOC-MCX516A-CCAT		1	4	\$849.00	\$3,396.00
Milan 75F3 DP/UP 32C/64T 2.95G 256M 280W SP3	PSE-MLN75F3-0313		1	16	\$4,834.00	\$77,344.00
Kioxia XG6 1TB NVMe M.2 22x80mm	HDS-TMNO-KXG60ZNV1T02		1	12	\$190.00	\$2,280.00
Micron 7300 PRO 960GB,PCIe NVMe,M.2 22x80mm,3D TLC,1DWPD	HDS-MMN-MTFDHBA960TDF1AW		1	4	\$178.50	\$714.00
Out of Band Firmware Management License-BIOS Flash /Setting	SFT-OOB-LIC		1	16	\$15.00	\$240.00
ASSEMBLY FEE	MC0037		1	16	\$25.00	\$400.00
0% 3 YRS LABOR, 3 YRS PARTS, 1 YR CRS UNDER LIMITED WRNTY	EWCS		1	16	(included)	(included)
On Site 4hrs 24x7x365 Support 3 Years with Extended Wrnty	OS4HR3		1	16	\$516.28	\$8,260.42
Master Node						
H12SSW-NT, CSV-116TS-R504WBP	AS-1114S-WTRT		1	1	\$1,304.00	\$1,304.00
32GB DDR4-3200 2Rx4 ECC REG DIMM	MEM-DR432L-HL01-ER32		1	8	\$184.60	\$1,476.80
Kioxia XG6 1TB NVMe M.2 22x80mm	HDS-TMNO-KXG60ZNV1T02		1	2	\$190.00	\$380.00
Mellanox ConnectX-5 EN network card 100GbE dual-port	AOC-MCX516A-CDAT		1	1	\$1,060.00	\$1,060.00
Milan 75F3 DP/UP 32C/64T 2.95G 256M 280W SP3	PSE-MLN75F3-0313		1	1	\$4,834.00	\$4,834.00
Out of Band Firmware Management License-BIOS Flash /Setting	SFT-OOB-LIC		1	1	\$15.00	\$15.00
ASSEMBLY FEE	MC0037		1	1	\$25.00	\$25.00
0% 3 YRS LABOR, 3 YRS PARTS, 1 YR CRS UNDER LIMITED WRNTY	EWCS		1	1	(included)	(included)
On Site 4hrs 24x7x365 Support 3 Years with Extended Wrnty	OS4HR3		1	1	\$454.36	\$454.36
Network and Cables						
E1031 48-port 1/10G Ethernet ToR switch	SSE-G3648BR		1	1	\$1,675.00	\$1,675.00
Cumulus-Linux SW 1G perpetual license with 3 yr Cumulus	SFT-CLSP1G-3Y		1	1	\$1,475.00	\$1,475.00
On Site 4hrs 24x7x365 Support 3 Years with Extended Wrnty	OS4HR3		1	1	\$315.00	\$315.00
32-port 100GbE QSFP28,B2F,2x800W R0872-F0004-01,HF	SSE-C3632SR		1	1	\$7,375.00	\$7,375.00
Cumulus-Linux Software 100G Perpetual License with 3 yr SnS	SFT-CLSNWPL-100G-3Y		1	1	\$6,399.00	\$6,399.00
On Site 4hrs 24x7x365 Support 3 Years with Extended Wrnty	OS4HR3		1	1	\$1,377.40	\$1,377.40
ETHERNET, QSFP28, 100GbE, PASSIVE, LSZH, 3m, Molex, RoHS	CBL-NTWK-0943-SQ28C30M		1	17	\$139.50	\$2,371.50
ETHERNET, CAT6, RJ45, SNAGLESS, YELLOW, 15FT (4.6M), 28AWG, Ro	CBL-C6-YL15FT-P		1	17	\$10.80	\$183.60
ETHERNET, CAT6, RJ45, SNAGLESS, GREEN, UTP, 15FT(4.5M), 28AWG, RoHS	CBL-C6-GN15FT-P		1	17	\$10.80	\$183.60
Infrastructure						
42U Enclosure system	SRK-42SE-11		1	1	\$1,516.30	\$1,516.30
Rack PDU, Switched, 2U, 30A, 208V, (16)C13	AP7911B		2	3	\$1,025.00	\$3,075.00
PWCD, US, IEC60320 C14 TO C13, 4FT, 16AWG, RoHS/REACH	CBL-PWCD-0373-IS		1	38	\$6.50	\$247.00
LONCEVON - 12 inch IPS 1920x1080p HDMI Monitor	N/A		3	3	\$99.99	\$299.97
Logitech MK200 Media Keyboard and Mouse Combo	920-002714		3	3	\$41.50	\$124.50
Spares, Accessories						
ETHERNET, QSFP28, 100GbE, PASSIVE, LSZH, 3m, Molex, RoHS	CBL-NTWK-0943-SQ28C30M		1	3	\$139.50	\$418.50
ETHERNET, CAT6, RJ45, SNAGLESS, YELLOW, 15FT (4.6M), 28AWG, Ro	CBL-C6-YL15FT-P		1	3	\$10.80	\$32.40
ETHERNET, CAT6, RJ45, SNAGLESS, GREEN, UTP, 15FT(4.5M), 28AWG, RoHS	CBL-C6-GN15FT-P		1	3	\$10.80	\$32.40
PWCD, US, IEC60320 C14 TO C13, 4FT, 16AWG, RoHS/REACH	CBL-PWCD-0373-IS		1	3	\$6.50	\$19.50
HARDWARE Subtotals					\$275,085.87	\$10,407.17

(continued next page)

	<h1>Supermicro Cluster</h1>	TPCx-HS 2.0.3																																		
		TPC Pricing 2.7.0																																		
		Report Date Sep. 16, 2021																																		
(continued from previous page)																																				
<p>SOFTWARE</p> <p>SUSE Linux Enterprise Server, x86 & x86-64, 1-2 Sockets or 1-2 Virtual Machines, Priority Subscription, 3 Year 874-006883</p> <p>Cloudera Data Platform Private Cloud Base Edition - Annual Subscription per Node for up to 16 Cores/128 GB RAM for compute and up to 48 TB for storage. BusinessLevel Support.</p> <p>COMPUTE: price per CCU per year for compute in excess of 16 cores/128GB RAM per Node, where 1 CCU = 1 core + 8 GB RAM</p> <p>SOFTWARE Subtotals</p>	<table border="1"> <thead> <tr> <th>Description</th> <th>Part Number</th> <th>Source</th> <th>Qty</th> <th>Unit Price</th> <th>Extended Price</th> <th>3 Yr. Maint. Price</th> </tr> </thead> <tbody> <tr> <td>SFT-NV-SU2P3YBAC</td> <td></td> <td></td> <td>1 17</td> <td>\$2,916.00</td> <td>\$49,572.00</td> <td></td> </tr> <tr> <td>SMC-CDP-PVCBASE-BUS</td> <td></td> <td></td> <td>1 51</td> <td>\$9,600.00</td> <td>\$489,600.00</td> <td></td> </tr> <tr> <td>SMC-CDP-COMPUTE</td> <td></td> <td></td> <td>1 816</td> <td>75</td> <td>\$61,200.00</td> <td></td> </tr> <tr> <td colspan="5"></td> <td>\$600,372.00</td> <td>\$0.00</td> </tr> </tbody> </table>	Description	Part Number	Source	Qty	Unit Price	Extended Price	3 Yr. Maint. Price	SFT-NV-SU2P3YBAC			1 17	\$2,916.00	\$49,572.00		SMC-CDP-PVCBASE-BUS			1 51	\$9,600.00	\$489,600.00		SMC-CDP-COMPUTE			1 816	75	\$61,200.00							\$600,372.00	\$0.00
Description	Part Number	Source	Qty	Unit Price	Extended Price	3 Yr. Maint. Price																														
SFT-NV-SU2P3YBAC			1 17	\$2,916.00	\$49,572.00																															
SMC-CDP-PVCBASE-BUS			1 51	\$9,600.00	\$489,600.00																															
SMC-CDP-COMPUTE			1 816	75	\$61,200.00																															
					\$600,372.00	\$0.00																														
<p>Pricing: 1 = Supermicro; 2 = APC; 3 = Amazon</p> <p>* Discount applies to all line items where Key = 1. Discount based upon total system cost as purchased by a regular customer.</p> <p>Audited by Doug Johnson, InfoSizing</p>	<p>Three-Year Cost of Ownership: \$885,866</p> <p>HSph@30TB: 45.36</p> <p>\$ / HSph@30TB: \$19,529.68</p>																																			
<p><i>Prices used in TPC benchmarks reflect the actual prices a customer would pay for a one-time purchase of the stated Line Items. Individually negotiated discounts are not permitted. Special prices based on assumptions about past or future purchases are not permitted. All discounts reflect standard pricing policies for the listed Line Items. For complete details, see the pricing section of the TPC Benchmark Standard. If you find that the stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.</i></p>																																				

	<h2>Supermicro Cluster</h2>	TPCx-HS 2.0.3
		TPC Pricing 2.7.0
		Report Date Sep. 16, 2021

Numerical Quantities	
----------------------	--

Performance Run – Run 2	
Scale Factor	30TB
Run Start Time	2021-08-25 08:25:01.000
Run End Time	2021-08-25 09:04:39.000
Run Elapsed Time	2,381.000
HSGen Start Time	2021-08-25 08:25:02.000
HSGen End Time	2021-08-25 08:33:38.000
HSGen Elapsed Time	517.915
HSSort Start Time	2021-08-25 08:33:53.000
HSSort End Time	2021-08-25 09:00:05.000
HSSort Elapsed Time	1,572.924
HSValidate Start Time	2021-08-25 09:00:19.000
HSValidate End Time	2021-08-25 09:04:39.000
HSValidate Elapsed Time	261.409
Repeatability Run – Run 1	
Scale Factor	30TB
Run Start Time	2021-08-25 07:19:05.000
Run End Time	2021-08-25 07:58:15.000
Run Elapsed Time	2,353.000
HSGen Start Time	2021-08-25 07:19:06.000
HSGen End Time	2021-08-25 07:27:40.000
HSGen Elapsed Time	515.927
HSSort Start Time	2021-08-25 07:27:58.000
HSSort End Time	2021-08-25 07:53:44.000
HSSort Elapsed Time	1,547.622
HSValidate Start Time	2021-08-25 07:54:00.000
HSValidate End Time	2021-08-25 07:58:15.000
HSValidate Elapsed Time	256.541



Supermicro Cluster

TPCx-HS	2.0.3
TPC Pricing	2.7.0
Report Date	Sep. 16, 2021

Run Reports

Run Report for Performance Run – Run 2

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 2 Details	Total Time =	2381
	Total Size =	300000000000
	Scale-Factor =	30

TPCx-HS Performance Metric (HSph@SF): 45.3651

Run Report for Repeatability Run – Run 1

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 1 Details	Total Time =	2353
	Total Size =	300000000000
	Scale-Factor =	30

TPCx-HS Performance Metric (HSph@SF): 45.8996


	<h1>Supermicro Cluster</h1>	TPCx-HS 2.0.3
		TPC Pricing 2.7.0
		Report Date Sep. 16, 2021
<h2>Revision History</h2>		
Date	Edition	Description
September 16, 2021	First	Initial Publication

Table of Contents

Abstract..... 3

Executive Summary 3

Table of Contents.....10

Clause 0 – Preamble11

 0.1 TPC Express Benchmark™ HS Overview.....11

Clause 1 – General Items12

 1.1 Test Sponsor12

 1.2 Parameter Settings12

 1.3 Configuration Diagrams12

 1.3.1 Measured Configuration.....13

 1.3.2 Priced Configuration13

 1.4 Dataset Distribution.....14

 1.5 Software Components Distribution.....14

Clause 2 – Workload Related Items.....15

 2.1 Hardware & Software Tunables15

 2.2 Run Report15

 2.3 Benchmark Kit Identification.....15

 2.4 Benchmark Kit Changes16

Clause 3 – SUT Related Items.....17

 3.1 Data Storage Ratio17

 3.2 Memory Ratio17

Clause 4 – Metrics Related Items18

 4.1 HSGen Time.....18

 4.2 HSSort Time18

 4.3 HSValidate Time18

 4.4 HSDataCheck Times18

 4.5 Performance & Price-Performance18

Auditor’s Information & Letter of Attestation19

Supporting Files Index22

Third-Party Price Quotes.....23

 APC.....23

 Amazon.....24

Clause 0 – Preamble

0.1 TPC Express Benchmark™ HS Overview

The TPC Express Benchmark™ HS (TPCx-HS) was developed to provide an objective measure of hardware, operating system and commercial Apache Hadoop File System API compatible software distributions, and to provide the industry with verifiable performance, price-performance and availability metrics. The benchmark models a continuous system availability of 24 hours a day, 7 days a week.

Even though the modeled application is simple, the results are highly relevant to hardware and software dealing with Big Data systems in general. TPCx-HS stresses both hardware and software including Hadoop run-time, Hadoop File-system API compatible systems and MapReduce layers. This workload can be used to assess a broad range of system topologies and implementation of Hadoop clusters. TPCx-HS can be used to assess a broad range of system topologies and implementation methodologies in a technically rigorous and directly comparable and vendor-neutral manner.

The TPCx-HS kit is available from the TPC (See www.tpc.org/tpcx-hs for more information). Users must sign-up and agree to the TPCx-HS User Licensing Agreement (ULA) to download the kit. Re-distribution of the kit is prohibited. All related work (such as collaterals, papers, derivatives) must acknowledge the TPC and include TPCx-HS copyright. The TPCx-HS Kit includes: TPCx-HS Specification document, TPCx-HS Users Guide documentation, shell scripts to set up the benchmark environment and Java code to execute the benchmark load.

The purpose of TPC benchmarks is to provide relevant, objective performance data to industry users. To achieve that purpose, TPC benchmark specifications require that benchmark tests be implemented with systems, products, technologies and pricing that:

- Are generally available to users;
- Are relevant to the market segment that the individual TPC benchmark models or represents (e.g., TPCx-HS models and represents Hadoop run-time and Hadoop File-system API compatible systems);
- Would plausibly be implemented by a significant number of users in the market segment the benchmark models or represents.

The use of new systems, products, technologies (hardware or software) and pricing is encouraged so long as they meet the requirements above. Specifically prohibited are benchmark systems, products, technologies or pricing (hereafter referred to as "implementations") whose primary purpose is performance optimization of TPC benchmark results without any corresponding applicability to real-world applications and environments. In other words, all "benchmark special" implementations that improve benchmark results but not real-world performance or pricing, are prohibited.

The rules for pricing are included in the TPC Pricing Specification and rules for energy measurement are included in the TPC Energy Specification. Further information is available at www.tpc.org.

Clause 1 – General Items

1.1 Test Sponsor

A statement identifying the benchmark sponsor(s) and other participating companies must be provided

This benchmark was sponsored by Super Micro Computer, Inc..

1.2 Parameter Settings

Settings must be provided for all customer-tunable parameters and options which have been changed from the defaults found in actual products, including by not limited to:

- *Configuration parameters and options for server, storage, network and other hardware component incorporated into the pricing structure;*
- *Configuration parameters and options for operating system and file system component incorporated into the pricing structure;*
- *Configuration parameters and options for any other software component incorporated into the pricing structure;*
- *Compiler optimization options.*

Comment 1: In the event that some parameters and options are set multiple times, it must be easily discernible by an interested reader when the parameter or option was modified and what new value it received each time.

Comment 2: This requirement can be satisfied by providing a full list of all parameters and options, as long as all those that have been modified from their default values have been clearly identified and these parameters and options are only set once.

The supporting files contain the parameters and options used to configure the components involved in this benchmark.

1.3 Configuration Diagrams

Diagrams of both measured and priced configurations must be provided, accompanied by a description of the differences. This includes, but is not limited to:

- *Total number of nodes used;*
- *Total number and type of processors used/total number of cores used/total number of threads used (including sizes of L2 and L3 caches);*
- *Size of allocated memory, and any specific mapping/partitioning of memory unique to the test;*
- *Number and type of disk units (and controllers, if applicable);*
- *Number of channels or bus connections to disk units, including their protocol type;*
- *Number of LAN (e.g., Ethernet) connections and speed for switches and other hardware components physically used in the test or are incorporated into the pricing structure;*
- *Type and the run-time execution location of software components.*

1.3.1 Measured Configuration

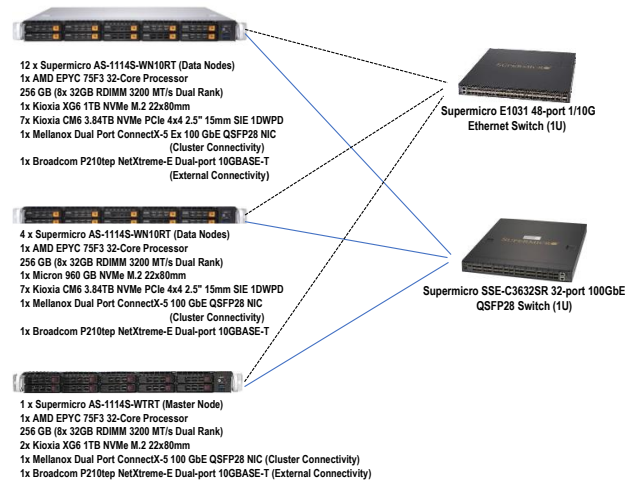


Figure 1-1 Measured Configuration

The measured configuration consisted of:

- Total Nodes: 17 (16x AS-1114S-WN10RT; 1x AS-1114S-WTRT)
- Total Processors/Cores/Threads: 17/544/1,088
- Total Memory: 4.25TiB
- Total Number of Storage Drives/Devices: 130
- Total Storage Capacity: 447.92TB

Server node details:

<p>16x AS-1114S-WN10RT Servers, each with:</p> <ul style="list-style-type: none"> • Processors/Cores/Threads: 1/32/64 • Processor Model: AMD EPYC 75F3 • Memory: 256 GiB • Drives: <ul style="list-style-type: none"> ○ 1x 1 TB NVMe (12 nodes) ○ 1x 960 GB NVMe (4 nodes) ○ 7x 3.84 TB NVMe (all nodes) • Network: <ul style="list-style-type: none"> ○ 1x Mellanox Dual-port ConnectX-5 Ex 100 GbE (12 nodes) ○ 1x Mellanox Dual-port ConnectX-5 100 GbE (4 nodes) ○ 1x Broadcom Dual-port 10 GbE (all nodes) 	<p>1x AS-1114S-WTRT Servers, each with:</p> <ul style="list-style-type: none"> • Processors/Cores/Threads: 1/32/64 • Processor Model: AMD EPYC 75F3 • Memory: 256 GiB • Drives: <ul style="list-style-type: none"> ○ 2x 1 TB NVMe • Network: <ul style="list-style-type: none"> ○ 1x Mellanox Dual-port ConnectX-5 Ex 100 GbE ○ 1x Broadcom Dual-port 10 GbE
--	--

Network connectivity detail:

- 1x SSE-C3632R 32-port 100 GbE
- 1x E1031 48-port 1/10 GbE

The distribution of software components over server nodes is detailed in section 1.5.

1.3.2 Priced Configuration

There are no differences between the priced configuration and the measured configuration.

1.4 Dataset Distribution

The distribution of dataset across all media must be explicitly described.

Table 1-1 describes the distribution of the dataset across all media in the system.

Server Node	Controller	Disk Drive	Description of Content
1	NVMe	nvme7n1	Operating System, Root, Swap, Hadoop Master
2-3	NVMe	nvme7n1	Operating System, Root, Swap, Hadoop Master
2-3	NVMe	nvme0n1, nvme1n1, nvme2n1, nvme3n1, nvme4n1, nvme5n1, nvme6n1	Data, Temp
4-17	NVMe	nvme7n1	Operating System, Root, Swap, Hadoop Master
4-17	NVMe	nvme0n1, nvme1n1, nvme2n1, nvme3n1, nvme4n1, nvme5n1, nvme6n1	Data, Temp

Table 1-1 Dataset Distribution

1.5 Software Components Distribution

The distribution of various software components across the system must be explicitly described.

Table 1-2 Describes the distribution of the software components across the system.

Node	Map/Reduce		HDFS		ZooKeeper
	Resource Manager	Node Manager	NameNode	DataNode	QuorumPeer
1	X		X		X
2-3		X		X	
4-17		X		X	

Table 1-2 Software Component Distribution

Distributed file system implementation and corresponding Hadoop File System API version must be disclosed.

CDP Private Cloud Base Edition 7.1.6 (fully HDFS compatible at the API level).

Map/Reduce implementation and corresponding version must be disclosed.

CDP Private Cloud Base Edition 7.1.6 (compatible equivalent to Hadoop 3.1.1.7.1.6.0-297).

Clause 2 – Workload Related Items

2.1 Hardware & Software Tunables

Script or text used to set for all hardware and software tunable parameters must be reported.

The Supporting File Archive contains all configuration scripts.

2.2 Run Report

The run report generated by TPCx-HS benchmark kit must be reported.

The Supporting File Archive contains the full run report. Following are extracts from the run report that lists the performance summary for both runs.

```

Run Report for Run 1 – Repeatability Run
=====
TPCx-HS Performance Metric (HSph@SF) Report

Test Run 1 Details      Total Time =                2353
                        Total Size =                300000000000
                        Scale-Factor =                30

TPCx-HS Performance Metric (HSph@SF):                45.8996

=====
    
```

```

Run Report for Run 2 – Performance Run
=====
TPCx-HS Performance Metric (HSph@SF) Report

Test Run 2 Details      Total Time =                2381
                        Total Size =                300000000000
                        Scale-Factor =                30

TPCx-HS Performance Metric (HSph@SF):                45.3651

=====
    
```

2.3 Benchmark Kit Identification

Version number of TPCx-HS kit and checksum for HSGen, HSSort and HSValidate Programs must be reported.

```

Kit Version                2.0.3

File                        MD5
-----
BigData_cluster_validate_suite.sh  57f7cd68251a9aba0feb6648630ff5da
HSDDataCheck.sh                bcf0b946a49d1249c9da174b5d9805f1
TPCx-HS-master_MR2.jar         492cbc51a1a60c28b43d96c79d08683d
TPCx-HS-master.sh              c619a0819571ecd00cd75d2b76ba8c64
    
```

2.4 Benchmark Kit Changes

The required data protection was provided by HDFS Erasure Coding rather than the default three-way data replication. A policy of RS-6-3-1024k was used. Therefore, each block group consisted of 6 data blocks and 3 parity blocks. Each block within a given block group was placed on a different node thus ensuring the required data protection.

To collect the necessary data for auditing, the HSDDataCheck.sh script was modified. In accordance with the TPCx-HS Standard Specification, this change received prior approval from the TPCx-HS subcommittee.

Clause 3 – SUT Related Items

3.1 Data Storage Ratio

The data storage ratio must be disclosed.

Table 3-1 describes the details of the storage devices configured on the system and their capacity.

Quantity	Capacity	Total (TB)
12	1 TB	12.00
4	960 GB	3.84
112	3.84 TB	430.08
2	1 TB	2.00
Total Storage (TB)		447.92

Table 3-1 Storage Device Capacities

Scale Factor = 30

Data Storage Ratio = (Total Storage (TB) / SF) = 14.93

3.2 Memory Ratio

The Scale Factor to memory ratio must be disclosed.

Total Configured Memory (TiB) = 4.25

Scale Factor to Memory Ratio = (SF / Total Memory(TiB)) = 7.06

Clause 4 – Metrics Related Items

4.1 HSGen Time

The HSGen time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSGen	515.927	517.915

Table 4-1 HSGen Times

4.2 HSSort Time

The HSSort time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSSort	1,547.622	1,572.924

Table 4-2 HSSort Times

4.3 HSValidate Time

The HSValidate time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSValidate	256.541	261.409

Table 4-3 HSValidate Times

4.4 HSDataCheck Times

Both HSDataCheck times must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSDataCheck (pre-sort)	18.000	15.000
HSDataCheck (post-sort)	16.000	14.000

Table 4-4 HSDataCheck Times

4.5 Performance & Price-Performance

The performance metric (HSph@SF) must be disclosed for Run 1 and Run 2. Price-performance metric (\$/HSph@SF) must be disclosed for the performance run.

	Run 1	Run 2
HSph@30TB	45.89	45.36

Table 4-5 Performance Metrics

Run 2 Price-Performance: 19,529.68 \$/ HSph@30TB

Auditor's Information & Letter of Attestation

The auditor's agency name, address, phone number, and Attestation letter must be included in the full disclosure report. A statement should be included specifying who to contact in order to obtain further information regarding the audit process.

This benchmark was audited by Doug Johnson, InfoSizing.

www.sizing.com
63 Lourdes Drive
Leominster, MA 10453
978-343-6562

This benchmark's Full Disclosure Report (FDR) can be downloaded from www.tpc.org.

A copy of the auditor's Letter of Attestation follows.



Srini Bala
 Super Micro Computer, Inc.
 980 Rock Avenue,
 San Jose, CA 95131
 USA

September 14, 2021

I verified the TPC Express Benchmark™ HS v2.0.3 performance of the following configuration:

Platform: Supermicro Cluster with:
 16x AS-1114S-WN10RT Servers (Data Nodes)
 1x AS-1114S-WTRT Server (Master Node)

Operating System: SUSE Linux Enterprise Server 12 SP5

Apache Hadoop CDP Private Cloud Base Edition 7.1.6

Compatible Software:

The results were:

Performance Metric 45.36 HSph@30TB
 Run Elapsed Time 2,381.00 Seconds

Cluster	16x AS-1114S-WN10RT, 1x AS-1114S-WTRT with:		
CPUs	1x AMD® EPYC 75F3 32-Core Processor (all nodes)		
Memory	256 GiB (all nodes)		
Storage	Qty	Size	Type
	2	1 TB	NVMe (master node)
	1	1 TB	NVMe (12 data nodes)
	1	960 GB	NVMe (4 data nodes)
	7	3.84 TB	NVMe (all data nodes)

In my opinion, these performance results were produced in compliance with the TPC requirements for the benchmark.

The following verification items were given special attention:

- All TPC-provided components were verified to be v2.0.3
- No modifications were made to any of the Java code
- Any and all modifications to shell scripts were reviewed for compliance
- All checksums were validated for compliance

63 Lourdes Dr. | Leominster, MA 01453 | 978-343-6562 | www.sizing.com

- The generated dataset was properly scaled to 30 TB
- The generated dataset and the sorted dataset were erasure coded with a policy of RS-6-3-1024k
- The elapsed times for all phases and runs were correctly measured and reported
- The Storage and Memory Ratios were correctly calculated and reported
- The system pricing was verified for major components and maintenance
- The major pages from the FDR were verified for accuracy

Additional Audit Notes:

None.

Respectfully Yours,

A handwritten signature in cursive script that reads "Doug Johnson". The signature is written in black ink and has a long, sweeping horizontal line extending to the right.

Doug Johnson, Certified TPC Auditor

63 Lourdes Dr. | Leominster, MA 01453 | 978-343-6562 | www.sizing.com

Supporting Files Index

Clause	Description	Archive File Pathname
Clause 1	Parameters and options used to configure the system	SupportingFiles/Clause1
Clause 2	Configuration scripts and Run Report	SupportingFiles/Clause2
Clause 3	System configuration details	SupportingFiles/Clause3

Third-Party Price Quotes

APC

The screenshot displays the APC website interface for the AP7911B product. At the top, there is a navigation bar with the Schneider Electric logo, location (UNITED STATES), and various utility links like 'Partner Login' and 'BECOME A PARTNER'. Below this is a search bar and a main navigation menu with categories like 'Products & Services', 'Solutions', 'Support', 'How to buy', and 'Partners & Alliances'. A green breadcrumb trail indicates the path: Home > Products & Services > ... > Rack Power Distribution > NetShelter Switched Rack PDUs > Rack PDU, Switched, 2U, 30A, 208V, (16)C13.

The main product section features a large image of the AP7911B Rack PDU. To its right, the product title is 'Rack PDU, Switched, 2U, 30A, 208V, (16)C13' with the model number 'AP7911B' and an 'R6HS' badge. The price is listed as '\$1,025.00'. Below the price is a quantity selector set to '1' and an 'Add To Cart' button. There are also links for 'Add to Favorites' and 'Compare'. A 'Get Support' section includes links for 'View User Manual', 'Update Software/Firmware', 'View Product FAQs', and 'Register Your Product'. A 'Our Partners' section lists several partners with 'Buy' buttons: Supermicro (In Stock), ZONES (In Stock), shi (In Stock), and another partner (In Stock). A 'Find a Reseller' link is also present.

Amazon

The screenshot shows the top of the Amazon website. It includes the Amazon logo, a search bar with 'Electronics' selected, and a navigation menu with categories like 'All', 'Best Sellers', 'Customer Service', 'Prime', 'New Releases', 'Pharmacy', 'Today's Deals', 'Books', 'Fashion', and 'Toys & Games'. Below the navigation menu, there are sub-categories for 'Computers', 'Laptops', 'Desktops', 'Monitors', 'Tablets', 'Computer Accessories', and 'PC Compon'. A banner for 'amazon live Beauty Haul Live with Jenni 'JWoww' Farley' is also visible.

Electronics › Computers & Accessories › Monitors

LONCEVON-12 inch IPS 1920x1080P Portable Small HDMI Monitor with VGA-AV-USB-BNC Port ; External Small HDMI Screen LCD Display Monitor for Laptop Raspberry PI/TV ; Build in Dual Speakers, Remote.

Visit the LONCEVON Store
 ★★★★★ 122 ratings
 | 27 answered questions

Price: **\$99.99** & FREE Returns

Get \$50 off instantly: Pay \$49.99 ~~\$99.99~~ upon approval for the Amazon Rewards Visa Card. No annual fee.

This screenshot is similar to the first one, showing the Amazon website header and navigation menu. It includes the Amazon logo, a search bar with 'Electronics' selected, and a navigation menu with categories like 'All', 'Best Sellers', 'Customer Service', 'Prime', 'New Releases', 'Pharmacy', 'Today's Deals', 'Books', 'Fashion', and 'Toys & Games'. Below the navigation menu, there are sub-categories for 'Computers', 'Laptops', 'Desktops', 'Monitors', 'Tablets', 'Computer Accessories', and 'PC'. A banner for 'amazon live Beauty Haul Live with Jenni 'JWoww' Farley' is also visible.

Electronics › Computers & Accessories › Computer Accessories & Peripherals › Keyboards, Mice & Accessories › Keyboard & Mouse Combos

Logitech Media Combo MK200 Full-Size Keyboard and High-Definition Optical Mouse (920-002714)

Visit the Logitech Store
 ★★★★★ 2,059 ratings

Amazon's Choice in Computer Keyboard & Mouse Combos by Logitech

Price: **\$41.50** & FREE Returns

Get \$10 off instantly: Pay \$31.50 upon approval for the Amazon Store Card.

- Instant media and internet access
- Eight shortcut keys
- Full-size keyboard
- Comfortable, quiet typing
- Instant access to applications

Click image to open expanded view