



Hewlett Packard
Enterprise

Hewlett Packard Enterprise

TPC Express Benchmark™ AI Full Disclosure Report

ProLiant DL325 Gen11

with 1x ProLiant DL325 Gen11
using

Anaconda Business

running on

Red Hat Enterprise Linux 8.6 (Ootpa)

TPCx-AI Version
Report Edition
Report Submitted

1.0.3.1
First
August 21, 2024

First Edition - August 2024

Hewlett Packard Enterprise (HPE), 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, TPC Express Benchmark™ AI 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.

HPE and the HPE Logo are trademarks of Hewlett Packard Enterprise 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 HPE and any other company.

TPC Express Benchmark™ AI, TPCx-AI, and AIUCpm@10, are registered certification marks of the Transaction Processing Performance Council.

The HPE 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 HPE business contact for information on the products or services available in your area. You can find additional information via HPE's web site at www.hpe.com. Actual performance and environmental costs of HPE products will vary depending on individual customer configurations and conditions.

Copyright© 2024 Hewlett Packard Enterprise

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

HPE conducted the TPC Express Benchmark™ AI (TPCx-AI) on the ProLiant DL325 Gen11. The software used included Anaconda. This report provides full disclosure of the results. All testing was conducted in conformance with the requirements of the TPCx-AI Standard Specification, Revision 1.0.3.1.

Configuration Overview


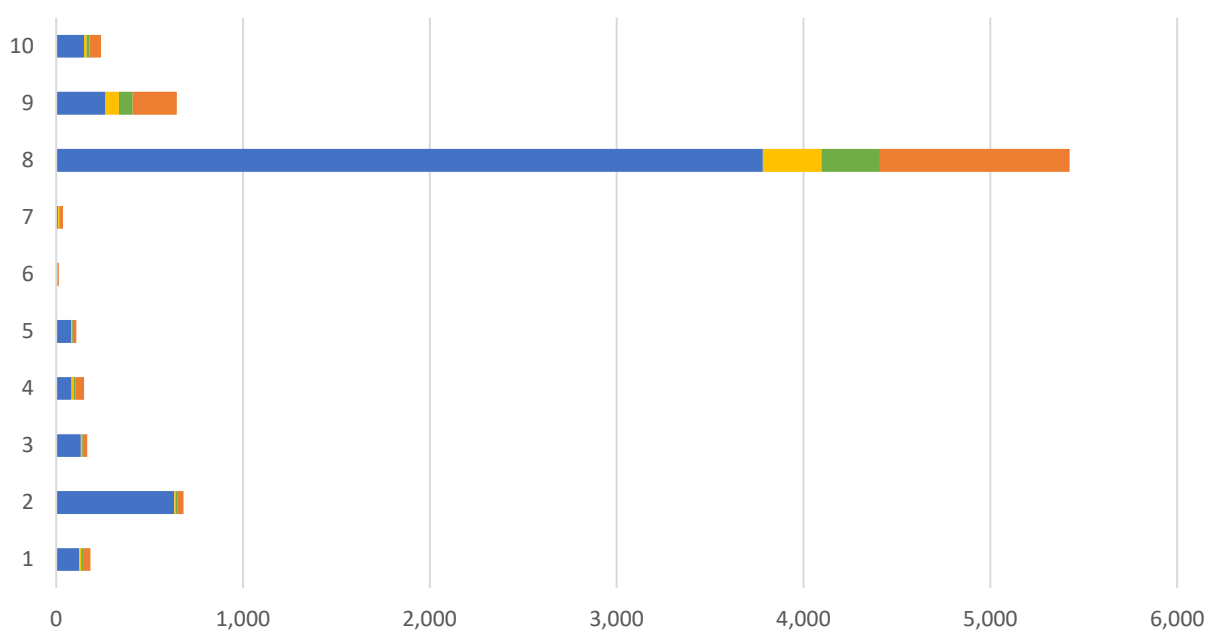
Test Sponsor	Node(s)	Operating System
HPE	1x ProLiant DL325 Gen11 (Server)	Red Hat Enterprise Linux 8.6 (Ootpa)

Metrics Overview

Total System Cost	Performance	Price/Performance	Availability Date
\$122,396 USD	714.04 AIUCpm@10	171.42 USD \$/AIUCpm@10	August 21, 2024

Executive Summary

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


 Hewlett Packard Enterprise	<h1>ProLiant DL325 Gen11</h1>		TPCx-AI 1.0.3.1 TPC Pricing 2.9.0 Report Date Aug. 21, 2024	
TPCx-AI Performance 714.04 AIUCpm@10	Total System Cost \$122,396 USD	Price/Performance \$171.42 USD/AIUCpm@10	Availability Date August 21, 2024	
Framework Anaconda Business	Operating System Red Hat Enterprise Linux 8.6 (Ootpa)	Other Software N/A	Scale Factor 10	Streams 100
<div data-bbox="235 651 690 693"> <h3>Use Case Time (sec.) by Phase</h3> </div> <div data-bbox="844 651 1445 693"> ■ Training ■ Serving 1 ■ Serving 2 ■ Throughput (Avg) </div> 				
Physical Storage / Scale Factor 96.00	Scale Factor / Physical Memory 0.01	Main Data Redundancy Model RAID 1		
Servers: Total Processors/Cores/Threads		1 1 / 32 / 64		
Server Type	1x ProLiant DL325 Gen11 (Server)			
Processors	1x AMD EPYC 9374F 32-Core Processor			
Memory	768 GiB			
Storage Controller	1x NS204i-u			
Storage Device	2x 480 GB M.2 NVMe SSD			
Network Controller	1x Intel I350-T4 1Gb 4-port			

 Hewlett Packard Enterprise	<h1>ProLiant DL325 Gen11</h1>	TPCx-AI	1.0.3.1
		TPC Pricing	2.9.0
		Report Date	Aug. 21, 2024

Description	Part Number	Source	List Price	Qty	Extended Price	1-Yr. Maintenance
Hardware						
HPE DL325 Gen11 8SFF CTO Server	P54199-B21	1	\$2,767.00	1	\$2,767.00	
AMD EPYC 9374F 3.85GHz 32-core 320W Processor for HPE	P53710-B21	1	\$11,019.00	1	\$11,019.00	
HPE ProLiant DL3X5 Gen11 CPU Performance 1U Heat Sink Kit	P58457-B21	1	\$225.00	1	\$225.00	
HPE ProLiant DL3XX Gen11 1U Performance Fan Kit	P58462-B21	1	\$113.00	1	\$113.00	
HPE 64GB (1x64GB) Dual Rank x4 DDR5-4800	P50312-B21	1	\$4,068.00	12	\$48,816.00	
HPE 800W FS Plat Ht Plug LH PS Kit	P38995-B21	1	\$685.00	2	\$1,370.00	
HPE USB US Keyboard/Mouse Kit	631341-B21	1	\$32.00	3	\$96.00	
Intel I350-T4 Ethernet 1Gb 4-port BASE-T Adapter for HPE	P21106-B21	1	\$1,102.00	1	\$1,102.00	
Add to Rack 48U 600mmx1075mm G2 Advanced Pallet	P9K19A	1	\$3,023.00	1	\$3,023.00	
HPE 3 Year Tech Care Essential DL325 Gen11 Service	H78S6E	1	\$2,958.00	1		\$2,958.00
Subtotal					\$68,531.00	\$2,958.00
Storage						
NS204i-u Gen11 OS Boot, RAID 1	P48183-B21	1	\$1,994.00	1	\$1,994.00	
HPE Gen 11 EI Rail Kit	P52351-B21	1	\$141.00	1	\$141.00	
HPE DL325G11 Standard Riser	standardDL385G11Riser	1	\$0.00	1	\$0.00	\$0.00
Subtotal					\$2,135.00	\$0.00
Other						
HP P22v G5 FHD Monitor	V9E67AA#ABA	3	\$109.00	3	\$327.00	
Subtotal					\$327.00	\$0.00
Server Software						
RHEL Svr Sckt/2 Gst 1yr 24x7 E-LTU	J8J36A	1	\$1,299.00	1	\$1,299.00	included
Anaconda Business, 100 users	NA	2	\$600.00	100	\$60,000.00	
Anaconda Business Premium Support, including 24x7 support	NA	2	\$35,000.00	1		\$35,000.00
Subtotal					\$61,299.00	\$35,000.00
Total Extended Price					\$132,292.00	\$37,958.00
Total Discounts					\$45,932.00	\$1,922.00

Pricing: 1 = HPE; 2 = Anaconda; 3 = Hewlett Packard Inc. * Discount applies to all line items where Key = 1. Discount based upon total system cost as purchased by a regular customer. Audited by Doug Johnson, InfoSizing	Total System Cost (USD):	\$122,396
	AIUCpm@10:	714.04
	\$/AIUCpm@10:	\$171.42

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.

 Hewlett Packard Enterprise	<h1>ProLiant DL325 Gen11</h1>	TPCx-AI 1.0.3.1 TPC Pricing 2.9.0 Report Date Aug. 21, 2024														
<u>Numerical Quantities</u>																
AIUCpm@10 Scale Factor Streams Kit Version Execution Status Accuracy Status	714.04 10 100 1.0.3.1 Pass Pass	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%;">T_{Load}</td> <td style="text-align: right;">2.26</td> </tr> <tr> <td>T_{LD}</td> <td style="text-align: right;">2.26</td> </tr> <tr> <td>T_{PTT}</td> <td style="text-align: right;">124.24</td> </tr> <tr> <td>T_{PST1}</td> <td style="text-align: right;">11.41</td> </tr> <tr> <td>T_{PST2}</td> <td style="text-align: right;">11.52</td> </tr> <tr> <td>T_{PST}</td> <td style="text-align: right;">11.52</td> </tr> <tr> <td>T_{TT}</td> <td style="text-align: right;">1.54</td> </tr> </table>	T_{Load}	2.26	T_{LD}	2.26	T_{PTT}	124.24	T_{PST1}	11.41	T_{PST2}	11.52	T_{PST}	11.52	T_{TT}	1.54
T_{Load}	2.26															
T_{LD}	2.26															
T_{PTT}	124.24															
T_{PST1}	11.41															
T_{PST2}	11.52															
T_{PST}	11.52															
T_{TT}	1.54															
<u>Test Times</u>																
Overall Run Start Time Overall Run End Time Overall Run Elapsed Time		2023-12-11 02:08:12.360 2023-12-11 04:21:01.262 7,968.902														
Load Test Start Time Load Test End Time Load Test Elapsed Time		2023-12-11 02:10:06.050 2023-12-11 02:10:08.319 2.269														
Power Training Start Time Power Training End Time Power Training Elapsed Time		2023-12-11 02:10:08.320 2023-12-11 03:37:51.970 5,263.650														
Power Serving 1 Start Time Power Serving 1 End Time Power Serving 1 Elapsed Time		2023-12-11 03:37:51.971 2023-12-11 03:45:18.417 446.446														
Power Serving 2 Start Time Power Serving 2 End Time Power Serving 2 Elapsed Time		2023-12-11 03:45:18.418 2023-12-11 03:52:45.718 447.300														
Scoring Start Time Scoring End Time Scoring Elapsed Time		2023-12-11 03:53:22.422 2023-12-11 03:55:12.316 109.894														
Throughput Start Time Throughput End Time Throughput Elapsed Time		2023-12-11 03:55:12.332 2023-12-11 04:21:01.260 1,548.928														

 Hewlett Packard Enterprise	<h1>ProLiant DL325 Gen11</h1>	TPCx-AI	1.0.3.1
		TPC Pricing	2.9.0
		Report Date	Aug. 21, 2024

Numerical Quantities (continued)

Use Case Times & Accuracy

Use Case	Training (sec)	Serving 1 (sec)	Serving 2 (sec)	Throughput (avg)	Accuracy
UC01	123.852	10.523	10.445	39.055	0.000
UC02	631.473	8.127	8.132	34.611	0.453
UC03	131.471	5.604	5.606	23.851	3.609
UC04	82.617	10.757	10.816	45.298	0.707
UC05	80.456	4.611	4.850	17.923	0.486
UC06	7.475	1.079	1.166	4.442	0.483
UC07	11.543	4.304	4.146	16.525	1.032
UC08	3,781.938	313.782	314.417	1,015.124	0.733
UC09	263.942	72.767	72.838	237.063	1.000
UC10	148.818	14.823	14.818	62.443	0.816

Use Case Serving Times (sec.)




Table of Contents

- Abstract..... 3
- Executive Summary 3
- Table of Contents..... 8
- Clause 0 – Preamble10
 - 0.1 TPC Express Benchmark™ AI Overview10
- 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 Differences Between the Measured and the Priced Configurations.....13
- Clause 2 – SW Components & Data Distribution14
 - 2.1 Roles and Dataset Distribution14
 - 2.2 File System Implementation.....14
 - 2.3 Execution Engine, Frameworks, Driver & Libraries14
 - 2.4 Applied Patches.....15
- Clause 3 – Workload Related Items.....16
 - 3.1 Hardware & Software Tuning.....16
 - 3.2 Kit Version & Modifications16
 - 3.3 Use Case Elapsed Times16
 - 3.4 SUT Validation Test Output22
 - 3.5 Configuration Parameters.....23
- Clause 4 – SUT Related Items.....24
 - 4.1 Specialized Hardware/Software24
 - 4.2 Configuration Files.....24
 - 4.3 SUT Environment Information.....24
 - 4.4 Data Storage to Scale Factor Ratio24
 - 4.5 Scale Factor to Memory Ratio24
 - 4.6 Output of Tests24
 - 4.7 Additional Sponsor Files24
 - 4.8 Model Optimizations24
- Clause 5 – Metrics and Scale Factor25
 - 5.1 Reported Performance Metrics25

5.2 Throughput Test Stream Times26

Auditor's Information27

Third-Party Price Quotes.....30

 Anaconda30

 Hewlett Packard Inc.....31

Supporting Files Index32

Clause 0 – Preamble

0.1 TPC Express Benchmark™ AI Overview

Artificial intelligence (AI) has become a key transformational technology of our times. Advances in neural networks and other machine learning techniques have made it possible to use AI on a variety of use cases. From the public sector to aerospace, defense and academia, new and improved ways to use AI techniques are changing the way we harness data and analytics. This along with advances in compute, interconnect and memory technologies have made possible to solve complicated challenges that will ultimately benefit customers in production datacenter and cloud environments.

Abundant volumes of rich data from text, images, audio and video are the essential starting point for creating a benchmark that would represent the myriad of use cases and customers. TPC Express Benchmark™ AI (TPCx-AI) is created in keeping with the TPC tradition of emulating real world AI scenarios and data science use cases. Unlike most other AI benchmarks, the TPCx-AI uses a diverse dataset and is able to scale across a wide range of scale factors. TPCx-AI may later expand with additional use cases and add additional flexibility for a greater variety of implementations.

The benchmark defines and provides a means to evaluate the System Under Test (SUT) performance as a general-purpose data science system that:

- Generates and processes large volumes of data.
- Trains preprocessed data to produce realistic machine learning models.
- Conducts accurate insights for real-world customer scenarios based on the generated models.
- Can scale to large scale distributed configurations.
- Allows for flexibility in configuration changes to meet the demands of the dynamic AI landscape.

The benchmark models real-life examples of companies and public-sector organizations that use a range of analytics techniques, both AI and more traditional machine learning approaches, as well as the potential application of these techniques in situations like those in which they have already been successfully deployed. In addition, the benchmark measures end to end time to provide insights for individual use cases, as well as throughput metrics to simulate multiuser environments for a given hardware, operating system, and data processing system configuration under a controlled, complex, multi-user AI or machine learning data science workload.

The purpose of TPC benchmarks is to provide relevant, objective performance data to industry users. To achieve that purpose, TPC benchmark specifications require benchmark runs 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-AI models and represents complex, high data volume, decision support environments).
- Would plausibly be implemented.

The TPCx-AI kit is available from the TPC website (see www.tpc.org/tpcx-ai/ for more information). Users must sign up and agree to the TPCx-AI End User Licensing Agreement (EULA) to download the kit. All related work (such as collaterals, papers, derivatives) must acknowledge the TPC and include the TPCx-AI copyright. The TPCx-AI kit includes: TPCx-AI Specification document (this document), TPCx-AI Users Guide (README.md) documentation, scripts to set up the benchmark environment, code to execute the benchmark workload, Data Generator, use case related files, and Benchmark Driver.

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.

Further information is available at www.tpc.org.

Clause 1 – General Items

1.1 Test Sponsor

This benchmark was sponsored by Hewlett Packard Enterprise.

1.2 Parameter Settings

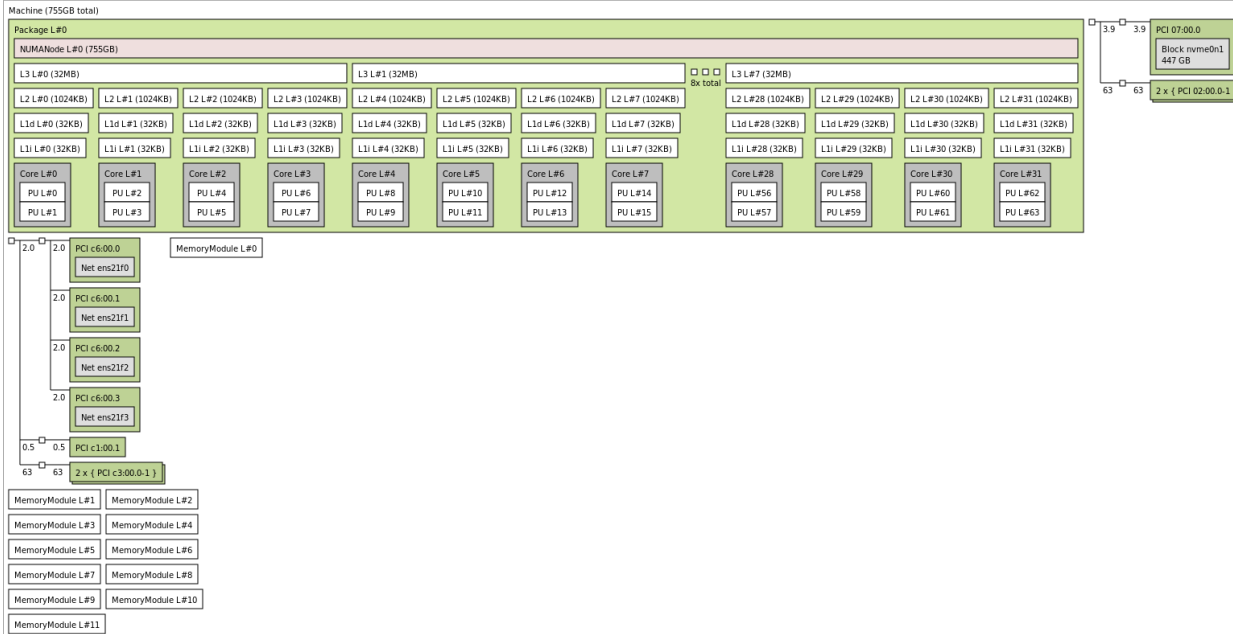
The [Supporting Files Archive](#) contains the parameters and options used to configure the components involved in this benchmark.

1.3 Configuration Diagrams

The measured configuration diagram is shown below. In addition, any differences between the measured and the priced configurations are described.

1.3.1 Measured Configuration

Nodes:	1		
Processors/Cores/Threads:	1/32/64	Storage Devices:	2
Total Memory:	768 GiB	Storage Capacity:	960 GB



Host: dl325g11-x4
 Indexes: logical
 Date: Mon 11 Dec 2023 03:21:31 AM CST

Server	<u>Server</u>
Procs/Cores/Threads:	1x ProLiant DL325 Gen11: 1/32/64
Processor Model:	1x AMD EPYC 9374F 32-Core Processor
Memory:	768 GiB
Storage Controller:	1x NS204i-u
Storage Devices:	2x 480 GB M.2 NVMe SSD
Network Controller:	1x Intel I350-T4 1Gb 4-port

The distribution of software components over server nodes is detailed in [Clause 2](#).

1.3.2 Differences Between the Measured and the Priced Configurations

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

Clause 2 – SW Components & Data Distribution

2.1 Roles and Dataset Distribution

Table 2-1 describes the distribution of the dataset across all media in the SUT.

Server	Host Name	SW Services	Storage	Contents
1x ProLiant DL325 Gen11	dl325g11-xAI	All	2x 480 GB M.2 NVMe SSD	OS, Data

Table 2-1 Software Components and Dataset Distribution

2.2 File System Implementation

A local file system provided by Red Hat Enterprise Linux 8.6 (Ootpa) / Anaconda Business was used for data generation and the Load Test. The data set was not relocated after generation and before the Load Test.

2.3 Execution Engine, Frameworks, Driver & Libraries

Anaconda Business consisted of the following components.

Component	Version
python	3.9
setuptools	59.8
pandas	1.5.3
scikit-learn	1.2.2
xgboost	1.7.4
numpy	1.23.5
nose	1.3.7
scipy	1.10.1
statsmodels	0.13.5
patsy	0.5.3
tqdm	4.65
keras	2.11.0
tensorflow	2.11.0
joblib	1.2.0
pyyaml	6.0
matplotlib	3.7.1
jinja2	3.1.2
pycryptodome	3.16
scikit-surprise	1.1.3
librosa	0.8.1
opencv	4.5.3
imbalanced-learn	0.10.1
tensorflow-addons	0.19.0

Table 2-2 Software Components

For a detailed listing of installed libraries, please see the envInfo logs in the [Supporting Files](#).

2.4 Applied Patches

No additional vendor-supported patches were applied to the SUT.

Clause 3 – Workload Related Items

3.1 Hardware & Software Tuning

The [Supporting Files](#) archive contains all hardware and software configuration scripts.

3.2 Kit Version & Modifications

Table 3-1 shows the version of the TPCx-AI used to produce this result along with any kit files that were modified to facilitate system, platform, and framework differences.

TPCx-AI Kit Version	1.0.3.1
<u>Modified File</u>	<u>Description of Changes</u>
No modifications	N/A

Table 3-1 Kit Version & Modifications

3.3 Use Case Elapsed Times

Below are the elapsed times for each use case. Use cases are grouped based on whether they use Deep Learning or Machine Learning techniques.

Type	UC ID	P1	P2	T1	T2	T3	T4
Deep Learning	2	8.127	8.132	31.070	27.001	28.951	33.410
	5	4.611	4.850	21.487	15.288	15.290	16.915
	9	72.767	72.838	220.256	245.356	298.129	236.120
Machine Learning	1	10.523	10.445	37.575	80.004	36.334	49.658
	3	5.604	5.606	54.438	17.398	15.044	22.112
	4	10.757	10.816	37.519	57.295	29.462	81.690
	6	1.079	1.166	3.491	2.691	3.725	2.759
	7	4.304	4.146	14.207	11.806	12.612	11.988
	8	313.782	314.417	995.343	972.582	1,011.022	1,004.934
	10	14.823	14.818	77.528	51.632	62.954	51.567

Type	UC ID	T5	T6	T7	T8	T9	T10
Deep Learning	2	32.668	30.130	29.073	35.284	29.388	31.375
	5	12.881	35.231	13.317	8.538	24.441	31.349
	9	300.236	209.090	236.868	214.681	230.624	254.740
Machine Learning	1	31.703	30.725	74.752	40.544	52.253	25.365
	3	21.541	18.789	24.631	48.873	49.234	41.487
	4	35.853	32.767	35.981	42.934	41.459	33.691
	6	5.607	8.204	5.777	4.679	5.212	6.446
	7	16.932	9.429	13.776	13.019	13.100	17.745
	8	977.043	997.368	1,013.850	1,062.859	969.434	976.495
	10	41.972	51.882	50.779	49.048	56.083	47.855

Type	UC ID	T11	T12	T13	T14	T15	T16
Deep Learning	2	26.510	37.711	16.889	30.021	25.847	37.019
	5	22.079	14.254	16.740	15.531	14.277	15.664
	9	212.422	238.345	292.587	234.403	221.069	209.417
Machine Learning	1	76.999	34.359	29.827	29.316	38.107	40.271
	3	19.585	11.510	23.292	16.335	22.821	22.406
	4	49.475	79.662	39.255	46.049	38.380	40.534
	6	5.707	5.089	1.421	3.452	5.594	3.132
	7	11.803	11.561	21.258	12.580	15.416	15.123
	8	1,041.571	1,028.289	1,041.751	1,009.825	1,017.149	972.100
	10	49.458	61.674	43.920	114.432	106.584	108.093

Type	UC ID	T17	T18	T19	T20	T21	T22
Deep Learning	2	30.825	35.684	30.521	33.446	31.892	28.567
	5	13.252	17.275	23.653	9.253	15.833	13.369
	9	214.803	222.226	199.136	289.446	297.490	303.298
Machine Learning	1	82.291	28.611	42.522	31.424	19.622	31.877
	3	24.465	18.719	51.785	21.012	20.325	13.321
	4	33.104	87.894	40.860	41.217	40.733	35.005
	6	5.177	3.319	6.807	4.428	6.610	3.060
	7	12.003	14.092	18.240	15.667	10.626	11.704
	8	1,036.348	1,027.057	1,037.775	1,026.802	1,016.307	966.999
	10	60.133	71.375	51.422	43.668	69.274	64.563

Type	UC ID	T23	T24	T25	T26	T27	T28
Deep Learning	2	35.041	27.641	27.458	31.505	27.397	28.538
	5	31.064	13.810	24.838	37.610	13.769	14.294
	9	199.152	292.042	210.048	216.574	256.602	206.220
Machine Learning	1	33.100	29.270	40.695	26.050	80.730	27.649
	3	19.692	21.016	17.839	16.108	12.812	50.153
	4	60.399	27.586	38.282	34.598	43.560	41.043
	6	5.174	5.262	3.511	5.663	5.359	5.894
	7	31.362	21.986	12.895	18.515	13.198	12.190
	8	1,003.352	1,011.535	1,037.364	1,040.357	962.288	1,027.019
	10	43.567	57.270	96.649	60.526	47.929	48.465

Type	UC ID	T29	T30	T31	T32	T33	T34
Deep Learning	2	35.230	26.970	29.012	36.635	25.822	71.356
	5	17.066	17.701	22.861	20.252	16.831	15.468
	9	245.713	277.726	299.865	248.595	212.205	244.264
Machine Learning	1	14.403	40.217	31.353	26.472	49.089	35.621
	3	6.003	17.184	21.625	22.037	18.258	15.275
	4	86.635	16.859	35.978	80.378	89.470	33.842
	6	1.302	4.743	3.950	2.872	2.941	3.319
	7	20.875	33.384	20.105	18.113	16.007	17.303
	8	1,038.686	1,050.395	969.266	969.803	1,038.170	1,005.127
	10	74.599	45.921	63.302	49.559	51.461	45.076

Type	UC ID	T35	T36	T37	T38	T39	T40
Deep Learning	2	64.262	31.391	35.334	45.121	27.416	30.049
	5	14.398	16.191	30.623	16.857	31.904	13.936
	9	202.043	212.615	243.069	201.234	270.113	215.933
Machine Learning	1	30.208	42.558	42.425	31.120	36.345	80.777
	3	23.795	16.640	23.155	17.870	15.615	17.940
	4	61.345	39.262	74.804	84.039	50.934	36.310
	6	4.524	5.364	5.518	5.376	6.421	2.509
	7	15.443	17.521	14.205	17.368	11.845	13.975
	8	1,037.243	1,011.883	938.833	962.692	1,037.605	1,014.116
	10	59.037	105.169	58.921	51.264	49.464	48.380

Type	UC ID	T41	T42	T43	T44	T45	T46
Deep Learning	2	36.896	33.094	32.889	24.729	25.068	24.898
	5	16.769	34.141	4.944	35.264	30.625	14.428
	9	198.113	234.895	235.990	205.920	263.357	199.462
Machine Learning	1	42.050	33.168	12.047	32.512	36.238	62.590
	3	16.469	6.192	20.618	21.567	24.763	45.106
	4	48.709	27.279	81.053	46.124	50.603	47.259
	6	3.516	4.597	4.076	2.702	6.050	4.182
	7	13.010	13.198	25.868	11.788	31.802	16.129
	8	996.740	1,055.097	1,069.749	1,024.339	970.048	1,032.010
	10	109.752	89.449	56.364	92.323	59.402	65.825

Type	UC ID	T47	T48	T49	T50	T51	T52
Deep Learning	2	26.653	23.197	28.617	27.477	31.593	34.642
	5	16.084	28.314	12.855	12.691	17.203	13.256
	9	210.102	235.253	231.617	232.193	260.965	197.738
Machine Learning	1	78.457	57.831	32.416	35.602	23.664	37.021
	3	18.208	23.860	19.292	17.194	44.723	17.685
	4	54.541	33.520	36.676	32.866	30.719	40.076
	6	3.607	4.990	3.031	3.370	4.930	6.191
	7	11.171	35.625	19.465	14.594	30.979	10.162
	8	1,000.813	1,016.366	966.788	1,042.313	1,031.370	1,071.458
	10	63.142	61.704	106.328	51.346	50.813	49.420

Type	UC ID	T53	T54	T55	T56	T57	T58
Deep Learning	2	68.715	66.751	36.130	25.983	21.875	68.636
	5	16.379	17.848	14.855	12.549	12.634	13.529
	9	262.309	199.044	206.356	222.420	248.311	207.920
Machine Learning	1	31.950	37.792	29.674	31.715	31.200	49.016
	3	19.242	35.223	18.484	22.293	17.773	32.223
	4	35.964	40.353	36.297	36.693	80.140	38.854
	6	4.459	6.032	3.101	3.394	3.817	3.446
	7	15.468	13.992	15.306	13.104	14.729	13.378
	8	969.295	985.380	1,032.963	1,034.676	970.847	994.397
10	38.351	60.810	70.707	69.236	73.182	53.888	

Type	UC ID	T59	T60	T61	T62	T63	T64
Deep Learning	2	29.215	37.816	27.024	37.108	35.107	30.571
	5	13.237	13.160	15.293	31.237	11.536	14.304
	9	207.004	218.563	230.117	196.892	275.038	228.989
Machine Learning	1	32.629	33.600	32.317	42.757	33.412	31.421
	3	21.298	21.934	29.454	22.467	47.357	16.677
	4	49.122	46.877	83.052	35.495	15.251	33.313
	6	2.660	3.964	6.014	3.417	2.642	3.875
	7	15.282	13.677	17.888	16.637	13.281	12.327
	8	1,029.542	975.285	1,017.725	1,038.881	1,050.793	1,055.714
10	108.690	107.638	52.506	54.412	49.329	47.734	

Type	UC ID	T65	T66	T67	T68	T69	T70
Deep Learning	2	25.360	41.953	37.404	55.420	29.623	28.967
	5	28.980	14.211	13.755	32.641	14.319	15.805
	9	202.767	211.979	288.995	236.201	204.445	248.889
Machine Learning	1	31.522	75.845	17.877	26.756	36.588	33.367
	3	18.303	15.862	17.254	24.109	17.024	47.542
	4	35.823	26.293	59.413	51.230	33.550	27.139
	6	3.046	3.433	2.508	4.202	6.333	3.374
	7	37.991	15.362	32.871	17.378	12.602	17.554
	8	967.351	981.354	1,006.263	1,008.591	1,015.204	1,017.635
10	49.439	48.150	52.716	49.276	47.979	83.543	

Type	UC ID	T71	T72	T73	T74	T75	T76
Deep Learning	2	27.333	28.859	70.147	71.523	28.612	26.186
	5	14.622	16.295	17.717	21.041	15.901	15.571
	9	246.619	204.798	226.842	200.042	224.066	268.584
Machine Learning	1	41.703	75.994	50.916	36.467	30.363	30.590
	3	12.000	25.214	21.493	28.575	52.888	20.269
	4	41.557	54.150	40.155	39.855	42.724	84.489
	6	6.083	4.316	6.447	3.604	6.055	3.740
	7	12.573	13.195	14.841	14.927	13.352	8.412
	8	1,008.735	1,029.479	984.200	1,036.808	1,032.338	996.970
	10	106.283	49.432	58.911	55.240	51.513	63.207

Type	UC ID	T77	T78	T79	T80	T81	T82
Deep Learning	2	33.766	38.015	30.358	26.436	35.595	43.978
	5	7.986	16.859	13.278	15.633	18.568	14.425
	9	225.784	304.445	302.869	218.562	236.309	248.590
Machine Learning	1	39.708	22.038	36.508	32.721	39.076	27.932
	3	50.565	18.981	25.779	16.955	17.470	16.308
	4	63.180	31.118	46.990	39.379	25.099	45.192
	6	6.127	5.012	1.991	4.070	3.900	6.286
	7	9.159	19.722	11.894	14.487	10.941	41.390
	8	1,052.376	1,010.565	1,004.174	1,047.601	1,034.143	992.305
	10	38.558	58.878	48.126	56.808	103.760	74.975

Type	UC ID	T83	T84	T85	T86	T87	T88
Deep Learning	2	69.164	25.079	28.698	35.052	67.190	30.296
	5	17.406	30.021	36.124	14.894	13.775	14.606
	9	266.529	223.959	255.789	259.446	260.290	196.723
Machine Learning	1	24.921	34.885	34.087	21.156	37.593	31.028
	3	20.487	18.732	36.608	43.036	20.017	19.137
	4	35.672	75.995	26.322	12.787	44.831	35.470
	6	4.981	3.806	6.390	5.728	5.832	4.078
	7	13.242	15.182	11.256	29.364	18.785	15.552
	8	1,019.058	1,002.633	960.944	1,059.470	964.533	1,127.864
	10	50.692	67.398	49.399	57.310	48.227	56.453

Type	UC ID	T89	T90	T91	T92	T93	T94
Deep Learning	2	14.105	39.226	24.844	27.889	29.814	26.209
	5	5.118	32.085	18.444	14.386	14.127	11.961
	9	262.767	246.996	221.708	219.926	209.456	293.994
Machine Learning	1	30.616	51.974	44.978	36.949	34.252	36.237
	3	6.234	15.699	20.689	16.749	20.700	32.134
	4	84.560	31.497	37.443	34.760	35.474	34.819
	6	8.906	5.413	3.635	2.979	3.745	4.139
	7	12.862	36.433	15.565	10.250	13.229	11.530
	8	1,055.359	994.780	992.454	1,081.813	1,022.481	1,001.530
	10	58.673	60.996	98.149	65.344	50.448	60.420

Type	UC ID	T95	T96	T97	T98	T99	T100
Deep Learning	2	13.453	25.975	68.421	65.112	28.957	17.911
	5	5.253	12.614	14.155	14.706	21.470	14.988
	9	243.510	242.691	209.458	222.704	222.732	303.508
Machine Learning	1	73.853	30.783	37.700	30.807	53.580	30.824
	3	19.409	16.707	22.548	24.758	54.344	18.282
	4	30.040	35.908	38.444	64.218	38.571	48.659
	6	1.268	2.409	5.869	4.182	4.771	5.378
	7	15.982	11.613	15.499	16.930	18.510	12.246
	8	1,057.803	1,089.713	988.685	1,032.633	1,024.152	1,019.145
	10	74.654	49.581	75.652	37.702	49.160	52.439

Table 3-2 Use Case Elapsed Times

3.4 SUT Validation Test Output

<u>Validation Run Report</u>			
AIUCpm@1	328.97	T _{Load}	0.41
Scale Factor	1	T _{LD}	0.41
Streams	100	T _{PTT}	23.32
Kit Version	1.0.3.1	T _{PST1}	3.46
Execution Status	Pass	T _{PST2}	3.45
Accuracy Status	Pass	T _{PST}	3.46
		T _{TT}	0.34
Test Times			
Overall Run Start Time	2023-12-11 01:38:51.916		
Overall Run End Time	2023-12-11 02:07:47.250		
Overall Run Elapsed Time	1,735.334		
Load Test Start Time	2023-12-11 01:40:06.326		
Load Test End Time	2023-12-11 01:40:06.746		
Load Test Elapsed Time	0.420		
Power Training Start Time	2023-12-11 01:40:06.747		
Power Training End Time	2023-12-11 01:57:08.522		
Power Training Elapsed Time	1,021.775		
Power Serving 1 Start Time	2023-12-11 01:57:08.523		
Power Serving 1 End Time	2023-12-11 01:58:27.017		
Power Serving 1 Elapsed Time	78.494		
Power Serving 2 Start Time	2023-12-11 01:58:27.019		
Power Serving 2 End Time	2023-12-11 01:59:45.546		
Power Serving 2 Elapsed Time	78.527		
Scoring Start Time	2023-12-11 02:00:20.534		
Scoring End Time	2023-12-11 02:02:02.926		
Scoring Elapsed Time	102.392		
Throughput Start Time	2023-12-11 02:02:02.942		
Throughput End Time	2023-12-11 02:07:47.247		
Throughput Elapsed Time	344.305		
(continued on next page)			

Validation Run Report (continued)

Accuracy Metrics					
Use Case	Metric Name	Metric	Criteria	Threshold	Status
1	N/A	0.000	N/A	0.00	Pass
2	word_error_rate	0.140	<=	0.50	Pass
3	mean_squared_log_error	4.582	<=	5.40	Pass
4	f1_score	0.701	>=	0.65	Pass
5	mean_squared_log_error	0.365	<=	0.50	Pass
6	matthews_corrcoef	0.439	>=	0.19	Pass
7	median_absolute_error	0.889	<=	1.80	Pass
8	accuracy_score	0.706	>=	0.65	Pass
9	accuracy_score	0.990	>=	0.90	Pass
10	accuracy_score	0.817	>=	0.70	Pass

3.5 Configuration Parameters

The [Supporting Files](#) archive contains all Global Benchmark Parameter and Use Case Specific Parameter settings.

Clause 4 – SUT Related Items

4.1 Specialized Hardware/Software

No Specialized Hardware/Software was used in the SUT.

4.2 Configuration Files

The [Supporting Files](#) archive contains all configuration files.

4.3 SUT Environment Information

All envInfo.log files are included in the [Supporting Files](#) archive.

4.4 Data Storage to Scale Factor Ratio

The details of the Data Storage Ratio are provided below.

Node Count	Disks	Size (GB)	Total (GB)
1	2	480	960

Total Storage (GB)	960
Scale Factor	10
Data Storage Ratio	96.00

4.5 Scale Factor to Memory Ratio

The details of the Memory to Scale Factor Ratio are provided below.

Nodes	Memory (GiB)	Total (GiB)
1	768	768

Scale Factor	10
Total Memory (GiB)	768
SF / Memory Ratio	0.01

4.6 Output of Tests

The [Supporting Files](#) archive contains the output files of all tests.

4.7 Additional Sponsor Files

The [Supporting Files](#) archive contains any additional files that were used.

4.8 Model Optimizations

The [Supporting Files](#) archive contains any model optimization files that were used.

Clause 5 – Metrics and Scale Factor

5.1 Reported Performance Metrics

Metric Overview

TPCx-AI Performance Metric	714.04	AIUCpm@10
TPCx-AI Price/Performance Metric	171.42	\$/AIUCpm@10
TPCx-AI Scale Factor	10	
TPCx-AI Stream Count	100	

Test Times


Overall Run Start Time	2023-12-11 02:08:12.360
Overall Run End Time	2023-12-11 04:21:01.262
Overall Run Elapsed Time	7,968.902
Load Test Start Time	2023-12-11 02:10:06.050
Load Test End Time	2023-12-11 02:10:08.319
Load Test Elapsed Time	2.269
Power Training Start Time	2023-12-11 02:10:08.320
Power Training End Time	2023-12-11 03:37:51.970
Power Training Elapsed Time	5,263.650
Power Serving 1 Start Time	2023-12-11 03:37:51.971
Power Serving 1 End Time	2023-12-11 03:45:18.417
Power Serving 1 Elapsed Time	446.446
Power Serving 2 Start Time	2023-12-11 03:45:18.418
Power Serving 2 End Time	2023-12-11 03:52:45.718
Power Serving 2 Elapsed Time	447.300
Scoring Start Time	2023-12-11 03:53:22.422
Scoring End Time	2023-12-11 03:55:12.316
Scoring Elapsed Time	109.894
Throughput Start Time	2023-12-11 03:55:12.332
Throughput End Time	2023-12-11 04:21:01.260
Throughput Elapsed Time	1,548.928

Accuracy Metrics

Use Case	Metric Name	Metric	Criteria	Threshold	Status
1	N/A	0.000	N/A	0.00	Pass
2	word_error_rate	0.453	<=	0.50	Pass
3	mean_squared_log_error	3.609	<=	5.40	Pass
4	f1_score	0.707	>=	0.65	Pass
5	mean_squared_log_error	0.486	<=	0.50	Pass
6	matthews_corrcoef	0.483	>=	0.19	Pass
7	median_absolute_error	1.032	<=	1.80	Pass
8	accuracy_score	0.733	>=	0.65	Pass
9	accuracy_score	1.000	>=	0.90	Pass
10	accuracy_score	0.816	>=	0.70	Pass

5.2 Throughput Test Stream Times

The following chart shows the minimum, 1st quartile, median, mean (X), 3rd quartile, and maximum stream times by use case for the Throughput Test. Outliers are marked with “o”.



Auditor's Information

This benchmark was audited by Doug Johnson, InfoSizing.

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

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

A copy of the auditor's attestation letter is included in the next two pages.



Rajesh Tadakamadla
 Hewlett Packard Enterprise
 Survey 192, Whitefield Road,
 Bangalore, India – 560 048

August 21, 2024

I verified the TPC Express Benchmark™ AI v1.0.3.1 performance of the following configuration:

Platform: 1x HPE ProLiant DL325 Gen11
 Operating System: Red Hat Enterprise Linux 8.6 (Ootpa)
 Additional Software: Anaconda Business

The results were:

Performance Metric 714.04 AIUCpm@10

Secondary Metrics	T _{LD}	2.26
	T _{PTT}	124.24
	T _{PST}	11.52
	T _{TT}	1.54

System Under Test 1x HPE ProLiant DL325 Gen11 with:

CPU	1x AMD EPYC 9374F 32-Core Processor		
Memory	768 GiB		
Storage	Qty	Size	Type
	2	480 GB	M.2 NVMe SSD

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 v1.0.3.1.
- All checksums were validated for compliance.
- Any modifications to shell scripts were reviewed for compliance.
- No modifications were made to any of the Java code.
- The generated dataset was properly scaled to 10 GB.
- The generated dataset used for testing was protected by RAID 1.

- 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 black ink that reads "Doug Johnson". The signature is written in a cursive style with a long horizontal flourish extending to the right.

Doug Johnson, Certified TPC Auditor

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

Third-Party Price Quotes

Anaconda



Anaconda Business: Support

To: Rajesh Tadakamadla, HPE and TPC Consortium | rajesh.tadakamadla@hpe.com
From: Nanette George, Staff Product Marketing Manager, Anaconda | ngeorge@anaconda.com
Re: Anaconda Support Quote for Publishing TPCx-AI Performance Benchmark Testing

Effective Date: August 2, 2024

Our reliable support team is one reason Anaconda is the platform of choice for one million organizations, including 93% of Fortune 500 companies.

The description of levels of service (below) applies to a **one-year subscription** to Anaconda Business, which includes support. Terms will remain valid for **120 days** following the above effective date.

Premium Support

Included Services	
Anaconda Business, 100 users, based on this configuration*	\$60,000
Anaconda Business Premium Support, including 24x7 support Includes acknowledgment within 4 hours of issues being reported. Once resolution is identified, it will be provided to the customer within 4 hours.	\$35,000
Named support contacts, including a dedicated technical account manager (TAM)	5

Software Packages Supported

This support quote applies to all packages and versions of open-source software packages included in the Anaconda repository: <https://repo.anaconda.com/pkgs/main>

HPE's TPCx-AI Benchmark Software Configuration

- Server Software
- RHEL Svr Sckt/2 Gst 1yr 24x7 E-LTU
- Anaconda Business 100 Subscription (100 users x 12 months x \$50)*

* This price is subject to change if the software configuration and/or number of users changes.



Anaconda Inc. • 901 E 6th St, Suite 102 • Austin, TX, 78702, USA
 Friedrichstrasse 123 • 10117 Berlin, Germany

© 2024 Anaconda Inc.


Hewlett Packard Inc.


EXPERT SALES HELP: 1-866-825-3906 ORDER STATUS CUSTOMER SERVICE DISCOUNT PROGRAMS HP FINANCING SIGN IN/REGISTER

hp Explore Shop Support


What are you looking for?  

HOME / BUSINESS / MONITORS / P-SERIES / HP P22V G5 FHD MONITOR





★★★★★ [Ask a question](#)

 Engineered for Sustainability

HP P22v G5 FHD Monitor


- FHD (1920 x 1080)
- 3000:1
- 5ms GTG (with overdrive)
- VA [See all Specs](#)

Can't find what you are looking for? [SEE SIMILAR PRODUCTS](#)

Price ~~\$109.99~~ **SAVE \$50.00 (45%)** [3% back in HP Rewards](#)

\$59.00

[ADD TO CART](#)

 **FREE SHIPPING** Estimated arrival date: Thu, Aug 08 - Wed, Aug 14

Supporting Files Index

The Supporting Files archive for this disclosure contains the following structure.

Supporting Files Directory	Description
CheckIntegrity/...	Output of CHECK_INTEGRITY test (if the phase is not done as part of the Validation and Performance Test).
PerformanceTest/...	Performance Test output files.
ValidationTest/...	Validation Test output files.
Additional files used by HPE	
Sponsor/ModelOptimization/...	Details of model optimization.
Sponsor/ModifiedKitFiles/...	0 modified file(s).
Sponsor/Tuning/...	All tuning files used.