Hewlett Packard Enterprise

Hewlett-Packard Enterprise

TPC Express Benchmark™ AI Full Disclosure Report

DL325 Gen11

with 1x ProLiant DL325 Gen11
using

Anaconda Pro

running on

Red Hat Enterprise Linux 8.6

TPCx-AI Version	1.0.2
Report Edition	First
Report Submitted	March 27, 2023

First Edition - March 2023

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 https://www.hpe.com/. Actual performance and environmental costs of HPE products will vary depending on individual customer configurations and conditions.

Abstract

HPE conducted the TPC Express Benchmark™ AI (TPCx-AI) on the DL325 Gen11. The software used included Anaconda Pro. This report provides full disclosure of the results. All testing was conducted in conformance with the requirements of the TPCx-AI Standard Specification, Revision 1.0.2.

Configuration Overview

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

Metrics Overview

Total System Cost	Performance	Price/Performance	Availability Date
$36,008 USD	517.77 AIUCpm@10	69.55 USD $/AIUCpm@10	March 27, 2023

Executive Summary

The Executive Summary follows on the next several pages.

Hewlett Packard Enterprise	DL325 Gen11	TPCx-AI 1.0.2 TPC Pricing 2.8.0 Report Date Mar. 27, 2023

TPCx-AI Performance	Total System Cost	Price/Performance	Availability Date
517.77 AIUCpm@10	**$36,008 USD**	**$69.55 USD/AIUCpm@10**	**March 27, 2023**

Framework	Operating System	Other Software	Scale Factor	Streams
Anaconda Pro	Red Hat Enterprise Linux 8.6	N/A	10	32

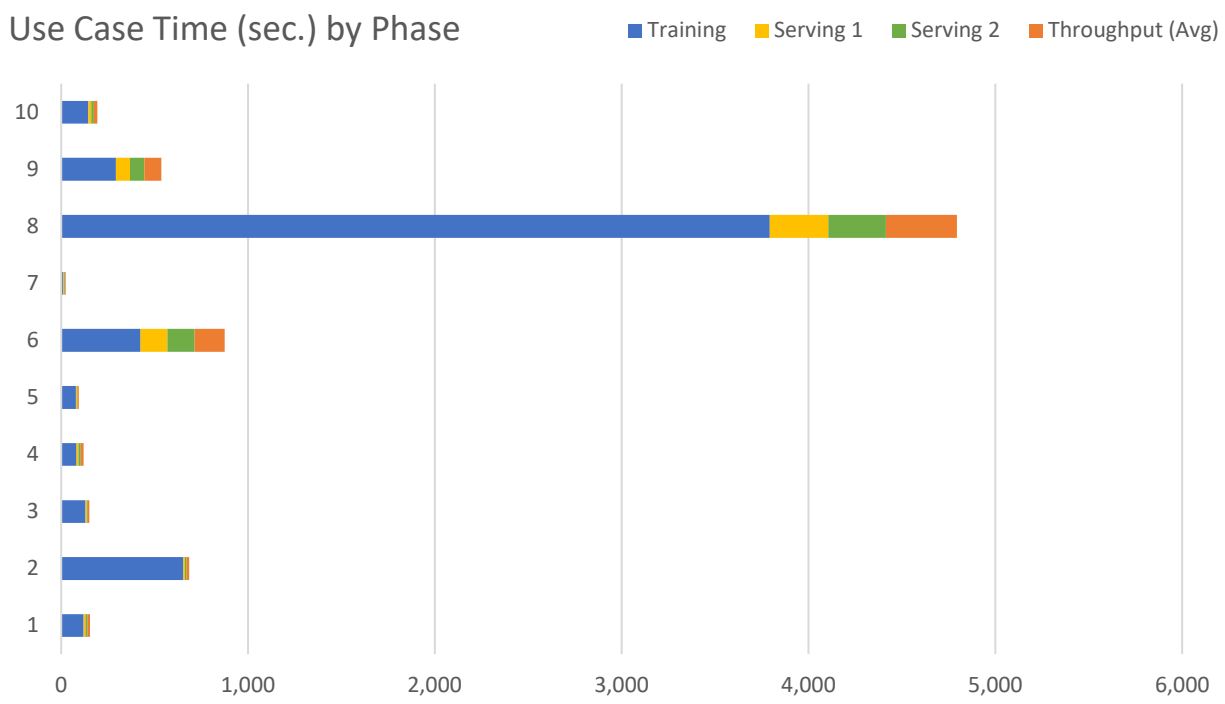

Physical Storage / Scale Factor	Scale Factor / Physical Memory	Main Data Redundancy Model
96.00	0.03	RAID 1

Servers: Total Processors/Cores/Threads	1 1 / 32 / 64

Server Type	1x ProLiant DL325 Gen11 (Server)
Processors	1x AMD EPYC 9374F 32-Core Processor GHz
Memory	384 GiB
Storage Controller	1x NS204i-u Gen11
Storage Device	2x 480 GB NVMe
Network Controller	1x Intel I350 4-port

Hewlett Packard Enterprise

DL325 Gen11

TPCx-AI	1.0.2
TPC Pricing	2.8.0
Report Date	Mar. 27, 2023

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	$11,087.00	1	$11,087.00	
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 32GB 2RX8 DDR5-4800 Smart kit	P50311-B21	1	$2,465.00	12	$29,580.00	
HPE 800W FS Plat Ht Plug LH PS Kit	P38995-B21	1	$524.00	2	$1,048.00	
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	$49,041.00	$2,958.00
Storage						
NS204i-u Gen11 OS Boot, RAID 1	P48183-B21	1	$2,005.00	1	$2,005.00	
HPE Gen 11 EI Rail Kit	P52351-B21	1	$34.48	1	$34.48	
HPE DL325G11 Standard Riser	standardDL385G11Riser	1		1	$0.00	
				Subtotal	$2,039.48	$0.00
Other						
HP V22v G5 FHD Monitor	65P56AA	3	$129.99	3	$389.97	
				Subtotal	**$389.97**	**$0.00**
Software Components						
RHEL Svr Sckt/2 Gst 1yr 24x7 E-LTU	J8J36A	1	$1,299.00	1	$1,299.00	
Anaconda Pro Subscription	N/A	2	$10,000.00	1		$10,000.00
				Subtotal	**$1,299.00**	**$10,000.00**
				Total Extended Price	$52,769.45	$12,958.00
				Total Discounts	$28,094.00	$1,626.00
				Grand Total	**$24,675.45**	**$11,332.00**

Pricing: 1 = HPE; 2 = Anaconda;3 = Hewlett Packard, Inc.

* Discount applies to all hardware 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):	**$36,008**
AIUCpm@10:	**517.77**
$/AIUCpm@10:	**$69.55**

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	DL325 Gen11	TPCx-AI 1.0.2 TPC Pricing 2.8.0 Report Date Mar. 27, 2023

Numerical Quantities

AIUCpm@10	**517.77**	T_{Load}	2.30
Scale Factor	10	T_{LD}	2.30
Streams	32	T_{PTT}	186.38
		T_{PST1}	18.46
Kit Version	1.0.2	T_{PST2}	18.41
Execution Status	Pass	T_{PST}	18.46
Accuracy Status	Pass	T_{TT}	2.28

Test Times

Overall Run Start Time	2023-02-16 11:38:00.964
Overall Run End Time	2023-02-16 13:50:01.515
Overall Run Elapsed Time	7,920.551
Load Test Start Time	2023-02-16 11:39:56.856
Load Test End Time	2023-02-16 11:39:59.168
Load Test Elapsed Time	2.312
Power Training Start Time	2023-02-16 11:39:59.169
Power Training End Time	2023-02-16 13:15:35.211
Power Training Elapsed Time	5,736.042
Power Serving 1 Start Time	2023-02-16 13:15:35.212
Power Serving 1 End Time	2023-02-16 13:25:26.333
Power Serving 1 Elapsed Time	591.121
Power Serving 2 Start Time	2023-02-16 13:25:26.334
Power Serving 2 End Time	2023-02-16 13:35:13.549
Power Serving 2 Elapsed Time	587.215
Scoring Start Time	2023-02-16 13:35:50.364
Scoring End Time	2023-02-16 13:37:51.334
Scoring Elapsed Time	120.970
Throughput Start Time	2023-02-16 13:37:51.340
Throughput End Time	2023-02-16 13:50:01.514
Throughput Elapsed Time	730.174

Hewlett Packard Enterprise	DL325 Gen11	TPCx-AI 1.0.2 TPC Pricing 2.8.0 Report Date Mar. 27, 2023

Numerical Quantities (continued)

Use Case Times & Accuracy

Use Case	Training (sec)	Serving 1 (sec)	Serving 2 (sec)	Throughput (avg)	Accuracy
UC01	120.225	10.382	10.332	13.605	0.000
UC02	654.090	8.617	8.655	14.603	0.425
UC03	131.280	5.670	5.582	8.084	3.609
UC04	82.538	10.747	10.884	15.894	0.707
UC05	79.180	4.409	4.382	6.639	0.465
UC06	424.810	144.534	144.927	161.298	0.548
UC07	10.852	3.824	3.792	5.171	1.031
UC08	3,793.418	312.561	308.296	381.139	0.733
UC09	293.481	75.727	75.702	91.608	1.000
UC10	146.105	14.586	14.600	19.272	0.816

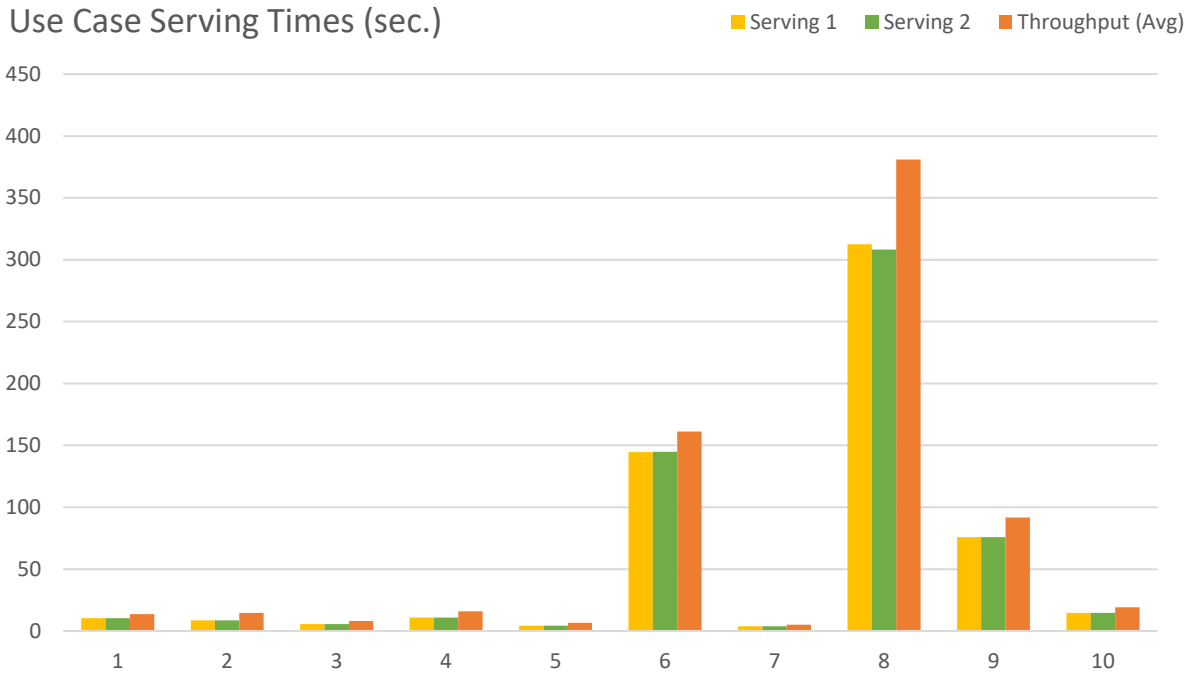

Table of Contents

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:	384 GiB	Storage Capacity:	960 GB

HPE ProLiant DL325 Gen11

HPE ProLiant DL325 Gen11
1 x AMD EPYC 9374F (3.85GHz/32-core)
384 GB MEM: 12x32GB DIMMs

1 x HPE NS204i-u Gen11 hot-plug NVMe boot device — Operation System, benchmark kit and data
2 x 480GB NVMe (RAID 1)

	Server
Server	1x ProLiant DL325 Gen11:
Procs/Cores/Threads:	1/32/64
Processor Model:	1x AMD EPYC 9374F 32-Core Processor
Memory:	384 GiB
Storage Controller:	1x NS204i-u Gen11
Storage Devices:	2x 480 GB NVMe
Network Controller:	1x Intel I350 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	tpcxai	All	2x 480 GB NVMe	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 / Anaconda Pro 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 Pro consisted of the following components.

Component	Version
python	3.9.13
setuptools	59.8.0
pandas	1.5.2
scikitlearn	1.2.0
xgboost	1.7.1
numpy	1.23.5
nose	1.3.7
scipy	1.10.0
statsmodels	0.13.5
patsy	0.5.2
tqdm	4.64.1
keras	2.10.0
tensorflow	2.10.0
joblib	1.1.0
pyyaml	6
jinja2	2.11.3
opencv	4.5.5

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 flies that were modified to facilitate system, platform, and framework differences.

TPCx-AI Kit Version	1.0.2
Modified File	Description of Changes
tools/python/dataRedundancyInformation.sh	Added platform-specific data collection.
tools/python/python-ks.yaml	Adjusted for software versions used.
See Auditor's Note	

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	8.617	8.655	21.386	9.890	12.644	10.408
	5	4.409	4.382	5.684	5.527	5.447	5.042
	9	75.727	75.702	108.437	90.716	87.077	89.901
Machine Learning	1	10.382	10.332	10.905	13.626	12.737	12.783
	3	5.670	5.582	7.461	5.845	22.654	6.581
	4	10.747	10.884	12.148	13.887	13.534	15.185
	6	144.534	144.927	156.821	157.690	155.881	157.512
	7	3.824	3.792	4.041	4.790	5.122	4.415
	8	312.561	308.296	383.778	406.595	392.254	406.838
	10	14.586	14.600	18.047	15.729	18.711	16.192

Type	UC ID	T5	T6	T7	T8	T9	T10
Deep Learning	2	34.067	11.262	11.813	15.320	11.730	11.610
	5	7.796	5.940	5.502	15.738	5.885	5.391
	9	87.542	96.318	95.666	98.246	86.503	86.420
Machine Learning	1	11.735	12.280	13.722	15.167	12.719	13.627
	3	6.109	10.114	6.896	7.214	7.091	5.837
	4	14.313	15.505	22.570	13.388	32.187	12.313
	6	167.247	166.060	163.339	160.712	158.537	158.669
	7	5.344	4.310	6.796	4.617	5.287	14.197
	8	364.918	354.933	356.812	360.454	379.134	395.192
	10	21.783	18.213	19.303	28.987	21.251	17.544

Type	UC ID	T11	T12	T13	T14	T15	T16
Deep Learning	2	11.546	9.453	10.897	9.985	10.577	33.284
	5	5.730	12.803	6.145	5.390	5.753	5.810
	9	115.733	86.959	86.728	89.944	97.228	88.893
Machine Learning	1	13.886	14.438	13.632	12.971	12.209	11.296
	3	5.801	21.552	6.353	7.620	6.944	8.254
	4	12.356	15.361	14.081	12.212	12.900	15.065
	6	156.086	157.926	178.449	157.945	173.573	157.514
	7	4.184	5.744	4.335	4.122	4.892	4.471
	8	377.413	374.115	371.131	404.284	366.816	377.410
	10	15.528	21.584	18.253	19.994	19.804	18.966

Type	UC ID	T17	T18	T19	T20	T21	T22
Deep Learning	2	13.666	11.888	10.891	12.208	12.849	13.399
	5	5.481	5.416	5.846	15.397	5.125	5.535
	9	94.611	87.591	90.015	87.144	88.088	87.994
Machine Learning	1	13.804	11.624	11.412	11.784	24.255	13.172
	3	6.032	7.627	7.535	6.626	8.198	6.212
	4	33.925	14.609	11.548	26.209	15.769	11.462
	6	164.788	157.007	158.893	165.506	166.147	157.203
	7	4.940	4.635	4.301	4.264	4.190	4.596
	8	367.598	408.372	392.175	375.966	367.240	409.107
	10	18.723	18.416	16.330	19.313	19.791	17.879

Type	UC ID	T23	T24	T25	T26	T27	T28
Deep Learning	2	14.719	11.042	11.046	38.227	12.167	9.963
	5	5.851	5.725	5.941	6.352	5.741	6.531
	9	96.941	85.317	90.260	96.040	86.418	95.141
Machine Learning	1	27.681	13.130	11.787	15.210	13.684	10.841
	3	6.070	6.309	5.828	6.073	6.692	6.976
	4	12.214	12.046	12.067	12.611	12.213	15.280
	6	158.331	158.739	156.749	167.721	157.545	166.614
	7	4.271	4.496	4.197	4.375	4.377	4.364
	8	381.525	376.939	410.331	361.747	403.174	363.567
	10	21.443	25.696	18.926	17.983	19.387	15.644

Type	UC ID	T29	T30	T31	T32
Deep Learning	2	11.048	12.345	25.141	10.816
	5	5.480	6.356	6.281	5.792
	9	89.063	87.866	87.682	88.964
Machine Learning	1	12.016	12.494	13.401	11.345
	3	13.608	7.956	8.385	6.234
	4	15.276	28.176	13.113	15.097
	6	164.208	156.205	164.840	157.065
	7	4.641	5.167	5.122	10.884
	8	369.059	380.251	363.724	393.601
	10	17.947	18.156	22.644	18.546

Table 3-2 Use Case Elapsed Times

3.4 SUT Validation Test Output

Validation Run Report

AIUCpm@1	309.49	T_{Load}	0.39
Scale Factor	1	T_{LD}	0.39
Streams	32	T_{PTT}	25.63
		T_{PST1}	3.39
Kit Version	1.0.2	T_{PST2}	3.37
Execution Status	Pass	T_{PST}	3.39
Accuracy Status	Pass	T_{TT}	0.42

Test Times

Overall Run Start Time	2023-02-16 11:07:44.381
Overall Run End Time	2023-02-16 11:35:19.158
Overall Run Elapsed Time	1,654.777
Load Test Start Time	2023-02-16 11:08:54.780
Load Test End Time	2023-02-16 11:08:55.175
Load Test Elapsed Time	0.395
Power Training Start Time	2023-02-16 11:08:55.176
Power Training End Time	2023-02-16 11:28:07.866
Power Training Elapsed Time	1,152.690
Power Serving 1 Start Time	2023-02-16 11:28:07.867
Power Serving 1 End Time	2023-02-16 11:29:26.517
Power Serving 1 Elapsed Time	78.650
Power Serving 2 Start Time	2023-02-16 11:29:26.518
Power Serving 2 End Time	2023-02-16 11:30:44.590
Power Serving 2 Elapsed Time	78.072
Scoring Start Time	2023-02-16 11:31:19.843
Scoring End Time	2023-02-16 11:33:02.834
Scoring Elapsed Time	102.991
Throughput Start Time	2023-02-16 11:33:02.840
Throughput End Time	2023-02-16 11:35:19.157
Throughput Elapsed Time	136.317

(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	0.338	<=	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	0.012	<=	0.50	Pass
6	matthews_corrcoef	0.462	>=	0.19	Pass
7	median_absolute_error	0.894	<=	1.80	Pass
8	accuracy_score	0.717	>=	0.65	Pass
9	accuracy_score	1.000	>=	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.

0	Memory (GiB)	Total (GiB)
1	384	384

Scale Factor	10
Total Memory (GiB)	384
SF / Memory Ratio	0.03

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	517.77	AIUCpm@10
TPCx-AI Price/Performance Metric	69.55	$/AIUCpm@10
TPCx-AI Scale Factor	10	
TPCx-AI Stream Count	32	

Test Times

Overall Run Start Time	2023-02-16 11:38:00.964
Overall Run End Time	2023-02-16 13:50:01.515
Overall Run Elapsed Time	7,920.551
Load Test Start Time	2023-02-16 11:39:56.856
Load Test End Time	2023-02-16 11:39:59.168
Load Test Elapsed Time	2.312
Power Training Start Time	2023-02-16 11:39:59.169
Power Training End Time	2023-02-16 13:15:35.211
Power Training Elapsed Time	5,736.042
Power Serving 1 Start Time	2023-02-16 13:15:35.212
Power Serving 1 End Time	2023-02-16 13:25:26.333
Power Serving 1 Elapsed Time	591.121
Power Serving 2 Start Time	2023-02-16 13:25:26.334
Power Serving 2 End Time	2023-02-16 13:35:13.549
Power Serving 2 Elapsed Time	587.215
Scoring Start Time	2023-02-16 13:35:50.364
Scoring End Time	2023-02-16 13:37:51.334
Scoring Elapsed Time	120.970
Throughput Start Time	2023-02-16 13:37:51.340
Throughput End Time	2023-02-16 13:50:01.514
Throughput Elapsed Time	730.174

Accuracy Metrics

Use Case	Metric Name	Metric	Criteria	Threshold	Status
1	N/A	0.000	N/A	0.00	Pass
2	word_error_rate	0.425	<=	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	0.465	<=	0.50	Pass
6	matthews_corrcoef	0.548	>=	0.19	Pass
7	median_absolute_error	1.031	<=	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”.

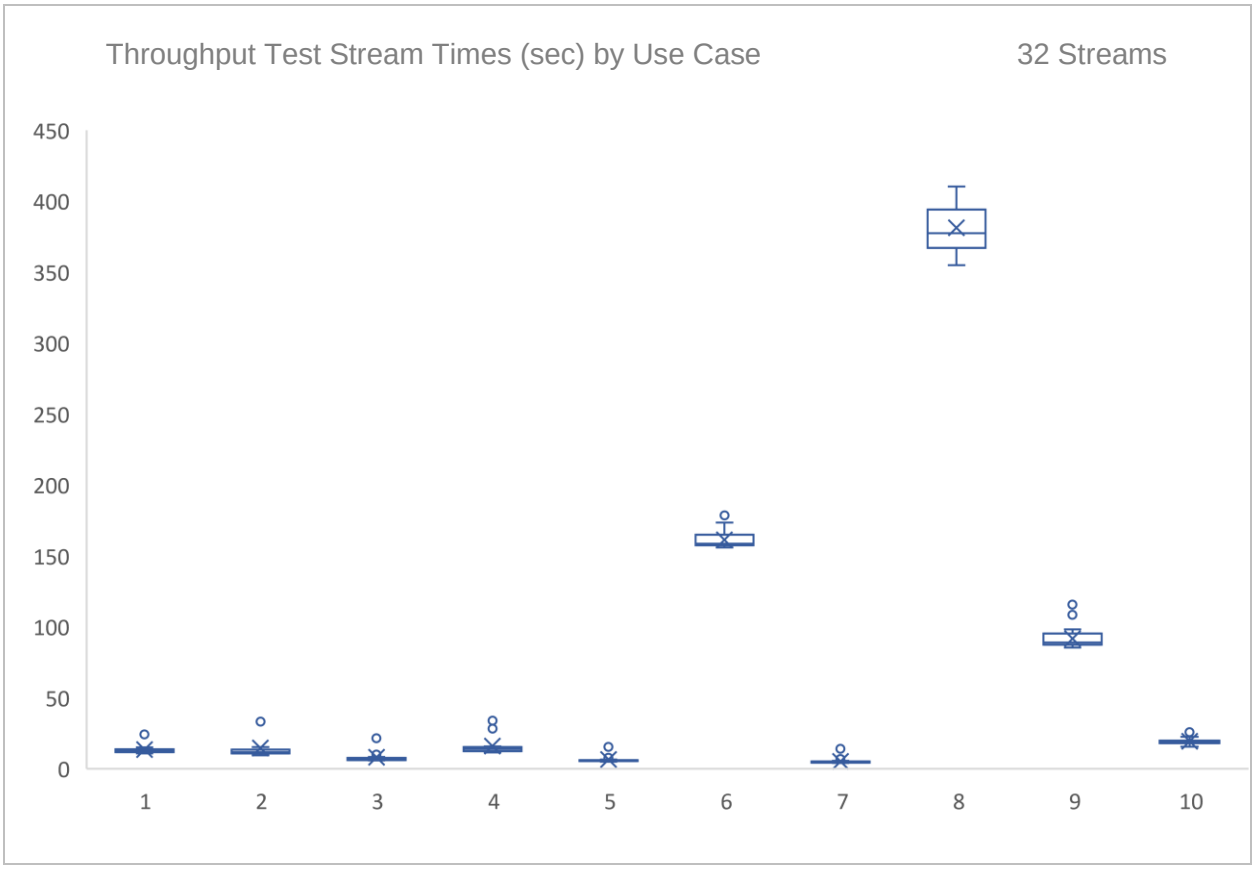

Auditor's Information

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.

InfoSizing
The Right Metric For Sizing IT

TPC
Certified Auditor

Ankit Chouksey
Hewlett-Packard Enterprise
192 Mahadevapura,
Whitefield Road
Bangalore, India 560048

March 23, 2023

I verified the TPC Express Benchmark™ AI v1.0.2 performance of the following configuration:

Platform:	1x DL325 Gen11
Operating System:	Red Hat Enterprise Linux 8.6
Additional Software:	Anaconda Pro

The results were:

Performance Metric	**517.77 AIUCpm@10**	
Secondary Metrics	T_{LD}	2.30
	T_{PTT}	186.38
	T_{PST}	18.46
	T_{TT}	2.28

System Under Test	**1x DL325 Gen11 with:**		
CPUs	1x AMD EPYC 9374F 32-Core Processor		
Memory	384 GiB		
Storage	***Qty***	***Size***	***Type***
	2	480 GB	NVMe

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 v1.0.2.
- 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.

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

- 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:

Two files were erroneously reported as having incorrect checksums. This is due to a minor issue in the TPC-provided kit. The TPCx-AI Subcommittee is aware of this and will correct it in a future release of the kit.

Respectfully Yours,

Doug Johnson, Certified TPC Auditor

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

Third-Party Price Quotes

Anaconda

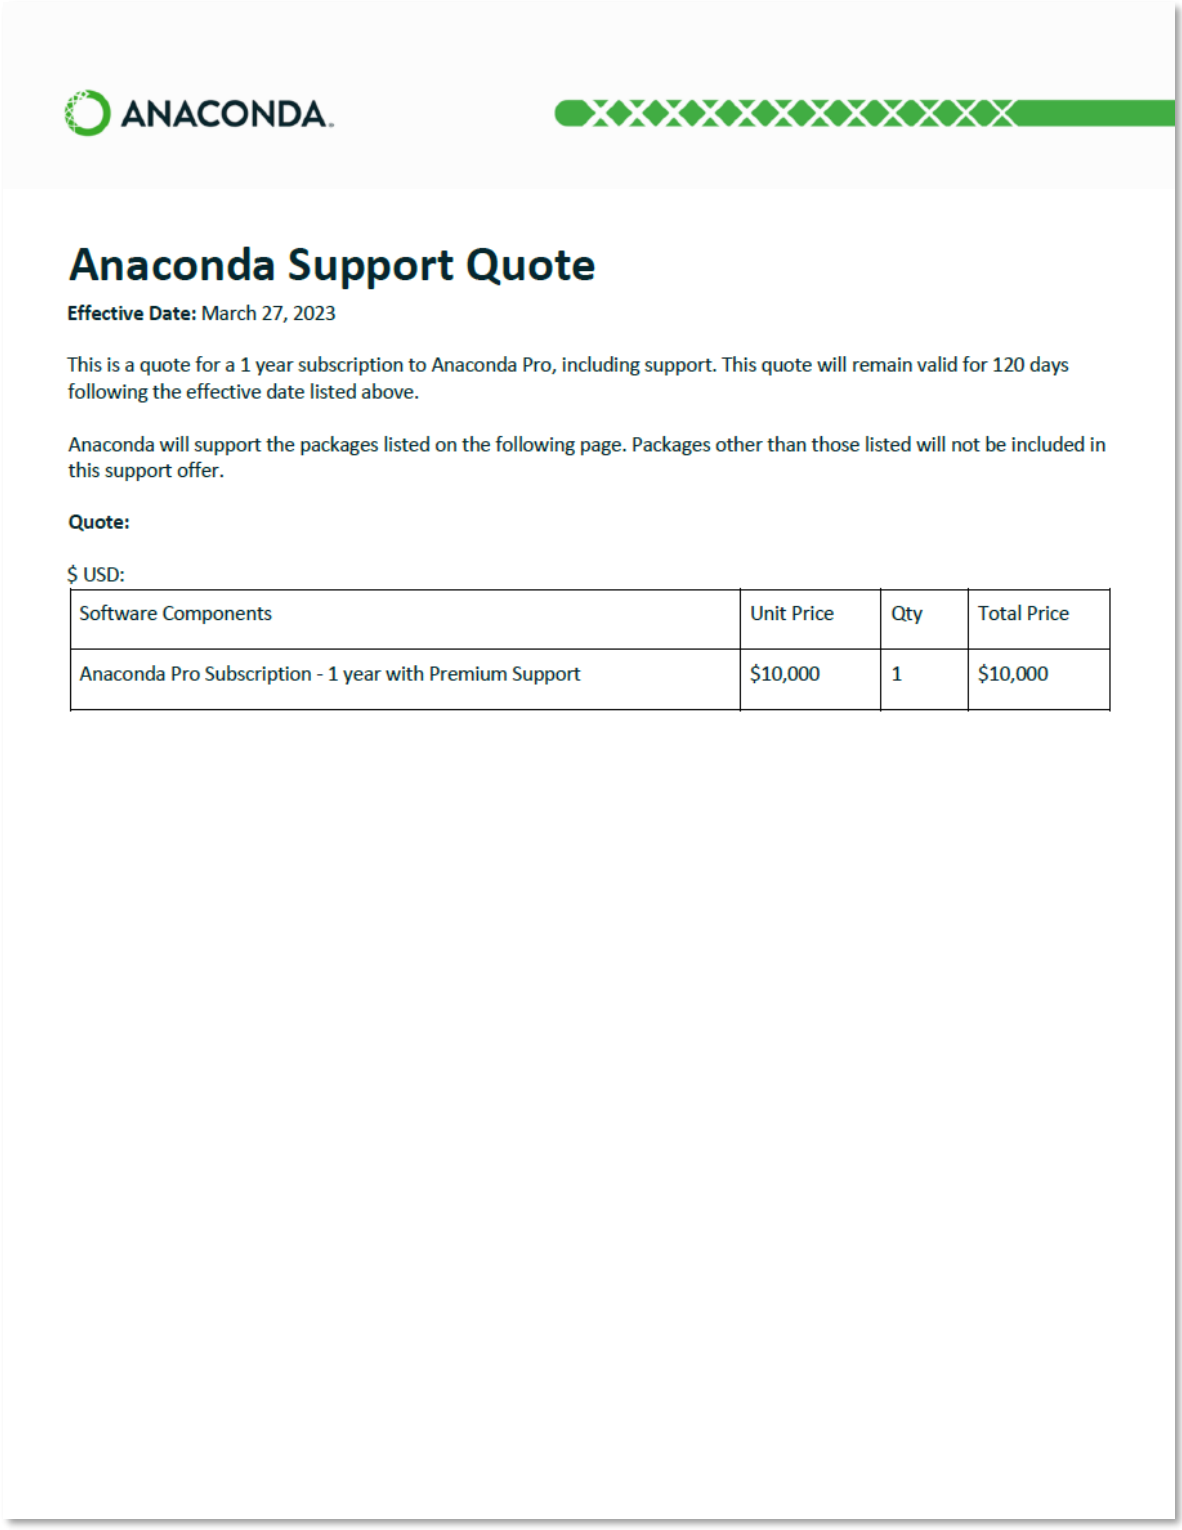

ANACONDA.

Anaconda Support Quote

Effective Date: March 27, 2023

This is a quote for a 1 year subscription to Anaconda Pro, including support. This quote will remain valid for 120 days following the effective date listed above.

Anaconda will support the packages listed on the following page. Packages other than those listed will not be included in this support offer.

Quote:

$ USD:

Software Components	Unit Price	Qty	Total Price
Anaconda Pro Subscription - 1 year with Premium Support	$10,000	1	$10,000

ANACONDA.

Included packages:

package name	source	version
python	main-anaconda	3.9.13
setuptools	main-anaconda	59.8.0
pandas	main-anaconda	1.5.2
scikitlearn	main-anaconda	1.2.0
xgboost	main-anaconda	1.7.1
numpy	main-anaconda	1.23.5
nose	main-anaconda	1.3.7
scipy	main-anaconda	1.10.0
statsmodels	main-anaconda	0.13.5
patsy	main-anaconda	0.5.2
tqdm	main-anaconda	4.64.1
keras	main-anaconda	2.10.0
tensorflow	main-anaconda	2.10.0
joblib	main-anaconda	1.1.0
pyyaml	main-anaconda	6
jinja2	main-anaconda	2.11.3
opencv	main-anaconda	4.5.5

Contact Sales: sales@anaconda.com | (512) 222-5440

Anaconda Inc.
1108 Lavaca Street Suite110-645
Austin, TX, 78701, USA

Hewlett Packard, Inc.

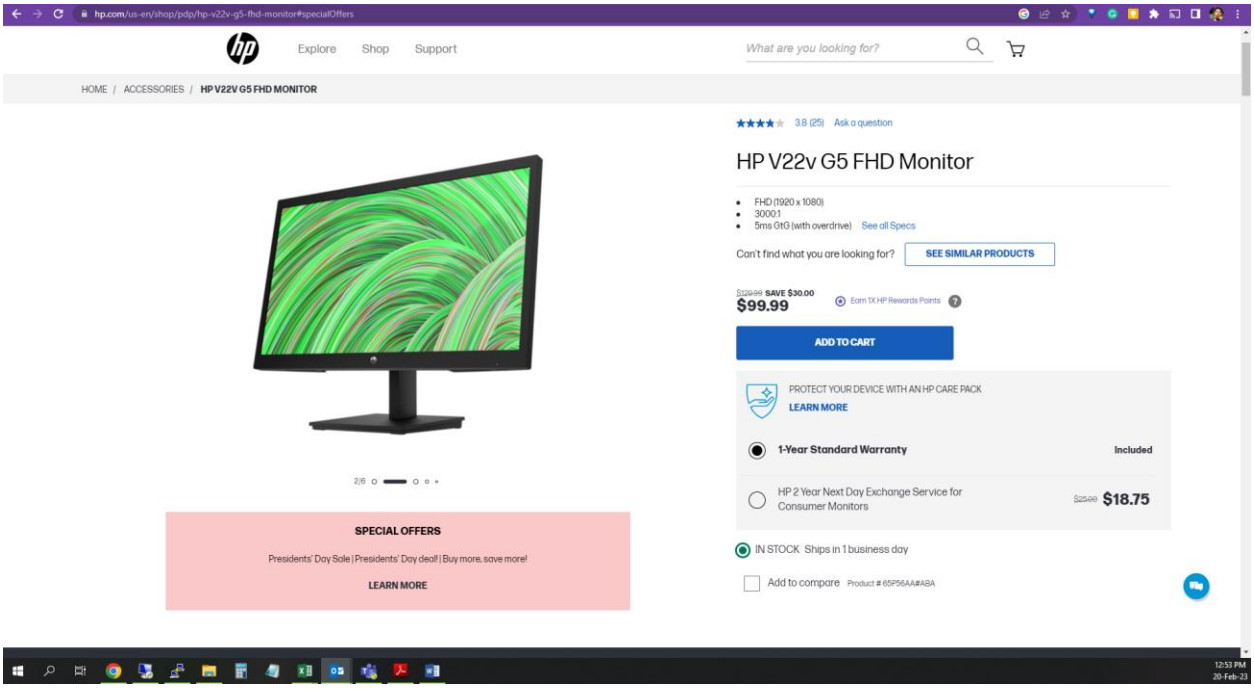

Supporting Files Index

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/…	2 modified file(s).
Sponsor/Tuning/…	All tuning files used.