

# TPC Express Benchmark™ AI Full Disclosure Report

## PowerEdge R6725

with 1x PowerEdge R6725  
using

### Anaconda Business

running on

## Red Hat Enterprise Linux 8.10 (Ootpa)

TPCx-AI Version  
Report Edition  
Report Submitted

2.0.0  
First  
October 10, 2024

**First Edition - October 2024**

*Dell Inc. (Dell), 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.*

*Dell and the Dell Logo are trademarks of Dell 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 Dell and any other company.*

*TPC Express Benchmark™ AI, TPCx-AI, and AIUCpm@30, are registered certification marks of the Transaction Processing Performance Council.*

*The Dell 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 Dell business contact for information on the products or services available in your area. You can find additional information via Dell's web site at [www.dell.com](http://www.dell.com). Actual performance and environmental costs of Dell products will vary depending on individual customer configurations and conditions.*

**Copyright© 2024 Dell 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

Dell conducted the TPC Express Benchmark™ AI (TPCx-AI) on the PowerEdge R6725. The software used included Anaconda Business. This report provides full disclosure of the results. All testing was conducted in conformance with the requirements of the TPCx-AI Standard Specification, Revision 2.0.0.

## Configuration Overview


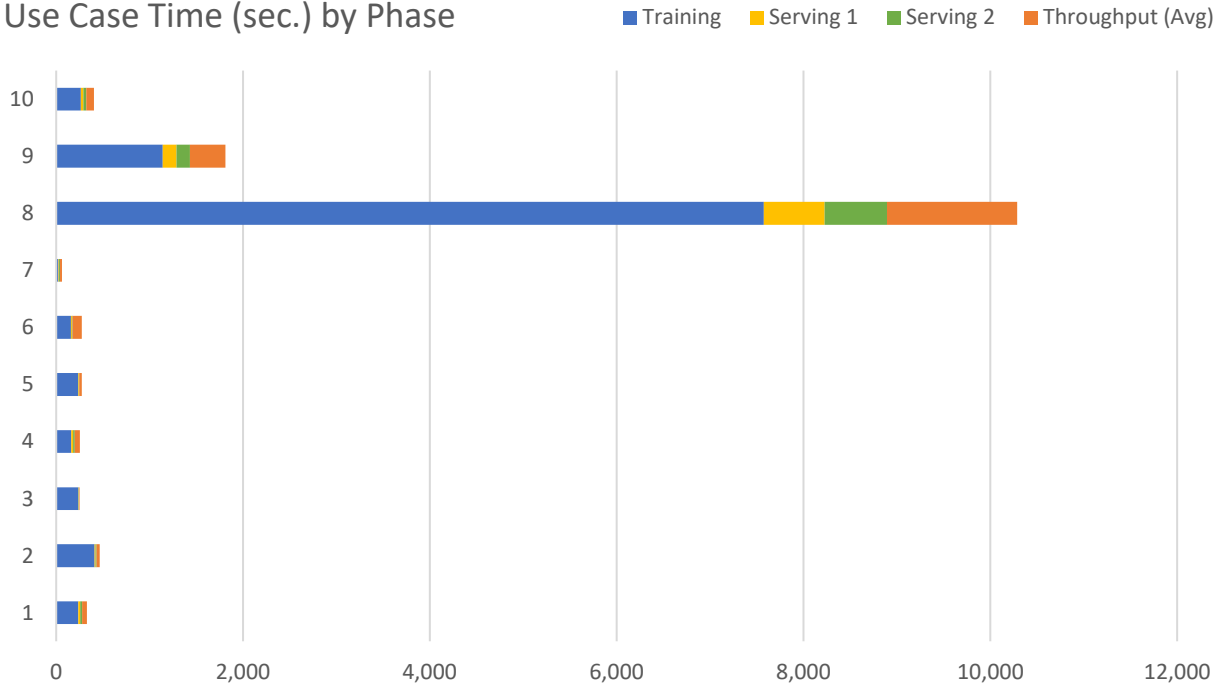
Test Sponsor	Node(s)	Operating System
Dell	1x PowerEdge R6725 (Server)	Red Hat Enterprise Linux 8.10 (Ootpa)


## Metrics Overview

Total System Cost	Performance	Price/Performance	Availability Date
\$140,934 USD	1,115.54 AIUCpm@30	126.34 USD \$/AIUCpm@30	November 9, 2024


# Executive Summary


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

		<h1>PowerEdge R6725</h1>		TPCx-AI 2.0.0 TPC Pricing 2.9.0 Report Date Oct. 10, 2024
TPCx-AI Performance <b>1,115.54 AIUCpm@30</b>	Total System Cost <b>\$140,934 USD</b>	Price/Performance <b>\$126.34 USD/AIUCpm@30</b>	Availability Date <b>November 9, 2024</b>	
Framework Anaconda Business	Operating System Red Hat Enterprise Linux 8.10 (Ootpa)	Other Software N/A	Scale Factor 30	Streams 100
<div style="display: flex; justify-content: space-between;"> <div> <h3>Use Case Time (sec.) by Phase</h3>  </div> <div> <span>■ Training</span> <span>■ Serving 1</span> <span>■ Serving 2</span> <span>■ Throughput (Avg)</span> </div> </div>				
Physical Storage / Scale Factor <b>32.00</b>	Scale Factor / Physical Memory <b>0.02</b>	Main Data Redundancy Model <b>RAID-1</b>		
Servers: Total Processors/Cores/Threads	1 2 / 64 / 128			
Server Type Processors Memory Storage Controller Storage Device Network Controller	1x PowerEdge R6725 (Server) 2x AMD EPYC 9355 32-Core Processor 1,536 GiB 1x BOSS-N1 2x 480 GB M.2 NVMe 1x Broadcom 57504 Quad Port 10/25GbE			

		<h1>PowerEdge R6725</h1>			TPCx-AI	2.0.0
					TPC Pricing	2.9.0
					Report Date	Oct. 10, 2024
Description	Part Number	Source	List Price	Qty	Extended Price	1-Yr. Maintenance
<b>Hardware</b>						
PowerEdge R6725 Server	210-BNNF	1	\$148,802.14	1	\$148,802.14	
High Performance Computing Cluster Information SKU	463-7922	1	0	1		
No Backplane	379-BDSY	1	0	1		
No Backplane	379-BDSV	1	0	1		
No Rear Storage	379-BDTE	1	0	1		
Trusted Platform Module 2.0 V3	461-AAIG	1	0	1		
No HD, No Backplane	321-BIGG	1	0	1		
AMD EPYC 9355 3.40GHz, 32C/64T, 256 Cache (280W) DDR5-6000	338-CRCD	1	0	2		
Additional Processor Selected	379-BDCO	1	0	1		
Heatsink for 2 CPU configuration	412-ABEE	1	0	1		
Performance Optimized	370-AHLL	1	0	1		
6400MT/s RDIMMs	370-BCCX	1	0	1		
64GB RDIMM, 6400MT/s, Dual Rank	370-BCCZ	1	0	24		
Diskless Configuration (No RAID, No Controller)	780-BCDH	1	0	1		
No Controller	405-AACD	1	0	1		
No Hard Drive	400-ABHL	1	0	1		
Performance BIOS Settings	384-BBBL	1	0	1		
No Energy Star	387-BBEY	1	0	1		
UEFI BIOS Boot Mode with GPTPartition	800-BBDM	1	0	1		
High Performance Fan x4	384-BD HQ	1	0	1		
Dual, Redundant (1+1),Hot-Plug MHS Power Supply, 1500W MM, Titanium	450-BCXC	1	0	1		
CORD,PWR,250V,13A,C15/14,NA,2M	450-AMKV	1	0	2		
Riser Config 2, 3 x16 LP (2xGen5)	330-BCCT	1	0	1		
PowerEdge R6725 Motherboard	329-BKMV	1	0	1		
Broadcom 57504 Quad Port 10/25GbE,SFP28, OCP NIC 3.0	540-BCRX	1	0	1		
Broadcom 5720 Dual Port 1GbE LOM	540-BDKD	1	0	1		
No Bezel	350-BBBW	1	0	1		
TRAY,W/LBLS,X8/X10,R6725	321-BIGJ	1	0	1		
BOSS-N1 controller card + with 2 M.2 480GB (RAID 1) (22x80)	403-BDMM	1	0	1		
No Operating System, No Utility Partition, BOSS	611-BBBX	1	0	1		
No Media Required	605-BBFN	1	0	1		
iDRAC10, Enterprise 17G	634-CSHV	1	0	1		
Secured Component Verification	528-COYT	1	0	1		
No Quick Sync	350-BBXM	1	0	1		
iDRAC,Legacy Password	379-BCSG	1	0	1		
iDRAC Service Module (ISM), NOT Installed	379-BCQX	1	0	1		
iDRAC Group Manager, Disabled	379-BCQY	1	0	1		
ReadyRails Sliding Rails Without Cable Management Arm or Strain Relief Bar	770-BECD	1	0	1		
No Systems Documentation, NoOpenManage DVD Kit	631-AACK	1	0	1		
PowerEdge R6725 Shipping	340-DDEC	1	0	1		
PowerEdge R6725 Shipping Material 8	340-DDBZ	1	0	1		
PowerEdge R6725 No CCC or CEMarking	470-AFOY	1	0	1		
None Required	817-BBBP	1	0	1		
Thank you choosing Dell ProSupport. For tech support, visit //www.dell.com/support or call 1-800- 945-3355	989-3439	1	0	1		
Dell Hardware Limited Warranty Plus On-Site Service	887-0748		\$249.00	1		\$249.00
3Yr ProSupport and 4hr Mission Critical - 3 Years	199-BONO		\$6,373.67	1		\$6,373.67
Dell 24 Monitor – S2425H	210-BMGX	1	\$99.99	1	\$99.99	
Keyboard/Mouse (Included with Server)	N/A	1	0	1		
			<b>Subtotal</b>		\$148,902.13	\$6,622.67

(continued on the next page)

	<h2>PowerEdge R6725</h2>	TPCx-AI 2.0.0
		TPC Pricing 2.9.0
		Report Date Oct. 10, 2024
(continued from the previous page)		
<b>Software</b>		
Anaconda Business, 85 users	N/A	2 \$51,000.00 1 \$51,000.00
Anaconda Business Premium Support, including 24x7 support	N/A	2 \$35,000.00 1 \$35,000.00
Red Hat Enterprise Linux, 2SKT, 1 Physical OR 2Guest, 1Yr PREMIUM SUB, No Media, CUS	528-BHPJ	1 \$1,428.90 1 \$1,428.90
		<b>Subtotal</b> \$52,428.90 \$35,000.00
		<b>Total</b> \$201,331.03 \$41,622.67
Large Purchase Discount (65%)*		-\$97,715.17 -\$4,304.74
Pricing: 1 = Dell; 2 = Anaconda	<b>Total System Cost (USD): \$140,934</b>	
* Discount applies to all line items where Key = 1. Discount based upon total system cost as purchased by a regular customer.	<b>AIUCpm@30: 1,115.54</b>	
<b>Audited by Doug Johnson, InfoSizing</b>	<b>\$/AIUCpm@30: \$126.34</b>	
<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 <a href="mailto:pricing@tpc.org">pricing@tpc.org</a>. Thank you.</i>		

	<h1 style="text-align: center;">PowerEdge R6725</h1>		TPCx-AI	2.0.0
			TPC Pricing	2.9.0
			Report Date	Oct. 10, 2024
<u>Numerical Quantities</u>				
<b>AIUCpm@30</b>	<b>1,115.54</b>	$T_{Load}$	4.50	
Scale Factor	30	$T_{LD}$	4.50	
Streams	100	$T_{PTT}$	310.45	
		$T_{PST1}$	20.75	
Kit Version	2.0.0	$T_{PST2}$	20.71	
Execution Status	Pass	$T_{PST}$	20.75	
Accuracy Status	Pass	$T_{TT}$	2.34	
Test Times				
Overall Run Start Time	2024-09-24 18:48:34.282			
Overall Run End Time	2024-09-24 22:58:26.260			
Overall Run Elapsed Time	14,991.978			
Load Test Start Time	2024-09-24 18:52:31.579			
Load Test End Time	2024-09-24 18:52:36.092			
Load Test Elapsed Time	4.513			
Power Training Start Time	2024-09-24 18:52:36.093			
Power Training End Time	2024-09-24 21:46:48.399			
Power Training Elapsed Time	10,452.306			
Power Serving 1 Start Time	2024-09-24 21:46:48.401			
Power Serving 1 End Time	2024-09-24 22:01:49.804			
Power Serving 1 Elapsed Time	901.403			
Power Serving 2 Start Time	2024-09-24 22:01:49.806			
Power Serving 2 End Time	2024-09-24 22:17:04.035			
Power Serving 2 Elapsed Time	914.229			
Scoring Start Time	2024-09-24 22:17:45.031			
Scoring End Time	2024-09-24 22:19:22.651			
Scoring Elapsed Time	97.620			
Throughput Start Time	2024-09-24 22:19:22.664			
Throughput End Time	2024-09-24 22:58:26.255			
Throughput Elapsed Time	2,343.591			



# PowerEdge R6725


TPCx-AI 2.0.0  
 TPC Pricing 2.9.0  
 Report Date Oct. 10, 2024

Numerical Quantities (continued)

Use Case Times & Accuracy

Use Case	Training (sec)	Serving 1 (sec)	Serving 2 (sec)	Throughput (avg)	Accuracy
UC01	238.819	20.249	20.181	51.833	0.000
UC02	410.288	10.846	10.838	34.377	0.445
UC03	236.639	3.203	3.246	8.873	3.553
UC04	162.736	18.751	18.613	53.225	0.706
UC05	236.712	5.886	5.753	27.051	0.036
UC06	159.143	8.936	8.936	96.813	0.624
UC07	25.692	7.528	7.467	20.257	1.002
UC08	7,577.223	650.287	665.046	1,394.051	0.754
UC09	1,141.745	145.657	144.166	380.100	1.000
UC10	263.225	29.968	29.897	82.927	0.817

Use Case Serving Times (sec.)





# Table of Contents

Abstract.....	3
Executive Summary .....	3
Table of Contents.....	9
Clause 0 – Preamble .....	11
0.1 TPC Express Benchmark™ AI Overview.....	11
Clause 1 – General Items .....	13
1.1 Test Sponsor .....	13
1.2 Parameter Settings.....	13
1.3 Configuration Diagrams .....	13
1.3.1 Measured Configuration.....	14
1.3.2 Differences Between the Measured and the Priced Configurations.....	14
Clause 2 – SW Components & Data Distribution .....	15
2.1 Roles and Dataset Distribution .....	15
2.2 File System Implementation.....	15
2.3 Execution Engine, Frameworks, Driver & Libraries .....	15
2.4 Applied Patches.....	15
Clause 3 – Workload Related Items.....	16
3.1 Hardware & Software Tuning.....	16
3.2 Kit Version & Modifications .....	16
3.3 Use Case Elapsed Times .....	16
3.4 SUT Validation Test Output .....	23
3.5 Configuration Parameters.....	24
Clause 4 – SUT Related Items.....	25
4.1 Specialized Hardware/Software .....	25
4.2 Configuration Files.....	25
4.3 SUT Environment Information.....	25
4.4 Data Storage to Scale Factor Ratio .....	25
4.5 Scale Factor to Memory Ratio .....	25
4.6 Output of Tests .....	25
4.7 Additional Sponsor Files .....	25
4.8 Model Optimizations .....	25
Clause 5 – Metrics and Scale Factor .....	26
5.1 Reported Performance Metrics .....	26

5.2 Throughput Test Stream Times .....27

Auditor's Information .....28

Third-Party Price Quotes.....31

    Anaconda .....31

Supporting Files Index .....32

# 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/](http://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](http://www.tpc.org).

# Clause 1 – General Items

## 1.1 Test Sponsor

This benchmark was sponsored by Dell Inc. and Advanced Micro Devices, Inc.

## 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:	2/64/128	Storage Devices: 2
Total Memory:	1,536 GiB	Storage Capacity: 960 GB

The diagram illustrates the hardware configuration of a server node. It features two AMD EPYC 9355 32-Core Processors (CPU 0 and CPU 1) connected via xGMI. CPU 0 is connected to DIMM 0, 1, 2, 3, 11, and storage (2x 480GB M.2 NVMe SSD, RAID1, DELL BOSS N1 (OS+Kit)). CPU 1 is connected to DIMM 12, 13, 14, 15, and DIMM 23. A Network Interface Card (NIC) is also connected to CPU 0.

<u>Server</u>	
Server	1x PowerEdge R6725
Procs/Cores/Threads:	2/32/64
Processor Model:	2x AMD EPYC 9355 32-Core Processor
Memory:	1,536 GiB
Storage Controller:	1x BOSS-N1
Storage Devices:	2x 480 GB M.2 NVMe
Network Controller:	1x Broadcom 57504 Quad Port 10/25GbE

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 PowerEdge R6725	idrac-cln2317-os	All	2x 480 GB M.2 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 8.10 (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

*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	2.0.0
<u>Modified File</u> tools/python/dataRedundancyInformation.sh	<u>Description of Changes</u> Added platform-specific data collection.

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	10.846	10.838	17.632	33.990	30.840	36.957
	5	5.886	5.753	34.490	34.379	16.557	17.050
	9	145.657	144.166	324.458	343.817	474.405	383.030
Machine Learning	1	20.249	20.181	46.820	64.351	44.202	57.389
	3	3.203	3.246	8.862	9.357	3.378	11.365
	4	18.751	18.613	52.138	57.643	52.170	61.003
	6	8.936	8.936	119.114	73.531	97.649	85.205
	7	7.528	7.467	21.129	18.896	17.120	17.977
	8	650.287	665.046	1,531.671	1,563.492	1,512.160	1,506.299
	10	29.968	29.897	107.023	106.190	76.111	99.308



Type	UC ID	T5	T6	T7	T8	T9	T10
Deep Learning	2	25.354	30.649	23.837	38.704	43.632	30.054
	5	32.987	35.527	11.367	12.459	41.143	29.545
	9	461.289	471.403	322.899	321.416	397.074	414.976
Machine Learning	1	54.253	39.330	58.344	51.443	64.257	40.368
	3	10.001	6.987	6.655	6.956	7.797	20.016
	4	52.412	39.333	41.354	49.680	50.275	53.221
	6	90.446	70.135	111.377	83.776	114.881	102.614
	7	20.455	19.137	35.653	18.252	24.787	8.445
	8	1,422.056	1,469.261	1,193.184	1,527.956	1,388.048	1,516.823
	10	64.515	87.941	59.957	92.838	64.301	98.022

Type	UC ID	T11	T12	T13	T14	T15	T16
Deep Learning	2	27.154	54.236	35.077	35.167	22.798	13.948
	5	60.471	16.415	17.117	37.903	18.541	40.557
	9	311.049	403.584	390.918	388.646	330.135	362.270
Machine Learning	1	55.659	45.579	49.915	41.523	49.436	56.049
	3	7.045	8.419	7.275	7.439	14.937	5.008
	4	54.434	68.978	57.304	64.586	54.278	34.542
	6	113.788	118.196	85.439	107.637	73.569	104.168
	7	13.634	15.276	21.596	19.845	16.900	23.489
	8	1,495.993	1,230.965	1,472.381	1,491.544	1,470.891	1,549.358
	10	65.955	71.789	75.232	101.990	95.107	113.415

Type	UC ID	T17	T18	T19	T20	T21	T22
Deep Learning	2	31.490	26.348	32.192	26.054	30.687	33.188
	5	15.202	13.579	41.762	7.866	11.753	18.238
	9	421.365	423.468	320.790	468.227	454.586	396.985
Machine Learning	1	62.856	43.378	54.877	56.486	44.275	56.386
	3	10.002	8.883	10.362	6.454	9.394	7.312
	4	44.521	61.779	63.964	43.308	50.726	54.903
	6	83.727	99.266	105.717	86.420	74.263	77.819
	7	16.077	18.710	25.578	20.581	21.072	19.197
	8	1,496.417	1,510.967	1,534.923	1,488.942	1,480.269	1,174.716
	10	80.069	117.922	79.885	70.077	79.912	72.739

Type	UC ID	T23	T24	T25	T26	T27	T28
Deep Learning	2	29.201	28.995	29.368	17.501	28.705	28.572
	5	41.762	22.032	20.357	11.611	13.494	11.561
	9	349.497	370.429	410.243	450.776	400.004	339.105
Machine Learning	1	42.638	57.337	64.427	56.138	67.126	42.464
	3	6.886	6.940	9.705	6.778	10.358	6.561
	4	57.830	55.377	50.771	40.072	49.998	38.461
	6	122.091	116.696	114.900	68.482	119.403	66.248
	7	19.968	19.816	16.861	16.801	14.958	17.964
	8	1,193.054	1,436.397	1,371.761	1,535.402	1,196.160	1,461.429
	10	61.017	87.932	91.920	85.038	55.922	73.035

Type	UC ID	T29	T30	T31	T32	T33	T34
Deep Learning	2	29.901	22.095	42.659	41.998	28.890	41.041
	5	31.234	16.448	49.831	47.102	34.700	22.831
	9	375.405	406.597	394.607	379.986	298.588	452.106
Machine Learning	1	42.268	42.702	37.450	41.376	57.072	43.065
	3	9.014	8.537	5.976	5.444	7.239	8.555
	4	59.135	32.959	43.727	61.467	61.094	54.054
	6	88.132	92.212	101.523	73.688	86.944	91.093
	7	20.482	27.145	15.032	20.171	18.160	26.335
	8	1,464.164	1,573.739	1,164.860	1,414.098	1,424.347	1,455.680
	10	93.513	73.723	68.875	70.351	91.158	61.220

Type	UC ID	T35	T36	T37	T38	T39	T40
Deep Learning	2	44.494	31.566	32.331	44.574	40.881	28.813
	5	17.104	24.632	26.164	33.752	24.045	11.808
	9	392.587	353.422	399.229	393.275	390.865	432.955
Machine Learning	1	48.885	59.632	62.629	82.605	45.541	62.738
	3	8.324	7.923	16.777	8.576	7.315	7.615
	4	71.969	68.324	57.139	73.031	44.276	46.937
	6	115.179	77.047	54.863	60.324	94.742	76.618
	7	20.002	24.016	27.275	30.734	20.350	17.741
	8	1,525.448	1,193.209	1,578.175	1,466.876	1,504.399	1,266.029
	10	93.933	115.792	66.132	97.250	82.074	66.474

Type	UC ID	T41	T42	T43	T44	T45	T46
Deep Learning	2	30.312	40.664	42.293	26.568	39.088	34.126
	5	47.479	30.118	8.544	31.309	31.039	19.952
	9	331.757	322.247	421.189	435.031	451.324	313.893
Machine Learning	1	41.303	44.338	36.593	54.578	49.751	69.717
	3	7.515	10.181	8.822	8.779	6.532	5.510
	4	72.594	48.128	58.169	70.800	45.960	41.278
	6	94.910	120.006	108.032	54.404	91.730	117.917
	7	16.924	19.679	18.408	17.417	28.424	14.417
	8	1,335.260	1,195.769	1,507.928	1,502.150	1,435.427	1,450.963
	10	113.849	92.838	69.501	91.552	72.225	90.767

Type	UC ID	T47	T48	T49	T50	T51	T52
Deep Learning	2	38.089	35.772	30.292	25.300	27.289	23.366
	5	23.893	34.899	13.683	10.948	13.163	17.116
	9	343.761	372.069	301.711	338.349	381.695	334.641
Machine Learning	1	57.742	65.141	62.484	37.643	56.638	44.304
	3	10.716	7.101	5.821	7.169	8.588	6.971
	4	62.757	61.248	42.759	33.318	50.310	40.854
	6	71.012	126.754	80.455	74.491	94.703	105.594
	7	20.345	29.891	17.707	17.832	20.836	16.724
	8	1,089.114	1,131.663	1,495.714	1,237.955	1,411.238	1,480.315
	10	97.395	110.787	88.032	69.135	87.742	76.278

Type	UC ID	T53	T54	T55	T56	T57	T58
Deep Learning	2	32.888	43.281	16.912	19.936	60.913	35.408
	5	49.720	56.275	15.970	11.206	17.167	53.381
	9	396.519	352.305	440.308	338.878	329.659	340.341
Machine Learning	1	50.766	29.391	45.311	39.686	50.046	67.230
	3	10.017	12.389	7.923	7.833	7.302	10.373
	4	65.612	53.728	41.070	40.606	57.289	71.829
	6	138.192	112.862	98.318	123.349	109.346	112.820
	7	18.120	16.225	28.002	18.226	14.903	21.359
	8	1,151.269	1,545.535	1,536.157	1,234.982	1,419.676	1,540.273
	10	76.310	80.048	62.271	65.046	96.067	48.807

Type	UC ID	T59	T60	T61	T62	T63	T64
Deep Learning	2	42.475	35.995	29.423	30.706	58.657	26.390
	5	18.853	54.716	32.400	11.935	24.667	31.389
	9	366.205	314.500	386.404	322.931	463.215	354.013
Machine Learning	1	62.276	51.349	46.041	49.176	70.659	34.515
	3	10.108	10.855	10.589	6.278	7.303	7.815
	4	72.948	64.890	65.136	40.520	43.814	45.961
	6	137.052	90.958	119.550	107.958	100.574	124.337
	7	29.195	28.722	22.562	15.139	15.289	17.735
	8	1,410.313	1,233.589	1,403.323	1,497.694	1,142.481	1,500.525
	10	90.989	88.382	91.462	72.714	71.939	97.832

Type	UC ID	T65	T66	T67	T68	T69	T70
Deep Learning	2	41.945	56.903	56.577	39.350	24.812	51.351
	5	34.686	15.737	16.247	29.170	18.432	42.937
	9	327.408	316.462	393.775	420.258	337.852	361.382
Machine Learning	1	44.108	72.289	48.259	57.048	47.132	55.636
	3	8.611	12.230	13.826	8.893	10.693	6.889
	4	51.525	45.248	50.735	66.284	47.741	45.822
	6	131.994	108.549	93.280	69.267	111.808	90.021
	7	22.182	19.173	20.535	25.012	18.169	21.095
	8	1,511.196	1,392.062	1,188.354	1,454.699	1,462.831	1,468.944
	10	75.715	77.396	116.834	96.808	84.237	107.311

Type	UC ID	T71	T72	T73	T74	T75	T76
Deep Learning	2	34.479	18.214	42.768	44.480	41.291	37.647
	5	17.842	34.143	20.355	66.394	11.213	55.736
	9	371.438	348.964	330.519	352.922	332.967	453.268
Machine Learning	1	62.894	66.582	65.783	50.930	41.274	71.781
	3	8.945	13.978	10.391	8.482	8.163	8.286
	4	46.930	66.618	50.369	59.000	46.568	59.047
	6	115.543	90.588	107.747	73.990	128.011	95.591
	7	19.444	21.757	21.800	28.112	17.303	20.372
	8	1,191.385	1,517.110	1,083.103	1,443.550	1,226.701	1,230.454
	10	95.944	85.924	79.596	94.641	70.752	61.655

Type	UC ID	T77	T78	T79	T80	T81	T82
Deep Learning	2	43.190	54.255	29.661	28.857	38.358	45.484
	5	14.136	14.398	27.283	11.407	36.784	60.796
	9	376.842	400.901	397.306	350.913	313.316	411.871
Machine Learning	1	43.607	36.181	45.602	46.540	61.418	60.865
	3	6.427	6.218	6.282	6.925	9.378	10.689
	4	75.345	43.867	53.368	41.346	55.554	44.168
	6	70.201	92.984	88.682	87.318	119.263	101.842
	7	21.054	21.709	18.791	15.683	14.659	22.839
	8	1,390.725	1,108.964	1,456.422	1,548.684	1,356.440	1,120.135
	10	81.125	65.091	89.420	82.124	112.365	61.108

Type	UC ID	T83	T84	T85	T86	T87	T88
Deep Learning	2	41.661	29.605	43.320	29.823	34.566	24.580
	5	19.081	24.196	34.066	16.629	13.549	10.364
	9	465.222	482.956	407.557	383.218	420.815	477.179
Machine Learning	1	39.490	46.856	49.792	43.056	47.108	37.306
	3	19.129	13.443	12.236	7.657	8.072	7.129
	4	42.930	55.982	50.373	47.816	43.724	41.206
	6	86.407	113.752	118.282	93.584	115.957	73.851
	7	19.659	10.968	23.692	26.656	13.661	12.748
	8	1,096.066	1,104.763	1,148.487	1,515.012	1,491.201	1,224.469
	10	66.727	113.122	82.447	77.220	69.209	69.957

Type	UC ID	T89	T90	T91	T92	T93	T94
Deep Learning	2	33.846	26.254	35.874	29.951	38.279	26.878
	5	14.424	39.596	33.824	21.077	15.622	98.127
	9	396.584	374.385	309.950	320.635	320.487	408.167
Machine Learning	1	49.529	58.641	60.524	54.328	46.403	55.204
	3	9.083	11.823	7.609	6.625	9.881	9.244
	4	66.952	51.006	45.395	43.840	51.079	43.783
	6	79.006	93.525	127.759	86.474	135.113	84.720
	7	19.806	25.010	21.994	14.966	17.361	20.292
	8	1,448.539	1,460.642	1,495.391	1,506.651	1,487.155	1,191.117
	10	80.977	59.730	111.341	61.398	66.050	67.618

Type	UC ID	T95	T96	T97	T98	T99	T100
Deep Learning	2	21.667	29.827	47.430	43.320	45.019	28.423
	5	13.456	26.508	21.858	14.850	53.316	21.665
	9	388.915	361.348	437.062	426.871	304.835	396.058
Machine Learning	1	55.307	44.903	44.786	51.630	59.524	51.043
	3	9.835	11.379	6.386	9.965	6.141	11.397
	4	41.423	82.098	53.921	74.418	61.195	55.107
	6	78.554	53.249	111.880	65.674	94.657	77.943
	7	26.552	22.956	27.730	17.442	25.329	17.450
	8	1,430.887	1,562.216	1,406.390	1,528.429	1,541.482	1,471.629
	10	111.961	102.995	104.427	66.959	65.264	59.586

Table 3-2 Use Case Elapsed Times

### 3.4 SUT Validation Test Output

<u>Validation Run Report</u>			
AIUCpm@1	405.96	T <sub>Load</sub>	0.39
Scale Factor	1	T <sub>LD</sub>	0.39
Streams	100	T <sub>PTT</sub>	20.82
Kit Version	2.0.0	T <sub>PST1</sub>	2.99
Execution Status	Pass	T <sub>PST2</sub>	2.99
Accuracy Status	Pass	T <sub>PST</sub>	2.99
		T <sub>TT</sub>	0.20
Test Times			
Overall Run Start Time		2024-09-24 18:20:44.705	
Overall Run End Time		2024-09-24 18:48:13.489	
Overall Run Elapsed Time		1,648.784	
Load Test Start Time		2024-09-24 18:22:10.154	
Load Test End Time		2024-09-24 18:22:10.558	
Load Test Elapsed Time		0.404	
Power Training Start Time		2024-09-24 18:22:10.559	
Power Training End Time		2024-09-24 18:39:46.130	
Power Training Elapsed Time		1,055.571	
Power Serving 1 Start Time		2024-09-24 18:39:46.133	
Power Serving 1 End Time		2024-09-24 18:41:05.488	
Power Serving 1 Elapsed Time		79.355	
Power Serving 2 Start Time		2024-09-24 18:41:05.489	
Power Serving 2 End Time		2024-09-24 18:42:24.590	
Power Serving 2 Elapsed Time		79.101	
Scoring Start Time		2024-09-24 18:43:07.246	
Scoring End Time		2024-09-24 18:44:50.343	
Scoring Elapsed Time		103.097	
Throughput Start Time		2024-09-24 18:44:50.357	
Throughput End Time		2024-09-24 18:48:13.485	
Throughput Elapsed Time		203.128	
(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	0.295	<=	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	0.013	<=	0.50	Pass
6	matthews_corrcoef	0.464	>=	0.19	Pass
7	median_absolute_error	0.891	<=	1.80	Pass
8	accuracy_score	0.715	>=	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			30
Data Storage Ratio			32.00

### 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	1,536	1,536
Scale Factor		30
Total Memory (GiB)		1,536
SF / Memory Ratio		0.02

### 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	1,115.54	AIUCpm@30
TPCx-AI Price/Performance Metric	126.34	\$/AIUCpm@30
TPCx-AI Scale Factor	30	
TPCx-AI Stream Count	100	

### Test Times


Overall Run Start Time	2024-09-24 18:48:34.282
Overall Run End Time	2024-09-24 22:58:26.260
Overall Run Elapsed Time	14,991.978
Load Test Start Time	2024-09-24 18:52:31.579
Load Test End Time	2024-09-24 18:52:36.092
Load Test Elapsed Time	4.513
Power Training Start Time	2024-09-24 18:52:36.093
Power Training End Time	2024-09-24 21:46:48.399
Power Training Elapsed Time	10,452.306
Power Serving 1 Start Time	2024-09-24 21:46:48.401
Power Serving 1 End Time	2024-09-24 22:01:49.804
Power Serving 1 Elapsed Time	901.403
Power Serving 2 Start Time	2024-09-24 22:01:49.806
Power Serving 2 End Time	2024-09-24 22:17:04.035
Power Serving 2 Elapsed Time	914.229
Scoring Start Time	2024-09-24 22:17:45.031
Scoring End Time	2024-09-24 22:19:22.651
Scoring Elapsed Time	97.620
Throughput Start Time	2024-09-24 22:19:22.664
Throughput End Time	2024-09-24 22:58:26.255
Throughput Elapsed Time	2,343.591

Accuracy Metrics

Use Case	Metric Name	Metric	Criteria	Threshold	Status
1	N/A	0.000	N/A	0.00	Pass
2	word_error_rate	0.445	<=	0.50	Pass
3	mean_squared_log_error	3.553	<=	5.40	Pass
4	f1_score	0.706	>=	0.65	Pass
5	mean_squared_log_error	0.036	<=	0.50	Pass
6	matthews_corrcoef	0.624	>=	0.19	Pass
7	median_absolute_error	1.002	<=	1.80	Pass
8	accuracy_score	0.754	>=	0.65	Pass
9	accuracy_score	1.000	>=	0.90	Pass
10	accuracy_score	0.817	>=	0.70	Pass

## 5.2 Throughput Test Stream Times

The following chart shows the minimum, 1<sup>st</sup> quartile, median, mean (X), 3<sup>rd</sup> 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](http://www.tpc.org).

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



Nicholas Wakou  
Dell Inc.  
701 E. Parmer Ln. Bld. 2  
Austin, TX 78753

October 1, 2024

I verified the TPC Express Benchmark™ AI v2.0.0 performance of the following configuration:

Platform: 1x Dell PowerEdge 6725  
Operating System: Red Hat Enterprise Linux 8.19 (Ootpa)  
Additional Software: Anaconda Business

The results were:

**Performance Metric 1,115.54 AIUCpm@100**

Secondary Metrics	T <sub>LD</sub>	4.50
	T <sub>PTT</sub>	310.45
	T <sub>PST</sub>	20.75
	T <sub>TT</sub>	2.34

**System Under Test 1x Dell PowerEdge 6725 with:**

CPU	2x AMD EPYC 9355 32-Core Processor		
Memory	1,536 GiB		
Storage	<b>Qty</b>	<b>Size</b>	<b>Type</b>
	2	480 GB	M.2 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 v2.0.0.
- 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 100 GB.
- The generated dataset used for testing was protected by RAID-1.

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

- 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](http://www.sizing.com)

# Third-Party Price Quotes

## Anaconda



### Anaconda Business: Support

**To:** Jesse Rangel, AMD and TPC Consortium | [jesse.rangel@amd.com](mailto:jesse.rangel@amd.com)  
**From:** Nanette George, Staff Product Marketing Manager, Anaconda | [ngeorge@anaconda.com](mailto:ngeorge@anaconda.com)  
**Re:** Anaconda Support Quote for Publishing TPCx-AI Performance Benchmark Testing

**Effective Date:** September 17, 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, 85 users*, based on this configuration**	\$51,000
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>

\* Premium Support requires a minimum annual contract of \$50,000 for Anaconda Business.  
 \*\* 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.

# 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 Dell	
Sponsor/ModelOptimization/...	Details of model optimization.
Sponsor/ModifiedKitFiles/...	1 modified file(s).
Sponsor/Tuning/...	All tuning files used.