

TPC Express Benchmark™ AI Full Disclosure Report

PowerEdge R6715

with 1x PowerEdge R6715
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@3, 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. 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 R6715. 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 R6715 (Server)	Red Hat Enterprise Linux 8.10 (Ootpa)


Metrics Overview


Total System Cost	Performance	Price/Performance	Availability Date
\$101,811 USD	720.11 AIUCpm@3	141.39 USD \$/AIUCpm@3	November 9, 2024


Executive Summary

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

		<h1>PowerEdge R6715</h1>		TPCx-AI 2.0.0 TPC Pricing 2.9.0 Report Date Oct. 10, 2024																																																							
TPCx-AI Performance 720.11 AIUCpm@3	Total System Cost \$101,811 USD	Price/Performance \$141.39 USD/AIUCpm@3	Availability Date November 9, 2024																																																								
Framework Anaconda Business	Operating System Red Hat Enterprise Linux 8.10 (Ootpa)	Other Software N/A	Scale Factor 3	Streams 100																																																							
<h3>Use Case Time (sec.) by Phase</h3>  <table border="1"> <caption>Approximate Use Case Time (sec.) by Phase</caption> <thead> <tr> <th>Phase</th> <th>Training (sec)</th> <th>Serving 1 (sec)</th> <th>Serving 2 (sec)</th> <th>Throughput (Avg) (sec)</th> </tr> </thead> <tbody> <tr><td>1</td><td>10</td><td>5</td><td>5</td><td>5</td></tr> <tr><td>2</td><td>380</td><td>5</td><td>5</td><td>5</td></tr> <tr><td>3</td><td>10</td><td>5</td><td>5</td><td>5</td></tr> <tr><td>4</td><td>10</td><td>5</td><td>5</td><td>5</td></tr> <tr><td>5</td><td>10</td><td>5</td><td>5</td><td>5</td></tr> <tr><td>6</td><td>5</td><td>5</td><td>5</td><td>5</td></tr> <tr><td>7</td><td>5</td><td>5</td><td>5</td><td>5</td></tr> <tr><td>8</td><td>900</td><td>50</td><td>50</td><td>100</td></tr> <tr><td>9</td><td>10</td><td>10</td><td>10</td><td>100</td></tr> <tr><td>10</td><td>10</td><td>5</td><td>5</td><td>5</td></tr> </tbody> </table>					Phase	Training (sec)	Serving 1 (sec)	Serving 2 (sec)	Throughput (Avg) (sec)	1	10	5	5	5	2	380	5	5	5	3	10	5	5	5	4	10	5	5	5	5	10	5	5	5	6	5	5	5	5	7	5	5	5	5	8	900	50	50	100	9	10	10	10	100	10	10	5	5	5
Phase	Training (sec)	Serving 1 (sec)	Serving 2 (sec)	Throughput (Avg) (sec)																																																							
1	10	5	5	5																																																							
2	380	5	5	5																																																							
3	10	5	5	5																																																							
4	10	5	5	5																																																							
5	10	5	5	5																																																							
6	5	5	5	5																																																							
7	5	5	5	5																																																							
8	900	50	50	100																																																							
9	10	10	10	100																																																							
10	10	5	5	5																																																							
Physical Storage / Scale Factor 320.00	Scale Factor / Physical Memory 0.02	Main Data Redundancy Model RAID-1																																																									
Servers: Total Processors/Cores/Threads	1 1 / 32 / 64																																																										
Server Type Processors Memory Storage Controller Storage Device Network Controller	1x PowerEdge R6715 (Server) 1x AMD EPYC 9355P 32-Core Processor 192 GiB 1x BOSS-N1 2x 480 GB M.2 NVMe 1x Broadcom 57504 Quad Port 10/25GbE																																																										

		<h1 style="text-align: center;">PowerEdge R6715</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
Hardware						
PowerEdge R6715 Server	210-BNRK		1 \$40,463.00	1	\$40,463.00	
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-BIFS		1 0	1		
AMD EPYC 9355P 3.40GHz, 32C/64T, 256 Cache (280W) DDR5-6000	338-CRCD		1 0	1		
High Performance Heatsink	412-BBGB		1 0	1		
Performance Optimized	370-AHLL		1 0	1		
6400MT/s RDIMMs	370-BBRX		1 0	1		
16GB RDIMM, 6400MT/s, Single Rank	370-BCGH		1 0	12		
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		
UEFI BIOS Boot Mode with GPTPartition	800-BBDM		1 0	1		
No Energy Star	387-BBEY		1 0	1		
4 Very High Performance Fan for 1 CPU	384-BDHS		1 0	1		
Dual, Redundant (1+1),Hot-Plug MHS Power Supply, 1500W MM, Titanium	450-BCXC		1 0	2		
C13 to C14, PDU Style, 12 AMP, 6.5 Feet (2m) Power Cord, North America	492-BBDI		1 0	1		
Riser Config 5, 2 x16 LP Slots (Gen5), 2nd OCP	330-BCWQ		1 0	1		
PowerEdge R6715 Motherboard	329-BKPB		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 for x8 /x10 Chassis, R6715	325-BETT		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		
Secured Component Verification	528-COYT		1 0	1		
iDRAC10, Enterprise 17G	634-CSHV		1 0	1		
No Quick Sync	350-BBXM		1 0	1		
iDRAC,Legacy Password	379-BCSQ		1 0	1		
iDRAC Service Module (ISM), NOT Installed	379-BCQX		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 R6715 Shipping	340-DCYK		1 0	1		
Shipping Material for No backplane chassis	340-DCYG		1 0	1		
PowerEdge R6715 No CCC or CEMarking	470-AFOI		1 0	1		
iDRAC Group Manager, Disabled	379-BCQY		1 0	1		
None Required	817-BBBP		1 0	1		
Dell Hardware Limited Warranty Plus On-Site Service	886-9770		1 \$249.00	1		\$249.00
3Yr ProSupport and 4hr Mission Critical - 3Years	199-BONO		1 \$2,931.26	1		\$2,931.26
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 24 Monitor – S2425H	210-BMGX		1 \$99.99	1	\$99.99	
Keyboard/Mouse (Included with Server)	N/A		1 0	1		
				Subtotal	\$40,562.99	\$3,180.26
(continued on the next page)						

	<h1 style="text-align: center;">PowerEdge R6715</h1>		TPCx-AI	2.0.0		
			TPC Pricing	2.9.0		
			Report Date	Oct. 10, 2024		
(continued from the previous page)						
Description	Part Number	Source	List Price	Qty	Extended Price	1-Yr. Maintenance
Software						
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	
				Subtotal	\$52,428.90	\$35,000.00
				Total	\$92,991.89	\$39,225.70
Large Purchase Discount (65%)*					-\$27,294.73	-\$2,746.71
Pricing: 1 = Dell; 2 = Anaconda * 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): \$101,811 AIUCpm@3: 720.11 \$/AIUCpm@3: \$141.39			
<i>Prices used in TPC benchmarks reflect the actual prices a customer would pay for a one-time purchase of the stated Line Items. Individually negotiated discounts are not permitted. Special prices based on assumptions about past or future purchases are not permitted. All discounts reflect standard pricing policies for the listed Line Items. For complete details, see the pricing section of the TPC Benchmark Standard. If you find that the stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.</i>						

	<h1>PowerEdge R6715</h1>		TPCx-AI	2.0.0
			TPC Pricing	2.9.0
			Report Date	Oct. 10, 2024
<u>Numerical Quantities</u>				
AIUCpm@3	720.11	T_{Load}	0.69	
Scale Factor	3	T_{LD}	0.69	
Streams	100	T_{PTT}	32.22	
Kit Version	2.0.0	T_{PST1}	4.25	
Execution Status	Pass	T_{PST2}	4.25	
Accuracy Status	Pass	T_{PST}	4.25	
		T_{TT}	0.42	
Test Times				
Overall Run Start Time	2024-09-24 19:05:41.560			
Overall Run End Time	2024-09-24 19:44:57.847			
Overall Run Elapsed Time	2,356.287			
Load Test Start Time	2024-09-24 19:06:56.439			
Load Test End Time	2024-09-24 19:06:57.142			
Load Test Elapsed Time	0.703			
Power Training Start Time	2024-09-24 19:06:57.143			
Power Training End Time	2024-09-24 19:31:46.064			
Power Training Elapsed Time	1,488.921			
Power Serving 1 Start Time	2024-09-24 19:31:46.066			
Power Serving 1 End Time	2024-09-24 19:33:50.140			
Power Serving 1 Elapsed Time	124.074			
Power Serving 2 Start Time	2024-09-24 19:33:50.141			
Power Serving 2 End Time	2024-09-24 19:35:53.271			
Power Serving 2 Elapsed Time	123.130			
Scoring Start Time	2024-09-24 19:36:25.897			
Scoring End Time	2024-09-24 19:37:59.290			
Scoring Elapsed Time	93.393			
Throughput Start Time	2024-09-24 19:37:59.303			
Throughput End Time	2024-09-24 19:44:57.843			
Throughput Elapsed Time	418.540			



PowerEdge R6715

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	24.168	2.379	2.376	7.381	0.000
UC02	377.655	5.929	5.944	22.477	0.482
UC03	29.743	2.798	2.815	10.142	3.507
UC04	27.267	4.333	4.347	15.309	0.700
UC05	43.637	2.513	2.501	12.571	0.435
UC06	1.063	0.516	0.516	3.258	0.493
UC07	3.066	1.256	1.260	4.316	0.953
UC08	910.996	76.263	75.305	219.047	0.751
UC09	39.185	24.276	24.251	88.003	1.000
UC10	32.064	3.724	3.732	12.684	0.816

Use Case Serving Times (sec.)




Table of Contents

- Abstract..... 3
- Executive Summary 3
- Table of Contents..... 9
- Clause 0 – Preamble11
 - 0.1 TPC Express Benchmark™ AI Overview11
- Clause 1 – General Items13
 - 1.1 Test Sponsor13
 - 1.2 Parameter Settings13
 - 1.3 Configuration Diagrams13
 - 1.3.1 Measured Configuration.....14
 - 1.3.2 Differences Between the Measured and the Priced Configurations.....14
- Clause 2 – SW Components & Data Distribution15
 - 2.1 Roles and Dataset Distribution15
 - 2.2 File System Implementation.....15
 - 2.3 Execution Engine, Frameworks, Driver & Libraries15
 - 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 Output23
 - 3.5 Configuration Parameters.....24
- Clause 4 – SUT Related Items.....25
 - 4.1 Specialized Hardware/Software25
 - 4.2 Configuration Files.....25
 - 4.3 SUT Environment Information.....25
 - 4.4 Data Storage to Scale Factor Ratio25
 - 4.5 Scale Factor to Memory Ratio25
 - 4.6 Output of Tests25
 - 4.7 Additional Sponsor Files25
 - 4.8 Model Optimizations25
- Clause 5 – Metrics and Scale Factor26
 - 5.1 Reported Performance Metrics26

5.2 Throughput Test Stream Times27

Auditor's Information28

Third-Party Price Quotes.....31

 Anaconda31

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 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:	1/32/64	Storage Devices:	2
Total Memory:	192 GiB	Storage Capacity:	960 GB

The diagram illustrates the hardware configuration of a single server node. At the center is CPU 0. It is connected to 12 DIMM modules, labeled DIMM 0 through DIMM 11, arranged in two vertical columns of six. Below the CPU, there is a storage configuration consisting of 2 x 480GB M.2 NVMe SSDs in RAID1 configuration, managed by a DELL BOSS N1 controller, containing the OS and Kit. A Network Interface Card (NIC) is also connected to the CPU.

	<u>Server</u>
Server	1x PowerEdge R6715
Procs/Cores/Threads:	1/32/64
Processor Model:	1x AMD EPYC 9355P 32-Core Processor
Memory:	192 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 R6715	idrac-c0n2408-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	5.929	5.944	19.272	20.343	22.656	19.665
	5	2.513	2.501	14.636	16.363	10.241	12.023
	9	24.276	24.251	88.498	95.621	87.926	80.017
Machine Learning	1	2.379	2.376	5.696	5.369	5.038	7.279
	3	2.798	2.815	7.346	10.827	4.255	15.384
	4	4.333	4.347	13.232	15.691	19.422	12.982
	6	0.516	0.516	3.586	3.196	4.740	2.874
	7	1.257	1.260	5.911	4.535	3.474	4.688
	8	76.263	75.305	231.734	219.606	235.408	220.281
	10	3.724	3.732	11.707	13.601	12.655	13.685

Type	UC ID	T5	T6	T7	T8	T9	T10
Deep Learning	2	17.090	14.300	29.686	24.192	22.294	24.730
	5	10.924	15.370	17.026	4.612	11.309	11.381
	9	89.948	97.042	92.931	101.103	86.530	84.424
Machine Learning	1	10.387	10.666	7.607	7.796	10.687	6.937
	3	11.630	11.404	8.225	9.536	7.679	11.632
	4	15.473	12.872	11.626	13.439	18.184	14.752
	6	3.683	1.092	3.549	3.808	3.821	4.667
	7	3.422	2.935	3.711	4.236	2.688	3.790
	8	227.711	218.400	216.208	215.803	202.544	196.169
	10	12.532	14.905	10.474	14.147	17.881	11.719

Type	UC ID	T11	T12	T13	T14	T15	T16
Deep Learning	2	25.861	29.404	23.287	27.352	7.558	21.849
	5	17.311	16.312	11.877	12.641	13.759	13.667
	9	85.166	83.399	85.283	84.852	97.785	82.091
Machine Learning	1	8.890	7.689	6.028	8.587	5.422	8.380
	3	11.714	8.111	10.370	14.757	17.270	10.443
	4	17.455	13.290	23.091	17.106	18.972	14.109
	6	4.390	3.164	2.998	2.453	0.717	3.246
	7	4.401	2.941	2.761	5.251	3.570	3.461
	8	218.087	217.796	215.963	221.417	228.420	221.795
	10	11.727	11.612	9.644	8.963	11.589	11.735

Type	UC ID	T17	T18	T19	T20	T21	T22
Deep Learning	2	26.504	21.010	21.864	20.937	20.258	18.273
	5	19.422	14.195	12.148	7.929	11.481	11.651
	9	102.504	96.881	97.926	96.102	92.135	82.698
Machine Learning	1	5.893	10.304	11.307	7.389	3.193	5.699
	3	10.357	9.463	7.954	9.165	4.913	8.959
	4	6.351	13.862	15.769	15.834	13.650	14.876
	6	1.462	3.600	3.462	2.773	2.991	2.836
	7	2.454	2.901	6.612	4.352	4.404	3.574
	8	214.372	211.447	216.022	213.687	239.559	204.851
	10	13.822	20.733	9.247	13.454	16.217	10.107

Type	UC ID	T23	T24	T25	T26	T27	T28
Deep Learning	2	6.464	24.949	22.929	20.138	25.253	25.035
	5	14.940	13.998	8.682	7.486	14.924	12.512
	9	92.181	92.352	91.024	96.341	101.398	89.863
Machine Learning	1	3.220	5.986	8.795	7.891	8.351	9.371
	3	15.130	12.872	8.599	8.654	13.751	8.059
	4	21.604	14.685	19.266	15.329	11.361	17.052
	6	2.977	8.828	2.797	3.249	5.351	1.973
	7	5.288	3.581	5.143	4.472	3.497	4.518
	8	234.717	230.143	198.562	217.399	207.505	216.685
	10	10.715	4.392	25.321	10.413	12.855	17.198

Type	UC ID	T29	T30	T31	T32	T33	T34
Deep Learning	2	23.286	20.434	23.058	20.670	27.744	17.725
	5	13.088	9.091	14.039	9.081	11.976	13.604
	9	87.884	88.540	86.281	96.870	77.811	84.083
Machine Learning	1	5.079	5.769	7.753	7.164	7.365	6.499
	3	4.009	12.275	10.804	9.343	10.366	8.535
	4	18.701	14.354	15.048	13.348	13.232	13.842
	6	1.500	3.249	3.508	2.692	2.741	2.589
	7	6.662	5.512	3.289	2.860	4.179	3.427
	8	223.156	226.426	207.988	210.561	211.662	234.084
	10	16.089	14.254	12.000	11.652	9.893	16.910

Type	UC ID	T35	T36	T37	T38	T39	T40
Deep Learning	2	23.040	27.115	20.905	26.777	28.877	29.933
	5	9.917	14.983	9.655	14.106	12.402	12.191
	9	85.003	85.825	84.882	87.896	86.506	90.623
Machine Learning	1	5.638	9.804	10.567	13.042	5.640	5.941
	3	7.753	9.532	9.829	7.720	12.714	9.703
	4	17.170	15.127	17.660	11.124	13.848	15.100
	6	4.006	2.774	2.271	1.433	2.468	2.061
	7	3.042	3.308	5.249	5.881	3.180	3.804
	8	232.499	218.726	225.572	228.481	216.713	206.985
	10	18.366	12.990	12.935	8.229	13.684	10.864

Type	UC ID	T41	T42	T43	T44	T45	T46
Deep Learning	2	21.136	24.242	24.014	21.259	17.966	20.129
	5	16.548	11.700	12.848	14.910	13.609	10.980
	9	92.484	84.542	87.121	86.087	92.617	97.384
Machine Learning	1	4.655	11.034	9.738	11.784	7.699	6.243
	3	12.016	9.704	13.284	11.716	9.673	5.859
	4	19.834	16.122	13.178	13.314	16.284	16.021
	6	1.276	3.683	3.827	2.382	1.883	3.828
	7	4.604	3.486	3.256	3.383	5.681	6.641
	8	222.282	201.526	202.788	205.751	232.612	219.345
	10	9.302	14.329	17.055	14.972	7.780	13.152

Type	UC ID	T47	T48	T49	T50	T51	T52
Deep Learning	2	28.091	18.291	29.787	24.775	18.299	22.603
	5	10.483	15.970	13.799	14.009	10.037	15.337
	9	87.926	89.672	92.343	90.961	87.706	83.329
Machine Learning	1	8.544	11.232	6.645	8.164	6.625	5.693
	3	11.648	11.226	9.024	6.974	11.952	9.257
	4	17.321	13.432	13.984	18.470	16.675	16.435
	6	2.917	2.895	3.343	2.840	3.370	4.319
	7	3.246	3.096	4.670	3.026	6.099	4.524
	8	209.474	209.464	203.558	197.461	221.959	186.681
	10	14.350	17.432	9.124	14.765	14.510	12.447

Type	UC ID	T53	T54	T55	T56	T57	T58
Deep Learning	2	26.141	22.485	8.517	23.997	19.470	32.530
	5	11.093	11.963	12.195	15.166	13.560	11.341
	9	81.380	80.719	101.224	86.324	79.640	81.582
Machine Learning	1	3.086	9.439	7.143	4.978	6.844	9.956
	3	8.685	8.051	12.709	7.072	13.880	9.728
	4	7.463	14.409	13.896	18.010	13.794	15.011
	6	2.255	5.479	2.371	3.175	3.543	3.117
	7	1.287	6.556	3.856	2.824	3.821	6.221
	8	246.643	216.399	236.317	233.030	220.209	234.188
	10	14.256	9.980	11.308	5.395	14.583	4.055

Type	UC ID	T59	T60	T61	T62	T63	T64
Deep Learning	2	25.572	16.642	19.527	21.337	21.689	29.227
	5	7.952	13.932	10.625	6.262	7.657	14.472
	9	84.306	86.466	88.588	93.402	92.757	82.105
Machine Learning	1	8.333	9.597	7.156	6.934	8.347	9.766
	3	13.621	10.624	8.627	15.609	12.638	7.379
	4	16.423	21.543	13.548	14.497	8.440	15.946
	6	4.488	2.631	2.154	3.146	2.868	3.029
	7	4.022	3.187	6.318	6.009	7.681	3.406
	8	222.621	216.418	207.900	203.602	226.700	213.574
	10	8.382	13.102	10.919	14.984	9.799	13.227

Type	UC ID	T65	T66	T67	T68	T69	T70
Deep Learning	2	28.315	23.294	21.092	22.923	19.812	27.195
	5	13.534	15.495	10.292	9.460	17.014	13.729
	9	54.839	83.720	94.575	92.148	96.735	87.242
Machine Learning	1	8.408	5.744	4.204	5.957	5.803	5.335
	3	13.223	6.798	8.863	14.133	11.185	6.136
	4	15.409	22.277	16.994	13.839	13.100	18.815
	6	2.580	4.125	2.012	1.937	2.705	3.237
	7	5.838	3.607	8.691	3.417	5.327	3.643
	8	253.438	214.892	224.951	223.096	209.038	209.201
	10	14.059	11.181	13.265	10.914	14.160	13.129

Type	UC ID	T71	T72	T73	T74	T75	T76
Deep Learning	2	24.027	22.212	29.914	17.853	17.039	13.886
	5	15.139	13.291	13.384	12.117	16.478	9.414
	9	90.991	79.096	69.941	82.850	97.350	94.301
Machine Learning	1	7.751	5.668	7.603	5.958	10.397	7.462
	3	7.814	13.307	11.605	11.205	7.445	12.003
	4	16.790	16.683	16.503	12.435	16.139	20.137
	6	8.426	3.223	8.565	2.238	4.132	2.827
	7	4.241	6.002	5.707	5.580	3.674	1.728
	8	214.349	212.612	231.038	221.422	216.422	227.225
	10	15.056	12.294	11.999	17.248	12.645	11.351

Type	UC ID	T77	T78	T79	T80	T81	T82
Deep Learning	2	26.719	21.043	27.096	25.882	25.076	27.771
	5	14.196	10.530	10.725	7.264	14.094	9.608
	9	79.813	86.569	85.652	87.368	75.009	93.300
Machine Learning	1	7.528	6.064	6.020	6.416	7.370	6.564
	3	6.502	12.682	6.630	9.705	8.938	8.251
	4	17.945	18.033	11.684	18.627	16.776	18.684
	6	2.686	2.222	2.431	2.555	4.539	5.696
	7	3.060	3.571	6.153	3.029	6.822	3.882
	8	212.300	222.385	234.081	232.116	223.589	215.894
	10	11.359	12.571	8.789	9.664	8.220	12.655

Type	UC ID	T83	T84	T85	T86	T87	T88
Deep Learning	2	16.531	21.682	23.546	23.020	16.775	24.377
	5	11.580	8.944	12.251	13.543	12.542	14.793
	9	93.757	83.360	87.914	92.951	85.308	79.246
Machine Learning	1	6.946	3.120	3.873	5.740	7.255	6.004
	3	12.136	15.253	12.858	8.737	11.498	8.199
	4	15.160	18.585	17.211	6.267	13.339	14.168
	6	3.626	4.193	3.867	2.587	3.309	2.799
	7	6.901	4.104	1.582	4.267	4.094	3.024
	8	218.174	228.224	225.473	232.872	215.855	200.182
	10	13.974	20.684	10.851	19.460	12.470	12.505

Type	UC ID	T89	T90	T91	T92	T93	T94
Deep Learning	2	22.148	19.102	26.369	20.732	25.221	19.521
	5	7.953	13.914	16.502	13.760	17.279	11.603
	9	100.799	78.548	66.283	95.209	76.195	88.741
Machine Learning	1	7.149	11.902	6.730	9.932	7.961	4.862
	3	8.306	12.814	7.866	10.020	7.286	9.756
	4	17.034	13.323	13.251	15.271	17.591	6.680
	6	2.667	2.012	3.201	3.394	2.984	2.665
	7	3.684	7.404	3.494	3.948	4.652	6.210
	8	220.617	231.687	233.780	213.684	190.839	238.774
	10	10.944	10.379	21.999	11.252	9.084	14.456

Type	UC ID	T95	T96	T97	T98	T99	T100
Deep Learning	2	24.283	26.044	27.349	17.525	26.254	15.676
	5	9.356	13.792	10.657	10.821	16.556	16.087
	9	89.551	82.525	86.677	86.562	87.067	101.216
Machine Learning	1	5.639	8.254	9.097	8.281	9.725	5.623
	3	10.821	7.126	11.509	11.735	9.451	10.383
	4	13.895	12.635	13.602	20.187	13.481	14.404
	6	2.993	3.318	8.768	3.212	0.989	2.656
	7	3.610	3.703	4.042	4.091	3.342	5.655
	8	203.970	198.476	212.504	234.932	239.338	221.634
	10	11.023	17.939	14.998	11.020	4.993	9.655

Table 3-2 Use Case Elapsed Times

3.4 SUT Validation Test Output

<u>Validation Run Report</u>			
AIUCpm@1	355.78	T _{Load}	0.33
Scale Factor	1	T _{LD}	0.33
Streams	100	T _{PTT}	21.62
Kit Version	2.0.0	T _{PST1}	3.06
Execution Status	Pass	T _{PST2}	3.04
Accuracy Status	Pass	T _{PST}	3.06
		T _{TT}	0.37
Test Times			
Overall Run Start Time	2024-09-24 18:27:14.834		
Overall Run End Time	2024-09-24 19:05:19.330		
Overall Run Elapsed Time	2,284.496		
Load Test Start Time	2024-09-24 18:28:18.407		
Load Test End Time	2024-09-24 18:28:18.748		
Load Test Elapsed Time	0.341		
Power Training Start Time	2024-09-24 18:28:18.749		
Power Training End Time	2024-09-24 18:52:32.382		
Power Training Elapsed Time	1,453.633		
Power Serving 1 Start Time	2024-09-24 18:52:32.383		
Power Serving 1 End Time	2024-09-24 18:54:20.425		
Power Serving 1 Elapsed Time	108.042		
Power Serving 2 Start Time	2024-09-24 18:54:20.426		
Power Serving 2 End Time	2024-09-24 18:56:08.367		
Power Serving 2 Elapsed Time	107.941		
Scoring Start Time	2024-09-24 18:56:39.958		
Scoring End Time	2024-09-24 18:59:03.478		
Scoring Elapsed Time	143.520		
Throughput Start Time	2024-09-24 18:59:03.491		
Throughput End Time	2024-09-24 19:05:19.327		
Throughput Elapsed Time	375.836		
(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	0.318	<=	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	0.011	<=	0.50	Pass
6	matthews_corrcoef	0.465	>=	0.19	Pass
7	median_absolute_error	0.893	<=	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	3
Data Storage Ratio	320.00

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	192	192

Scale Factor	3
Total Memory (GiB)	192
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	720.11	AIUCpm@3
TPCx-AI Price/Performance Metric	141.39	\$/AIUCpm@3
TPCx-AI Scale Factor	3	
TPCx-AI Stream Count	100	

Test Times


Overall Run Start Time	2024-09-24 19:05:41.560
Overall Run End Time	2024-09-24 19:44:57.847
Overall Run Elapsed Time	2,356.287
Load Test Start Time	2024-09-24 19:06:56.439
Load Test End Time	2024-09-24 19:06:57.142
Load Test Elapsed Time	0.703
Power Training Start Time	2024-09-24 19:06:57.143
Power Training End Time	2024-09-24 19:31:46.064
Power Training Elapsed Time	1,488.921
Power Serving 1 Start Time	2024-09-24 19:31:46.066
Power Serving 1 End Time	2024-09-24 19:33:50.140
Power Serving 1 Elapsed Time	124.074
Power Serving 2 Start Time	2024-09-24 19:33:50.141
Power Serving 2 End Time	2024-09-24 19:35:53.271
Power Serving 2 Elapsed Time	123.130
Scoring Start Time	2024-09-24 19:36:25.897
Scoring End Time	2024-09-24 19:37:59.290
Scoring Elapsed Time	93.393
Throughput Start Time	2024-09-24 19:37:59.303
Throughput End Time	2024-09-24 19:44:57.843
Throughput Elapsed Time	418.540

Accuracy Metrics

Use Case	Metric Name	Metric	Criteria	Threshold	Status
1	N/A	0.000	N/A	0.00	Pass
2	word_error_rate	0.482	<=	0.50	Pass
3	mean_squared_log_error	3.507	<=	5.40	Pass
4	f1_score	0.700	>=	0.65	Pass
5	mean_squared_log_error	0.435	<=	0.50	Pass
6	matthews_corrcoef	0.493	>=	0.19	Pass
7	median_absolute_error	0.953	<=	1.80	Pass
8	accuracy_score	0.751	>=	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.



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 R6715
 Operating System: Red Hat Enterprise Linux 8.19 (Ootpa)
 Additional Software: Anaconda Business

The results were:

Performance Metric 720.11 AIUCpm@3

Secondary Metrics	T _{LD}	0.69
	T _{PTT}	32.22
	T _{PST}	4.25
	T _{TT}	0.42

System Under Test 1x Dell PowerEdge R6715 with:

CPU	1x AMD EPYC 9355P 32-Core Processor		
Memory	192 GiB		
Storage	Qty	Size	Type
	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 3 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 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: Jesse Rangel, AMD and TPC Consortium | jesse.rangel@amd.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: 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/pkg/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.