

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@10, 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


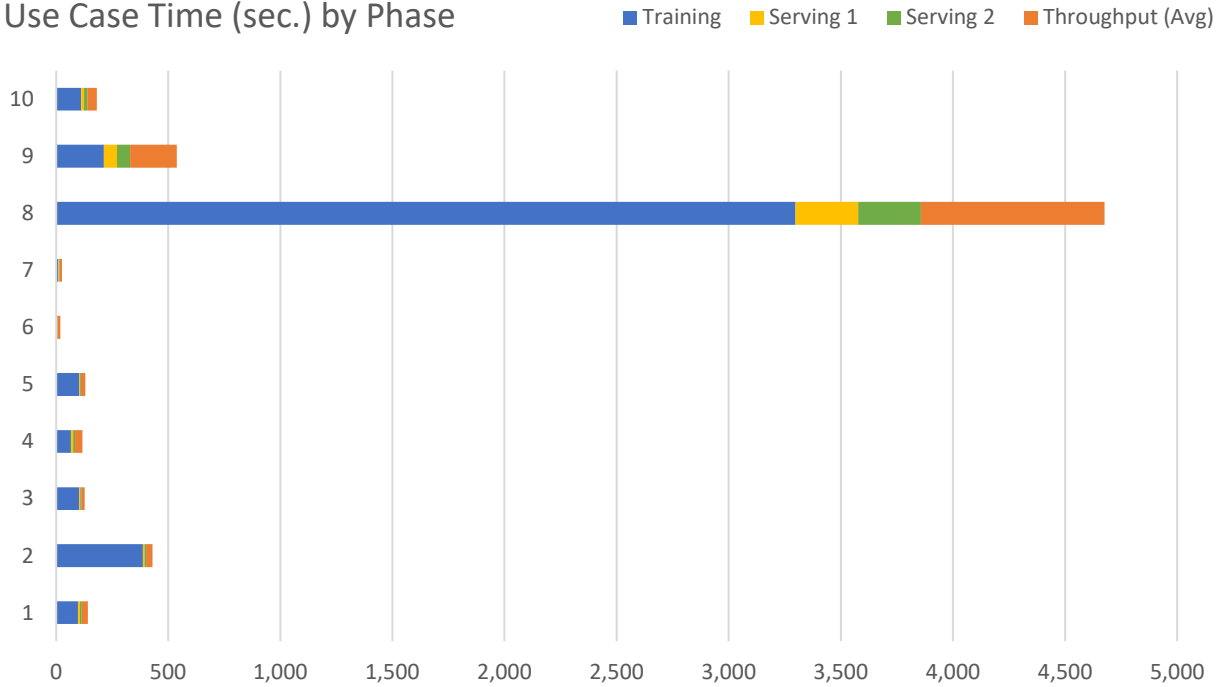
Test Sponsor	Node(s)	Operating System
Dell	1x PowerEdge R6715 (Server)	Red Hat Enterprise Linux 8.10 (Ootpa)


Metrics Overview

Total System Cost	Performance	Price/Performance	Availability Date
\$118,344 USD	864.53 AIUCpm@10	136.89 USD \$/AIUCpm@10	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 864.53 AIUCpm@10	Total System Cost \$118,344 USD	Price/Performance \$136.89 USD/AIUCpm@10	Availability Date November 9, 2024	
Framework Anaconda Business	Operating System Red Hat Enterprise Linux 8.10 (Ootpa)	Other Software N/A	Scale Factor 10	Streams 100
<div style="display: flex; justify-content: space-between;"> <div> <h3>Use Case Time (sec.) by Phase</h3>  </div> <div> <p>■ Training ■ Serving 1 ■ Serving 2 ■ Throughput (Avg)</p> </div> </div>				
Physical Storage / Scale Factor 96.00		Scale Factor / Physical Memory 0.01		Main Data Redundancy Model RAID-1
Servers: Total Processors/Cores/Threads		1 1 / 32 / 64		
Server Type		1x PowerEdge R6715 (Server)		
Processors		1x AMD EPYC 9355P 32-Core Processor		
Memory		768 GiB		
Storage Controller		1x BOSS-N1		
Storage Device		2x 480 GB NVMe		
Network Controller		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 \$84,527.00	1	\$84,527.00	
High Performance Computing Cluster Information SKU	463-7922		1 \$0.00	1		
No Backplane	379-BDSY		1 \$0.00	1		
No Backplane	379-BDSV		1 \$0.00	1		
No Rear Storage	379-BDTE		1 \$0.00	1		
Trusted Platform Module 2.0V3	461-AAIG		1 \$0.00	1		
No HD, No Backplane	321-BIFS		1 \$0.00	1		
AMD EPYC 9355P 3.40GHz, 32C/64T, 256 Cache (280W) DDR5-6000	338-CRCD		1 \$0.00	1		
High Performance Heatsink	412-BBGB		1 \$0.00	1		
Performance Optimized	370-AHLL		1 \$0.00	1		
6400MT/s RDIMMs	370-BBRX		1 \$0.00	1		
64GB RDIMM, 6400MT/s, Dual Rank	370-BCCY		1 \$0.00	12		
Diskless Configuration (No RAID, No Controller)	780-BCDH		1 \$0.00	1		
No Controller	405-AACD		1 \$0.00	1		
No Hard Drive	400-ABHL		1 \$0.00	1		
Performance BIOS Settings	384-BBBL		1 \$0.00	1		
UEFI BIOS Boot Mode with GPTPartition	800-BBDM		1 \$0.00	1		
No Energy Star	387-BBEY		1 \$0.00	1		
4 Very High Performance Fan for 1 CPU	384-BDHS		1 \$0.00	1		
Dual, Redundant (1+1), Hot-Plug MHS Power Supply, 1500W MM, Titanium	450-BCXC		1 \$0.00	1		
C13 to C14, PDU Style, 12 AMP, 6.5 Feet (2m) Power Cord, North America	492-BBDI		1 \$0.00	2		
Riser Config 5, 2 x16 LP Slots (Gen5), 2nd OCP	330-BCWQ		1 \$0.00	1		
PowerEdge R6715 Motherboard	329-BKPB		1 \$0.00	1		
Broadcom 57504 Quad Port 10/25GbE, SFP28, OCP NIC 3.0	540-BCRX		1 \$0.00	1		
Broadcom 5720 Dual Port 1GbE LOM	540-BDKD		1 \$0.00	1		
No Bezel for x8 /x10 Chassis, R6715	325-BETT		1 \$0.00	1		
BOSS-N1 controller card + with 2 M.2 480GB (RAID 1) (22x80)	403-BDMM		1 \$0.00	1		
No Operating System, No Utility Partition, BOSS	611-BBBX		1 \$0.00	1		
No Media Required	605-BBFN		1 \$0.00	1		
Secured Component Verification	528-COYT		1 \$0.00	1		
iDRAC10, Enterprise 17G	634-CSHV		1 \$0.00	1		
No Quick Sync	350-BBXN		1 \$0.00	1		
iDRAC, Legacy Password	379-BCSG		1 \$0.00	1		
iDRAC Service Module (ISM), NOT Installed	379-BCQX		1 \$0.00	1		
ReadyRails Sliding Rails Without Cable Management Arm or Strain Relief Bar	770-BECD		1 \$0.00	1		
No Systems Documentation, No OpenManage DVD Kit	631-AAACK		1 \$0.00	1		
PowerEdge R6715 Shipping	340-DCYK		1 \$0.00	1		
Shipping Material for No backplane chassis	340-DCYG		1 \$0.00	1		
PowerEdge R6715 No CCC or CEMarking	470-AFOI		1 \$0.00	1		
iDRAC Group Manager, Disabled	379-BCQY		1 \$0.00	1		
None Required	817-BBBD		1 \$0.00	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 \$6,103.87	1		\$6,103.87
Thank you choosing Dell ProSupport. For tech support, visit //www.dell.com/support or call 1-800- 945-3355	989-3439		1 \$0.00	1		
Dell 24 Monitor – S2425H	210-BMGX		1 \$99.99	1	\$99.99	
Keyboard/Mouse (Included with Server)	N/A		1 0	1		
				Subtotal	\$84,626.99	\$6,352.87
(continued on the next page)						


	<h1>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 Med 528-BHPJ		1	1428.9	1	\$1,428.90	
				Subtotal	\$52,428.90	\$35,000.00
Large Purchase Discount (65%)*				Total	\$137,055.89	\$41,352.87

Pricing: 1 = Dell; 2 = Anaconda	Total System Cost (USD):	\$118,344
* Discount applies to all line items where Key = 1. Discount based upon total system cost as purchased by a regular customer.	AIUCpm@10:	864.53
Audited by Doug Johnson, InfoSizing	\$/AIUCpm@10:	\$136.89

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.

	<h1 style="text-align: center;">PowerEdge R6715</h1>		TPCx-AI	2.0.0
			TPC Pricing	2.9.0
			Report Date	Oct. 10, 2024
<u>Numerical Quantities</u>				
AIUCpm@10	864.53	T_{Load}		1.95
Scale Factor	10	T_{LD}		1.95
Streams	100	T_{PTT}		101.01
		T_{PST1}		9.35
Kit Version	2.0.0	T_{PST2}		9.37
Execution Status	Pass	T_{PST}		9.37
Accuracy Status	Pass	T_{TT}		1.26
Test Times				
Overall Run Start Time				2024-09-18 17:45:11.991
Overall Run End Time				2024-09-18 19:36:09.605
Overall Run Elapsed Time				6,657.614
Load Test Start Time				2024-09-18 17:46:46.551
Load Test End Time				2024-09-18 17:46:48.518
Load Test Elapsed Time				1.967
Power Training Start Time				2024-09-18 17:46:48.519
Power Training End Time				2024-09-18 19:00:03.956
Power Training Elapsed Time				4,395.437
Power Serving 1 Start Time				2024-09-18 19:00:03.957
Power Serving 1 End Time				2024-09-18 19:06:31.536
Power Serving 1 Elapsed Time				387.579
Power Serving 2 Start Time				2024-09-18 19:06:31.538
Power Serving 2 End Time				2024-09-18 19:12:59.561
Power Serving 2 Elapsed Time				388.023
Scoring Start Time				2024-09-18 19:13:32.070
Scoring End Time				2024-09-18 19:15:07.621
Scoring Elapsed Time				95.551
Throughput Start Time				2024-09-18 19:15:07.634
Throughput End Time				2024-09-18 19:36:09.601
Throughput Elapsed Time				1,261.967

	<h2>PowerEdge R6715</h2>	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	97.948	8.397	8.432	26.751	0.000
UC02	386.901	7.247	7.267	28.836	0.483
UC03	103.161	4.201	4.194	15.099	3.609
UC04	66.647	8.941	9.028	33.185	0.707
UC05	103.267	3.364	3.400	19.961	0.082
UC06	6.147	0.987	0.984	10.946	0.548
UC07	8.501	3.139	3.136	10.611	1.026
UC08	3,297.876	279.549	279.826	818.372	0.735
UC09	212.812	58.644	58.668	207.675	1.000
UC10	112.092	13.023	12.999	43.659	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 Overview.....11
- Clause 1 – General Items13
 - 1.1 Test Sponsor13
 - 1.2 Parameter Settings.....13
 - 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:	768 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 DIMMs, labeled DIMM 0 through DIMM 11, arranged in two vertical columns of six. Below the CPU, there are two 480GB M.2 NVMe SSDs configured in RAID1, managed by a Dell BOSS N1 controller. A Network Interface Card (NIC) is also connected to the system.

	<u>Server</u>
Server	1x PowerEdge R6715
Procs/Cores/Threads:	1/32/64
Processor Model:	1x AMD EPYC 9355P 32-Core Processor
Memory:	768 GiB
Storage Controller:	1x BOSS-N1
Storage Devices:	2x 480 GB 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-c0n2409-os	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 8.10 (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	7.247	7.267	18.139	26.801	29.062	31.100
	5	3.364	3.400	21.246	18.639	20.013	14.918
	9	58.644	58.668	218.231	202.754	213.450	188.638
Machine Learning	1	8.397	8.432	33.645	31.042	27.078	28.272
	3	4.201	4.194	27.298	13.020	7.867	12.703
	4	8.941	9.028	29.868	38.009	33.326	43.399
	6	0.987	0.984	9.215	9.151	11.179	11.973
	7	3.139	3.136	8.904	10.134	11.052	10.220
	8	279.549	279.826	815.976	793.331	829.813	775.379
	10	13.023	12.999	42.184	48.269	47.518	45.874

Type	UC ID	T5	T6	T7	T8	T9	T10
Deep Learning	2	21.352	28.459	37.277	33.552	34.674	39.246
	5	20.470	18.634	26.343	18.996	24.391	25.083
	9	225.051	241.941	209.408	223.844	197.840	194.947
Machine Learning	1	34.756	27.510	31.342	29.789	24.117	22.909
	3	17.094	18.151	11.499	12.568	11.731	22.555
	4	24.862	33.723	21.686	26.599	27.339	23.676
	6	7.968	10.758	11.403	8.417	12.466	9.184
	7	13.394	7.308	10.646	11.252	11.429	12.200
	8	821.558	773.770	830.310	780.786	839.347	780.717
	10	38.264	39.552	45.372	47.182	40.445	48.453

Type	UC ID	T11	T12	T13	T14	T15	T16
Deep Learning	2	27.564	36.544	28.709	28.169	28.039	14.713
	5	21.501	14.084	15.267	26.292	21.102	22.711
	9	199.278	216.399	215.705	196.307	222.620	199.635
Machine Learning	1	28.495	29.922	22.667	29.380	21.905	29.957
	3	11.607	10.131	15.997	16.529	21.749	12.683
	4	35.870	37.300	33.917	32.064	25.872	29.515
	6	11.906	14.558	12.031	9.373	10.247	10.800
	7	7.314	7.048	13.446	11.687	10.273	14.980
	8	839.196	821.737	812.701	815.814	783.956	852.417
	10	43.160	38.654	38.148	46.149	61.311	49.929

Type	UC ID	T17	T18	T19	T20	T21	T22
Deep Learning	2	33.820	30.519	28.073	29.853	29.666	26.279
	5	26.832	22.082	25.583	12.531	18.086	13.211
	9	227.554	202.809	207.716	224.997	221.618	204.990
Machine Learning	1	31.552	28.413	35.573	27.503	16.247	22.530
	3	13.312	12.371	14.406	19.285	11.578	13.835
	4	33.393	39.742	34.052	31.197	33.154	31.724
	6	15.231	8.136	12.601	9.144	12.102	14.373
	7	9.281	8.408	18.944	8.957	11.239	13.167
	8	762.213	758.066	813.220	820.336	839.238	791.303
	10	54.020	44.067	35.296	39.122	43.387	43.991

Type	UC ID	T23	T24	T25	T26	T27	T28
Deep Learning	2	22.452	28.654	25.114	32.493	29.554	34.445
	5	26.291	18.893	18.415	18.789	26.533	20.245
	9	220.095	221.733	221.855	221.295	204.471	209.498
Machine Learning	1	24.455	23.640	27.204	27.898	31.399	31.991
	3	17.087	12.893	15.764	12.579	10.754	22.453
	4	42.588	34.196	29.342	34.898	33.515	38.064
	6	13.197	12.314	11.369	11.782	13.479	6.183
	7	13.328	16.089	16.228	9.528	13.664	11.151
	8	796.208	833.494	819.629	755.706	790.882	803.890
	10	41.816	31.905	59.388	47.930	47.224	43.485

Type	UC ID	T29	T30	T31	T32	T33	T34
Deep Learning	2	29.001	28.879	11.365	27.533	29.297	27.763
	5	20.393	15.837	8.225	19.603	20.588	18.111
	9	196.288	226.673	211.383	210.127	199.755	210.049
Machine Learning	1	23.634	23.188	22.946	22.416	28.900	24.600
	3	9.546	13.866	12.036	18.220	16.446	14.085
	4	38.838	28.160	38.392	34.317	36.487	29.131
	6	8.250	12.244	7.911	9.253	10.793	10.613
	7	16.714	7.445	9.159	14.740	12.238	9.733
	8	846.864	832.786	863.839	771.150	823.101	796.528
	10	40.694	35.515	50.465	44.799	42.011	51.127

Type	UC ID	T35	T36	T37	T38	T39	T40
Deep Learning	2	31.991	30.536	30.240	32.770	31.293	36.319
	5	14.627	16.931	25.598	21.372	21.835	23.402
	9	202.301	197.710	201.810	192.416	205.470	198.289
Machine Learning	1	20.057	31.991	20.326	21.102	24.400	30.467
	3	13.003	11.221	11.236	15.099	10.731	13.829
	4	43.057	32.637	39.835	36.542	33.604	46.135
	6	17.751	9.731	10.053	14.822	12.533	9.025
	7	8.545	18.879	12.201	8.282	9.397	12.785
	8	784.056	821.215	760.227	780.236	857.216	813.551
	10	42.000	48.921	41.609	45.615	38.545	38.435

Type	UC ID	T41	T42	T43	T44	T45	T46
Deep Learning	2	22.192	28.157	35.417	28.479	30.500	27.092
	5	22.719	17.346	8.313	16.857	24.676	13.342
	9	228.598	214.429	189.273	194.019	223.047	195.648
Machine Learning	1	16.140	27.512	20.373	29.501	21.637	47.411
	3	19.540	14.714	19.226	11.938	11.899	14.482
	4	27.961	33.561	35.275	40.357	39.504	25.614
	6	3.753	9.024	10.362	8.042	10.229	13.690
	7	11.357	12.998	8.706	9.711	7.549	8.713
	8	840.152	782.111	860.697	822.234	841.674	814.609
	10	49.858	52.140	45.198	53.739	27.873	42.963

Type	UC ID	T47	T48	T49	T50	T51	T52
Deep Learning	2	29.973	21.311	32.354	36.464	35.647	32.091
	5	15.700	23.929	18.075	25.973	19.427	23.879
	9	207.061	208.387	208.497	188.888	198.835	214.388
Machine Learning	1	29.715	31.110	27.064	30.755	22.338	17.685
	3	19.458	21.332	15.449	10.717	19.229	13.671
	4	34.184	33.325	40.296	29.112	32.213	27.848
	6	10.150	8.935	10.945	13.083	10.917	15.064
	7	8.995	9.030	8.222	11.274	6.852	9.291
	8	822.402	821.422	793.080	809.989	797.694	830.447
	10	53.448	40.688	57.134	39.970	45.156	42.623

Type	UC ID	T53	T54	T55	T56	T57	T58
Deep Learning	2	28.465	34.573	13.775	33.433	26.263	32.639
	5	22.796	22.650	24.809	22.974	20.867	15.844
	9	204.013	205.540	225.454	225.074	196.018	198.160
Machine Learning	1	20.028	21.412	20.712	32.002	25.067	40.510
	3	14.508	20.044	22.928	15.114	16.356	16.054
	4	32.268	26.205	39.025	30.052	34.126	28.510
	6	18.556	12.354	9.863	10.874	10.546	11.243
	7	6.319	8.732	12.921	9.116	9.786	7.974
	8	837.629	815.288	829.043	810.286	788.638	864.994
	10	44.959	41.030	37.214	37.048	53.877	26.138

Type	UC ID	T59	T60	T61	T62	T63	T64
Deep Learning	2	26.677	21.369	12.771	8.103	31.667	31.899
	5	14.963	19.521	19.300	14.632	24.139	21.159
	9	180.481	197.721	216.352	231.916	210.892	220.177
Machine Learning	1	27.298	22.364	20.797	30.251	37.135	19.360
	3	13.714	17.323	13.932	18.300	12.161	14.264
	4	37.399	41.122	41.121	17.271	27.509	24.657
	6	11.156	7.545	14.149	10.099	8.836	10.480
	7	8.475	13.145	9.113	8.827	12.388	11.860
	8	864.020	834.588	835.442	860.097	820.915	772.628
	10	53.760	52.557	47.627	47.120	39.544	41.381

Type	UC ID	T65	T66	T67	T68	T69	T70
Deep Learning	2	30.674	30.723	33.018	38.684	30.789	28.724
	5	17.066	21.532	14.886	25.272	30.780	24.264
	9	174.924	195.077	207.159	199.292	246.744	206.695
Machine Learning	1	29.869	31.650	21.581	29.785	20.853	27.789
	3	13.178	15.790	16.663	18.935	13.619	27.991
	4	43.128	37.679	28.273	29.210	33.266	22.018
	6	9.309	11.565	8.736	9.325	10.293	8.383
	7	15.914	11.554	7.855	7.034	14.129	7.852
	8	845.550	784.676	845.545	800.963	791.712	825.522
	10	53.483	39.017	41.732	46.953	40.163	50.399

Type	UC ID	T71	T72	T73	T74	T75	T76
Deep Learning	2	31.516	23.348	32.928	34.394	19.983	8.854
	5	12.958	21.693	24.223	17.649	19.724	17.582
	9	196.119	205.517	207.986	203.543	213.789	230.057
Machine Learning	1	22.255	32.406	36.928	24.462	30.237	36.177
	3	5.616	17.110	11.942	15.345	19.047	8.578
	4	29.772	30.880	29.681	29.721	40.877	44.719
	6	9.341	14.323	11.750	10.435	8.731	11.168
	7	10.033	9.810	10.956	12.842	14.706	3.140
	8	876.913	801.923	818.617	795.331	806.913	848.475
	10	47.139	55.823	36.530	36.246	42.311	39.208

Type	UC ID	T77	T78	T79	T80	T81	T82
Deep Learning	2	35.752	37.036	32.850	34.381	27.981	34.433
	5	16.355	16.188	20.566	18.893	23.116	21.384
	9	193.790	213.950	217.056	229.254	197.903	207.428
Machine Learning	1	26.814	27.689	19.574	25.068	27.031	9.975
	3	10.470	18.177	12.753	10.495	11.333	7.716
	4	35.219	40.973	28.926	38.676	34.925	25.591
	6	10.524	9.936	3.640	12.013	10.678	11.700
	7	11.019	9.769	11.102	9.262	8.257	13.957
	8	798.220	790.747	865.394	761.039	787.830	874.026
	10	41.996	34.690	30.586	43.172	43.885	41.445

Type	UC ID	T83	T84	T85	T86	T87	T88
Deep Learning	2	32.863	29.164	25.041	26.523	26.934	35.232
	5	17.579	19.361	28.818	15.942	16.503	16.601
	9	203.937	218.176	196.102	199.031	210.773	193.973
Machine Learning	1	20.854	32.953	23.292	29.528	26.730	26.341
	3	16.571	16.977	21.530	22.338	10.192	19.286
	4	33.675	33.388	30.657	17.420	37.085	39.529
	6	14.425	11.204	10.221	14.894	13.137	11.101
	7	10.671	7.976	6.229	10.016	12.484	11.533
	8	825.627	782.872	844.586	856.443	840.622	832.658
	10	44.227	56.133	42.113	39.881	36.707	42.939

Type	UC ID	T89	T90	T91	T92	T93	T94
Deep Learning	2	25.446	29.323	31.440	24.881	30.966	24.741
	5	12.463	22.161	18.975	28.501	27.895	18.059
	9	208.947	202.366	202.132	218.557	205.228	206.885
Machine Learning	1	31.843	24.846	27.816	32.689	30.185	22.644
	3	13.093	17.073	14.499	16.535	12.877	16.474
	4	34.101	29.447	38.476	33.076	37.597	23.136
	6	13.636	12.660	10.918	11.318	11.911	13.408
	7	9.555	7.462	11.160	11.594	11.443	7.594
	8	839.697	829.577	804.648	803.223	816.067	845.723
	10	36.635	43.020	45.073	38.244	39.527	50.145

Type	UC ID	T95	T96	T97	T98	T99	T100
Deep Learning	2	27.907	32.865	31.747	34.513	33.181	12.106
	5	14.142	21.531	25.734	15.902	21.023	13.197
	9	197.304	195.032	136.418	219.615	209.740	207.232
Machine Learning	1	33.668	34.760	8.556	25.102	27.589	22.153
	3	15.537	15.017	4.240	20.604	25.605	11.477
	4	29.856	25.912	45.027	34.807	33.933	29.259
	6	11.090	5.971	17.402	7.955	9.475	8.599
	7	9.307	9.367	10.682	13.030	7.745	11.282
	8	803.404	847.596	929.330	825.386	835.229	883.277
	10	48.888	40.801	41.815	41.453	29.070	39.584

Table 3-2 Use Case Elapsed Times

3.4 SUT Validation Test Output

<u>Validation Run Report</u>			
AIUCpm@1	357.71	T _{Load}	0.35
Scale Factor	1	T _{LD}	0.35
Streams	100	T _{PTT}	20.51
Kit Version	2.0.0	T _{PST1}	2.98
Execution Status	Pass	T _{PST2}	2.99
Accuracy Status	Pass	T _{PST}	2.99
		T _{TT}	0.37
Test Times			
Overall Run Start Time	2024-09-18 17:13:05.253		
Overall Run End Time	2024-09-18 17:44:51.382		
Overall Run Elapsed Time	1,906.129		
Load Test Start Time	2024-09-18 17:14:09.054		
Load Test End Time	2024-09-18 17:14:09.422		
Load Test Elapsed Time	0.368		
Power Training Start Time	2024-09-18 17:14:09.424		
Power Training End Time	2024-09-18 17:33:17.958		
Power Training Elapsed Time	1,148.534		
Power Serving 1 Start Time	2024-09-18 17:33:17.959		
Power Serving 1 End Time	2024-09-18 17:34:45.666		
Power Serving 1 Elapsed Time	87.707		
Power Serving 2 Start Time	2024-09-18 17:34:45.667		
Power Serving 2 End Time	2024-09-18 17:36:13.911		
Power Serving 2 Elapsed Time	88.244		
Scoring Start Time	2024-09-18 17:36:45.480		
Scoring End Time	2024-09-18 17:38:40.726		
Scoring Elapsed Time	115.246		
Throughput Start Time	2024-09-18 17:38:40.739		
Throughput End Time	2024-09-18 17:44:51.378		
Throughput Elapsed Time	370.639		
(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	0.196	<=	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	0.464	>=	0.19	Pass
7	median_absolute_error	0.894	<=	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			10
Data Storage Ratio			96.00

4.5 Scale Factor to Memory Ratio

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

Nodes	Memory (GiB)	Total (GiB)
1	768	768
Scale Factor		10
Total Memory (GiB)		768
SF / Memory Ratio		0.01

4.6 Output of Tests

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

4.7 Additional Sponsor Files

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

4.8 Model Optimizations

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

Clause 5 – Metrics and Scale Factor

5.1 Reported Performance Metrics

Metric Overview

TPCx-AI Performance Metric	864.53	AIUCpm@10
TPCx-AI Price/Performance Metric	136.89	\$/AIUCpm@10
TPCx-AI Scale Factor	10	
TPCx-AI Stream Count	100	

Test Times


Overall Run Start Time	2024-09-18 17:45:11.991
Overall Run End Time	2024-09-18 19:36:09.605
Overall Run Elapsed Time	6,657.614
Load Test Start Time	2024-09-18 17:46:46.551
Load Test End Time	2024-09-18 17:46:48.518
Load Test Elapsed Time	1.967
Power Training Start Time	2024-09-18 17:46:48.519
Power Training End Time	2024-09-18 19:00:03.956
Power Training Elapsed Time	4,395.437
Power Serving 1 Start Time	2024-09-18 19:00:03.957
Power Serving 1 End Time	2024-09-18 19:06:31.536
Power Serving 1 Elapsed Time	387.579
Power Serving 2 Start Time	2024-09-18 19:06:31.538
Power Serving 2 End Time	2024-09-18 19:12:59.561
Power Serving 2 Elapsed Time	388.023
Scoring Start Time	2024-09-18 19:13:32.070
Scoring End Time	2024-09-18 19:15:07.621
Scoring Elapsed Time	95.551
Throughput Start Time	2024-09-18 19:15:07.634
Throughput End Time	2024-09-18 19:36:09.601
Throughput Elapsed Time	1,261.967

Accuracy Metrics

Use Case	Metric Name	Metric	Criteria	Threshold	Status
1	N/A	0.000	N/A	0.00	Pass
2	word_error_rate	0.483	<=	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	0.082	<=	0.50	Pass
6	matthews_corrcoef	0.548	>=	0.19	Pass
7	median_absolute_error	1.026	<=	1.80	Pass
8	accuracy_score	0.735	>=	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

September 30, 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 864.53 AIUCpm@10

Secondary Metrics	T _{LD}	1.95
	T _{PTT}	101.01
	T _{PST}	9.37
	T _{TT}	1.26

System Under Test 1x Dell PowerEdge R6715 with:

CPU	1x AMD EPYC 9355P 32-Core Processor		
Memory	768 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 10 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 includes 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/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.