

# TPC Express Benchmark™ AI Full Disclosure Report

## PowerEdge C6615

with 1x PowerEdge C6615  
using

### Anaconda Pro

running on

## Red Hat Enterprise Linux 8.6 (Ootpa)

TPCx-AI Version  
Report Edition  
Report Submitted

1.0.3.1  
First  
September 18, 2023

**First Edition - September 2023**

*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](http://www.dell.com). Actual performance and environmental costs of Dell products will vary depending on individual customer configurations and conditions.*

**Copyright© 2023 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 C6615. The software used included Anaconda Pro. This report provides full disclosure of the results. All testing was conducted in conformance with the requirements of the TPCx-AI Standard Specification, Revision 1.0.3.1.

## Configuration Overview

Test Sponsor	Node(s)	Operating System
Dell	1x PowerEdge C6615 (Server)	Red Hat Enterprise Linux 8.6 (Ootpa)


## Metrics Overview


Total System Cost	Performance	Price/Performance	Availability Date
\$38,688 USD	506.30 AIUCpm@10	76.42 USD \$/AIUCpm@10	September 18, 2023

# Executive Summary

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

		<h1>PowerEdge C6615</h1>		TPCx-AI 1.0.3.1 TPC Pricing 2.8.0 Report Date Sep. 18, 2023
TPCx-AI Performance <b>506.30 AIUCpm@10</b>	Total System Cost <b>\$38,688 USD</b>	Price/Performance <b>\$76.42 USD/AIUCpm@10</b>	Availability Date <b>September 18, 2023</b>	
Framework Anaconda Pro	Operating System Red Hat Enterprise Linux 8.6 (Ootpa)	Other Software N/A	Scale Factor 10	Streams 100
<h3>Use Case Time (sec.) by Phase</h3> 				
Physical Storage / Scale Factor <b>192.00</b>	Scale Factor / Physical Memory <b>0.03</b>	Main Data Redundancy Model <b>RAID 1</b>		
Servers: Total Processors/Cores/Threads	1 1 / 32 / 64			
Server Type Processors Memory Storage Controller Storage Device Network Controller	1x PowerEdge C6615 (Server) 1x AMD EPYC 8324P 32-Core Processor 384 GiB 1x Dell BOSS-N1 2x 960 GB M.2 NVMe SSD 1x Broadcom NetXtreme BCM5720 Single Port 1 GbE			

	<h1 style="text-align: center;">PowerEdge C6615</h1>			TPCx-AI	1.0.3.1	
				TPC Pricing	2.8.0	
				Report Date	Sep. 18, 2023	
Description	Part Number	Source	List Price	Qty	Extended Price	1-Yr. Maintenance
<b>Hardware</b>						
Dell PowerEdge C6615 Server	210-BHYX	1	\$77,577.68	1	\$77,577.68	
2.5 Chassis	379-BDTF	1	\$0.00	1	\$0.00	
No Backplane	379-BDSY	1	\$0.00	1	\$0.00	
No Backplane	379-BDSV	1	\$0.00	1	\$0.00	
Trusted Platform Module 2.0 V3	461-AAIG	1	\$0.00	1	\$0.00	
PowerEdge C6600 - C6XXX Shipping	340-DCEG	1	\$0.00	1	\$0.00	
PowerEdge C6615, Diskless with No Backplane	379-BFGD	1	\$0.00	1	\$0.00	
AMD EPYC 8324P 2.35GHz, 32C/64T, 64M Cache (185W) DDR5-4800	338-CMJF	1	\$0.00	1	\$0.00	
PowerEdge C6615 CPU Heatsink	412-BBFV	1	\$0.00	1	\$0.00	
Performance Optimized	370-AHLL	1	\$0.00	1	\$0.00	
4800MT/s RDIMMs	370-AHCL	1	\$0.00	1	\$0.00	
64GB RDIMM, 4800MT/s Dual Rank	370-AGZR	1	\$0.00	6	\$0.00	
BOSS-N1 controller card + with 2 M.2 960GB (RAID 1)	403-BCSB	1	\$0.00	1	\$0.00	
Diskless Configuration (No RAID, No Controller)	780-BCDH	1	\$0.00	1	\$0.00	
No Controller	405-AACD	1	\$0.00	1	\$0.00	
No Hard Drive	565-BBBC	1	\$0.00	1	\$0.00	
Performance BIOS Settings	384-BBBL	1	\$0.00	1	\$0.00	
UEFI BIOS Boot Mode with GPT Partition	800-BBDM	1	\$0.00	1	\$0.00	
Dual, Hot-plug, Redundant Power Supply (1+1), 1800W	379-BFCW	1	\$0.00	1	\$0.00	
Redundant PSU Mode	384-BCIT	1	\$0.00	1	\$0.00	
PowerEdge C6615 Riser Config 2, Riser 1A+2C	406-BBTV	1	\$0.00	1	\$0.00	
PowerEdge C6615, Motherboard, WMX	321-BKCD	1	\$0.00	1	\$0.00	
RHEL, 1-2SKT, Physical Node, 3YR Premium Sub, 1 Virtual Guest, Digitally Fulfilled	528-CHFH	1	\$0.00	1	\$0.00	
No Media Required	605-BBFN	1	\$0.00	1	\$0.00	
iDRAC9, Enterprise 16G	528-CTIC	1	\$0.00	1	\$0.00	
iDRAC Service Module (ISM), NOT Installed	379-BCQX	1	\$0.00	1	\$0.00	
iDRAC Group Manager, Disabled	379-BCQY	1	\$0.00	1	\$0.00	
C6xxx/C6600 Complex Tied Order Ship	340-CZIS	1	\$0.00	1	\$0.00	
PowerEdge C65xx No CE	340-CQZE	1	\$0.00	1	\$0.00	
Broadcom 5720 Single Port 1GbE BASE-T Adapter, OCP NIC 3.0	540-BDRK	1	\$0.00	1	\$0.00	
Basic Next Business Day 12 Months	709-BBFC	1	\$100.00	1		\$100.00
ProSupport Plus and 4Hr Mission Critical Initial, 12 Month(s)	865-BBKW	1	\$2,798.29	1		\$2,798.29
Dell 24 Monitor –S2421HN	210-AXHJ	1	\$158.49	1	\$158.49	
Dell Wireless Keyboard and Mouse - KM3322W	580-AKCW	1	\$29.99	1	\$29.99	
			<b>Subtotal</b>		\$77,766.16	\$2,898.29
<b>Software</b>						
Anaconda Pro 1 Subscription (1-Year)		2	\$10,000.00	1	\$10,000.00	
Red Hat Enterprise Linux Server, Premium 24/7 (Physical or Virtual Nodes)		1	\$1,299.00	1	\$1,299.00	
			<b>Subtotal</b>		\$11,299.00	\$0.00
			<b>Total</b>		\$89,065.16	\$2,898.29
Large Purchase Discount (65%)*					-\$51,392.35	-\$1,883.89
Pricing: 1 = Dell; 2 = Anaconda			<b>Total System Cost (USD):</b>		<b>\$38,688</b>	
* Discount applies to all line items where Key = 1. Discount based upon total system cost as purchased by a regular customer.			<b>AIUCpm@10:</b>		<b>506.30</b>	
<b>Audited by Doug Johnson, InfoSizing</b>			<b>\$/AIUCpm@10:</b>		<b>\$76.42</b>	
<p><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></p>						

	<h1>PowerEdge C6615</h1>		TPCx-AI	1.0.3.1
			TPC Pricing	2.8.0
			Report Date	Sep. 18, 2023
<u>Numerical Quantities</u>				
<b>AIUCpm@10</b>	<b>506.30</b>	$T_{Load}$	3.23	
Scale Factor	10	$T_{LD}$	3.23	
Streams	100	$T_{PTT}$	181.16	
Kit Version	1.0.3.1	$T_{PST1}$	15.98	
Execution Status	Pass	$T_{PST2}$	16.03	
Accuracy Status	Pass	$T_{PST}$	16.03	
		$T_{TT}$	2.10	
Test Times				
Overall Run Start Time	2023-09-07 08:25:43.099			
Overall Run End Time	2023-09-07 11:30:52.416			
Overall Run Elapsed Time	11,109.317			
Load Test Start Time	2023-09-07 08:28:22.567			
Load Test End Time	2023-09-07 08:28:25.818			
Load Test Elapsed Time	3.251			
Power Training Start Time	2023-09-07 08:28:25.819			
Power Training End Time	2023-09-07 10:31:13.971			
Power Training Elapsed Time	7,368.152			
Power Serving 1 Start Time	2023-09-07 10:31:13.972			
Power Serving 1 End Time	2023-09-07 10:41:45.351			
Power Serving 1 Elapsed Time	631.379			
Power Serving 2 Start Time	2023-09-07 10:41:45.352			
Power Serving 2 End Time	2023-09-07 10:52:16.551			
Power Serving 2 Elapsed Time	631.199			
Scoring Start Time	2023-09-07 10:53:08.853			
Scoring End Time	2023-09-07 10:55:46.128			
Scoring Elapsed Time	157.275			
Throughput Start Time	2023-09-07 10:55:46.150			
Throughput End Time	2023-09-07 11:30:52.413			
Throughput Elapsed Time	2,106.263			



# PowerEdge C6615


TPCx-AI 1.0.3.1  
 TPC Pricing 2.8.0  
 Report Date Sep. 18, 2023

## Numerical Quantities (continued)

### Use Case Times & Accuracy

Use Case	Training (sec)	Serving 1 (sec)	Serving 2 (sec)	Throughput (avg)	Accuracy
UC01	155.795	13.893	13.936	45.576	0.000
UC02	822.398	11.784	11.683	50.689	0.480
UC03	175.030	8.033	7.966	27.215	3.609
UC04	115.182	15.200	15.097	54.459	0.707
UC05	236.443	6.407	6.339	33.775	0.079
UC06	11.806	1.658	1.661	7.351	0.505
UC07	15.674	5.588	5.948	17.668	1.033
UC08	5,289.639	445.800	445.574	1,395.636	0.735
UC09	377.226	102.986	103.045	329.613	1.000
UC10	168.860	19.940	19.856	73.288	0.816

Use Case Serving Times (sec.)



# Table of Contents

- Abstract..... 3
- Executive Summary ..... 3
- Table of Contents..... 8
- Clause 0 – Preamble .....10
  - 0.1 TPC Express Benchmark™ AI Overview .....10
- Clause 1 – General Items .....12
  - 1.1 Test Sponsor .....12
  - 1.2 Parameter Settings .....12
  - 1.3 Configuration Diagrams .....12
    - 1.3.1 Measured Configuration.....13
    - 1.3.2 Differences Between the Measured and the Priced Configurations.....13
- Clause 2 – SW Components & Data Distribution .....14
  - 2.1 Roles and Dataset Distribution.....14
  - 2.2 File System Implementation.....14
  - 2.3 Execution Engine, Frameworks, Driver & Libraries .....14
  - 2.4 Applied Patches.....14
- Clause 3 – Workload Related Items.....15
  - 3.1 Hardware & Software Tuning.....15
  - 3.2 Kit Version & Modifications .....15
  - 3.3 Use Case Elapsed Times.....15
  - 3.4 SUT Validation Test Output .....22
  - 3.5 Configuration Parameters .....23
- Clause 4 – SUT Related Items.....24
  - 4.1 Specialized Hardware/Software .....24
  - 4.2 Configuration Files .....24
  - 4.3 SUT Environment Information.....24
  - 4.4 Data Storage to Scale Factor Ratio.....24
  - 4.5 Scale Factor to Memory Ratio.....24
  - 4.6 Output of Tests .....24
  - 4.7 Additional Sponsor Files .....24
  - 4.8 Model Optimizations .....24
- Clause 5 – Metrics and Scale Factor .....25
  - 5.1 Reported Performance Metrics .....25



5.2 Throughput Test Stream Times.....26

Auditor's Information .....27

Third-Party Price Quotes.....30

    Anaconda .....30

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..

## 1.2 Parameter Settings

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

## 1.3 Configuration Diagrams

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

### 1.3.1 Measured Configuration

Nodes:	1	
Processors/Cores/Threads:	1/32/64	Storage Devices: 2
Total Memory:	384 GiB	Storage Capacity: 1,920 GB

The diagram illustrates the hardware configuration of a single node. At the center is 'CPU 0'. It is connected via bidirectional arrows to a stack of six DIMMs: DIMM 0, DIMM 1, DIMM 2 on the left, and DIMM 3, DIMM 4, DIMM 5 on the right. Below the CPU, a single arrow points to a storage block labeled '2 x 960GB M.2 NVME SSD, RAID1, DELL BOSS N1 (OS + Kit)'. To the right of the CPU, another arrow points to a 'NIC' (Network Interface Card).

	<u>Server</u>
Server	1x PowerEdge C6615:
Procs/Cores/Threads:	1/32/64
Processor Model:	1x AMD EPYC 8324P 32-Core Processor
Memory:	384 GiB
Storage Controller:	1x Dell BOSS-N1
Storage Devices:	2x 960 GB M.2 NVMe SSD
Network Controller:	1x Broadcom NetXtreme BCM5720 Single Port 1 GbE

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 C6615	localhost	All	2x 960 GB M.2 NVMe SSD	OS, Data

*Table 2-1 Software Components and Dataset Distribution*

### 2.2 File System Implementation

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

### 2.3 Execution Engine, Frameworks, Driver & Libraries

Anaconda Pro consisted of the following components.

Component	Version
conda	23.7.2
python	3.9.18
setuptools	59.8.0
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	4.65.2
keras	2.11.0
tensorflow	2.11.0
joblib	1.2.0
opencv	4.5.3
pyyaml	6.0.1
matplotlib	3.7.1
jinja2	3.1.2

*Table 2-2 Software Components*

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

### 2.4 Applied Patches

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

## Clause 3 – Workload Related Items

### 3.1 Hardware & Software Tuning

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

### 3.2 Kit Version & Modifications

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

TPCx-AI Kit Version	1.0.3.1
<u>Modified File</u> 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	11.784	11.683	35.060	43.950	40.568	48.346
	5	6.407	6.339	45.581	40.665	23.553	25.879
	9	102.986	103.045	338.742	321.069	341.040	346.413
Machine Learning	1	13.893	13.936	39.485	64.651	38.026	50.168
	3	8.033	7.966	35.262	32.926	20.223	35.134
	4	15.200	15.097	39.619	61.348	51.862	75.311
	6	1.658	1.661	6.838	6.312	6.609	5.768
	7	5.588	5.948	16.095	17.311	14.179	17.473
	8	445.800	445.574	1,431.375	1,395.397	1,443.953	1,415.598
	10	19.940	19.856	72.677	54.909	66.794	57.001

Type	UC ID	T5	T6	T7	T8	T9	T10
Deep Learning	2	26.775	44.539	44.338	42.552	40.099	38.599
	5	27.896	41.090	33.660	29.034	52.303	40.858
	9	340.052	329.253	355.456	325.609	305.395	304.585
Machine Learning	1	49.570	51.694	68.565	32.402	46.357	47.969
	3	29.541	22.824	19.901	34.549	33.265	33.450
	4	57.409	64.692	29.531	38.433	45.425	44.781
	6	4.116	8.522	11.770	5.698	9.333	9.739
	7	18.517	15.827	21.183	16.288	16.502	17.407
	8	1,453.550	1,278.979	1,424.182	1,295.033	1,421.026	1,394.617
	10	66.653	71.225	68.625	73.327	86.999	60.093

Type	UC ID	T11	T12	T13	T14	T15	T16
Deep Learning	2	47.744	50.081	52.109	38.872	59.782	18.936
	5	35.464	27.198	34.514	32.074	36.115	33.959
	9	329.931	321.193	332.478	298.744	328.845	344.816
Machine Learning	1	75.101	37.633	33.740	51.159	42.536	32.869
	3	29.862	32.694	19.990	26.551	29.272	17.978
	4	53.495	67.110	52.476	45.665	65.826	55.400
	6	7.368	7.136	5.107	8.000	6.206	6.081
	7	19.049	19.755	22.546	20.596	18.191	22.242
	8	1,414.619	1,386.486	1,431.368	1,407.421	1,343.920	1,449.878
	10	54.789	76.157	59.946	107.789	100.552	97.408

Type	UC ID	T17	T18	T19	T20	T21	T22
Deep Learning	2	46.143	48.814	46.504	56.371	47.149	40.686
	5	35.946	34.573	43.486	15.276	31.607	22.697
	9	339.928	311.742	302.838	356.740	362.392	350.147
Machine Learning	1	77.487	41.280	59.651	32.754	53.978	40.638
	3	24.453	23.270	39.460	30.163	28.970	29.621
	4	53.066	69.299	48.870	53.807	55.761	51.884
	6	9.504	7.811	11.965	6.835	7.348	6.360
	7	21.801	15.217	16.583	15.200	15.469	12.176
	8	1,373.239	1,419.771	1,402.525	1,431.806	1,371.796	1,395.226
	10	68.060	75.554	57.697	54.653	57.658	59.469



Type	UC ID	T23	T24	T25	T26	T27	T28
Deep Learning	2	60.337	50.684	45.816	61.947	42.716	47.360
	5	36.215	26.943	33.779	29.000	42.374	33.360
	9	327.705	337.108	301.605	328.347	324.269	384.178
Machine Learning	1	45.471	49.304	46.315	37.707	61.814	37.856
	3	23.832	22.039	25.864	25.570	26.871	40.170
	4	55.837	44.160	42.536	36.438	57.340	47.357
	6	8.128	8.741	5.793	8.159	8.268	9.510
	7	22.371	12.768	20.868	20.992	14.815	15.926
	8	1,376.040	1,414.252	1,443.559	1,409.627	1,395.896	1,321.120
	10	74.790	75.817	91.377	75.601	82.893	81.532

Type	UC ID	T29	T30	T31	T32	T33	T34
Deep Learning	2	43.694	52.244	58.391	44.292	50.780	82.545
	5	39.485	23.474	29.670	35.699	30.132	22.645
	9	334.900	333.758	355.561	364.627	358.095	326.675
Machine Learning	1	40.027	37.883	33.367	30.336	51.158	44.535
	3	14.111	32.140	24.755	26.289	30.116	20.592
	4	78.120	58.526	58.332	78.007	70.626	49.207
	6	6.241	6.513	8.278	6.290	6.595	6.668
	7	19.855	22.682	21.640	12.718	19.924	17.575
	8	1,425.579	1,392.201	1,374.471	1,388.716	1,355.382	1,441.370
	10	63.225	67.120	68.789	61.076	67.385	65.111

Type	UC ID	T35	T36	T37	T38	T39	T40
Deep Learning	2	73.281	46.787	63.041	54.525	41.528	47.286
	5	33.344	25.209	42.209	35.234	41.667	37.726
	9	350.874	304.051	301.070	329.897	325.802	357.811
Machine Learning	1	32.094	49.586	58.539	37.005	41.571	76.625
	3	26.188	20.313	23.329	35.155	29.111	23.823
	4	57.912	54.047	70.043	69.658	46.805	57.678
	6	10.003	6.420	3.678	6.410	7.139	6.057
	7	18.864	21.351	16.716	16.187	12.709	15.316
	8	1,345.604	1,398.637	1,438.646	1,406.662	1,420.040	1,361.775
	10	65.224	101.042	56.170	59.946	62.250	60.890

Type	UC ID	T41	T42	T43	T44	T45	T46
Deep Learning	2	57.833	57.892	47.216	15.456	63.515	53.935
	5	33.614	45.445	11.749	41.746	39.306	41.332
	9	302.001	338.769	312.683	321.890	358.226	320.350
Machine Learning	1	53.043	33.536	31.796	47.570	24.250	64.179
	3	21.071	12.481	28.535	27.476	38.342	33.915
	4	50.969	35.453	85.500	53.270	48.082	70.653
	6	6.557	9.088	6.835	1.686	6.931	8.422
	7	14.677	18.051	12.703	15.553	19.383	14.302
	8	1,374.277	1,448.020	1,480.600	1,478.318	1,481.476	1,358.028
	10	110.121	75.160	57.019	79.658	21.239	71.168

Type	UC ID	T47	T48	T49	T50	T51	T52
Deep Learning	2	50.652	26.412	43.011	47.491	42.319	43.491
	5	22.753	39.325	31.332	40.682	29.635	46.884
	9	321.899	342.986	359.253	321.496	320.312	327.506
Machine Learning	1	71.523	51.488	44.091	49.296	51.535	37.675
	3	18.109	32.195	17.576	18.207	43.009	21.720
	4	57.587	51.559	49.559	51.041	46.207	41.680
	6	7.030	6.729	5.415	6.247	10.829	8.459
	7	16.227	18.313	14.950	16.417	25.087	18.199
	8	1,408.926	1,435.939	1,375.213	1,334.228	1,387.750	1,255.560
	10	71.590	71.456	87.056	65.403	71.487	60.664

Type	UC ID	T53	T54	T55	T56	T57	T58
Deep Learning	2	77.671	76.218	51.702	46.766	48.726	77.998
	5	23.460	27.478	58.364	34.045	33.235	24.312
	9	340.116	316.993	319.119	327.897	332.041	314.525
Machine Learning	1	25.229	54.467	48.423	53.064	41.687	50.082
	3	20.340	28.218	22.585	26.224	26.964	30.137
	4	48.975	50.886	54.475	41.108	75.587	54.755
	6	10.040	8.890	7.523	8.927	6.449	6.551
	7	6.256	13.772	14.892	22.680	20.924	18.049
	8	1,460.198	1,341.353	1,394.722	1,418.788	1,335.043	1,361.289
	10	71.642	75.901	70.105	70.523	75.595	79.260

Type	UC ID	T59	T60	T61	T62	T63	T64
Deep Learning	2	48.189	31.484	41.140	53.319	44.091	50.822
	5	21.986	33.676	25.729	31.803	35.407	35.981
	9	322.608	320.056	316.617	314.830	316.436	330.015
Machine Learning	1	44.858	36.728	55.464	29.393	36.091	38.183
	3	18.268	27.506	24.515	32.485	37.208	17.660
	4	53.745	55.498	69.238	54.592	44.903	50.457
	6	6.968	6.901	9.257	6.074	9.045	6.906
	7	18.529	17.596	18.521	25.625	20.311	16.484
	8	1,430.534	1,425.947	1,416.749	1,375.858	1,343.860	1,410.325
	10	90.959	106.454	77.082	74.388	53.835	75.419

Type	UC ID	T65	T66	T67	T68	T69	T70
Deep Learning	2	48.219	53.567	51.265	54.687	48.831	41.343
	5	50.050	33.120	21.022	64.097	34.443	35.803
	9	320.723	378.620	335.724	302.007	325.953	293.489
Machine Learning	1	51.567	77.135	29.851	43.127	55.326	35.428
	3	27.092	21.727	27.670	36.501	25.652	29.844
	4	47.515	53.049	51.693	55.962	36.594	57.881
	6	5.760	10.221	6.815	9.479	10.695	8.129
	7	18.182	20.091	24.456	17.634	14.566	17.232
	8	1,322.761	1,340.191	1,435.629	1,341.613	1,420.383	1,361.799
	10	68.193	55.975	89.488	76.995	70.125	92.561

Type	UC ID	T71	T72	T73	T74	T75	T76
Deep Learning	2	43.846	54.859	79.432	84.672	52.100	64.620
	5	21.383	54.784	35.988	35.222	31.909	26.541
	9	306.685	341.026	326.058	319.772	315.408	316.860
Machine Learning	1	46.109	64.968	44.099	33.495	45.613	34.234
	3	21.972	32.360	32.116	28.969	30.794	39.251
	4	60.798	56.068	46.145	43.555	43.150	81.646
	6	8.729	6.734	8.666	6.160	8.029	7.379
	7	18.502	17.353	12.346	19.839	15.764	13.912
	8	1,401.520	1,329.517	1,412.105	1,349.142	1,320.313	1,387.592
	10	114.692	63.542	61.676	61.286	64.210	63.523

Type	UC ID	T77	T78	T79	T80	T81	T82
Deep Learning	2	51.892	47.460	52.751	49.592	50.609	82.559
	5	15.837	28.217	24.126	11.944	33.299	31.547
	9	306.154	333.148	350.521	301.370	312.152	342.495
Machine Learning	1	37.884	36.348	36.935	50.609	36.359	27.116
	3	31.223	40.344	20.766	23.409	17.572	19.962
	4	70.003	55.341	51.779	60.729	59.627	59.486
	6	6.624	7.596	5.920	9.004	7.365	8.911
	7	16.755	16.254	17.683	16.644	13.062	18.855
	8	1,449.512	1,419.742	1,392.976	1,481.007	1,348.011	1,423.737
	10	73.359	75.227	84.767	64.383	106.941	60.595

Type	UC ID	T83	T84	T85	T86	T87	T88
Deep Learning	2	74.191	47.010	42.457	41.627	71.764	48.074
	5	27.119	34.440	59.787	23.026	24.612	36.674
	9	336.331	307.124	312.747	322.924	324.851	342.079
Machine Learning	1	30.488	48.542	35.485	60.813	46.834	41.906
	3	26.840	26.560	24.872	31.833	34.400	22.028
	4	50.317	67.135	43.122	43.685	56.936	55.780
	6	6.935	6.238	8.699	7.825	7.588	5.978
	7	14.419	14.570	9.218	32.124	18.123	16.285
	8	1,416.573	1,360.813	1,472.613	1,414.483	1,358.359	1,401.813
	10	85.521	80.767	58.930	82.345	79.148	56.987

Type	UC ID	T89	T90	T91	T92	T93	T94
Deep Learning	2	55.929	43.629	46.225	48.367	45.046	46.278
	5	11.973	48.341	37.783	43.076	53.193	33.766
	9	376.930	328.133	341.063	338.825	314.468	359.616
Machine Learning	1	52.381	44.412	45.873	43.170	40.200	39.106
	3	17.870	28.408	31.457	20.701	35.140	40.033
	4	84.999	61.020	56.991	44.348	44.092	23.531
	6	9.394	6.460	7.530	7.607	5.948	5.959
	7	17.122	24.026	20.310	22.033	18.061	13.117
	8	1,372.804	1,370.198	1,370.376	1,327.938	1,357.710	1,443.795
	10	74.436	88.876	116.340	77.611	70.154	73.694

Type	UC ID	T95	T96	T97	T98	T99	T100
Deep Learning	2	57.896	47.364	76.680	74.748	47.737	18.950
	5	20.071	34.241	34.085	28.791	70.981	15.338
	9	302.816	323.149	334.102	309.564	332.031	374.020
Machine Learning	1	66.598	39.744	49.337	41.556	55.613	36.799
	3	26.412	22.160	19.965	27.918	35.185	18.148
	4	45.808	44.584	42.293	66.584	50.543	63.678
	6	5.670	3.919	9.155	6.100	5.848	4.945
	7	15.775	20.144	22.418	15.421	17.455	16.652
	8	1,418.318	1,458.029	1,389.691	1,400.373	1,418.102	1,449.745
	10	89.020	68.834	79.560	82.335	45.746	79.545

Table 3-2 Use Case Elapsed Times

### 3.4 SUT Validation Test Output

<u>Validation Run Report</u>			
AIUCpm@1	215.00	T <sub>Load</sub>	0.52
Scale Factor	1	T <sub>LD</sub>	0.52
Streams	100	T <sub>PTT</sub>	34.16
Kit Version	1.0.3.1	T <sub>PST1</sub>	4.84
Execution Status	Pass	T <sub>PST2</sub>	4.86
Accuracy Status	Pass	T <sub>PST</sub>	4.86
		T <sub>TT</sub>	0.70
Test Times			
Overall Run Start Time	2023-09-07 07:34:05.218		
Overall Run End Time	2023-09-07 08:25:12.481		
Overall Run Elapsed Time	3,067.263		
Load Test Start Time	2023-09-07 07:35:42.332		
Load Test End Time	2023-09-07 07:35:42.867		
Load Test Elapsed Time	0.535		
Power Training Start Time	2023-09-07 07:35:42.868		
Power Training End Time	2023-09-07 08:05:30.342		
Power Training Elapsed Time	1,787.474		
Power Serving 1 Start Time	2023-09-07 08:05:30.343		
Power Serving 1 End Time	2023-09-07 08:07:38.652		
Power Serving 1 Elapsed Time	128.309		
Power Serving 2 Start Time	2023-09-07 08:07:38.654		
Power Serving 2 End Time	2023-09-07 08:09:46.953		
Power Serving 2 Elapsed Time	128.299		
Scoring Start Time	2023-09-07 08:10:35.723		
Scoring End Time	2023-09-07 08:13:23.492		
Scoring Elapsed Time	167.769		
Throughput Start Time	2023-09-07 08:13:23.514		
Throughput End Time	2023-09-07 08:25:12.478		
Throughput Elapsed Time	708.964		
(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	0.247	<=	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	0.443	>=	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	960	1,920

Total Storage (GB)	1,920
Scale Factor	10
Data Storage Ratio	192.00

### 4.5 Scale Factor to Memory Ratio

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

Node Count	Memory (GiB)	Total (GiB)
1	384	384

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

### 4.6 Output of Tests

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

### 4.7 Additional Sponsor Files

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

### 4.8 Model Optimizations

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



## Clause 5 – Metrics and Scale Factor

### 5.1 Reported Performance Metrics

#### Metric Overview

TPCx-AI Performance Metric	506.30	AIUCpm@10
TPCx-AI Price/Performance Metric	76.42	\$/AIUCpm@10
TPCx-AI Scale Factor	10	
TPCx-AI Stream Count	100	

#### Test Times


Overall Run Start Time	2023-09-07 08:25:43.099
Overall Run End Time	2023-09-07 11:30:52.416
Overall Run Elapsed Time	11,109.317
Load Test Start Time	2023-09-07 08:28:22.567
Load Test End Time	2023-09-07 08:28:25.818
Load Test Elapsed Time	3.251
Power Training Start Time	2023-09-07 08:28:25.819
Power Training End Time	2023-09-07 10:31:13.971
Power Training Elapsed Time	7,368.152
Power Serving 1 Start Time	2023-09-07 10:31:13.972
Power Serving 1 End Time	2023-09-07 10:41:45.351
Power Serving 1 Elapsed Time	631.379
Power Serving 2 Start Time	2023-09-07 10:41:45.352
Power Serving 2 End Time	2023-09-07 10:52:16.551
Power Serving 2 Elapsed Time	631.199
Scoring Start Time	2023-09-07 10:53:08.853
Scoring End Time	2023-09-07 10:55:46.128
Scoring Elapsed Time	157.275
Throughput Start Time	2023-09-07 10:55:46.150
Throughput End Time	2023-09-07 11:30:52.413
Throughput Elapsed Time	2,106.263

Accuracy Metrics

Use Case	Metric Name	Metric	Criteria	Threshold	Status
1	N/A	0.000	N/A	0.00	Pass
2	word_error_rate	0.480	<=	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	0.079	<=	0.50	Pass
6	matthews_corrcoef	0.505	>=	0.19	Pass
7	median_absolute_error	1.033	<=	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, 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

September 13, 2023

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

Platform: 1x Dell PowerEdge C6615  
 Operating System: Red Hat Enterprise Linux 8.6 (Ootpa)  
 Additional Software: Anaconda Pro

The results were:

**Performance Metric 506.30 AIUCpm@10**

Secondary Metrics	T <sub>LD</sub>	3.23
	T <sub>PTT</sub>	181.16
	T <sub>PST</sub>	16.03
	T <sub>TT</sub>	2.10

**System Under Test 1x Dell PowerEdge C6615 with:**

CPU's	1x AMD EPYC 8324P 32-Core Processor		
Memory	384 GiB		
Storage	<b>Qty</b>	<b>Size</b>	<b>Type</b>
	2	960 GB	M.2 NVMe SSD

In my opinion, these performance results were produced in compliance with the TPC requirements for the benchmark.

The following verification items were given special attention:

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

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 Support Quote

Effective Date: September 7, 2023

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

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

**Quote:**

\$ USD:

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



Included packages:

package name	source	version
conda	main-anaconda	23.7.2
python	main-anaconda	3.9.18
setuptools	main-anaconda	59.8.0
pandas	main-anaconda	1.5.3
scikit-learn	main-anaconda	1.2.2
xgboost	main-anaconda	1.7.4
numpy	main-anaconda	1.23.5
nose	main-anaconda	1.3.7
scipy	main-anaconda	1.10.1
statsmodels	main-anaconda	0.13.5
patsy	main-anaconda	0.5.3
tqdm	main-anaconda	4.65.2
keras	main-anaconda	2.11.0
tensorflow	main-anaconda	2.11.0
joblib	main-anaconda	1.2.0
opencv	main-anaconda	4.5.3
pyyaml	main-anaconda	6.0.1
matplotlib	main-anaconda	3.7.1
jinja2	main-anaconda	3.1.2



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

Anaconda Inc.  
 1108 Lavaca Street Suite 110-645  
 Austin, TX, 78701, USA

# 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.