Nettrix 宁畅 **Nettrix Information Industry Co., LTD.**

TPC Express Benchmark™ AI Full Disclosure Report

R620 G40

using

Anaconda Pro

running on

Red Hat Enterprise Linux v8.2

TPCx-AI Version	1.0.1
Report Edition	First
Report Submitted	May 31, 2022

First Edition - May 2022

Nettrix Information Industry Co., LTD. (Nettrix Information Industry Co., LTD.)*, 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.

Nettrix Information Industry Co., LTD. and the Nettrix Information Industry Co., LTD. Logo are trademarks of Nettrix Information Industry Co., LTD. and/or its affiliates in China 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 Nettrix Information Industry Co., LTD. and any other company.

TPC Express Benchmark™ AI, TPCx-AI, and AIUCpm@10, are registered certification marks of the Transaction Processing Performance Council.

The Nettrix Information Industry Co., LTD. 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 Nettrix Information Industry Co., LTD. business contact for information on the products or services available in your area. You can find additional information via Nettrix Information Industry Co., LTD.'s web site at https://www.nettrix.com.cn/en. Actual performance and environmental costs of Nettrix Information Industry Co., LTD. products will vary depending on individual customer configurations and conditions.

Abstract

Nettrix Information Industry Co., LTD. conducted the TPC Express Benchmark™ AI (TPCx-AI) on the R620 G40. 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 1.0.1.

***Configuration** Overview*

Test Sponsor	Node(s)	Operating System
Nettrix Information Industry Co., LTD.	1x R620G40 (Server)	Red Hat Enterprise Linux v8.2

***Metrics** Overview*

Total System Cost	Performance	Price/Performance	Availability Date
$24,548 USD	291.35 AIUCpm@10	84.26 USD $/AIUCpm@10	May 31, 2022

Executive Summary

The Executive Summary follows on the next several pages.

Nettrix宁畅	R620 G40		TPCx-AI 1.0.1 TPC Pricing 2.8.0 Report Date May 31, 2022
TPCx-AI Performance **291.35 AIUCpm@10**	Total System Cost **$24,548 USD**	Price/Performance **$84.26 USD/AIUCpm@10**	Availability Date **May 31, 2022**

Framework	Operating System	Other Software	Scale Factor	Streams
Anaconda Pro	Red Hat Enterprise Linux v8.2	N/A	10	100

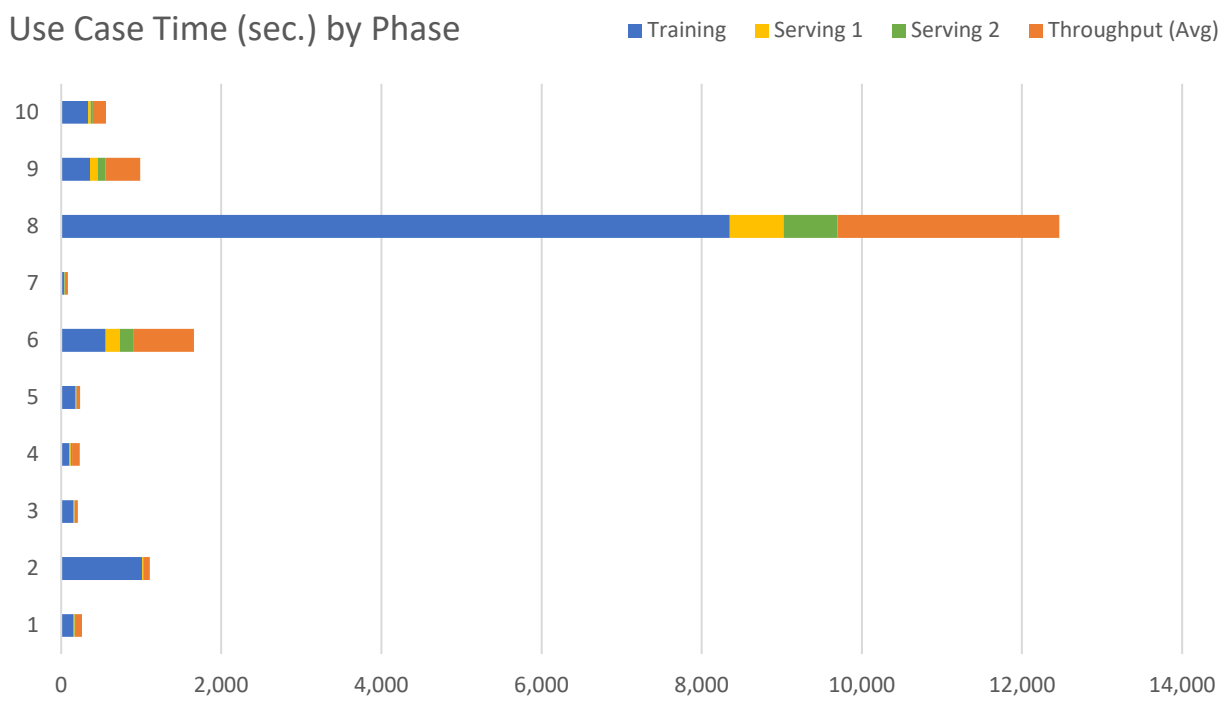

Physical Storage / Scale Factor	Scale Factor / Physical Memory	Main Data Redundancy Model
1,176.00	0.02	RAID 5

Servers: Total Processors/Cores/Threads	1 2 / 32 / 64
Server Type	1x R620G40 (Server)
Processors	2x Intel(R) Xeon(R) Gold 6346 CPU @ 3.10GHz GHz
Memory	512 GiB
Storage Controller	1x AVAGO MegaRAID SAS 9361-8i
Storage Devices	1x 240G M.2 SATA 6G SSD 3x 3.84TB SATA SSD
Network Controller	1x INTEL BLD I350 1G RJ45 2-Ports PCIe NIC

Nettrix宁畅	R620 G40	TPCx-AI	1.0.1
		TPC Pricing	2.8.0
		Report Date	May 31, 2022

Description	Part Number	Source	List Price	Qty	Extended Price	1-Yr. Maintenance
Server						
R620G40	6101550402481540	1	$13,749.00	1	$13,749.00	
Intel® Xeon® Gold 6346 CPU @ 3.10GHz	CD8068904570201	1	(included)	2		
Hynix Hynix HMA84GR7DJR4N-XN 32 GiB	HMA84GR7DJR4N-XN	1	(included)	32		
IT SSDSCKKB240GB 240G M.2 SATA6G R SSD	SSDSCKKB240G801	1	(included)	1		
INTEL BLD I350 1G RJ45 2-Ports PCIe NIC	I350T4G2P20	1	(included)	1		
AVAGO MegaRAID SAS 9361-8i	032542008A	1	(included)	1		
SAMSUNG MZ7L33T8 3.84TB	MZ7L33T8	1	(included)	3		
Hardware Support (24x7x4hr)		1	(included)	1		(included)
Red Hat Enterprise Linux 8.2		1	(included)	1		(included)
RHEL Support (24x7x4hr)		1	$799.00	1		$799.00
				Subtotal	$13,749.00	$799.00
Software						
Anaconda Pro Subscription - 1 year with Premium Support		2	$10,000.00	1		$10,000.00
				Subtotal	$0.00	$10,000.00

Pricing: 1 = Nettrix; 2 = Anaconda Inc.

Audited by Doug Johnson, InfoSizing

Total System Cost (USD):	**$24,548**
AIUCpm@10:	**291.35**
$/AIUCpm@10:	**$84.26**

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.

Nettrix宁畅	R620 G40	TPCx-AI	1.0.1
		TPC Pricing	2.8.0
		Report Date	May 31, 2022

Numerical Quantities

AIUCpm@10	**291.35**	T_{Load}	4.53
Scale Factor	10	T_{LD}	4.53
Streams	100	T_{PTT}	314.80
		T_{PST1}	26.89
Kit Version	1.0.1	T_{PST2}	26.65
Execution Status	Pass	T_{PST}	26.89
Accuracy Status	Pass	T_{TT}	4.69

Test Times

Overall Run Start Time	2022-04-14 22:10:35.264
Overall Run End Time	2022-04-15 03:18:01.119
Overall Run Elapsed Time	18,445.855
Load Test Start Time	2022-04-14 22:13:08.514
Load Test End Time	2022-04-14 22:13:13.055
Load Test Elapsed Time	4.541
Power Training Start Time	2022-04-14 22:13:13.057
Power Training End Time	2022-04-15 01:20:45.169
Power Training Elapsed Time	11,252.112
Power Serving 1 Start Time	2022-04-15 01:20:45.172
Power Serving 1 End Time	2022-04-15 01:37:59.775
Power Serving 1 Elapsed Time	1,034.603
Power Serving 2 Start Time	2022-04-15 01:37:59.778
Power Serving 2 End Time	2022-04-15 01:55:18.756
Power Serving 2 Elapsed Time	1,038.978
Scoring Start Time	2022-04-15 01:56:13.008
Scoring End Time	2022-04-15 01:59:43.287
Scoring Elapsed Time	210.279
Throughput Start Time	2022-04-15 01:59:43.314
Throughput End Time	2022-04-15 03:18:01.116
Throughput Elapsed Time	4,697.802

Nettrix宁畅	R620 G40	TPCx-AI	1.0.1
		TPC Pricing	2.8.0
		Report Date	May 31, 2022

Numerical Quantities (continued)

Use Case Times & Accuracy

Use Case	Training (sec)	Serving 1 (sec)	Serving 2 (sec)	Throughput (avg)	Accuracy
UC01	156.444	14.325	14.221	74.152	0.000
UC02	1,012.214	10.488	9.484	73.545	0.451
UC03	157.823	6.293	6.254	37.705	3.614
UC04	105.676	15.602	15.571	94.434	0.707
UC05	179.304	5.965	6.074	43.860	0.385
UC06	557.347	176.172	176.270	750.221	0.548
UC07	35.886	6.397	6.433	34.853	1.031
UC08	8,351.830	670.424	676.107	2,769.088	0.738
UC09	358.789	98.687	98.529	431.751	1.000
UC10	336.704	30.151	29.940	162.004	0.816

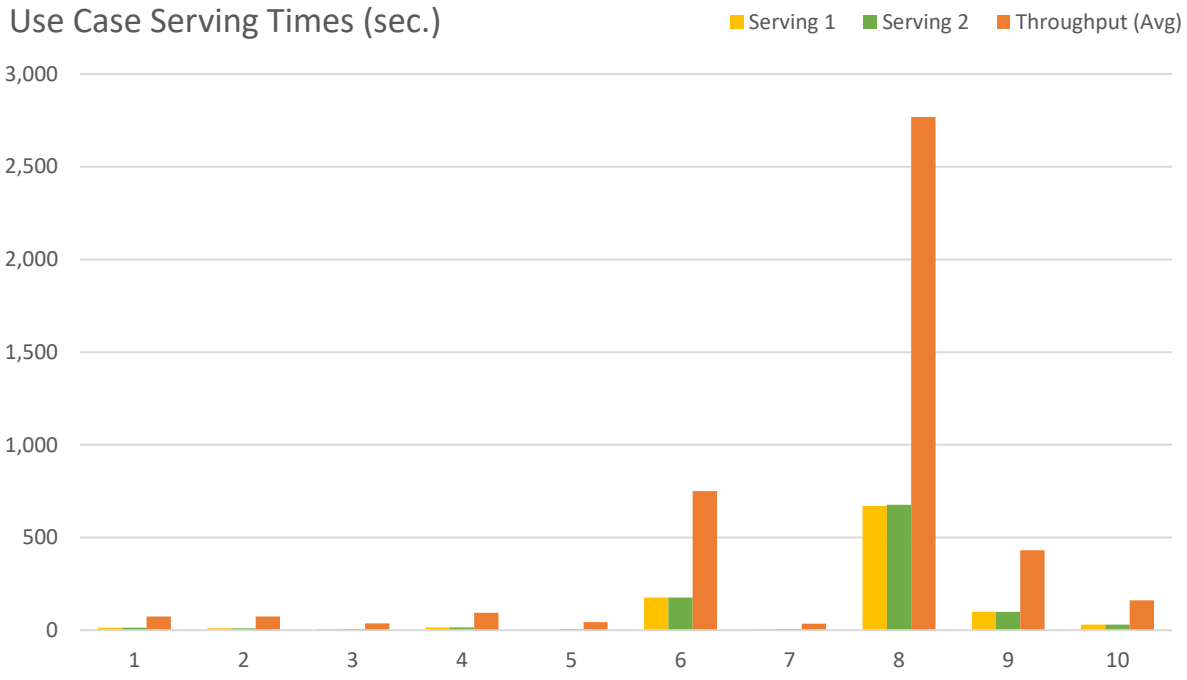

Table of Contents

Clause 0 – Preamble

0.1 TPC Express Benchmark™ AI Overview

Artificial intelligence (AI) has become a key transformational technology of our times. Advances in neural networks and other machine learning techniques have made it possible to use AI on a variety of use cases. From the public sector to aerospace, defense and academia, new and improved ways to use AI techniques are changing the way we harness data and analytics. This along with advances in compute, interconnect and memory technologies have made possible to solve complicated challenges that will ultimately benefit customers in production datacenter and cloud environments.

Abundant volumes of rich data from text, images, audio and video are the essential starting point for creating a benchmark that would represent the myriad of use cases and customers. TPC Express Benchmark™ AI (TPCx-AI) is created in keeping with the TPC tradition of emulating real world AI scenarios and data science use cases. Unlike most other AI benchmarks, the TPCx-AI uses a diverse dataset and is able to scale across a wide range of scale factors. TPCx-AI may later expand with additional use cases and add additional flexibility for a greater variety of implementations.

The benchmark defines and provides a means to evaluate the System Under Test (SUT) performance as a general-purpose data science system that:

- Generates and processes large volumes of data.
- Trains preprocessed data to produce realistic machine learning models.
- Conducts accurate insights for real-world customer scenarios based on the generated models.
- Can scale to large scale distributed configurations.
- Allows for flexibility in configuration changes to meet the demands of the dynamic AI landscape.

The benchmark models real-life examples of companies and public-sector organizations that use a range of analytics techniques, both AI and more traditional machine learning approaches, as well as the potential application of these techniques in situations like those in which they have already been successfully deployed. In addition, the benchmark measures end to end time to provide insights for individual use cases, as well as throughput metrics to simulate multiuser environments for a given hardware, operating system, and data processing system configuration under a controlled, complex, multi-user AI or machine learning data science workload.

The purpose of TPC benchmarks is to provide relevant, objective performance data to industry users. To achieve that purpose, TPC benchmark specifications require benchmark runs be implemented with systems, products, technologies and pricing that:

- Are generally available to users.
- Are relevant to the market segment that the individual TPC benchmark models or represents (e.g., TPCx-AI models and represents complex, high data volume, decision support environments).
- Would plausibly be implemented.

The TPCx-AI kit is available from the TPC website (see www.tpc.org/tpcx-ai/ for more information). Users must sign up and agree to the TPCx-AI End User Licensing Agreement (EULA) to download the kit. All related work (such as collaterals, papers, derivatives) must acknowledge the TPC and include the TPCx-AI copyright. The TPCx-AI kit includes: TPCx-AI Specification document (this document), TPCx-AI Users Guide (README.md) documentation, scripts to set up the benchmark environment, code to execute the benchmark workload, Data Generator, use case related files, and Benchmark Driver.

The use of new systems, products, technologies (hardware or software) and pricing is encouraged so long as they meet the requirements above. Specifically prohibited are benchmark systems, products, technologies or pricing (hereafter referred to as "implementations") whose primary purpose is performance optimization of TPC benchmark results without any corresponding applicability to real-world applications and environments. In other words, all "benchmark special" implementations that improve benchmark results but not real-world performance or pricing, are prohibited.

The rules for pricing are included in the TPC Pricing Specification.

Further information is available at www.tpc.org.

Clause 1 – General Items

1.1 Test Sponsor

This benchmark was sponsored by Nettrix Information Industry Co., LTD..

1.2 Parameter Settings

The Supporting Files Archive contains the parameters and options used to configure the components involved in this benchmark.

1.3 Configuration Diagrams

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

1.3.1 Measured Configuration

Nodes:	1		
Processors/Cores/Threads:	2/32/64	Storage Devices:	4
Total Memory:	512 GiB	Storage Capacity:	11,760 GB

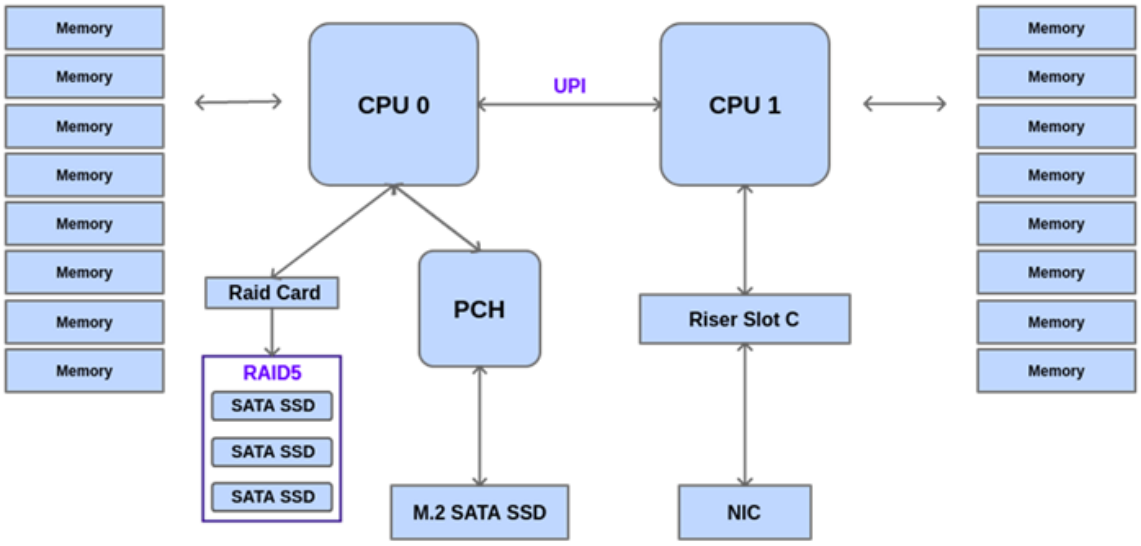

	Server
Server	1x R620G40
Procs/Cores/Threads:	2/16/32
Processor Model:	2x Intel(R) Xeon(R) Gold 6346 CPU @ 3.10GHz
Memory:	512 GiB
Storage Controller:	1x AVAGO MegaRAID SAS 9361-8i
Storage Devices:	1x 240G M.2 SATA 6G SSD 3x 3.84TB SATA SSD
Network Controller:	1x INTEL BLD I350 1G RJ45 2-Ports PCIe NIC

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 R620G40	tpcx	All	1x 240G M.2 SATA 6G SSD 3x 3.84TB SATA 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 v8.2 was used for data generation and the Load Test. The data set was not relocated after generation and before the Load Test.

2.3 Execution Engine, Frameworks, Driver & Libraries

Anaconda Pro consisted of the following components.

Component	Version
python	3.7
setuptools	58
pandas	1.2.4
scikitlearn	1.0.2
xgboost	1.5.0
numpy	1.19.2
nose	1.3.7
scipy	1.7.3
statsmodels	0.12.2
patsy	0.5.2
tqdm	4.62
keras	2.3.1
tensorflow	2.1
joblib	1.1.0
pyyaml	6
jinja2	3.0.2
opencv	3.4.2

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

TPCx-AI Kit Version	1.0.1
Modified File	Description of Changes
No modifications	n/a

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	T5	T6	T7
Deep Learning	2	10.488	9.4836	86.28457	54.56962	76.51767	54.75372	85.61199	81.00112	67.43781
	5	5.9646	6.0736	46.92587	32.17855	45.59675	42.27636	63.99611	78.15637	42.00244
	9	98.687	98.529	506.7399	415.1166	478.1851	391.158	487.3929	562.1389	431.0178
Machine Learning	1	14.325	14.221	57.22639	112.9629	52.42596	68.31676	65.53	43.7396	114.0909
	3	6.2928	6.2544	69.81558	30.47719	50.69181	28.10899	36.94587	21.43443	60.0896
	4	15.602	15.571	59.93313	80.73138	71.31531	119.4117	137.2967	68.01685	119.9383
	6	176.17	176.27	650.6448	589.4529	598.5928	580.3097	852.893	902.2419	885.3835
	7	6.3965	6.4329	18.19124	25.46839	22.17869	26.80679	31.52506	36.06706	26.53437
	8	670.42	676.11	2979.214	3199.929	3106.983	3225.256	2727.546	2219.494	2475.909
	10	30.151	29.94	143.0473	110.7069	119.1508	137.2929	180.0398	165.6316	146.7739

Type	UC ID	T8	T9	T10	T11	T12	T13	T14	T15	T16
Deep Learning	2	69.43774	59.44314	56.3041	74.88924	61.54225	53.03213	56.59834	88.51073	78.57805
	5	57.01991	44.74422	42.29559	45.48583	41.43825	35.83872	33.73819	43.6692	37.6735
	9	418.5367	372.2703	363.9924	537.8915	383.8438	477.4005	360.2963	385.4303	326.9481
Machine Learning	1	72.62709	71.30178	52.48083	111.8638	102.664	72.04217	50.54771	98.22743	54.98185
	3	67.63539	63.15983	28.58977	26.60063	20.09738	37.8085	31.09227	35.479	51.81137
	4	109.424	55.49187	80.15611	74.64062	120.095	89.89961	57.9807	93.36647	125.5833
	6	840.7484	676.9597	703.375	648.3609	669.8591	883.2097	624.9785	881.1129	652.0876
	7	32.25859	17.43061	19.47953	36.00613	55.50646	27.29072	22.77172	33.25687	33.50861
	8	2485.083	3157.917	3013.389	2945.14	3059.365	2878.295	3030.549	2406.051	2867.021
	10	196.0151	159.2301	102.9095	121.2248	155.8375	136.7148	188.5189	196.3193	197.4672

Type	UC ID	T17	T18	T19	T20	T21	T22	T23	T24	T25
Deep Learning	2	86.94281	60.93657	58.81367	103.8208	79.79735	73.89796	82.18826	62.79167	60.33818
	5	45.73379	33.79046	47.82365	17.50287	58.78643	41.04732	58.49291	41.82662	68.50283
	9	430.7004	372.0382	364.2344	467.5174	467.7625	477.3488	469.5817	376.0761	334.9205
Machine Learning	1	113.2901	51.48213	49.58095	55.14361	94.7032	53.15532	102.5363	53.1045	47.65108
	3	46.2704	30.28731	60.39026	56.81725	50.09299	24.39099	35.53376	31.64547	23.81341
	4	124.3957	130.366	70.93872	130.7622	52.55535	72.62099	62.02836	74.13418	69.7959
	6	834.076	608.6109	662.9307	844.8452	839.0466	564.5273	645.8651	702.4923	645.1828
	7	52.91461	25.45151	31.72248	37.9737	53.06827	36.71891	34.13933	25.61873	24.79539
	8	2355.128	3229.903	3207.711	2592.662	2802.502	2957.73	2703.484	2898.156	3055.804
	10	195.4001	129.2876	118.527	154.7392	160.7588	135.4404	170.3645	194.4865	157.3483

Type	UC ID	T26	T27	T28	T29	T30	T31	T32	T33	T34
Deep Learning	2	101.2855	67.31137	75.32466	58.95637	92.41488	88.91864	65.92239	68.96063	106.9154
	5	70.42074	34.0957	39.88457	32.51221	42.13802	31.05047	38.5667	35.22824	39.9644
	9	440.679	356.4105	466.3308	382.7233	415.0399	463.5994	362.8036	481.2809	397.787
Machine Learning	1	83.26185	110.7304	65.72432	62.75392	61.8184	70.43617	57.40147	62.09485	65.70651
	3	59.71428	22.26517	60.64156	61.90924	57.30649	15.70692	27.91611	28.86168	33.36684
	4	57.12512	63.50035	143.8507	129.5675	114.8977	111.8502	133.5417	127.4639	78.23115
	6	851.025	636.2451	893.7902	872.9064	573.5676	872.6843	587.0954	816.9967	569.0213
	7	56.43828	25.63632	75.22263	27.38215	72.3333	30.54687	38.39934	21.00091	30.79475
	8	2374.071	2969.749	2294.795	2823.909	2883.234	2713.15	2942.148	2517.805	3193.919
	10	192.401	115.3377	136.7657	138.9092	138.776	221.016	133.6886	112.9225	120.3496

Type	UC ID	T35	T36	T37	T38	T39	T40	T41	T42	T43
Deep Learning	2	100.2507	61.47479	95.62219	59.75323	55.87355	65.71064	62.89182	62.15642	67.95694
	5	45.74058	28.8383	70.62766	41.17787	77.13246	57.2093	35.50385	57.93221	13.00653
	9	354.4897	379.3367	371.8117	522.7883	417.6049	469.8751	385.3416	550.8135	372.5549
Machine Learning	1	59.51747	54.3733	65.26183	47.5812	62.72141	114.5752	116.7785	57.86473	53.78415
	3	30.92097	28.20792	21.28485	32.83271	29.5078	34.1621	28.2096	41.68343	23.60739
	4	69.91478	69.10667	95.26349	115.5259	76.23885	141.3935	63.6059	118.5668	131.7555
	6	649.2536	577.0274	903.4352	890.386	581.2585	784.1808	857.1041	647.6382	611.9526
	7	27.50117	21.42582	24.6393	27.03227	32.82947	49.04853	36.56016	20.3493	27.11303
	8	3155.677	2952.318	2785.132	2519.89	3214.348	2361.598	2399.935	2966.839	3248.153
	10	123.5617	183.3869	166.4255	126.9883	123.9506	153.3698	191.1122	145.1016	137.1092

Type	UC ID	T44	T45	T46	T47	T48	T49	T50	T51	T52
Deep Learning	2	87.15627	73.23444	92.90626	58.54782	80.75107	59.04325	97.33718	79.99764	90.73959
	5	43.56423	49.87739	31.72647	48.10165	50.23637	35.84524	31.80259	34.45649	48.36592
	9	424.4776	395.1132	523.5807	535.8054	352.6086	498.226	494.1863	371.2562	487.3725
Machine Learning	1	50.38222	102.24	79.99441	110.1166	61.66346	81.34908	106.2177	40.8908	71.31835
	3	33.58163	34.75141	65.58413	29.82877	31.43491	27.09213	32.57528	48.79476	29.89655
	4	68.82748	96.97153	91.23618	77.52658	62.42001	81.31209	95.5937	67.31097	92.25172
	6	875.2769	616.8828	652.6812	885.7713	916.4373	813.7473	811.5288	668.1941	666.2813
	7	22.90634	58.24841	22.87986	35.10195	60.61912	46.04907	54.93091	57.02667	25.58494
	8	2463.284	2856.57	2932.854	2593.445	2582.085	2513.927	2461.47	2882.743	2894.832
	10	170.2976	178.6491	118.3465	114.0402	234.0858	186.7731	163.8016	237.2967	159.5406

Type	UC ID	T53	T54	T55	T56	T57	T58	T59	T60	T61
Deep Learning	2	119.8396	58.12699	93.88273	75.9306	55.21915	97.7126	58.91734	92.80743	43.91141
	5	48.70412	28.39417	35.64238	38.95665	45.55093	40.83977	47.17911	33.76915	39.02764
	9	359.2169	353.2513	463.6523	381.0954	485.496	354.9035	554.7161	330.859	372.1851
Machine Learning	1	54.66785	92.16343	79.95998	82.88765	54.32808	64.66045	52.59137	50.805	101.4974
	3	26.88736	24.33746	46.49327	38.02952	39.63615	39.3294	23.71109	25.46687	27.0896
	4	92.46085	70.26774	137.3839	101.0084	132.0314	76.78835	70.27381	74.62449	128.3943
	6	880.0803	702.0877	800.1343	892.5685	808.2289	623.6873	868.4967	623.842	645.7964
	7	51.7586	22.27492	25.56283	40.30948	31.78029	22.80007	24.12282	27.53392	28.33597
	8	2789.19	2928.835	2441.662	2525.159	2408.764	3142.783	2650.886	2892.638	3130.637
	10	148.5321	110.6029	162.0295	207.1912	126.954	186.8417	189.7938	193.4024	157.9339

Type	UC ID	T62	T63	T64	T65	T66	T67	T68	T69	T70
Deep Learning	2	75.60876	83.79407	68.70786	60.90296	68.34224	92.26897	76.29088	66.45723	58.76344
	5	61.81266	38.37984	39.49336	46.14254	46.10733	40.57026	70.82581	50.89748	42.77573
	9	432.5545	427.5942	548.4531	472.7915	541.8668	434.1445	389.6991	493.8003	378.2905
Machine Learning	1	74.53289	124.8754	63.08748	76.8026	115.2878	79.64137	65.63355	92.11368	49.83683
	3	44.5606	67.08964	23.0652	29.28387	31.58193	29.22459	29.37535	52.27422	63.92081
	4	119.1941	53.30605	104.2564	167.8778	78.75805	82.42446	82.47002	95.46672	98.67276
	6	859.8952	915.5203	865.3763	819.1243	609.8864	847.6357	836.1338	678.2138	879.8168
	7	31.41819	27.67061	48.75321	67.05678	30.8116	58.15633	19.33899	36.44563	25.06279
	8	2481.237	2617.537	2518.631	2229.658	2675.773	2718.682	2773.176	2665.712	2835.197
	10	165.0999	141.4717	117.64	160.9323	169.4446	198.3305	151.7356	216.0559	153.054

Type	UC ID	T71	T72	T73	T74	T75	T76	T77	T78	T79
Deep Learning	2	52.565	82.49747	68.67275	105.5766	93.02889	46.8799	64.83046	72.3628	66.0911
	5	47.40484	39.17186	45.31869	33.19828	41.45785	39.55737	54.80721	36.13192	40.54341
	9	348.2592	410.5135	535.3722	412.2238	449.8481	389.1452	351.3802	483.3742	489.2797
Machine Learning	1	57.04268	108.5808	56.07114	55.12793	59.78399	93.87575	57.24756	88.38782	64.98452
	3	38.97616	31.09051	24.25993	23.23395	76.79023	22.0603	63.68369	44.16785	32.04009
	4	66.88072	85.13494	69.3146	72.44908	71.50769	118.8646	100.7341	94.24476	82.22894
	6	693.0946	565.4605	703.9011	578.5532	880.7173	875.6897	600.5495	631.2714	881.0042
	7	28.11207	22.58371	22.61883	32.64807	53.86132	59.73234	23.78969	24.61618	24.89754
	8	3185.739	3176.499	2816.863	3017.518	2390.439	2763.278	3184.734	2756.927	2646.4
	10	130.0285	130.0628	123.2276	125.5466	175.0927	181.4202	139.908	226.1619	178.2393

Type	UC ID	T80	T81	T82	T83	T84	T85	T86	T87	T88
Deep Learning	2	65.40562	61.50412	59.51835	101.7495	59.9328	58.62607	59.66792	94.03665	57.00684
	5	56.05362	41.10766	38.32051	43.52586	63.71765	32.32365	34.84997	33.36091	43.92229
	9	402.9203	485.0796	379.8993	408.8393	543.2768	378.3388	400.7752	348.6395	504.0043
Machine Learning	1	81.17594	64.76972	45.52011	92.64266	78.70411	102.2989	52.49406	71.54427	134.2603
	3	29.69846	25.58669	18.35832	38.93816	23.87876	37.32668	67.40266	19.20722	21.16408
	4	94.25358	127.5523	115.078	137.9247	98.96143	77.9865	133.6536	88.52198	96.731
	6	871.3352	641.1331	690.9517	892.8673	631.5367	683.8208	841.9513	667.0102	835.3315
	7	33.75345	60.00121	29.90661	36.60177	32.9006	39.54656	37.39998	37.27033	44.21627
	8	2516.884	2722.427	2912.935	2502.602	2797.591	3130.857	2886.391	2853.15	2327.996
	10	188.3345	174.8938	218.9474	175.7338	131.8047	117.0977	152.2933	250.417	168.9439

Type	UC ID	T89	T90	T91	T92	T93	T94	T95	T96	T97
Deep Learning	2	56.93946	72.8269	66.34875	74.53743	57.71452	61.92975	83.09312	97.87329	59.83445
	5	53.44869	37.71941	41.75049	52.40043	38.91819	44.10869	38.25769	37.48366	35.34993
	9	375.9545	362.0957	494.1179	432.5745	547.6139	478.8272	364.4599	505.3005	508.3365
Machine Learning	1	68.55675	62.01906	59.13502	72.61408	56.53627	91.62328	110.8855	69.52281	99.14864
	3	58.61528	50.62151	31.8205	39.32956	29.78826	62.44994	24.82837	27.31834	52.62115
	4	128.982	130.5938	84.12344	146.3092	58.62863	106.1361	91.58743	69.1761	68.20643
	6	890.9486	828.1246	799.1225	884.5104	880.6676	863.9349	872.5924	862.8358	697.6087
	7	33.93352	66.04576	24.76286	32.17801	49.09204	33.06658	26.53733	29.31824	27.80383
	8	2739.757	2591.666	2441.744	2276.311	2579.608	2543.345	2673.476	2429.153	2765.937
	10	230.8108	234.3853	186.985	169.073	160.2496	176.6136	130.0322	159.0718	121.9053

Type	UC ID	T98	T99	T100
Deep Learning	2	102.1924	64.90473	87.39458
	5	39.22836	50.04497	45.86081
	9	388.3939	422.9163	466.3714
Machine Learning	1	55.58883	60.19916	71.22927
	3	26.77447	56.60015	19.83591
	4	85.07332	65.94996	85.25729
	6	587.2175	860.3554	601.3518
	7	29.20105	32.45426	23.00012
	8	3248.39	2506.707	3019.247
	10	127.3176	230.1897	231.3035

Table 3-2 Use Case Elapsed Times

3.4 SUT Validation Test Output

Validation Run Report

AIUCpm@1	191.85	T_{Load}	0.69
Scale Factor	1	T_{LD}	0.69
Streams	100	T_{PTT}	34.49
		T_{PST1}	4.90
Kit Version	1.0.1	T_{PST2}	4.92
Execution Status	Pass	T_{PST}	4.92
Accuracy Status	Pass	T_{TT}	0.82

Test Times

Overall Run Start Time	2022-04-14 21:16:29.485
Overall Run End Time	2022-04-14 22:07:59.105
Overall Run Elapsed Time	3,089.620
Load Test Start Time	2022-04-14 21:18:11.038
Load Test End Time	2022-04-14 21:18:11.743
Load Test Elapsed Time	0.705
Power Training Start Time	2022-04-14 21:18:11.746
Power Training End Time	2022-04-14 21:45:43.485
Power Training Elapsed Time	1,651.739
Power Serving 1 Start Time	2022-04-14 21:45:43.487
Power Serving 1 End Time	2022-04-14 21:48:00.897
Power Serving 1 Elapsed Time	137.410
Power Serving 2 Start Time	2022-04-14 21:48:00.900
Power Serving 2 End Time	2022-04-14 21:50:18.764
Power Serving 2 Elapsed Time	137.864
Scoring Start Time	2022-04-14 21:51:10.203
Scoring End Time	2022-04-14 21:54:14.731
Scoring Elapsed Time	184.528
Throughput Start Time	2022-04-14 21:54:14.758
Throughput End Time	2022-04-14 22:07:59.101
Throughput Elapsed Time	824.343

(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	0.486	<=	0.50	Pass
3	mean_squared_log_error	4.579	<=	5.40	Pass
4	f1_score	0.701	>=	0.65	Pass
5	mean_squared_log_error	0.333	<=	0.50	Pass
6	matthews_corrcoef	0.462	>=	0.19	Pass
7	median_absolute_error	0.894	<=	1.80	Pass
8	accuracy_score	0.710	>=	0.65	Pass
9	accuracy_score	0.980	>=	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	1	240	240
1	3	3,840	11,520

Total Storage (GB)	11,760
Scale Factor	1,000
Data Storage Ratio	1,176.00

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	512	512

Scale Factor	1,000
Total Memory (GiB)	512
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	291.35	AIUCpm@10
TPCx-AI Price/Performance Metric	84.26	$/AIUCpm@10
TPCx-AI Scale Factor	10	
TPCx-AI Stream Count	100	

Test Times

Overall Run Start Time	2022-04-14 22:10:35.264
Overall Run End Time	2022-04-15 03:18:01.119
Overall Run Elapsed Time	18,445.855
Load Test Start Time	2022-04-14 22:13:08.514
Load Test End Time	2022-04-14 22:13:13.055
Load Test Elapsed Time	4.541
Power Training Start Time	2022-04-14 22:13:13.057
Power Training End Time	2022-04-15 01:20:45.169
Power Training Elapsed Time	11,252.112
Power Serving 1 Start Time	2022-04-15 01:20:45.172
Power Serving 1 End Time	2022-04-15 01:37:59.775
Power Serving 1 Elapsed Time	1,034.603
Power Serving 2 Start Time	2022-04-15 01:37:59.778
Power Serving 2 End Time	2022-04-15 01:55:18.756
Power Serving 2 Elapsed Time	1,038.978
Scoring Start Time	2022-04-15 01:56:13.008
Scoring End Time	2022-04-15 01:59:43.287
Scoring Elapsed Time	210.279
Throughput Start Time	2022-04-15 01:59:43.314
Throughput End Time	2022-04-15 03:18:01.116
Throughput Elapsed Time	4,697.802

Accuracy Metrics

Use Case	Metric Name	Metric	Criteria	Threshold	Status
1	N/A	0.000	N/A	0.00	Pass
2	word_error_rate	0.451	<=	0.50	Pass
3	mean_squared_log_error	3.614	<=	5.40	Pass
4	f1_score	0.707	>=	0.65	Pass
5	mean_squared_log_error	0.385	<=	0.50	Pass
6	matthews_corrcoef	0.548	>=	0.19	Pass
7	median_absolute_error	1.031	<=	1.80	Pass
8	accuracy_score	0.738	>=	0.65	Pass
9	accuracy_score	1.000	>=	0.90	Pass
10	accuracy_score	0.816	>=	0.70	Pass

5.2 Throughput Test Stream Times

The following chart shows the minimum, 1st quartile, median, mean (X), 3rd quartile, and maximum stream times by use case for the Throughput Test. Outliers are marked with “o”.

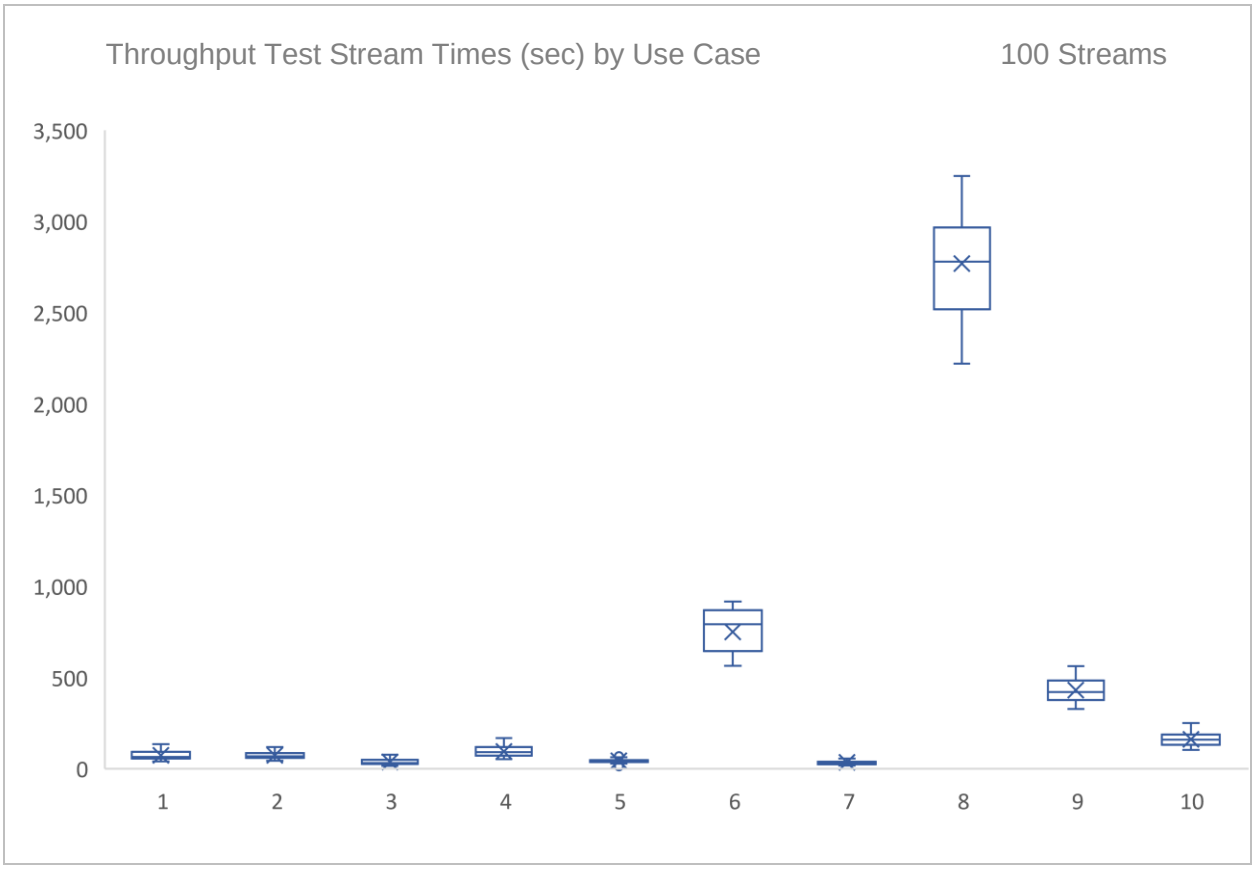

Auditor's Information

This benchmark was audited by Doug Johnson, InfoSizing.

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

This benchmark's Full Disclosure Report can be downloaded from www.tpc.org.

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

InfoSizing
The Right Metric For Sizing IT

TPC
Certified Auditor

Gabby Chen
Nettrix Information Industry Co., LTD.
Validation Center
No.25 Zhongguancun Software Park,
8 Dongbeiwang West Road,
Haidian District, Beijing

May 31, 2022

I verified the TPC Express Benchmark™ AI v1.0.1 performance of the following configuration:

Platform:	R620 G40
Operating System:	Red Hat Enterprise Linux v8.2
Additional Software:	Anaconda Pro

The results were:

Performance Metric	**291.35 AIUCpm@10**	
Secondary Metrics	T_{LD}	4.53
	T_{PTT}	314.80
	T_{PST}	26.89
	T_{TT}	4.69

SUT	**1x R620 G40**		
CPUs	2x Intel(R) Xeon(R) Gold 6346 CPU @ 3.10GHz		
Memory	512 GiB		
Storage	***Qty***	***Size***	***Type***
	1	240 GB	M.2 SATA 6G SSD
	3	3.84 TB	SATA SSC

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 v1.0.1.
- 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.

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

- The generated dataset was properly scaled to 10 GB.
- The generated dataset used for testing was protected by RAID 5.
- 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,

Doug Johnson, Certified TPC Auditor

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

Third-Party Price Quotes

Anaconda Inc.

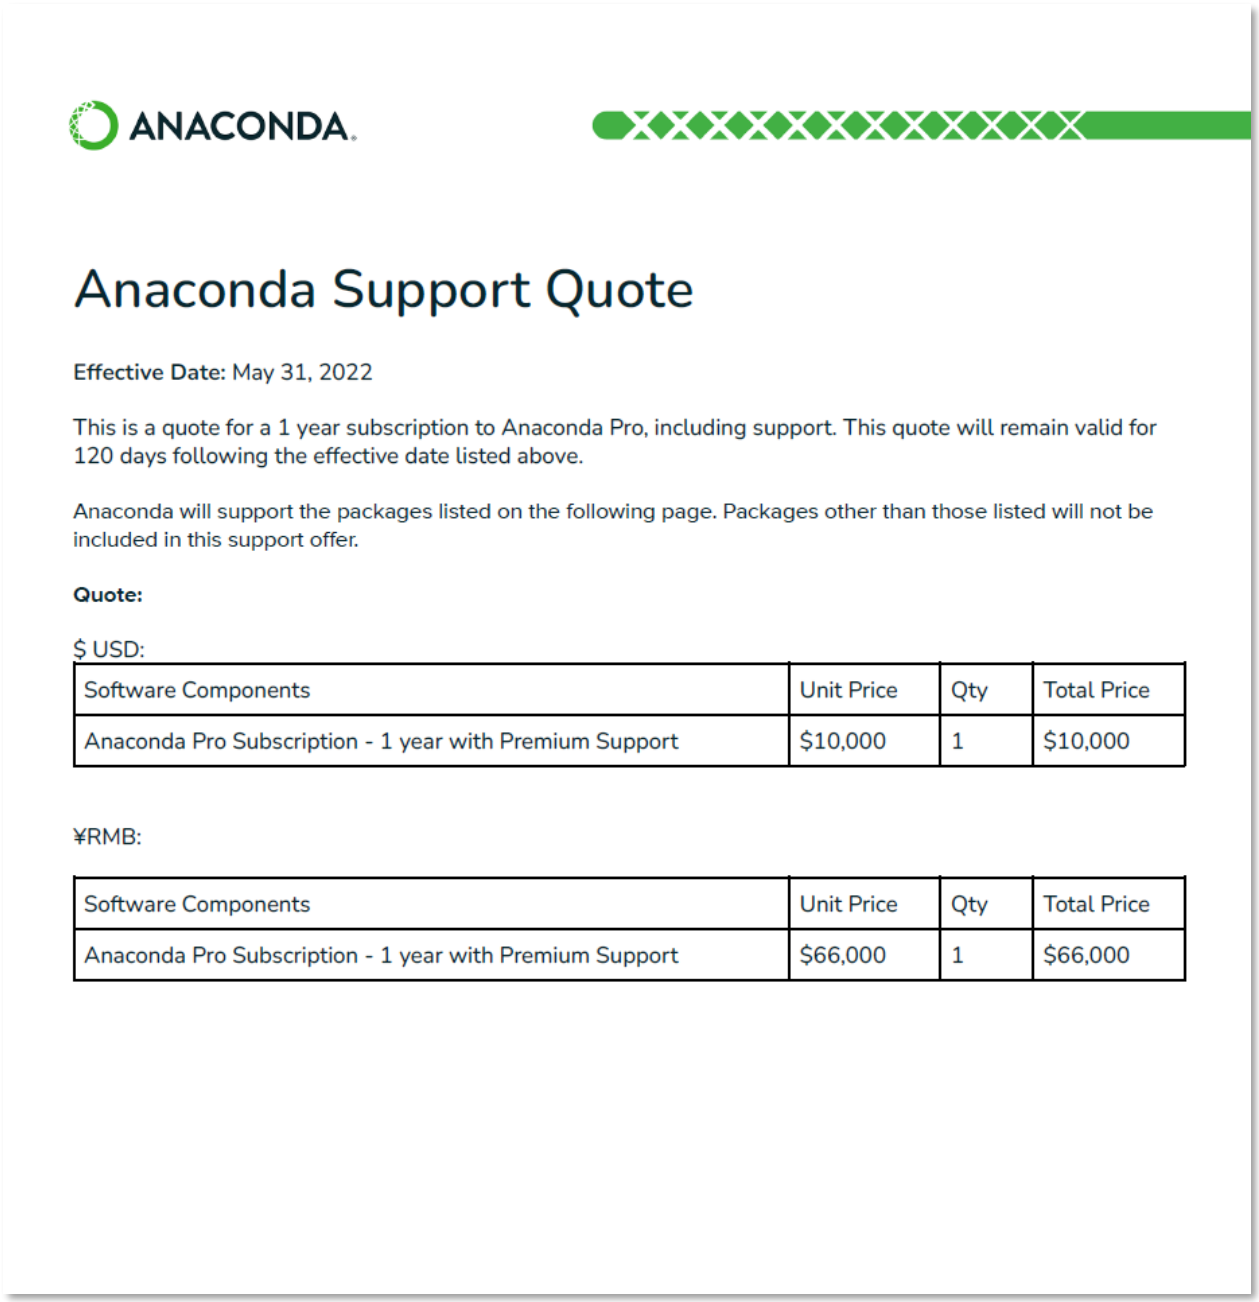

ANACONDA.

Anaconda Support Quote

Effective Date: May 31, 2022

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

¥RMB:

Software Components	Unit Price	Qty	Total Price
Anaconda Pro Subscription - 1 year with Premium Support	$66,000	1	$66,000

ANACONDA.

Included packages:

package name	source	version
python	main-anaconda	3.7
setuptools	main-anaconda	58
pandas	main-anaconda	1.2.4
scikitlearn	main-anaconda	1.0.2
xgboost	main-anaconda	1.5.0
numpy	main-anaconda	1.19.2
nose	main-anaconda	1.3.7
scipy	main-anaconda	1.7.3
statsmodels	main-anaconda	0.12.2
patsy	main-anaconda	0.5.2
tqdm	main-anaconda	4.62
keras	main-anaconda	2.3.1
tensorflow	main-anaconda	2.1
joblib	main-anaconda	1.1.0
pyyaml	main-anaconda	6
jinja2	main-anaconda	3.0.2
opencv	main-anaconda	3.4.2

Contact Sales: 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 Nettrix Information Industry Co., LTD.

Sponsor/ModelOptimization/…	Details of model optimization (if used).
Sponsor/ModifiedKitFiles/…	0 modified file(s).
Sponsor/Tuning/…	All tuning files used.