



Lenovo

TPC Express Benchmark™ AI Full Disclosure Report

ThinkSystem SR655 V3

with 1x ThinkSystem SR655V3
using

Anaconda Business
running on

Red Hat Enterprise Linux 8.10 (Ootpa)

TPCx-AI Version
Report Edition
Report Submitted

2.0.0
First
January 29, 2026

First Edition - January 2026

Lenovo (Lenovo), 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.

Lenovo and the Lenovo Logo are trademarks of Lenovo 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 Lenovo and any other company.

TPC Express Benchmark™ AI, TPCx-AI, and AIUCpm@3, are registered certification marks of the Transaction Processing Performance Council.

The Lenovo 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 Lenovo business contact for information on the products or services available in your area. You can find additional information via Lenovo's web site at <https://www.lenovo.com/us/en/>. Actual performance and environmental costs of Lenovo products will vary depending on individual customer configurations and conditions.

Copyright© 2026 Lenovo

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

Lenovo conducted the TPC Express Benchmark™ AI (TPCx-AI) on the ThinkSystem SR655 V3. 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

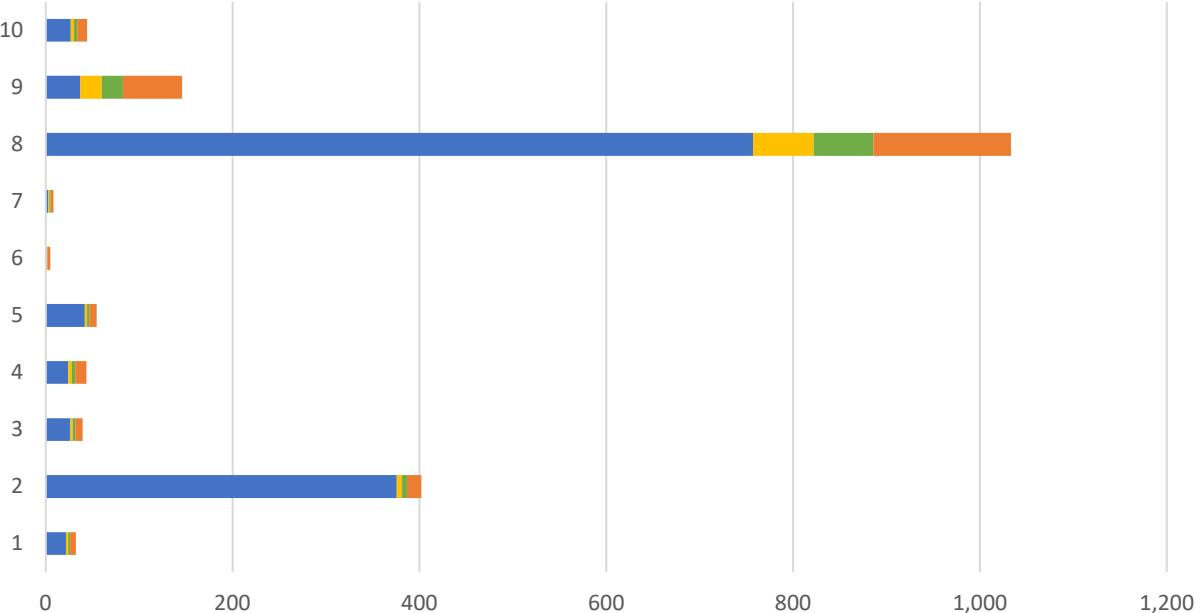
Test Sponsor	Node(s)	Operating System
Lenovo	1x ThinkSystem SR655V3 (Server)	Red Hat Enterprise Linux 8.10 (Ootpa)

Metrics Overview

Total System Cost	Performance	Price/Performance	Availability Date
\$120,776 USD	893.92 AIUCpm@3	135.11 USD \$/AIUCpm@3	January 29, 2026

Executive Summary

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

		ThinkSystem SR655 V3		TPCx-AI	2.0.0						
				TPC Pricing	2.9.0						
				Report Date	Jan. 29, 2026						
TPCx-AI Performance	893.92 AIUCpm@3	Total System Cost	Price/Performance	Availability Date							
Framework	Operating System	Other Software	Scale Factor	Streams							
Anaconda Business	Red Hat Enterprise Linux 8.10 (Ootpa)	N/A	3	100							
Use Case Time (sec.) by Phase		■ Training ■ Serving 1 ■ Serving 2 ■ Throughput (Avg)									
											
Physical Storage / Scale Factor	640.00	Scale Factor / Physical Memory	0.01	Main Data Redundancy Model							
Servers:	1	RAID1									
Total Processors/Cores/Threads	1 / 64 / 128										
Server Type	1x ThinkSystem SR655V3 (Server)										
Processors	1x AMD EPYC 9575F 64-Core Processor										
Memory	384 GiB										
Storage Controller	1x M.2 NVMe 2-Bay RAID Adapter										
Storage Device	2x M.2 VA 960GB Read Intensive NVMe SSD										
Network Controller	1x Broadcom 5719 1GbE RJ45 4-port Adapter										

 ThinkSystem SR655 V3		TPCx-AI	2.0.0	
		TPC Pricing	2.9.0	
Report Date Jan. 29, 2026				
Description	Part Number	Source	List Price	
Server Hardware				
ThinkSystem SR655 V3 Configure-To-Order, includes:	7D9ECTO1WW	1	\$52,675.00	
ThinkSystem V3 2U 24x2.5" Chassis	BLKK	1		
AMD EPYC 9575F 64C 400W 3.3GHz Processor	C4H8	1		
ThinkSystem SR655 V3 2U High Performance Heatsink	BQF9	1		
ThinkSystem 32GB TruDDR5 6400MHz (1Rx4) RDIMM-A	C1PL	12		
ThinkSystem M.2 NVMe 2-Bay RAID Adapter	B8P9	1		
ThinkSystem M.2 VA 960GB Read Intensive NVMe PCIe 4.0x4 NHS SSD	CBT0	2		
ThinkSystem Broadcom 5719 1GbE RJ45 4-port OCP Ethernet Adapter	B5T1	1		
ThinkSystem 1800W 230V Platinum Hot-Swap Gen2 Power Supply v2	BMUF	2		
2.8m, 13A/100-250V, C13 to C14 Jumper Cord	6400	2		
ThinkSystem 2U V3 Performance Fan Module	BLL6	6		
ThinkSystem Toolless Slide Rail Kit v2	B8LA	1		
ThinkSystem 2U V3 EIA Latch Standard	BQQ2	1		
ThinkSystem SR655 V3 MB w/Io+PIB+FB,2U	C2AC	1		
ThinkSystem 8x1 2.5" HDD Filler	AVEQ	3		
MS 2FH Riser Filler	BC4X	1		
ThinkSystem 2U V3 EIA right with FIO	BQQ6	1		
ThinkSystem SR655 V3 Main Airduct	BQ81	1		
ThinkSystem SR635 V3/SR655 V3 Root of Trust Module Low Voltage-RoW V2	C1PT	1		
ThinkSystem MS 2U common Airduct Filler	BQ31	1		
ThinkSystem 2U MS 3FH Riser Filler	B8MM	2		
Red Hat Linux w/Lenovo Support	750FCTO1WW	1		
RHEL Server Physical or Virtual Node, 2 Skt Standard Subscription w/Lenovo Support 1Yr	SON5	1		
3Yr Premier 24x7 4Hr Resp SR655 V3	5WS7C00976	1	\$1,997.00	
		Subtotal	\$52,675.00	
			\$1,997.00	
Server Software				
Anaconda Business Subscription (1 year)		2	\$60,000.00	
Anaconda Business Premium Support, including 24x7 support		2	\$30,000.00	
		Subtotal	\$60,000.00	
			\$30,000.00	
Infrastructure				
ThinkSystem Pref. Pro II USB Keyboard - US English (2 spares)	7ZB7A05522	1	\$79.00	
ThinkSystem Optical Wheel Mouse - USB (2 spares)	7M57A04698	1	\$109.00	
Lenovo L22i-40 21.5-inch IPS FHD Monitor (2 spares)	67AEKCCBUS	1	\$95.00	
		Subtotal	\$849.00	
			\$0.00	
Dollar Volume Discount (See Note 1)	44.57%	1	Total	\$113,524.00
				\$31,997.00
				\$24,745.00
Pricing: 1 = Lenovo; 2 = Anaconda		Total System Cost (USD): \$120,776		
* Discount applies to all line items where Source = 1. Discount based upon total system cost as purchased by a regular customer.		AIUCpm@3: 893.92		
		\$/AIUCpm@3: \$135.11		
Audited by Doug Johnson, InfoSizing				
<i>Prices used in TPC benchmarks reflect the actual prices a customer would pay for a one-time purchase of the stated Line Items. Individually negotiated discounts are not permitted. Special prices based on assumptions about past or future purchases are not permitted. All discounts reflect standard pricing policies for the listed Line Items. For complete details, see the pricing section of the TPC Benchmark Standard. If you find that the stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.</i>				

 ThinkSystem SR655 V3		TPCx-AI	2.0.0
		TPC Pricing	2.9.0
		Report Date	Jan. 29, 2026
<u>Numerical Quantities</u>			
AIUCpm@3	893.92	T_{Load}	0.50
Scale Factor	3	T_{LD}	0.50
Streams	100	T_{PTT}	29.33
Kit Version	2.0.0	T_{PST1}	3.83
Execution Status	Pass	T_{PST2}	3.85
Accuracy Status	Pass	T_{PST}	3.85
		T_{TT}	0.29
Test Times			
Overall Run Start Time		2026-01-15 19:18:04.839	
Overall Run End Time		2026-01-15 19:51:49.712	
Overall Run Elapsed Time		2,024.873	
Load Test Start Time		2026-01-15 19:19:26.062	
Load Test End Time		2026-01-15 19:19:26.574	
Load Test Elapsed Time		0.512	
Power Training Start Time		2026-01-15 19:19:26.575	
Power Training End Time		2026-01-15 19:41:21.674	
Power Training Elapsed Time		1,315.099	
Power Serving 1 Start Time		2026-01-15 19:41:21.676	
Power Serving 1 End Time		2026-01-15 19:43:10.300	
Power Serving 1 Elapsed Time		108.624	
Power Serving 2 Start Time		2026-01-15 19:43:10.302	
Power Serving 2 End Time		2026-01-15 19:44:58.702	
Power Serving 2 Elapsed Time		108.400	
Scoring Start Time		2026-01-15 19:45:34.333	
Scoring End Time		2026-01-15 19:46:56.190	
Scoring Elapsed Time		81.857	
Throughput Start Time		2026-01-15 19:46:56.203	
Throughput End Time		2026-01-15 19:51:49.709	
Throughput Elapsed Time		293.506	

Lenovo	ThinkSystem SR655 V3	TPCx-AI	2.0.0
		TPC Pricing	2.9.0
Report Date Jan. 29, 2026			

Numerical Quantities (continued)

Use Case Times & Accuracy

Use Case	Training (sec)	Serving 1 (sec)	Serving 2 (sec)	Throughput (avg)	Accuracy
UC01	21.693	2.224	2.205	6.072	0.000
UC02	375.870	5.344	5.394	15.654	0.465
UC03	26.504	2.461	2.504	7.979	3.507
UC04	24.221	3.697	3.778	11.734	0.700
UC05	41.806	2.402	2.384	7.805	0.451
UC06	1.009	0.476	0.480	2.998	0.492
UC07	2.824	1.126	1.138	3.377	0.955
UC08	757.113	64.705	64.324	147.212	0.751
UC09	37.113	22.794	22.811	63.233	1.000
UC10	26.891	3.337	3.325	10.661	0.816

Use Case Serving Times (sec.)

■ Serving 1 ■ Serving 2 ■ Throughput (Avg)

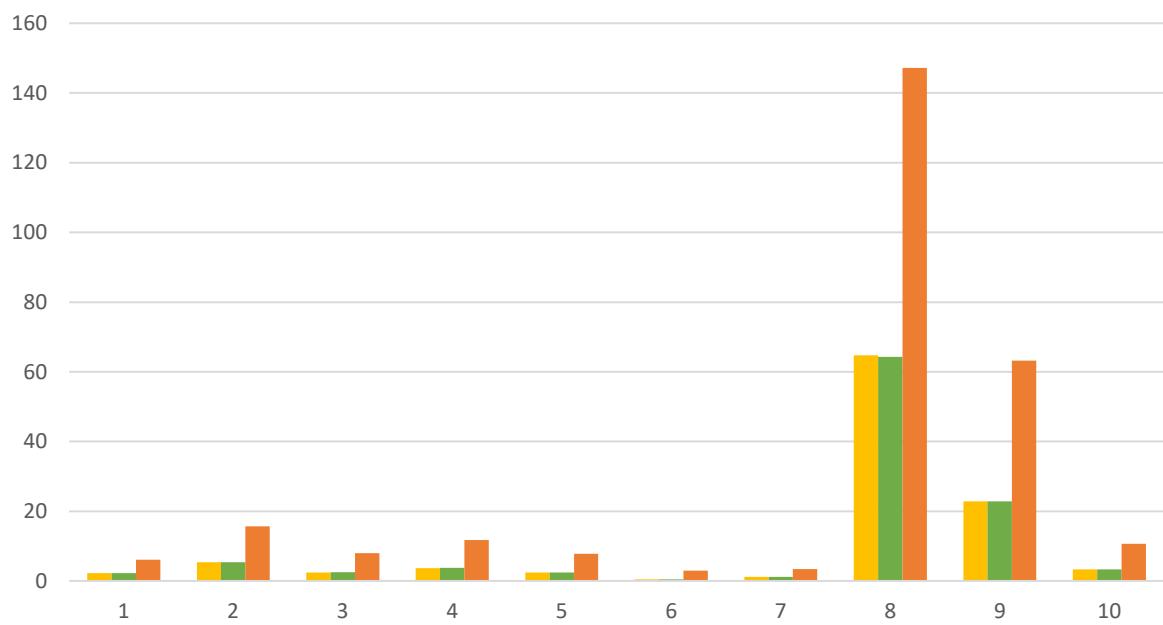


Table of Contents

Abstract.....	3
Executive Summary	3
Table of Contents.....	1
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.....	15
Clause 3 – Workload Related Items.....	16
3.1 Hardware & Software Tuning.....	16
3.2 Kit Version & Modifications	16
3.3 Use Case Elapsed Times	16
3.4 SUT Validation Test Output	23
3.5 Configuration Parameters	24
Clause 4 – SUT Related Items.....	25
4.1 Specialized Hardware/Software	25
4.2 Configuration Files.....	25
4.3 SUT Environment Information.....	25
4.4 Data Storage to Scale Factor Ratio	25
4.5 Scale Factor to Memory Ratio	25
4.6 Output of Tests	25
4.7 Additional Sponsor Files	25
4.8 Model Optimizations	25
Clause 5 – Metrics and Scale Factor	26
5.1 Reported Performance Metrics	26

5.2 Throughput Test Stream Times	27
Auditor's Information	28
Third-Party Price Quotes.....	31
Anaconda	31
Supporting Files Index	32

Clause 0 – Preamble

0.1 TPC Express Benchmark™ AI Overview

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

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

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

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

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

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

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

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

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

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

Further information is available at [www\(tpc.org](http://www(tpc.org).

Clause 1 – General Items

1.1 Test Sponsor

This benchmark was sponsored by Lenovo.

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	Storage Devices:	2
Processors/Cores/Threads:	1/64/128	Storage Capacity:	1,920 GB
Total Memory:	384 GiB		
			
<p><u>Server</u></p> <p>Server</p> <p>1x ThinkSystem SR655V3:</p> <p>Procs/Cores/Threads: 1/64/128</p> <p>Processor Model: 1x AMD EPYC 9575F 64-Core Processor</p> <p>Memory: 384 GiB</p> <p>Storage Controller: 1x M.2 NVMe 2-Bay RAID Adapter</p> <p>Storage Devices: 2x M.2 VA 960GB Read Intensive NVMe SSD</p> <p>Network Controller: 1x Broadcom 5719 1GbE RJ45 4-port Adapter</p>			

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 ThinkSystem SR655V3	SR655V3	All	2x M.2 VA 960GB Read Intensive 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 8.10 (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	3.9.18
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
pyyaml	6.0.2
matplotlib	3.7.1
jinja2	3.1.2
pycryptodome	3.16.0
scikit-surprise	1.1.3
librosa	0.8.1
imbalanced-learn	0.10.1
tensorflow-addons	0.19.0

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
Modified File tools/python/dataRedundancyInformation.sh	Description of Changes Added platform-specific data collection.

Table 3-1 Kit Version & Modifications

3.3 Use Case Elapsed Times

Below are the elapsed times for each use case. Use cases are grouped based on whether they use Deep Learning or Machine Learning techniques.

Type	UC ID	P1	P2	T1	T2	T3	T4
Deep Learning	2	5.344	5.394	17.685	11.907	15.550	14.258
	5	2.402	2.384	10.292	8.701	7.970	8.243
	9	22.794	22.811	67.839	68.710	71.662	61.622
Machine Learning	1	2.224	2.205	5.614	6.511	4.602	7.253
	3	2.461	2.504	9.104	7.685	3.754	8.671
	4	3.697	3.778	10.845	14.560	11.834	13.179
	6	0.476	0.480	3.101	2.971	2.707	3.223
	7	1.126	1.138	3.037	3.024	2.930	3.140
	8	64.705	64.324	141.642	139.219	153.850	151.696
	10	3.337	3.325	11.752	9.770	9.940	9.839

Type	UC ID	T5	T6	T7	T8	T9	T10
Deep Learning	2	12.730	20.800	12.748	15.199	16.186	16.624
	5	5.946	9.005	5.302	11.167	11.793	10.270
	9	70.507	57.538	55.867	56.793	58.232	59.335
Machine Learning	1	6.332	4.903	6.390	5.134	5.871	4.757
	3	12.915	6.810	9.511	9.712	8.760	10.262
	4	10.307	9.846	14.634	10.109	11.742	10.576
	6	3.303	3.405	3.060	3.038	3.217	3.426
	7	4.251	4.718	4.046	2.842	4.621	2.004
	8	145.384	144.790	146.922	142.938	144.074	151.537
	10	10.943	12.052	8.335	13.610	14.251	10.371

Type	UC ID	T11	T12	T13	T14	T15	T16
Deep Learning	2	12.652	17.863	13.883	13.345	18.703	19.013
	5	9.252	5.820	8.348	10.181	5.186	7.874
	9	66.163	57.353	70.698	65.825	55.991	62.642
Machine Learning	1	6.495	8.103	6.261	4.699	7.081	5.898
	3	6.918	3.670	9.886	6.541	9.730	11.162
	4	11.597	13.853	13.814	11.883	10.200	14.283
	6	2.991	3.845	1.077	2.939	2.752	3.335
	7	3.138	3.948	3.716	3.056	3.158	4.420
	8	148.597	156.606	145.095	150.343	136.289	138.295
	10	9.143	14.745	11.999	10.956	11.165	12.100

Type	UC ID	T17	T18	T19	T20	T21	T22
Deep Learning	2	12.674	13.924	13.392	14.743	17.379	15.384
	5	6.190	6.389	9.329	3.028	8.273	8.540
	9	51.943	59.031	60.108	72.345	71.052	71.769
Machine Learning	1	6.152	6.576	6.543	6.136	2.974	5.638
	3	10.428	9.164	7.592	10.434	5.115	6.911
	4	15.235	12.520	10.281	6.959	13.922	8.986
	6	3.583	3.241	3.861	2.905	2.989	3.103
	7	4.856	2.461	3.838	4.796	3.608	2.944
	8	155.660	128.300	135.352	154.719	147.034	147.918
	10	9.274	13.072	9.683	11.846	13.031	11.191

Type	UC ID	T23	T24	T25	T26	T27	T28
Deep Learning	2	16.142	18.361	14.013	14.443	13.745	13.234
	5	7.359	7.695	8.291	9.677	4.977	7.300
	9	55.840	71.892	59.648	59.091	72.926	66.365
Machine Learning	1	7.546	5.551	6.795	5.523	6.655	5.008
	3	7.276	8.241	6.786	4.815	7.078	8.856
	4	12.970	10.061	9.642	10.503	12.961	8.797
	6	3.666	2.118	3.120	3.030	3.142	2.804
	7	3.388	3.550	2.971	3.430	3.707	2.827
	8	153.343	148.424	150.970	131.565	148.617	142.551
	10	11.644	7.544	13.642	9.434	10.890	8.972

Type	UC ID	T29	T30	T31	T32	T33	T34
Deep Learning	2	15.630	17.529	6.923	15.157	13.696	17.923
	5	10.305	8.240	5.901	9.483	8.123	6.565
	9	70.663	70.494	72.027	62.825	66.835	64.787
Machine Learning	1	5.482	6.153	5.894	5.456	7.764	6.414
	3	3.770	10.596	12.043	6.574	6.328	8.144
	4	12.407	6.082	13.160	13.988	12.354	11.008
	6	1.948	2.444	3.463	2.966	2.488	2.949
	7	4.795	3.393	3.906	3.338	3.228	3.886
	8	151.156	149.959	154.381	141.286	135.293	145.490
	10	9.521	10.340	11.588	12.182	9.317	8.290

Type	UC ID	T35	T36	T37	T38	T39	T40
Deep Learning	2	16.793	12.019	21.145	16.514	11.548	13.264
	5	5.251	7.290	8.341	7.023	9.105	5.517
	9	64.518	57.279	67.714	63.970	70.027	69.583
Machine Learning	1	4.975	7.187	6.118	5.227	6.296	6.764
	3	8.219	6.863	7.659	6.380	8.774	6.804
	4	12.785	11.158	13.426	12.523	12.907	10.205
	6	3.062	3.093	1.302	2.871	3.039	3.164
	7	3.386	4.937	3.335	3.388	3.523	2.322
	8	148.953	150.246	146.943	149.272	149.807	147.470
	10	9.742	11.957	8.407	8.855	8.819	8.369

Type	UC ID	T41	T42	T43	T44	T45	T46
Deep Learning	2	16.375	16.430	15.303	17.666	17.974	20.696
	5	9.949	8.291	6.002	8.064	10.386	10.624
	9	53.808	54.284	59.568	55.509	65.493	49.962
Machine Learning	1	8.604	5.583	6.843	5.393	5.378	7.034
	3	6.691	7.149	7.990	6.675	12.328	9.462
	4	12.457	15.009	13.408	11.592	14.155	12.926
	6	2.554	3.654	3.201	3.452	2.987	3.182
	7	2.850	2.749	3.865	2.828	3.777	3.071
	8	141.941	154.369	154.352	139.024	147.077	143.197
	10	13.191	12.972	11.736	12.755	5.960	10.954

Type	UC ID	T47	T48	T49	T50	T51	T52
Deep Learning	2	14.852	16.107	13.355	12.858	16.122	19.438
	5	6.890	11.181	7.419	5.961	8.426	5.536
	9	61.337	63.230	67.011	65.233	63.284	57.534
Machine Learning	1	6.323	6.813	5.544	9.046	8.001	5.245
	3	5.578	10.294	7.562	6.781	7.716	5.836
	4	15.292	13.297	10.037	10.170	12.385	14.259
	6	2.638	3.120	2.575	2.504	4.051	3.565
	7	3.771	4.603	2.546	2.051	3.885	3.480
	8	155.890	135.294	153.335	146.270	135.377	154.969
	10	11.341	13.462	9.295	8.359	13.142	8.590

Type	UC ID	T53	T54	T55	T56	T57	T58
Deep Learning	2	17.933	19.046	18.210	13.709	12.810	18.605
	5	9.557	7.035	6.107	5.782	6.650	8.347
	9	60.107	62.361	55.565	67.244	67.882	60.017
Machine Learning	1	4.793	9.013	4.664	6.866	4.554	5.713
	3	7.618	8.255	6.590	10.246	6.935	8.394
	4	14.039	9.957	8.308	9.659	13.036	11.530
	6	2.949	3.304	3.218	2.769	2.606	2.936
	7	2.256	3.643	2.943	3.043	2.786	3.113
	8	151.850	135.694	144.155	158.770	142.599	150.996
	10	11.424	9.765	10.523	6.844	13.392	9.841

Type	UC ID	T59	T60	T61	T62	T63	T64
Deep Learning	2	14.855	15.821	8.439	16.609	17.935	12.548
	5	7.279	6.002	6.952	8.466	8.259	6.404
	9	64.119	60.421	62.202	58.853	64.347	63.675
Machine Learning	1	5.309	5.328	6.224	5.485	7.148	5.054
	3	5.881	7.240	8.557	8.980	7.962	6.701
	4	15.212	15.102	15.199	12.734	6.677	10.411
	6	3.839	2.842	3.351	2.866	1.673	3.223
	7	3.330	3.160	4.013	2.795	5.026	2.993
	8	143.957	153.627	158.127	147.613	155.997	148.997
	10	11.252	10.885	13.198	8.531	11.430	9.005

Type	UC ID	T65	T66	T67	T68	T69	T70
Deep Learning	2	11.571	19.433	14.775	20.345	20.028	15.261
	5	10.026	6.460	8.386	9.542	6.122	7.487
	9	66.648	64.732	69.017	59.212	57.743	54.341
Machine Learning	1	5.386	6.576	8.322	5.570	6.541	4.696
	3	7.249	6.905	10.415	11.980	6.703	7.692
	4	9.280	11.712	12.278	10.587	10.694	11.969
	6	2.996	3.200	3.904	1.383	2.945	2.685
	7	4.380	3.118	3.164	2.735	3.542	3.989
	8	138.196	143.961	132.204	151.539	150.792	159.318
	10	8.525	8.079	10.176	9.683	12.377	11.280

Type	UC ID	T71	T72	T73	T74	T75	T76
Deep Learning	2	14.868	12.165	19.425	17.287	18.388	16.012
	5	7.105	7.419	7.748	9.820	6.725	6.875
	9	60.607	67.900	62.766	64.896	59.616	64.632
Machine Learning	1	7.187	6.607	6.389	5.234	5.005	5.881
	3	8.374	8.549	8.060	6.933	8.574	11.873
	4	11.505	13.940	9.504	10.563	8.964	12.640
	6	3.375	3.075	3.314	2.788	3.013	3.252
	7	3.286	3.234	2.611	3.488	2.357	2.964
	8	149.963	150.806	151.091	148.763	145.420	141.941
	10	11.805	8.823	9.963	8.970	7.503	10.689

Type	UC ID	T77	T78	T79	T80	T81	T82
Deep Learning	2	16.356	15.571	15.541	10.258	15.537	18.640
	5	9.855	7.924	5.671	8.873	9.820	7.063
	9	61.651	72.764	71.236	56.925	56.629	61.233
Machine Learning	1	6.255	2.587	5.703	5.092	5.006	3.166
	3	7.238	4.464	9.428	6.260	6.930	6.299
	4	14.623	12.520	13.005	10.560	15.959	13.598
	6	3.291	3.290	1.137	3.359	3.591	3.440
	7	2.810	3.774	3.646	2.805	3.866	3.519
	8	143.411	154.017	147.685	134.836	139.334	152.099
	10	9.620	9.833	9.672	9.500	12.425	15.633

Type	UC ID	T83	T84	T85	T86	T87	T88
Deep Learning	2	20.043	13.315	15.579	14.626	18.862	13.347
	5	8.139	8.239	9.143	8.056	8.501	6.256
	9	63.357	59.680	59.032	64.447	59.225	65.476
Machine Learning	1	2.920	6.478	5.514	8.636	7.149	4.905
	3	10.041	7.690	11.019	7.842	5.602	6.721
	4	9.199	13.778	11.461	15.550	14.322	9.215
	6	2.967	3.365	2.892	2.815	3.943	3.053
	7	4.011	2.907	2.691	4.088	3.006	1.734
	8	150.462	150.746	155.404	138.182	149.666	147.703
	10	15.149	10.456	9.348	10.448	13.195	9.133

Type	UC ID	T89	T90	T91	T92	T93	T94
Deep Learning	2	15.953	17.721	13.882	20.199	14.549	18.446
	5	4.309	8.992	7.656	6.124	5.409	10.098
	9	74.229	65.816	62.259	58.374	63.640	71.648
Machine Learning	1	8.009	7.038	8.104	7.809	5.202	4.707
	3	7.656	5.913	7.300	6.078	10.775	10.110
	4	12.704	5.476	9.976	7.506	10.513	6.378
	6	3.171	2.649	2.960	3.031	2.628	2.586
	7	3.514	3.128	2.602	3.525	2.340	3.624
	8	141.901	157.342	143.678	152.585	151.661	148.155
	10	11.188	13.636	13.090	11.605	7.978	10.878

Type	UC ID	T95	T96	T97	T98	T99	T100
Deep Learning	2	15.718	12.756	20.179	19.298	14.175	6.894
	5	9.060	6.607	7.305	8.049	11.545	4.951
	9	57.693	67.445	54.701	59.231	66.021	75.918
Machine Learning	1	5.972	5.221	8.256	5.428	7.166	6.839
	3	8.799	6.996	8.276	7.396	8.679	9.727
	4	10.855	10.610	10.145	11.590	10.148	15.122
	6	3.385	1.958	3.174	3.185	2.953	3.223
	7	3.718	3.233	3.101	2.374	3.124	5.300
	8	138.195	156.852	151.058	146.641	148.469	151.626
	10	11.610	7.849	10.691	9.323	9.673	9.889

Table 3-2 Use Case Elapsed Times

3.4 SUT Validation Test Output

<u>Validation Run Report</u>			
AIUCpm@1	438.29	T_{Load}	0.27
Scale Factor	1	T_{LD}	0.27
Streams	100	T_{PTT}	19.65
		T_{PST1}	2.67
Kit Version	2.0.0	T_{PST2}	2.67
Execution Status	Pass	T_{PST}	2.67
Accuracy Status	Pass	T_{TT}	0.25
Test Times			
Overall Run Start Time		2026-01-15 18:54:24.480	
Overall Run End Time		2026-01-15 19:17:46.103	
Overall Run Elapsed Time		1,401.623	
Load Test Start Time		2026-01-15 18:55:36.277	
Load Test End Time		2026-01-15 18:55:36.556	
Load Test Elapsed Time		0.279	
Power Training Start Time		2026-01-15 18:55:36.557	
Power Training End Time		2026-01-15 19:09:35.216	
Power Training Elapsed Time		838.659	
Power Serving 1 Start Time		2026-01-15 19:09:35.217	
Power Serving 1 End Time		2026-01-15 19:10:36.823	
Power Serving 1 Elapsed Time		61.606	
Power Serving 2 Start Time		2026-01-15 19:10:36.824	
Power Serving 2 End Time		2026-01-15 19:11:38.055	
Power Serving 2 Elapsed Time		61.231	
Scoring Start Time		2026-01-15 19:12:12.780	
Scoring End Time		2026-01-15 19:13:32.552	
Scoring Elapsed Time		79.772	
Throughput Start Time		2026-01-15 19:13:32.564	
Throughput End Time		2026-01-15 19:17:46.100	
Throughput Elapsed Time		253.536	

(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	0.272	<=	0.50	Pass
3	mean_squared_log_error	4.582	<=	5.40	Pass
4	f1_score	0.701	>=	0.65	Pass
5	mean_squared_log_error	0.013	<=	0.50	Pass
6	matthews_corrcoef	0.462	>=	0.19	Pass
7	median_absolute_error	0.893	<=	1.80	Pass
8	accuracy_score	0.715	>=	0.65	Pass
9	accuracy_score	1.000	>=	0.90	Pass
10	accuracy_score	0.817	>=	0.70	Pass

3.5 Configuration Parameters

The [Supporting Files](#) archive contains all Global Benchmark Parameter and Use Case Specific Parameter settings.

Clause 4 – SUT Related Items

4.1 Specialized Hardware/Software

No Specialized Hardware/Software was used in the SUT.

4.2 Configuration Files

The [Supporting Files](#) archive contains all configuration files.

4.3 SUT Environment Information

All envInfo.log files are included in the [Supporting Files](#) archive.

4.4 Data Storage to Scale Factor Ratio

The details of the Data Storage Ratio are provided below.

Node Count	Disks	Size (GB)	Total (GB)
1	2	960	1,920
Total Storage (GB)			1,920
Scale Factor			3
Data Storage Ratio			640.00

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	384	384
Scale Factor		3
Total Memory (GiB)		384
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

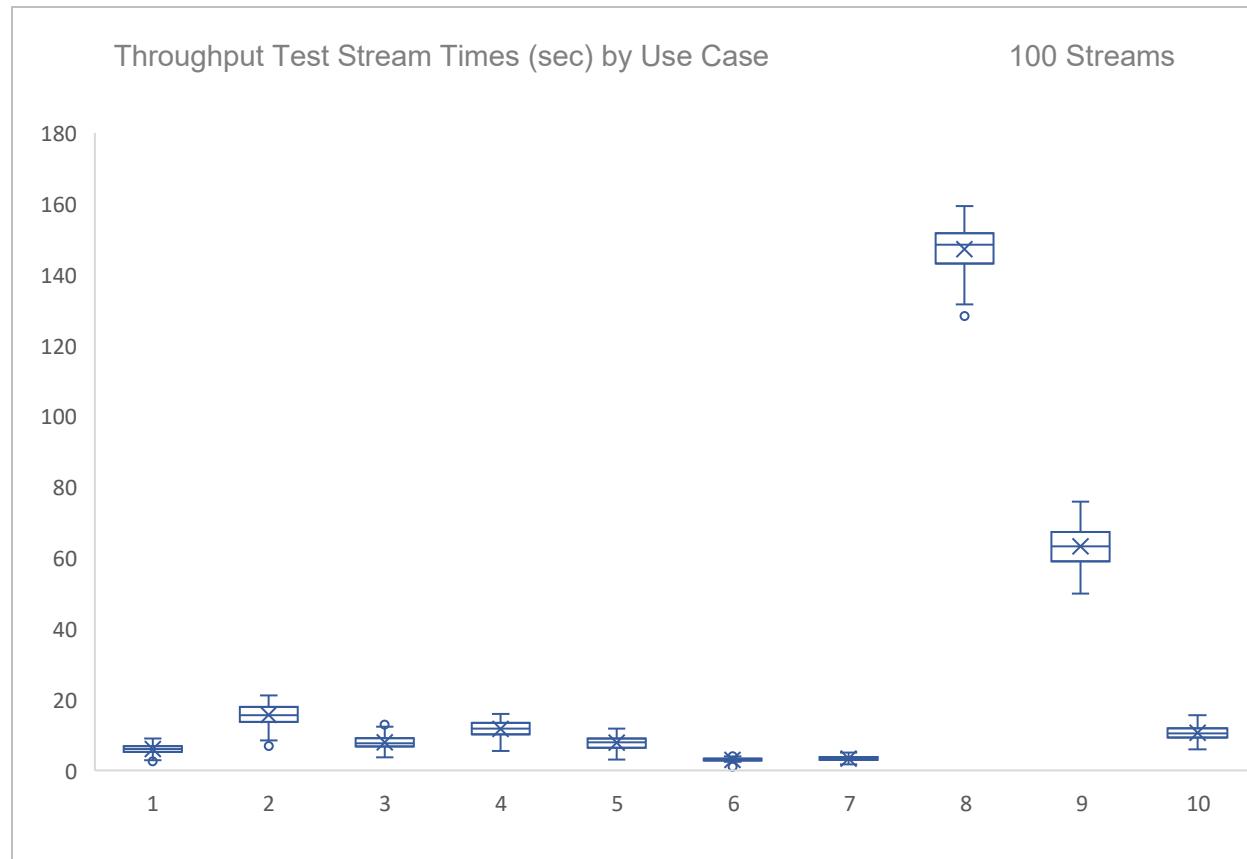
<u>Metric Overview</u>			
TPCx-AI Performance Metric	893.92	AIUCpm@3	
TPCx-AI Price/Performance Metric	135.11	\$/AIUCpm@3	
TPCx-AI Scale Factor	3		
TPCx-AI Stream Count	100		
<u>Test Times</u>			
Overall Run Start Time	2026-01-15 19:18:04.839		
Overall Run End Time	2026-01-15 19:51:49.712		
Overall Run Elapsed Time	2,024.873		
Load Test Start Time	2026-01-15 19:19:26.062		
Load Test End Time	2026-01-15 19:19:26.574		
Load Test Elapsed Time	0.512		
Power Training Start Time	2026-01-15 19:19:26.575		
Power Training End Time	2026-01-15 19:41:21.674		
Power Training Elapsed Time	1,315.099		
Power Serving 1 Start Time	2026-01-15 19:41:21.676		
Power Serving 1 End Time	2026-01-15 19:43:10.300		
Power Serving 1 Elapsed Time	108.624		
Power Serving 2 Start Time	2026-01-15 19:43:10.302		
Power Serving 2 End Time	2026-01-15 19:44:58.702		
Power Serving 2 Elapsed Time	108.400		
Scoring Start Time	2026-01-15 19:45:34.333		
Scoring End Time	2026-01-15 19:46:56.190		
Scoring Elapsed Time	81.857		
Throughput Start Time	2026-01-15 19:46:56.203		
Throughput End Time	2026-01-15 19:51:49.709		
Throughput Elapsed Time	293.506		

Accuracy Metrics

Use Case	Metric Name	Metric	Criteria	Threshold	Status
1	N/A	0.000	N/A	0.00	Pass
2	word_error_rate	0.465	<=	0.50	Pass
3	mean_squared_log_error	3.507	<=	5.40	Pass
4	f1_score	0.700	>=	0.65	Pass
5	mean_squared_log_error	0.451	<=	0.50	Pass
6	matthews_corrcoef	0.492	>=	0.19	Pass
7	median_absolute_error	0.955	<=	1.80	Pass
8	accuracy_score	0.751	>=	0.65	Pass
9	accuracy_score	1.000	>=	0.90	Pass
10	accuracy_score	0.816	>=	0.70	Pass

5.2 Throughput Test Stream Times

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



Auditor's Information

This benchmark was audited by Doug Johnson, InfoSizing.

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

This benchmark's Full Disclosure Report can be downloaded from [www.tpc.org](http://www(tpc.org).

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



The Right Metric For Sizing IT



Danny Zhang
Performance Development Manager
7001 Development Drive
Morrisville, NC 27560

January 28, 2026

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

Platform: 1x ThinkSystem SR655 V3
Operating System: Red Hat Enterprise Linux 8.10 (Ootpa)
Additional Software: Anaconda Business

The results were:

Performance Metric 893.92 AIUCpm@3

Secondary Metrics	T_{LD}	0.50
	T_{PTT}	29.33
	T_{PST}	3.85
	T_{TT}	0.29

System Under Test 1x ThinkSystem SR655 V3 with:

CPUs	1x AMD EPYC 9575F 64-Core Processor		
Memory	384 GiB		
Storage	Qty	Size	Type
	2	960 GB	M.2 VA Read Intensive 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 v2.0.0.
- All checksums were validated for compliance.
- Any modifications to shell scripts were reviewed for compliance.
- No modifications were made to any of the Java code.
- The generated dataset was properly scaled to 3 GB.
- The generated dataset used for testing was protected by RAID1.

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,



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: Ray Engler, Lenovo and TPC Consortium | rengler@lenovo.com
 From: Anaconda Support | support@anaconda.com
 Re: Anaconda Support Quote for Publishing TPCx-AI Performance Benchmark Testing

Effective Date: January 13, 2026

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 remain valid for **120 days** from the effective date above.

Premium Support

Included Services	
Anaconda Business, 100 users, based on this configuration*	\$60,000
Anaconda Business Gold, including P1 24x7 support	\$30,000
Includes acknowledgment within 4 hours of issues being reported. Once resolution is identified, it will be provided to the customer within 4 hours.	
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>

Lenovo's TPCx-AI Benchmark Software Configuration

- Anaconda Business 100 Subscription (100 users x 12 months x \$50)*
- 1 server = Lenovo ThinkSystem SR655 V3
- 1 processor = AMD EPYC 9575F 64C 400W 3.3GHz Processor
 - 12 x 32GB DIMMs
 - 2 SSDs, ThinkSystem M.2 VA 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD (configured RAID-1)
- 1 ThinkSystem Broadcom 5719 1GbE RJ45 4-port OCP Ethernet Adapter
- RHEL

* 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

© 2026 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 Lenovo

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