Hewlett Packard Enterprise

TPC Express Benchmark™ AI
Full Disclosure Report

ProLiant DL380 Gen11

with 8x ProLiant DL380 Gen11; 2x ProLiant DL380 Gen10 Plus
using

Cloudera 7.1.9

running on

Red Hat Enterprise Linux 8.9 (Ootpa)

TPCx-AI Version	2.0.0
Report Edition	First
Report Submitted	March 7, 2025

First Edition - March 2025

Hewlett Packard Enterprise (HPE)*, 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.

HPE and the HPE Logo are trademarks of Hewlett Packard Enterprise 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 HPE and any other company.

TPC Express Benchmark™ AI, TPCx-AI, and AIUCpm@100, are registered certification marks of the Transaction Processing Performance Council.

The HPE 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 HPE business contact for information on the products or services available in your area. You can find additional information via HPE's web site at www.hpe.com. Actual performance and environmental costs of HPE products will vary depending on individual customer configurations and conditions.

Abstract

HPE conducted the TPC Express Benchmark™ AI (TPCx-AI) on the ProLiant DL380 Gen11. The software used included Cloudera 7.1.9. 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
HPE	8x ProLiant DL380 Gen11 (Compute) 2x ProLiant DL380 Gen10 Plus (Controller)	Red Hat Enterprise Linux 8.9 (Ootpa)

Metrics Overview

Total System Cost	Performance	Price/Performance	Availability Date
$867,492 USD	1,600.16 AIUCpm@100	542.13 USD $/AIUCpm@100	March 7, 2025

Executive Summary

The Executive Summary follows on the next several pages.

Hewlett Packard Enterprise	ProLiant DL380 Gen11	TPCx-AI 2.0.0 TPC Pricing 2.9.0 Report Date Mar. 07, 2025

TPCx-AI Performance	Total System Cost	Price/Performance	Availability Date
1,600.16 AIUCpm@100	**$867,492 USD**	**$542.13 USD/AIUCpm@100**	**March 7, 2025**

Framework	Operating System	Other Software	Scale Factor	Streams
Cloudera 7.1.9	Red Hat Enterprise Linux 8.9 (Ootpa)	N/A	100	54

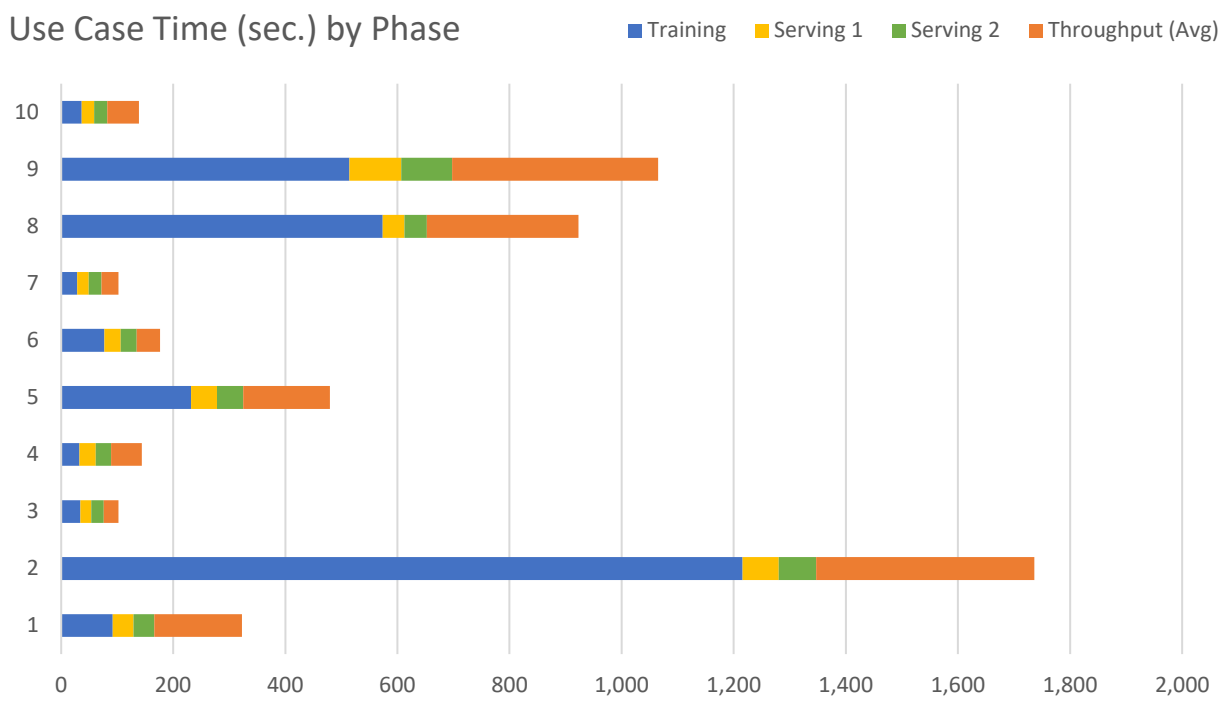

Physical Storage / Scale Factor	Scale Factor / Physical Memory	Main Data Redundancy Model
879.20	0.01	Replication 3

Servers: Total Processors/Cores/Threads	10 20 / 624 / 1,248

Server Type	8x ProLiant DL380 Gen11 (Compute)	2x ProLiant DL380 Gen10 Plus (Controller)
Processors	2x Intel(R) Xeon(R) Platinum 8562Y+ CPU @ 2.80GHz	2x Intel(R) Xeon(R) Gold 6330 CPU @ 2.00GHz
Memory	1,024 GiB	512 GiB
Storage Controller	1x NS204i-u Gen11 (7 nodes)	1x MR416i-p Gen10+ or P408i-a SR Gen10
Storage Device	2x 480 GB NVMe SSD (7 nodes) 3x 3.2 TB NVMe SSD (7 nodes) 4x 3.2 TB NVMe SSD (1 node)	2x 300 GB 15K SAS
Network Controller	1x InfiniBand NDR200/Ethernet 200GbE 2-port	1x MCX623106AS CDAT Ethernet 100Gb 2-port
Connectivity	1x CX 8360-12C v2 12-port 100G (Compute); 1x CX 5420 24p 10M/100M/1G (Manage)	

Hewlett Packard Enterprise

ProLiant DL380 Gen11

TPCx-AI	2.0.0
TPC Pricing	2.9.0
Report Date	Mar. 07, 2025

Description	Part Number	Source	List Price	Qty	Extended Price	1-Yr. Maintenance
Server Hardware						
HPE ProLiant DL380 Gen11 8SFF CPU1/2 NVMe Cable Kit	P48825-B21	1	$2,127.00	8	$17,016.00	
Intel Xeon-Platinum 8562Y+ 2.8GHz 32-core 300W Processor for HPE	P67085-B21	1	$11,296.00	16	$180,736.00	
HPE 64GB (1x64GB) Dual Rank x4 DDR5-5600 CAS-46-45-45 EC8 Registered Smart Memory Kit	P43331-B21	1	$4,960.00	128	$634,880.00	
HPE NS204i-u Gen11 NVMe Hot Plug Boot Optimized Storage Device	P48183-B21	1	$1,994.00	7	$13,958.00	
HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 CM7 SSD	P63849-B21	1	$6,555.00	25	$163,875.00	
HPE InfiniBand NDR200/Ethernet 200GbE 2-port QSFP112 PCIe5 x16 MCX755106AC-HEAT Adapter	P65333-B21	1S	$7,331.00	8	$58,648.00	
Intel E810 2-port for HPE	P08443-B21	1	$756.00	8	$6,048.00	
HPE ProLiant DL360 Gen11 CPU1 to OCP2 x8 Enablement Kit	P51911-B21	1	$120.00	8	$960.00	
					$0.00	
HPE ProLiant DL380 Gen10 Plus 2U 8SFF x1 Tri-Mode 24G U.3 BC Front Drive Cage Kit	P27194-B21	1	$737.00	2	$1,474.00	
Intel Xeon-Gold 6330 2.0GHz 28-core 205W Processor for HPE	P36927-B21	1	$3,758.00	4	$15,032.00	
HPE 32GB (1x32GB) Dual Rank x4 DDR4-3200 CAS-22-22-22 Registered Smart Memory Kit	P06033-B21	1	$1,524.00	32	$48,768.00	
HPE 300GB 15K 12G SFF SAS SC	870753-B21	1	$231.00	4	$924.00	
HPE ProLiant DL380 Gen10 Plus x8/x16/x8 Primary FIO Riser Kit	P37038-B21	1	$215.00	2	$430.00	
HPE MR416i-a Gen10 Plus x16 Lanes 4GB Cache NVMe/SAS 12G Controller	P26279-B21	1	$2,299.00	1	$2,299.00	
HPE Smart Array P408i-a SR Gen10 12G SAS Controller	804331-B21	1	$1,128.00	1	$1,128.00	
Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T OCP3 Adapter for HPE	P51181-B21	1	$485.00	2	$970.00	
HPE MCX623106AS CDAT Ethernet 100Gb 2 port QSFP56 Adapter	P25960-B21	1	$4,373.00	2	$8,746.00	
Network Hardware						
HPE Aruba Networking CX 8360-12C v2 12-port 100G QSFP+/QSFP28	JL709C	1	$32,735.00	1	$32,735.00	
HPE Aruba Networking 100G QSFP28 to QSFP28 3m Direct Attach Copper Cable	JL307A	1	$734.00	10	$7,340.00	
HPE Aruba Networking CX 5420 24p 10M/100M/1G Class4 PoE Module - switch	S0U62A	1	$6,149.00	1	$6,149.00	
HPE 1.2m/4ft CAT5 RJ45 M/M Ethernet Cable	C7533A	1	16	10	$160.00	
HPE 3 Year Tech Care Essential wDMR DL380a Gen11 HW Service	H38YLE	1	22626	1		$22,626.00
HPE 42U 600mmx1200mm G2 Enterprise Shock Rack	P9K40A	1	$6,200.00	1	$6,200.00	
				Subtotal	**$1,208,476.00**	**$22,626.00**
Other						
HPE USB US Keyboard/Mouse Kit	631341-B21	1	24.99	3	$74.97	
HP M22f FHD Monitor	2D9J9AA#ABA	2	144.99	3	$434.97	
				Subtotal	**$509.94**	**$0.00**
Software						
RHEL Svr Sckt/2 Gst 1yr 24x7 E-LTU	J8J36A	1	$1,214.00	10	$12,140.00	
CDP Private Cloud Base Edition - Business Select	CDP-PVC-BASE-SLT	1	$5,462.50	10	$54,625.00	
				Software Subtotal	**$66,765.00**	**$0.00**
				Total Extended Price	**$1,275,750.94**	**$22,626.00**
				Total Discounts	**$422,966.00**	**$7,919.00**

Pricing: 1 = HPE; 2 = HP.com

* Discount applies to Server/Network line items where Key = 1. Discount based upon total system cost as purchased by a regular customer.

One or more components of the Measured Configuration have been substituted in the Priced Configuration. See the FDR for details.

Audited by Doug Johnson, InfoSizing

Total System Cost (USD):	**$867,492**
AIUCpm@100:	**1,600.16**
$/AIUCpm@100:	**$542.13**

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.

Hewlett Packard Enterprise	ProLiant DL380 Gen11	TPCx-AI	2.0.0
		TPC Pricing	2.9.0
		Report Date	Mar. 07, 2025

Numerical Quantities

AIUCpm@100	**1,600.16**	T_{Load}	141.69
Scale Factor	100	T_{LD}	141.69
Streams	54	T_{PTT}	121.33
		T_{PST1}	35.39
Kit Version	2.0.0	T_{PST2}	36.20
Execution Status	Pass	T_{PST}	36.20
Accuracy Status	Pass	T_{TT}	3.18

Test Times

Overall Run Start Time	2025-01-29 06:57:55.787
Overall Run End Time	2025-01-29 08:47:16.266
Overall Run Elapsed Time	6,560.479
Load Test Start Time	2025-01-29 03:04:52.203
Load Test End Time	2025-01-29 03:07:16.791
Load Test Elapsed Time	144.588
Power Training Start Time	2025-01-29 03:07:16.859
Power Training End Time	2025-01-29 03:54:49.762
Power Training Elapsed Time	2,852.903
Power Serving 1 Start Time	2025-01-29 03:54:49.848
Power Serving 1 End Time	2025-01-29 04:01:45.492
Power Serving 1 Elapsed Time	415.644
Power Serving 2 Start Time	2025-01-29 04:01:45.570
Power Serving 2 End Time	2025-01-29 04:08:46.560
Power Serving 2 Elapsed Time	420.990
Scoring Start Time	2025-01-29 04:12:02.561
Scoring End Time	2025-01-29 04:18:04.459
Scoring Elapsed Time	361.898
Throughput Start Time	2025-01-29 04:18:04.586
Throughput End Time	2025-01-29 04:47:16.263
Throughput Elapsed Time	1,751.677

Hewlett Packard Enterprise	ProLiant DL380 Gen11	TPCx-AI	2.0.0
		TPC Pricing	2.9.0
		Report Date	Mar. 07, 2025

Numerical Quantities (continued)

Use Case Times & Accuracy

Use Case	Training (sec)	Serving 1 (sec)	Serving 2 (sec)	Throughput (avg)	Accuracy
UC01	92.237	36.724	37.077	156.562	0.000
UC02	1,215.930	64.606	66.461	389.176	0.457
UC03	34.131	19.842	21.984	26.136	3.885
UC04	32.885	29.183	27.259	54.669	0.694
UC05	231.837	46.398	46.763	154.738	0.084
UC06	77.189	28.865	28.939	41.268	0.204
UC07	28.484	20.680	22.729	30.157	1.307
UC08	573.633	38.795	39.899	270.592	0.757
UC09	514.635	91.932	91.131	367.508	0.990
UC10	36.497	22.678	23.421	56.341	0.816

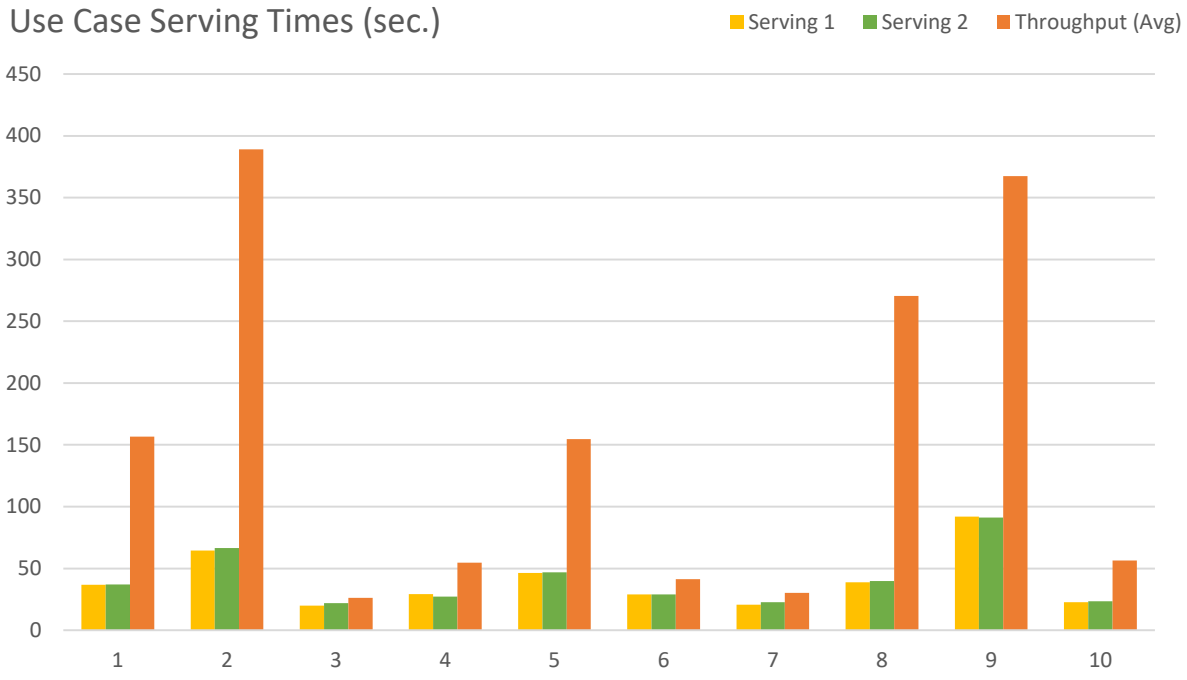

Table of Contents

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 Hewlett Packard Enterprise.

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:	10		
Processors/Cores/Threads:	20/624/1,248	Storage Devices:	43
Total Memory:	9,216 GiB	Storage Capacity:	87,920 GB

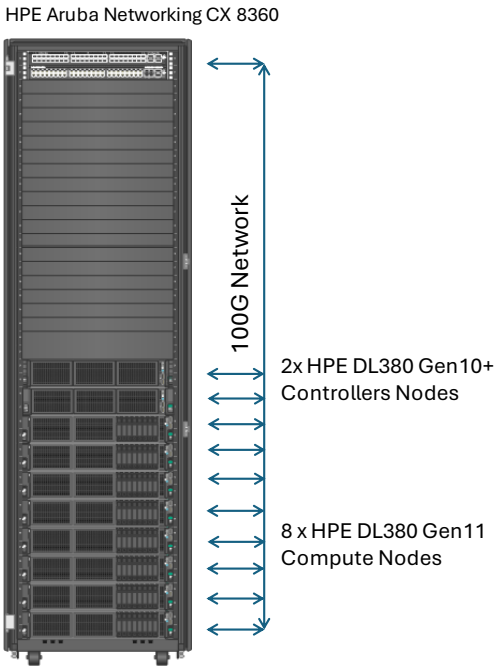

	Compute	Controller
Server	8x ProLiant DL380 Gen11	2x ProLiant DL380 Gen10 Plus
Procs/Cores/Threads:	2/32/64	2/28/56
Processor Model:	2x Intel(R) Xeon(R) Platinum 8562Y+ CPU @ 2.80GHz	2x Intel(R) Xeon(R) Gold 6330 CPU @ 2.00GHz
Memory:	1,024 GiB	512 GiB
Storage Controller:	1x NS204i-u Gen11	1x MR416i-p Gen10+ or P408i-a SR Gen10
Storage Devices:	2x 480 GB NVMe SSD (7 nodes) 3x 3.2 TB NVMe SSD (7 nodes) 4x 3.2 TB NVMe SSD (1 node)	2x 300 GB 15K SAS
Network Controller:	1x InfiniBand NDR200/Ethernet 200GbE 2-port	1x MCX623106AS CDAT Ethernet 100Gb 2-port
Network:	1x CX 8360-12C v2 12-port 100G (Compute); 1x CX 5420 24p 10M/100M/1G (Manage)	

The distribution of software components over server nodes is detailed in Clause 2.

1.3.2 Differences Between the Measured and the Priced Configurations

The Measured Configuration had four Nvidia Mellanox Technologies MT28908 Family [ConnectX-6] 100 Gb 2-port adapters that were substituted with four HPE InfiniBand NDR200/Ethernet 200GbE 2-port QSFP112 PCIe5 x16 MCX755106AC-HEAT adapters.

Additionally, the Measured Configuration contained 90 extra storage devices, as follows.

Server	Extra Storage Devices
tpcxai-01.rdma.ai.hpe.local	3x 300 GB 15K SAS
tpcxai-02.rdma.ai.hpe.local	3x 300 GB 15K SAS
tpcxai-03.rdma.ai.hpe.local	11x 3.2 TB NVMe SSD
tpcxai-04.rdma.ai.hpe.local	11x 3.2 TB NVMe SSD
tpcxai-05.rdma.ai.hpe.local	11x 3.2 TB NVMe SSD
tpcxai-06.rdma.ai.hpe.local	11x 3.2 TB NVMe SSD
tpcxai-07.rdma.ai.hpe.local	10x 3.2 TB NVMe SSD
tpcxai-08.rdma.ai.hpe.local	10x 3.2 TB NVMe SSD
tpcxai-09.rdma.ai.hpe.local	10x 3.2 TB NVMe SSD
tpcxai-10.rdma.ai.hpe.local	10x 3.2 TB NVMe SSD

These 90 extra devices were not utilized during testing and are not included in 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
ProLiant DL380 Gen10 Plus	tpcxai-01.rdma.ai.hpe.local	Balancer, Gateway, NameNode, Activity Processor, Event Server, Host Monitor, ResourceManager, Service Monitor, YARN Application Timeline Server, YARN Application Timeline Server, History Server, JobHistory Server, Sentry Server	2x 300 GB 15K SAS	OS, Kit
ProLiant DL380 Gen10 Plus	tpcxai-02.rdma.ai.hpe.local	Gateway, Secondary NameNode, Gateway	2x 300 GB 15K SAS	OS, Kit
8x ProLiant DL380 Gen11	tpcxai-03.rdma.ai.hpe.local - tpcxai-10.rdma.ai.hpe.local	DataNode, Gateway, NodeManager	2x 480 GB NVMe SSD 3x 3.2 TB NVMe SSD	OS, Kit Models, Data

Table 2-1 Software Components and Dataset Distribution

2.2 File System Implementation

A distributed file system provided by Red Hat Enterprise Linux 8.9 (Ootpa) / Cloudera 7.1.9 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

Cloudera 7.1.9 consisted of the following components.

Component	Version
HDFS	3.1.1
YARN	3.1.1
MapReduce2	3.1.1
Spark	2.4.8

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	2.0.0
Modified File	Description of Changes
Tools/parallel-data-load.sh	group commands into a single background job, mkdir completes before proceeding, reduce SSH connections per file
Tools/spark/load_hdfs.sh	add “&” to make it run as the background process

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	64.606	66.461	412.247	376.886	389.264	379.339
	5	46.398	46.763	139.941	163.515	146.771	169.244
	9	91.932	91.131	317.247	306.776	444.323	379.375
Machine Learning	1	36.724	37.077	155.232	150.149	152.493	131.940
	3	19.842	21.984	29.674	27.483	16.890	26.411
	4	29.183	27.259	64.289	62.045	53.606	47.883
	6	28.865	28.939	43.034	45.306	45.070	42.625
	7	20.680	22.729	26.972	29.408	24.861	26.197
	8	38.795	39.899	279.659	282.877	283.406	293.251
	10	22.678	23.421	65.574	53.526	45.619	58.935

Type	UC ID	T5	T6	T7	T8	T9	T10
Deep Learning	2	383.068	396.439	362.441	393.930	409.158	367.186
	5	155.923	143.025	136.500	153.907	131.214	151.004
	9	365.762	391.534	380.536	365.136	289.193	445.518
Machine Learning	1	150.015	159.365	146.122	150.379	157.770	151.441
	3	24.672	25.571	27.296	35.812	28.072	25.125
	4	46.404	67.049	51.968	53.204	51.067	59.720
	6	39.526	34.405	39.067	44.524	41.082	36.117
	7	31.287	31.005	31.916	28.214	26.212	25.051
	8	271.968	283.083	284.908	264.379	237.962	298.474
	10	50.304	57.369	57.042	65.848	63.710	65.914

Type	UC ID	T11	T12	T13	T14	T15	T16
Deep Learning	2	380.193	429.156	372.722	392.241	366.448	348.383
	5	168.968	95.057	157.815	150.274	133.070	212.005
	9	389.086	465.320	333.317	331.702	341.671	361.032
Machine Learning	1	159.450	196.938	183.077	154.976	152.067	180.973
	3	21.995	22.646	26.914	24.275	32.471	26.694
	4	56.937	45.438	54.854	59.263	50.238	54.639
	6	37.916	69.437	41.367	51.456	39.944	45.447
	7	28.676	24.100	29.043	25.045	33.826	38.318
	8	263.737	301.397	251.227	237.129	340.712	303.953
	10	48.744	44.912	52.869	51.002	45.721	48.217

Type	UC ID	T17	T18	T19	T20	T21	T22
Deep Learning	2	382.049	366.154	395.476	400.080	389.581	363.255
	5	135.608	162.750	154.570	182.157	162.688	138.691
	9	346.629	361.152	477.976	360.575	367.982	328.268
Machine Learning	1	149.865	156.175	166.788	150.966	142.940	138.232
	3	25.086	25.591	24.913	22.214	24.091	21.731
	4	61.291	45.572	79.783	63.432	51.769	62.693
	6	44.383	36.309	40.897	38.372	44.078	38.307
	7	26.491	26.906	29.325	25.460	26.163	29.206
	8	259.012	284.117	271.216	227.845	305.565	257.834
	10	60.045	88.105	51.777	53.603	58.383	59.422

Type	UC ID	T23	T24	T25	T26	T27	T28
Deep Learning	2	419.269	398.258	395.719	395.761	409.104	399.551
	5	142.786	149.615	136.030	152.133	141.536	159.878
	9	370.086	358.976	339.693	387.276	386.079	428.814
Machine Learning	1	131.132	138.467	149.088	133.976	145.663	150.599
	3	26.981	21.313	25.470	31.653	25.568	24.163
	4	62.619	59.251	65.514	53.711	61.933	46.638
	6	44.169	41.858	38.171	37.336	40.769	32.919
	7	30.611	37.030	39.234	33.142	34.004	29.408
	8	229.917	258.473	272.992	257.962	259.495	262.584
	10	49.144	50.031	60.974	53.413	65.994	55.821

Type	UC ID	T29	T30	T31	T32	T33	T34
Deep Learning	2	403.961	411.095	369.512	396.722	390.517	402.198
	5	189.906	159.396	159.281	150.219	205.030	165.742
	9	370.075	370.283	477.813	318.798	320.452	379.656
Machine Learning	1	120.404	168.587	135.135	127.126	145.404	150.477
	3	20.411	24.842	23.462	27.205	40.427	21.595
	4	49.360	40.484	46.914	55.305	47.251	51.945
	6	33.459	39.991	35.944	41.408	44.315	42.896
	7	49.276	27.565	27.035	24.838	27.546	29.526
	8	285.719	261.184	257.491	281.510	273.974	276.731
	10	55.100	54.234	57.465	53.622	55.165	58.405

Type	UC ID	T35	T36	T37	T38	T39	T40
Deep Learning	2	322.737	389.780	424.718	366.510	370.917	387.170
	5	166.043	173.741	141.619	166.269	136.550	157.572
	9	355.464	334.157	362.937	409.576	395.945	370.089
Machine Learning	1	158.066	179.271	173.009	130.703	180.294	141.804
	3	31.454	24.319	23.458	24.975	25.040	23.085
	4	48.148	55.320	56.796	41.427	65.893	58.285
	6	45.279	39.597	35.285	37.368	43.156	48.818
	7	29.711	37.512	29.026	32.261	30.889	31.266
	8	236.775	283.957	264.565	295.939	250.087	254.356
	10	56.705	49.655	53.281	63.981	50.061	56.511

Type	UC ID	T41	T42	T43	T44	T45	T46
Deep Learning	2	410.000	376.927	440.389	407.704	402.042	381.821
	5	125.508	149.575	106.070	150.004	138.216	164.932
	9	340.537	417.245	352.807	372.857	370.302	311.913
Machine Learning	1	173.763	168.382	201.305	176.330	154.287	134.019
	3	22.865	19.579	29.311	27.047	28.200	25.797
	4	63.081	39.798	51.568	59.760	47.754	62.186
	6	34.575	43.683	35.513	37.576	37.213	42.812
	7	26.530	23.515	36.783	30.587	38.127	32.339
	8	289.186	276.527	293.713	270.950	251.021	243.025
	10	49.311	56.695	57.124	55.155	52.624	63.123

Type	UC ID	T47	T48	T49	T50	T51	T52
Deep Learning	2	398.034	371.782	390.425	382.869	403.362	390.987
	5	166.616	151.567	175.485	170.119	145.653	168.927
	9	383.690	324.062	341.078	311.113	368.122	364.147
Machine Learning	1	137.132	182.047	137.837	168.131	203.057	170.557
	3	25.018	23.279	24.450	29.317	28.653	41.155
	4	52.876	53.329	47.958	50.441	49.805	45.847
	6	40.782	38.381	40.929	37.536	55.119	38.471
	7	29.399	33.334	29.433	27.311	27.433	32.955
	8	271.585	312.794	308.190	275.601	256.368	259.658
	10	68.902	49.907	51.337	59.954	57.422	52.763

Type	UC ID	T53	T54
Deep Learning	2	388.148	361.841
	5	179.109	166.037
	9	349.659	351.614
Machine Learning	1	182.241	168.718
	3	32.760	22.874
	4	50.197	67.566
	6	41.715	42.739
	7	27.202	29.956
	8	222.379	233.543
	10	60.917	60.960

Table 3-2 Use Case Elapsed Times

3.4 SUT Validation Test Output

Validation Run Report

AIUCpm@1	25.89	T_{Load}	107.01
Scale Factor	1	T_{LD}	107.01
Streams	54	T_{PTT}	83.56
		T_{PST1}	30.60
Kit Version	2.0.0	T_{PST2}	30.76
Execution Status	Pass	T_{PST}	30.76
Accuracy Status	Pass	T_{TT}	1.05

Test Times

Overall Run Start Time	2025-01-29 05:53:05.970
Overall Run End Time	2025-01-29 06:56:57.381
Overall Run Elapsed Time	3,831.411
Load Test Start Time	2025-01-29 01:56:57.368
Load Test End Time	2025-01-29 01:58:47.266
Load Test Elapsed Time	109.898
Power Training Start Time	2025-01-29 01:58:47.343
Power Training End Time	2025-01-29 02:24:47.244
Power Training Elapsed Time	1,559.901
Power Serving 1 Start Time	2025-01-29 02:24:47.325
Power Serving 1 End Time	2025-01-29 02:30:51.757
Power Serving 1 Elapsed Time	364.432
Power Serving 2 Start Time	2025-01-29 02:30:51.839
Power Serving 2 End Time	2025-01-29 02:36:57.078
Power Serving 2 Elapsed Time	365.239
Scoring Start Time	2025-01-29 02:40:12.060
Scoring End Time	2025-01-29 02:46:57.983
Scoring Elapsed Time	405.923
Throughput Start Time	2025-01-29 02:46:58.098
Throughput End Time	2025-01-29 02:56:57.378
Throughput Elapsed Time	599.280

(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	0.461	<=	0.50	Pass
3	mean_squared_log_error	5.778	<=	5.40	Fail*
4	f1_score	0.697	>=	0.65	Pass
5	mean_squared_log_error	0.314	<=	0.50	Pass
6	matthews_corrcoef	0.223	>=	0.19	Pass
7	median_absolute_error	1.715	<=	1.80	Pass
8	accuracy_score	0.696	>=	0.65	Pass
9	accuracy_score	1.000	>=	0.90	Pass
10	accuracy_score	0.817	>=	0.70	Pass

*Because of the small dataset size used for the Validation Test, Spark-based implementations may not be able to satisfy the accuracy threshold for Use Case 3. The TPCx-AI Subcommittee is aware of this issue and has decided that this failure does not invalidate the test.

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)
2	?	300	1,200
7	2	480	6,720
7	3	3,200	67,200
1	4	3,200	12,800

Total Storage (GB)	87,920
Scale Factor	100
Data Storage Ratio	879.20

4.5 Scale Factor to Memory Ratio

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

Nodes	Memory (GiB)	Total (GiB)
2	512	1,024
8	1,024	8,192

Scale Factor	100
Total Memory (GiB)	9,216
SF / Memory Ratio	0.01

4.6 Output of Tests

The Supporting Files archive contains the output files of all tests.

4.7 Additional Sponsor Files

The Supporting Files archive contains any additional files that were used.

4.8 Model Optimizations

The Supporting Files archive contains any model optimization files that were used.

Clause 5 – Metrics and Scale Factor

5.1 Reported Performance Metrics

Metric Overview

TPCx-AI Performance Metric	1,600.16	AIUCpm@100
TPCx-AI Price/Performance Metric	542.13	$/AIUCpm@100
TPCx-AI Scale Factor	100	
TPCx-AI Stream Count	54	

Test Times

Overall Run Start Time	2025-01-29 06:57:55.787
Overall Run End Time	2025-01-29 08:47:16.266
Overall Run Elapsed Time	6,560.479
Load Test Start Time	2025-01-29 03:04:52.203
Load Test End Time	2025-01-29 03:07:16.791
Load Test Elapsed Time	144.588
Power Training Start Time	2025-01-29 03:07:16.859
Power Training End Time	2025-01-29 03:54:49.762
Power Training Elapsed Time	2,852.903
Power Serving 1 Start Time	2025-01-29 03:54:49.848
Power Serving 1 End Time	2025-01-29 04:01:45.492
Power Serving 1 Elapsed Time	415.644
Power Serving 2 Start Time	2025-01-29 04:01:45.570
Power Serving 2 End Time	2025-01-29 04:08:46.560
Power Serving 2 Elapsed Time	420.990
Scoring Start Time	2025-01-29 04:12:02.561
Scoring End Time	2025-01-29 04:18:04.459
Scoring Elapsed Time	361.898
Throughput Start Time	2025-01-29 04:18:04.586
Throughput End Time	2025-01-29 04:47:16.263
Throughput Elapsed Time	1,751.677

Accuracy Metrics

Use Case	Metric Name	Metric	Criteria	Threshold	Status
1	N/A	0.000	N/A	0.00	Pass
2	word_error_rate	0.457	<=	0.50	Pass
3	mean_squared_log_error	3.885	<=	5.40	Pass
4	f1_score	0.694	>=	0.65	Pass
5	mean_squared_log_error	0.084	<=	0.50	Pass
6	matthews_corrcoef	0.204	>=	0.19	Pass
7	median_absolute_error	1.307	<=	1.80	Pass
8	accuracy_score	0.757	>=	0.65	Pass
9	accuracy_score	0.990	>=	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”.

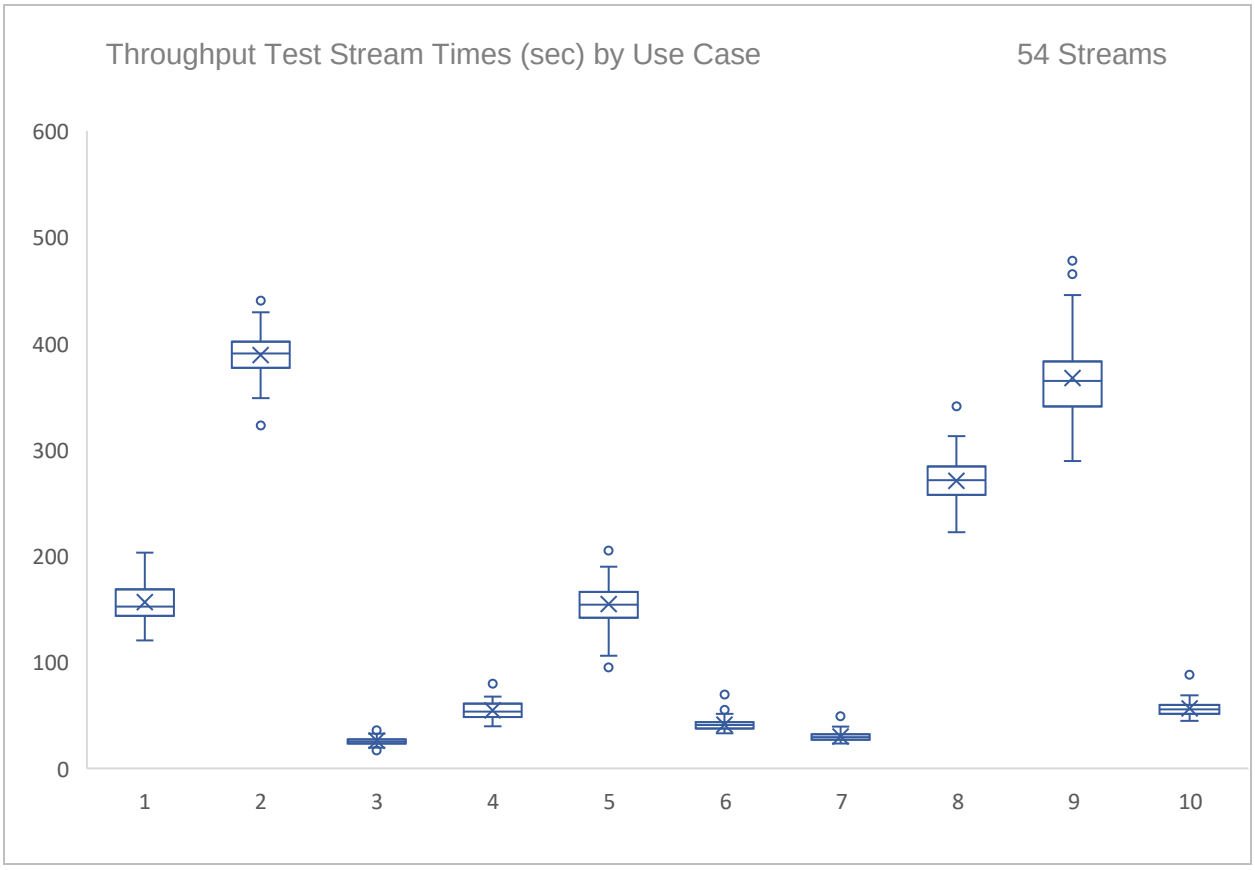

Auditor's Information

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

Paul Cao
Hewlett Packard Enterprise
3-West.103
1701 East Mossy Oaks Road
Spring, TX 77389

March 6, 2025

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

Platform:	8x ProLiant DL380 Gen11; 2x ProLiant DL380 Gen10 Plus
Operating System:	Red Hat Enterprise Linux 8.9 (Ootpa)
Additional Software:	Cloudera 7.1.9

The results were:

Performance Metric	**1,600.16 AIUCpm@100**	
Secondary Metrics	T_{LD}	141.69
	T_{PTT}	121.33
	T_{PST}	36.20
	T_{TT}	3.18

System Under Test	**8x ProLiant DL380 Gen11; 2x ProLiant DL380 Gen10 Plus with:**		
CPUs	2x Intel® Xeon® Platinum 8562Y+ (8 compute nodes)		
	2z Intel® Xeon® Gold 6330 (2 controller nodes)		
Memory	1,024 GiB (8 compute nodes)		
	512 GiB (2 controller nodes)		
Storage	***Qty***	***Size***	***Type***
	2	480 GB	NVMe SSD (7 compute nodes)
	3	3.2 TB	NVMe SSD (7 compute nodes)
	4	3.2 TB	NVMe SSD (1 compute node)
	2	300 GB	300 GB 15K SAS (2 controller nodes)

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:

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

- 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 100 GB.
- The generated dataset used for testing was protected by Replication 3.
- 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:

Because of the small dataset size used for the Validation Test, this Spark-based implementation was not able to satisfy the accuracy threshold for Use Case 3. The TPCx-AI Subcommittee is aware of this issue and has decided that this failure does not invalidate the test.

Respectfully Yours,

Doug Johnson, Certified TPC Auditor

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

Third-Party Price Quotes

HP.com

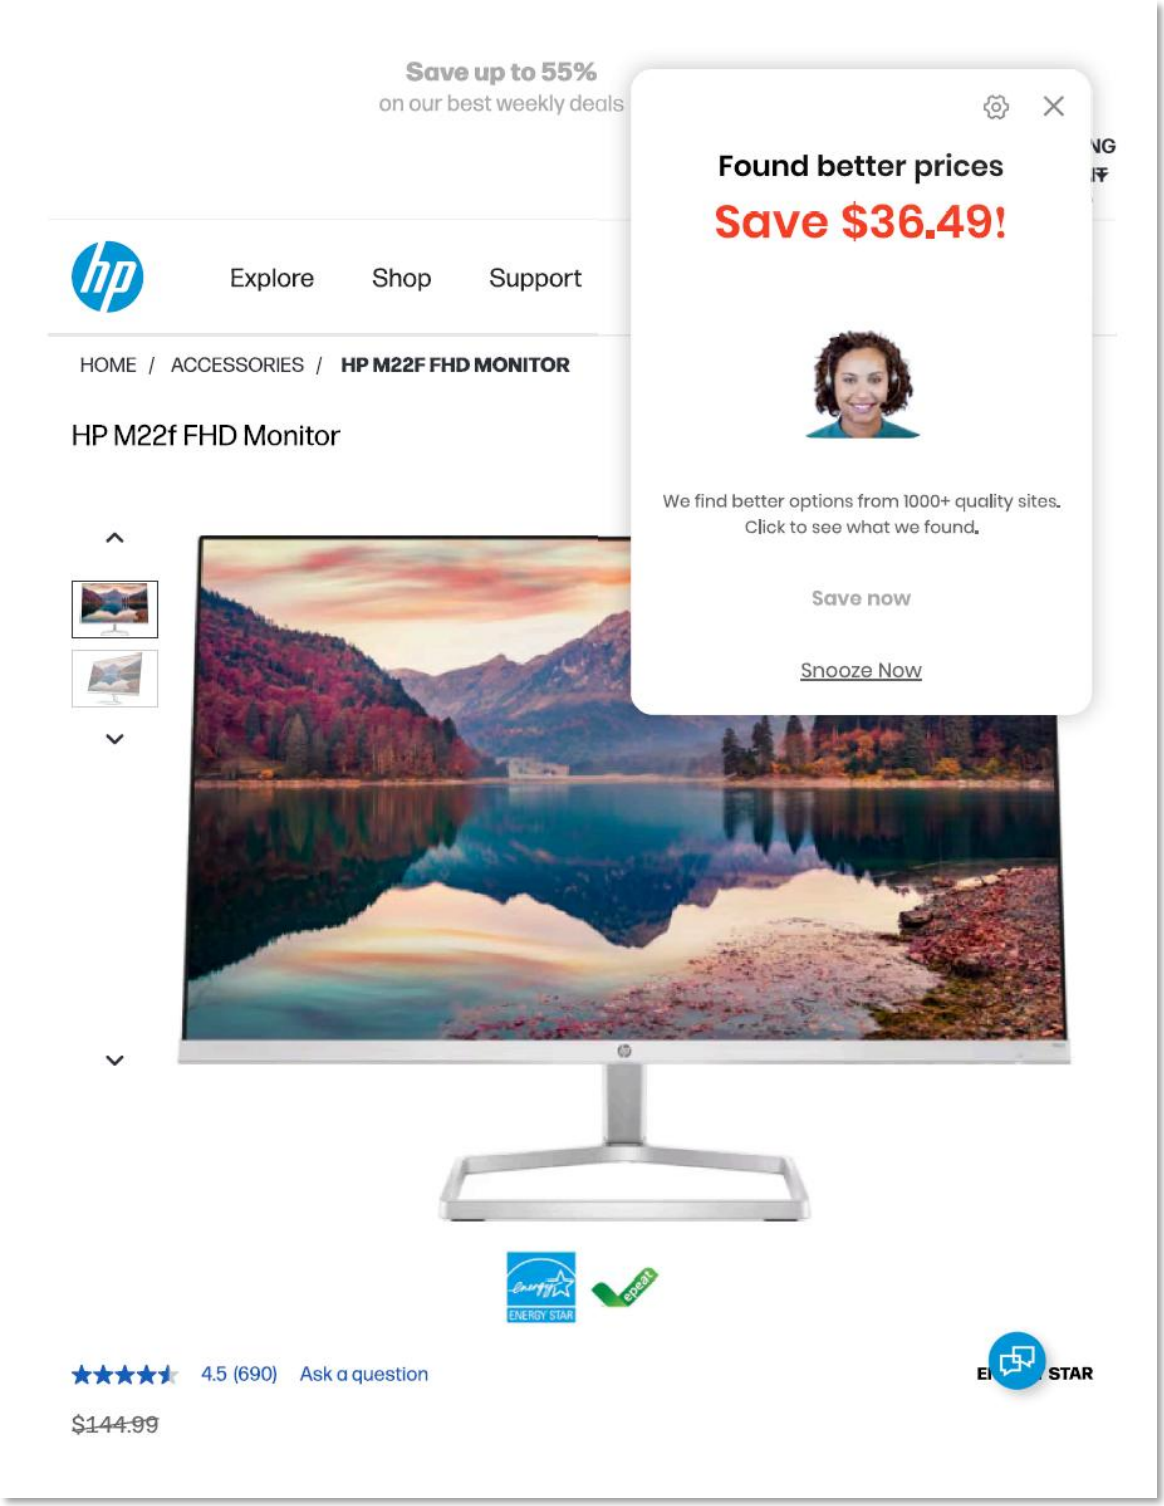

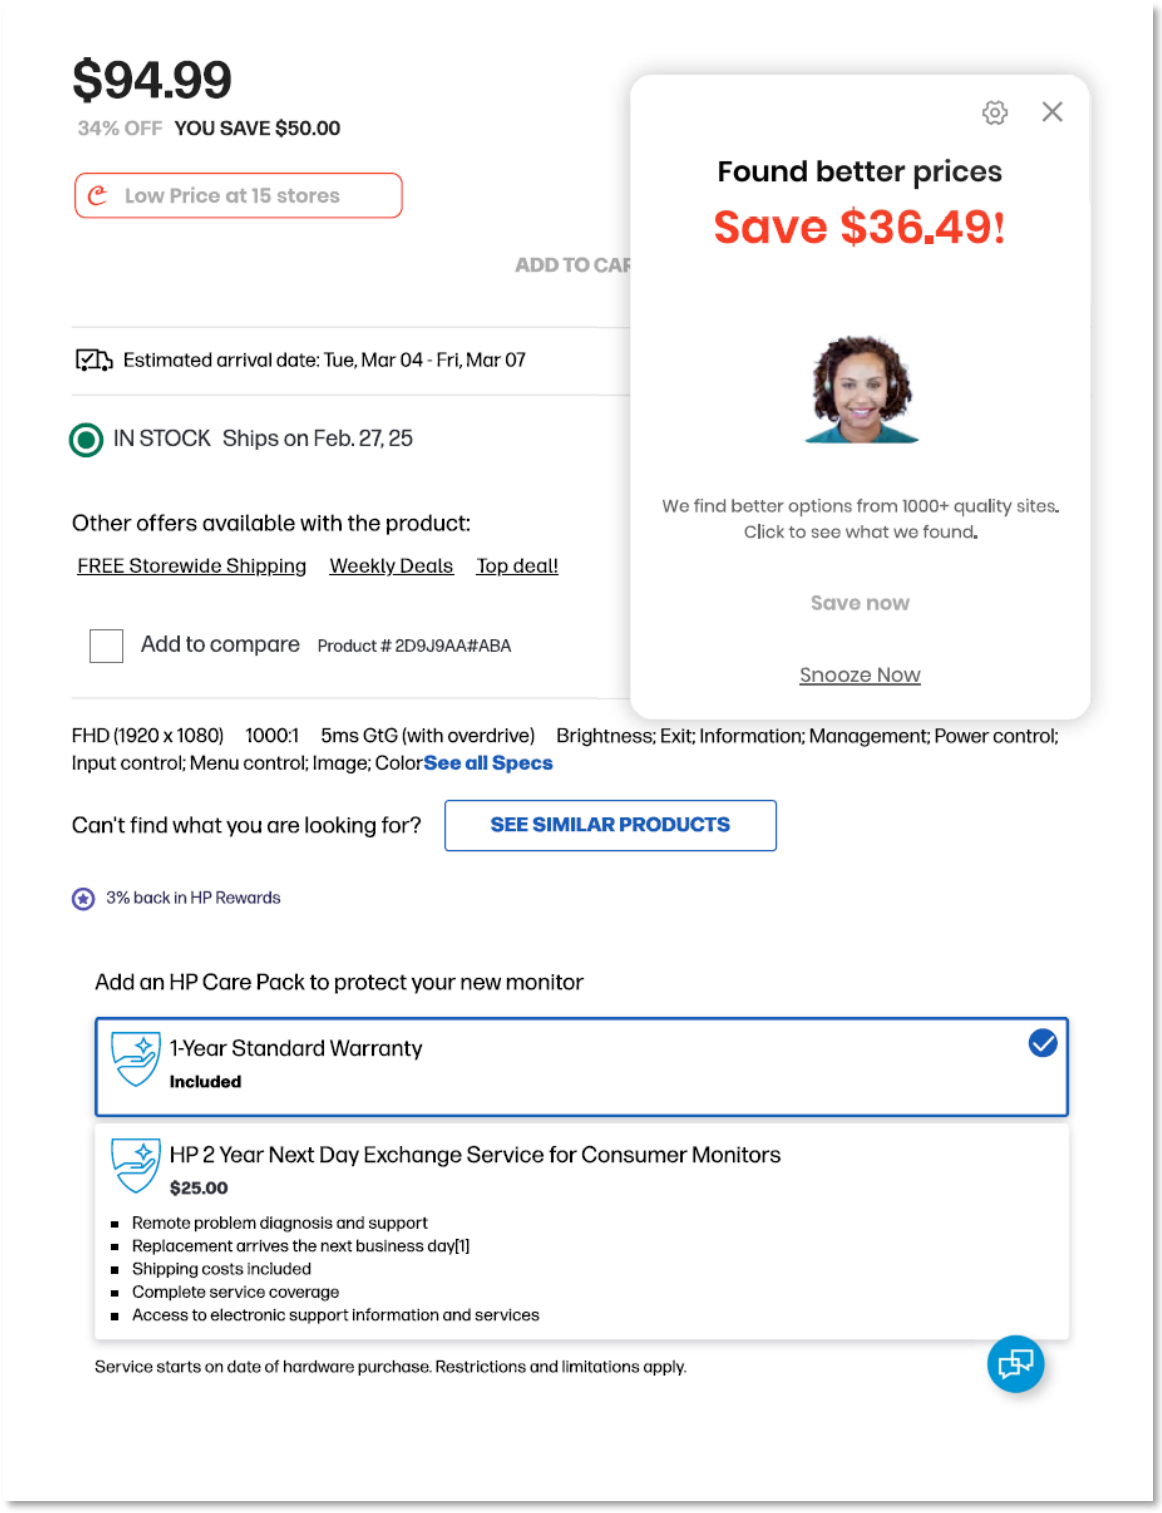
$94.99
34% OFF YOU SAVE $50.00
Low Price at 15 stores
ADD TO CAR
Estimated arrival date: Tue, Mar 04 - Fri, Mar 07
IN STOCK Ships on Feb. 27, 25
Other offers available with the product:
FREE Storewide Shipping
Weekly Deals
Top deal!
Add to compare
Product # 2D9J9AA#ABA
FHD (1920 x 1080) 1000:1 5ms GtG (with overdrive) Brightness; Exit; Information; Management; Power control; Input control; Menu control; Image; Color
See all Specs
Can't find what you are looking for?
SEE SIMILAR PRODUCTS
3% back in HP Rewards
Add an HP Care Pack to protect your new monitor
1-Year Standard Warranty
Included
HP 2 Year Next Day Exchange Service for Consumer Monitors
$25.00
Remote problem diagnosis and support
Replacement arrives the next business day[1]
Shipping costs included
Complete service coverage
Access to electronic support information and services
Service starts on date of hardware purchase. Restrictions and limitations apply.
Found better prices
Save $36.49!
We find better options from 1000+ quality sites.
Click to see what we found.
Save now
Snooze Now

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 HPE	
Sponsor/ModelOptimization/…	Details of model optimization.
Sponsor/ModifiedKitFiles/…	2 modified file(s).
Sponsor/Tuning/…	All tuning files used.