

TPC Express Benchmark[™] IoT Full Disclosure Report for

Lenovo[®] ThinkSystemTM SR655

(With 5x Lenovo SR655 Servers)

Using

HBase 2.1.4 on Cloudera Distribution for Apache Hadoop Enterprise Edition 6.3.0

and

Red Hat Enterprise Linux Server Release 7.6

TPCx-IoT Version V2.0.0 Report Edition Second Report Submitted April 2, 2021

Full Disclosure Report

Second Edition - April 2021

THE INFORMATION CONTAINED IN THIS DOCUMENT IS DISTRIBUTED ON AN AS-IS BASIS WITHOUT ANY WARRANTY EITHER EXPRESSED OR IMPLIED. The use of this information or the implementation of any of these techniques is the customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item has been reviewed by Lenovo for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environment do so at their own risk. 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, the TPC Express BenchmarkTM IoT 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.

In this document, any references made to a Lenovo licensed program are not intended to state or imply that only Lenovo's licensed program may be used; any functionally equivalent program may be used.

This publication was produced in the United States. Lenovo may not offer the products, services, or features discussed in this document in other countries, and the information is subject to change without notice. Consult your local Lenovo representative for information on products and services available in your area.

Lenovo, System x, ThinkSystem, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol ($^{\odot}$ or TM), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at https://www.lenovo.com/us/en/legal/copytrade/. 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 IoT, TPCx-IoT, and IoTps, 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 www.lenovo.com. Actual performance and environmental costs of Lenovo products will vary depending on individual customer configurations and conditions.

Copyright © Lenovo 2021.

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.

				TPCx-IoT: v2.0.0
Lenovo Total System Cost		Lenovo ThinkSystem SR655		TPC-Pricing: v2.7.0
				Report Date: Apr 2, 2021
		TPCx-IoT Per	formance Metric	Price/Performance
\$193,642 USD		742,256	5.79 IoTps	260.89 USD \$/kIoTps
Number of Records	DBMS Software	Operating System	Other Software	Availability Date
1,224 Million	HBase 2.1.4 on Cloudera Distribution for Apache Hadoop 6.3.0	System Red Hat Enterprise Linux Server Release 7.6		Dec 18, 2019
		System Configura	tion	
	1 X AI 8 X TI 1 X TI 1 X TI 2 X TI 2 X TI	MD EPYC ^{IIII} 7502P 32-Core ninkSystem 32GB TruDDR4 ninkSystem M.2 5100 4800 ninkSystem Mellanox Conr ninkSystem U.2 Toshiba Cl	2.5GHz Processor 4 3200MHz (2Rx4 1.2V) RDIM GB SATA 6Gbps Non-Hot Swa 1ectX-4 Lx 10/25GbE SFP28 2 M5-V 1.6TB Mainstream NVM	M-A (256 GB Total) p SSD -port PCIe Ethernet Adapter le PCIe 3.0 x4 Hot Swap SSD
	1x Al 8x Th 1x Th 2x Th 2x Th 1x Ler 48x 2	ND EPYC ^{IIII} /502P 32-Core ninkSystem 32GB TruDDR4 ninkSystem M.2 5100 4800 ninkSystem U.2 Toshiba Cf Novo ThinkSystem NI 25GbE ports + 6x 40GbE/1	2.5GHZ Processor 3200MHz (2Rx4 1.2V) RDIM GB SATA 6Gbps Non-Hot Swa hectX-4 Lx 10/25GbE SFP28 2- M5-V 1.6TB Mainstream NVM E2572 RackSwitch 00GbE ports	M-A (256 GB Total) p SSD -port PCIe Ethernet Adapter le PCIe 3.0 x4 Hot Swap SSD
Image: Second	1x A 8x Ti 1x Ti 1x Ti 2x Ti 2x Ti 1x Ler 48x 2 ers:	MD EPIC ^{IIII} /502P 32-Core ninkSystem 32GB TruDDR4 ninkSystem M.2 5100 4800 ninkSystem U.2 Toshiba Cf novo ThinkSystem NI 25GbE ports + 6x 40GbE/1	2.5GHZ Processor 3.200MHz (2Rx4 1.2V) RDIM GB SATA 6Gbps Non-Hot Swa hectX-4 Lx 10/25GbE SFP28 2- M5-V 1.6TB Mainstream NVW E2572 RackSwitch 00GbE ports 5x Lenovo ThinkSy	M-A (256 GB Total) p SSD -port PCIe Ethernet Adapter le PCIe 3.0 x4 Hot Swap SSD
Total Processors/Core	IX AI 8x Ti 1x Ti 2x Ti 2x Ti 2x Ti 1x Ler 48x 2 ers: •s/Threads:	MD EPIC ^{IIII} /502P 32-Core ninkSystem 32GB TruDDR4 ninkSystem M.2 5100 4800 ninkSystem Mellanox Conr ninkSystem U.2 Toshiba Cf	2.5GH2 Processor 3 2000Hz (2Rx4 1.2V) RDIM GB SATA 6Gbps Non-Hot Swa hectX-4 Lx 10/25GbE SFP28 2- M5-V 1.6TB Mainstream NVW E2572 RackSwitch 00GbE ports 5x Lenovo ThinkSy 5/160/320	M-A (256 GB Total) p SSD -port PCIe Ethernet Adapter le PCIe 3.0 x4 Hot Swap SSD
Total Processors/Core	ers: s/Threads: Processors	ND EPC. ^{III} /302P 32-Core ninkSystem 32GB TruDDR4 ninkSystem M.2 5100 4800 ninkSystem Mellanox Conr ninkSystem U.2 Toshiba Cf Novo ThinkSystem NI 25GbE ports + 6x 40GbE/1	2.5GH2 Processor 3 2000Hz (2Rx4 1.2V) RDIM GB SATA 6Gbps Non-Hot Swa hectX-4 Lx 10/25GbE SFP28 2- M5-V 1.6TB Mainstream NVW E2572 RackSwitch 00GbE ports 5x Lenovo ThinkSy 5/160/320 1x AMD EPYC 750	M-A (256 GB Total) p SSD -port PCIe Ethernet Adapter le PCIe 3.0 x4 Hot Swap SSD sstem SR655
Total Processors/Core	ers: s/Threads: Processors Memory	inkSystem 32GB TruDDR4 hinkSystem M.2 5100 4800 hinkSystem Mellanox Conr hinkSystem U.2 Toshiba Cf hovo ThinkSystem NI 25GbE ports + 6x 40GbE/1	2.5GH2 Processor 3 2000Hz (2Rx4 1.2V) RDIM GB SATA 6Gbps Non-Hot Swa hectX-4 Lx 10/25GbE SFP28 2- M5-V 1.6TB Mainstream NVW E2572 RackSwitch 00GbE ports 5x Lenovo ThinkSy 5/160/320 1x AMD EPYC 750 256GB	M-A (256 GB Total) p SSD port PCIe Ethernet Adapter le PCIe 3.0 x4 Hot Swap SSD ystem SR655 D2P 2.5GHz 32-Core
Total Number of Server Server Configuration (ers: s/Threads: (each):	Vice	2.5GH2 Processor 3.200MHz (2Rx4 1.2V) RDIM GB SATA 6Gbps Non-Hot Swa hectX-4 Lx 10/25GbE SFP28 2- M5-V 1.6TB Mainstream NVW E2572 RackSwitch 00GbE ports 5x Lenovo ThinkSy 5/160/320 1x AMD EPYC 750 256GB 1x 480GB M.2 SAT 2x 1.6TP U.2 NV2	M-A (256 GB Total) p SSD port PCIe Ethernet Adapter le PCIe 3.0 x4 Hot Swap SSD rstem SR655 D2P 2.5GHz 32-Core FA SSD (All Nodes) Le PCIe 2.0 (Data Nation)
Total Number of Server Server Configuration (ers: (each): IX A 8x Ti 1x Ti 2x Th 1x Ler 48x 2 Processors Memory Storage De Network	ND EPIC ^{III} /S02P 32-Cofe ninkSystem 32GB TruDDR4 ninkSystem M.2 5100 4800 ninkSystem Mellanox Conr ninkSystem U.2 Toshiba Cf novo ThinkSystem NI 25GbE ports + 6x 40GbE/1 Vice	 2.5GHZ Processor 3200MHz (2Rx4 1.2V) RDIM GB SATA 6Gbps Non-Hot SwanectX-4 Lx 10/25GbE SFP28 2- M5-V 1.6TB Mainstream NVW E2572 RackSwitch 00GbE ports 5x Lenovo ThinkSy 5/160/320 1x AMD EPYC 750 256GB 1x 480GB M.2 SAT 2x 1.6TB U.2 NVW 1x Mellanox Conneport NIC 	M-A (256 GB Total) p SSD port PCIe Ethernet Adapter le PCIe 3.0 x4 Hot Swap SSD rstem SR655 D2P 2.5GHz 32-Core TA SSD (All Nodes) Ie PCIe 3.0 (Data Nodes) setX-4 Lx 10/25GbE SFP28 2
Total Number of Server Server Configuration (Connectivity:	ers: s/Threads: (each): IX A Sx Ti 1x Ti 2x Ti Ix Ler 48x Z Processors Memory Storage De Network	ND EPIC ^{III} 7302P 32-Core ninkSystem 32GB TruDDR4 ninkSystem M.2 5100 4800 ninkSystem Wellanox Conr ninkSystem U.2 Toshiba Cf novo ThinkSystem NI 25GbE ports + 6x 40GbE/1 Vice	 2.5GHZ Processor 3200MHz (2Rx4 1.2V) RDIM 3B SATA 6Gbps Non-Hot SwanectX-4 Lx 10/25GbE SFP28 2- M5-V 1.6TB Mainstream NVV E2572 RackSwitch 00GbE ports 5x Lenovo ThinkSy 5/160/320 1x AMD EPYC 750 256GB 1x 480GB M.2 SAT 2x 1.6TB U.2 NVN 1x Mellanox Conneport NIC Lenovo ThinkSyste 	M-A (256 GB Total) p SSD -port PCIe Ethernet Adapter le PCIe 3.0 x4 Hot Swap SSD //stem SR655 D2P 2.5GHz 32-Core FA SSD (All Nodes) le PCIe 3.0 (Data Nodes) le PCIe 3.0 (Data Nodes) setX-4 Lx 10/25GbE SFP28 2 m NE2572 Rack Switch

Lenovo

Lenovo ThinkSystem SR655

TPCx-IoT: v2.0.0

TPC-Pricing: v2.7.0

Report Date: Apr 2, 2021

Description	Part	Price	Unit	Qt	Extended	3-Yr. Maint.
Description	Number	Key	Price	У	Price	Price
Server Hardware						
Lenovo SR655 Server : ThinkSystem SR655 - 3yr Warranty	7Z01CTO1WW	1	\$23,268	5	\$116,340	
-ThinkSystem SR655 24x2.5" Chassis	B5VJ			5		
-ThinkSystem AMD EPYC 7502P 32C 180W 2.5GHz Processor	B6VX			5		
-ThinkSystem 32GB TruDDR4 3200MHz (2Rx4 1.2V) RDIMM-A	B5XE			40		
-ThinkSystem 2U 2.5" NVMe 8-Bay Backplane	B4PC			5		
-ThinkSystem U.2 Toshiba CM5-V 1.6TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD	B21X			8		
-ThinkSystem M.2 SATA/NVMe 2-Bay Enablement Kit	B5XJ			5		
-ThinkSystem M.2 5100 480GB SATA 6Gbps Non-Hot Swap SSD	B11V			5		
-ThinkSystem SR655 x16/x8/x8 PCIe Riser2	B5VU			5		
-ThinkSystem Mellanox ConnectX-4 Lx 10/25GbE SFP28 2-port PCIe Ethernet Adapter	B653			5		
-ThinkSystem 1600W (230V) Platinum Hot-Swap Power Supply	AVWG			10		
-2.8m, 13A/100-250V, C13 to C14 Jumper Cord	6400			10		
-ThinkSystem Toolless Slide Rail	AXCA			5		
-ThinkSystem 2U left EIA Latch Standard	AURD			5		
-ThinkSystem SR655 8x PCIE Rear Wall BKT	B5WY			5		
-ThinkSystem SR655 Airduct	B5W5			5		
-ThinkSystem SR655 Riser Filler 1 (3FH)	B5WX			5		
-ThinkSystem SR655 Riser Filler 3 (2FH)	B5WS			5		
-ThinkSystem SR655 Performance Fan Module	B5W3			30		
-ThinkSystem Rome Internal M2 BKT	B5WK			5		
-ThinkSystem SR655 EIA Latch with FIO (Right, 2 USB3.0)	B5WR			5		
-ThinkSystem SR655 MS 2FH+1LP Riser BKT	B5WP			5		
-ThinkSystem OCP3 Filler	B5WJ			5		
-ThinkSystem 2x Supercap BKT	B5WM			5		
-ThinkSystem SR635/SR655 Motherboard	B5VH			5		
Essential Service - 3Yr 24x7 4Hr Response + YourDrive YourData	5PS7A34674	1	\$1,559	5		\$7,795
			Subtotal		\$116,340	\$7,795

Lenovo

Lenovo ThinkSystem SR655

TPCx-IoT: v2.0.0

TPC-Pricing: v2.7.0

Report Date: Apr 2, 2021

Description	Part	Price	Unit	Qty	Extended	3-Yr. Maint.
	Number	Key	Price		Price	Price
Server Software						
RHEL Server Physical w/up to 1 Virtual Node, 2 Skt	7505070110/00/	4	¢0 700	F	¢10 515	
Cloudera Enterprise Operational Database Edition,	CEODN-	I	φ3,703	5	φ10,010	
Node License, Gold Support 1yr 24x7	GOLD-NK	2	\$6,000	15	\$90,000	
			Subtotal		\$108,515	\$0
Network Lenovo ThinkSystem NE2572 RackSwitch (Rear to						
Front)	7159HE3	1	\$23,312	1	\$23,312	
Front)	AV19	1		1		
-Adjustable 19" 4 Post Rail Kit	A3KP	1		1		
-2m Passive DAC SFP+ Cable	A51P	1		5		
-2.8m, 10A/100-250V, C13 to C14 Jumper Cord	6311	1		2		
Essential Service - 3Yr 24x7 4Hr Response (Switch)	5WS7A16480	1	\$3,199	1		\$3,199
			Subtotal		\$23,312	\$3,199
Infrastructure						
Infrastructure Rack : NetBAY S2 25U Rack Cabinet	9307RC2	1	\$1,330	1	\$1,330	
0U 36 C13/6 C19 24A 1 Phase PDU	00YJ776	1	\$479	1	\$479	
ThinkSystem Pref. Pro II USB Keyboard - US English	AXTL	1	\$29	1	\$29	
ThinkSystem Optical Wheel Mouse - USB	BOLN	1	\$19	1	\$19	
Lenovo D22-10 21.5inch LED Backlit LCD Monitor	61D3KCR6US	1	\$89	1	\$89	
Essential Service - 3Yr 24x7 4Hr Response (Rack)	41L2760	1	\$315	1		\$315
			Subtotal		\$1,946	\$315
			Total		\$250,113	\$11,309
Dollar Volume Discount (See Note 1)	39.54%	1			\$67,780	
Pricing: 1 - Lenovo 1-877-782-7134; 2 - C	Cloudera (See		Three-	Year	Cost of	¢102 640
the FDR for details); Note 1: Discount applies to all line			Own	ersh	ip USD:	\$193,04Z
quantities Discounts for similarly sized or	onfigurations					
will be similar to what is guoted here, but	may vary				loTps:	742,256.79
based on the specific components priced.	Audited by					
Pre-Publication Board.				\$/k	doTps:	\$260.89
				0	•	
Prices used in TPC benchmarks reflect the ac	tual prices a cus	stomer	would pay	for a	one-time pu	rchase 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.

Lenovo	Lenovo ThinkSystem SR655	TPCx-IoT: v2.0.0 TPC-Pricing: v2.7.0 Report Date: Apr 2, 2021					
Measurement Results for Performance Run							
Total Number of Records	1,224 Million						
Warmup Run - Start Time Warmup Run - End Time Warmup Run Elapsed Time in Secon	2019-09-03 10:3 2019-09-03 11:1 ds 2,134.338	34:55 10:34					
Measured Run Start Time Measured Run End Time Total Time In Seconds	2019-09-03 11:1 2019-09-03 11:3 1,628.390	10:34 37:51					
M	easurement Results for Repeatability	Run					
Total Number of Records	1,224 Million						
Warmup Run Start Time Warmup Run End Time Warmup Run Elapsed Time in Secon	2019-09-03 11:5 2019-09-03 12:3 ds 2,225.802	56:36 33:42					
Measured Run Start Time Measured Run End Time Total Time In Seconds	2019-09-03 12:3 2019-09-03 13:0 1,649.025	33:42 01:17					



Lenovo ThinkSystem SR655

TPCx-IoT: v2.0.0

TPC-Pricing: v2.7.0

Report Date: Apr 2, 2021

Run Report for Performance Run

```
_____
```

```
TPCx-IoT Performance Metric (IoTps) Report
```

```
Total Time For Warmup Run In Seconds = 2,134.338
Total Time In Seconds = 1,628.390
Total Number of Records = 1,224 Million
```

TPCx-IoT Performance Metric (IoTps): 751,662.6852

Run Report for Repeatability Run

```
TPCx-IoT Performance Metric (IoTps) Report
Total Time For Warmup Run In Seconds = 2,225.802
Total Time In Seconds = 1,649.025
Total Number of Records = 1,224 Million
TPCx-IoT Performance Metric (IoTps): 742,256.7881
```

Table of Contents

Contents

ABSTRACT	9
PREFACE	10
CLAUSE 1: GENERAL ITEMS	11
1.1 TEST SPONSOR	11
1.2 PARAMETER SETTINGS	11
1.3 CONFIGURATION DIAGRAMS	11
1.4 DATASET DISTRIBUTION	13
1.5 SOFTWARE COMPONENTS DISTRIBUTION	13
CLAUSE 2: WORKLOAD RELATED ITEMS	14
2.1 HARDWARE & SOFTWARE TUNABLE	14
2.2 Run Report	14
2.3 BENCHMARK KIT IDENTIFICATION	15
2.4 BENCHMARK KIT CHANGES	15
CLAUSE 3: SCALE FACTORS AND METRICS	17
3.1 TOTAL RUN TIME	17
3.2 PERFORMANCE AND PRICE PERFORMANCE	17
3.3 System Configuration Information	17
SUPPORTING FILE INDEX	
THIRD PARTY PRICE QUOTES	19

Abstract

This document contains the methodology and results of the TPC Express Benchmark IoT (TPCx-IoT) test conducted in conformance with the requirements of the TPCx-IoT Standard Specification, Revision 2.0.0.

The test was conducted for a Scale Factor of 1,224 million records with Lenovo ThinkSystem SR655 servers running HBase 2.1.4 on Cloudera Distribution for Apache Hadoop Edition 6.3.0 on Red Hat Enterprise Linux Server Release 7.6.

Measured Configuration

Company Name	Cluster Node	Virtualization	Operating System
Lenovo	Lenovo ThinkSystem SR655	Not Used	Red Hat Enterprise Linux Server Release 7.6

TPC Express Benchmark IoT Metrics

Total System Cost	IoTps	Price/Performance	Availability Date
\$193,642 USD	742,256.79	\$260.89 USD / kIoTps	Dec 18, 2019

Preface

TPC Express Benchmark IoT Overview

TPC Express Benchmark IoT (TPCx-IoT) was developed to provide an objective measure of hardware, operating system and commercial NoSQL database software distributions, and to provide the industry with verifiable performance, price-performance and availability metrics. The benchmark models a continuous system availability of 24 hours a day, 7 days a week.

Even though the modeled application is simple, the results are highly relevant to hardware and software dealing with IoT Gateway systems in general. The TPCx-IoT stresses both hardware and software including database APIs and network connections to the database. This workload can be used to assess a broad range of NoSQL databases. The TPCx-IoT can be used to assess a range of NoSQL implementations in a technically rigorous and directly comparable and vendor-neutral manner. The metric effectively represents the total number of records that can be inserted into a NoSQL database per second while running queries against the database.

The TPCx-IoT kit is available from the TPC (See www.tpc.org/tpcx-iot for more information). Users must sign- up and agree to the TPCx-IoT User Licensing Agreement (ULA) to download the kit. Re-distribution of the kit is prohibited. All related work (such as collaterals, papers, derivatives) must acknowledge the TPC and include TPCx-IoT copyright. The TPCx-IoT Kit includes: TPCx-IoT Specification document, TPCx-IoT Users Guide documentation, shell scripts to set up the benchmark environment and Java code to execute the benchmark load.

The purpose of TPC benchmarks is to provide relevant, objective performance data to industry users. To achieve that purpose, TPC benchmark specifications require that benchmark tests 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- IoT models and represents a NoSQL database mimicking an IoT gateway system)

• Would plausibly be implemented by a significant number of users in the market segment the benchmark models or represents.

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 atwww.tpc.org

Clause 1: General Items

1.1 Test Sponsor

A statement identifying the benchmark sponsor(s) and other participating companies must be provided.

This benchmark was sponsored by Lenovo.

1.2 Parameter Settings

Settings must be provided for all customer-tunable parameters and options which have been changed from the defaults found in actual products, including by not limited to:

- Configuration parameters and options for server, storage, network and other hardware component incorporated into the pricing structure;
- Configuration parameters and options for operating system and file system component incorporated into the pricing structure;
- Configuration parameters and options for any other software component incorporated into the pricing structure;
- Compiler optimization options.

Comment 1: In the event that some parameters and options are set multiple times, it must be easily discernible by an interested reader when the parameter or option was modified and what new value it received each time.

Comment 2: This requirement can be satisfied by providing a full list of all parameters and options, as long as all those that have been modified from their default values have been clearly identified and these parameters and options are only set once.

The supporting files contain the parameters and options used to configure the components involved in this benchmark.

1.3 Configuration Diagrams

Diagrams of both measured and priced configurations must be provided, accompanied by a description of the differences. This includes, but is not limited to:

- Total number of nodes used;
- Total number and type of processors used/total number of cores used/total number of threads used (including sizes of L2 and L3 caches);

- Size of allocated memory, and any specific mapping/partitioning of memory unique to the test;
- Number and type of disk units (and controllers,) if applicable;
- Number of channels or bus connections to disk units, including their protocol type;
- Number of LAN (e.g., Ethernet) connections and speed for switches and other hardware components physically used in the test or are incorporated into the pricing structure;
- *Type and the run-time execution location of software components.*

Measured Configuration:



The measured configuration consisted of:

- Total Nodes: 5
- Total Processor/Cores/Threads: 5/160/320
- L1/L2/L3 Processor Cache (MB): 2/16/128
- Total Memory: 1.28TB
- Total Number of Storage Devices: 13
- Total Storage Capacity: 15.2 TB

Server nodes details:

- Lenovo ThinkSystem SR655 servers, each with:
 - o Processors/Cores/Threads: 1/32/64
 - Processor Model: 1x AMD EPYCTM 7502P 2.5GHz 32-core 128MB L3
 - Memory: 256GB (8 x 32GB 2R RDIMM-A 3200MB/s), 8 Channels
 - Drives: 1x 480GB M.2 SATA SSD (for all Servers)
 - 2x 1.6TB U.2 NVMe PCIe Gen3 (for all Data Node Servers)
 - Network: 1x Mellanox ConnectX-4 Lx 10/25GbE SFP28 2-port NIC

Priced Configuration:

There are no differences between the priced and measured configurations.

1.4 Dataset Distribution

The distribution of dataset across all media must be explicitly described.

Table 1.4 describes the distribution of the dataset across all media in the system.

Table 1.4: Dataset Distribution

Server Node	Storage	Disk Drive	Description of Content
1	M.2 SATA	1xSSD	Operating System, Swap, Hadoop Master, Root, Temp
2-4	M.2 SATA	1xSSD	Operating System, Swap, Hadoop Master, Root, Temp
	U.2 PCIe Gen3	NVMe0n1, NVMe1n1	Data, Temp

1.5 Software Components Distribution

The distribution of various software components across the system must be explicitly described.

Table 1.5 describes the distribution of the software components across the system.

Table 1.5: Software Component Distribution

	HD	FS	HB	ase	YA	RN	Zoo Keeper
Node	NameNode	DataNode	Master	Region	Resource	Node	
				Server	Manager	Manager	
1	Х		Х		Х		Х
2-3		Х		Х		Х	Х
4-5		Х		X		Х	

NoSQL Database version must be disclosed.

HBase -2.1.4 on Cloudera Distribution for Apache Hadoop 6.3.0

Full Disclosure Report

Clause 2: Workload Related Items

2.1 Hardware & Software Tunable

Script or text used to set for all hardware and software tunable parameters must be reported.

The Supporting File Archive contains all configuration scripts.

2.2 Run Report

The run report generated by TPCx-IoT benchmark kit must be reported.

The Supporting File Archive contains the full run report. Following are extracts from the run report that lists the performance summary for both runs.

Run Report for Performance Run

TPCx-IoT Performance Metric (IoTps) Report Total Time For Warmup Run In Seconds = 2,134.338 Total Time In Seconds = 1,628.390 Total Number of Records = 1,224 Million TPCx-IoT Performance Metric (IoTps): 751,662.6852

Run Report for Repeatability Run

TPCx-IoT Performance Metric (IoTps) Report Total Time For Warmup Run In Seconds = 2,225.802 Total Time In Seconds = 1,649.025 Total Number of Records = 1,224 Million TPCx-IoT Performance Metric (IoTps): 742,256.7881

2.3 Benchmark Kit Identification

Version number of TPCx-IoT kit and checksum for the jar file and master Programs must be reported.

Table 2.3

Kit Version	1.0.3
0fc49d9c668d22d6a394279fd27474b7	./TPC-IoT-master.sh
68379f9375c7b584fc3253dfe9c4f7a6	./tpcx-iot/lib/core-0.13.0-SNAPSHOT.jar
7bebf1e17d5c2b380df575fad160d7f8	./loT_cluster_validate_suite.sh

2.4 Benchmark Kit Changes

No modifications were made to the TPC provided kit other than the following fixes made in these three scripts:

- TPC-IoT-master.sh
- TPC-IoT-client.sh
- TPC-IoT-instances.sh

TPCx-IoT kit did not generate tpc_iot_instanceNN_workload files with correct values for benchmark tests that involved multiple client (driver) nodes.

```
# diff
         org/TPC-IoT-master.sh TPC-IoT-master.sh
130c130
      insertstart=$(echo "$insertstart+$num_records_per_client+1" | bc)
<
      insertstart=$(echo "$insertstart+$num_records_per_client" | bc)
>
# diff org/TPC-IoT-client.sh TPC-IoT-client.sh
35c35
< DATABASE RECORDS COUNT=$ (echo $operation count string | cut -d'=' -f2)
> OPERATION_COUNT=$ (echo $operation_count_string | cut -d'=' -f2)
38,39c38,39
< echo ">>>>>>> $PWD/TPC-IoT-instances.sh $DATABASE_RECORDS_COUNT $NUM_INSTANCES
$NUM_THREADS $INSERT_START $clientId $DATABASE_CLIENT $LOGFILE_NAME"
< $PWD/TPC-IoT-instances.sh $DATABASE_RECORDS_COUNT $NUM_INSTANCES $NUM_THREADS
$INSERT_START $clientId $DATABASE_CLIENT $PWD $SUT_PARAMETERS $LOGFILE_NAME
> echo ">>>>>>>> $PWD/TPC-IoT-instances.sh $DATABASE RECORDS COUNT
$OPERATION_COUNT $NUM_INSTANCES $NUM_THREADS $INSERT_START $clientId
$DATABASE CLIENT $LOGFILE NAME"
> $PWD/TPC-IoT-instances.sh $DATABASE_RECORDS_COUNT $OPERATION_COUNT
```

\$NUM_INSTANCES \$NUM_THREADS \$INSERT_START \$clientId \$DATABASE_CLIENT \$PWD
\$SUT_PARAMETERS \$LOGFILE_NAME

```
# diff org/TPC-IoT-instances.sh TPC-IoT-instances.sh
10,17c10,20
< numInstances=$2
< threadCount=$3
< start=$4
< clientID=$5
< DATABASE_CLIENT=$6
< PWD=$7
< SUT_PARAMETERS=$8
< RUN_TYPE=$9
___
> totalOperationCount=$2
> numInstances=$3
> threadCount=$4
> start=$5
> clientID=$6
> DATABASE_CLIENT=$7
> PWD=$8
> SUT_PARAMETERS=$9
> RUN_TYPE=${10}
>
> operationCount=$((totalOperationCount / numInstances)) # Improve this to be
total of record count
19d21
< operationCount=$((recordCount / numInstances)) # Improve this to be total of
record count
52c54
< start=$((operationCount * counter))
> start=$((operationCount + start))
```

A separate document is provided for the justification and details of the fixes. See details in the SUPPORTING_FILES_ARCHIVE\clause2\additional-files\IssuesFixed-TPCx-IoT-scripts-documentation.rtf file.

Clause 3: Scale Factors and Metrics

3.1 Total Run Time

	Run 1	Run 2
Total Run Time	1,628.390	1,649.025

3.2 Performance and Price Performance

The performance metric (IoTps) must be disclosed for Run1 and Run2. The price-performance metric (\$/kIoTps) must be disclosed for the performance run.

	Run 1	Run 2
IoTps	751,662.68	742,256.79

\$/kloTps \$260.89

3.3 System Configuration Information

Storage System Software	Operating System	Other Software	System Availability Date
	Red Hat Enterprise Linux Server Release 7.6		Dec 18, 2019

Cloudera 6.3.0				
Component	Package Version			
Apache Hadoop	hadoop-3.0.0+cdh6.3.0			
HBase	hbase-2.1.4+cdh6.3.0			
YARN	yarn-3.0.0+cdh6.3.0			
Zookeeper	zookeeper-3.4.5+cdh6.3.0			

Supporting File Index

Clause	Contents	Location		
Clause 1	Parameters and options used to configure and tune the SUT	Supporting Files Archive/Clause1		
Clause 2	Configuration Scripts and Run Report	Supporting Files Archive/Clause2		
Clause 3	System Configuration Details	Supporting Files Archive/Clause3		

The following index outlines the information included in the supporting file archive.

Third-Party Price Quotes

Cloudera Pricing



395 Page Mill Road Palo Alto, CA 94306 US Phone: +1 (650) 843-0595 Email: dwilson@cloudera.com Quote #: Date: Offer Expires On: End User: Subscription Bill Terms:

Bill To

XXXXX

Q-XXXXX X-X--20XX X-X--20XX XXXXXXXX

Ship To XXXXXX

Product Name	Product Code	Description	List Price	Qty	Start Date	End Date	Discount (%)	Net Total
Cloudera Enterprise Operational Database Edition, Node License, Gold Support	CEODN- GOLD-NK	Subscription (per node) for the Operational Database Edition of the Cloudera Enterprise platform. Includes a commercial license and Gold Support for "Core Hadoop" (i.e. HDFS, YARN/ Mapreduce, Hive, Pig, Hue, Sentry, Flume, Sqoop), Kudu, HBase, Accumulo, Search, Spark, Cloudera Manager, Cloudera Director, and Cloudera Navigator (Audit & Lineage, Encryption & Key Trustee). Gold-Level Support provides technical support 24 hours per day, 7 days per week. The price listed is for an annual (12-month) term.	USD 6,000.00	5	8-20-2019	8-19-2020	0.00	USD 30,000.00
	-	·	-				TOTAL:	USD

30,000.00

The resale of the Cloudera Products listed in this Order Form to the End User identified above is governed by the current, valid, mutually executed reseller agreement between Cloudera and the reseller identified above in the Bill To section ("Reseller") dated prior to or on or about the Effective Date of this Order Form ("Reseller Agreement"). Per the terms of the Reseller Agreement, Reseller will ensure that the End User: (i) receives instructions on how to access and accept the applicable Cloudera End Customer Agreement; or (ii) verifies that there is an applicable executed agreement in place with Cloudera for the Cloudera Products provided herein.

The Reseller Agreement shall prevail over any additional, conflicting or inconsistent terms and conditions which may appear or any purchase order furnished by Reseller, and any additional terms and conditions in any such purchase order shall have no force or effect, notwithstanding Cloudera's acceptance or execution of such purchase order.

Cloudera will invoice Reseller upon receipt of this signed Order Form, and fees will be due as set forth in the agreement between Cloudera and Reseller governing the terms of this Order Form.

This Order Form may be executed by exchange of signature by electronic means through facsimile or scanned and emailed signature, or by electronic signature service where legally permitted. For clarity, an electronic, digital, machine-generated or image of a signature will create a valid and binding obligation of Customer.

Signature:

Effective Date:

Name (Print):

Title: