

# Transwarp Technology (Shanghai) Co., Ltd

---

TPC Express Benchmark™ Big Bench (TPCx-BB)

Full Disclosure Report

for

Transwarp Big Data Appliance

(with 20x TxData-2L Servers)

using

Transwarp Data Hub 9.1

and

CentOS Linux release 8

---

**First Edition**

**September 19, 2023**

**Transwarp Technology (Shanghai) Co., Ltd (Transwarp)**, 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 as a result of these and other factors. Therefore, the TPC Express Benchmark™ BB 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.





Transwarp and the Transwarp Logo are trademarks of Transwarp Technology (Shanghai) Co., Ltd and/or its affiliates in China and other countries. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Transwarp and any other company.

TPC Benchmark™, TPCx-BB and BBQpm, are registered certification marks of the Transaction Processing Performance Council.

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

**Copyright © 2023 Transwarp Technology (Shanghai) Co., Ltd**

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.

 星 环 科 技		<h2>Transwarp Big Data Appliance</h2>			TPCx-BB Rev. v1.6.1 TPC-Pricing Rev. v2.8.0	
					Report Date: September 19, 2023	
Total System Cost		TPCx-BB Performance Metric			Price/Performance	
<b>3,760,954 CNY</b>		<b>5,230.08</b> BBQpm@3000			<b>719.11 CNY</b> ¥/BBQpm@3000	
Framework	Operating System	Other Software	Availability Date	Scale Factor	Streams	
Transwarp Data Hub 9.1	CentOS Linux release 8	Spark 2.3.3	September 19, 2023	3000	8	
<h3>System Configuration</h3>						
<div style="display: flex; justify-content: space-around; align-items: center;"> <div style="text-align: center;">  <p><b>20 x Transwarp Big Data Appliance TxData-2L server</b>            2 x Intel® Xeon® Gold 6330 Processor            256GB (8 x Samsung 32G DDR4 RECC 3200MT/s)            2 x Intel P4510 U.2 1T SSD            6 x Intel P5510 U.2-3.84T NVMe SSD            1 x Intel I350 2-port 1Gb Network Adapter            1 x Mellanox 100G EDR IB Dual-Port QSFP28 Network Adapter</p> </div> <div style="text-align: center;">  <p>1x Mellanox MSB7800-ES2F</p> </div> <div style="text-align: center;">  <p>1x H3C S5110-52P-SI</p> </div> </div>						
Physical Storage/Scale Factor: 166.93		Scale Factor/Physical Memory: 0.59		Main Data Redundancy Model: 3-way Replication		
Servers:		20x TxData-2L				
Total Processors/Cores/Threads		40/1,120/2,240				
Server Configuration:		<b>20x TxData-2L (Cluster Node):</b>				
Processors		2x Intel(R) Xeon(R) Gold 6330 CPU @ 2.00GHz				
Memory		256 GiB				
Storage Device		2x 1 TB NVMe SSD 6x 3.84 TB NVMe SSD				
Network Controller		1x Intel I350 2-port 1 Gb 1x Mellanox 100 G EDR IB Dual-Port QSFP28				
Connectivity:		1x Mellanox MSB7800-ES2F Switch 1x H3C S5110-52P-SI Switch				



# Transwarp Big Data Appliance

TPCx-BB Rev. v1.6.1  
TPC-Pricing Rev. v2.8.0

Report Date:  
September 19, 2023

Description	Part Number	Source	List Price	Qty	Extended Price	3-Yr. Maintenance
<b>Transwarp Big Data Appliance TxData-2L server</b>						
Transwarp Big Data Appliance TxData-2L server, 2U, single node	PRO-993		1 ¥24,450	20	¥489,000	
Intel® Xeon® Gold 6330 Processor, 42M Cache, 2.00 GHz, 28C, 56T	PRO-881		1 ¥22,560	40	¥902,400	
Samsung 32G DDR4 RECC 3200MT/s	PRO-764		1 ¥1,670	160	¥267,200	
Intel P4510 U.2 1T SSD	PRO-633		1 ¥2,100	40	¥84,000	
Intel P5510 U.2 3.84T NVMe SSD	PRO-668		1 ¥6,200	120	¥744,000	
Intel I350 2-port 1Gb Network Adapter	PRO-518		1 ¥700	20	¥14,000	
Mellanox 100G EDR IB Dual-Port QSFP28 Network Adapter	PRO-556		1 ¥8,400	20	¥168,000	
Transwarp 4-hour 7x24 On-site Service, 3 years	PRO-100		1 ¥133,430	1		¥133,430
<b>Network, Cables, Infrastructure</b>						
Mellanox MSB7800-ES2F Switch-IB 2 Based EDR InfiniBand 1U Switch 36 QSFP28 Ports	MSB7800-ES2F		1 ¥159,840	1	¥159,840	
Support and Warranty - 3 Year for MSB7800-ES2F	SUP-MSB7800-ES2F		1 ¥15,984	1		¥15,984
Mellanox® Passive Copper cable, IB EDR, up to 100Gb/s, QSFP28, 5m, Black, 26AWG			1 ¥2,795	20	¥55,900	
H3C S5110-52P-SI 52-port Gigabit Ethernet Switch,with 1 year Support and Warranty	S5110-52P-SI		1 ¥2,800	3	¥8,400	
42U Enclosure system			1 ¥4,000	1	¥4,000	
24" LED Monitor			1 ¥1,500	3	¥4,500	
Keyboard and Mouse			1 ¥100	3	¥300	
<b>Hardware Subtotals</b>					<b>¥2,901,540</b>	<b>¥149,414</b>
<b>Software</b>						
Transwarp Data Hub 9.1 Subscription Edition - 3 Years			1 ¥30,000	20		¥600,000
Support Service - 3 Years for CentOS 8			1 ¥5,500	20		¥110,000
<b>Software Subtotals</b>					<b>¥0</b>	<b>¥710,000</b>
Pricing:1 = {Source 1}; 2 = {Source 2}  <sup>(1)</sup> All discounts are based on US list prices and for similar quantities and configurations. The discounts are based on the overall specific components pricing from respective vendors in this single quotation. Discounts for similarly sized configurations will be similar to those quoted here, but may vary based on the components in the configuration.  <b>Audited by Doug Johnson, InfoSizing</b>					<b>Three-Year Cost of Ownership ¥3,760,954</b>  <b>BBQpm@3000 5,230.08</b>  <b>¥/BBQpm@3000 ¥ 719.11</b>	
Prices used in TPC benchmarks reflect the actual prices a customer would pay for a one-time purchase of the stated components. 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 components. For complete details, see the pricing sections of the TPC benchmark specifications. If you find that the stated prices are not available according to these terms, please inform at <a href="mailto:pricing@tpc.org">pricing@tpc.org</a> . Thank you.						

Numerical Quantities

Scale Factor	3000
Streams	8
SUT Validation Test	PASS

Performance Run (Run 2)

Overall Run Start Time	2023-09-04 14:38:04.898
Overall Run End Time	2023-09-04 16:50:48.853
Overall Run Elapsed Time	7,963.955
Load Test Start Time	2023-09-04 14:38:04.900
Load Test End Time	2023-09-04 14:46:45.226
Load Test Elapsed Time	520.326
Power Test Start Time	2023-09-04 14:46:45.252
Power Test End Time	2023-09-04 15:16:14.590
Power Test Elapsed Time	1,769.338
Throughput Test Start Time	2023-09-04 15:16:14.590
Throughput Test End Time	2023-09-04 16:50:48.853
Throughput Test Elapsed Time	5,674.263
Performance Metric (BBQpm@ 3000)	5,230.08

Repeatability Run (Run 1)

Overall Run Start Time	2023-09-04 11:58:26.919
Overall Run End Time	2023-09-04 14:09:51.204
Overall Run Elapsed Time	7,884.285
Load Test Start Time	2023-09-04 11:58:26.919
Load Test End Time	2023-09-04 12:07:23.459
Load Test Elapsed Time	536.540
Power Test Start Time	2023-09-04 12:07:23.460
Power Test End Time	2023-09-04 12:36:17.853
Power Test Elapsed Time	1,734.393
Throughput Test Start Time	2023-09-04 12:36:17.854
Throughput Test End Time	2023-09-04 14:09:51.203
Throughput Test Elapsed Time	5,613.349
Performance Metric (BBQpm@ 3000)	5,273.03

## Performance Run Report (Run 2)

\*\*\*\*\*

TPCx-BB

Result

v1.6.1

\*\*\*\*\*

INFO: T\_LOAD = 520.326

INFO: T\_LD = 0.1 \* T\_LOAD: 52.0326

INFO: T\_PT = 1355.30495689559

INFO: T\_T\_PUT = 5674.263

INFO: T\_TT = 709.282875

INFO: === Checking validity of the final result ===

INFO: OK: All required BigBench phases were performed.

INFO: OK: All 30 queries were running in the power test.

INFO: OK: All 30 queries were running in the first throughput test.

INFO: OK: Pretend mode was inactive. All commands were executed.

INFO: === Final result ===

INFO: VALID BBQpm@3000 = 5230.08033596877

## Repeatability Run Report (Run 1)

\*\*\*\*\*

TPCx-BB

Result

v1.6.1

\*\*\*\*\*

INFO: T\_LOAD = 536.54

INFO: T\_LD = 0.1 \* T\_LOAD: 53.6539999999999

INFO: T\_PT = 1342.11765250356

INFO: T\_T\_PUT = 5613.349

INFO: T\_TT = 701.668625

INFO: === Checking validity of the final result ===

INFO: OK: All required BigBench phases were performed.

INFO: OK: All 30 queries were running in the power test.

INFO: OK: All 30 queries were running in the first throughput test.

INFO: OK: Pretend mode was inactive. All commands were executed.

INFO: === Final result ===

INFO: VALID BBQpm@3000 = 5273.03820932679

Summary details of the run reports are shown above. For the complete run reports, see the Support Files Archive.

# Table of Contents

---

ABSTRACT .....	8
PREFACE .....	9
CLAUSE 1: GENERAL ITEMS .....	10
1.1 TEST SPONSOR.....	10
1.2 PARAMETER SETTINGS .....	10
1.3 CONFIGURATION DIAGRAMS.....	10
CLAUSE 2: SOFTWARE COMPONENTS AND DATASET DISTRIBUTION .....	12
2.1 ROLES AND DATASET DISTRIBUTION.....	12
2.2 DISTRIBUTED FILE SYSTEM IMPLEMENTATION .....	12
2.3 ENGINE IMPLEMENTATION .....	12
2.4 FRAMEWORKS .....	13
2.5 APPLIED PATCHES .....	13
CLAUSE 3: WORKLOAD RELATED ITEMS .....	14
3.1 HARDWARE & SOFTWARE TUNABLE .....	14
3.2 KIT VERSION .....	14
3.3 RUN REPORT.....	14
3.4 QUERY ELAPSED TIMES .....	15
3.5 VALIDATION TEST OUTPUT.....	16
3.6 GLOBAL FRAMEWORK PARAMETERS.....	16
3.7 KIT MODIFICATIONS.....	16
CLAUSE 4: SUT RELATED ITEMS.....	17
4.1 SPECIALIZED HARDWARE/SOFTWARE .....	17
4.2 FRAMEWORK CONFIGURATION FILES .....	17
4.3 SUT ENVIRONMENT INFORMATION .....	17
4.4 DATA STORAGE TO SCALE FACTOR RATIO.....	17
4.5 SCALE FACTOR TO MEMORY RATIO .....	17
CLAUSE 5: METRICS AND SCALE FACTORS.....	18
5.1 PERFORMANCE RUN METRIC .....	18
5.2 REPEATABILITY RUN METRIC .....	18
5.3 PRICE-PERFORMANCE METRIC .....	18
5.4 SCALE FACTOR.....	18
5.5 STREAM COUNT.....	18
5.6 ELAPSED RUN TIMES.....	19
5.7 ELAPSED TEST TIMES.....	19
AUDITORS' INFORMATION AND ATTESTATION LETTER.....	20
THIRD PARTY PRICE QUOTES.....	23
SUPPORTING FILE INDEX.....	24

# Abstract

This document contains the methodology and results of the TPC Express Benchmark™ Big Bench (TPCx-BB) test conducted in conformance with the requirements of the TPCx-BB Standard Specification, Revision v1.6.1.

The test was conducted at a Scale Factor of 3000 with 20 nodes (20x TxData-2L, [Click or tap here to enter text.x](#) [Click or tap here to enter text.](#)) running Transwarp Data Hub 9.1 on CentOS Linux release 8.

## Measured Configuration

Company Name	Cluster Node	Virtualization	Operating System
Transwarp Technology (Shanghai) Co., Ltd	20x TxData-2L	n/a	CentOS Linux release 8

## TPC Express Benchmark© Big Bench Metrics

Total System Cost	BBQpm@3000	Price/Performance	Availability Date
3,760,954 CNY	5,230.08	719.11 CNY	September 19, 2023



# Preface

## TPC Express Benchmark™ Big Bench Overview

*Big data analytics is a growing field of research and business. The significant decrease in the overall cost of hardware, the emergence of Open Source based analytics frameworks, along with the greater depth of data mining capabilities allows new types of data sources to be correlated with traditional data sources. For example, online retailers used to record only successful transactions on their website, whereas modern systems are capable of recording every interaction. The former allowed for simple shopping basket analysis techniques, while the current level of detail in monitoring makes detailed user modeling possible. The growing demands on data management systems and the new forms of analysis have led to the development of a new type of **Big Data Analytics Systems (BDAS)**.*

*Similar to the advent of **Database Management Systems**, there is a vastly growing ecosystem of diverse approaches to enabling Big Data Analytics Systems. This leads to a dilemma for customers of **BDAS**, as there are no realistic and proven measures to compare different **BDAS** solutions. To address this, TPC has developed TPCx-BB (BigBench), which is an express benchmark for comparing **BDAS** solutions. The TPCx-BB Benchmark was developed to cover essential functional and business aspects of big data use cases. The benchmark allows for an objective measurement of **BDAS** System under Test, and provides the industry with verifiable performance, price/performance, and availability metrics.*

*The TPCx-BB kit is available from the TPC website (see [www.tpc.org](http://www.tpc.org) for more information). Users must sign-up and agree to the TPCx-BB 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-BB copyright. The TPCx-BB kit includes: TPCx-BB Specification document (this document), TPCx-BB Users Guide documentation, shell scripts to set up the benchmark environment, Java code to execute the benchmark workload, Data Generator, **Query** files, and Benchmark Driver.*

*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-BB models and represents a Big Data Analytics System such as Hadoop ecosystem or Hadoop File-system API compatible systems);*
- *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 and rules for energy measurement are included in the TPC Energy Specification.*

*Further information is available at [www.tpc.org](http://www.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 Transwarp Technology (Shanghai) Co., Ltd

## 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 components used by the SUT.*
- *Configuration parameters and options for Operating System and file system components used by the SUT.*
- *Configuration parameters and options for any other software components (e.g compiler optimization options) used by the SUT.*

*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.*

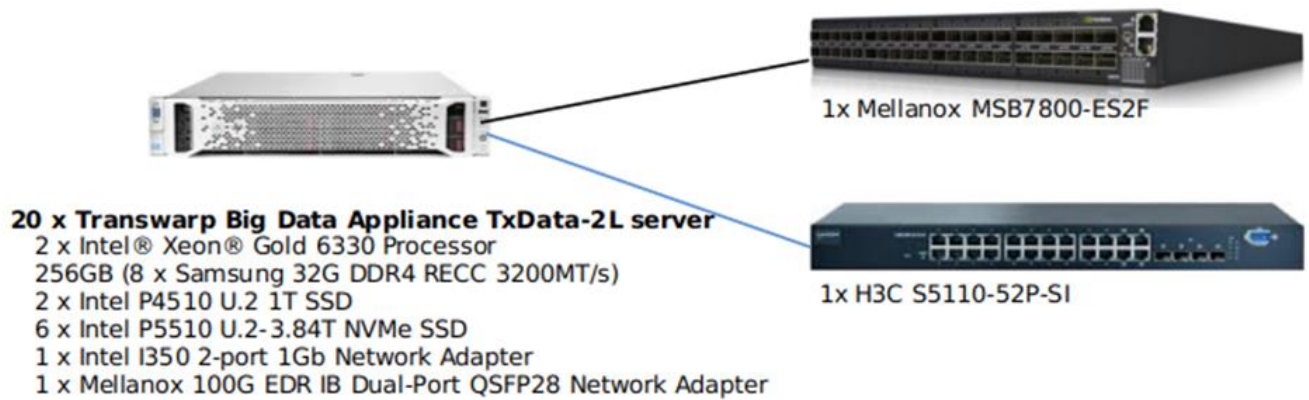
The Supporting Files Archive contains the parameters and options used to configure the components involved in this benchmark.

## 1.3 Configuration Diagrams

*7.4.4 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:	20
Total Processors/Cores/Threads:	40/1,120/2,240
Total Memory:	5,120
Total Number of Storage Devices:	160
Total Storage Capacity:	500,800

Network:	1x Mellanox MSB7800-ES2F Switch 1x H3C S5110-52P-SI Switch
----------	---

	<b>20x TxData-2L (Cluster Node):</b>
Processors/Cores/Threads:	2/56/112
Processor Model:	2x Intel(R) Xeon(R) Gold 6330 CPU @ 2.00GHz
Memory:	256 GiB
Storage Devices:	2x 1 TB NVMe SSD 6x 3.84 TB NVMe SSD
Network Controller:	Intel I350 2-port 1 Gb 1x Mellanox 100 G EDR IB Dual-Port QSFP28

The distribution of software components over server nodes is detailed in section 2.1.

## Priced Configuration

There are no differences between the priced and measured configurations.

# Clause 2: Software Components and Dataset Distribution

## 2.1 Roles and Dataset Distribution

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

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

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

**Table 1.4: Software Components and Dataset Distribution**

Server	Role(s)	Count	Virtual	Host Names	HW/SW Configuration	Storage Setup
Cluster Node	See Below	20	N	fw-perf[11-30]	<b>TxData-2L</b> 2x Intel(R) Xeon(R) Gold 6330 CPU @ 2.00GHz 256 GiB 2x 1 TB NVMe SSD 6x 3.84 TB NVMe SSD 1x Intel I350 2-port 1 Gb 1x Mellanox 100 G EDR IB Dual-Port QSFP28  CentOS Linux release 8 TDH 9.1 Spark 2.3.3	OS: 2x 1 TB NVMe SSD, RAID 1  6x 3.84 TB NVMe SSD Shuffle, Intermediate, Temp Data and Distributed File System

Node	Inceptor		HDFS		Yarn		Transwarp-manager	
	Driver	Worker	NN & JN	DN	RM	Node Manager	Master	Agent
fw-perf11		X		X		X	X	X
fw-perf[12-16]		X		X		X		X
fw-perf[17-18]		X	X	X		X		X
fw-perf[19-20]		X		X	X	X		X
fw-perf[21-29]		X		X		X		X
fw-perf30	X	X		X		X		X

## 2.2 Distributed File System Implementation

*Distributed file system implementation and corresponding Hadoop File System API version must be disclosed.*

Transwarp Data Hub 9.1 (fully HDFS compatible at the API level).

## 2.3 Engine Implementation

*The Engine implementation and corresponding version must be disclosed.*

Component	Version
HDFS	transwarp-9.0.0
YARN	transwarp-9.0.0
Inceptor	quark-8.35
Spark	2.3.3

## 2.4 Frameworks

*Frameworks and Engine used in the benchmark should be disclosed.*

Framework	Version
TDH	9.1
HDFS	transwarp-9.0.0
YARN	transwarp-9.0.0
Inceptor	quark-8.35
Spark	2.3.3

## 2.5 Applied Patches

*Any additional vendor supported patches applied to the SUT should be disclosed.*

No additional patches were applied.

# Clause 3: Workload Related Items

## 3.1 Hardware & Software Tunable

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

The Supporting Files Archive contains all configuration scripts.

## 3.2 Kit Version

*Version number of the TPCx-BB kit must be included in the Report.*

<b>TPCx-BB Kit Version</b>
----------------------------

v1.6.1
--------

## 3.3 Run Report

*The run report generated by TPCx-BB benchmark kit must be included in the Report.*

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

- **Run1 Report Summary (Repeatability Run)**

```
*****
TPCx-BB
Result
v1.6.1
*****
INFO: T_LOAD = 536.54
INFO: T_LD = 0.1 * T_LOAD: 53.6539999999999
INFO: T_PT = 1342.11765250356
INFO: T_T_PUT = 5613.349
INFO: T_TT = 701.668625
INFO: === Checking validity of the final result ===
INFO: OK: All required BigBench phases were performed.
INFO: OK: All 30 queries were running in the power test.
INFO: OK: All 30 queries were running in the first throughput test.
INFO: OK: Pretend mode was inactive. All commands were executed.
INFO: === Final result ===
INFO: VALID BBQpm@3000 = 5273.03820932679
```

- **Run2 Report Summary (Performance Run)**

```
*****
TPCx-BB
Result
v1.6.1
*****
INFO: T_LOAD = 520.326
INFO: T_LD = 0.1 * T_LOAD: 52.0326
INFO: T_PT = 1355.30495689559
INFO: T_T_PUT = 5674.263
INFO: T_TT = 709.282875
INFO: === Checking validity of the final result ===
INFO: OK: All required BigBench phases were performed.
INFO: OK: All 30 queries were running in the power test.
INFO: OK: All 30 queries were running in the first throughput test.
INFO: OK: Pretend mode was inactive. All commands were executed.
INFO: === Final result ===
INFO: VALID BBQpm@3000 = 5230.08033596877
```

### 3.4 Query Elapsed Times

*Elapsed times of all power and throughput Queries needs to be reported from the Performance Run, grouped respectively as Structured, semi-structured and unstructured buckets.*

Type	Query	Power	Stream 1	Stream 2	Stream 3	Stream 4	Stream 5	Stream 6	Stream 7	Stream 8
Structured	1	25.959	310.831	36.171	169.079	173.165	215.799	85.034	326.635	226.159
	6	43.719	120.073	60.738	230.771	59.863	80.376	132.292	233.058	59.870
	7	24.722	101.795	377.510	122.058	368.411	48.223	118.694	33.876	106.270
	9	25.683	58.551	55.633	194.491	62.708	24.726	59.919	35.089	54.015
	11	21.481	89.363	29.647	54.945	94.864	183.364	51.249	29.300	90.995
	13	27.431	66.319	95.280	35.261	40.133	156.122	48.368	59.312	44.358
	14	17.219	252.787	303.096	27.628	291.661	260.270	26.914	220.838	84.788
	15	19.712	71.777	102.758	51.017	332.613	23.397	432.327	22.414	49.262
	16	30.609	128.762	99.473	58.278	63.828	78.274	41.664	79.442	47.815
	17	25.963	120.388	30.503	39.313	29.312	179.758	50.419	392.179	34.959
	20	83.817	175.532	172.199	328.741	147.880	189.963	220.195	163.606	298.277
	21	33.243	223.836	239.299	87.596	217.218	47.042	207.840	85.865	83.494
	22	19.867	34.863	371.768	68.057	49.132	327.716	177.709	44.990	87.107
	23	38.050	217.596	67.943	162.428	95.980	301.429	75.157	175.880	164.004
	24	25.950	390.063	55.005	118.208	31.067	44.227	30.063	33.098	281.315
25	121.855	358.738	231.948	383.482	547.013	385.124	382.263	208.430	575.892	
26	89.482	190.531	441.816	198.941	148.759	333.042	137.152	168.903	145.815	
29	45.038	105.792	379.245	74.612	91.270	112.850	352.060	98.470	109.117	
Semi-structured	2	237.751	280.016	284.642	341.396	546.547	476.755	371.342	541.753	315.429
	3	45.435	103.289	302.949	86.842	437.242	237.974	176.220	85.500	172.375
	4	154.584	229.105	361.182	184.025	176.893	391.370	285.860	321.112	213.875
	5	116.433	207.233	252.768	202.861	217.547	532.778	347.192	160.178	198.332
	8	89.127	447.297	172.475	411.388	350.689	144.735	208.830	518.104	182.929
	12	37.217	83.816	79.303	319.170	113.572	56.046	136.334	160.330	570.329
	30	87.789	209.876	117.560	221.669	284.936	165.379	107.639	112.260	167.966
Unstructured	10	24.222	329.317	70.510	87.499	28.509	36.442	70.403	81.926	132.946
	18	88.879	367.041	303.457	159.919	316.726	147.415	212.894	129.849	176.006
	19	51.553	155.234	113.609	474.790	110.336	189.027	121.817	114.068	108.382
	27	25.181	30.754	45.521	51.657	69.614	38.552	157.539	40.505	67.205
	28	91.296	164.948	236.341	221.222	152.180	266.052	363.751	132.486	275.248

### 3.5 Validation Test Output

*Output report from successful SUT Validation test must be included in the Report.*

Query Number	Query Execution	Output Validation
1	PASS	PASS
2	PASS	PASS
3	PASS	PASS
4	PASS	PASS
5	PASS	PASS
6	PASS	PASS
7	PASS	PASS
8	PASS	PASS
9	PASS	PASS
10	PASS	PASS
11	PASS	PASS
12	PASS	PASS
13	PASS	PASS
14	PASS	PASS
15	PASS	PASS
16	PASS	PASS
17	PASS	PASS
18	PASS	PASS
19	PASS	PASS
20	PASS	PASS
21	PASS	PASS
22	PASS	PASS
23	PASS	PASS
24	PASS	PASS
25	PASS	PASS
26	PASS	PASS
27	PASS	PASS
28	PASS	PASS
29	PASS	PASS
30	PASS	PASS

### 3.6 Global Framework Parameters

*Global Framework parameter settings files must be included in the Report.*

The Supporting File Archive contains the global framework parameter settings files.

### 3.7 Kit Modifications

*Test Sponsor kit modifications files must be included in the Report.*

The following files were modified by the Test Sponsor to facilitate system, platform and Framework differences.

- bigBench-configs/conf/userSettings.conf



# Clause 4: SUT Related Items

## 4.1 Specialized Hardware/Software

*Specialized Hardware/Software used in the SUT must be included.*

No specialized hardware or software was used.

## 4.2 Framework Configuration Files

*All Framework configuration files from SUT, for the performance run.*

All Framework configuration files are included in the Supporting Files Archive.

## 4.3 SUT Environment Information

*SUT environment info in form of envinfo.log from a representative worker node from every role in the server.*

All envinfo.log files are included in the Supporting Files Archive.

## 4.4 Data Storage to Scale Factor Ratio

*The data storage ratio must be disclosed.*

Nodes	Disks	Size (GB)	Total (GB)
20	2	1,000	40,000
20	6	3,840	460,800

Total Storage (GB) 500,800

Scale Factor 3000

Data Storage Ratio 166.93

## 4.5 Scale Factor to Memory Ratio

*The Scale Factor to memory ratio must be disclosed.*

Nodes	Memory (GB)	Total (GB)
20	256	5,120

Scale Factor 3000

Total Memory (GB) 5,120

SF / Memory Ratio 0.59

# Clause 5: Metrics and Scale Factors

## 5.1 Performance Run Metric

*The Reported Performance Metric (BBQpm@SF for the Performance Run) must be disclosed in the Report.*

Performance Run
BBQpm@3000 5,230.08

## 5.2 Repeatability Run Metric

*The Performance Metric (BBQpm@SF) for the Repeatability Run must be disclosed in the Report..*

Repeatability Run
BBQpm@3000 5,273.03

## 5.3 Price-Performance Metric

*The Reported Performance Metric (BBQpm@SF for the Performance Run) must be disclosed in the Report.*

Price / Performance
¥BBQpm@3000 719.11

## 5.4 Scale Factor

*The Scale Factor used for the Result must be disclosed in the Report.*

Scale Factor
3000

## 5.5 Stream Count

*The number of streams in the throughput run used for the Result must be disclosed in the Report.*

Streams
8

## 5.6 Elapsed Run Times

*The total elapsed time for the execution of the Performance Run and Repeatability Run must be disclosed in the Report.*

<b>Run</b>	<b>Elapsed Time</b>	<b>Seconds</b>
Run 1	00 02:11:24.285	7,884.285
Run 2	00 02:12:43.955	7,963.955

## 5.7 Elapsed Test Times

*The total time for each of the three tests must be disclosed for the Performance Run and the Repeatability Run.*

<b>Test</b>	<b>Performance Run</b>	<b>Repeatability Run</b>
Load Test	520.326	536.540
Power Test	1,769.338	1,734.393
Throughput Test	5,674.263	5,613.349

# Auditors' Information and Attestation Letter

*The auditor's agency name, address, phone number, and Attestation letter must be included in the full disclosure report. A statement should be included specifying who to contact in order to obtain further information regarding the audit process.*

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 (FDR) 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.

Jun Zheng  
Transwarp Technology (Shanghai) Co., Ltd.  
Floor 11 & 12, Block B, No. 88 Hongcao Road  
Xuhui District, Shanghai  
China

September 20, 2023

I verified the TPC Express Benchmark™ BB v1.6.1 performance of the following configuration:

Platform: Transwarp Big Data Appliance (with 20x TxData-2L Servers)  
Operating System: CentOS Linux release 8  
Apache Hadoop: Transwarp Data Hub 9.1  
Compatible Software:

The results were:

**Performance Metric** **5,230.08 BBQpm@3000GB**  
Run Elapsed Time 00 02:12:43.955 (7,963.955 Seconds)

<b>Cluster</b>	<b>20x TxData-2L Servers, each with:</b>		
CPUs	2x Intel® Xeon® Gold 6330 Processors (3.10 GHz, 28-Core)		
Memory	256 GB		
Storage	<b>Qty</b>	<b>Size</b>	<b>Type</b>
	2	1 TB	NVMe SSD (RAID1)
	6	3.84 TB	NVMe SSD

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

The following verification items were given special attention:

- All TPC-provided components were verified to be v1.6.1.
- No modifications were made to any of the Java code.
- Any and all modifications to shell scripts were reviewed for compliance.
- The tested Scale Factor (3000GB) was confirmed to be valid for publication.
- All validation queries executed successfully and produced compliant results.

- No errors were reported during the run.
- The elapsed times for all phases and runs were correctly measured and reported.
- The Storage and Memory Ratios were correctly calculated and reported.
- The system pricing was verified for major components and maintenance.
- The major pages from the FDR were verified for accuracy.

Additional Audit Notes:

None.

Respectfully Yours,

A handwritten signature in cursive script that reads "Doug Johnson". The signature is written in black ink and includes a long horizontal flourish extending to the right.

Doug Johnson, TPC Auditor

63 Lourdes Dr. | Leominster, MA 01453 | 978-343-6562 | [www.sizing.com](http://www.sizing.com)

# Third Party Price Quotes

All components are available directly through the Test Sponsor (Transwarp Technology (Shanghai) Co., Ltd).

# Supporting File Index

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

Description	Archive File Pathname
<b>Clause 1 - General Items</b>	
The Supporting Files Archive contains the parameters and options used to configure the components involved in this benchmark	Supporting-Files-3TB-Transwarp-TDH-09-2023\
Validation Run Files	Supporting-Files-3TB-Transwarp-TDH-09-2023\Validation-Run logs
Performance Run Files	Supporting-Files-3TB-Transwarp-TDH-09-2023\Performance-run logs
Repeatability Run Files	Supporting-Files-3TB-Transwarp-TDH-09-2023\Repeatability-run logs
<b>Clause 3 - Workload Related Items</b>	
Benchmark Generic Parameters	Supporting-Files-3TB-Transwarp-TDH-09-2023\Performance-run logs\bigBench-configs\conf\userSettings.conf
Query Parameters used in the benchmark execution Settings	Supporting-Files-3TB-Transwarp-TDH-09-2023\Performance-run logs\bigBench-configs\inceptor\conf\queryParameters.sql
Benchmark Global Framework Parameters Settings	Supporting-Files-3TB-Transwarp-TDH-09-2023\Performance-run logs\bigBench-configs\inceptor\conf\engineSettings.sql
Benchmark Global Framework Parameters Settings	Supporting-Files-3TB-Transwarp-TDH-09-2023\Performance-run logs\bigBench-configs\inceptor\conf\engineSettings.conf
Load Test script	Supporting-Files-3TB-Transwarp-TDH-09-2023\Performance-run logs\bigBench-configs\inceptor\population\hiveCreateLoad.sql
Queries specific optimization parameters settings	Supporting-Files-3TB-Transwarp-TDH-09-2023\Performance-run logs\bigBench-configs\inceptor\queries\q[01-30]\engineLocalSettings.conf
Queries specific optimization parameters settings	Supporting-Files-3TB-Transwarp-TDH-09-2023\Performance-run logs\bigBench-configs\inceptor\queries\q[01-30]\engineLocalSettings.sql
<b>Clause 4 - SUT Related Items</b>	
Data Redundancy report	Supporting-Files-3TB-Transwarp-TDH-09-2023\Performance-run logs\run-logs\data_redundancy_report.log
Benchmark execution script	Supporting-Files-3TB-Transwarp-TDH-09-2023\TPCxBB_FullBenchmark_sequence_run.sh
Hardware and Software Report from a representative node	Supporting-Files-3TB-Transwarp-TDH-09-2023\Performance-run logs\envInfo-fw-perf29\envInfo.log
All Framework configuration files are included in the Supporting Files Archive	Supporting-Files-3TB-Transwarp-TDH-09-2023\Performance-run logs\envInfo-fw-perf29\hdfs1
	Supporting-Files-3TB-Transwarp-TDH-09-2023\Performance-run logs\envInfo-fw-perf29\quark1
	Supporting-Files-3TB-Transwarp-TDH-09-2023\spark-conf
<b>Clause 5 - Metric and Scale Factor Related Items</b>	
Benchmark Performance Report	Supporting-Files-3TB-Transwarp-TDH-09-2023\Performance-run logs\run-logs\BigBenchResult.log
Validation Test Report	Supporting-Files-3TB-Transwarp-TDH-09-2023\Validation-Run logs\run-logs\BigBenchResult.log