



Super Micro Computer, Inc.

TPC Benchmark™ DS
Full Disclosure Report
for
Supermicro A+ Server 2123BT-HNC0R
using
Transwarp ArgoDB v1.2.1
and
Red Hat Enterprise Linux Server Release 7.6

First Edition

07-Aug-2019

First Edition – 07-Aug-2019

Super Micro Computer, Inc (Supermicro), a global leader in high performance, high efficiency server technology and innovation, we develop and provide end-to-end green computing solutions to the data center, cloud computing, enterprise IT, big data, HPC, and embedded markets. Our Building Block Solutions® approach allows us to provide a broad range of SKUs, and enables us to build and deliver application-optimized solutions based upon your requirements.

Table of Contents

Abstract.....	5
Preface.....	11
General Items	12
0.1 Test Sponsor	12
0.2 Parameter Settings.....	12
0.3 Configuration Diagrams	12
Clause 2: Logical Database Design Related Items.....	15
2.1 Database Definition Statements.....	15
2.2 Physical Organization.....	15
2.3 Horizontal Partitioning	15
2.4 Replication.....	15
Clause 3: Scaling and Database Population.....	16
3.1 Initial Cardinality of Tables.....	16
3.2 Distribution of Tables and Logs Across Media.....	16
3.3 Mapping of Database Partitions/Replications.....	18
3.4 Implementation of RAID	18
3.5 DBGEN Modifications	18
3.6 Database Load time.....	18
3.7 Data Storage Ratio	18
3.8 Database Load Mechanism Details and Illustration	18
3.9 Qualification Database Configuration.....	19
Clause 4 and 5: Query and Data Maintenance Related Items.....	20
4.1 Query Language.....	20
4.2 Verifying Method of Random Number Generation	20
4.3 Generating Values for Substitution Parameters	20
4.4 Query Text and Output Data from Qualification Database	20
4.5 Query Substitution Parameters and Seeds Used	20
4.6 Refresh Setting	21
4.7 Source Code of Refresh Functions.....	21
4.8 staging Area	21

Clause 6: Data Persistence Properties Related Items	22
Clause 7: Performance Metrics and Execution Rules Related Items	23
7.1 System Activity	23
7.2 Test Steps	23
7.3 Timing Intervals for Each Query and Refresh Function.....	23
7.4 Throughput Test Result	23
7.5 Time for Each Stream.....	23
7.6 Time for Each Refresh Function	23
7.7 Performance Metrics.....	23
Clause 8: SUT and Driver Implementation Related Items	24
8.1 Driver.....	24
8.2 Implementation Specific Layer (ISL)	24
8.3 Profile-Directed Optimization	24
Clause 9: Pricing Related Items	25
9.1 Hardware and Software Used	25
9.2 Availability Date	25
9.3 Country-Specific Pricing	25
Clause 11: Audit Related Items.....	26
Supporting Files Index	30
Appendix A: Price Quotes.....	31
Third party Price Quotes	31
Mellanox SN2410 Switch.....	31
Miscellaneous Items	32
Transwarp ArgoDB.....	37

Abstract

This document contains the methodology and results of the TPC Benchmark™ DS (TPC-DS) test conducted in conformance with the requirements of the TPC-DS Standard Specification, Revision 2.11.0.

The test was conducted at a Scale Factor of 10000GB with 16 Supermicro A+ Server 2123BT-HNC0R Servers running Transwarp ArgoDB v1.2.1 on Red Hat Enterprise Linux Server Release 7.6.

Measured Configuration

Company Name	Cluster Node	Database Software	Operating System
Super Micro Computer, Inc.	Supermicro A+ Server 2123BT-HNC0R	Transwarp ArgoDB v1.2.1	Red Hat Enterprise Linux Server Release 7.6

TPC Benchmark™ DS Metrics

Total System Cost	TPC-DS Throughput	Price/Performance	Availability Date
\$487,232 USD	4,418,054	\$0.12 USD	31-Aug-2019



Supermicro A+ Server 2123BT-HNC0R

TPC-DS 2.11.0
TPC-Pricing 2.4.0

Report Date:
07-Aug-2019

Total System Cost

Composite Query per Hour Metric

Price / Performance

**\$487,232
USD**

4,418,054
QphDS@10000GB

\$0.12 USD
\$ / QphDS@10000GB

Database Size ¹

Database Manager

Operating System

Other Software

Availability date

Clustered Or Not

10,000GB

Transwarp
ArgoDB v1.2.1

Red Hat Enterprise
Linux Server
Release 7.6

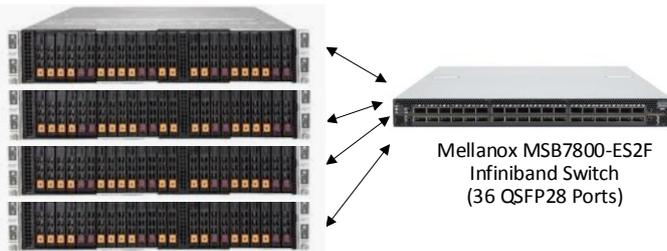
No

31-Aug-2019

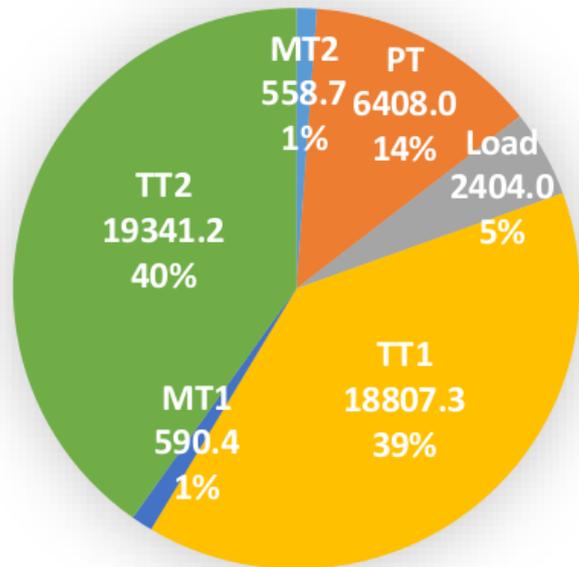
Yes

4x Supermicro A+ Server AS-2123BT-HNC0R

BigTwin™ with 4x 217BH-22H12 Nodes, each with:
2x AMD EPYC 7502 32-Core Processor
256 GB (16x16GB RDIMM 2666MT/s Dual Rank)
1x 960GB M.2 PCIe Gen3
2x 3.5TB NVMe Drive
2x 2TB NVMe Drive
1x Mellanox MT27800 ConnectX-5 25GbE SFP28 NIC



Benchmarked Configuration



Elapsed Time

Load includes backup = No

RAID = No

System Configuration:

Supermicro A+ Server Cluster

Servers:

16 x Supermicro A+ Server 2123BT-HNC0R

Total Processors/Cores/Threads:

16/1024/2048

Total Memory:

4096GB

Total Storage²:

195,584 GB

Storage Ratio³:

19.56

Connectivity:

1 x Mellanox MSB7800-ES2F Switch-IB 2 Based EDR Infiniband 1U Switch 36 QSFP28 Ports

Server Configuration:

Per node

Processors:

2 x AMD EPYC 7502 32-Core Processor

Memory:

256GB

Network:

1 x Mellanox MT27800 Family [ConnectX-5] Network Adapter 25GbE

Storage Device:

2 x 3.5 TB NVMe Drive(data nodes and tablet servers)
2 x 2 TB NVMe Drive (data nodes and tablet servers)
1 x 960 GB M.2 PCIe Gen3 (all nodes, boot disk)

¹ Database Size includes only raw data (i.e., no temp, index, redundant storage space, etc.).

² Total Storage = 16 x DataNode & TabletServer + 16 x BootDisk (i.e. 16 x (2 x 3.5 x 1024 + 2 x 2 x 1024 + 1 x 960) GB)

³ Storage Ratio = TotalStorage / SF (i.e. 195,584 GB / 10000 GB)



Supermicro A+ Server 2123BT-HNC0R

TPC-DS 2.11.0
TPC-Pricing 2.4.0

Report Date:
07-Aug-2019

Description	Part Number	Src	Unit Price	Qty	Extended Price	3 Year Maint. Price
HARDWARE COMPONENTS						
Supermicro Model 217BH-22H12 - Rail, Power Cord (included)	AS-2123BT-HNC0R		\$4,778.00	4	\$19,112.00	
AMD EPYC™ 7502 2.5GHz, 32C/64T 180W		1	\$1,850.00	32	\$59,200.00	
Hynix 16GB PC4-2666 HMA82GR7CJR8N-VK	HMA82GR7CJR8N-VK	1	\$67.90	256	\$17,382.40	
M.2 Model : MZ-1LW9600 (SAMSUNG MZ1LW960HMJP-00003)	MZ1LW960HMJP-00003	1	\$145.00	16	\$2,320.00	
Micron_9200_MTFDHAL3T8TCT 3.5TB	MTFDHAL3T8TCT	1	\$272.50	32	\$8,720.00	
Intel DC P4510 2TB NVMe PCI-E 3.1 x4 3D TLC 2.5" 15mm 1DWPD	HDS-IUN2-SSDPE2KX020T8	1,S	\$331.50	32	\$10,608.00	
Mellanox MT27800 Family [ConnectX-5] Network Adapter 25GbE	AOC-MH25G-m2S2TM	1	\$249.90	16	\$3,998.40	
Supermicro 4-hour 7x24 On-site Service 3years	OS4HR3	1	\$1,310.92	4		\$5,243.68
Mellanox MSB7800-ES2F Switch-IB 2 Based EDR Infiniband 1U Switch 36 QSFP28 Ports	MPN: MSB7800-ES2F	2	\$25,633.00	1	\$25,633.00	
Mellanox SUP-SB7800-3S Technical Support and Warranty - Silver 3 Year for SB7800 Series Switch	SUP-SB7800-3S	2	\$2,815.00	1	\$2,815.00	
Mellanox MCP2M00-A001E30N Ethernet Passive Copper Cable 25GbE SFP28 1m Black 30AWG CA-N	MCP2M00-A001E30N	2	\$52.00	20	\$1,040.00	
12U Open Frame Server Rack - Adjustable Depth - 4-Post Data Rack	4POSTRACK12U	2	\$195.99	1	\$195.99	
NEMA 5-15P to IEC320 C15 14AWG 125V/15A	P515C15-15ABE-10	2	\$10.00	4	\$40.00	
24" LED Monitor		2	\$79.00	1	\$79.00	
Keyboard and Mouse		2	\$14.56	1	\$14.56	
Subtotal					\$151,158.35	\$5,243.68

HARDWARE COMPONENTS Subtotal \$151,158.35 \$5,243.68

Description	Part Number	Src	Unit Price	Qty	Extended Price	3 Year Maint. Price
SOFTWARE COMPONENTS						
Transwarp high performance Data Mart for Online Analysis Version 1.2.1 with 3years 24x7 support		3	\$17,400.00	16	\$278,400.00	
Red Hat Enterprise Linux 7.6, 3yr Premium Subscription		1	\$3,276.86	16		\$52,429.76
SOFTWARE COMPONENTS					Subtotal \$278,400.00	\$52,429.76
Total					\$429,558.35	\$57,673.44

Pricing : 1 Supermicro (based upon total system cost as purchased by a regular customer). Pricing : 2 Obtained from third-party or online. Pricing : 3 Transwarp. Pricing : S One or more components of the Measured Configuration have been substituted in the Priced Configuration. See FDR for details.	Three-Year Cost of Ownership:	\$487,232
	QphDS@10000GB:	4,418,054
	\$ / QphDS@10000GB:	\$0.12 USD
Audited by Francois Raab, InfoSizing		

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 pricing@tpc.org. Thank you.



Supermicro A+ Server 2123BT-HNC0R

TPC-DS 2.11.0
TPC-Pricing 2.4.0

Report Date:
07-Aug-2019

Metrics Details:

Name	Value	Unit
Scale Factor (SF)	10,000	GB
Streams	4	Stream
Queries (Q)	396	Queries
T_load	2,404.0	Second
T_ld	0.0268	Hour
T_power	6,408.0	Second
T_pt	7.1200	Hour
T_tt1	18,807.3	Second
T_tt2	19,341.2	Second
T_dm1	590.4	Second
T_dm2	558.7	Second
T_tt	10.5969	Hour
T_dm	0.3192	Hour

Load Step	Start	End	(sec.)	(hh:mm:ss)
Build	07/25/19 19:01:39.	07/25/19 19:38:29.	2,209.58	0:36:50
Audit	07/25/19 19:38:29.	07/25/19 20:02:19.	1,430.45	0:23:50
Get Stats	07/25/19 20:02:19.	07/25/19 20:05:33.	194.33	0:03:14
Reported	07/25/19 19:01:39.	07/25/19 20:05:33.	2,403.91	0:40:04

Test	Start	End	(sec.)	(hh:mm:ss)
Power	07/25/19 20:14:36.	07/25/19 22:01:24.	6,407.98	1:46:48
Thruput-1	07/25/19 22:01:24.	07/26/19 03:14:52.	18,807.21	5:13:27
Thruput-2	07/26/19 03:24:42.	07/26/19 08:47:03.	19,341.12	5:22:21
DM-1	07/26/19 03:14:52.	07/26/19 03:24:42.	590.35	0:09:50
DM-2	07/26/19 08:47:03.	07/26/19 08:56:22.	558.70	0:09:19

Stream	Start	End	(sec.)	(hh:mm:ss)
Pt - 0	07/25/19 20:14:36.	07/25/19 22:01:24.	6,407.98	1:46:48
Tt1 - 1	07/25/19 22:01:24.	07/26/19 03:13:35.	18,730.34	5:12:10
Tt1 - 2	07/25/19 22:01:24.	07/26/19 03:14:52.	18,807.20	5:13:27
Tt1 - 3	07/25/19 22:01:24.	07/26/19 03:05:30.	18,245.34	5:04:05
Tt1 - 4	07/25/19 22:01:24.	07/26/19 03:10:22.	18,537.68	5:08:58
Tt2 - 5	07/26/19 03:24:42.	07/26/19 08:44:58.	19,216.49	5:20:16
Tt2 - 6	07/26/19 03:24:42.	07/26/19 08:47:03.	19,341.12	5:22:21
Tt2 - 7	07/26/19 03:24:42.	07/26/19 08:43:28.	19,125.82	5:18:46
Tt2 - 8	07/26/19 03:24:42.	07/26/19 08:43:31.	19,129.18	5:18:49
DMt1 - 1	07/26/19 03:14:52.	07/26/19 03:20:02.	310.85	0:05:11
DMt1 - 2	07/26/19 03:20:02.	07/26/19 03:24:42.	279.49	0:04:39
DMt2 - 3	07/26/19 08:47:03.	07/26/19 08:51:50.	287.07	0:04:47
DMt2 - 4	07/26/19 08:51:50.	07/26/19 08:56:22.	271.62	0:04:32

Timing Intervals for Each Query (In Seconds)

Query	Stream0	Stream1	Stream2	Stream3	Stream4	Min	25%tile	Median	75%tile	Max	Stream5	Stream6	Stream7	Stream8	Min	25%tile	Median	75%tile	Max
1	18.8	368.4	65.1	104.1	56.4	56.4	62.9	84.6	170.2	368.4	46.7	84.1	303.0	54.4	46.7	52.5	69.3	138.8	303.0
2	29.7	48.1	48.5	418.7	590.4	48.1	48.4	233.6	461.6	590.4	72.3	133.0	83.7	241.4	72.3	80.9	108.4	160.1	241.4
3	34.0	67.9	135.8	98.4	49.0	49.0	63.2	83.2	107.8	135.8	44.7	86.1	190.0	178.7	44.7	75.8	132.4	181.5	190.0
4	336.4	395.1	601.4	396.1	545.5	395.1	395.9	470.8	559.5	601.4	562.0	448.6	429.8	454.8	429.8	443.9	451.7	481.6	562.0
5	40.5	58.4	63.0	72.1	81.8	58.4	61.9	67.6	74.5	81.8	141.6	63.4	206.8	77.3	63.4	73.8	109.5	157.9	206.8
6	14.1	128.4	19.8	50.5	40.2	19.8	35.1	45.4	70.0	128.4	23.9	19.5	86.4	118.6	19.5	22.8	55.2	94.5	118.6
7	27.8	60.7	38.1	135.2	109.3	38.1	55.1	85.0	115.8	135.2	145.3	32.0	46.2	50.8	32.0	42.7	48.5	74.4	145.3
8	7.5	52.1	339.2	81.5	59.7	52.1	57.8	70.6	145.9	339.2	27.1	210.3	332.1	138.3	27.1	110.5	174.3	240.8	332.1
9	95.8	265.9	185.6	495.4	159.2	159.2	179.0	225.8	323.3	495.4	430.9	563.9	314.7	538.0	314.7	401.9	484.5	544.5	563.9
10	28.7	93.8	145.9	67.8	75.7	67.8	73.7	84.8	106.8	145.9	30.3	135.8	789.2	79.1	30.3	66.9	107.5	299.2	789.2
11	200.3	418.0	230.6	450.6	368.2	230.6	333.8	393.1	426.2	450.6	495.5	249.1	539.1	280.1	249.1	272.4	387.8	506.4	539.1
12	3.2	143.6	8.7	17.0	56.3	8.7	14.9	36.7	78.1	143.6	8.5	60.0	51.7	11.5	8.5	10.8	31.6	53.8	60.0
13	143.6	214.2	340.2	130.2	360.0	130.2	193.2	277.2	345.2	360.0	200.8	152.4	224.2	446.7	152.4	188.7	212.5	279.8	446.7
14	206.2	676.8	722.7	1,110.8	1,010.8	676.8	711.2	866.8	1,035.8	1,110.8	1,158.9	831.3	964.3	660.2	660.2	788.5	897.8	1,013.0	1,158.9
15	9.0	13.8	326.2	16.5	73.6	13.8	15.8	45.1	136.8	326.2	123.9	15.2	17.1	47.2	15.2	16.6	32.2	66.4	123.9
16	74.1	168.1	351.7	146.1	239.8	146.1	162.6	204.0	267.8	351.7	99.9	122.9	87.0	91.8	87.0	90.6	95.9	105.7	122.9
17	50.2	243.0	174.2	110.5	191.9	110.5	158.3	183.1	204.7	243.0	76.9	89.2	119.0	295.3	76.9	86.1	104.1	163.1	295.3
18	24.3	355.8	91.4	38.3	67.6	38.3	60.3	79.5	157.5	355.8	202.6	65.8	612.5	102.9	65.8	93.6	152.8	305.1	612.5
19	6.1	71.9	23.3	152.6	140.3	23.3	59.8	106.1	143.4	152.6	21.3	358.4	28.3	98.3	21.3	26.6	63.3	163.3	358.4
20	6.2	65.5	5.8	15.0	325.5	5.8	12.7	40.3	130.5	325.5	169.5	15.7	229.4	16.1	15.7	16.0	92.8	184.5	229.4
21	4.2	8.1	15.0	6.9	10.3	6.9	7.8	9.2	11.5	15.0	158.2	291.6	29.5	41.8	29.5	38.7	100.0	191.6	291.6
22	23.3	42.6	67.5	333.2	123.3	42.6	61.3	95.4	175.8	333.2	44.8	36.4	31.3	162.3	31.3	35.1	40.6	74.2	162.3
23	456.8	940.3	804.7	995.2	1,010.8	804.7	906.4	967.8	999.1	1,010.8	1,118.1	882.8	830.3	1,255.0	830.3	869.7	1,000.5	1,152.3	1,255.0
24	143.9	513.1	364.0	573.1	173.1	173.1	316.3	438.6	528.1	573.1	421.6	834.8	263.6	479.4	263.6	382.1	450.5	568.3	834.8
25	22.5	62.6	76.3	80.1	124.6	62.6	72.9	78.2	91.2	124.6	24.5	67.2	125.9	83.4	24.5	56.5	75.3	94.0	125.9
26	12.6	130.1	140.7	32.5	21.0	21.0	29.6	81.3	132.8	140.7	15.8	167.4	48.2	55.2	15.8	40.1	51.7	83.3	167.4
27	25.3	155.8	231.6	275.4	366.1	155.8	212.7	253.5	298.1	366.1	68.3	75.1	45.5	104.7	45.5	62.6	71.7	82.5	104.7
28	132.4	276.4	176.0	202.5	176.2	176.0	176.2	189.4	221.0	276.4	425.3	178.7	218.5	280.6	178.7	208.6	249.6	316.8	425.3
29	60.2	316.6	90.7	103.6	148.4	90.7	100.4	126.0	190.5	316.6	83.6	157.4	113.9	169.6	83.6	106.3	135.7	160.5	169.6
30	24.4	179.1	39.3	42.4	69.7	39.3	41.6	56.1	97.1	179.1	183.1	417.8	45.3	367.5	45.3	148.7	275.3	380.1	417.8
31	40.8	427.9	453.8	166.0	397.0	166.0	339.3	412.5	434.4	453.8	193.3	357.6	90.1	184.6	90.1	161.0	189.0	234.4	357.6
32	14.0	58.1	157.0	90.2	161.9	58.1	82.2	123.6	158.2	161.9	70.7	27.7	34.4	40.5	27.7	32.7	37.5	48.1	70.7
33	73.1	393.8	19.6	343.6	21.2	19.6	20.8	182.4	356.2	393.8	35.0	237.7	32.7	50.7	32.7	34.4	42.9	97.5	237.7
34	28.6	77.8	141.1	51.2	53.1	51.2	52.6	65.5	93.6	141.1	49.9	110.5	55.5	620.4	49.9	54.1	83.0	238.0	620.4
35	28.8	164.3	428.1	109.4	135.9	109.4	129.3	150.1	230.3	428.1	124.6	38.9	249.8	221.4	38.9	103.2	173.0	228.5	249.8
36	63.7	114.6	355.4	49.3	241.7	49.3	98.3	178.2	270.1	355.4	227.1	68.5	182.2	87.3	68.5	82.6	134.8	193.4	227.1
37	5.4	73.0	106.6	23.3	30.9	23.3	29.0	52.0	81.4	106.6	305.4	18.6	119.1	84.7	18.6	68.2	101.9	165.7	305.4
38	46.0	57.3	143.8	277.7	229.2	57.3	122.2	186.5	241.3	277.7	160.4	173.9	321.0	77.7	77.7	139.7	167.2	210.7	321.0
39	38.8	321.9	180.0	140.5	260.2	140.5	170.1	220.1	275.6	321.9	188.4	139.6	271.5	229.7	139.6	176.2	209.1	240.2	321.9
40	9.5	30.5	280.4	344.6	76.3	30.5	64.9	178.4	296.5	344.6	14.9	87.3	67.3	32.1	14.9	27.8	49.7	72.3	344.6
41	5.1	10.7	21.7	17.3	11.8	10.7	11.5	14.6	18.4	21.7	84.6	18.0	40.9	8.0	8.0	15.5	29.5	51.8	84.6
42	3.5	14.2	25.6	13.3	41.2	13.3	14.0	19.9	29.5	41.2	6.8	163.8	96.0	7.6	6.8	7.4	51.8	113.0	163.8
43	43.7	126.4	95.6	116.6	29.4	29.4	79.1	106.1	119.1	126.4	47.4	64.3	79.4	365.3	47.4	60.1	71.9	150.9	365.3

44	44.3	382.8	502.1	121.4	380.2	121.4	315.5	381.5	412.6	502.1	152.2	84.0	58.3	165.3	58.3	77.6	118.1	155.5	165.3
45	15.1	43.3	114.3	42.6	184.6	42.6	43.1	78.8	131.9	184.6	136.0	78.1	27.2	142.5	27.2	65.4	107.1	137.6	142.5
46	99.4	82.8	45.6	61.6	273.8	45.6	57.6	72.2	130.6	273.8	55.0	66.0	194.6	50.6	50.6	53.9	60.5	98.2	194.6
47	113.3	502.7	147.7	242.4	145.7	145.7	147.2	195.1	307.5	502.7	676.0	120.0	454.8	399.3	120.0	329.5	427.1	510.1	676.0
48	87.3	432.1	95.9	213.8	193.5	95.9	169.1	203.7	268.4	432.1	341.0	162.0	115.8	196.2	115.8	150.5	179.1	232.4	341.0
49	25.7	75.2	53.9	46.1	101.7	46.1	52.0	64.6	81.8	101.7	53.4	73.6	295.6	55.9	53.4	55.3	64.8	129.1	295.6
50	83.1	164.6	143.5	127.6	93.0	93.0	119.0	135.6	148.8	164.6	774.6	122.8	89.3	341.2	89.3	114.4	232.0	449.6	774.6
51	57.4	122.8	325.4	161.6	247.3	122.8	151.9	204.5	266.8	325.4	87.8	108.8	146.8	107.4	87.8	102.5	108.1	118.3	146.8
52	3.3	7.8	3.3	5.1	98.4	3.3	4.7	6.5	30.5	98.4	97.7	102.6	12.6	79.5	12.6	62.8	88.6	98.9	102.6
53	11.6	195.8	109.6	174.0	168.7	109.6	153.9	171.4	179.5	195.8	72.8	23.7	20.0	116.6	20.0	22.8	48.3	83.8	116.6
54	24.1	155.2	239.0	123.2	128.2	123.2	127.0	141.7	176.2	239.0	75.0	33.1	249.6	249.7	33.1	64.5	162.3	249.6	249.7
55	2.8	16.0	4.8	6.3	16.9	4.8	5.9	11.2	16.2	16.9	287.0	6.9	52.5	83.3	6.9	41.1	67.9	134.2	287.0
56	9.1	53.3	64.3	47.8	151.9	47.8	51.9	58.8	86.2	151.9	27.1	53.1	104.9	47.7	27.1	42.6	50.4	66.1	104.9
57	75.8	385.3	151.5	433.3	170.8	151.5	166.0	278.1	397.3	433.3	583.5	220.7	122.2	128.3	122.2	126.8	174.5	311.4	583.5
58	21.0	90.6	229.6	183.7	155.0	90.6	138.9	169.4	195.2	229.6	148.8	113.9	226.2	235.4	113.9	140.1	187.5	228.5	235.4
59	50.5	95.1	56.1	208.6	143.5	56.1	85.4	119.3	159.8	208.6	433.5	192.7	101.3	185.5	101.3	164.5	189.1	252.9	433.5
60	14.5	26.9	238.3	105.6	51.9	26.9	45.7	78.8	138.8	238.3	26.3	13.1	51.7	61.9	13.1	23.0	39.0	54.3	61.9
61	8.1	251.9	49.8	75.4	155.6	49.8	69.0	115.5	179.7	251.9	375.3	152.2	17.7	107.0	17.7	84.7	129.6	208.0	375.3
62	22.0	49.6	96.6	58.2	61.1	49.6	56.1	59.7	70.0	96.6	76.0	43.2	105.7	339.9	43.2	67.8	90.9	164.3	339.9
63	10.6	61.9	40.1	69.3	21.5	21.5	35.5	51.0	63.8	69.3	75.8	99.8	70.7	91.5	70.7	74.5	83.7	93.6	99.8
64	212.8	412.1	504.5	664.3	444.2	412.1	436.2	474.4	544.5	664.3	646.4	801.0	481.3	545.1	481.3	529.2	595.8	685.1	801.0
65	59.9	196.4	130.0	124.6	173.7	124.6	128.7	151.9	179.4	196.4	125.6	115.0	139.4	307.5	115.0	123.0	132.5	181.4	307.5
66	19.6	29.5	27.4	107.9	29.6	27.4	29.0	29.6	49.2	107.9	28.2	31.5	144.4	26.7	26.7	27.8	29.9	59.7	144.4
67	338.1	645.1	577.9	498.7	453.6	453.6	487.4	538.3	594.7	645.1	476.2	389.8	620.9	530.0	389.8	454.6	503.1	552.7	620.9
68	49.4	142.8	63.5	93.7	70.1	63.5	68.5	81.9	106.0	142.8	148.8	79.3	437.2	179.7	79.3	131.4	164.3	244.1	437.2
69	39.1	264.0	320.8	67.8	213.5	67.8	177.1	238.8	278.2	320.8	164.8	43.3	34.6	50.8	34.6	41.1	47.1	79.3	164.8
70	45.9	529.5	412.0	344.4	381.9	344.4	372.5	397.0	441.4	529.5	388.4	120.3	118.0	375.0	118.0	119.7	247.7	378.4	388.4
71	19.5	55.0	57.3	85.8	38.2	38.2	50.8	56.2	64.4	85.8	26.3	414.0	221.2	43.3	26.3	39.1	132.3	269.4	414.0
72	84.8	178.1	69.0	53.3	46.7	46.7	51.7	61.2	96.3	178.1	115.5	199.0	233.1	67.8	67.8	103.6	157.3	207.5	233.1
73	72.2	97.8	95.4	38.9	131.5	38.9	81.3	96.6	106.2	131.5	100.0	85.8	66.6	66.4	66.4	66.6	76.2	89.4	100.0
74	182.2	184.1	264.3	428.9	288.3	184.1	244.3	276.3	323.5	428.9	249.6	455.2	256.2	299.3	249.6	254.6	277.8	338.3	455.2
75	129.4	279.5	255.4	323.4	274.5	255.4	269.7	277.0	290.5	323.4	245.8	466.1	357.2	201.7	201.7	234.8	301.5	384.4	466.1
76	53.5	273.9	121.7	285.0	435.2	121.7	235.9	279.5	322.6	435.2	415.2	152.7	297.9	137.4	137.4	148.9	225.3	327.2	415.2
77	19.9	64.7	212.5	96.4	117.0	64.7	88.5	106.7	140.9	212.5	153.6	118.6	104.7	134.6	104.7	115.1	126.6	139.4	153.6
78	355.4	608.0	551.3	443.0	547.5	443.0	521.4	549.4	565.5	608.0	551.9	692.9	475.1	1,317.4	475.1	532.7	622.4	849.0	1,317.4
79	109.7	127.9	111.5	246.6	65.0	65.0	99.9	119.7	157.6	246.6	82.2	84.7	167.4	74.2	74.2	80.2	83.5	105.4	167.4
80	40.9	128.7	605.1	116.8	510.9	116.8	125.7	319.8	534.5	605.1	45.1	285.2	168.8	111.4	45.1	94.8	140.1	197.9	285.2
81	22.1	69.7	74.1	279.1	65.5	65.5	68.7	71.9	125.4	279.1	133.7	132.5	41.1	181.0	41.1	109.7	133.1	145.5	181.0
82	9.8	88.1	113.9	189.9	15.5	15.5	70.0	101.0	132.9	189.9	97.9	95.9	111.2	16.9	16.9	76.2	96.9	101.2	111.2
83	7.5	97.5	424.0	125.7	69.1	69.1	90.4	111.6	200.3	424.0	305.7	159.7	127.7	31.0	31.0	103.5	143.7	196.2	305.7
84	7.5	24.7	87.3	143.3	41.4	24.7	37.2	64.4	101.3	143.3	17.5	157.6	35.4	22.8	17.5	21.5	29.1	66.0	157.6
85	50.1	61.7	168.1	145.4	371.6	61.7	124.5	156.8	219.0	371.6	140.3	287.5	73.2	109.8	73.2	100.7	125.1	177.1	287.5
86	16.0	112.4	57.0	137.0	300.1	57.0	98.6	124.7	177.8	300.1	39.8	105.2	35.8	61.8	35.8	38.8	50.8	72.7	105.2
87	42.5	71.3	111.7	80.3	56.9	56.9	67.7	75.8	88.2	111.7	68.6	69.7	251.4	80.2	68.6	69.4	75.0	123.0	251.4
88	108.3	263.0	290.2	524.2	391.0	263.0	283.4	340.6	424.3	524.2	198.0	620.3	437.9	509.9	198.0	377.9	473.9	537.5	620.3
89	17.3	53.8	16.3	25.2	91.3	16.3	23.0	39.5	63.2	91.3	21.4	27.0	81.3	48.4	21.4	25.6	37.7	56.6	81.3
90	10.5	283.8	319.0	194.2	31.5	31.5	153.5	239.0	292.6	319.0	85.0	297.0	34.7	125.1	34.7	72.4	105.1	168.1	297.0
91	7.8	111.6	76.3	13.3	55.9	13.3	45.3	66.1	85.1	111.6	190.2	11.8	56.0	12.4	11.8	12.3	34.2	89.6	190.2
92	6.1	134.9	32.6	73.8	19.9	19.9	29.4	53.2	89.1	134.9	121.0	27.3	399.9	28.9	27.3	28.5	75.0	190.7	399.9
93	122.9	102.4	339.7	182.4	140.0	102.4	130.6	161.2	221.7	339.7	131.1	182.0	163.8	199.7	131.1	155.6	172.9	186.4	199.7
94	68.9	262.4	343.1	105.5	156.2	105.5	143.5	209.3	282.6	343.1	200.8	64.5	94.5	107.7	64.5	87.0	101.1	131.0	200.8
95	434.9	708.4	414.7	513.1	478.0	414.7	462.2	495.6	561.9	708.4	457.1	741.3	848.2	566.6	457.1	539.2	654.0	768.0	848.2
96	15.8	41.3	76.0	164.0	100.2	41.3	67.3	88.1	116.2	164.0	27.4	306.5	35.2	32.7	27.4	31.4	34.0	103.0	306.5
97	142.5	352.4	312.1	167.3	317.6	167.3	275.9	314.9	326.3	352.4	236.6	591.7	183.7	159.8	159.8	177.7	210.2	325.4	591.7
98	10.0	11.2	37.4	16.5	95.1	11.2	15.2	27.0	51.8	95.1	25.3	325.6	39.9	25.4	25.3	25.4	32.7	111.3	325.6
99	34.9	64.3	93.0	144.8	63.1	63.1	64.0	78.7	106.0	144.8	68.5	383.2	189.5	167.4	68.5	142.7	178.5	237.9	383.2

Timing Intervals for Each Refresh Function (In Seconds)

DM Fx	R-Run1	R-Run2	R-Run3	R-Run4	Min	25%tile	Median	75%tile	Max
LF_CR	35.0	34.1	37.1	40.2	34.1	34.8	36.1	37.9	40.2
LF_CS	59.4	58.4	64.3	48.1	48.1	55.8	58.9	60.6	64.3
LF_I	27.0	26.7	19.0	20.9	19.0	20.4	23.8	26.8	27.0
LF_SR	32.2	33.1	35.8	37.3	32.2	32.9	34.5	36.2	37.3
LF_SS	87.5	60.4	64.1	51.7	51.7	58.2	62.3	70.0	87.5
LF_WR	32.5	29.5	36.0	34.7	29.5	31.8	33.6	35.0	36.0
LF_WS	36.8	36.4	43.6	43.9	36.4	36.7	40.2	43.7	43.9
DF_CS	217.0	212.0	215.4	212.9	212.0	212.7	214.2	215.8	217.0
DF_SS	187.3	172.2	188.6	189.6	172.2	183.5	188.0	188.9	189.6
DF_WS	124.6	135.8	131.1	152.5	124.6	129.5	133.5	140.0	152.5
DF_I	64.9	95.1	102.5	62.8	62.8	64.4	80.0	97.0	102.5

Preface

TPC Benchmark™ DS Overview

The TPC Benchmark™ DS (TPC-DS) is a decision support benchmark that models several generally applicable aspects of a decision support system, including queries and data maintenance. The benchmark provides a representative evaluation of performance as a general purpose decision support system.

This benchmark illustrates decision support systems that:

- Examine large volumes of data;
- Give answers to real-world business questions;
- Execute queries of various operational requirements and complexities (e.g., ad-hoc, reporting, iterative OLAP, data mining);
- Are characterized by high CPU and IO load;
- Are periodically synchronized with source OLTP databases through database maintenance functions.
- Run on “Big Data” solutions, such as RDBMS as well as Hadoop/Spark based systems.

A benchmark result measures query response time in single user mode, query throughput in multi user mode and data maintenance performance for a given hardware, operating system, and data processing system configuration under a controlled, complex, multi-user decision support 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 tests be implemented with systems, products, technologies and pricing that:

- a) Are generally available to users;
- b) Are relevant to the market segment that the individual TPC benchmark models or represents (e.g., TPC-DS models and represents complex, high data volume, decision support environments);
- c) Would plausibly be implemented by a significant number of users in the market segment modeled or represented by the benchmark.

In keeping with these requirements, the TPC-DS database must be implemented using commercially available data processing software, and its queries must be executed via SQL interface. 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, which improve benchmark results but not real-world performance or pricing, are prohibited.

TPC benchmark results are expected to be accurate representations of system performance. Therefore, there are specific guidelines that are expected to be followed when measuring those results. The approach or methodology to be used in the measurements are either explicitly described in the specification or left to the discretion of the test sponsor.

When not described in the specification, the methodologies and approaches used must meet the following requirements:

- The approach is an accepted engineering practice or standard;
- The approach does not enhance the result;
- Equipment used in measuring the results is calibrated according to established quality standards;
- Fidelity and candor is maintained in reporting any anomalies in the results, even if not specified in the benchmark requirements.

Further information is available at www.tpc.org

General Items

0.1 Test Sponsor

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

This benchmark was sponsored by Super Micro Computer, Inc.

0.2 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:

- *Database Tuning Options*
- *Optimizer/Query execution options*
- *Query processing tool/language configuration parameters*
- *Recovery/commit options*
- *Consistency/locking options*
- *Operating system and configuration parameters*
- *Configuration parameters and options for any other software component incorporated into the pricing structure*
- *Compiler optimization options*

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

The Supporting File Archive contains the Operating System and DBMS parameters used in this benchmark.

0.3 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:

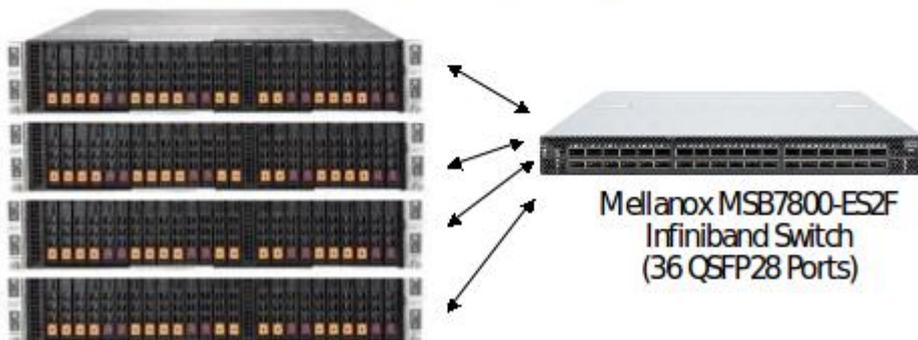
- *Number and type of processors*
- *Size of allocated memory, and any specific mapping/partitioning of memory unique to the test.*
- *Number and type of disk units (and controllers, if applicable).*
- *Number of channels or bus connections to disk units, including their protocol type.*
- *Number of LAN (e.g. Ethernet) Connections, including routers, workstations, terminals, etc., that were physically used in the test or are incorporated into the pricing structure.*
- *Type and the run-time execution location of software components (e.g., DBMS, query processing tools/languages, middle-ware components, software drivers, etc.).*

Measured Configuration

4x Supermicro A+ Server AS-2123BT-HNC0R

BigTwin™ with 4x 217BH-22H12 Nodes, each with:

- 2x AMD EPYC 7502 32-Core Processor
- 256 GB (16x16GB RDIMM 2666MT/s Dual Rank)
- 1x 960GB M.2 PCIe Gen3
- 2x 3.5TB NVMe Drive
- 2x 1TB NVMe Drive
- 1x Mellanox MT27800 ConnectX-5 25GbE SFP28 NIC



The measured configuration consisted of:

- Total Nodes: 16
- Total Processors/Cores/Threads: 32/1024/2048
- Total Memory: 4096GB
- Total Number of Storage Drives/Devices: 80
- Total Storage Capacity: 162816 GB

Server nodes details:

- 16 x Supermicro A+ Servers, each with:
 - ◆ Processors/Cores/Threads: 2/64/128
 - ◆ Processor Model: 2 x AMD EPYC 7502 32-Core Processor
 - ◆ Memory: 256GB
 - ◆ Drives:
 - ✓ 2 x 3.5 TB NVMe Drive (data nodes and tablet servers)
 - ✓ 2 x 1 TB NVMe Drive (data nodes and tablet servers)
 - ✓ 1 x 960 GB M.2 PCIe Gen3 (all nodes, boot disk)
 - ◆ Network: 1 x Mellanox MT27800 Family [ConnectX-5] Network Adapter 25GbE

Transwarp System Components Configuration

Node	ArgoDB Storage		HDFS		ZooKeeper	ArgoDB computing	
	Shiva Master	Shiva tabletserver	NameNode	DataNode	QuorumPeer	Server	Executor
1	x	x	x	x		x	x
2-3	x	x		x			x
4	x	x		x	x		x
5	x	x	x	x	x		x
6		x		x	x		x
7-11		x		x			x
12		x		x	x		x
13-16		x		x			x

Note:

1) In the table, NameNode, DataNode and others are software services, not hardware nodes. HA (High Availability) is enabled for NameNode service. The service on Node 1 is the active NameNode service and the one on Node 5 is the standby one. In all other part of the documents, NameNode is hardware Node 1, on which the active NameNode service runs.

2) HDFS is mainly used for raw data. The data of the test DB is managed by storage service (Shiva), which is a distributed storage engine based on the distributed consistency protocol Raft to ensure the stability and reliability of the data platform. For more detail of ArgoDB, please refer: <http://www.transwarp.io/en/transwarp/product-ArgoDB.html?categoryId=77>

Priced Configuration

In the measured configurations, each node included 2 x 1TB Intel® NVMe DC P4500 Series. In the priced configuration, these were substituted by 2 x 2.0TB NVMe DC P4510 Series. This substitution was based on documented performance for the components. As required by the TPC Pricing specification, the priced component shows equal or better performance than the measured component. The characteristics of the components are as follows:

Characteristic	Priced	Measured
Disk Model #	SSDPE2KX020T8	SSDPE2KX010T7
Intel Series	Intel P4510 Series	Intel P4500 Series
Capacity, Form Factor and Interface	2.0TB, 2.5in PCIe* 3.1 x4, 3D2, TLC	1.0TB, 2.5in PCIe 3.1 x4, 3D1, TLC
Launch Date	Q3'17	Q2'17
Operation	Performance	Performance
Sequential Read (up to)	3200 MB/s	3200 MB/s
Sequential Write (up to)	2000 MB/s	600 MB/s
Random Read (100% Span)	637000 IOPS	279500 IOPS
Random Write (100% Span)	81500 IOPS	30500 IOPS
Latency - Read	77 μs	85 μs
Latency - Write	18 μs	20 μs

Clause 2: Logical Database Design Related Items

2.1 Database Definition Statements

Listings must be provided for the DDL scripts and must include all table definition statements and all other statements used to set up the test and qualification databases

The Supporting File Archive contains the table definitions and all other statements used to set up the test and qualification databases.

2.2 Physical Organization

The physical organization of tables and indices within the test and qualification databases must be disclosed. If the column ordering of any table is different from that specified in Clause 2.3 or 2.4., it must be noted.

The store_sales, store_returns, catalog_sales, catalog_returns, web_sales, web_returns and inventory are partitioned. The partition column: ss_sold_date_sk, sr_returned_date_sk, cs_sold_date_sk, cr_returned_date_sk, ws_sold_date_sk, wr_returned_date_sk and inv_date_sk will be the last column in the tables.

2.3 Horizontal Partitioning

If any directives to DDLs are used to horizontally partition tables and rows in the test and qualification databases, these directives, DDLs, and other details necessary to replicate the partitioning behavior must be disclosed.

Horizontal partitioning is used on store_sales, store_returns, catalog_sales, catalog_returns, web_sales, web_returns and inventory tables and the partitioning columns are ss_sold_date_sk, sr_returned_date_sk, cs_sold_date_sk, cr_returned_date_sk, ws_sold_date_sk, wr_returned_date_sk and inv_date_sk.

For *_sales and inventory, the partition granularity is by month.

For *_returns, the partition granularity is by quarter.

2.4 Replication

Any replication of physical objects must be disclosed and must conform to the requirements of Clause 2.5.3.

All the objects are replicated by ArgoDB storage engine (Shiva) in 3 replications.

Clause 3: Scaling and Database Population

3.1 Initial Cardinality of Tables

The cardinality (e.g., the number of rows) of each table of the test database, as it existed at the completion of the database load (see Clause 7.1.2) must be disclosed.

Table 3.1 lists the TPC Benchmark DS defined tables and the row count for each table as they existed upon completion of the build.

Table 3.1 Initial Number of Rows	
Table Name	Row count
call_center	54
catalog_page	40,000
catalog_returns	1,440,033,112
catalog_sales	14,399,964,710
customer	65,000,000
customer_address	32,500,000
customer_demographics	1,920,800
date_dim	73,049
household_demographics	7,200
income_band	20
inventory	1,311,525,000
item	402,000
promotion	2,000
reason	70
ship_mode	20
store	1,500
store_returns	2,879,512,807
store_sales	28,799,795,112
time_dim	86,400
warehouse	25
web_page	4,002
web_returns	720,020,485
web_sales	7,199,963,324
web_site	78

3.2 Distribution of Tables and Logs Across Media

The distribution of tables and logs across all media must be explicitly described using a format similar to that shown in the following example for both the tested and priced systems.

Server node	Controller	Disk driver	Description of Content
1-16		1-2 (SSD)	Raw data, Table, Temp
1-16		3-4 (SSD)	Raw data, Table, Temp
1 - 16		0 (SSD)	Operating system, root, swap, Transwarp data hub, logs

All the raw data were on HDFS, which are distributed among the disks in datanodes:

/data/disk1/hadoop/data, /data/disk2/hadoop/data,
/data/disk3/hadoop/data, /data/disk4/hadoop/data

After load, the table data will be put in Shiva, which is configured to use 3 replications by default. The data dir for each node is:

/mnt/disk1/shiva-tabletserver/data, /mnt/disk2/shiva-tabletserver/data,
/mnt/disk3/shiva-tabletserver/data, /mnt/disk4/shiva-tabletserver/data

We can read the table sizes in inceptor-server-argodbcomputing1 UI: 10.216.91.101:8888/argodb/set4040 -> holodesk, then select a table, it will show the details which include the TotalSize. Here is the list of table size:

Tables	size
store_sales	1720.0G
catalog_sales	1249.2G
web_sales	622.6G
web_returns	64.3G
store_returns	210.5G
catalog_returns	136.9G
inventory	14.7G
customer	5.5G
customer_address	1585.3M
item	61.1M
customer_demographics	21.9M
date_dim	2.4M
time_dim	1777.5K
web_site	19.1K
web_page	128.4K
call_center	16.8K
promotion	118.2K
reason	2.8k
household_demographics	68.9K
ship_mode	2.4k
income_band	842
store	172.7K
warehouse	5.0K
catalog_page	2.6M

Temp data from Transwarp data hub are also distributed among the disks in datanodes:

/mnt/disk1/hadoop/ngmr/argodbcomputing1, /mnt/disk2/hadoop/ngmr/argodbcomputing1,
/mnt/disk3/hadoop/ngmr/argodbcomputing1, /mnt/disk4/hadoop/ngmr/argodbcomputing1

Logs are at /var/log on each node.

3.3 Mapping of Database Partitions/Replications

The mapping of database partitions/replications must be explicitly described.

Horizontal partitioning is used on store_sales, store_returns, catalog_sales, catalog_returns, web_sales, web_returns and inventory tables and the partitioning columns are ss_sold_date_sk, sr_returned_date_sk, cs_sold_date_sk, cr_returned_date_sk, ws_sold_date_sk, wr_returned_date_sk and inv_date_sk.

For *_sales and inventory, the partition granularity is by month.

For *_returns, the partition granularity is by quarter.

3.4 Implementation of RAID

Implementations may use some form of RAID. The RAID level used must be disclosed for each device. If RAID is used in an implementation, the logical intent of its use must be disclosed

RAID is not used. The database tables were on top of Shiva which does distributed 3-way replications based on the distributed consistency protocol Raft to ensure the stability and reliability.

3.5 DBGEN Modifications

The version number (i.e., the major revision number, the minor revision number, and third tier number) of dsdgen must be disclosed. Any modifications to the dsdgen source code (see Appendix B:) must be disclosed. In the event that a program other than dsdgen was used to populate the database, it must be disclosed in its entirety.

Dsdgen version 2.11.0 was used. No modifications for dsdgen were made except a map/reduce wrapper to reduce the time. The wrapper with source codes were included in the Supporting Files.

3.6 Database Load time

The database load time for the test database (see Clause 7.4.3.7) must be disclosed.

The database load time was 2404.0 seconds.

3.7 Data Storage Ratio

The data storage ratio must be disclosed. It is computed by dividing the total data storage of the priced configuration (expressed in GB) by SF corresponding to the scale factor chosen for the test database as defined in Clause 3.1. The ratio must be reported to the nearest 1/100th, rounded up. For example, a system configured with 96 disks of 2.1 GB capacity for a 100GB test database has a data storage ratio of 2.02.

Size (GB)	Disks	Nodes	Total disks	Total (GB)
960	1	16	16	15,360
3,584	2	16	32	114,688
2,048	2	16	32	65,536
Grand Total (GB)				195,584

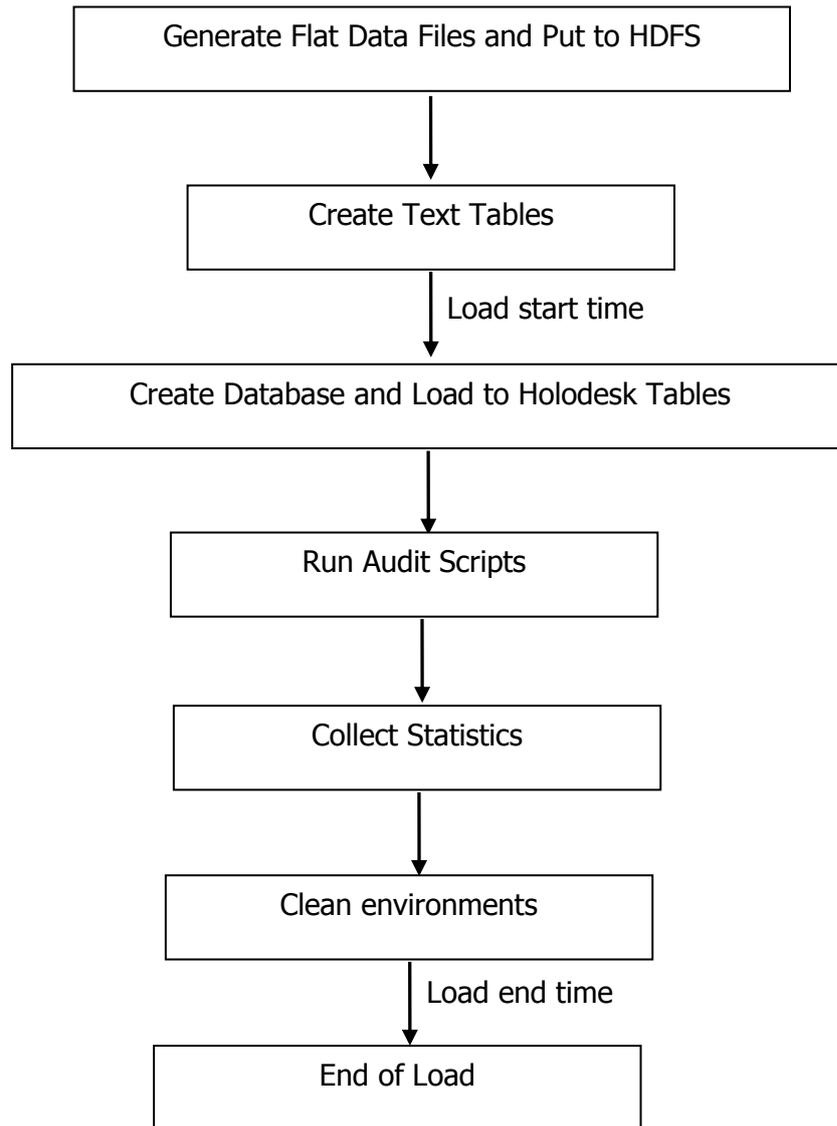
SF	10,000	Ratio	19.56
-----------	--------	--------------	-------

3.8 Database Load Mechanism Details and Illustration

The details of the database load must be disclosed, including a block diagram illustrating the overall process. Disclosure of the load procedure includes all steps, scripts, input and configuration files required to completely reproduce the test and qualification databases.

The tables were loaded as shown in Figure 3.8. All the sql and scripts were included in the Supporting Files.

Figure 3.8: Block Diagram of Database Load Process



The final database load time is (load end time – load start time) – auditor time.

3.9 Qualification Database Configuration

Any differences between the configuration of the qualification database and the test database must be disclosed

The qualification database used identical scripts to create and load the data with changes to adjust for the database scale factor.

Clause 4 and 5: Query and Data Maintenance Related Items

4.1 Query Language

The query language used to implement the queries must be identified

SQL was the query language used to implement the queries.

4.2 Verifying Method of Random Number Generation

The method of verification for the random number generation must be described unless the supplied dsdgen and dsqgen were used.

A map/reduce wrapper based on TPC-supplied dsdgen version 2.11.0 and dsqgen version 2.11.0 were used.

4.3 Generating Values for Substitution Parameters

The method used to generate values for substitution parameters must be disclosed. The version number (i.e., the major revision number, the minor revision number, and third tier number) of dsqgen must be disclosed.

TPC supplied dsqgen version 2.11.0 was used to generate the substitution parameters.

```
./dsqgen -input ../../tdh_templates/templates.lst -directory ../../tdh_templates/ -dialect tdh -scale 10000 -streams 9 -  
output ../../SQLs/ -RNGSEED ${seed}
```

4.4 Query Text and Output Data from Qualification Database

The executable query text used for query validation must be disclosed along with the corresponding output data generated during the execution of the query text against the qualification database. If minor modifications have been applied to any functional query definitions or approved variants in order to obtain executable query text, these modifications must be disclosed and justified. The justification for a particular minor query modification can apply collectively to all queries for which it has been used. The output data for the power and Throughput Tests must be made available electronically upon request.

Supporting Files Archive contains the actual query text and query output. Following are the modifications to the query.

Query variants are used for the following queries:

- 10, 27, 35, 36, 67, 70, 86

The following MQM are used:

- Add and use column alias. (MQM e.5)
 - Q32
 - Q92

Query results are inserted in a file (clause 4.2.5):

- Use of an external table with column delimiter
 - Q64
 - using an external table named q64_resultfile_[s] (stream [s])

Supporting Files Archive contains the full query template used in tests.

4.5 Query Substitution Parameters and Seeds Used

All the query substitution parameters used during the performance test must be disclosed in tabular format, along with the seeds used to generate these parameters.

Supporting Files Archive contains the query substitution parameters and seed used.

4.6 Refresh Setting

All query and refresh session initialization parameters, settings and commands must be disclosed

Supporting Files Archive contains the query and scripts.

4.7 Source Code of Refresh Functions

The details of how the data maintenance functions were implemented must be disclosed (including source code of any non-commercial program used).

Supporting Files Archive contains the Source Code of refresh functions.

4.8 staging Area

Any object created in the staging area (see Clause 5.1.8 for definition and usage restrictions) used to implement the data maintenance functions must be disclosed. Also, any disk storage used for the staging area must be priced, and any mapping or virtualization of disk storage must be disclosed.

Not used.

Clause 6: Data Persistence Properties Related Items

The results of the data accessibility tests must be disclosed along with a description of how the data accessibility requirements were met. This includes disclosure of the code written to implement the data accessibility Query.

The data accessibility test was performed by failing the data dirs of raw data and table data during the first throughput test.

In SupportingFiles/Clause_4/validation/test.sh, commands to make the two data dir on tpcds15 not accessible:

```
pdsh -w tpcds[15] "chmod -rw -R /mnt/disk4/hadoop/data"
```

```
pdsh -w tpcds[15] "chmod -rw -R /mnt/disk4/shiva-tabletserver/data"
```

Commands to capture the status of a durable medium:

```
pdsh -w tpcds[15] "ls -l /mnt/disk4/hadoop "
```

```
pdsh -w tpcds[15] "ls -l /mnt/disk4/shiva-tabletserver "
```

The test continued to run without interruptions. Supporting Files Archive contains the disk status before and after the failure.

Clause 7: Performance Metrics and Execution Rules Related Items

7.1 System Activity

Any system activity on the SUT that takes place between the conclusion of the load test and the beginning of the performance test must be fully disclosed including listings of scripts or command logs.

The page caches of each node is cleared between the end of *load test* and the beginning of the *performance test*.

Supporting Files Archive contains the scripts and logs.

7.2 Test Steps

The details of the steps followed to implement the performance test must be disclosed.

Supporting Files Archive contains the scripts and logs.

7.3 Timing Intervals for Each Query and Refresh Function

The timing intervals defined in Clause 7 must be disclosed.

See the Executive Summary at the beginning of this report.

7.4 Throughput Test Result

For each Throughput Test, the minimum, the 25th percentile, the median, the 75th percentile, and the maximum times for each query shall be reported.

See the Executive Summary at the beginning of this report.

7.5 Time for Each Stream

The start time and finish time for each query stream must be reported.

See the Executive Summary at the beginning of this report.

7.6 Time for Each Refresh Function

The start time and finish time for each data maintenance function in the refresh run must be reported for the Throughput Tests

See the Executive Summary at the beginning of this report.

7.7 Performance Metrics

The computed performance metric, related numerical quantities and the price/performance metric must be reported.

QphDS@10000GB = 4,418,054

See the Executive Summary at the beginning of this report for more detail.

Clause 8: SUT and Driver Implementation Related Items

8.1 Driver

A detailed textual description of how the driver performs its functions, how its various components interact and any product functionalities or environmental settings on which it relies must be provided. All related source code, scripts and configuration files must be disclosed. The information provided should be sufficient for an independent reconstruction of the driver.

Transwarp ArgoDB is the test DBMS. The data of the test DB is managed by storage service (Shiva), which is a distributed storage engine based on the distributed consistency protocol Raft to ensure the stability and reliability of the data platform. All the objects are replicated by ArgoDB storage engine (Shiva) in 3 replications. For more detail, please refer: <http://www.transwarp.io/en/transwarp/product-ArgoDB.html?categoryId=77>

beeline is the client in Transwarp ArgoDB. It connects to the server by JDBC. The command is:

```
beeline -u jdbc:hive2://localhost:10000 --maxWidth=1000 -f sqlfile
```

(--maxWidth is a parameter to control the output width)

There is a argodbcomputing server to compile SQL. After compilation, the execution plan will be sent to executors. The result will be on HDFS. Then the client will get the result from HDFS.

In the test, tpcds1 is configured as the server, all nodes are configured as workers.

Supporting Files Archive contains all the command, scripts and logs.

8.2 Implementation Specific Layer (ISL)

If an implementation specific layer is used, then a detailed description of how it performs its functions, how its various components interact and any product functionalities or environmental setting on which it relies must be provided. All related source code, scripts and configuration files must be disclosed. The information provided should be sufficient for an independent reconstruction of the implementation specific layer.

No Specific Layer is used.

Supporting Files Archive contains all the command, scripts and logs.

8.3 Profile-Directed Optimization

If profile-directed optimization as described in Clause 7.2.10 is used, such use must be disclosed. In particular, the procedure and any scripts used to perform the optimization must be disclosed.

Profile-directed optimization was not used.

Clause 9: Pricing Related Items

9.1 Hardware and Software Used

A detailed list of hardware and software used in the priced system must be reported. The rules for pricing are included in the current revision of the TPC Pricing Specification located on the TPC website (<http://www.tpc.org>)

A detailed list of all hardware and software, is provided in the Executive Summary of this report.

9.2 Availability Date

The System Availability Date (see Clause 7.6.5) must be the single availability date reported on the first page of the executive summary. The full disclosure report must report Availability Dates individually for at least each of the categories for which a pricing subtotal must be. All Availability Dates required to be reported must be disclosed to a precision of 1 day, but the precise format is left to the test sponsor.

The total system availability date is 31-Aug-2019.

9.3 Country-Specific Pricing

Additional Clause 7 related items may be included in the full disclosure report for each country specific priced configuration..

The configuration is priced for US.

Clause 11: Audit Related Items

Auditors' Information and Attestation Letter

The auditor's agency name, address, phone number, and attestation letter with a brief audit summary report indicating compliance must be included in the full disclosure report. A statement should be included specifying whom to contact in order to obtain further information regarding the audit process.

This benchmark was audited by:

Francois Raab,

InfoSizing,

20 Kreg Ln.

Manitou Springs, CO 80829.

Phone Number: 719-473-7555.

Benchmark sponsor: Super Micro Computer, Inc.
980 Rock Avenue
San Jose, CA 95131, USA

August 1, 2019

I verified the TPC Benchmark™ DS (TPC-DS™ v2.11.0) performance of the following configuration:

Platform: Supermicro A+ Server 2123BT-HNC0R (16-node cluster)
Operating System: Red Hat Enterprise Linux Server Release 7.6
Database Manager: Transwarp ArgoDB V1.2.1

The results were:

Performance Metric **4,418,054 QphDS@10000GB**
Database Load Time 0h 40m 04s

Servers

16 x Supermicro A+ Server 2123BT-HNC0R, each with:

CPU	2 x AMD EPYC™ 7502 3.35GHz processors (32 cores, 64 threads)		
Memory	256 GB		
Storage	Qty	Size	Type
	1	960 GB	M2 PCIe Gen3
	2	1.0 TB	NVMe
	2	3.5 TB	NVMe

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:

- The database records were defined with the proper layout and size
- The database population was generated using Dsdgen
- The database was properly scaled to 10,000GB and populated accordingly
- The database load time was correctly measured and reported
- The query templates were produced using approved minor query modifications and query variants
- The query input variables were generated by Dsqgen
- The execution of the queries against the qualification database produced compliant output
- The tests were driven and sequenced according to the requirements
- The throughput tests involved 4 query streams

- The execution times for queries and data maintenance functions were correctly measured and reported
- The data accessibility test was performed and verified
- The system pricing was verified for major components and maintenance
- The major pages from the FDR were verified for accuracy

Additional Audit Notes:

In the measured configurations, each node included 2 x 1TB Intel® NVMe DC P4500 Series. In the priced configuration, these were substituted by 2 x 2.0TB NVMe DC P4510 Series. This substitution was based on documented performance for the components. As required by the TPC Pricing specification, the priced component shows equal or better performance than the measured component. The characteristics of the components are listed in the Full Disclosure Report.

Respectfully Yours,

A handwritten signature in black ink, appearing to read "François Raab", with a long horizontal flourish extending to the right.

François Raab, TPC Certified Auditor

Supporting Files Index

Clause	Description	Pathname
Introduction	Database Tunable	SupportingFiles/Introduction/DBtune.txt
	Parameters	SupportingFiles/Introduction/Transwarp-ArgoDB/
	OS Tunable	SupportingFiles/Introduction/OSTune.txt
	Parameters	SupportingFiles/Introduction/OperatingSystem/
Clause 2		Index is not used.
Clause 3	Load transaction scripts	SupportingFiles/Clause3/doLoad.sh
Clause 4	Qualification test	SupportingFiles/Clause_4/QUALIFICATION-test/
	Performance Test scripts/output	SupportingFiles/Clause_4/
Clause 5	Data maintenance scripts	SupportingFiles/Clause5/doRefresh.sh
Clause 6	Data Accessibility	
	Scripts	In SupportingFiles/Clause_4/test.sh
	Logs	SupportingFiles/Clause_4/validation/

Appendix A: Price Quotes

Third party Price Quotes

Mellanox SN2410 Switch

The screenshot shows the Mellanox store product page for the MSB7800-ES2F switch. The page features a blue header with the Mellanox logo, contact information, and navigation links. The main content area includes a product image, a detailed title, and a list of specifications. The price is listed as \$25,633.00, and the availability is marked as 'Limited'. There are buttons for 'Add to Cart' and 'Get a Quote'.

https://store.mellanox.com/products/mellanox-msb7800-es2f-switch-ib-2-based-edr-infiniband-1u-switch-36-qsfp28-ports-2-power-supplies-ac-x86-dual-core-standard-dep...

Mellanox TECHNOLOGIES
THE OFFICIAL STORE

Call toll-free: (855) 897-1098
International +01-512-220-1515

Currency ▾

Sign In Or Register Return to Mellanox.com

Search the store: [input] [Q]

Get a Quote [CART 1]

ADAPTERS ▾ INTERCONNECT ▾ SWITCHES ▾ CERTIFIED REFURBISHED ▾ SUPPORT AND SERVICES ▾ CUSTOMER SERVICE ▾

Home | Switches | InfiniBand and VPI Switches | SB7800 | Mellanox MSB7800-ES2F Switch-IB 2 Based EDR InfiniBand 1U Switch 36 QSFP...

Mellanox MSB7800-ES2F Switch-IB 2 Based EDR InfiniBand 1U Switch 36 QSFP28 Ports 2 Power Supplies AC x86 Dual Core Standard Depth P2C Airflow Rail Kit RoHS6

MPN: MSB7800-ES2F Condition: New

Availability: **Limited** 📌
MSRP: **\$25,633.00**

Switch Family: SB7800
Condition: New
Max Ports: 36
Connector Type: QSFP28
Max Speed: EDR
ECCN: 5A991
Technology: InfiniBand
Product Brief: [Download MSB7800-ES2F](#)

Add Mellanox Support?:

Quantity:

[Add to Cart](#)

[Get a Quote](#)

Need it fast? These similar items **ship today**.

Mellanox Certified

<https://store.mellanox.com/products/mellanox-sup-sb7800-3s-technical-support-and-warranty-silver-3-year-for-sb7800-series-switch.html>


Mellanox
 TECHNOLOGIES
 THE OFFICIAL STORE

Call toll-free: (855) 897-1098
 International +01-512-220-1515

Currency ▾

Sign In Or Register Return to Mellanox.com

Get a Quote **CART 1**

Search the store

ADAPTERS ▾ INTERCONNECT ▾ SWITCHES ▾ CERTIFIED REFURBISHED ▾ SUPPORT AND SERVICES ▾ CUSTOMER SERVICE ▾

Home | Support and Services | Hardware Support | InfiniBand Hardware Support | InfiniBand Switch Support | SB7800 Switch Support

| Mellanox SUP-SB7800-3S Technical Support and Warranty - Silver 3 Year for SB7...



Mellanox
 TECHNOLOGIES
 Connect. Accelerate. Outperform.™

In stock

Mellanox SUP-SB7800-3S Technical Support and Warranty - Silver 3 Year for SB7800 Series Switch

Availability: In stock 📦 \$2,815.00	Quantity: <input type="text" value="1"/>
Technology: InfiniBand Program Level: Silver Contract Length: 3 Years	<div style="background-color: #008000; color: white; padding: 5px; margin: 5px auto; width: 80%; border-radius: 5px;">Add to Cart</div> <div style="background-color: #ccc; padding: 5px; margin: 5px auto; width: 80%; border-radius: 5px;">Get a Quote</div>



One-Year
Hardware Warranty



Services & Support
Call (855) 897-1098



Mellanox Specialists
Send us an Email

Miscellaneous Items

Ships same day **Mellanox MCP2M00-A001E30N Ethernet Passive Copper Cable 25GbE SFP28 1m Black 30AWG CA-N**

MPN: MCP2M00-A001E30N Condition: **New**



Availability:	Ships same day 📦
Price:	\$52.00
Condition:	New
Technology:	Ethernet
Max Speed:	25GbE
Material:	Copper
Connector Type:	SFP28
Passive/Active:	Passive
Length:	2.0m & under
ECCN:	EAR99
Product Brief:	Download MCP2M00-A001E30N

Quantity:

[Add to Cart](#)

[Get a Quote](#)



10ft (3m) NEMA 5-15P to IEC320 C15 14AWG 125V/15A Power Cord, Blue #36113

★★★★★ 56 Reviews | 0 Question | [Request Stock](#)

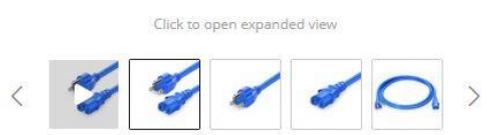
US\$ 10.00
FS P/N: [P515C15-15ABE-10](#)

Length:

Color:

Size:

37 In Stock, U.S. Warehouse ⓘ
Deliver to New York, 10010
Free Shipping via FedEx Ground® on orders over US\$ 79
Get it by Tue, Jul. 16 ⓘ



[Share](#)

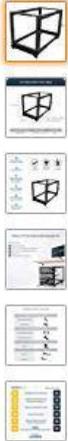
[Add to Cart](#)



Save up to 40% on groceries

Shop now

Electronics > Accessories & Supplies > Audio & Video Accessories > Speaker Parts & Components



Roll over image to zoom in

StarTech.com 12U Open Frame Server Rack - Adjustable Depth - 4-Post Data Rack - w/ Casters/Levelers/Cable Management Hooks (4POSTRACK12U)

by StarTech

★★★★☆ 123 customer reviews

| 46 answered questions

Amazon's Choice for "server rack"

Price: \$195.99 **prime**

Thank you for being a Prime member. Get \$80 off instantly: Pay \$115.99 upon approval for the Amazon Prime Rewards Visa Card.

Free Amazon tech support included

Size: 12U

12U Adjustable Posts Slide-Out

Style: with Casters

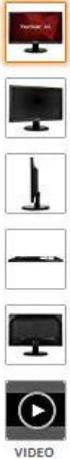
with Casters No Casters



\$15 off your order of \$35+
FREE 2-hour delivery

Shop now

← Back to results



ViewSonic VA2055SA 20 Inch 1080p LED Monitor with VGA Input and Enhanced Viewing Comfort

by ViewSonic

★★★★☆ 207 customer reviews
| 85 answered questions

Price: **\$79.00** ✓prime FREE One-Day & FREE Returns

Get \$80 off instantly: Pay \$0.00 upon approval for the Amazon Prime Rewards Visa Card.

Free Amazon tech support included

Style: **Regular**

Extra Ports and Built-in Speakers
\$84.99
✓prime

Regular

\$79.00
✓prime



Save up to 40% on groceries

Shop now

Back to results



AmazonBasics USB Wired Computer Keyboard and Wired Mouse Bundle Pack

by AmazonBasics

★★★★☆ 1,989 customer reviews

| 209 answered questions

Amazon's Choice for "keyboard and mouse"

Price: **\$14.56** ✓prime FREE One-Day & FREE Returns

Thank you for being a Prime member. Get \$80 off instantly: Pay \$0.00 upon approval for the Amazon Prime Rewards Visa Card.

Style: **1-Pack**

1-Pack \$14.56 ✓prime	10-Pack \$122.30 ✓prime
------------------------------------	--------------------------------------

Transwarp ArgoDB



Super Micro Computer, Inc

980 Rock Avenue

San Jose, CA 95131

Here is the information you requested regarding pricing for Transwarp products to be used in conjunction with your TPC-DS benchmark testing.

All pricing shown is in US Dollars (\$).

Part Code	Description	Unit Price	Quantity	Total Price
TDB-ARGO-DM	Transwarp high performance Data Mart for Online Analysis Version 1.2.1 with 3 years 24x7 support	\$17,400 per 1 node	16	\$278,400

Jun Zheng (Business Development Director)

Jun.zheng@transwarp.io