



Cisco Systems, Inc.

TPC Benchmark™ DS
Full Disclosure Report
for
Cisco UCS Integrated Infrastructure for Big Data
(with 17 Cisco UCS C240 M4 Servers)
using
Transwarp Data Hub V5.1
and
Red Hat Enterprise Linux Server Release 6.7

First Edition

March 05, 2018

First Edition – March 05, 2018

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.

The Cisco 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 Cisco business contact for information on the products or services available in your area. You can find additional information via Cisco's World Wide Web server at <http://www.cisco.com>. Actual performance and environmental costs of Cisco products will vary depending on individual customer configurations and conditions.

Table of Contents

Abstract	5
Preface.....	12
General Items.....	13
0.1 Test Sponsor.....	13
0.2 Parameter Settings.....	13
0.3 Configuration Diagrams	13
Third Party Price Quotes.....	15
Clause 2: Logical Database Design Related Items	17
2.1 Database Definition Statements.....	17
2.2 Physical Organization.....	17
2.3 Horizontal Partitioning.....	17
2.4 Replication.....	17
Clause 3: Scaling and Database Population.....	18
3.1 Initial Cardinality of Tables.....	18
3.2 Distribution of Tables and Logs Across Media.....	18
3.3 Mapping of Database Partitions/Replications	20
3.4 Implementation of RAID	20
3.5 DBGEN Modifications	20
3.6 Database Load time.....	20
3.7 Data Storage Ratio	20
3.8 Database Load Mechanism Details and Illustration	20
3.9 Qualification Database Configuration.....	21
Clause 4 and 5: Query and Data Maintenance Related Items	22
4.1 Query Language.....	22
4.2 Verifying Method of Random Number Generation	22
4.3 Generating Values for Substitution Parameters.....	22
4.4 Query Text and Output Data from Qualification Database	22
4.5 Query Substitution Parameters and Seeds Used	23
4.6 Refresh Setting	23
4.7 Source Code of Refresh Functions.....	23

4.8 staging Area	23
Clause 6: Data Persistence Properties Related Items	24
Clause 7: Performance Metrics and Execution Rules Related Items	25
7.1 System Activity.....	25
7.2 Test Steps	25
7.3 Timing Intervals for Each Query and Refresh Function.....	25
7.4 Throughput Test Result	25
7.5 Time for Each Stream	25
7.6 Time for Each Refresh Function	25
7.7 Performance Metrics.....	25
Clause 8: SUT and Driver Implementation Related Items	26
8.1 Driver.....	26
8.2 Implementation Specific Layer (ISL).....	26
8.3 Profile-Directed Optimization	26
Clause 9: Pricing Related Items	27
9.1 Hardware and Software Used	27
9.2 Availability Date	27
9.3 Country-Specific Pricing	27
Clause 11: Audit Related Items	28
Supporting Files Index.....	31

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 2.7.0.

The test was conducted at a Scale Factor of 10000GB with 17 Cisco UCS C240 M4 Servers running Transwarp Data Hub V5.1 on Red Hat Enterprise Linux Server Release 6.7.

Measured Configuration

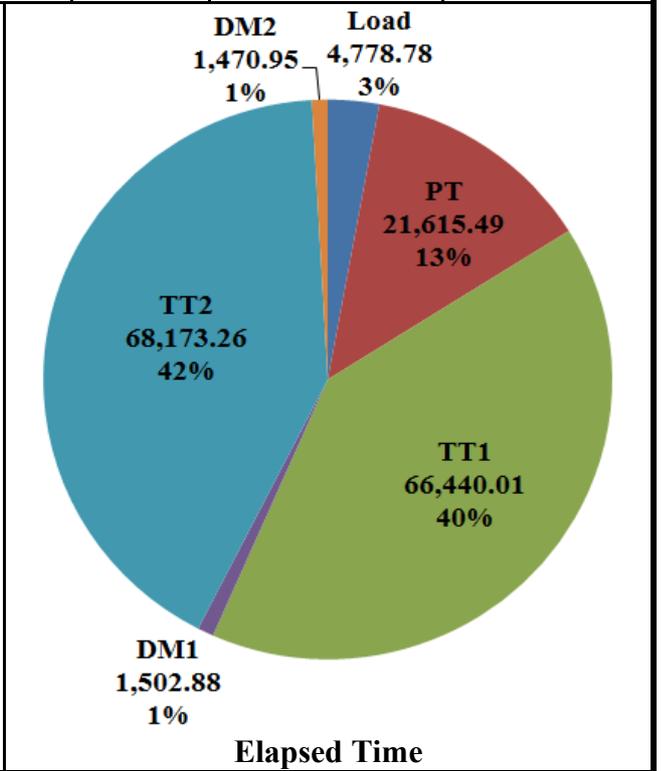
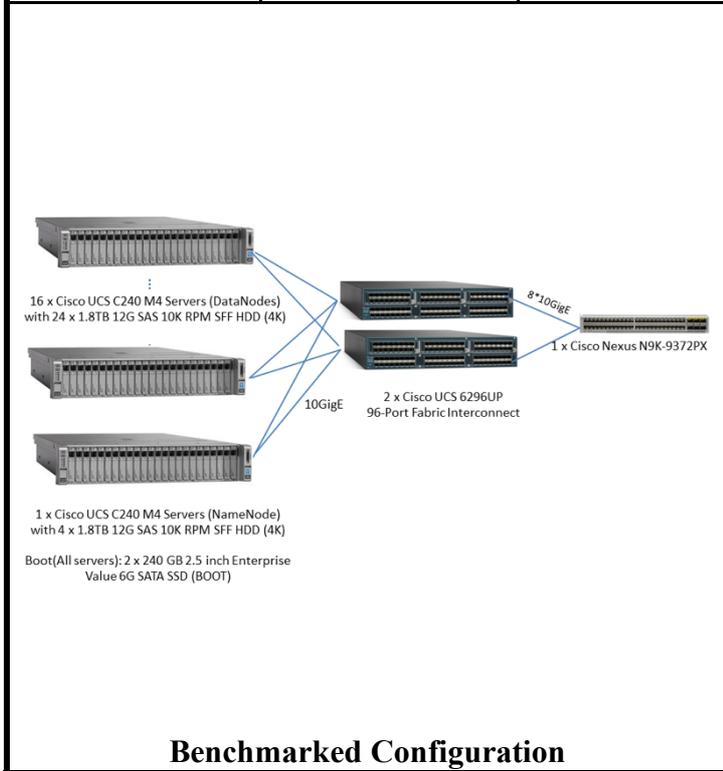
Company Name	Cluster Node	Database Software	Operating System
Cisco Systems, Inc	Cisco UCS C240 M4 Server	Transwarp Data Hub V5.1	Red Hat Enterprise Linux Server Release 6.7

TPC Benchmark™ DS Metrics

Total System Cost	TPC-DS Throughput	Price/Performance	Availability Date
\$1,008,066 USD	1,580,649 QphDS@10000GB	\$0.64 USD \$ / QphDS@10000GB	05-Mar-2018

	<h2 style="text-align: center;">Cisco UCS Integrated Infrastructure for Big Data</h2>	TPC-DS 2.7.0 TPC-Pricing 2.1.1
		Report Date: 05-Mar-2018
Total System Cost	Composite Query per Hour Metric	Price / Performance
\$1,008,066 USD	1,580,649 QphDS@10000GB	\$0.64 USD \$/ QphDS@10000GB

Database Size ¹	Database Manager	Operating System	Other Software	Availability date	Clustered Or Not
10,000GB	Transwarp Data Hub V5.1	Red Hat Enterprise Linux Server Release 6.7	No	05-Mar-2018	Yes



Load includes backup = No	RAID = No (using HDFS with 3-way replication)	
System Configuration:	Cisco UCS Integrated Infrastructure for Big Data	
Servers:	17 x Cisco UCS C240 M4 Server	
Total Processors/Cores/Threads:	34/476/952	
Total Memory:	8704 GB	
Total Storage ² :	723321.6 GB	
Storage Ratio ³ :	72.34	
Connectivity:	1 x Cisco Nexus 9372PX Switch 2 x Cisco UCS 6296UP 96-Port Fabric Interconnect	
Server Configuration:	Per node	
Processors:	2 x Intel® Xeon® Processor E5-2680 v4, 2.40 GHz, 35 MB L3	
Memory:	512 GB	
Network:	1 x Cisco VIC 1227 Dual Port 10Gb SFP+	
Storage Device:	24 x 1.8 TB 12G SAS 10K RPM SFF HDD (data nodes) 4 x 1.8 TB 12G SAS 10K RPM SFF HDD (name node) 2 x 240 GB 2.5" 6G SATA SSD (all nodes, boot disk)	
<small>¹ Database Size includes only raw data (i.e., no temp, index, redundant storage space, etc.). ² Total Storage = 16 x DataNode + 1 x NameNode + 17 x BootDisk (i.e. 16 x 24 x 1.8 x 1024 GB + 1 x 4 x 1.8 x 1024 GB + 17 x 2 x 240 GB) ³ Storage Ratio = TotalStorage / SF (i.e. 723321.6 GB / 10000 GB)</small>		
	Cisco UCS Integrated Infrastructure for Big Data	TPC-DS 2.7.0 TPC-Pricing 2.1.1
		Report Date: 05-Mar-2018

Description	Part Number	Brand	Src	Unit Price (\$)	Qty	Extended Price (\$)	3 Year Maint. Price(\$)
UCS C240 M4 SFF 24 HD w/o CPU,mem,HD,PCIe,PS,railkt w/expndr	UCSC-C240-M4SX	Cisco	1	3,995.00	16	63,920.00	
2.40 GHz E5-2680 v4/120W 14C/35MB Cache/DDR4 2400MHz	UCS-CPU-E52680E	Cisco	1	5,259.00	32	168,288.00	
32GB DDR4-2400-MHz RDIMM/PC4-19200/dual rank/x4/1.2v	UCS-MR-1X322RV-A	Cisco	1	2,160.00	256	552,960.00	
240 GB 2.5 inch Enterprise Value 6G SATA SSD (boot)	UCS-SD240GBKS4-EB	Cisco	1	917.00	32	29,344.00	
Right PCI Riser Bd (Riser 1) 2onbd SATA bootdrvs+ 2PCI slts	UCSC-PCI-1C-240M4	Cisco	1	148.00	16	2,368.00	
1.8 TB 12G SAS 10K RPM SFF HDD (4K)	UCS-HD18TB10KS4K	Cisco	1	1,918.00	384	736,512.00	
Cisco UCS VIC1227 VIC MLOM - Dual Port 10Gb SFP+	UCSC-MLOM-CSC-02	Cisco	1	1,499.00	16	23,984.00	
1200W / 800W V2 AC Power Supply for 2U C-Series Servers	UCSC-PSU2V2-1200W	Cisco	1	749.00	32	23,968.00	
AC Power Cord - 250V, 10A - PRC	CAB-250V-10A-CN	Cisco	1	0.00	32	0.00	
Ball Bearing Rail Kit for C220 & C240 M4 & M5 rack servers	UCSC-RAILB-M4	Cisco	1	220.00	16	3,520.00	
IMC SW (Recommended) latest release for C-Series Servers.	CIMC-LATEST	Cisco	1	0.00	16	0.00	
Cisco M4 - v4 CPU asset tab ID label (Auto-Expand)	UCS-M4-V4-LBL	Cisco	1	0.00	16	0.00	
Supercap cable 250mm	UCSC-SCCBL240	Cisco	1	0.00	16	0.00	
Heat sink for UCS C240 M4 rack servers	UCSC-HS-C240M4	Cisco	1	0.00	32	0.00	
Cisco 12G SAS Modular Raid Controller	UCSC-MRAID12G	Cisco	1	656.00	16	10,496.00	
Cisco 12Gbps SAS 2GB FBWC Cache module (Raid 0/1/5/6)	UCSC-MRAID12G-2GB	Cisco	1	1,405.00	16	22,480.00	
Cisco ONE Data Center Compute Opt Out Option	C1UCS-OPT-OUT	Cisco	1	0.00	16	0.00	
UCS SP Select 6296 FI w/18p LIC 2Pk	UCS-SP-FI96-2X	Cisco	1	0.00	1	0.00	
(Not sold Standalone)UCS SP Select 6296 FI w/18p LIC	UCS-SP-FI96P	Cisco	1	18,575.00	2	37,150.00	
Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	CAB-9K12A-NA	Cisco	1	0.00	4	0.00	
10GBASE-SR SFP Module	SFP-10G-SR	Cisco	1	0.00	8	0.00	
10GBASE-CU SFP+ Cable 3 Meter	SFP-H10GB-CU3M	Cisco	1	0.00	8	0.00	
8 Gbps Fibre Channel SW SFP+, LC	DS-SFP-FC8G-SW	Cisco	1	0.00	8	0.00	
UCS 6296UP Chassis Accessory Kit	UCS-ACC-6296UP	Cisco	1	0.00	2	0.00	
UCS 6200 Series Expansion Module Blank	UCS-BLKE-6200	Cisco	1	0.00	6	0.00	
UCS 6296UP Fan Module	UCS-FAN-6296UP	Cisco	1	0.00	8	0.00	
UCS Manager v3.1	N10-MGT014	Cisco	1	0.00	2	0.00	
UCS 6296UP Power Supply/100-240VAC	UCS-PSU-6296UP-AC	Cisco	1	0.00	4	0.00	
UCS C240 M4 SFF 24 HD w/o CPU,mem,HD,PCIe,PS,railkt w/expndr	UCSC-C240-M4SX	Cisco	1	3,995.00	1	3,995.00	
2.40 GHz E5-2680 v4/120W 14C/35MB Cache/DDR4 2400MHz	UCS-CPU-E52680E	Cisco	1	5,259.00	2	10,518.00	
32GB DDR4-2400-MHz RDIMM/PC4-19200/dual rank/x4/1.2v	UCS-MR-1X322RV-A	Cisco	1	2,160.00	16	34,560.00	
240 GB 2.5 inch Enterprise Value 6G SATA SSD (boot)	UCS-SD240GBKS4-EB	Cisco	1	917.00	2	1,834.00	
Right PCI Riser Bd (Riser 1) 2onbd SATA bootdrvs+ 2PCI slts	UCSC-PCI-1C-240M4	Cisco	1	148.00	1	148.00	
1.8 TB 12G SAS 10K RPM SFF HDD (4K)	UCS-HD18TB10KS4K	Cisco	1	1,918.00	4	7,672.00	
Cisco UCS VIC1227 VIC MLOM - Dual Port 10Gb SFP+	UCSC-MLOM-CSC-02	Cisco	1	1,499.00	1	1,499.00	
1200W / 800W V2 AC Power Supply for 2U C-Series Servers	UCSC-PSU2V2-1200W	Cisco	1	749.00	2	1,498.00	
AC Power Cord - 250V, 10A - PRC	CAB-250V-10A-CN	Cisco	1	0.00	2	0.00	
Ball Bearing Rail Kit for C220 & C240 M4 & M5 rack servers	UCSC-RAILB-M4	Cisco	1	220.00	1	220.00	
IMC SW (Recommended) latest release for C-Series Servers.	CIMC-LATEST	Cisco	1	0.00	1	0.00	
Cisco M4 - v4 CPU asset tab ID label (Auto-Expand)	UCS-M4-V4-LBL	Cisco	1	0.00	1	0.00	
Supercap cable 250mm	UCSC-SCCBL240	Cisco	1	0.00	1	0.00	
Heat sink for UCS C240 M4 rack servers	UCSC-HS-C240M4	Cisco	1	0.00	2	0.00	
Cisco 12G SAS Modular Raid Controller	UCSC-MRAID12G	Cisco	1	656.00	1	656.00	
Cisco 12Gbps SAS 2GB FBWC Cache module (Raid 0/1/5/6)	UCSC-MRAID12G-2GB	Cisco	1	1,405.00	1	1,405.00	
Cisco ONE Data Center Compute Opt Out Option	C1UCS-OPT-OUT	Cisco	1	0.00	1	0.00	
UCS 2.5 inch HDD blanking panel	N20-BBLKD	Cisco	1	0.00	20	0.00	
Cisco Smart Net 24X7X4 3Y UCS C240 M4S BD SP Server	CON-OSP-C240V4SP	Cisco	1	1,284.99	17		21,845
2rd Gen FI License to connect C-direct only	UCS-L-6200-10G-C	Cisco	1	925.00	30	27,750.00	
10GBASE-CU SFP+ Cable 3 Meter	SFP-H10GB-CU3M	Cisco	1	100.00	16	1,600.00	
Cisco R42610 standard rack w/side panels	RACK-UCS2	Cisco	1	3,429.00	1	3,429.00	
Cisco Smart Net 24X7X4 3Y UCS 6296UP 2RU Fabre Int/2 PSU/4 Fans	CON-SNTP-FI6296UP	Cisco	1	5,781.00	2		11,562
Cisco Nexus 9372PX Switch	N9K-C9372PX	Cisco	1	22,500.00	1	22,500.00	
Cisco SNTP Cisco Nexus 9372PX Switch, 36 24x7	CON-SNTP-9372PX	Cisco	1	5,256.00	1		5,256
Red Hat Enterprise Linux Server, 3Y 24x7	CON-ISV1-EL2S2V3A	Cisco	1	3,897.00	17	66,249.00	
					Total	1,860,523.00	38,663
Large Purchase Discount ¹	61% for products and 35% for service	Cisco	1			-1,134,919.03	-13,532
Acer V206HQL-LED monitor - 20" (Inc 2 spares)	UM.1V6AA.A02	CDW	2	78.99	3	236.97	
Logitech USB Corded Keyboard/Mouse Combo MK120 (Inc 2 spares)	920-002565	CDW	2	17.99	3	53.97	
Transwarp Data Hub Professional Version 5.1 with 3 years 24x7 support	TDH-BD-SUITE-P	Transwarp	3	21,600.00	17	367,200.00	
Large Purchase Discount ¹	30% for Transwarp Data Hub	Transwarp	3			-110,160.00	

1= Cisco, 2 = DW.com, 3 = TDH

¹All 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.

Audited by Francois Raab, InfoSizing

Three-Year Cost of Ownership: \$1,008,066 USD

QphDS@10000GB: 1,580,649

\$ / QphDS@10000GB: \$0.64 USD

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.



Cisco UCS Integrated Infrastructure for Big Data

TPC-DS 2.7.0
TPC-Pricing 2.1.1

Report Date:
05-Mar-2018

Metrics Details:

Name	Value	Unit
Scale Factor (SF)	10,000	GB
Streams	4	Stream
Queries (Q)	396	Queries
T_load	4,778.8	Second
T_ld	0.0531	Hour
T_pt	24.0173	Hour
T_tt1	66,440.1	Second
T_tt2	68,173.3	Second
T_dm1	1,502.9	Second
T_dm2	1,471.0	Second
T_tt	37.3927	Hour
T_dm	0.8261	Hour

Load Step	Start	End	(sec.)	(hh:mm:ss)
Build	01/23/18 19:22:46.21	01/23/18 20:38:19.77	4,533.56	1:15:34
Audit	01/23/18 20:38:19.77	01/23/18 22:12:48.93	5,669.16	1:34:29
Get Stats	01/23/18 22:12:48.93	01/23/18 22:16:54.15	245.22	0:04:05
Reported	01/23/18 19:22:46.21	01/23/18 22:16:54.15	4,778.78	1:19:39

Test	Start	End	(sec.)	(hh:mm:ss)
Power	01/23/18 22:28:21.92	01/24/18 04:28:37.41	21,615.49	6:00:15
Thruput-1	01/24/18 04:28:37.41	01/24/18 22:55:57.42	66,440.01	18:27:20
Thruput-2	01/24/18 23:21:00.31	01/25/18 18:17:13.57	68,173.26	18:56:13
DM-1	01/24/18 22:55:57.42	01/24/18 23:21:00.30	1,502.88	0:25:03
DM-2	01/25/18 18:17:13.57	01/25/18 18:41:44.52	1,470.95	0:24:31

Stream	Start	End	(sec.)	(hh:mm:ss)
Pt - 0	01/23/18 22:28:21.92	01/24/18 04:28:37.41	21,615.49	6:00:15
Tt1 - 1	01/24/18 04:28:37.41	01/24/18 22:43:57.06	65,719.65	18:15:20
Tt1 - 2	01/24/18 04:28:37.42	01/24/18 22:55:57.42	66,440.00	18:27:20
Tt1 - 3	01/24/18 04:28:37.42	01/24/18 22:35:12.92	65,195.50	18:06:35
Tt1 - 4	01/24/18 04:28:37.41	01/24/18 22:34:14.87	65,137.46	18:05:37
Tt2 - 5	01/24/18 23:21:00.31	01/25/18 18:09:14.08	67,693.77	18:48:14
Tt2 - 6	01/24/18 23:21:00.31	01/25/18 18:15:04.29	68,043.98	18:54:04
Tt2 - 7	01/24/18 23:21:00.31	01/25/18 18:13:57.73	67,977.42	18:52:57
Tt2 - 8	01/24/18 23:21:00.31	01/25/18 18:17:13.57	68,173.26	18:56:13
DMt1 - 1	01/24/18 22:55:57.42	01/24/18 23:09:58.30	840.88	0:14:01
DMt1 - 2	01/24/18 23:09:58.30	01/24/18 23:21:00.30	662.00	0:11:02
DMt2 - 3	01/25/18 18:17:13.57	01/25/18 18:29:41.85	748.28	0:12:28
DMt2 - 4	01/25/18 18:29:41.85	01/25/18 18:41:44.52	722.67	0:12:03

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	29.5	454.2	551.4	482.5	769.8	454.2	475.4	517.0	606.0	769.8	585.8	666.6	732.1	468.1	468.1	556.4	626.2	683.0	732.1
2	374.5	465.7	528.9	541.2	978.4	465.7	513.1	535.1	650.5	978.4	662.5	1,141.3	1,434.4	680.9	662.5	676.3	911.1	1,214.6	1,434.4
3	96.4	611.1	128.1	390.3	172.5	128.1	161.4	281.4	445.5	611.1	175.8	418.8	130.2	528.5	130.2	164.4	297.3	446.2	528.5
4	1,111.7	2,472.0	2,234.9	1,843.1	1,625.7	1,625.7	1,788.8	2,039.0	2,294.2	2,472.0	1,856.4	1,562.3	1,460.8	1,673.4	1,460.8	1,536.9	1,617.9	1,719.2	1,856.4
5	89.0	118.1	253.1	113.5	252.9	113.5	117.0	185.5	253.0	253.1	328.4	668.3	472.7	431.4	328.4	405.7	452.1	521.6	668.3
6	10.4	135.8	256.8	127.5	568.7	127.5	133.7	196.3	334.8	568.7	648.0	30.7	225.4	923.6	30.7	176.7	436.7	716.9	923.6
7	61.6	80.4	139.7	89.0	349.0	80.4	86.9	114.4	192.0	349.0	88.5	598.6	508.1	307.4	88.5	252.7	407.8	530.7	598.6
8	12.3	83.5	402.2	33.9	260.0	33.9	71.1	171.8	295.6	402.2	690.3	396.9	80.5	201.8	80.5	171.5	299.4	470.3	690.3
9	1,079.2	3,876.2	3,307.6	4,324.8	4,512.5	3,307.6	3,734.1	4,100.5	4,371.7	4,512.5	5,277.3	4,382.1	4,312.2	5,028.8	4,312.2	4,364.6	4,705.5	5,090.9	5,277.3
10	13.7	25.2	68.6	27.2	25.8	25.2	25.7	26.5	37.6	68.6	122.9	546.6	358.4	290.1	122.9	248.3	324.3	405.5	546.6
11	436.5	821.1	921.5	1,608.9	544.6	544.6	752.0	871.3	1,093.4	1,608.9	933.1	659.1	507.3	1,072.6	507.3	621.2	796.1	968.0	1,072.6
12	4.2	551.2	32.5	672.9	566.8	32.5	421.5	559.0	593.3	672.9	15.5	7.9	15.4	31.3	7.9	13.5	15.5	19.5	31.3
13	101.9	209.6	153.5	346.9	183.1	153.5	175.7	196.4	243.9	346.9	203.8	256.6	346.9	330.5	203.8	243.4	293.6	334.6	346.9
14	1,094.0	2,865.6	4,468.6	4,071.4	3,224.8	2,865.6	3,135.0	3,648.1	4,170.7	4,468.6	2,592.8	2,915.7	5,077.0	4,139.3	2,592.8	2,835.0	3,527.5	4,373.7	5,077.0
15	11.3	224.1	115.8	72.3	36.0	36.0	63.2	94.1	142.9	224.1	106.9	136.2	55.4	201.0	55.4	94.0	121.6	152.4	201.0
16	186.5	195.7	570.6	491.0	755.2	195.7	417.2	530.8	616.8	755.2	488.3	236.2	301.1	1,267.1	236.2	284.9	394.7	683.0	1,267.1
17	67.6	529.8	346.7	98.0	350.5	98.0	284.5	348.6	395.3	529.8	359.7	100.5	290.0	88.4	88.4	97.5	195.3	307.4	359.7
18	104.3	1,518.2	391.4	168.6	598.7	168.6	335.7	495.1	828.6	1,518.2	293.4	329.8	325.4	634.2	293.4	317.4	327.6	405.9	634.2
19	23.3	520.6	1,542.3	76.2	781.6	76.2	409.5	651.1	971.8	1,542.3	49.0	173.3	141.2	614.5	49.0	118.2	157.3	283.6	614.5
20	5.9	216.0	613.1	138.3	89.0	89.0	126.0	177.2	315.3	613.1	14.7	92.0	849.4	333.3	14.7	72.7	212.7	462.3	849.4
21	4.1	152.6	4.3	135.0	79.1	4.3	60.4	107.1	139.4	152.6	174.5	607.9	365.6	100.0	100.0	155.9	270.1	426.2	607.9
22	20.3	19.3	94.7	340.2	443.5	19.3	75.9	217.5	366.0	443.5	791.6	409.6	107.9	48.0	48.0	92.9	258.8	505.1	791.6
23	1,608.9	3,752.8	3,588.1	5,314.3	2,467.6	2,467.6	3,308.0	3,670.5	4,143.2	5,314.3	4,146.5	3,171.6	4,205.5	3,299.2	3,171.6	3,267.3	3,722.9	4,161.3	4,205.5
24	460.3	2,092.8	769.1	1,041.1	644.6	644.6	738.0	905.1	1,304.0	2,092.8	1,568.0	2,554.0	1,076.6	2,128.3	1,076.6	1,445.2	1,848.2	2,234.7	2,554.0
25	43.9	83.4	333.0	1,066.1	145.0	83.4	129.6	239.0	516.3	1,066.1	778.0	470.1	460.9	441.8	441.8	456.1	465.5	547.1	778.0
26	35.9	147.8	412.7	57.1	186.0	57.1	125.1	166.9	242.7	412.7	300.9	158.4	228.5	427.9	158.4	211.0	264.7	332.5	427.9
27	69.7	982.5	923.3	340.7	203.8	203.8	306.5	632.0	938.1	982.5	373.7	164.9	250.4	632.7	164.9	229.0	312.1	438.5	632.7
28	976.4	2,248.0	1,790.1	1,413.7	2,079.9	1,413.7	1,696.0	1,935.0	2,121.9	2,248.0	2,357.1	1,792.2	1,359.4	1,129.2	1,129.2	1,301.9	1,575.8	1,933.4	2,357.1
29	86.7	104.2	745.3	171.9	336.7	104.2	155.0	254.3	438.9	745.3	115.3	390.3	127.6	152.9	115.3	124.5	140.3	212.3	390.3
30	23.0	486.3	439.1	841.2	108.7	108.7	356.5	462.7	575.0	841.2	455.6	379.4	52.2	153.1	52.2	127.9	266.3	398.5	455.6
31	72.3	1,670.9	297.2	320.4	330.4	297.2	314.6	325.4	665.5	1,670.9	789.5	1,009.9	1,124.8	76.4	76.4	61.2	899.7	1,038.6	1,124.8
32	15.4	23.6	48.0	22.3	51.5	22.3	23.3	35.8	48.9	51.5	334.2	969.3	55.8	573.9	55.8	264.6	454.1	672.8	969.3
33	13.3	1,089.7	467.5	220.4	763.3	220.4	405.7	615.4	844.9	1,089.7	191.0	347.6	274.0	15.8	15.8	147.2	232.5	292.4	347.6
34	106.3	118.2	887.8	187.0	277.9	118.2	169.8	232.5	430.4	887.8	289.6	385.3	276.3	257.7	257.7	271.7	283.0	313.5	385.3
35	36.5	302.6	1,037.9	110.7	382.5	110.7	254.6	342.6	546.4	1,037.9	100.1	79.9	316.5	282.4	79.9	95.1	191.3	290.9	316.5
36	73.0	336.2	353.0	300.5	456.8	300.5	327.3	344.6	379.0	456.8	429.3	200.2	393.6	429.5	200.2	345.3	411.9	429.5	456.8
37	48.6	164.1	305.0	78.9	62.9	62.9	74.9	121.5	199.3	305.0	173.5	207.8	217.6	206.3	173.5	198.1	207.1	210.3	217.6
38	147.1	185.3	485.1	411.9	268.3	185.3	247.6	340.1	430.2	485.1	206.8	1,063.1	447.9	200.1	200.1	205.1	327.4	601.7	1,063.1
39	21.4	163.9	775.8	453.2	527.4	163.9	380.9	490.3	589.5	775.8	155.8	137.9	526.6	345.1	137.9	151.3	250.5	390.5	526.6
40	21.9	642.6	62.6	136.8	285.8	62.6	118.3	211.3	375.0	642.6	57.4	527.9	549.2	449.4	57.4	351.4	488.7	533.2	549.2
41	3.3	219.8	39.8	106.3	5.8	5.8	31.3	73.1	134.7	219.8	375.2	4.0	79.0	206.9	4.0	60.3	143.0	249.0	375.2
42	15.3	177.8	199.6	50.5	205.1	50.5	146.0	188.7	201.0	205.1	36.2	53.7	50.3	282.1	36.2	46.8	52.0	110.8	282.1
43	79.1	521.1	100.1	231.3	486.1	100.1	198.5	358.4	494.9	521.1	884.2	102.1	77.9	239.2	77.9	96.1	170.7	400.5	884.2
44	349.7	1,106.3	581.9	1,198.8	836.8	581.9	773.1	971.6	1,129.4	1,198.8	609.3	771.4	798.2	587.3	587.3	603.8	690.4	778.1	978.2
45	10.7	12.2	68.7	85.3	322.6	12.2	54.6	77.0	144.6	322.6	416.5	124.5	245.7	157.1	124.5	149.0	201.4	288.4	416.5
46	125.6	325.6	122.9	518.9	221.4	122.9	196.8	273.5	373.9	518.9	429.3	184.5	403.4	182.2	182.2	183.9	294.0	409.9	429.3
47	215.5	713.0	218.1	492.1	355.5	218.1	321.2	423.8	547.3	713.0	888.7	838.0	546.0	515.3	515.3	538.3	692.0	850.7	888.7
48	100.8	200.1	99.1	781.1	2,176.7	99.1	174.9	490.6	1,130.0	2,176.7	785.5	223.8	412.6	124.4	124.4	199.0	318.2	505.8	785.5
49	68.4	457.3	132.6	287.0	97.0	97.0	123.7	209.8	329.6	457.3	141.5	398.7	1,144.3	449.7	141.5	334.4	424.2	623.4	1,144.3
50	669.2	1,107.6	1,210.4	1,212.3	1,275.4	1,107.6	1,184.7	1,211.4	1,228.1	1,275.4	1,874.9	883.0	926.3	1,826.2	883.0	915.5	1,376.3	1,838.4	1,874.9
51	116.4	188.2	217.0	426.2	273.1	188.2	209.8	245.1	311.4	426.2	260.8	498.9	174.7	548.9	174.7	239.3	379.9	511.4	548.9
52	16.0	976.3	188.7	411.4	88.4	88.4	163.6	300.1	552.6	976.3	22.7	94.5	185.6	621.4	22.7	76.6	140.1	294.6	621.4
53	41.8	57.9	356.0	337.2	337.9	57.9	267.4	337.6	342.4	356.0	371.0	615.4	485.7	179.7	179.7	323.2	428.4	518.1	615.4
54	14.9	641.7	182.2	16.1	206.7	16.1	140.7	194.5	315.5	641.7	1,265.3	24.6	226.3	406.6	24.6	175.9	316.5	621.3	1,265.3
55	14.3	54.6	691.8	72.7	217.1	54.6	68.2	144.9	335.8	691.8	235.2	548.5	397.8	45.5	45.5	187.8	316.5	435.5	548.5
56	11.3	751.9	166.0	265.0	260.8	166.0	237.1	262.9	386.7	751.9	50.6	40.0	411.1	352.1	40.0	48.0	201.4	366.9	411.1
57	105.9	478.0	309.9	313.2	212.5	212.5	285.6	311.6	354.4	478.0	677.9	255.6	279.3	285.0	255.6	273.4	282.2	383.2	677.9
58	257.4	337.5	602.6	487.2	579.9	337.5	449.8	533.6	585.6	602.6	377.6	502.2	1,055.1	615.1	377.6	471.1	558.7	725.1	1,055.1
59	435.5	543.1	667.9	766.4	932.9	543.1	636.7	717.2	808.0	932.9	1,166.8	790.0	1,262.0	593.7	593.7	740.9	978.4	1,190.6	1,262.0
60	22.3	126.7	383.5	255.8	193.4	126.7	176.7												

88	740.4	2,610.0	1,152.7	2,537.5	1,277.9	1,152.7	1,246.6	1,907.7	2,555.6	2,610.0	1,563.1	1,903.7	1,294.4	2,022.6	1,294.4	1,495.9	1,733.4	1,933.4	2,022.6
89	50.8	82.9	52.7	179.7	135.5	52.7	75.4	109.2	146.6	179.7	710.5	275.9	70.6	105.4	70.6	96.7	190.7	384.6	710.5
90	47.4	106.1	817.9	59.7	362.7	59.7	94.5	234.4	476.5	817.9	74.0	447.3	244.7	331.3	74.0	202.0	288.0	360.3	447.3
91	7.3	9.5	547.6	20.0	96.5	9.5	17.4	58.3	209.3	547.6	50.7	8.8	364.3	552.3	8.8	40.2	207.5	411.3	552.3
92	8.9	229.1	10.6	176.4	23.3	10.6	20.1	99.9	189.6	229.1	330.1	320.1	619.4	239.1	239.1	299.9	325.1	402.4	619.4
93	292.7	843.5	344.7	378.4	332.5	332.5	341.7	361.6	494.7	843.5	693.2	544.9	457.5	924.7	457.5	523.1	619.1	751.1	924.7
94	108.4	311.0	503.4	153.3	259.2	153.3	232.7	285.1	359.1	503.4	329.3	132.8	248.4	470.8	132.8	219.5	288.9	364.7	470.8
95	720.2	1,280.9	2,588.6	1,029.2	1,770.5	1,029.2	1,218.0	1,525.7	1,975.0	2,588.6	1,814.8	892.4	1,091.9	1,707.3	892.4	1,042.0	1,399.6	1,734.2	1,814.8
96	82.9	416.1	298.2	376.7	254.7	254.7	287.3	337.5	386.6	416.1	538.0	206.6	519.1	102.5	102.5	180.6	362.9	523.8	538.0
97	317.9	972.7	784.5	480.6	585.8	480.6	559.5	685.2	831.6	972.7	603.6	655.0	817.5	563.3	563.3	593.5	629.3	695.6	817.5
98	15.7	214.5	85.1	488.2	748.6	85.1	182.2	351.4	553.3	748.6	366.6	306.2	296.3	1,054.9	296.3	303.7	336.4	538.7	1,054.9
99	128.0	428.3	627.9	229.8	438.3	229.8	378.7	433.3	485.7	627.9	341.1	285.7	258.9	159.3	159.3	234.0	272.3	299.6	341.1

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	98.3	105.9	104.5	102.0	98.3	101.1	103.3	104.9	105.9
LF_CS	90.3	81.0	102.0	90.4	81.0	88.0	90.4	93.3	102.0
LF_I	87.2	103.1	125.9	85.1	85.1	86.7	95.2	108.8	125.9
LF_SR	81.1	75.3	76.9	78.4	75.3	76.5	77.7	79.1	81.1
LF_SS	102.9	82.8	77.0	111.3	77.0	81.4	92.9	105.0	111.3
LF_WR	70.1	66.0	90.0	63.0	63.0	65.3	68.1	75.1	90.0
LF_WS	56.5	62.1	71.0	57.7	56.5	57.4	59.9	64.3	71.0
DF_CS	341.2	275.2	331.4	361.1	275.2	317.4	336.3	346.2	361.1
DF_SS	562.8	338.1	379.8	360.8	338.1	355.1	370.3	425.6	562.8
DF_WS	439.1	276.3	332.7	313.3	276.3	304.1	323.0	359.3	439.1
DF_I	31.1	25.9	30.1	30.9	25.9	29.1	30.5	31.0	31.1

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 Cisco Systems, 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.

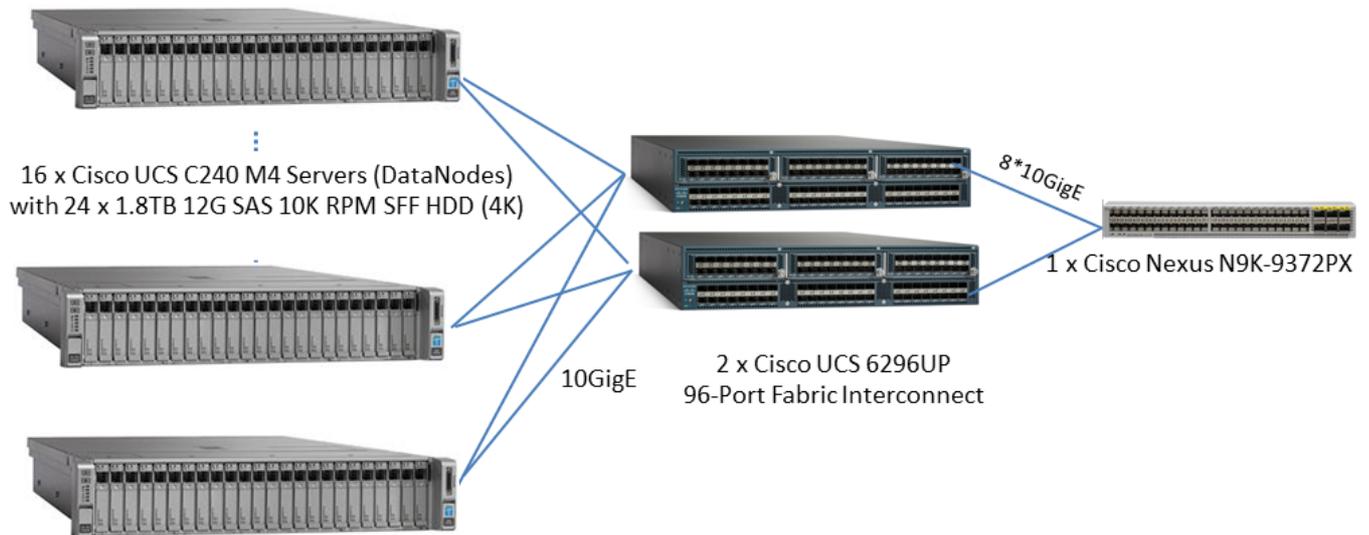
The Supporting File Archive contains the Operating System and DBMS parameters used in this benchmark.

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



1 x Cisco UCS C240 M4 Servers (NameNode)
with 4 x 1.8TB 12G SAS 10K RPM SFF HDD (4K)

Boot(All servers): 2 x 240 GB 2.5 inch Enterprise
Value 6G SATA SSD (BOOT)

The measured configuration consisted of:

- Total Nodes: 17
- Total Processors/Cores/Threads: 34/476/952
- Total Memory: 8704 GB
- Total Number of Storage Drives/Devices: 422
- Total Storage Capacity: 723321.6 GB

Server nodes details:

- 17 x Cisco UCS C240 M4 Servers, each with:
 - ◆ Processors/Cores/Threads: 2/28/56
 - ◆ Processor Model: 2 x Intel® Xeon® Processor E5-2680 v4, 2.40 GHz, 35 MB L3
 - ◆ Memory: 512GB
 - ◆ Controller: 1 x Cisco 12G SAS Modular Raid Controller
 - ◆ Drives:
 - ✓ 24 x 1.8 TB 10K RPM SAS HDD (data nodes)
 - ✓ 4 x 1.8 TB 10K RPM SAS HDD (name node)
 - ✓ 2 x 240 GB 2.5" 6G SATA SSD (all nodes, boot disk)
 - ◆ Network: 1 x Cisco VIC 1227 Dual Port 10Gb SFP+

Network connectivity detail:

1. 1 x Cisco Nexus 9372PX Switch
2. 2 x Cisco UCS 6296UP 96-Port Fabric Interconnect

Transwarp System Components Configuration

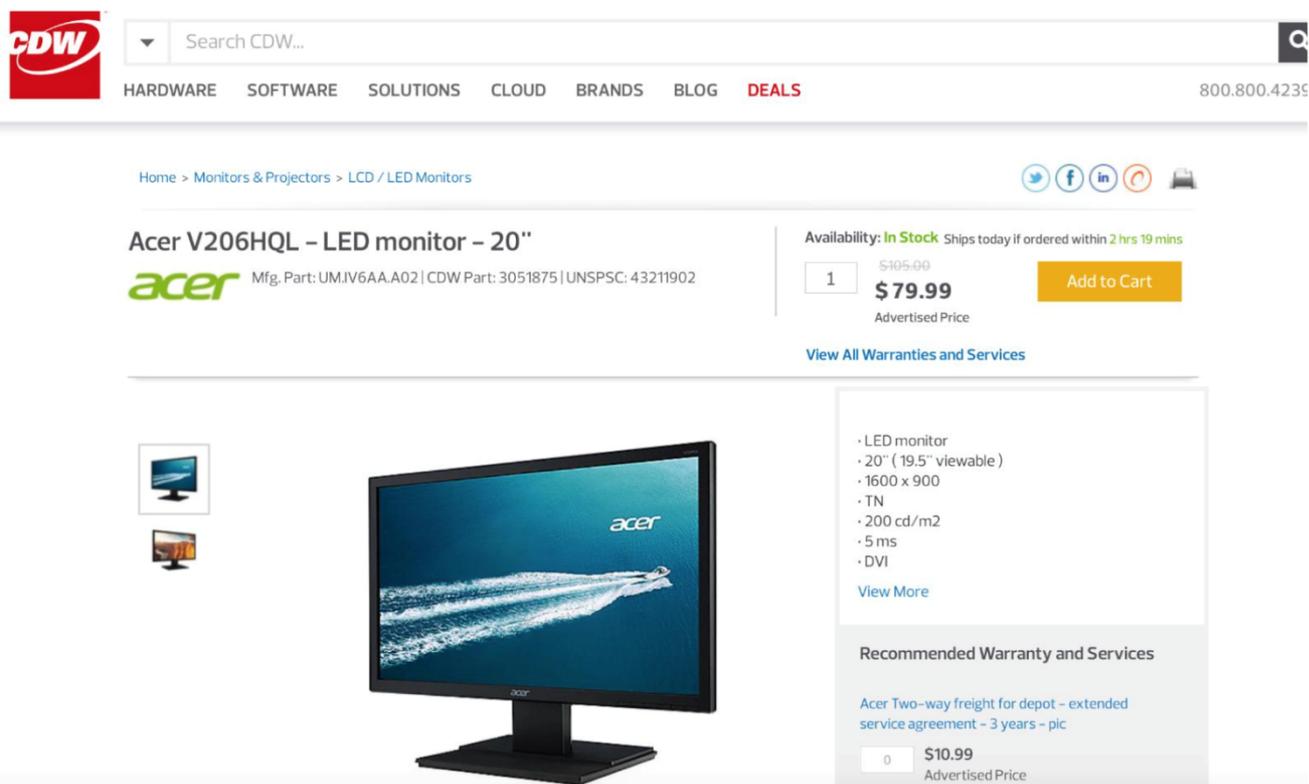
Node	Map/Reduce		HDFS		ZooKeeper	Transwarp Inceptor	
	Resource Manager	Node Manager	NameNode	DataNode	QuorumPeer	Server	Executor
1	x		x		x	x	
2-9		x		x			x
10		x	x	x	x		x
11		x		x	x		x
12-17		x		x			x

Note: 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 10 is the standby one. In all other part of the documents, NameNode is hardware Node 1, on which the active NameNode service runs.

Priced Configuration

There are no differences between the priced and measured configurations.

Third Party Price Quotes



CDW

Search CDW...

HARDWARE SOFTWARE SOLUTIONS CLOUD BRANDS BLOG DEALS 800.800.4235

Home > Monitors & Projectors > LCD / LED Monitors

Acer V206HQL - LED monitor - 20"

acer Mfg. Part: UM.IV6AA.A02 | CDW Part: 3051875 | UNSPSC: 43211902

Availability: **In Stock** Ships today if ordered within 2 hrs 19 mins

1 ~~\$105.00~~
\$79.99
Advertised Price

Add to Cart

View All Warranties and Services

- LED monitor
- 20" (19.5" viewable)
- 1600 x 900
- TN
- 200 cd/m2
- 5 ms
- DVI

View More

Recommended Warranty and Services

Acer Two-way freight for depot - extended service agreement - 3 years - pic

0 ~~\$10.99~~
\$10.99
Advertised Price



Search CDW...



HARDWARE SOFTWARE SOLUTIONS CLOUD BRANDS BLOG DEALS

800.800.4239

Home > Computer Accessories > Keyboards & Keypads > Keyboard and Mouse Bundle



Logitech MK120 USB Wired Keyboard/Mouse Set – \$4.50 svgs while supplies last

logitech Mfg. Part: 920-002565 | CDW Part: 2124292 | UNSPSC: 43211706

★★★★★ [Read 1 review](#) | [Write a review](#)

Availability: **In Stock** Ships today if ordered within 2 hrs 20 mins

1

~~\$19.99~~
\$17.99
Advertised Price

[Add to Cart](#)



- Keyboard
- mouse
- Wired
- PC compatible
- USB interface
- standard version

[View More](#)

[Log On to Email this page or Save as Favorite](#)

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. The partition granularity is by month.

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 HDFS 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	2,879,513,368
store_sales	28,800,426,268
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
2-17	Cisco 12G MegaRAID SAS	1-24 (HDD)	Table, Temp
1 - 17	Intel Chipset Embedded SATA RAID	0 (2 SSD, RAID-1, Internal)	Operating system, root, swap, Transwarp data hub, logs

All the Table contents were on HDFS, which are distributed among the disks in datanodes:

/data/disk1/hadoop/data,/data/disk10/hadoop/data, /data/disk11/hadoop/data,/data/disk12/hadoop/data,
/data/disk13/hadoop/data,/data/disk14/hadoop/data,/data/disk15/hadoop/data,/data/disk16/hadoop/data,
/data/disk17/hadoop/data,/data/disk18/hadoop/data,/data/disk19/hadoop/data,/data/disk2/hadoop/data,
/data/disk20/hadoop/data,/data/disk21/hadoop/data,/data/disk22/hadoop/data,/data/disk23/hadoop/data,
/data/disk24/hadoop/data,/data/disk3/hadoop/data,/data/disk4/hadoop/data,/data/disk5/hadoop/data,
/data/disk6/hadoop/data,/data/disk7/hadoop/data,/data/disk8/hadoop/data,/data/disk9/hadoop/data

Table size on HDFS:

9.8 K /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/call_center
1.4 M /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/catalog_page
80.3 G /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/catalog_returns
743.0 G /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/catalog_sales
2.8 G /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/customer
525.7 M /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/customer_address
9.9 M /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/customer_demographics
260.5 K /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/date_dim
2.7 K /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/household_demographics
902 /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/income_band
4.5 G /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/inventory
31.3 M /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/item
60.4 K /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/promotion
1.6 K /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/reason
1.6 K /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/ship_mode
104.2 K /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/store
114.8 G /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/store_returns
986.1 G /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/store_sales
130.2 K /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/time_dim
3.1 K /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/warehouse
45.5 K /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/web_page
38.1 G /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/web_returns
367.4 G /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/web_sales
11.6 K /inceptorsql1/user/hive/warehouse/tpcds_torc_10000.db/hive/web_site

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

/data/disk1/hadoop/ngmr/inceptorsql1, /data/disk10/hadoop/ngmr/inceptorsql1,
/data/disk11/hadoop/ngmr/inceptorsql1, /data/disk12/hadoop/ngmr/inceptorsql1,
/data/disk13/hadoop/ngmr/inceptorsql1, /data/disk14/hadoop/ngmr/inceptorsql1,
/data/disk15/hadoop/ngmr/inceptorsql1, /data/disk16/hadoop/ngmr/inceptorsql1,
/data/disk17/hadoop/ngmr/inceptorsql1, /data/disk18/hadoop/ngmr/inceptorsql1,
/data/disk19/hadoop/ngmr/inceptorsql1, /data/disk2/hadoop/ngmr/inceptorsql1,
/data/disk20/hadoop/ngmr/inceptorsql1, /data/disk21/hadoop/ngmr/inceptorsql1,
/data/disk22/hadoop/ngmr/inceptorsql1, /data/disk23/hadoop/ngmr/inceptorsql1,
/data/disk24/hadoop/ngmr/inceptorsql1, /data/disk3/hadoop/ngmr/inceptorsql1,
/data/disk4/hadoop/ngmr/inceptorsql1, /data/disk5/hadoop/ngmr/inceptorsql1,
/data/disk6/hadoop/ngmr/inceptorsql1, /data/disk7/hadoop/ngmr/inceptorsql1,
/data/disk8/hadoop/ngmr/inceptorsql1, /data/disk9/hadoop/ngmr/inceptorsql1

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. The partition granularity is by month.

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

The database tables were on top of Hadoop file system (HDFS) , HDFS does distributed 3-way replication.

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 2.6.0 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.

No changes were made to Dsdgen between version 2.6.0 and 2.7.0.

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 4778.8 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.

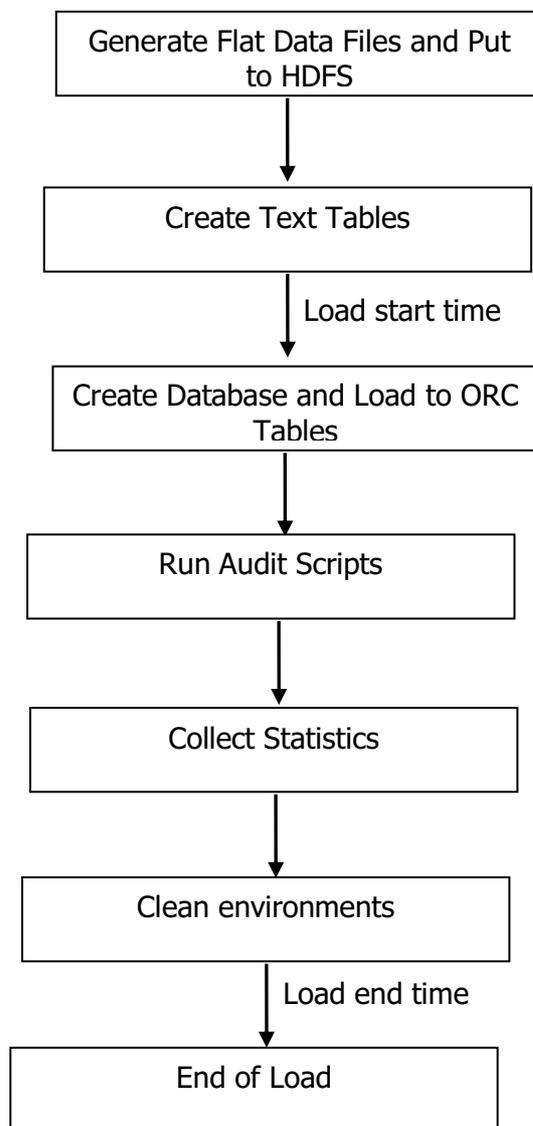
The data storage ratio is $723321.6/10000 = 72.34$.

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 2.6.0 and dsqgen version 2.7.0 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 2.7.0 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
- Add Parentheses around union all (MQM g.2)
 - Q2,
- Ordinal position in ORDER BY replaced by column name (MQM g.3)
Issue: dsqgen can generate a random 3rd column.
Auditor's interpretation: Use 3rd column from template, regardless of columns generated by dsqgen.
 - Q47
 - Q57

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])

TPC template has typo (correction approved for version 2.8.0):

- Q70a
 - a "(" is missing in front of "select total_sum ,s_state"

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 a disk drive during the first throughput test. The disk was failed by disabling RW access. 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 = 1,580,649

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 Inceptor is the SQL engine in Transwarp Data Hub (TDH). The TDH is a Hadoop based system. All data are saved in Hadoop HDFS. And any data on HDFS is 3 replications by default. All the files are managed by HDFS.

beeline is the client in Transwarp Inceptor. It connects to the Inceptor 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 server (Inceptor) 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, rhell is configured as the Inceptor server, all other 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 March 02, 2018.

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: Manankumar Trivedi
Cisco Systems Inc.
3800 Zanker Road
San Jose, CA 95134

March 2, 2018

I verified the TPC Benchmark™ DS (TPC-DS™ v2.7.0) performance of the following configuration:

Platform: Cisco UCS Integrated Infrastructure for Big Data
Operating System: Red Hat Enterprise Linux Server Release 6.7
Database Manager: Transwarp Data Hub V5.1

The results were:

Performance Metric 1,580,649 QphDS@10000GB
Database Load Time 1h 19m 39s

Servers **Cisco UCS C240 M4 (16 data nodes + 1 name node)**

CPUs	2 x Intel Xeon Processor E5-2680 v4 (2.4GHz, 35MB L3)		
Memory	512 GB		
Disks	Qty	Size	Type
	2	240 GB	2.5" 6G SATA SSD (all nodes)
	4	1.8 TB	10K RPM SAS HDD (name nodes only)
	24	1.8 TB	10K RPM SAS HDD (data node only)

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:

This is the first publication of a TPC-DS result. In the course of the benchmark execution and the independent audit process, a number of issues were raised with the benchmark maintenance subcommittee. These issues were resolved, sometimes resulting in changes to the benchmark specification. While this result was audited against version 2.7.0 of the benchmark, it also takes advantage of some already approved changes that are intended for release in version 2.8.0 of the benchmark.

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-Inceptor/
	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/

Transwarp Data Hub



Cisco

3800 Zanker Road,

San Jose, CA 95134

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	Discount	Extended Price
TDH-BD-SUITE-P	Transwarp Data Hub Professional Version 5.1 with 3 years 24x7 support	\$21,600 per node	17	30% (Large Purchase Discount)	\$257,040

Jun Zheng (Transwarp Business Development Director)

jun.zheng@transwarp.io