TPC Benchmark™ C

Full Disclosure Report

First Edition

18–Aug–2020

Using

Goldilocks v3.1 Standard Edition

on

Supermicro A+ Server 1124US-TNRP

First Edition: 18-Aug-2020

TTA, Telecommunications Technology Association, believes that all the information in this document is accurate as of the publication date. The information in this document is subject to change without notice. TTA, the sponsor of this benchmark test, assumes no responsibility for any errors that may appear in this document. The pricing information in this document is believed to accurately reflect the current prices as of the publication date. However, the sponsor provides no warranty of the pricing information in this document.

Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, the TPC Benchmark™ C should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or product evaluation decisions are contemplated.

All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly. No warranty of system performance or price/performance is expressed or implied in this report.

Trademarks

The following terms used in this publication are trademarks of other companies as follows:

- *TPC Benchmark, TPC-C, and tpmC are trademarks of the Transaction Processing Performance Council*
- *TTA is a registered trademark of Telecommunications Technology Association*
- *Goldilocks is a registered trademark of SUNJESOFT, Inc.*
- *JBoss is a registered trademark of RedHat, Inc.*
- *AMD and EPYC are trademarks of AMD Corporation..*
- *All other trademarks and copyrights are properties of their respective owners.*

Table of Contents

Abstract

This report documents the methodology and results of the TPC Benchmark™ C (TPC-C) test conducted by TTA on the Goldilocks v3.1 Standard Edition on Supermicro A+ Server 1124US-TNRP

Goldilocks v3.1 Standard Edition on Supermicro A+ Server 1124US-TNRP

Company Name	System Name	Database Software	Operating System
Telecommunications Technology Association	Supermicro A+ Server 1124US-TNRP	Goldilocks v3.1 Standard Edition	RedHat Enterprise Linux 7.7

TPC Benchmark™ C Metrics

Total System Cost	TPC-C Throughput	Price/Performance	Availability Date
₩ 248,570,400 (KRW)	380,475 tpmC	654 KRW/tpmC	1-Sep-2020

Preface

The Transaction Processing Performance Council (TPC™) is a non-profit corporation founded to define transaction processing and database benchmarks and to disseminate objective, verifiable TPC performance data to the industry. The TPC Benchmark© C is an on-line transaction processing benchmark (OLTP) developed by the TPC.

TPC Benchmark™ C Overview

TPC Benchmark™ C (TPC-C) simulates a complete computing environment where a population of users executes transactions against a database. The benchmark is centered around the principal activities (transactions) of an order-entry environment. These transactions include entering and delivering orders, recording payments, checking the status of orders, and monitoring the level of stock at the warehouses. While the benchmark portrays the activity of a wholesale supplier, TPC-C is not limited to the activity of any particular business segment, but, rather represents any industry that must manage, sell, or distribute a product or service.

TPC-C consists of a mixture of read-only and update intensive transactions that simulate the activities found in complex OLTP application environments. It does so by exercising a breadth of system components associated with such environments, which are characterized by:

- *The simultaneous execution of multiple transaction types that span a breadth of complexity*
- *On-line and deferred transaction execution modes*
- *Multiple on-line terminal sessions*
- *Moderate system and application execution time*
- *Significant disk input/output*
- *Transaction integrity (ACID properties)*
- *Non-uniform distribution of data access through primary and secondary keys*
- *Databases consisting of many tables with a wide variety of sizes, attributes, and relationships*
- *Contention of data access and update*

The performance metric reported by TPC-C is a "business throughput" measuring the number of orders processed per minute. Multiple transactions are used to simulate the business activity of processing an order, and each transaction is subject to a response time constraint. The performance metric for this benchmark is expressed in transactions-per-minute-C (tpmC). To be compliant with the TPC-C standard, all references to tpmC results must include the tpmC rate, the associated price-per-tpmC, and the availability date of the priced configuration.

TPC-C uses terminology and metrics that are similar to other benchmarks, originated by the TPC or others. Such similarity in terminology does not in any way imply that TPC-C results are comparable to other benchmarks. The only benchmark results comparable to TPC-C are other TPC-C results conformant with the same revision.

Despite the fact that this benchmark offers a rich environment that emulates many OLTP applications, this benchmark does not reflect the entire range of OLTP requirements. In addition, the extent to which a customer can achieve the results reported by a vendor is highly dependent on how closely TPC-C approximates the customer application. The relative performance of systems derived from this benchmark does not necessarily hold for other workloads or environments. Extrapolations to other environments are not recommended.

Benchmark results are highly dependent upon workload, specific application requirements, and systems design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, TPC-C should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or product evaluation decisions are contemplated.

Further information is available at www.tpc.org

	Goldilocks v3.1 Standard Edition on Supermicro A+ Server 1124US-TNRP			**TPC-C Version 5.11.0** **TPC Pricing 2.6.0** Report Date **18-Aug-2020**
Total System Cost	TPC-C Throughput	Price/Performance		Availability Date
₩248,570,400 (KRW)	**380,475 tpmC**	**654 KRW/tpmC**		**1-Sep-2020**
Server Processors/Cores/Threads	Database Manager	Operating System	Other Software	Number of Users
2/48/96	**Goldilocks v3.1 Standard Edition**	**RHEL 7.7**	**JBoss Web Server**	**300,000**

100/25Gb Ethernet Switch

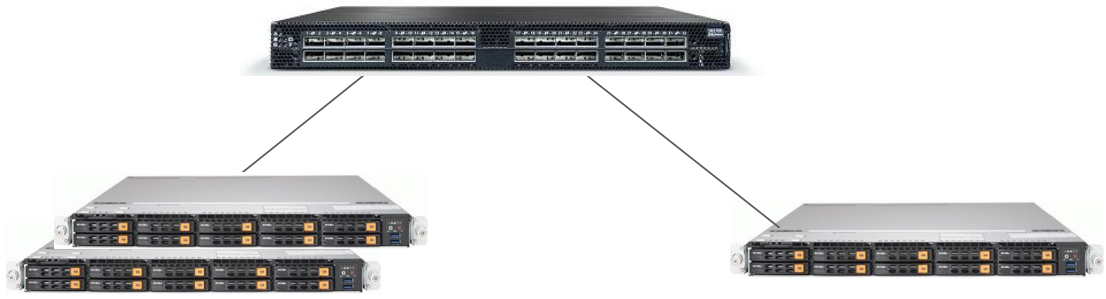

Web Application Server

2 x Supermicro A+ Server 1123US-TN10RT
- 2 x AMD EPYC 7262 8-Core Processor
- 8 x 16GB (128GB) Memory
- 1 x 1TB NVMe SSD
- 1 x 25Gb 2-Port Ethernet Adaptor

Database Server

1 x Supermicro A+ Server 1124US-TNRP
- 2 x AMD EPYC 7F72 24-Core Processor
- 16 x 256GB (4TB) Memory
- 2 x 7.68TB NVMe SSD
- 6 x 3.8TB SATA SSD
- 1 x 25Gb 2-Port Ethernet Adaptor

System Components	DB Server		WAS Server	
	Quantity	Description	Quantity	Description
Processors/Cores/Threads	2/48/96	AMD EPYC 7F72 3.2GHz	1/8/16	AMD EPYC 7262 3.2GHz
Memory	16	256GB	8	16GB
Storage Controller	1	SAS 12G Raid		
Storage Device	6 2	3.8TB SATA SSD 7.68TB NVMe SSD	1	1TB NVMe SSD
Total Storage Capacity		38.16TB		

	Goldilocks v3.1 Standard Edition on Supermicro A+ Server 1124US-TNRP	TPC-C Version 5.11.0 TPC Pricing 2.6.0 Report Date 18-Aug-2020 Availability Date 1-Sep-2020

Description	Part Number	Source	Unit Price	Qty	Price	3-Yr. Maint. Price
Server Hardware						
1 x DB Server						
Supermicro A+ Server 1124US-TNRP	AS-1124US-TNRP	1	108,500,000	1	**108,500,000**	
AMD EPYC 7F72 24-Core Processor	PSE-ROM7F72-0141	1	(included)	2		
Samsung M386ABG40M5B-CYF 256GB LRDIMM Memory	MEM-DR425L-SL01-LR29	1	(included)	16		
Micron 9300 MTFDHAL7T6TDP 7.68TB NVMe SSD	HDS-MUN-MTFDHALT6TDP1AT	1	(included)	2		
Samsung SAMSUNG MZ7LH3T8HMLT-00005 3.84TB SATA SSD	HDS-S2T1-MZ7LH3T8HMLT05	1	(included)	6		
Broadcom/LSI MegaRAID SAS-3 3108	AOC-S3108L-H8IR	1	(included)	1		
2-port 25GbE SFP28 Mellanox CX-4 Lx EN	AOC-MCX4121A-ACAT	1	(included)	1		
ASSEMBLY FEE	MC0037	1	(included)	1		
Maintenance - 7x24x4 Care Pack (3-yrs)	OS4HR3	1	(included)	1		
2 x WAS Servers (per server)						
Supermicro A+ Server 1123US-TN10RT	AS-1123US-TN10RT	1	9,850,000	2	**19,700,000**	
AMD EPYC 7262 8-Core	PSE-ROM7262-0041	1	(included)	2		
Supermicro 16GB 288-Pin DDR4 2666MHz	MEM-DR416L-HL03-ER26	1	(included)	4		
Intel DC P4500 1TB NVMe PCIe 3.0 SSD	HDS-IUN0-SSDPE2KX010T7	1	(included)	1		
2-port 25GbE SFP28 Mellanox CX-4 Lx EN and 2-port 10GbE RJ45 Intel X550	AOC-MH25G-M2S2TM-O	1	(included)	1		
ASSEMBLY FEE	MC0037	1	(included)	1		
Maintenance - 7x24x4 Care Pack (3-yrs)	OS4HR3	1	(included)	1		
Server Hardware Sub Total					**128,200,000**	
Network						
Mellanox MSN2700-CS2F Spectrum 100GbE 1U Open Ethernet Switch	MSN2700-CS2F	2	29,000,000	1	29,000,000	
Mellanox ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 1m, Colored, 30AWG, CA-N	MCP7F00-A001R30N	2	156,000	1	156,000	
Technical Support and Warranty - 24x7, 4 Hours On-Site Support 1yr	Maintenance	2	2,900,000	3		8,700,000
Network Sub Total					**29,156,000**	**8,700,000**
Client/Server Software						
Red Hat Enterprise Linux Server Standard 3yrs	RH00004F3	3	4,098,000	3	12,294,000	
RHEL Server Standard Maintenance - 3yrs 24x7x4hrs	RP-CPS(OS)	3	6,000,000	3		18,000,000
Red Hat JBoss Web Server 16-Core Standard 3Year	MW0232248F3	3	7,799,000	2	15,598,000	
JBoss Web Server per 16Core 3Year Maintenance	RP-CPS(WAS)	3	12,000,000	2		24,000,000
Goldilocks v3.1 Standard Edition		4	96,000,000	1	96,000,000	
Goldilocks v3.1 Standard Edition Technical Supports		4	10,000,000	3		30,000,000
Software Sub Total					**123,892,000**	**72,000,000**

Infrastructure						
HP EliteDisplay E243 23.8-inch Monitor (w/ spares)	1FH48AA	5	249,900	3	749,700	
HP Slim USB Keyboard and Mouse (w/ spares)	T6L04AA	5	50,000	3	164,700	
Infrastructure Sub Total					**914,400**	
Discounts*						
Red Hat Enterprise Linux Server Standard					(5,094,000)	(12,000,000)
Red Hat Jboss Web Server 16 Core Standard					(5,598,000)	(12,000,000)
Goldilocks v3.1 Standard Edition					(64,000,000)	(15,600,000)
Discounts Sub Total					**(74,692,000)**	**(39,600,000)**
Total					**207,470,400**	**41,100,000**

Pricing Notes

1)Super Solution Co.,LTD
2) ICRAFT Co., LTD
3) Rockplace Inc.
4) Sunjesoft Inc.
5) Hewlett Packard Inc.

All of the prices are based on South Korea's currency, KRW (₩, Korean Won) and excluded VAT.

* All discounts are based on Korea list prices and for similar quantities and configurations. Discounts for similarly sized configurations will be similar to those quoted here, but may vary based on the components in the configuration.

Three year cost of ownership KRW(₩): 248,570,400

TPC-C throughput: 380,475 tpmC

Price/Performance: 654 ₩ / tpmC

Benchmark implementation and results independantly audited by Francois Raab of InfoSizing (www.sizing.com)

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 pricing specifications. If you find that the stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.

Goldilocks v3.1 Standard Edition on Supermicro A+ Server 1124US-TNRP

TPC-C Version 5.11.0
TPC Pricing 2.6.0
Report Date
18-Aug-2020
Availability Date
1-Sep-2020

MQTh, computed Maximum Qualified Throughput	**380,475 tpmC**

Response Times (seconds)	**Min**	**Average**	**90th**	**Max**
New-Order	0.102	0.116	0.112	5.326
Payment	0.102	0.115	0.112	5.325
Order-Status	0.101	0.113	0.111	5.274
Delivery (interactive portion)	0.101	0.110	0.110	5.128
Delivery (deferred portion)	0.002	0.249	0.543	8.502
Stock-Level	0.102	0.114	0.112	5.274
Menu	0.101	0.111	0.110	5.129

Emulated Display Delay: 0.1 sec.

Transaction Mix	**Percent**	**Number**
New-Order	44.980%	159,799,583
Payment	43.011%	152,806,136
Order-Status	4.003%	14,219,897
Delivery	4.003%	14,222,472
Stock-Level	4.003%	14,222,400

Keying Times (seconds)	**Min**	**Average**	**Max**
New-Order	18.001	18.001	18.020
Payment	3.001	3.001	3.023
Order-Status	2.001	2.001	2.018
Delivery	2.001	2.001	2.014
Stock-Level	2.001	2.001	2.020

Think Times (seconds)	**Min**	**Average**	**Max**
New-Order	0.001	12.043	120.501
Payment	0.001	12.044	120.501
Order-Status	0.001	10.048	100.501
Delivery	0.001	5.031	50.301
Stock-Level	0.001	5.030	50.301

Test Duration	
Ramp-up time	65 min
Measurement Interval (MI)	420 min
Checkpoints in MI	15
Checkpoint Interval (Average / Max)	28.12 min / 28.15 min
Number of Transactions in MI (all types)	355,270,488

General Items

0.1 Application Code and Definition Statements

The application program (as defined in clause 2.1.7) must be disclosed. This includes, but is not limited to, the code implementing the five transactions and the terminal input output functions.

Appendix A contains the application source code for the transactions.

0.2 Benchmark Sponsor

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

This benchmark was sponsored by TTA, Telecommunications Technology Association. The implementation was developed and engineered with SUNJESOFT Inc. and Super Micro Computer, Inc.

0.3 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 options*
- *Recover/commit options*
- *Consistency locking options*
- *Operating system and application configuration parameters*

This requirement can be satisfied by providing a full list of all parameters.

Appendix B contains the tunable parameters for the database, the operating system, and the transaction monitor.

0.4 Configuration Diagrams

Diagrams of both measured and priced configurations must be provided, accompanied by a description of the differences.

The configuration diagram for both the tested and priced system is depicted in Figure 0.1. There was no difference between the priced and tested configurations.

Figure 0.1: Benchmarked and Priced Configuration

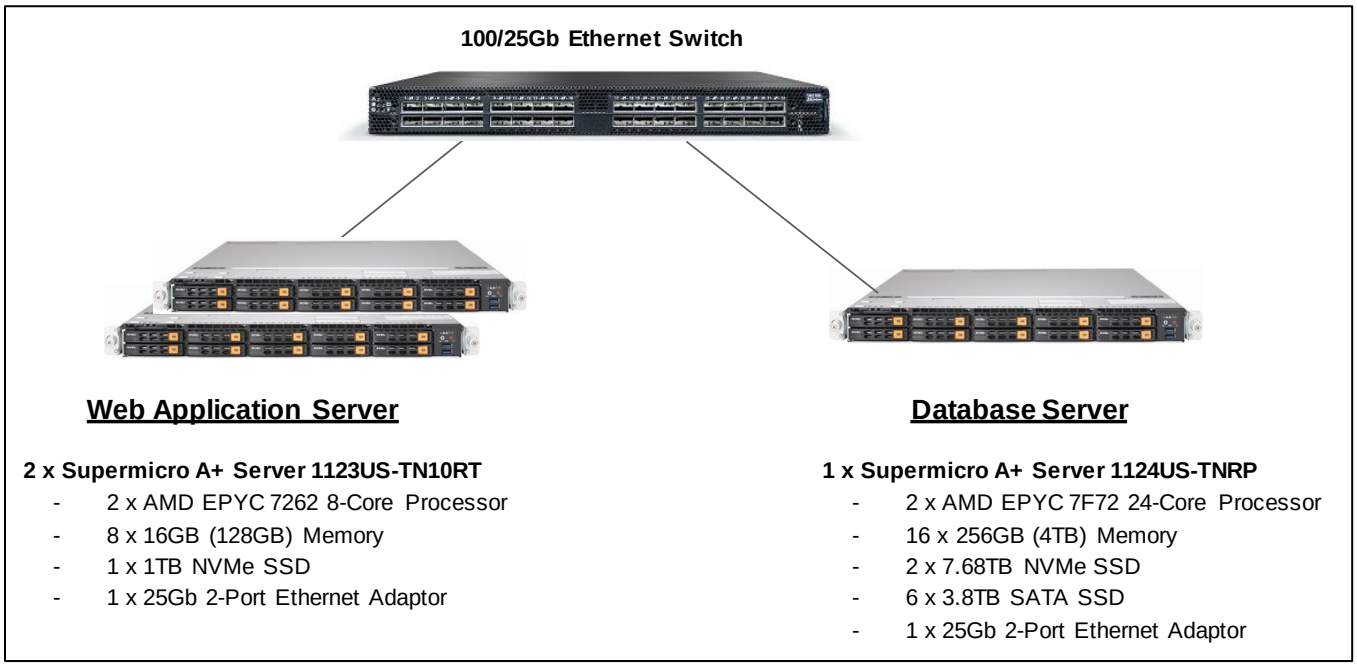

Clause 1: Logical Database Design

1.1 Table Definitions

Listing must be provided for all table definition statements and all other statements used to set up the database.

Appendix A contains the code used to define and load the database tables.

1.2 Physical Organization of Database

The physical organization of tables and indices within the database must be disclosed.

The physical organization of the database is shown in Table 1.2.

Table 1.2: Physical Organization of the Database

Controller	Array	RAID Array	Drives	Content
MegaRAID SAS-3 3108	Internal	RAID 1	2 x 3.8TB SATA SSD	OS
MegaRAID SAS-3 3108	Internal	RAID 10	4 x 3.8TB SATA SSD	Redo Logs
Software RAID	Internal	RAID 1	2 x 7.68TB NVMe SSD	Data Files

1.3 Insert and Delete Operations

It must be ascertained that insert and/or delete operations to any of the tables can occur concurrently with the TPC-C transaction mix. Furthermore, any restrictions in the SUT database implementation that precludes inserts beyond the limits defined in Clause 1.4.11 must be disclosed. This includes the maximum number of rows that can be inserted and the minimum key value for these new rows.

All insert and delete functions were verified to be fully operational during the entire benchmark.

1.4 Horizontal or Vertical Partitioning

While there are a few restrictions placed upon horizontal or vertical partitioning of tables and rows in the TPC-C benchmark, any such partitioning must be disclosed.

No horizontal or vertical partitioning was used in this benchmark.

1.5 Replication or Duplication

Replication of tables, if used, must be disclosed. Additional and/or duplicated attributes in any table must be disclosed along with a statement on the impact on performance.

No replications, duplications or additional attributes were used in this benchmark.

Clause 2: Transaction and Terminal Profiles

2.1 Random Number Generation

The method of verification for the random number generation must be described.

Random numbers were generated using ‘SysVr4 rand_r()’ call. The seed value for ‘rand_r()’ was collected and reviewed by the auditor.

2.2 Input/Output Screens

The actual layout of the terminal input/output screens must be disclosed.

All screen layouts were verified by the auditor to validate that they followed the requirements of the specifications.

2.3 Priced Terminal Feature

The method used to verify that the emulated terminals provide all the features described in Clause 2.2.2.4 must be explained. Although not specifically priced, the type and model of the terminals used for the demonstration in 8.1.3.3 must be disclosed and commercially available (including supporting software and maintenance).

The terminal attributes were manually verified by the auditor by verifying that each required feature was implemented.

2.4 Presentation Managers

Any usage of presentation managers or intelligent terminals must be explained.

Application code running on the client systems implemented the TPC-C user interface. No presentation manager software or intelligent terminal features were used. The source code for the user interface is listed in Appendix A.

2.5 Transaction Statistics

Table 2.1 lists the transaction statistics defined in Clauses 8.1.3.5 to 8.1.3.11 and observed during the Measurement Interval.

Table 2.1: Transaction Statistics

Statistic		Value
New Order	Home warehouse order lines Remote warehouse order lines Rolled back transactions Average items per order	99.001% 0.999% 1.001% 10.000
Payment	Home warehouse Remote warehouse Accessed by last name	84.999% 15.001% 59.996%
Order Status	Accessed by last name	59.971%
Delivery	Skipped transactions	0
Transaction Mix	New Order Payment Order status Delivery Stock level	44.980% 43.011% 4.003% 4.003% 4.003%

2.6 Queuing Mechanism

The queuing mechanism used to defer the execution of the Delivery transaction must be disclosed.

The queuing mechanism was implemented using ‘BlockingQueue’ provided by Java.

Clause 3: Transaction and System Properties

The results of the ACID tests must be disclosed along with a description of how the ACID requirements were met. This includes disclosing which case was followed for the execution of Isolation Test 7.

All ACID property tests were conducted according to the specification.

3.1 Atomicity

The system under test must guarantee that the database transactions are atomic; the system will either perform all individual operations on the data or will assure that no partially completed operations leave any effects on the data.

3.1.1 Atomicity of Completed Transactions

Perform the Payment transaction for a randomly selected warehouse, district, and customer (by customer number) and verify that the records in the CUSTOMER, DISTRICT, and WAREHOUSE tables have been changed appropriately.

A row was randomly selected from the CUSTOMER, DISTRICT, and WAREHOUSE tables, and the balances noted. A payment transaction was started with the same Customer, District, and Warehouse identifiers and a known amount. The payment transaction was committed and the rows were verified to contain correctly updated balances.

3.1.2 Atomicity of Aborted Transactions

Perform the Payment transaction for a randomly selected warehouse, district, and customer (by customer number) and substitute a ROLLBACK of the transaction for the COMMIT of the transaction. Verify that the records in the CUSTOMER, DISTRICT, and WAREHOUSE tables have NOT been changed.

A row was randomly selected from the CUSTOMER, DISTRICT, and WAREHOUSE tables, and the balances noted. A payment transaction was started with the same Customer, District, and Warehouse identifiers and a known amount. The payment transaction was rolled back and the rows were verified to contain the original balances.

3.2 Consistency

Consistency is the property of the application that requires any execution of a data base transaction to take the database from one consistent state to another, assuming that the data base is initially in a consistent state.

Verify that the data base is initially consistent by verifying that it meets the consistency conditions defined in Clauses 3.3.2.1 to 3.3.2.4. Describe the steps used to do this in sufficient detail so that the steps are independently repeatable.

The specification defines 12 consistency conditions, of which Consistency conditions 1 through 4 were demonstrated as follows:

1. The sum of balances (d_ytd) for all Districts within a specific Warehouse is equal to the balance (w_ytd) of that Warehouse.
2. For each District within a Warehouse, the next available Order ID (d_next_o_id) minus one is equal to the most recent Order ID [max(o_id)] for the ORDER table associated with the preceding District and Warehouse. Additionally, that same relationship exists for the most recent Order ID [max(o_id)] for the NEW-ORDER table associated with the same District and Warehouse. Those relationships can be illustrated as:
   d_next_o_id – 1 = max(o_id) = max(no_o_id)
   where (d_w_id = o_w_id = no_w_id) and (d_id = o_d_id = no_d_id)

3. For each District within a Warehouse, the value of the most recent Order ID [max(no_o_id)] minus the first Order ID [min(no_o_id)] plus one, for the NEW-ORDER table associated with the District and Warehouse, equals the number of rows in that NEW-ORDER table.
   That relationship can be illustrated as:
   max(no_o_id) – min(no_o_id) + 1 = rows in NEW-ORDER
   where (o_w_id = no_w_id) and (o_d_id = no_d_id)
4. For each District within a Warehouse, the sum of Order-Line counts [sum(o_ol_cnt)] for the Orders associated with the District equals the number of rows in the ORDER-LINE table associated with the same District.
   That relationship can be illustrated as:
   sum(o_ol_cnt) = rows in the ORDER-LINE table for the Warehouse and District

To test consistency, the following steps were executed:

1. The consistency conditions 1 through 4 were tested by running queries against the database. All queries showed that the database was in a consistent state.
2. An RTE run was executed at full load for a duration sufficient to include at least one completed checkpoint.
3. The consistency conditions 1 through 4 were tested again. All queries showed that the database was still in a consistent state.

3.3 Isolation

Sufficient conditions must be enabled at either the system or application level to ensure the required isolation defined above (clause 3.4.1) is obtained.

The benchmark specification defines nine tests to demonstrate the property of transaction isolation. The tests, described in Clauses 3.4.2.1 – 3.4.2.9, were all successfully executed using a series of scripts. Each included timestamps to demonstrate the concurrency of operations. The results of the queries were logged. The captured logs were verified to demonstrate the required isolation had been met.

Isolation Test 1

This test demonstrates isolation for read-write conflicts of Order-Status and New-Order transactions when the New-Order transaction is committed.

The test proceeds as follows:

1. An Order-Status transaction T0 was executed and committed for a randomly selected Customer, and the Order returned was noted.
2. A New-Order transaction T1 was started for the same Customer used in T0. T1 was stopped prior to COMMIT.
3. An Order-Status transaction T2 was started for the same Customer used in T1. T2 completed and was committed without being blocked by T1. T2 returned the same Order that T0 had returned.
4. T1 was allowed to complete and was committed.
5. An Order-Status transaction T3 was started for the same Customer used in T1. T3 returned the Order inserted by T1.

Isolation Test 2

This test demonstrates isolation for read-write conflicts of Order-Status and New-Order transactions when the New-Order transaction is rolled back.

The test proceeds as follows:

1. An Order-Status transaction T0 was executed and committed for a randomly selected Customer and the Order returned was noted.
2. A New-Order transaction T1 with an invalid item number was started for the same Customer used in T0. T1 was stopped immediately prior to ROLLBACK.
3. An Order-Status transaction T2 was started for the same Customer used in T1. T2 completed and was committed without being blocked by T1. T2 returned the same Order that T0 had returned.
4. T1 was allowed to ROLLBACK.
5. An Order-Status transaction T3 was started for the same Customer used in T1. T3 returned the same Order that T0 had returned.

Isolation Test 3

This test demonstrates isolation for write-write conflicts of two New-Order transactions when both transactions are committed.

The test proceeds as follows:

1. The D_NEXT_O_ID of a randomly selected district was retrieved.
2. A New-Order transaction T1 was started for a randomly selected customer within the District used in step 1. T1 was stopped immediately prior to COMMIT.
3. Another New-Order transaction T2 was started for the same customer used in T1. T2 waited.
4. T1 was allowed to complete. T2 completed and was committed.
5. The order number returned by T1 was the same as the D_NEXT_O_ID retrieved in step 1. The order number returned by T2 was one greater than the order number returned by T1.
6. The D_NEXT_O_ID of the same District was retrieved again. It had been incremented by two (i.e. it was one greater than the order number returned by T2).

Isolation Test 4

This test demonstrates isolation for write-write conflicts of two New-Order transactions when one transaction is rolled back.

The test proceeds as follows:

1. The D_NEXT_O_ID of a randomly selected District was retrieved.
2. A New-Order transaction T1, with an invalid item number, was started for a randomly selected customer within the district used in step 1. T1 was stopped immediately prior to ROLLBACK.
3. Another New-Order transaction T2 was started for the same customer used in T1. T2 waited.
4. T1 was allowed to roll back, and T2 completed and was committed.
5. The order number returned by T2 was the same as the D_NEXT_O_ID retrieved in step 1.

6. The D_NEXT_O_ID of the same District was retrieved again. It had been incremented by one (i.e. one greater than the order number returned by T2).

Isolation Test 5

This test demonstrates isolation for write-write conflicts of Payment and Delivery transactions when Delivery transaction is committed.

The test proceeds as follows:

1. A query was executed to find out the Customer who is to be updated by the next Delivery transaction for a randomly selected Warehouse and District.
2. The C_BALANCE of the Customer found in step 1 was retrieved.
3. A Delivery transaction T1 was started for the same Warehouse used in step 1. T1 was stopped immediately prior to COMMIT.
4. A Payment transaction T2 was started for the same Customer found in step 1. T2 waited.
5. T1 was allowed to complete. T2 completed and was committed.
6. The C_BALANCE of the Customer found in step 1 was retrieved again. The C_BALANCE reflected the results of both T1 and T2.

Isolation Test 6

This test demonstrates isolation for write-write conflicts of Payment and Delivery transactions when the Delivery transaction is rolled back.

The test proceeds as follows:

1. A query was executed to find out the Customer who is to be updated by the next delivery transaction for a randomly selected Warehouse and District.
2. The C_BALANCE of the Customer found in step 1 was retrieved.
3. A Delivery transaction T1 was started for the same Warehouse used in step 1. T1 was stopped immediately prior to COMMIT.
4. A Payment transaction T2 was started for the same customer found in step 1. T2 waited.
5. T1 was forced to execute a ROLLBACK. T2 completed and was committed. The C_BALANCE of the Customer found in step 1 was retrieved again. The C_BALANCE reflected the results of only T2.

Isolation Test 7

This test demonstrates repeatable reads for the New-Order transaction while an interactive transaction updates the prices of some items.

The test proceeds as follows:

1. The I_PRICE of two randomly selected items X and Y were retrieved.
2. A New-Order transaction T1 with a group of Items including Items X and Y was started. T1 was stopped immediately after retrieving the prices of all items. The prices of Items X and Y retrieved matched those retrieved in step 1.

3. A transaction T2 was started to increase the price of Items X and Y by 10%.
4. T2 did not stall and was committed.
5. T1 was resumed, and the prices of all Items were retrieved again within T1. The prices of Items X and Y matched those retrieved in step 1.
6. T1 was committed.
7. The prices of Items X and Y were retrieved again. The values matched the values set by T2.

The Execution followed Case D, where T3 does not stall and no transaction is rolled back. Query T4 verifies the price change made by T3.

Isolation Test 8

This test demonstrates isolation for phantom protection between New-Order and Order-Status transactions.

The test proceeds as follows:

1. An Order-Status transaction T1 was started for a randomly selected Customer.
2. T1 was stopped immediately after reading the ORDER table for the selected Customer to find the most recent Order for that Customer.
3. A New-Order transaction T2 was started for the same Customer. T2 completed and was committed without being blocked by T1.
4. T1 was resumed and the ORDER table was read again to determine the most recent Order for the same Customer. The Order found was the same as the one found in step 2.
5. T1 completed and was committed.

Isolation Test 9

This test demonstrates isolation for phantom protection between New-Order and Delivery transactions.

The test proceeds as follows:

1. The NO_D_ID of all NEW_ORDER rows for a randomly selected Warehouse and District was changed to 11. The changes were committed.
2. A Delivery transaction T1 was started for the selected Warehouse.
3. T1 was stopped immediately after reading the NEW_ORDER table for the selected Warehouse and District. No qualifying row was found.
4. A New-Order transaction T2 was started for the same Warehouse and District. T2 completed and was committed without being blocked by T1.
5. T1 was resumed and the NEW_ORDER table was read again. No qualifying row was found.
6. T1 completed and was committed.
7. The NO_D_ID of all NEW_ORDER rows for the selected Warehouse and District was restored to the original value. The changes were committed.

3.4 Durability

The tested system must guarantee durability: the ability to preserve the effects of committed transactions and ensure data base consistency after recovery from any one of the failures listed in Clause 3.5.3

- *Permanent irrecoverable failure of any single durable medium containing TPC-C database tables or recovery log data (this test includes failure of all or part of memory)*
- *Instantaneous interruption (system crash/system hang) in processing that requires system reboot to recover*
- *Failure of all or part of memory (loss of contents)*

3.4.1 Durable Media Failure

3.4.1.1 Loss of Log Media and Data Media

This test was conducted on a fully scaled database. To demonstrate recovery from a permanent failure of durable medium containing TPC-C Log Media and Data Media, the following steps were executed:

1. The total number of Orders is determined by the sum of D_NEXT_O_ID of all rows in the DISTRICT table; giving count-1.
2. The consistency is verified.
3. The RTE is started with full user load.
4. The test is allowed to run for a minimum of 5 minutes after ramp-up.
5. A first checkpoint is initiated and completed.
6. The test is allowed to run for a minimum of 2 more minutes.
7. A second checkpoint is initiated.
8. Before the second checkpoint completes, one data disk is disabled by removing it logically. Since the data disks are configured with redundancy, the transactions continued to run without interruption.
9. The test is allowed to run until the completion of the second checkpoint and for at least 5 minutes
10. A third checkpoint is initiated.
11. Before the third checkpoint completes, one log device is disabled by removing it logically. Since the log devices are configured with redundancy, the transactions continued to run without interruption.
12. The test is allowed to run until the third checkpoint has completed, but no less than 5 more minutes.
13. The RTE run is completed.
14. The consistency is verified.
15. Step 1 is repeated, giving count-2.
16. The RTE result file is used to determine the number of New-Order transactions successfully completed during the full run.
17. The difference between the count-1 and count-2 is compared with the number of New-Order transactions successfully completed during the full run. The difference indicated that no committed transactions had been lost.
18. Data from the success file is used to query the database to demonstrate that the last 500 successful New-Orders have corresponding rows in the ORDER table.

3.4.2 Instantaneous Interruption, Loss of Memory

As the loss of power erases the contents of memory, the instantaneous interruption and the loss of memory tests were combined into a single test. This test was executed on a fully scaled database. The following steps were executed:

1. The total number of Orders is determined by the sum of D_NEXT_O_ID of all rows in the DISTRICT table; giving count-1.
2. The consistency is verified.
3. The RTE is started with full user load.
4. The test is allowed to run for a minimum of 5 minutes at full load (after ramp-up).
5. A first checkpoint is initiated and completed.
6. The test is allowed to run for a minimum of 2 more minutes.
7. A second checkpoint is initiated.
8. Before the second checkpoint completes, the primary power is shut off logically using BMC
9. The RTE is shutdown.
10. Power is restored to the database server and the system performs an automatic recovery.
11. GOLDILOCKS is restarted and performs an automatic recovery.
12. Step 1 is repeated, giving count-2.
13. The consistency is verified.
14. The RTE result file is used to determine the number of New-Order transactions successfully completed during the full run.
15. The difference between the count-1 and count-2 is compared with the number of New-Order transactions successfully completed during the full run. The difference indicated that all committed transactions had been successfully recovered.
16. Data from the success file is used to query the database to demonstrate that the last 500 successful New-Orders have corresponding rows in the ORDER table.

Clause 4: Scaling and Database Population

4.1 Cardinality of Tables

The cardinality (e.g. number of rows) of each table, as it existed at the start of the benchmark run, must be disclosed. If the database was over-scaled and inactive rows of the WAREHOUSE table were deleted, the cardinality of the WAREHOUSE table as initially configured and the number of rows

Table 4.1 shows that number of rows for each table as they were initially populated.

Table 4.1: Number of Rows for Server

Table	Cardinality
Warehouse	30,000
District	300,000
Customer	900,000,000
History	900,000,000
Order	900,000,000
New Order	270,000,000
Order Line	8,998,142,067
Stock	3,000,000,000
Item	100,000
Unused Warehouses	0

4.2 Database Implementation

A statement must be provided that describes: The data model implemented by DBMS used (e.g. relational, network, hierarchical). The database interfaces (e.g. embedded, call level) and access language (e.g. SQL, DL/1, COBOL read/write used to implement the TPC-C transaction. If more than one interface/access language is used to implement TPC-C, each interface/access language must be described and a list of which interface/access language is used with which transaction type must be disclosed.

Goldilocks v3.1 is an in-memory DBMS, implementing the relational model.

The transactions are implemented in SQL via JDBC calls to the database engine.

All application code and procedures are listed in Appendix A.

4.3 Distribution of Database Files

The distribution of tables and logs across all media must be explicitly depicted for tested and priced systems.

The database files are stored on a set of two 7.68TB disks configured as RAID1. The database log files are stored on four 3.8TB disks configured as RAID10.

Table 4.3: Database file locations

Name	Location	Description
system_XXX.dbf	/data/db/db1	System tables and dictionary
tpcc_data_XXX.dbf	/data/db/db1 /data/db/db2 /data/db/db3 /data/db/db4 /data/db/db5	Database data files
redo_X_X.log	/wal	Database log files

The distribution of tables and logs across storage media is shown in Table 1.2.

4.4 60-Day Space

Details of the 60-day space computations along with proof that the database is configured to sustain 8 hours of growth for the dynamic tables (Order, Order-Line, and History) must be disclosed.

A test run of over 8 hours was executed to demonstrate that the configuration is capable of sustaining 8 hours of growth at the reported throughput. The computation of the 60-day storage requirements is shown in Table 4.4.

Table 4.4: 60-Day Space Calculations

Base Unit (KBytes)	1
tpmC	380,475.198

Table	Rows	Data	Index	Initial Population	5% Growth	8-Hour Growth	Required Runtime Space
WAREHOUSE	30,000	241,872	816	242,688	12,134	0	254,822
DISTRICT	300,000	38,256	12,648	50,904	2,545	0	53,449
CUSTOMER	900,000,000	579,642,632	75,115,080	654,757,712	32,737,886	0	687,495,598
NEW_ORDER	270,000,000	17,051,872	9,714,440	26,766,312	1,338,316	0	28,104,628
ITEM	100,000	10,808	2,736	13,544	677	0	14,221
STOCK	3,000,000,000	1,108,169,448	100,255,824	1,208,425,272	60,421,264	0	1,268,846,536
HISTORY	900,000,000	75,435,976	0	75,435,976	0	15,307,476	90,743,452
ORDERS	900,000,000	57,461,200	69,300,504	126,761,704	0	11,660,033	138,421,737
ORDER_LINE	8,998,142,067	857,262,552	357,526,848	1,214,789,400	0	173,955,808	1,388,745,208
Total		2,695,314,616	611,928,896	3,307,243,512	94,512,822	200,923,317	3,602,679,650

60-Day Requirements	
Dynamic-Space	990,159,728
Free-Space	6,118,168
Static-Space	2,317,083,784
Daily-Growth	200,923,317
Daily-Spread	0
60-Day Space	**14,372,482,778**

Memory Requirements	
Final Allocation	3,668,042,008
Non-Growing 5%	94,512,822
1-Day Memory	**3,762,554,830**

Storage Requirements	
Total Disk Space	18,573,915,548
Log space used	524,288,000
60-Day Space	14,372,482,778
Remaining Space	**3,677,144,770**

Clause 5: Performance Metrics

5.1 TPC Benchmark C Metrics

The TPC-C Metrics are reported in the front of this report as part of the executive summary.

5.2 Response Times

Ninetieth percentile, maximum and average response times must be reported for all transaction types as well as for the menu response time.

During the performance run transactions are submitted by the RTE in accordance with the required mix, Keying Times and Think Times of the benchmark Specification. Transactions are submitted by emulated users via HTTP. All timings are recorded by the RTE. The response time is measured from the submission of the transaction until the last byte of response is received by the RTE.

The details of the response times are reported in the front of this report as part of the Executive Summary.

5.3 Keying and Think Times

The minimum, the average, and the maximum keying and think times must be reported for each transaction type.

The details of the keying and think times are reported in the front of this report as part of the Executive Summary.

5.4 Distribution and Performance Curves

5.4.1 Response Time frequency distribution curves

Response Time frequency distribution curves must be reported for each transaction type.

Figure 5.4.1.1 shows the Response Time frequency distribution curves for the New-Order transaction.

Figure 5.4.1.2 shows the Response Time frequency distribution curves for the Payment transaction.

Figure 5.4.1.3 shows the Response Time frequency distribution curves for the Order-Status transaction.

Figure 5.4.1.4 shows the Response Time frequency distribution curves for the interactive portion of the Delivery transaction.

Figure 5.4.1.5 shows the Response Time frequency distribution curves for the Stock-Level transaction.

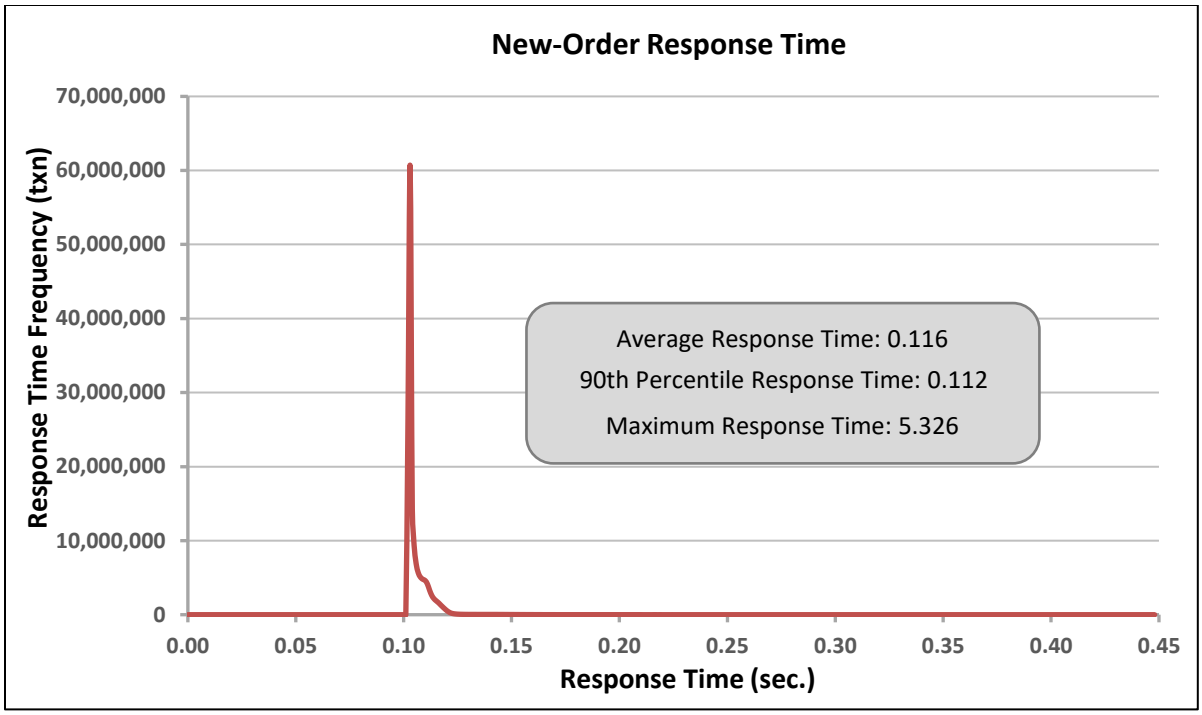

Figure 5.4.1.1: New-Order RT Frequency Distribution

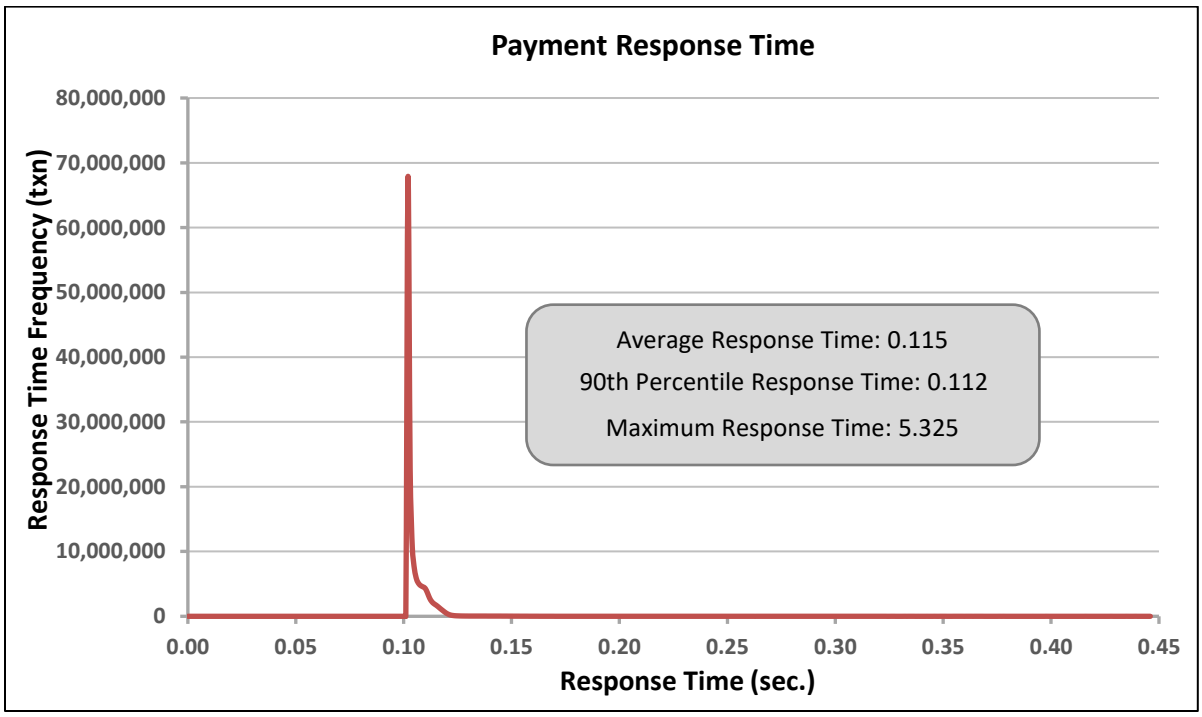

Figure 5.4.1.2: Payment RT Frequency Distribution

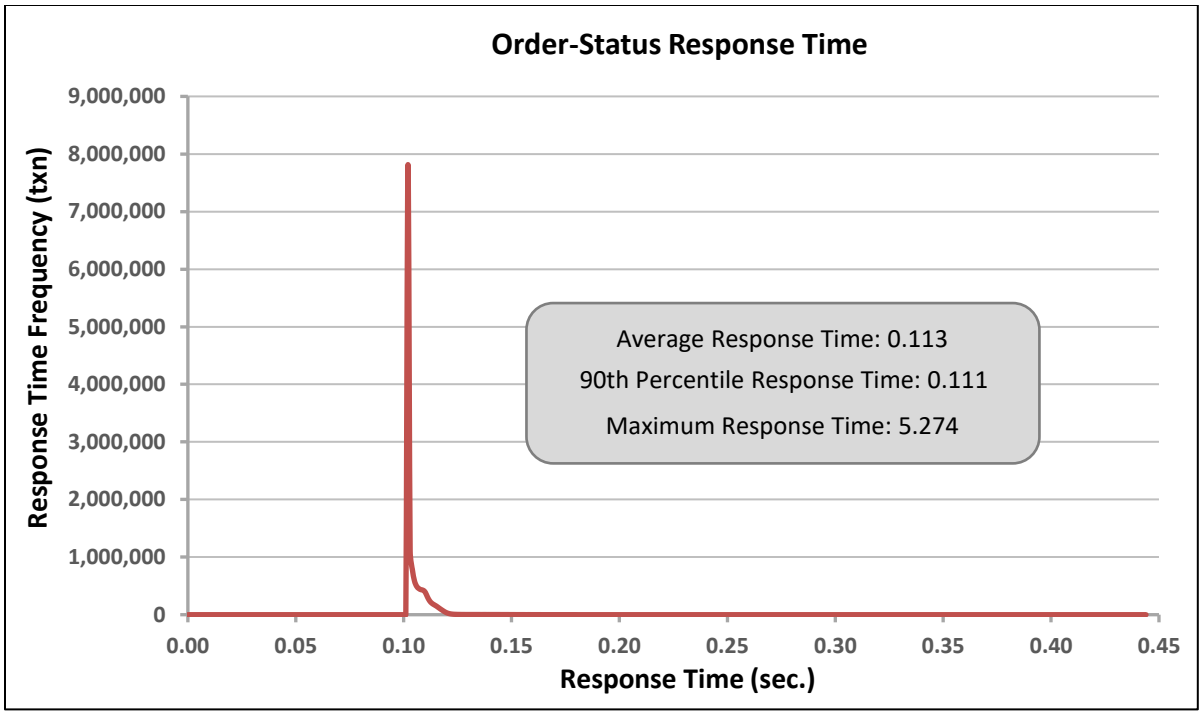

Figure 5.4.1.3: Order-Status RT Frequency Distribution

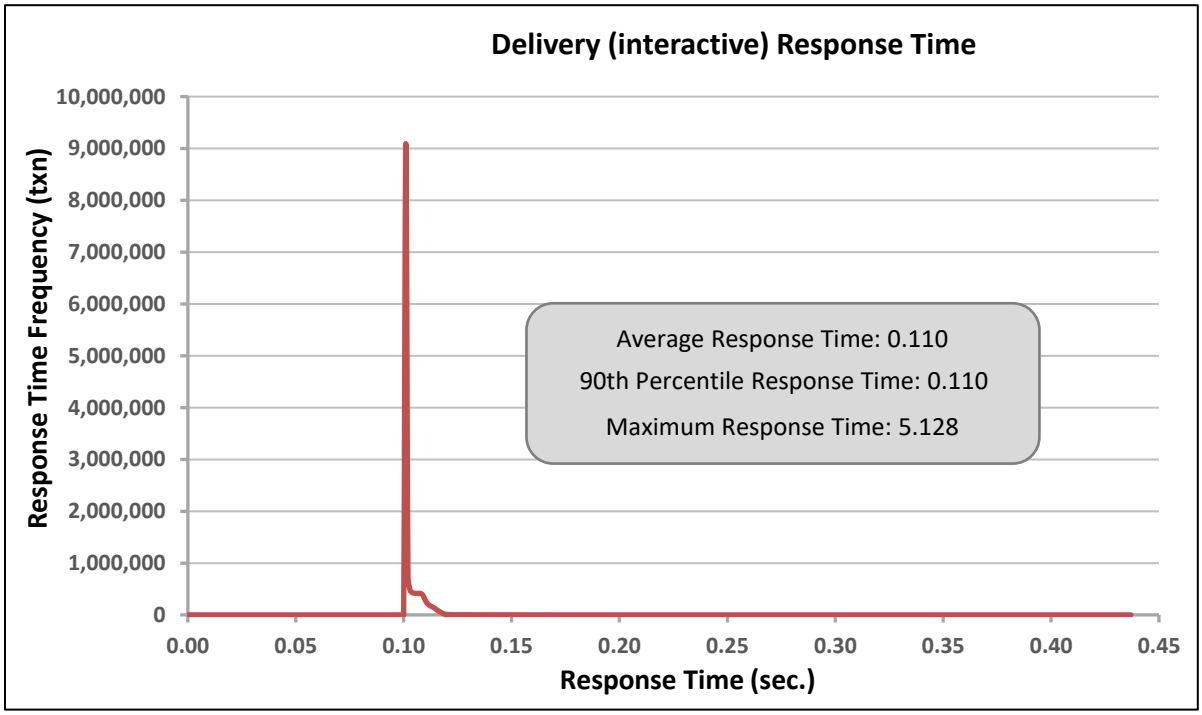

Figure 5.4.1.4: Delivery (Interactive) RT Frequency Distribution

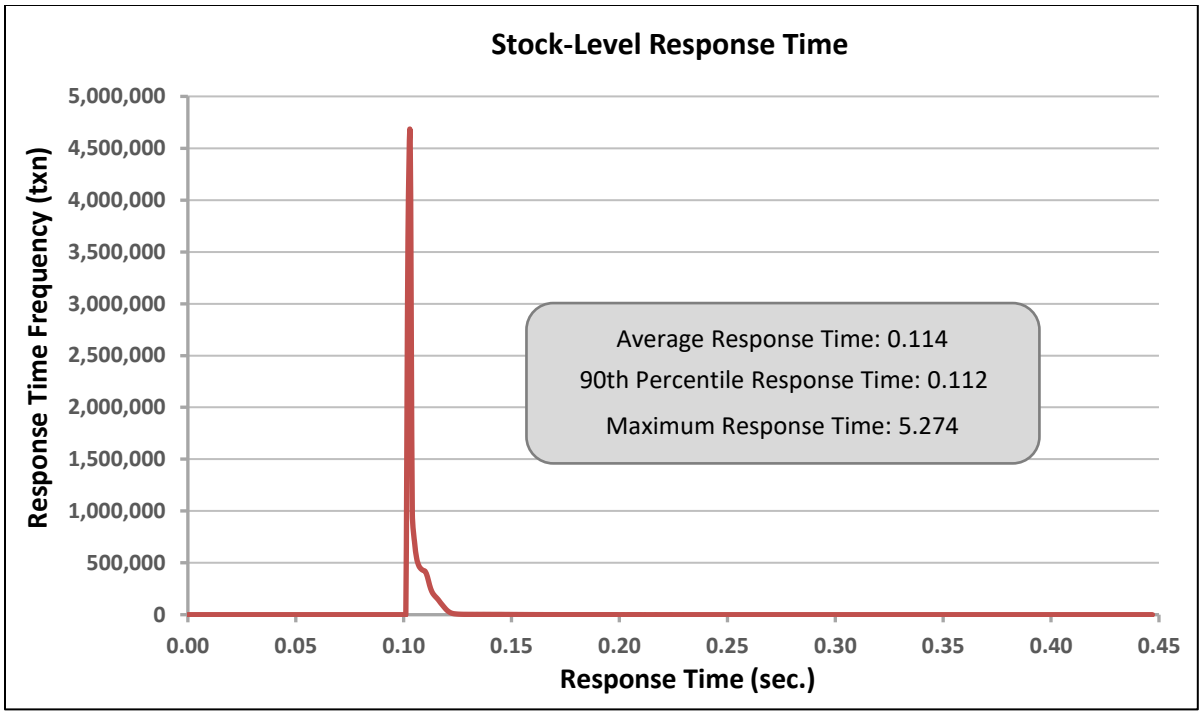

Figure 5.4.1.5: Stock-Level RT Frequency Distribution

5.4.2 Response Time versus throughput

The performance curve for response times versus throughput must be reported for the New-Order transaction.

Figure 5.4.2 shows the Response Time versus throughput curves for the New-Order transaction.

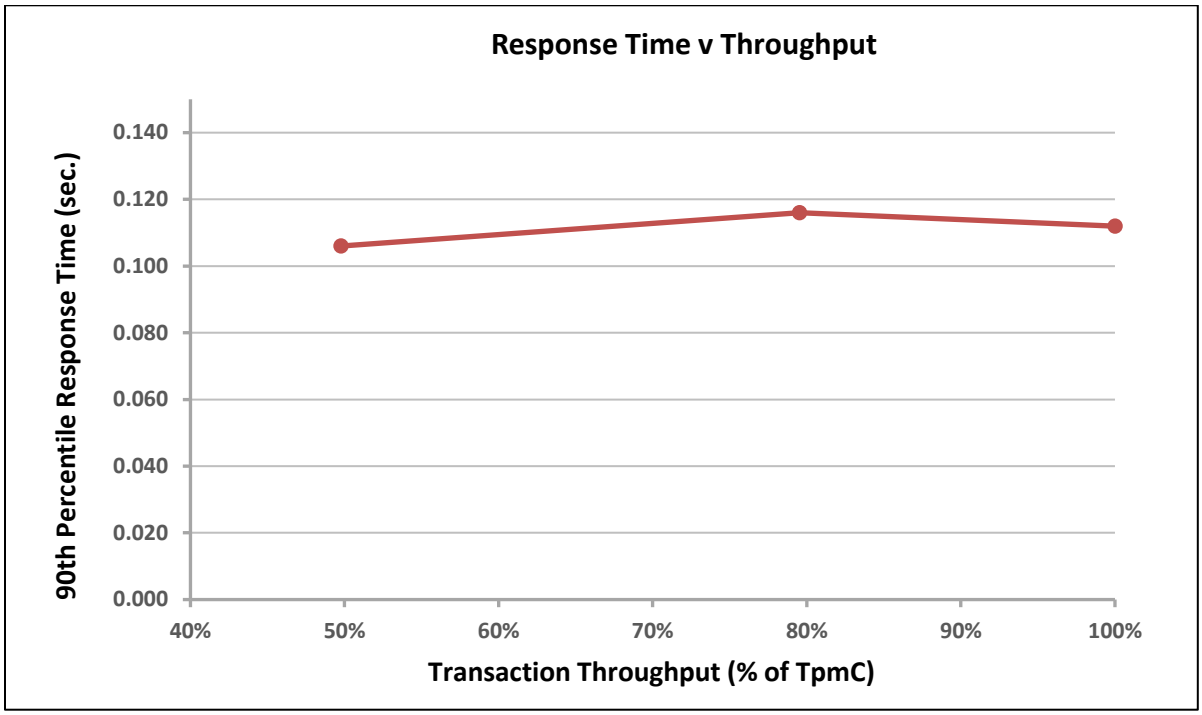

Figure 5.4.2: New-Order RT versus Throughput

5.4.3 Think Time frequency distribution

Think Time frequency distribution curves (see Clause 5.6.3) must be reported for the New-Order transaction.

Figure 5.4.3 shows the Think Time frequency distribution curves for the New-Order transaction.

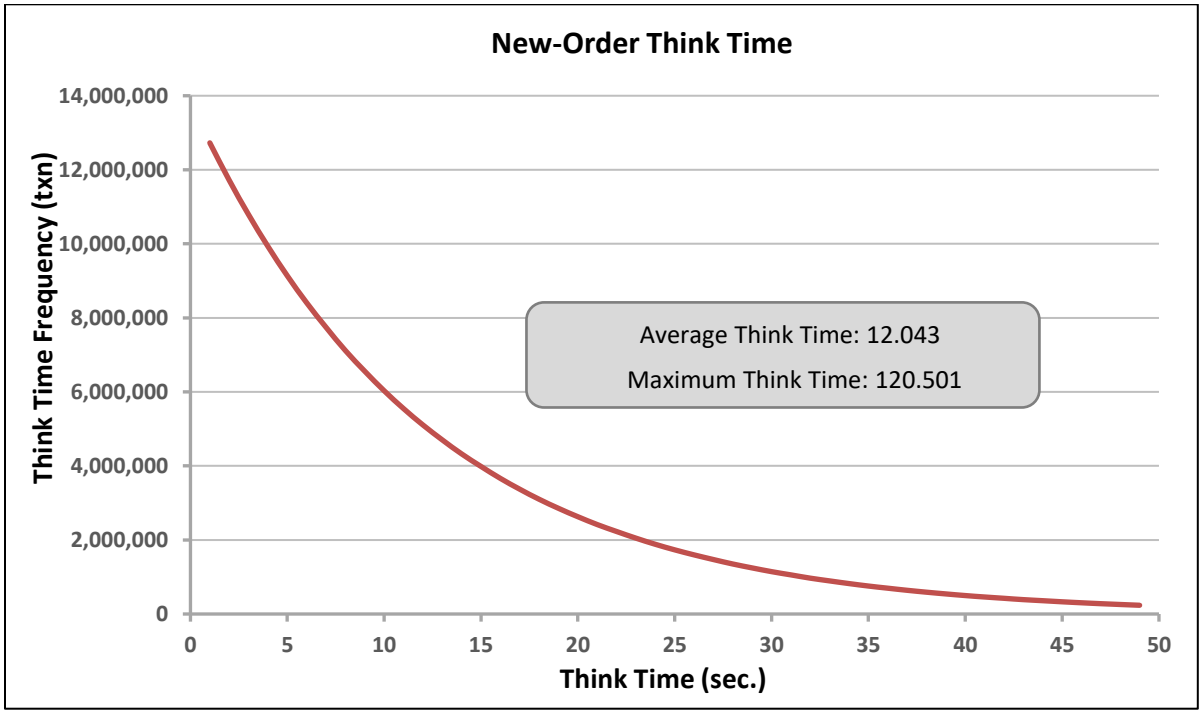

Figure 5.4.3: New-Order Think Time Frequency Distribution

5.4.4 Throughput versus elapsed time

A graph of throughput versus elapsed time must be reported for the New-Order transaction.

Figure 5.4.4 shows the throughput versus elapsed time for the New-Order transaction. The start and end of the Measurement Interval is included on the figure.

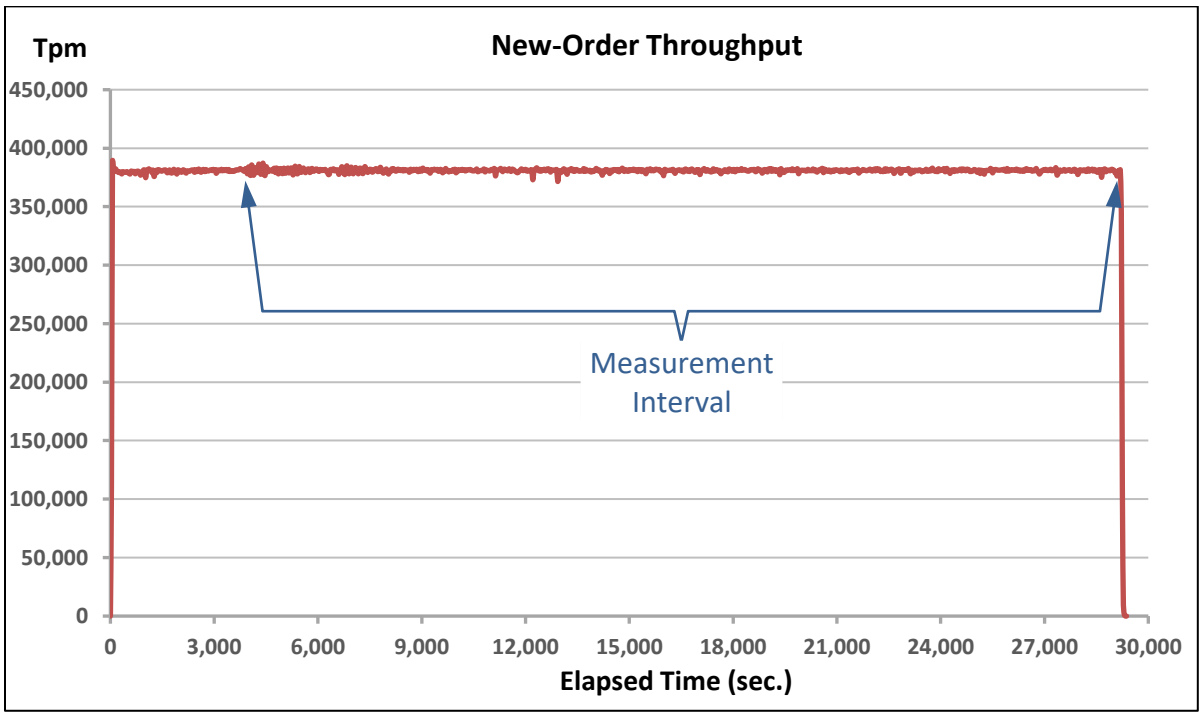

Figure 5.4.4: New-Order Throughput versus Elapsed Time

5.5 Steady State Determination

The method used to determine that the SUT had reached a steady state prior to commencing the measurement interval must be disclosed.

Steady state was determined using real time monitor utilities from the RTE. Steady state was further confirmed by a visual analysis of the throughput graph.

5.6 Work Performed During Steady State

A description of how the work normally performed during a sustained test (for example checkpointing, writing redo/undo log records, etc.) actually occurred during the measurement interval must be reported.

During the test, Goldilocks satisfied all of the ACID properties required by the benchmark specification. Committed transactions write a Redo record in the transaction log, to be used in case of system failure. The Redo records are used for roll-forward recovery during a re-start following a failure. This prevents the system from losing any committed transactions. Checkpoints periodically occurred about every 28 min. and are completed in about 14 min.

5.7 Measurement Period Duration

A statement of the duration of the measurement interval for the reported Maximum Qualified Throughput (tpmC) must be included.

The duration of the reported measured interval was 7 hours (7hr = 420min = 25,200sec).

5.8 Transaction Statistics

The percentage of the total mix for each transaction type must be disclosed. The percentage of New-Order transactions rolled back as a result of invalid item number must be disclosed. The average number of order-lines entered per New-Order transaction must be disclosed. The percentage of remote order lines per New-Order transaction must be disclosed. The percentage of remote Payment transactions must be disclosed. The percentage of customer selections by customer last name in the Payment and Order-Status transactions must be disclosed. The percentage of skipped Delivery transactions must be disclosed.

The details of the transaction statistics are reported in the front of this report as part of the Executive Summary.

5.9 Checkpoints

The number of checkpoints in the Measurement Interval, the time in seconds from the start of the Measurement Interval to the first checkpoint, and the Checkpoint Interval must be disclosed.

Two full checkpoints occurred before the Measurement Interval. 15 full checkpoints occurred during the Measurement Interval. The checkpoints' start and end times and durations during the Measurement Interval are listed in table 5.6.

Table 5.6: Checkpoints

Event	Event time	Execution time	Interval
Measurement Interval Begin	2020-04-14 06:42:22	-	-
Checkpoint3 Begin	2020-04-14 06:49:12		00:28:15
Checkpoint3 End	2020-04-14 07:02:57	00:13:45	
Checkpoint4 Begin	2020-04-14 07:17:22		00:28:10
Checkpoint4 End	2020-04-14 07:31:11	00:13:49	
Checkpoint5 Begin	2020-04-14 07:45:34		00:28:12
Checkpoint5 End	2020-04-14 07:59:22	00:13:48	
Checkpoint6 Begin	2020-04-14 08:13:44		00:28:10
Checkpoint6 End	2020-04-14 08:27:36	00:13:51	
Checkpoint7 Begin	2020-04-14 08:41:56		00:28:11
Checkpoint7 End	2020-04-14 08:55:50	00:13:54	
Checkpoint8 Begin	2020-04-14 09:10:07		00:28:11
Checkpoint8 End	2020-04-14 09:24:00	00:13:53	
Checkpoint9 Begin	2020-04-14 09:38:21		00:28:14
Checkpoint9 End	2020-04-14 09:52:15	00:13:55	
Checkpoint10 Begin	2020-04-14 10:06:32		00:28:12
Checkpoint10 End	2020-04-14 10:20:25	00:13:53	
Checkpoint11 Begin	2020-04-14 10:34:45		00:28:13
Checkpoint11 End	2020-04-14 10:48:40	00:13:55	
Checkpoint12 Begin	2020-04-14 11:02:56		00:28:11
Checkpoint12 End	2020-04-14 11:16:49	00:13:53	
Checkpoint13 Begin	2020-04-14 11:31:08		00:28:11
Checkpoint13 End	2020-04-14 11:45:04	00:13:56	
Checkpoint14 Begin	2020-04-14 11:59:20		00:28:12
Checkpoint14 End	2020-04-14 12:13:13	00:13:54	
Checkpoint15 Begin	2020-04-14 12:27:31		00:28:12
Checkpoint15 End	2020-04-14 12:41:25	00:13:54	
Checkpoint16 Begin	2020-04-14 12:55:42		00:28:11
Checkpoint16 End	2020-04-14 13:09:37	00:13:55	
Checkpoint17 Begin	2020-04-14 13:23:55		00:28:13
Checkpoint17 End	2020-04-14 13:37:45	00:13:49	
Measurement Interval End	2020-04-14 13:42:22	-	-

Clause 6: SUT, Driver and Communication

6.1 Remote Terminal Emulator (RTE)

If the RTE is commercially available, then its inputs must be specified. Otherwise, a description must be supplied of what inputs (e.g., scripts) to the RTE had been used.

The RTE software used was internally developed. The RTE simulated web users. It generated random input data based on the benchmark requirements and recorded response times and other statistics for each transaction cycle.

6.2 Emulated Components

It must be demonstrated that the functionality and performance of the components being emulated in the Driver System are equivalent to the priced system. The results of the test described in Clause 6.6.3.4 must be disclosed.

No components were emulated by the driver system.

6.3 Functional Diagrams

A complete functional diagram of both the benchmark configuration and the configuration of the proposed (target) system must be disclosed. A detailed list of all hardware and software functionality being performed on the Driver System and its interface to the SUT must be disclosed.

The diagram in Figure 0.1 shows the tested and priced benchmark configurations.

6.4 Networks

The network configuration of both the tested services and proposed (target) services which are being represented and a thorough explanation of exactly which parts of the proposed configuration are being replaced with the Driver System must be disclosed.

The bandwidth of the networks used in the tested/priced configuration must be disclosed.

The diagram in Figure 0.1 shows the network configuration between the components of the tested configuration. The RTE and the SUT are connected through a 25/100Gbit switch.

The network bandwidths are listed in Figure 0.1.

6.5 Operator Intervention

If the configuration requires operator intervention (see Clause 6.6.6), the mechanism and the frequency of this intervention must be disclosed.

No operator intervention is required to sustain eight hours at the reported throughput.

Clause 7: Pricing

7.1 Hardware and Software Pricing

A detailed list of hardware and software used in the priced system must be reported. Each separately orderable item must have vendor part number, description, and release/revision level, and either general availability status or committed delivery date. If package-pricing is used, vendor part number of the package and a description uniquely identifying each of the components of the package must be disclosed. Pricing source and effective date(s) of price(s) must also be reported.

The details of the hardware and software are reported in the front of this report as part of the Executive Summary.

7.2 Three Year Price

The total 3-year price of the entire configuration must be reported, including: hardware, software, and maintenance charges. Separate component pricing is recommended. The basis of all discounts used must be disclosed.

The pricing details for this TPC-C result are reported in the front of this report as part of the Executive Summary.

7.3 Availability Dates

The committed delivery date for general availability (availability date) of products used in the price calculations must be reported. When the priced system includes products with different availability dates, the reported availability date for the priced system must be the date at which all components are committed to be available.

Category	Available
Server Hardware	Sep, 2020
Server Software	Now (date of publication)
Goldilocks v3.1 Standard Edition	Now (date of publication)

Clause 8: Reporting

8.1 Full Disclosure Report

A Full Disclosure report is required in order for results to be considered compliant with the TPC-C benchmark specification

This document constitute the Full Disclosure Report for the TPC-C benchmark result describes within.

Clause 9: Auditor Attestation

9.1 Auditor Information

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 who to contact in order to obtain further information regarding the audit process.

This benchmark was audited by:

InfoSizing

Francois Raab
20 Kreg Ln
Manitou Springs, CO 80829

Phone: +1 (719) 473-7555

www.sizing.com

9.2 Attestation Letter

The auditor's attestation letter is included in the following pages.

InfoSizing

TPC Certified Auditor

Kihan Choi
Principal Research Engineer
Telecommunications Technology Association (TTA)
Bundang-ro 47, Bundang-gu, Seongnam-city
Gyeonggi-do, 13591, Republic of Korea

August 14, 2020

I verified the TPC Benchmark C (TPC-C™ v5.11.0) performance of the following configuration:

Platform:	Supermicro A+ Server 1124US-TNRP
Operating System:	Red Hat Enterprise Linux Server 7.7
Database Manager:	Goldilocks v3.1 Standard Edition

The results were:

Performance Metric	**380,475 tpmC**
Number of Users	300,000

Server	**Supermicro A+ Server 1124US-TNRP**		
CPUs	2 x AMD EPYC 7F72 (3.2GHz, 192MB L3)		
Memory	4 TB		
Disks	***Qty***	***Size***	***Type***
	2	7.68TB	NVMe SSD (internal)
	6	3.8TB	SATA SSD (internal)

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 transactions were correctly implemented
- The database records were the proper size
- The database was properly scaled and populated
- The ACID properties were met

20 Kreg Lane • Manitou Springs, CO 80829 • 719-473-7555 • www.sizing.com

- Input data was generated according to the specified percentages
- The transaction cycle times included the required keying and think times
- The reported response times were correctly measured.
- At least 90% of all delivery transactions met the 80 Second completion time limit
- All 90% response times were under the specified maximums
- The measurement interval was representative of steady state conditions
- The reported measurement interval was over 120 minutes
- Checkpoints intervals were under 30 minutes
- The 60-day storage requirement was correctly computed
- The system pricing was verified for major components and maintenance

Additional Audit Notes:

None.

Respectfully Yours,

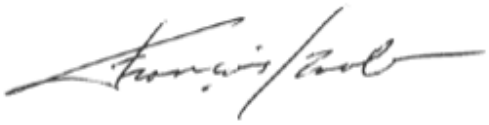

François Raab, TPC Certified Auditor

20 Kreg Lane • Manitou Springs, CO 80829 • 719-473-7555 • www.sizing.com

Appendix A: Source Code

The source code and scripts used to implement the benchmark is provided as a soft appendix. This soft appendix includes the following files:

```
\ACID
    \ACID\include
    \ACID\src
    \ACID\include\acid.h
    \ACID\src\atom.c
    \ACID\src\compare.c
    \ACID\src\consist.c
    \ACID\src\Delivery.c
    \ACID\src\isol1.c
    \ACID\src\isol2.c
    \ACID\src\isol3.c
    \ACID\src\isol4.c
    \ACID\src\isol5.c
    \ACID\src\isol6.c
    \ACID\src\isol7.c
    \ACID\src\isol8.c
    \ACID\src\isol9.c
    \ACID\src\Makefile
    \ACID\src\NewOrder.c
    \ACID\src\OrderStatus.c
    \ACID\src\Payment.c
    \ACID\src\support.c
\bin
    \bin\load.sh
\html
    \html\DeliveryInput.html
    \html\MainMenu.html
    \html\NewOrderInput.html
    \html\OrderStatusInput.html
    \html\PaymentInput.html
    \html\StockLevelInput.html
\include
    \include\spt_proc.h
    \include\support.h
\java
    \java\Common.java
    \java\Delivery.java
    \java\NewOrder.java
    \java\OrderStatus.java
    \java\Payment.java
    \java\StockLevel.java
\scripts
    \scripts\analyze_system.sql
    \scripts\analyze_table.sql
    \scripts\analyze_table_district.sql
    \scripts\analyze_table_item.sql
    \scripts\analyze_table_new_order.sql
    \scripts\analyze_table_orders.sql
    \scripts\analyze_table_order_line.sql
```

```
    \scripts\analyze_table_stock.sql
    \scripts\analyze_table_warehouse.sql
    \scripts\audit.sql
    \scripts\checkpoint.py
    \scripts\count.sql
    \scripts\create_audit_table.sql
    \scripts\create_index.sql
    \scripts\create_procedure.sql
    \scripts\create_table.sql
    \scripts\create_tablespace.sql
    \scripts\dbcheck.sql
    \scripts\dbtables.sql
    \scripts\runcheck.sql
    \scripts\sys
         \scripts\sys\be
             \scripts\sys\be\part_info.sh
             \scripts\sys\be\reboot_info.sh
             \scripts\sys\be\sw_info.sh
             \scripts\sys\be\sys_info.sh
\src
    \src\free_space.c
    \src\load.c
    \src\load_new.c
    \src\Makefile
    \src\support.c
```

Appendix B: Tunable Parameters

goldilocks.properties.conf

```
TRANSACTION_COMMIT_WRITE_MODE = 1
TRANSACTION_TABLE_SIZE = 1024
UNDO_RELATION_COUNT = 1024
LOG_BUFFER_SIZE = 3G
LOG_FILE_SIZE = 100G
LOG_GROUP_COUNT = 5
PENDING_LOG_BUFFER_COUNT = 8
SPIN_COUNT = 1
BUSY_WAIT_COUNT = 1000
SYSTEM_TABLESPACE_DIR = '/data/db/db1'
SYSTEM_MEMORY_UNDO_TABLESPACE_SIZE = 2G
SYSTEM_MEMORY_TEMP_TABLESPACE_SIZE = 2G
SHARED_MEMORY_STATIC_SIZE = 4G
PARALLEL_IO_FACTOR = 5
PARALLEL_IO_GROUP_1 = '/data/db/db1'
PARALLEL_IO_GROUP_2 = '/data/db/db2'
PARALLEL_IO_GROUP_3 = '/data/db/db3'
PARALLEL_IO_GROUP_4 = '/data/db/db4'
PARALLEL_IO_GROUP_5 = '/data/db/db5'
LOG_DIR = '/wal'
CLIENT_MAX_COUNT = 1024
PROCESS_MAX_COUNT = 1024
PARALLEL_LOAD_FACTOR = 16
SHARED_SESSION = NO
CONTROL_FILE_COUNT = 2
CONTROL_FILE_0 = '/wal/control_0.ctl'
CONTROL_FILE_1 = '/wal/control_1.ctl'
```

limit.conf

```
# /etc/security/limits.conf
#
#This file sets the resource limits for the users logged in via PAM.
#It does not affect resource limits of the system services.
#
#Also note that configuration files in /etc/security/limits.d
directory,
#which are read in alphabetical order, override the settings in this
#file in case the domain is the same or more specific.
#That means for example that setting a limit for wildcard domain here
#can be overriden with a wildcard setting in a config file in the
#subdirectory, but a user specific setting here can be overriden only
#with a user specific setting in the subdirectory.
#
#Each line describes a limit for a user in the form:
#
#<domain>        <type>  <item>  <value>
#
#Where:
#<domain> can be:
#        - a user name
#        - a group name, with @group syntax
#        - the wildcard *, for default entry
#        - the wildcard %, can be also used with %group syntax,
#                 for maxlogin limit
#
#<type> can have the two values:
#        - "soft" for enforcing the soft limits
#        - "hard" for enforcing hard limits
#
#<item> can be one of the following:
#        - core - limits the core file size (KB)
#        - data - max data size (KB)
#        - fsize - maximum filesize (KB)
#        - memlock - max locked-in-memory address space (KB)
#        - nofile - max number of open file descriptors
#        - rss - max resident set size (KB)
#        - stack - max stack size (KB)
#        - cpu - max CPU time (MIN)
#        - nproc - max number of processes
#        - as - address space limit (KB)
#        - maxlogins - max number of logins for this user
#        - maxsyslogins - max number of logins on the system
#        - priority - the priority to run user process with
#        - locks - max number of file locks the user can hold
#        - sigpending - max number of pending signals
#        - msgqueue - max memory used by POSIX message queues (bytes)
#        - nice - max nice priority allowed to raise to values: [-20,
19]
#        - rtprio - max realtime priority
#
#<domain>      <type>  <item>         <value>
#

#*               soft    core            0
#*               hard    rss             10000
#@student        hard    nproc           20
#@faculty        soft    nproc           20
#@faculty        hard    nproc           50
#ftp             hard    nproc           0
#@student        -       maxlogins       4

tpcc    soft    nofile  1000000
tpcc    hard    nofile  1000000
tpcc    soft    nproc   unlimited
tpcc    hard    nproc   unlimited
```

server.xml

```
<?xml version='1.0' encoding='utf-8'?>
<!--
  Licensed to the Apache Software Foundation (ASF) under one or more
  contributor license agreements.  See the NOTICE file distributed
with
  this work for additional information regarding copyright ownership.
  The ASF licenses this file to You under the Apache License, Version
2.0
  (the "License"); you may not use this file except in compliance with
  the License.  You may obtain a copy of the License at

      http://www.apache.org/licenses/LICENSE-2.0

  Unless required by applicable law or agreed to in writing, software
  distributed under the License is distributed on an "AS IS" BASIS,
  WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or
implied.
  See the License for the specific language governing permissions and
  limitations under the License.
-->
<!-- Note:  A "Server" is not itself a "Container", so you may not
     define subcomponents such as "Valves" at this level.
     Documentation at /docs/config/server.html
 -->
<Server port="8005" shutdown="SHUTDOWN">
  <Listener
className="org.apache.catalina.mbeans.JmxRemoteLifecycleListener"
        rmiRegistryPortPlatform="9840" rmiServerPortPlatform="9841"/>
  <Listener
className="org.apache.catalina.startup.VersionLoggerListener" />
  <!-- Security listener. Documentation at /docs/config/listeners.html
  <Listener className="org.apache.catalina.security.SecurityListener"
/>
  -->
  <!--APR library loader. Documentation at /docs/apr.html -->
  <Listener className="org.apache.catalina.core.AprLifecycleListener"
SSLEngine="on" />
  <!--Initialize Jasper prior to webapps are loaded. Documentation at
/docs/jasper-howto.html -->
  <Listener className="org.apache.catalina.core.JasperListener" />
  <!-- Prevent memory leaks due to use of particular java/javax APIs--
>
  <Listener
className="org.apache.catalina.core.JreMemoryLeakPreventionListener"
/>
  <Listener
className="org.apache.catalina.mbeans.GlobalResourcesLifecycleListener
" />
  <Listener
className="org.apache.catalina.core.ThreadLocalLeakPreventionListener"
/>

  <!-- Global JNDI resources
       Documentation at /docs/jndi-resources-howto.html
  -->
  <GlobalNamingResources>
    <!-- Editable user database that can also be used by
         UserDatabaseRealm to authenticate users
    -->
    <Resource name="UserDatabase" auth="Container"
              type="org.apache.catalina.UserDatabase"
              description="User database that can be updated and
saved"

factory="org.apache.catalina.users.MemoryUserDatabaseFactory"
              pathname="conf/tomcat-users.xml" />
  </GlobalNamingResources>
```

```
  <!-- A "Service" is a collection of one or more "Connectors" that
share
       a single "Container" Note:  A "Service" is not itself a
"Container",
       so you may not define subcomponents such as "Valves" at this
level.
       Documentation at /docs/config/service.html
   -->
  <Service name="Catalina">

    <!--The connectors can use a shared executor, you can define one
or more named thread pools-->
    <!--
    <Executor name="tomcatThreadPool" namePrefix="catalina-exec-"
        maxThreads="150" minSpareThreads="4"/>
    -->


    <!-- A "Connector" represents an endpoint by which requests are
received
         and responses are returned. Documentation at :
         Java HTTP Connector: /docs/config/http.html (blocking & non-
blocking)
         Java AJP  Connector: /docs/config/ajp.html
         APR (HTTP/AJP) Connector: /docs/apr.html
         Define a non-SSL HTTP/1.1 Connector on port 8080
    -->
          <Connector port="8080"
                        acceptCount="150000"
                        maxConnections="141000"
                        connectionTimeout="20000000"
                        maxThreads="1024"
                        maxKeepAliveRequests="-1" keepAliveTimeout="-
1"

protocol="org.apache.coyote.http11.Http11NioProtocol"
redirectPort="8443"
         />

          <!--
    <Connector port="8080" protocol="HTTP/1.1"
               connectionTimeout="20000"
               redirectPort="8443" />
          -->
    <!-- A "Connector" using the shared thread pool-->
    <!--
    <Connector executor="tomcatThreadPool"
               port="8080" protocol="HTTP/1.1"
               connectionTimeout="20000"
               redirectPort="8443" />
    -->
    <!-- Define a SSL HTTP/1.1 Connector on port 8443
         This connector uses the BIO implementation that requires the
JSSE
         style configuration. When using the APR/native
implementation, the
         OpenSSL style configuration is required as described in the
APR/native
         documentation -->
    <!--
    <Connector port="8443"
protocol="org.apache.coyote.http11.Http11Protocol"
               maxThreads="150" SSLEnabled="true" scheme="https"
secure="true"
               clientAuth="false" sslProtocol="TLS" />
    -->

    <!-- Define an AJP 1.3 Connector on port 8009 -->
    <Connector port="8009" protocol="AJP/1.3" redirectPort="8443" />


    <!-- An Engine represents the entry point (within Catalina) that
processes
         every request.  The Engine implementation for Tomcat stand
alone
         analyzes the HTTP headers included with the request, and
passes them
         on to the appropriate Host (virtual host).
         Documentation at /docs/config/engine.html -->

    <!-- You should set jvmRoute to support load-balancing via AJP
ie :
    <Engine name="Catalina" defaultHost="localhost" jvmRoute="jvm1">
    -->
    <Engine name="Catalina" defaultHost="localhost">

      <!--For clustering, please take a look at documentation at:
          /docs/cluster-howto.html  (simple how to)
          /docs/config/cluster.html (reference documentation) -->
      <!--
      <Cluster
className="org.apache.catalina.ha.tcp.SimpleTcpCluster"/>
      -->

      <!-- Use the LockOutRealm to prevent attempts to guess user
passwords
           via a brute-force attack -->
      <Realm className="org.apache.catalina.realm.LockOutRealm">
        <!-- This Realm uses the UserDatabase configured in the global
JNDI
             resources under the key "UserDatabase".  Any edits
             that are performed against this UserDatabase are
immediately
             available for use by the Realm.  -->
        <Realm className="org.apache.catalina.realm.UserDatabaseRealm"
               resourceName="UserDatabase"/>
      </Realm>

      <Host name="localhost"  appBase="webapps"
            unpackWARs="true" autoDeploy="true">

        <!-- SingleSignOn valve, share authentication between web
applications
             Documentation at: /docs/config/valve.html -->
        <!--
        <Valve
className="org.apache.catalina.authenticator.SingleSignOn" />
        -->

        <!-- Access log processes all example.
             Documentation at: /docs/config/valve.html
             Note: The pattern used is equivalent to using
pattern="common"
       <Valve className="org.apache.catalina.valves.AccessLogValve"
directory="logs"
              prefix="localhost_access_log." suffix=".txt"
              pattern="%h %l %u %t &amp;quot;%r&amp;quot; %s %b" />-->

      </Host>
    </Engine>
  </Service>
</Server>
```

Sysctl fe.conf

```
# sysctl settings are defined through files in
# /usr/lib/sysctl.d/, /run/sysctl.d/, and /etc/sysctl.d/.
#
# Vendors settings live in /usr/lib/sysctl.d/.
# To override a whole file, create a new file with the same in
# /etc/sysctl.d/ and put new settings there. To override
# only specific settings, add a file with a lexically later
# name in /etc/sysctl.d/ and put new settings there.
#
# For more information, see sysctl.conf(5) and sysctl.d(5).


net.core.netdev_max_backlog=300000
net.ipv4.tcp_max_syn_backlog=300000
net.ipv4.ip_local_port_range=10000 65535
net.core.somaxconn=65535
net.ipv4.tcp_tw_reuse=1
```

Appendix C: Price Quotations

DB Server

견 적 서

견적일자 : 2020년 8월 10일 월요일 Quotation No : SS2020-0810-001

수신 : 한국정보통신기술협회
참조 : 최기한 님

아래와 같이 견적합니다.

㈜슈퍼솔루션 대표이사:김성현
서울시 금천구 벚꽃로 244 (가산동, 벽산디지털밸리5차 1107호~1108호,707호)
Tel : (국번없이) 1661-8560 | FAX : 02-2082-5543 | 등록번호:113-86-20325

견적금액	일금 일억일천구백삼십오만 원정	₩ 119,350,000	(VAT포함가)

NO.	구분	제품명 & 제품사양	수량	세트	소비자단가	제안단가 (VAT별도)	합계금액 (VAT별도)
A	System	A+ Server 1124US-TNRP (출시 예정)	식	1		108,500,000	108,500,000
	플랫폼 SKUs	AS-1124US-TNRP	1				
	Socket	Dual Socket SP3 (LGA 4094) support; System on Chip (SoC)					
	Processor	Dual AMD EPYC™ 7002 Series Processors					
	Memory	32x DIMM slots, Supports up to 8TB Registered ECC DDR4 3200MHz, 8-channel memory bus					
	Network	Dual 10GBase-T LAN ports via Intel® X710-AT2 + Dual 10G SFP+ via Intel® X710-AT2					
	VGA	Aspeed AST2500 BMC					
	IPMI	Support for Intelligent Platform Management Interface v.2.0					
	PCI-Express	2x PCI-E 4.0 x16 (FH/HL, 9.5"L), 1x PCI-E 4.0 x16 (LP), 1x PCI-E 4.0 x16 (internal proprietary LP)					
	Hot-swap	12x Hot-swap 2.5" U.2 NVMe drive bays or SAS3/SATA3 (via optional kits)					
	Power Supply	1+1 1200W Redundant Power Supplies Titanium Level (96%)					
	Form Factor	1U Rackmount					
	Standard Parts List (Items Included)						
	Processor	AMD EPYC™ 7F72 Processor (24C/48T, 192M Cache, 3.20 GHz, 240W)	2				
	Memory	256GB DDR4 ECC-LRDIMM	16				
	NVMe	7.6TB NVMe U.2 2.5"	2				
	SSD	3.84TB SATA3 2.5"	6				
	RAID	Broadcom 3108 8-port PCI-E SAS3(12Gb/s) controller; 2GB cache; Supports RAID 0, 1, 5, 6, 10, 50, 60	1				
	NETWORK	MCX4121A-ACAT ConnectX-4 Lx EN NIC, 25Gbe, 2x Port SFP28	1				
	Warranty	OS4HR3: 3년 무상, 24시간 On-Site Service	1				

소 계	₩	108,500,000
부 가 세	₩	10,850,000
합 계 금 액	₩	119,350,000

* REMARK 영업담당자 : 이선기 부장 010-8029-6564 / sklee@supersolution.co.kr

▷ 유효기간 : 견적일후, 90일 이내 ▷ 결제조건 : 협의

▷ 하자보증기간: 검수 후 3년 ▷ 유지보수요율: 공급가의 7%

* 이 견적서 검토 하신 후, 가격 및 모든조건들을 수용하는 조건으로 발주 진행을 원하실 경우, 아래 내용을 작성하셔서 주시거나 귀사 발주서를 보내주시기 바랍니다.(이 견적서에 날인한 경우, 발주서와 동일한 효력을 가집니다)

1. 발주 담당자:
2. 담당자연락처:
3. 배송지상세정보 :
4. 결제 조건 및 예정일 :

(명판 또는 직인 날인)

WAS Server

견 적 서

견적일자 : 2020년 7월 7일 문서번호 SW2020-0707-05

아래와 같이 견적합니다.

TO.	한국정보통신기술협회	귀중
CC.	최기한님	귀하
Tel.	010-5110-7276	
FAX		
Mail	kihanc@tta.or.kr	

상 호.	㈜슈퍼솔루션	대 표	김성현
주 소.	서울시 금천구 가산동 60-73 벽산디지털밸리5차 1107호		
전 화.	02-2082-5540	팩 스	02-2082-5543
등록번호	113-86-20325		

㈜슈퍼솔루션 차장 이성운
swlee@supersolution.co.kr
Mobile : 010-8029-6563

1. 대금결재 : 협의
2. 납품장소 : 귀사 지정장소(서울,경지 지역 외 운임비 별도)

[단위 : 원. VAT별도]

NO.	모델명	제품사양	수량	세트	소비자단가	제안단가	금 액
A	Part NO.	**A+ Server 1123US-TN10RT**		2		9,850,000	19,700,000
1	**Product SKUs**	**AS-1123US-TN10RT**					
	Processor	Dual AMD EPYC™ 7001/7002 series Processors					
	System Memory	32 DIMM slots, Supports up to 8TB Registered ECC DDR4 3200MHz, 8-channel memory bus					
	Chipset	System on Chip (SoC)					
	NVMe	10 U.2 NVMe support					
	Network Controllers	Intel® X540 Controller					
	Onboard VGA	Aspeed AST2500 BMC					
	Expansion Slots	1 PCI-E 3.0 x16 (FH, 9.5"L) slot 1 PCI-E 3.0 x8 (LP) slot, 1 M.2 PCI-E 3.0 x4 M-Key NVMe					
	IPMI	Support for Intelligent Platform Management Interface v.2.0					
	Form Factor	1U Rackmount					
	Dimensions	Height 3.5" (89mm) Width 17.2" (437mm) Depth 25.5" (647mm)					
	Hot-swap	10 hot-swap 2.5" U.2 NVMe drive bays					
	Fans	8 heavy-duty PWM fans with optimal fan speed control					
	Power	1+1 1000W Redundant Power Supplies with PMBus					
		Additional Parts					
	Processor	AMD EPYC™ 7262 3.2GHz/3.4GHz, 8C/16T, 128M	2				
	Memory	16GB DDR4 ECC REG RDIMM	8				
	NVMe	1TB U.2 PCIe NVMe 3.1 X4 2.5 in	1				
	AOC	AOC-MH25G-M2S2TM-O SIOM 2+2-port 25G, 10G	1				
	Warranty	3년 무상, 24시간 On-Site Service	1				

* REMARK

※ 상기 제품에 대한 OS 및 Software는 별도입니다.

※ 견적유효기간 : 견적후 90일(직인생략)

소 액	₩19,700,000
세 액	₩1,970,000
총 액(포함가)	₩21,670,000

RHEL/JWS

rockPLACE
(주)락플레이스 www.rockplace.co.kr

㈜락플레이스
135-120 서울시 강남구 신사동 634-10 윤당빌딩 3층 Tel.02)6251.7788 Fax.02)6251.6677

rockPLACE, Inc.
3F, Yundang bldg, 634-10, Shinsa-dong, Gangnam-gu,Seoul, Korea Tel : 822-6251-7788 Fax: 822-6251-6677

견 적 서

REF No.	: 2020RP08-0703		TERMS AND CONDITION
DATE	: 2020. 08. 07.		
COMPANY	: TTA	납 기	: 발주후 4주이내
ATTN	: 최기한 선임연구원님 귀하 TEL : 031-780-9256	유지보수	: 납품일로부터 1년
Email	: kihanc@tta.or.kr	결제조건	: 익월말 현금
FROM	: ㈜락플레이스 이왕모 과장 TEL : 010-9116-4680	유효기간	: 견적일로부터 3개월

下記와 같이 見積합니다.

㈜ 락플레이스
대표이사 서 동 식

ITEM DESCRIPTION

(VAT 별도, 단위 : 원)

Part No.	Description	수량	소비자가	공급단가	공급합계
OS	Red Hat Enterprise Linux Operating System Platform				
RH00004F3	Red Hat Enterprise Linux Server, Standard (Physical or Virtual Nodes) 3Year support : Easy ISOs: OS, Source, Documentation ISO Images 가상화 Guest OS : 2guests Red Hat Network 서비스 : 3년 Phone,email Support : 09:00 ~ 17:00 Scope of Coverage : Standard Maximum Memory Support: Unlimited	3	4,098,000	2,400,000	7,200,000
연간기술지원	연간 방문 기술지원 (옵션)				
RP-CPS(OS)	rockPLACE Support Carepack - Linux Standard (3년) per Server 3 Year, 24x7, 4hr response 이메일, 전화, 원격지원, 현장지원 서비스 On Site Support - Total 연간 10회 Support (아래 지원내역에 준함) - Installation & Startup Service Included - Problem tracking/Emergency assistance - Update, Patch 작업 지원 - 서비스, 시스템 환경, 네트워크 환경 설정 변경 지원 - 인수 시험, 성능 시험, 비상 복구 훈련 지원	3	6,000,000	2,000,000	6,000,000
소 계 금 액					13,200,000

Part No.	Description	수량	소비자가	공급단가	공급합계
WEB	**Red Hat JBoss Web Server**				
MW0232248F3	Red Hat JBoss Web Server, 16-Core Standard 3Year - 전화/웹 지원 : 월-금, 9 a.m. - 5 p.m. 4시간내 응답 - unlimited incidents,	2	7,799,000	5,000,000	10,000,000
연간기술지원	연간 방문 기술지원 (옵션)				
RP-CPS(WAS)	rockPLACE Support Carepack - JBoss Standard (3년) per 16Core 3 Year, 24x7, 4hr response 이메일, 전화, 원격지원, 현장지원 서비스 On Site Support - Total 10회 Support (아래 지원내역에 준함) - Installation & Startup Service Included - Problem tracking/Emergency assistance - Update, Patch 작업 지원 - 서비스, 시스템 환경, 네트워크 환경 설정 변경 지원 - 인수 시험, 성능 시험, 비상 복구 훈련 지원	2	12,000,000	6,000,000	12,000,000
소 계 금 액					22,000,000

합 계	35,200,000
부가세	3,520,000
합 계(부가세포함)	38,720,000

Remarks

1. Red Hat 제품은 년간 Subscription 제품이며, 기간이 만료되실 경우 Renewal을 하셔야합니다.
2. 발주 시에는 반드시 고객정보(엔드유져명, 담당자, 연락처, Email)가 있어야 합니다.
3. OnSite 방문지원이 필요하실 경우에는 케어팩을 구매하셔야 합니다.

Network

ICRAFT 아이크래프트㈜ www.icraft21.com
06057 서울특별시 강남구 언주로 726, 12F(논현동, 두산빌딩) Tel. 02-541-0474(대표) Fax. 02-543-1442

견 적 서

수 신 TTA
사 업 명 Goldilocks Standard Edition for LINUX
견 적 일 2020년 8월 12일
담 당 자 황세진 선임연구원

등록번호	220-81-75308
상 호	아이크래프트㈜
대표이사	박 우 진 (직인생략)
주 소	서울시 강남구 언주로 726, 12F
종 목	소프트웨어 개발, 컨설팅 통신장비 외

견 적 금 액 : ₩ 37,856,000 (VAT별도) [단위: 원, VAT별도]

No.	모델명	제 품 상 세	수량	단위	공급가		비고
					공급단가	공급금액	
Mellanox Ethernet Solution							
1	MSN2700-CS2F	Mellanox Spectrum based 100GbE 1U Open Ethernet Switch with Onyx, 32 QSFP28 ports, 2 Power Supplies (AC), x86 CPU, standard depth, P2C airflow, Rail Kit	1	EA	29,000,000	29,000,000	
2	유지보수	장애지원, 12개월, 24x7, 4시간 이내 접수	3	년	2,900,000	8,700,000	
3	MCP7F00-A001R30N	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 1m, Colored, 30AWG, CA-N	1	EA	156,000	156,000	
합 계 (VAT별도)						₩ 37,856,000	

작성자. 엔터프라이즈사업본부 D1팀 문지연 차장 Moblie. 010-7766-8854 Email. moon7@icraft21.com NVIDIA.

※ Remark
1 본 견적은 Goldilocks Standard Edition for LINUX 사업에 한합니다
2 결제조건 : 납품 후 현금결제
3 유효기간 : 견적 후 120일 이내
4 납품일정 : 발주 후 60일 예정 (제품에 따라 지연될 수 있습니다.)
5 설치지원 : 별도 (서울/경기 : 50만원, 지방 : 70만원, 케이블링 작업 별도 협의)
6 발주 전 스위치 Airflow 를 확인해 주세요 (발주 후 변경 불가)

DBMS

Quotation

㈜TTA 貴中

Title : TPC-C Performance&Quality Authentication

참 조 : 황세진 선임님
견적일자 : 2020년 08월 10일
유효기간 : 견적일로부터 4개월

대표이사 : 김 기 완 (인)
주 소 : 서울시 영등포구 당산로
171 금강펜테리움IT타워 604호
영업대표 : 사업본부 최승렬 이사
전화번호 : 010-9312-0188
e-mail : slchoi@sunjesoft.com

* Goldilocks Standard Edition for LINUX 1식 (단위 : 원)

No.	Description	Unit List Price	Q'ty	Total Amount Price	Offer Price
1	Goldilocks Ver 3.1 DBMS Standard Edition	₩96,000,000	1 Set(s)	₩96,000,000	₩32,000,000
	- Query Processes Module				
	- Storage Management Module				
	Goldilocks DBMS License Fee	License Proposal Price			₩32,000,000
2	DBMS Implementaion & Supports	₩10,000,000	3 Set(s)	₩30,000,000	₩14,400,000
	Goldilocks Technical Supports Fee(3yr)	Support Proposal Price			₩14,400,000
Total Amount(VAT Exclude)				₩126,000,000	₩46,400,000

Goldilocks Total Amount (Offer Price)	₩46,400,000

* For Technical supports, it indicates 24 x 7 x 4 hours of support

Keyboard/Monitor