

TPC Benchmark™ C

Full Disclosure Report



First Edition
25-Sep-2022

Using
IBM DB2 Advanced Edition
on
LTechKorea LKG2312

First Edition: 25-Sep-2023

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*
- *DB2 is a registered trademark of IBM.*
- *JBoss is a registered trademark of RedHat, Inc.*
- *Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation.*
- *All other trademarks and copyrights are properties of their respective owners.*

Table of Contents

- TABLE OF CONTENTS.....3**
- ABSTRACT5**
- PREFACE6**
- GENERAL ITEMS.....11**
- 0.1 APPLICATION CODE AND DEFINITION STATEMENTS 11
- 0.2 BENCHMARK SPONSOR..... 11
- 0.3 PARAMETER SETTINGS..... 11
- 0.4 CONFIGURATION DIAGRAMS 12
- CLAUSE 1: LOGICAL DATABASE DESIGN.....13**
- 1.1 TABLE DEFINITIONS..... 13
- 1.2 PHYSICAL ORGANIZATION OF DATABASE 13
- 1.3 INSERT AND DELETE OPERATIONS 13
- 1.4 HORIZONTAL OR VERTICAL PARTITIONING..... 13
- 1.5 REPLICATION OR DUPLICATION..... 13
- CLAUSE 2: TRANSACTION AND TERMINAL PROFILES.....14**
- 2.1 RANDOM NUMBER GENERATION..... 14
- 2.2 INPUT/OUTPUT SCREENS 14
- 2.3 PRICED TERMINAL FEATURE 14
- 2.4 PRESENTATION MANAGERS..... 14
- 2.5 TRANSACTION STATISTICS 15
- 2.6 QUEUING MECHANISM..... 15
- CLAUSE 3: TRANSACTION AND SYSTEM PROPERTIES.....16**
- 3.1 ATOMICITY 16
 - 3.1.1 *Atomicity of Completed Transactions* 16
 - 3.1.2 *Atomicity of Aborted Transactions* 16
- 3.2 CONSISTENCY 16
- 3.3 ISOLATION..... 17
- 3.4 DURABILITY 21
 - 3.4.1 *Durable Media Failure* 21
 - 3.4.2 *Instantaneous Interruption, Loss of Memory*..... 22
- CLAUSE 4: SCALING AND DATABASE POPULATION.....23**
- 4.1 CARDINALITY OF TABLES 23
- 4.2 DATABASE IMPLEMENTATION 23
- 4.3 DISTRIBUTION OF DATABASE FILES 24
- 4.4 60-DAY SPACE 24
- CLAUSE 5: PERFORMANCE METRICS.....24**
- 5.1 TPC BENCHMARK C METRICS..... 25
- 5.2 RESPONSE TIMES 25
- 5.3 KEYING AND THINK TIMES 25

5.4 DISTRIBUTION AND PERFORMANCE CURVES.....25

 5.4.1 *Response Time frequency distribution curves*.....25

 5.4.2 *Response Time versus throughput*.....29

 5.4.3 *Think Time frequency distribution*.....30

 5.4.4 *Throughput versus elapsed time*.....31

5.5 STEADY STATE DETERMINATION32

5.6 WORK PERFORMED DURING STEADY STATE32

5.7 MEASUREMENT PERIOD DURATION34

5.8 TRANSACTION STATISTICS34

5.9 CHECKPOINTS34

CLAUSE 6: SUT, DRIVER AND COMMUNICATION.....34

6.1 REMOTE TERMINAL EMULATOR (RTE).....34

6.2 EMULATED COMPONENTS34

6.3 FUNCTIONAL DIAGRAMS35

6.4 NETWORKS.....35

6.5 OPERATOR INTERVENTION35

CLAUSE 7: PRICING35

7.1 HARDWARE AND SOFTWARE PRICING35

7.2 THREE YEAR PRICE.....35

7.3 AVAILABILITY DATES.....35

CLAUSE 8: REPORTING36

8.1 FULL DISCLOSURE REPORT36

CLAUSE 9: AUDITOR ATTESTATION36

9.1 AUDITOR INFORMATION36

9.2 ATTESTATION LETTER36

APPENDIX A: SOURCE CODE39

APPENDIX B: TUNABLE PARAMETERS40

APPENDIX C: PRICE QUOTATIONS.....43

Abstract

This report documents the methodology and results of the TPC Benchmark™ C (TPC-C) test conducted by TTA on the IBM DB2 11.5.8 Advanced Edition on LtechKorea LKG2312

IBM DB2 11.5.8 on LtechKorea LKG2312

Company Name	System Name	Database Software	Operating System
Telecommunications Technology Association	LTechKorea LKG2312	IBM DB2 11.5.8 Advanced Edition	RedHat Enterprise Linux 7.9

TPC Benchmark™ C Metrics

Total System Cost	TPC-C Throughput	Price/Performance	Availability Date
₩ 1,351,171,600 (KRW)	279,185 tpmC	4,840 KRW/tpmC	Available Now

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



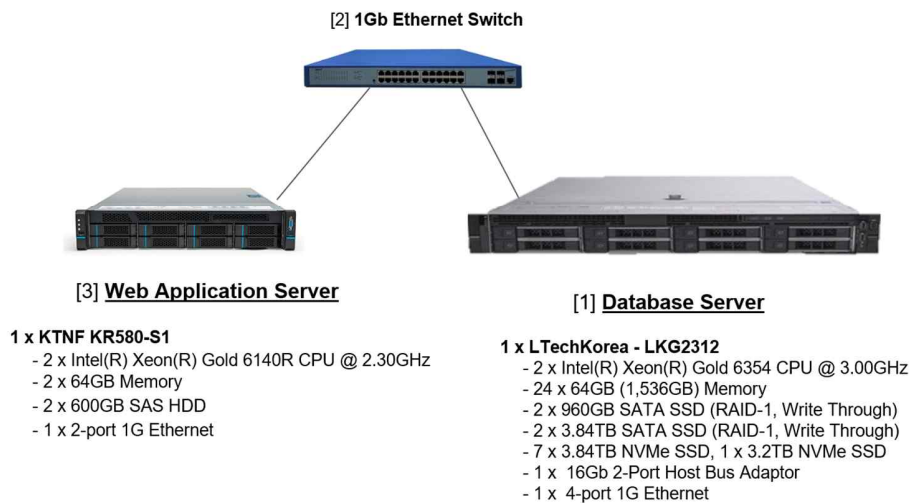
IBM DB2 11.5.8 Advanced Edition on LtechKorea LKG2312

TPC-C Version 5.11.0
TPC Pricing 2.8.0

Report Date
25-Sep-2023

Total System Cost	TPC-C Throughput	Price/Performance		Availability Date
₩ 1,351,171,600 (KRW)	279,185 tpmC	4,840 KRW/tpmC		Available Now
Server Processors/Cores/Threads	Database Manager	Operating System	Other Software	Number of Users
2/36/72	IBM DB2 11.5.8 Advanced Edition	RHEL 7.9	JBoss Web Server	220,000

Priced Configuration (LTK)



System Components	DB Server		WAS Server	
	Quantity	Description	Quantity	Description
Processors/Cores/Threads	2/36/72	Intel(R) Xeon(R) Gold 6354 CPU @ 3.00GHz	2/36/72	Intel(R) Xeon(R) Gold 6140 CPU @ 2.30GHz
Memory	24	64 GB	2	64GB
Storage Controller	1	Broadcom / LSI MegaRAID SAS-3 3108	1	Broadcom / LSI MegaRAID SAS-3 3108
Storage Device	2 2 7 1	960GB SAS HDD (Write Through) 3.84TB SAS SSD (Write Through) 3.84TB NVMe SSD (Write Through) 3.2 TB NVMe SSD (Write Through)	2	600GB SAS SSD
Total Storage Capacity		39.68TB		



IBM DB2 11.5.8 Advanced Edition on LtechKorea LKG2312

**TPC-C Version 5.11.0
TPC Pricing 2.8.0**

**Report Date
25-Sep-2023**

Available Now

Description	Part Number	Source	Unit Price	Qty	Price	3-Yr. Maint. Price
Server Hardware						
DB Server - LTeckKorea(LKG2312)	LKG2312	1	30,180,000	1	30,180,000	
Intel® Xeon® Gold 6354 18C/36T (Up To 3.6Ghz)	CPU	1	(included)	2		
64GB DDR4 Memory	RAM	1	(included)	24		
SSD 960GB, Enterprise	DISK1	1	(included)	2		
SSD 3.84TB, Enterprise	DISK2	1	(included)	2		
SSD 3.2TB NVME SSD	DISK3	1	(included)	1		
SSD 3.84, Enterprise NVMe, U.2, Gen4	DISK4	1	(included)	7		
SAS3108(2GB) 12Gbps Controller	RAID	1	(included)	1		
RJ45 1G 4port	NIC	1	(included)	1		
16Gb/s 2port (include GBIC)	HBA	1	(included)	1		
LKG-2312 Server(power 1600Watt) 3.5/2.5 inch disk bay * 12EA	CHASSIS	1	(included)	1		
3year, 24x7x4hr Onsite Support Service 1	Maintenance	1	(included)	1		
WAS Servers - KTNF(KR580S1)						
WAS Servers - KTNF(KR580S1)	KR580S1	2	16,700,000	1	16,700,000	
Intel® Xeon® Scalable Gold 6140 (2.30Ghz, 18core)	CPU	2	(included)	2		
64GB DDR4 ECC RDIMM Memory	Memory	2	(included)	2		
600GB SAS 12Gb/s 15K RPM (128MB)	HDD	2	(included)	2		
3year, 24x7x4hr Onsite Support Service 1	Maintenance	2	(included)	1		
Server Hardware Sub Total					46,880,000	
Client/Server Software						
Red Hat Enterprise Linux Server Standard 3yrs	RH00004F3	3	4,089,000	2	8,178,000	
RHEL Server Standard Maintenance - 3yrs 24x7x4hrs	RP-CPS(OS)	3	8,000,000	2		16,000,000
Red Hat JBoss Web Server 4-Core Standard 3Year	MW00123F3	3	2,138,000	9	19,242,000	
JBoss Web Server per 4Core 3Year Maintenance	RP-CPS(WAS)	3	12,000,000	9		108,000,000
IBM DB2 Advanced Edition Cartridge for IBM Cloud Pak for Data Virtual Processor Core License + SW Subscription & Support 12 Months (Ver.11.5.8.0)	-	4	128,307,000	36	4,619,052,000	
IBM Db2 License Renewal (2 years)	-	4	189,625,800	2		379,251,600
IBM Db2 Implementaion & Support	-	4	20,000,000	3		60,000,000
Software Sub Total					4,646,472,000	563,251,600

Other Hardware						
DASAN Networks access switch - D2224GP, 24port POE	network switch	5	1,900,000	3	5,700,000	
Other Hardware Sub Total					5,700,000	
Discounts*						
Red Hat OS Discount					-3,378,000	-12,000,000
Red Hat JBoss Discount					-7,902,000	-54,000,000
SW Discount - DB2					-3,833,852,000	0
Discounts Sub Total					-3,845,132,000	-66,000,000
Total					853,920,000	497,251,600

Pricing Notes 1) LTechKorea, Inc. 4) D.BtoE Inc. 2) KTNF, Inc. 5) DASAN Networks 3) Rockplace Inc.		Three year cost of ownership KRW(W): 1,351,171,600 TPC-C throughput: 279,185 tpmC Price/Performance: 4,840 ₩ / tpmC
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.		

Benchmark implementation and results independantly audited by Doug Johnson 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.



IBM DB2 11.5.8 Advanced Edition on LtechKorea LKG2312

TPC-C Version 5.11.0
TPC Pricing 2.8.0

Report Date
25-Sep-2023

Available Now

MQTh, computed Maximum Qualified Throughput **279,185 tpmC**

Response Times (seconds)	Min	Average	90 th	Max
New-Order	0.102	0.112	0.104	35.353
Payment	0.101	0.111	0.102	43.442
Order-Status	0.101	0.110	0.102	2.524
Delivery (interactive portion)	0.101	0.101	0.101	0.262
Delivery (deferred portion)	0.001	0.019	0.005	3.290
Stock-Level	0.102	0.111	0.103	2.523
Menu	0.101	0.101	0.102	2.196

Emulated Display Delay: 0.1 sec.

Transaction Mix	Percent	Number
New-Order	44.980%	117,262,770
Payment	43.011%	112,129,869
Order-Status	4.002%	10,434,208
Delivery	4.003%	10,436,393
Stock-Level	4.003%	10,436,555

Keying Times (seconds)	Min	Average	Max
New-Order	18.001	18.001	18.059
Payment	3.001	3.001	3.081
Order-Status	2.001	2.001	2.012
Delivery	2.001	2.001	2.025
Stock-Level	2.001	2.001	2.052

Think Times (seconds)	Min	Average	Max
New-Order	0.001	12.045	120.501
Payment	0.001	12.045	120.501
Order-Status	0.001	10.046	100.501
Delivery	0.001	5.025	50.301
Stock-Level	0.001	5.028	50.301

Test Duration	
Ramp-up time	65 min
Measurement Interval (MI)	420 min
Checkpoints in MI	-
Checkpoint Interval (Average / Max)	-
Number of Transactions in MI (all types)	260,669,795

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 IBM. and LTechKorea 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.

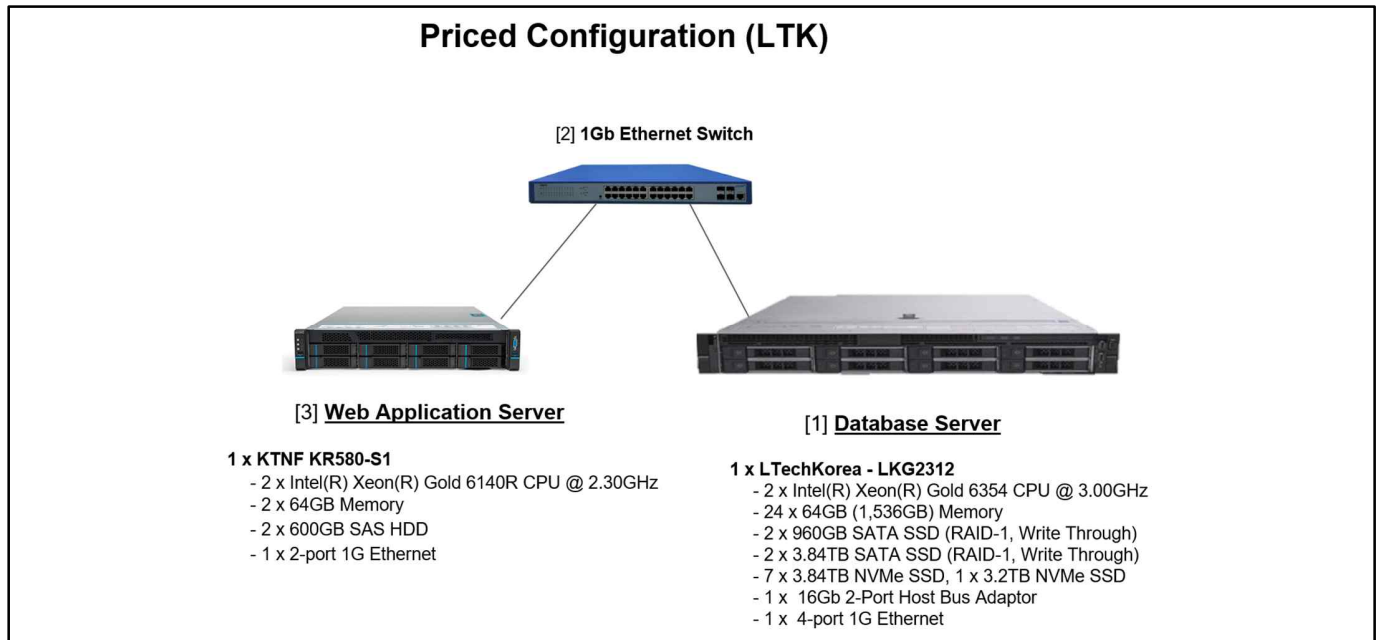
Appendix B contains the tunable parameters for the database, the operating system, and the transaction monitor.

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
LSI MegaRAID SAS-3 3108	Internal	RAID 1	2 x SAS 960 GB SSD	OS
LSI MegaRAID SAS-3 3108	Internal	RAID 1	7 x 3.84TB NVMe 1 x 3.2 TB NVMe	Database files
LSI MegaRAID SAS-3 3108	Internal	RAID 1	2 x 3.84TB SSD	Redo Logs

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	99.001%
	Remote warehouse order lines	0.999%
	Rolled back transactions	1.001%
	Average items per order	10.000
Payment	Home warehouse	84.998%
	Remote warehouse	15.002%
	Accessed by last name	59.998%
Order Status	Accessed by last name	60.020%
Delivery	Skipped transactions	0
Transaction Mix	New Order	44.980%
	Payment	43.011%
	Order status	4.002%
	Delivery	4.003%
	Stock level	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 [$\max(\text{no_o_id})$] minus the first Order ID [$\min(\text{no_o_id})$] 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:

$$\begin{aligned} \max(\text{no_o_id}) - \min(\text{no_o_id}) + 1 &= \text{rows in NEW-ORDER} \\ &\text{where } (\text{o_w_id} = \text{no_w_id}) \text{ and } (\text{o_d_id} = \text{no_d_id}) \end{aligned}$$

4. For each District within a Warehouse, the sum of Order-Line counts [$\text{sum}(\text{o_ol_cnt})$] 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:

$$\text{sum}(\text{o_ol_cnt}) = \text{rows in the ORDER-LINE table for the Warehouse and District}$$

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 multiple log file switches.
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 has been blocked and wait until T1 to be committed.
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.
6. T3 and T2 returned the same Order that 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 has been blocked and wait until T1 to be roll-backed.
4. T1 was allowed to ROLLBACK. T2 returned the same Order that T0 had returned
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 was blocked and did not committed until T1 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. **Then, T2 was committed as soon as T1 was committed.**
7. The prices of Items X and Y were retrieved again(**T3**). The values matched the values set by T2.

The Execution followed Case **A**, where **T2 does stall** and no transaction is rolled back. Query **T3** verifies the price change made by **T2**.

Isolation Test 8

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 was blocked and did not committed until T1 was committed.
5. T1 was resumed and the NEW_ORDER table was read again. No qualifying row was found.
6. T1 completed and was committed. Then, T2 was committed as soon as T1 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.

Isolation Test 9

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 was blocked and did not committed until T1 was committed.
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. Then, T2 was committed as soon as T1 was 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 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. One of the data disks containing the TPC-C tables is disabled by removing it physically. Since the data disks are configured with redundancy, the transactions continued to run without interruption.
6. The test is allowed to run for a minimum of 5 minutes after data disk failed.
7. One of the log disks containing the transaction log is disabled by removing it physically. Since the log disks are configured with redundancy, the transactions continued to run without interruption.
8. The test is allowed to run until the test ends, the Measurement Interval would be at least an hour.
9. The RTE run is completed.
10. The consistency is verified.
11. Step 1 is repeated, giving count-2.
12. The RTE result file is used to determine the number of New-Order transactions successfully completed during the full run.
13. 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.
14. 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 30 minutes at full load (after ramp-up).
5. The primary power to the back-end server is shut off (removing both power cords), which removed power from all system components, including memory.
6. The RTE is shutdown.
7. Power is restored to the database server and the system performs an automatic recovery.
8. IBM DB2 is restarted and performs an automatic recovery.
9. Step 1 is repeated, giving count-2.
10. The consistency is verified.
11. The RTE result file is used to determine the number of New-Order transactions successfully completed during the full run.
12. 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.
13. 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	22,000
District	220,000
Customer	660,000,000
History	660,000,000
Order	660,000,000
New Order	198,000,000
Order Line	6,598,479,460
Stock	2,200,000,000
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.

The database manager used for this testing was IBM DB2 11.5.8. DB2 is relational DBMS

DB2 remote stored procedures and embedded SQL statements were used.

The DB2 stored procedures were invoked 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 seven 3.84TB disks and one 3.2TB disk configured as RAID1. The database log files are stored on two 3.84TB disks configured as RAID1.

Table 4.3: Database file locations

Name	Location	Description
C000000.LRG	/opt/db2inst1/db2inst1/NODE0000/TESTDB/	System tables and dictionary
C000000X.LRG	/db2data1/db2inst1/NODE0000/TESTDB/ /db2data2/db2inst1/NODE0000/TESTDB/ /db2data3/db2inst1/NODE0000/TESTDB/ /db2data4/db2inst1/NODE0000/TESTDB/	Database data files
S00000XX.LOG	/dblogs/NODE0000/LOGSTREAM0000/	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

TABLE_NAME	Record Count	TOTAL(KBYTE)	USED(KBYTE)	TBSP_NAME	TOTAL(KBYTE)	USED(KBYTE)
WAREHOUSE	22,000	2,512	2,512	TPCC_TB_W	5,180,416	3,076,352
WAREHOUSE_IDX		368	368			
DISTRICT	220,000	27,076	27,076	TPCC_TB_D	204,800	33,024
DISTRICT_IDX		4,660	4,660			
CUSTOMER	660,000,000	431,341,264	431,341,264	TPCC_TB_C	926,206,976	729,403,264
CUSTOMER_IDX		16,923,080	16,923,080			
NEW_ORDER	198,000,000	4,330,048	4,330,048	TPCC_TB_NO	20,971,520	11,744,256
NEW_ORDER_IDX		5,076,924	5,076,924			
ORDERS	660,000,000	33,862,652	33,862,652	TPCC_TB_O	126,877,696	93,597,952
ORDERS_IDX		16,923,080	16,923,080			
ORDERS_IDX		20,625,004	20,625,004			
NEW_ORDER	198,000,000	4,330,048	4,330,048	TPCC_TB_NO	20,971,520	11,744,256
NEW_ORDER_IDX		5,076,924	5,076,924			
ORDER_LINE	6,598,479,460	528,153,000	528,153,000	TPCC_TB_OL	1,549,175,808	1,107,954,944
ORDER_LINE_IDX		198,450,512	198,450,512			
ORDER_LINE_IDX		58,784,672	58,784,672			
ITEM	100,000	9,044	9,044	TPCC_TB_I	204,800	21,248
ITEM_IDX		1,668	1,668			
STOCK	2,200,000,000	772,583,172	772,583,172	TPCC_TB_S	1,758,524,416	1,251,917,696
STOCK_IDX		46,560,848	46,560,848			
HISTORY	660,000,000	48,905,440	48,905,440	TPCC_TB_H	104,857,600	97,815,168
Total Before 8-hour run					4,492,204,032	3,295,563,904

TABLE_NAME	Record Count	TOTAL(KBYTE)	USED(KBYTE)	TBSP_NAME	TOTAL(KBYTE)	USED(KBYTE)
WAREHOUSE	22,000	2,512	2,512	TPCC_TB_W	5,180,416	3,076,352
WAREHOUSE_IDX		368	368			
DISTRICT	220,000	27,076	27,076	TPCC_TB_D	204,800	33,024
DISTRICT_IDX		4,660	4,660			
CUSTOMER	660,000,000	432,257,236	431,341,264	TPCC_TB_C	926,206,976	729,396,864
CUSTOMER_IDX		16,923,080	16,923,080			
NEW_ORDER	211,946,577	33,554,728	33,554,728	TPCC_TB_NO	20,971,520	15,906,048
NEW_ORDER_IDX		7,303,680	6,329,568			
ORDERS	796,893,007	8,523,432	8,518,172	TPCC_TB_O	126,877,696	98,346,368
ORDERS_IDX		40,838,888	40,838,832			
ORDERS_IDX		23,321,036	23,321,036			
NEW_ORDER	211,946,577	33,554,728	33,554,728	TPCC_TB_NO	20,971,520	15,906,048
NEW_ORDER_IDX		7,303,680	6,329,568			
ORDER_LINE	7,967,454,159	636,946,304	636,946,180	TPCC_TB_OL	1,549,175,808	1,107,840,896
ORDER_LINE_IDX		270,519,648	270,519,648			
ORDER_LINE_IDX		78,418,140	78,418,140			
ITEM	100,000	9,044	9,044	TPCC_TB_I	204,800	21,248
ITEM_IDX		1,668	1,668			
STOCK	2,200,000,000	772,583,172	772,583,172	TPCC_TB_S	1,758,524,416	1,251,917,696
STOCK_IDX		46,560,848	46,560,848			
HISTORY	792,153,047	58,805,548	58,805,496	TPCC_TB_H	104,857,600	97,476,096
Total After 8-hour run					4,492,204,032	3,304,014,592

Table	Rows	Data	Index	Initial Population	5% Growth	8-Hour Growth	Required Runtime Space
WAREHOUSE	22000	2,512	368	2,880	144	0	3,024
DISTRICT	220000	27,076	4,660	31,736	1,587	0	33,323
CUSTOMER	660000000	431,341,264	29,148,468	460,489,732	23,024,487	0	483,514,219
NEW_ORDER	198000000	4,330,048	5,076,924	9,406,972	470,349	0	9,877,321
ITEM	100000	9,044	1,668	10,712	536	0	11,248
STOCK	2200000000	772,583,172	46,560,848	819,144,020	40,957,201	0	860,101,221
HISTORY	660000000	48,905,440	0	48,905,440	0	9,929,938	58,835,378
ORDERS	660000000	33,862,652	37,548,084	71,410,736	0	6,875,596	78,286,332
ORDER_LINE	6,598,479,460	528,153,000	257,235,184	785,388,184	0	107,238,106	892,626,290
Total		1,819,214,208	375,576,204	2,194,790,412	64,454,303	124,043,640	2,383,288,355

60-Day Requirements		Storage configuration	
Dynamic-Space	610,921,092	Disk Type	Capacity
Free-Space	0	NVMe(Data)	2,628,306,712
Static-Space	1,583,869,320	# of Disk(LUN)	3
Daily-Growth	124,043,640	Capacity	7,884,920,136
Daily-Spread	0	NVMe(Data)	3,749,290,960
60-Day Space	9,026,487,723	SSD(for logs)	2,928,290,816
		Total Disk Space	14,562,501,912

Storage Requirements	
Total Disk Space	14,562,501,912
Log space used	20,480,200
60-Day Space	9,026,487,723
Remaining Space	5,515,533,989

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

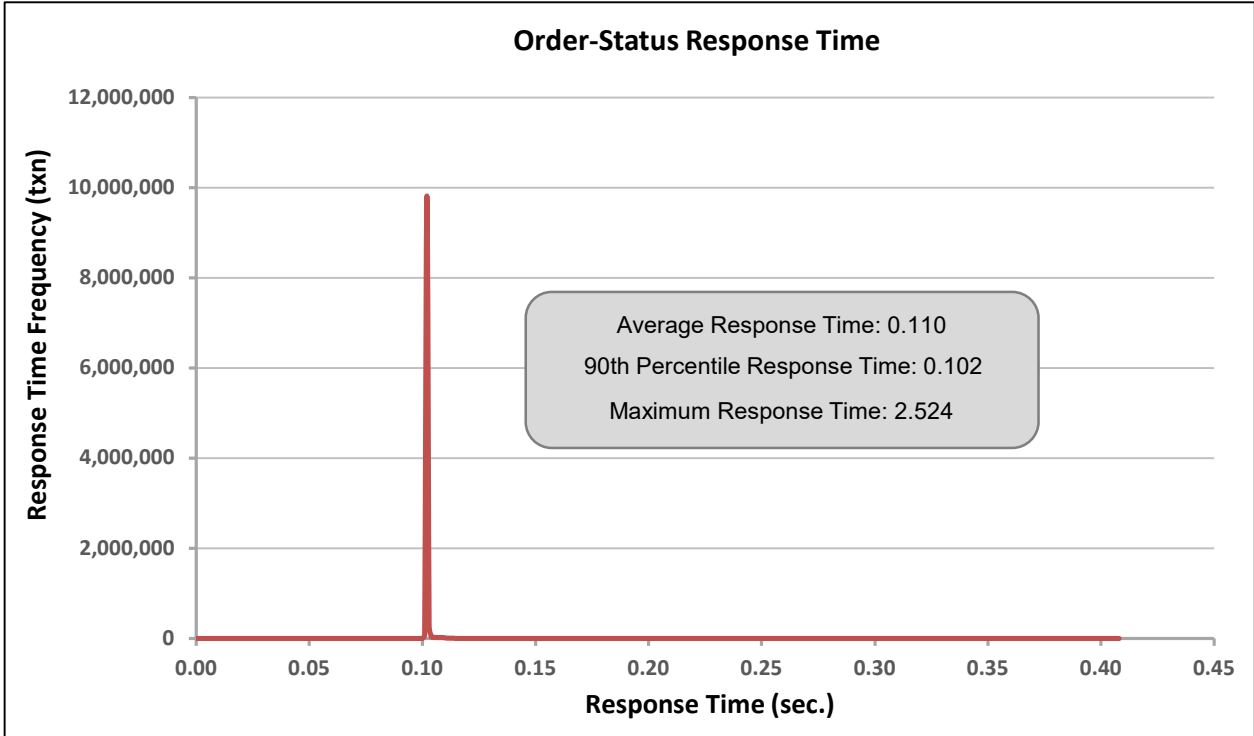


Figure 5.4.1.3: Order-Status RT Frequency Distribution




Figure 5.4.1.4: Delivery (Interactive) RT Frequency Distribution

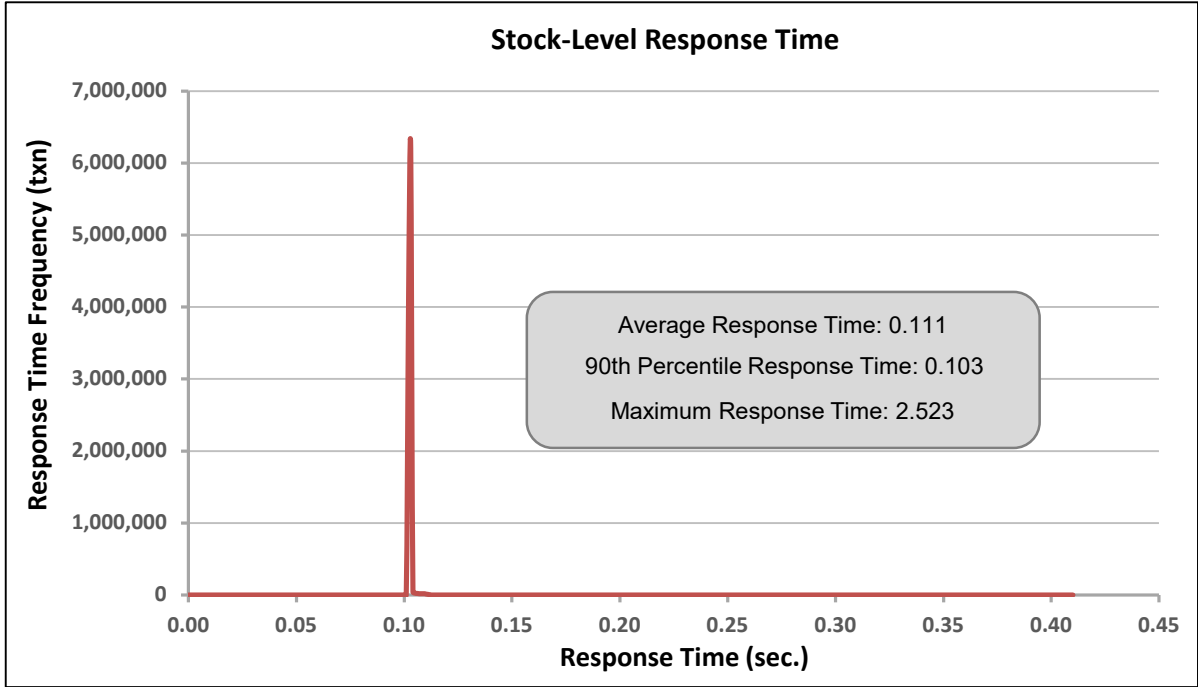


Figure 5.4.1.5: Stock-Level RT Frequency Distribution

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

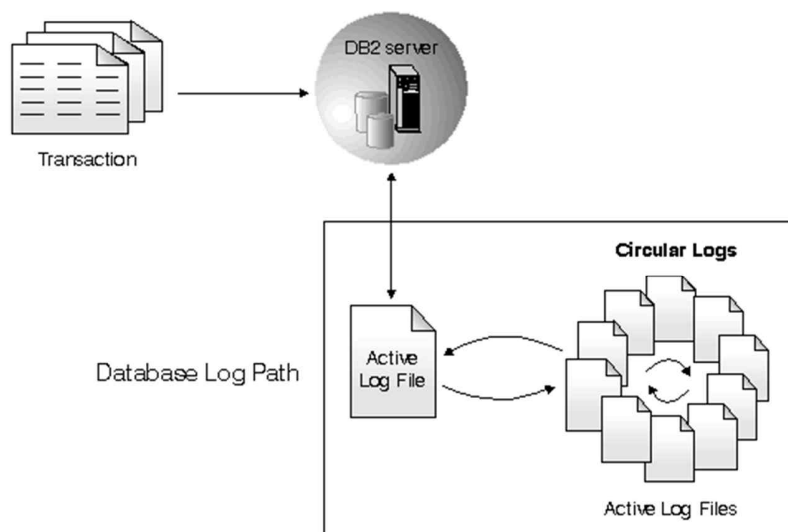
The method used to determine that the SUT had reached a steady state prior to commencing the measurement interval must be disclosed.

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, IBM DB2 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. DB2 uses a write-ahead-logging protocol to guarantee recovery. This prevents the system from losing any committed transactions. In this test, dbms logging method is set to circular logging. Active logs are used during crash recovery to prevent a failure (system power or application error) from leaving a database in an inconsistent state.



When DB2 Server changes a database table with an update, insert, or delete operation, the change is initially made in memory, not on disk. When there is not enough space in the memory buffer to read in or write additional data pages, DB2 Server will make space by flushing some modified pages to disk. Modified pages are also written to disk as part of the “Soft” checkpoint to ensure that no updates remain unflushed for longer than the allowed time. The soft checkpoint interval time can be set by the ‘page_age_trgt_mcr’ parameter which is set to 240 seconds in this test. The following table 5.6 shows the soft checkpoint interval measured at 5-minute intervals during the 8-hour full run test.

Table 5.6: soft Checkpoint intervals

No.	Interval time	No.	Interval time	No.	Interval time
1	04:04.8	35	04:31.0	69	04:13.8
2	04:16.2	36	04:09.1	70	04:15.6
3	04:16.8	37	04:11.0	71	04:17.2
4	04:09.4	38	04:12.7	72	04:07.1
5	04:10.6	39	04:14.7	73	04:08.6
6	04:10.7	40	04:16.5	74	04:11.2
7	04:10.0	41	04:06.8	75	04:12.8
8	04:10.1	42	04:09.3	76	04:14.8
9	04:10.2	43	04:12.0	77	04:16.6
10	04:10.3	44	04:14.7	78	04:17.8
11	04:10.6	45	04:18.0	79	04:06.6
12	04:22.6	46	04:09.6	80	04:08.2
13	04:10.5	47	04:12.5	81	04:08.5
14	04:11.0	48	04:14.9	82	04:09.9
15	04:12.2	49	04:17.8	83	04:11.4
16	04:14.8	50	04:08.3	84	04:12.7
17	04:07.5	51	04:10.3	85	04:14.5
18	04:12.2	52	04:12.5	86	04:16.0
19	04:16.3	53	04:14.8	87	04:17.5
20	04:08.2	54	04:16.8	88	04:18.8
21	04:10.5	55	04:18.7	89	04:08.4
22	04:12.4	56	04:21.3	90	04:22.2
23	04:14.4	57	04:13.1	91	04:12.0
24	04:16.2	58	04:16.1	92	04:12.9
25	04:17.6	59	04:07.1	93	04:14.4
26	04:07.2	60	04:09.3	94	04:15.3
27	04:09.0	61	04:11.9	95	04:16.7
28	04:11.2	62	04:13.9	96	04:18.8
29	04:13.3	63	04:15.9	97	04:08.5
30	04:16.2	64	04:17.4	98	-
31	04:08.1	65	04:18.8	99	-
32	04:11.9	66	04:08.6	100	-
33	04:14.3	67	04:10.5	101	-
34	04:04.9	68	04:12.0	102	-

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

DB2 uses a write-ahead-logging protocol to guarantee recovery. This protocol uses “Soft” checkpoint to write least-recently-used database pages to disk independent of transaction commit. However, enough log information to redo/undo the change to a database pages is committed to disk before the database page itself is written. This protocol therefore renders checkpoint unnecessary for DB2 . For a more detailed description of the general principles of the write-ahead-logging protocol, see the IBM research paper, “ARIES: A Transaction Recovery Method Supporting Fine Granularity Locking and Partial Rollbacks Using Write-Ahead Logging,” by C. Mohan, Database Technology Institute, IBM Almaden Research Center. (<https://dl.acm.org/doi/10.1145/128765.128770>)

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 1Gbps 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.

All components of the priced system are available as of the date of this 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

Doug Johnson

63 Lourdes Drive

Leominster, MA, 01453 USA

Phone: +1 (978) 343-6562

www.sizing.com

9.2 Attestation Letter

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

Sejin Hwang
Senior Research Engineer
Telecommunications Technology Association (TTA)
Bundang-ro 47, Bundang-gu, Seongnam-city
Gyeonggi-do, 13591, Republic of Korea

September 25, 2023

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

Platform: LTechKorea LKG2312
Operating System: Red Hat Enterprise Linux 7.9
Database Manager: IBM DB2 Advanced Edition

The results were:

Performance Metric **279,185 tpmC**
Number of Users 220,000

Server

LTechKorea LKG2312

CPU	2x Intel® Xeon® Gold 6354 (3.00 GHz, 18-core, 39 MB Cache)		
Memory	1,536 GB		
Storage	Qty	Size	Type
	2	960 GB	SAS HDD
	2	3.84 TB	SAS SSD
	7	3.84 TB	NVMe SSD
	1	3.2 TB	NVMe SSD

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

The following verification items were given special attention:

- The transactions were correctly implemented.
- The database records were the proper size.
- The database was properly scaled and populated.
- The ACID properties were met.
- 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 at least 120 minutes.
- Write-ahead logging was active during the measurement interval.
- The 60-day storage requirement was correctly computed.
- The system pricing was verified for major components and maintenance.

Additional Audit Notes:

None.

Respectfully Yours,

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

Doug Johnson, Certified TPC Auditor

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\procedure
  \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
  \ACID\procedure\getCustAll.sql
  \ACID\procedure\getCustBalance.sql
  \ACID\procedure\getCustBalanceName.sql
  \ACID\procedure\getCustomerCount.sql
  \ACID\procedure\getCustomerID.sql
  \ACID\procedure\getDistrictCount.sql
  \ACID\procedure\getDistrictID.sql
  \ACID\procedure\getItemData.sql
  \ACID\procedure\getItemPrice.sql
  \ACID\procedure\getMaxWarehouseID.sql
  \ACID\procedure\getMinWarehouseID.sql
  \ACID\procedure\getNewOrderCount.sql
  \ACID\procedure\getNextDistrictOrderID.sql
  \ACID\procedure\getOrderCarrierEntry.sql
  \ACID\procedure\getOrderCount.sql
  \ACID\procedure\getOrderCustomerID.sql
  \ACID\procedure\getOrderCustomerIDByNO.sql
  \ACID\procedure\getOrderEx.sql
  \ACID\procedure\getOrderID.sql
  \ACID\procedure\getOrderlineAmount.sql
  \ACID\procedure\getOrderLineCount.sql
  \ACID\procedure\getStockData.sql
  \ACID\procedure\getSysdate.sql
  \ACID\procedure\getWarehouseCount.sql
  \ACID\procedure\getWarehouseCustomer.sql
  \ACID\procedure\getWarehouseID.sql
  \ACID\procedure\updateCustomerData.sql
  \ACID\procedure\updateDistrictGetOrderline.sql
  \ACID\procedure\updateDistrictYTD.sql
  \ACID\procedure\updateWarehouseYTD.sql
  \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\audit.sql
  \scripts\count.sql
  \scripts\create_procedure_delivery.sql
  \scripts\create_procedure_neworder.sql
  \scripts\create_procedure_orderstatus.sql
  \scripts\create_procedure_payment.sql
  \scripts\create_procedure_sotcklevel.sql
  \scripts\create_table.sql
  \scripts\create_tablespace.sql
  \scripts\dbcheck.sql
  \scripts\dbtables.sql
  \scripts\runcheck.sql
  \scripts\tunning.sql
  \src
  \src\free_space.c
  \src\load.c
  \src\Makefile
  \src\support.c
```

Appendix B: Tunable Parameters

Database Configuration

Database configuration release level	= 0x1500	Lifetime of cached credentials (AUTHN_CACHE_DURATION)	= 3
Database release level	= 0x1500	Max number of users in the cache (AUTHN_CACHE_USERS)	= 0
Update to database level pending	= NO (0x0)	Max DB files open per database (MAXFILOP)	= 61440
Database territory	= US	Active log space disk capacity (MB) (LOG_DISK_CAP)	= 0
Database code page	= 1208	Log file size (4KB) (LOGFILSIZ)	= 204800
Database code set	= UTF-8	Number of primary log files (LOGPRIMARY)	= 13
Database country/region code	= 1	Number of secondary log files (LOGSECOND)	= 12
Database collating sequence	= IDENTITY	Changed path to log files (NEWLOGPATH)	=
Alternate collating sequence (ALT_COLLATE)	=	Path to log files	=
Number compatibility	= OFF	/dblogs/NODE0000/LOGSTREAM0000/	
Varchar2 compatibility	= OFF	Overflow log path (OVERFLOWLOGPATH)	=
Date compatibility	= OFF	Mirror log path (MIRRORLOGPATH)	=
Database page size	= 4096	First active log file	=
Statement concentrator (STMT_CONC)	= OFF	Block log on disk full (BLK_LOG_DSK_FUL)	= NO
Discovery support for this database (DISCOVER_DB)	= ENABLE	Block non logged operations (BLOCKNONLOGGED)	= NO
Restrict access	= NO	Percent max primary log space by transaction (MAX_LOG)	= 0
Default query optimization class (DFT_QUERYOPT)	= 5	Num. of active log files for 1 active UOW (NUM_LOG_SPAN)	= 0
Degree of parallelism (DFT_DEGREE)	= ANY	Percent log file reclaimed before soft chkpt (SOFTMAX)	= 0
Continue upon arithmetic exceptions (DFT_SQLMATHWARN)	= NO	Target for oldest page in LBP (PAGE_AGE_TRGT_MCR)	= 240
Default refresh age (DFT_REFRESH_AGE)	= 0	HADR database role	= STANDARD
Default maintained table types for opt (DFT_MTTB_TYPES)	= SYSTEM	HADR local host name (HADR_LOCAL_HOST)	=
Number of frequent values retained (NUM_FREQVALUES)	= 10	HADR local service name (HADR_LOCAL_SVC)	=
Number of quantiles retained (NUM_QUANTILES)	= 20	HADR remote host name (HADR_REMOTE_HOST)	=
Decimal floating point rounding mode (DECFLT_ROUNDING)	= ROUND_HALF_EVEN	HADR remote service name (HADR_REMOTE_SVC)	=
DECIMAL arithmetic mode (DEC_ARITHMETIC)	=	HADR instance name of remote server (HADR_REMOTE_INST)	=
Large aggregation (LARGE_AGGREGATION)	= NO	HADR timeout value (HADR_TIMEOUT)	= 120
Backup pending	= NO	HADR target list (HADR_TARGET_LIST)	=
All committed transactions have been written to disk	= NO	HADR log write synchronization mode (HADR_SYNCMODE)	= NEARSYNC
Rollforward pending	= NO	HADR spool log data limit (4KB) (HADR_SPOOL_LIMIT)	=
Restore pending	= NO	AUTOMATIC(0)	
Upgrade pending	= NO	HADR log replay delay (seconds) (HADR_REPLAY_DELAY)	= 0
Multi-page file allocation enabled	= YES	HADR peer window duration (seconds) (HADR_PEER_WINDOW)	= 0
Log retain for recovery status	= NO	First log archive method (LOGARCHMETH1)	= OFF
User exit for logging status	= NO	Archive compression for logarchmeth1 (LOGARCHCOMPR1)	= OFF
Self tuning memory (SELF_TUNING_MEM)	= ON	Options for logarchmeth1 (LOGARCHOPT1)	=
Size of database shared memory (4KB) (DATABASE_MEMORY)	=	Second log archive method (LOGARCHMETH2)	= OFF
AUTOMATIC(320568320)		Archive compression for logarchmeth2 (LOGARCHCOMPR2)	= OFF
Database memory threshold (DB_MEM_THRESH)	= 100	Options for logarchmeth2 (LOGARCHOPT2)	=
Max storage for lock list (4KB) (LOCKLIST)	= 12500000	Failover log archive path (FAILARCHPATH)	=
Percent. of lock lists per application (MAXLOCKS)	= AUTOMATIC(98)	Number of log archive retries on error (NUMARCHRETRY)	= 5
Package cache size (4KB) (PCKCACHESZ)	= 128000	Log archive retry Delay (secs) (ARCHRETRYDELAY)	= 20
Sort heap thres for shared sorts (4KB) (SHEAPTHRES_SHR)	=	Vendor options (VENDOROPT)	=
AUTOMATIC(41472)		Auto restart enabled (AUTORESTART)	= ON
Sort list heap (4KB) (SORTHEAP)	= AUTOMATIC(8294)	Index re-creation time and redo index build (INDEXREC)	= SYSTEM
Database heap (4KB) (DBHEAP)	=	(RESTART)	
AUTOMATIC(13256)		Log pages during index build (LOGINDEXBUILD)	= OFF
Catalog cache size (4KB) (CATALOGCACHE_SZ)	= 2500000	Default number of loadrec sessions (DFT_LOADREC_SES)	= 1
Log buffer size (4KB) (LOGBUFFSZ)	= 10240	Number of database backups to retain (NUM_DB_BACKUPS)	= 12
Utilities heap size (4KB) (UTIL_HEAP_SZ)	= 50000	Recovery history retention (days) (REC_HIS_RETENTN)	= 90
SQL statement heap (4KB) (STMTHRAP)	= AUTOMATIC(8192)	Auto deletion of recovery objects (AUTO_DEL_REC_OBJ)	= OFF
Default application heap (4KB) (APPLHEAPSZ)	= AUTOMATIC(256)	TSM management class (TSM_MGMTCLASS)	=
Application Memory Size (4KB) (APPL_MEMORY)	=	TSM node name (TSM_NODENAME)	=
AUTOMATIC(40000)		TSM owner (TSM_OWNER)	=
Statistics heap size (4KB) (STAT_HEAP_SZ)	= AUTOMATIC(4384)	TSM password (TSM_PASSWORD)	=
Interval for checking deadlock (ms) (DLCHKTIME)	= 60000	Automatic maintenance (AUTO_MAINT)	= OFF
Lock timeout (sec) (LOCKTIMEOUT)	= 50	Automatic database backup (AUTO_DB_BACKUP)	= OFF
Changed pages threshold (CHNGPGS_THRESH)	= 80	Automatic table maintenance (AUTO_TBL_MAINT)	= OFF
Number of asynchronous page cleaners (NUM_IOCLEANERS)	= AUTOMATIC(36)	Automatic runstats (AUTO_RUNSTATS)	= OFF
Number of I/O servers (NUM_IOSERVERS)	= AUTOMATIC(72)	Real-time statistics (AUTO_STMT_STATS)	= OFF
Sequential detect flag (SEQDETECT)	= YES	Statistical views (AUTO_STATS_VIEWS)	= OFF
Default prefetch size (pages) (DFT_PREFETCH_SZ)	= AUTOMATIC	Automatic sampling (AUTO_SAMPLING)	= OFF
Track modified pages (TRACKMOD)	= NO	Automatic column group statistics (AUTO_CG_STATS)	= OFF
Default number of containers	= 1	Automatic reorganization (AUTO_REORG)	= OFF
Default tablespace extentsize (pages) (DFT_EXTENT_SZ)	= 32	Auto-Revalidation (AUTO_REVAL)	= DISABLED
Max number of active applications (MAXAPPLS)	= 10000	Currently Committed (CUR_COMMIT)	= ON
Average number of active applications (AVG_APPLS)	= AUTOMATIC(1)	CHAR output with DECIMAL input (DEC_TO_CHAR_FMT)	= NEW
		Enable XML Character operations (ENABLE_XMLCHAR)	= YES
		Enforce Constraint (DDL_CONSTRAINT_DEF)	= YES
		Enable row compression by default (DDL_COMPRESSION_DEF)	= NO
		Replication site ID (REPL_SITE_ID)	= 0
		Monitor Collect Settings	
		Request metrics (MON_REQ_METRICS)	= NONE
		Activity metrics (MON_ACT_METRICS)	= NONE
		Object metrics (MON_OBJ_METRICS)	= NONE
		Routine data (MON_RTN_DATA)	= NONE
		Routine executable list (MON_RTN_EXECLIST)	= OFF
		Unit of work events (MON_UOW_DATA)	= NONE
		UOW events with package list (MON_UOW_PKGLIST)	= OFF
		UOW events with executable list (MON_UOW_EXECLIST)	= OFF


```

SMTP Server (SMTP_SERVER) =
SQL conditional compilation flags (SQL_CCFLAGS) =
Section actuals setting (SECTION_ACTUALS) = NONE
Connect procedure (CONNECT_PROC) =
Adjust temporal SYSTEM_TIME period (SYSTIME_PERIOD_ADJ) = NO
Log DDL Statements (LOG_DDL_STMTS) = NO
Log Application Information (LOG_APPL_INFO) = NO
Default data capture on new Schemas (DFT_SCHEMAS_DCC) = NO
Strict I/O for EXTBL_LOCATION (EXTBL_STRICT_IO) = NO
Allowed paths for external tables (EXTBL_LOCATION) = /opt/db2inst1
Default table organization (DFT_TABLE_ORG) = ROW
Default string units (STRING_UNITS) = SYSTEM
National character string mapping (NCHAR_MAPPING) = CHAR_CU32
Database is in write suspend state = NO
Extended row size support (EXTENDED_ROW_SZ) = ENABLE
Encryption Library for Backup (ENCRLIB) =
Encryption Options for Backup (ENCROPTS) =

WLM Collection Interval (minutes) (WLM_COLLECT_INT) = 0
Target agent load per CPU core (WLM_AGENT_LOAD_TRGPT) = AUTOMATIC(21)
WLM admission control enabled (WLM_ADMISSION_CTRL) = NO
Allocated share of CPU resources (WLM_CPU_SHARES) = 1000
CPU share behavior (hard/soft) (WLM_CPU_SHARE_MODE) = HARD
Maximum allowable CPU utilization (%) (WLM_CPU_LIMIT) = 0
Activity Sort Memory Limit (ACT_SORTMEM_LIMIT) = NONE
Control file recovery path (CTRL_FILE_RECOV_PATH) =
Encrypted database = NO
Procedural language stack trace (PL_STACK_TRACE) = NONE
HADR SSL certificate label (HADR_SSL_LABEL) =
HADR SSL Hostname Validation (HADR_SSL_HOST_VAL) = OFF

BUFFPAGE size to be used by optimizer (OPT_BUFFPAGE) = 0
LOCKLIST size to be used by optimizer (OPT_LOCKLIST) = 0
MAXLOCKS size to be used by optimizer (OPT_MAXLOCKS) = 0
SORTHEAP size to be used by optimizer (OPT_SORTHEAP) = 0

```

Database Manager Configuration

Node type = Enterprise Server Edition with local and remote clients

```

Database manager configuration release level = 0x1500

CPU speed (millisec/instruction) (CPUSPEED) = 9.840506e-08
Communications bandwidth (MB/sec) (COMM_BANDWIDTH) = 1.000000e+02

Max number of concurrently active databases (NUMDB) = 1
Federated Database System Support (FEDERATED) = NO
Transaction processor monitor name (TP_MON_NAME) =

Default charge-back account (DFT_ACCOUNT_STR) =

Java Development Kit installation path (JDK_PATH) =
/opt/db2inst1/sqllib/java/jdk64

Diagnostic error capture level (DIAGLEVEL) = 3
Notify Level (NOTIFYLEVEL) = 3
Diagnostic data directory path (DIAGPATH) =
/opt/db2inst1/sqllib/db2dump/$m
Current member resolved DIAGPATH =
/opt/db2inst1/sqllib/db2dump/DIAG0000/
Alternate diagnostic data directory path (ALT_DIAGPATH) =
Current member resolved ALT_DIAGPATH =
Size of rotating db2diag & notify logs (MB) (DIAGSIZE) = 0

Default database monitor switches
Buffer pool (DFT_MON_BUFPOOL) = OFF
Lock (DFT_MON_LOCK) = OFF
Sort (DFT_MON_SORT) = OFF
Statement (DFT_MON_STMT) = OFF
Table (DFT_MON_TABLE) = OFF
Timestamp (DFT_MON_TIMESTAMP) = OFF
Unit of work (DFT_MON_UOW) = OFF
Monitor health of instance and databases (HEALTH_MON) = OFF

SYSADM group name (SYSADM_GROUP) = DB2IADM1
SYSCTRL group name (SYSCTRL_GROUP) =
SYSMAINT group name (SYSMAINT_GROUP) =
SYSMON group name (SYSMON_GROUP) =

Client Userid-Password Plugin (CLNT_PW_PLUGIN) =
Client Kerberos Plugin (CLNT_KRB_PLUGIN) =
Group Plugin (GROUP_PLUGIN) =
GSS Plugin for Local Authorization (LOCAL_GSSPLUGIN) =
Server Plugin Mode (SRV_PLUGIN_MODE) = UNFENCED
Server List of GSS Plugins (SRVCON_GSSPLUGIN_LIST) =
Server Userid-Password Plugin (SRVCON_PW_PLUGIN) =
Server Connection Authentication (SRVCON_AUTH) = NOT_SPECIFIED
Cluster manager =

Database manager authentication (AUTHENTICATION) = SERVER
Alternate authentication (ALTERNATE_AUTH_ENC) = NOT_SPECIFIED
Cataloging allowed without authority (CATALOG_NOAUTH) = NO

```

```

Trust all clients (TRUST_ALLCLNTS) = YES
Trusted client authentication (TRUST_CLNTAUTH) = CLIENT
Bypass federated authentication (FED_NOAUTH) = NO

Default database path (DFTDBPATH) =
/opt/db2inst1

Database monitor heap size (4KB) (MON_HEAP_SZ) =
AUTOMATIC(256)
Java Virtual Machine heap size (4KB) (JAVA_HEAP_SZ) = 65536
Audit buffer size (4KB) (AUDIT_BUF_SZ) = 0
Global instance memory (% or 4KB) (INSTANCE_MEMORY) =
AUTOMATIC(376365531)
Member instance memory (% or 4KB) = GLOBAL
Agent stack size (AGENT_STACK_SZ) = 1024
Sort heap threshold (4KB) (SHEAPTHRES) = 0

Directory cache support (DIR_CACHE) = YES

Application support layer heap size (4KB) (ASLHEAPSZ) = 15
Max requester I/O block size (bytes) (RQIOBLK) = 65535
Workload impact by throttled utilities (UTIL_IMPACT_LIM) = 10

Priority of agents (AGENTPRI) = SYSTEM
Agent pool size (NUM_POOLAGENTS) =
AUTOMATIC(10000)
Initial number of agents in pool (NUM_INITAGENTS) = 1500
Max number of coordinating agents (MAX_COORDAGENTS) =
AUTOMATIC(10000)
Max number of client connections (MAX_CONNECTIONS) =
AUTOMATIC(MAX_COORDAGENTS)

Keep fenced process (KEEPFENCED) = YES
Number of pooled fenced processes (FENCED_POOL) =
AUTOMATIC(MAX_COORDAGENTS)
Initial number of fenced processes (NUM_INITFENCED) = 0

Index re-creation time and redo index build (INDEXREC) = RESTART

Transaction manager database name (TM_DATABASE) = 1ST_CONN
Transaction resync interval (sec) (RESYNC_INTERVAL) = 180

SPM name (SPM_NAME) = localhos
SPM log size (SPM_LOG_FILE_SZ) = 256
SPM resync agent limit (SPM_MAX_RESYNC) = 20
SPM log path (SPM_LOG_PATH) =

TCP/IP Service name (SVCCNAME) =
db2c_db2inst1
Discovery mode (DISCOVER) = SEARCH
Discover server instance (DISCOVER_INST) = ENABLE

SSL server keydb file (SSL_SVR_KEYDB) =
SSL server stash file (SSL_SVR_STASH) =
SSL server certificate label (SSL_SVR_LABEL) =
SSL service name (SSL_SVCCNAME) =
SSL cipher specs (SSL_CIPHERSPECS) =
SSL versions (SSL_VERSIONS) =
SSL client keydb file (SSL_CLNT_KEYDB) =
SSL client stash file (SSL_CLNT_STASH) =

Maximum query degree of parallelism (MAX_QUERYDEGREE) = ANY
Enable intra-partition parallelism (INTRAPARTALLEL) = NO

Maximum Asynchronous TQs per query (FEDERATED_ASYNC) = 0

Number of FCM buffers (FCM_NUM_BUFFERS) =
AUTOMATIC(4096)
FCM buffer size (FCM_BUFFER_SIZE) = 32768
Number of FCM channels (FCM_NUM_CHANNELS) =
AUTOMATIC(2048)
FCM parallelism (FCM_PARALLELISM) =
AUTOMATIC(36)
Node connection elapse time (sec) (CONN_ELAPSE) = 10
Max number of node connection retries (MAX_CONNRETRIES) = 5
Max time difference between nodes (min) (MAX_TIME_DIFF) = 60

db2start/db2stop timeout (min) (START_STOP_TIME) = 10

WLM dispatcher enabled (WLM_DISPATCHER) = NO
WLM dispatcher concurrency (WLM_DISP_CONCUR) = COMPUTED
WLM dispatcher CPU shares enabled (WLM_DISP_CPU_SHARES) = NO
WLM dispatcher min. utilization (%) (WLM_DISP_MIN_UTIL) = 5

Communication buffer exit library list (COMM_EXIT_LIST) =
Current effective arch level (CUR_EFF_ARCH_LVL) = V:11 R:5
M:8 F:0 I:0 SB:0
Current effective code level (CUR_EFF_CODE_LVL) = V:11 R:5
M:8 F:0 I:0 SB:0

Keystore type (KEYSTORE_TYPE) = NONE
Keystore location (KEYSTORE_LOCATION) =

Path to python runtime (PYTHON_PATH) =
Path to R runtime (R_PATH) =

Multioart upload part size (MULTIPARTSIZEMB) = 100

```

limit.conf

```
# /etc/security/limits.conf
#
#<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

db2inst1      soft   nofile          1000000
db2inst1      hard   nofile          1000000
db2inst1      soft   nproc           unlimited
db2inst1      hard   nproc           unlimited

# End of file
```

server.xml

```
<?xml version='1.0' encoding='utf-8'?>
<Server port="808#" shutdown="SHUTDOWN">
  <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 -->
  <!--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" />

  <GlobalNamingResources>
    <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>

  <Service name="Catalina#">

    <Connector port="808#"

    acceptCount="150000"

    maxConnections="60000"

    connectionTimeout="2000000"                                maxThreads="60000"

    maxKeepAliveRequests="-1" keepAliveTimeout="-1"

    protocol="org.apache.coyote.http11.Http11NioProtocol" redirectPort="844#"
    />

    <Connector port="8014" protocol="AJP/1.3" redirectPort="844#" />

    <Engine name="Catalina#" defaultHost="localhost">

      <Realm className="org.apache.catalina.realm.LockOutRealm">
        <Realm className="org.apache.catalina.realm.UserDatabaseRealm"
          resourceName="UserDatabase"/>
      </Realm>

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


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

Sysctl fe.xml

```
# 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).
## 256 * <size of RAM in GB>
kernel.shmni=32768
## size of RAM in bytes
kernel.shmmax=131737212000
## 2 * <size of RAM in the default system page size> = 2 *
Memroy(GB) * ( 1024*1024*1024)/PAGE_SIZE
kernel.shmall=67108864
#kernel.sem=<SEMMSL> <SEMMNS> <SEMOPM> <SEMMNI>
# kernel.sem (SEMMNI) 256 * <size of RAM in GB>
# kernel.sem (SEMMSL) 250
# kernel.sem (SEMMNS) 256000
# kernel.sem (SEMOPM) 32
kernel.sem=250 256000 32 32768
# kernel.msgmni (MSGMNI) 1 024 * <size of RAM in GB>
# kernel.msgmax (MSGMAX) 65 536
# kernel.msgmnb (MSGMNB) 65 536 3
kernel.msgmni=131072
kernel.msgmax=65536
kernel.msgmnb=65536
```

Appendix C: Price Quotations

DB Server



엘텍코리아

견 적 서
(QUOTATION)

수 신 : TTA

참 조 : 황세진 선임 연구원님.

견 적 번 호 : LTK-230816-01

견 적 일 자 : 2023. 08. 16.

견 적 내 용 : LKG2312 납품의견.

견적 유효기간 : 발행 후 4달

납 품 일 자 : 발주시 협의사항

지 물 조 건 : 현금

제품 보증기간 : 납품 후 3년 보증 3-year 24x7 w/4-hour maintenance

상 호 : 엘텍코리아㈜

사업자 번호: 119 - 86 - 44723

대 표 이 사 : 김 세 호

주 소 : 서울특별시 금천구 빛꽃로 234, 1904호
(에이스하이엔드타워 6차)

연 락 처 : Tel. 1599-7230 | Fax. 02-515-2263

담 당 자 : 박 장 순 대리

연 락 처 : 010- 6269- 2745

E-mail : jspark@ltechkorea.com

일금 : 삼천삼백일십구만팔천 원정 (₩33,198,000) VAT 포함

※ 견적 요청에 감사 드리며, 귀사의 일익번창을 기원합니다. [단위 : 원]

구 분	품 목	상 세 사 양	수 량	공 급 가		비 고
		Specification		Q'ty	공급단가	
x86	2U	LKG2312	1	30,180,000	30,180,000	
1	CPU	인텔® 제온® 골드 6354 프로세서 18C/36Th 3Ghz (Up To 3.6Ghz)	2			
2	RAM	64GB DDR4 Memory	24			
3	DISK1	SSD 960GB, Enterprise	2			
4	DISK2	SSD 3.84TB, Enterprise	2			
5	DISK3	SSD 3.2TB NVMe SSD	1			
6	DISK4	SSD 3.84TB, Enterprise NVMe, U2, Gen4	7			
7	RAID	SAS3108(2GB) 12Gbps Controller	1			
8	NIC	RJ45 1G 4port	1			
9	HBA	16Gb/s 2port (GBIC 포함)	1			
10	CHASSIS	LKG-2312 Server(power 1600watt) 3.5/2.5 inch disk bay * 12EA	1			
<p>※ 비 고</p> <p>1) 위 견적은 발주시점의 파트 수급상황에 따라 딜리버리 및 가격변동 있을수 있음.</p> <p>2) OS미포함.</p> <p>3) 서울, 경기지역 등 수도권 외 설치 시 별도 요금 추가.</p>				공급가 합계		30,180,000
				V.A.T		3,018,000
				총 합계		33,198,000

※ 견적서를 검토하신 후 제품 발주를 원하시는 경우 하단의 내용을 작성하여 명판/직인 후 스캔본을 회신 주시면 발주서의 효력을 갖습니다.
또는, 귀사 양식의 발주서를 팩스나 메일로 보내주시기 바랍니다. Fax. 02-515-2263

- 1) 발 주 담 당 :
- 2) 연 락 번 호 :
- 3) 배 송 일 자 :
- 4) 배 송 장 소 :

WAS Server

견 적 서

TEL : 010-5110-5598
 국제공인 시험인증(TPC-C) 서버

TTA 貴中

참 조 : 서 병 준 책임님

귀사의 성공적인 사업을 전심으로 지원하며 아래와 같이 견적드립니다.
 In compliance with your inquiry, we hereby submit our quotation as mentioned here under.

견적 금액 합계 : ₩18,370,000 (부가가치세 포함)

납품 예정 일자 : 발주 후 4~6주 이내 (별도 협의)

견적 유효 기간 : 견적일로부터 1개월

지 불 조 건 : 납품 기준 익월말 현금

견 적 담 당 : 김 상현 부장 (Mobile : 010-3720-5644)


견 적 일 자 : 2023년 8월 17일



주식회사 케이티엔에프
 서울시 강서구 마곡중앙 8로 3길 21
 (마곡동, KTNF빌딩)
 TEL : 02-865-5200 / FAX : 02-855-8814



등록번호	106-86-07697
대표이사	이 증 연
업 태	제 조
종 목	컴퓨터 주변기기

항목 Item	품 명 Description	규 격 Specification	수량 Q'ty	단 가 Unit Price	합 계 Amount	비 고 Remark
1	KR580S1 CPU Memory HDD Maintenance	2U, 8Bay, 1G 2port copper, 2x 800W PSU, Rack rail - Intel Xeon Scalable Gold 6140 (2.30GHz, 18core) - 64GB DDR4 ECC RDIMM Memory - 600GB SAS 12Gb/s 10K RPM (128MB) - 무상 3년	1 2 2 2	16,700,000	16,700,000	

Other Comment/Remarks

- 상기 단가는 부가세 별도입니다.
- 자세한 사항은 전화주시기 바랍니다.

합 계 : ₩16,700,000
 세 액 : ₩1,670,000
합 계 : ₩18,370,000

- 상기 제품은 대외무역법 제19조 제1항에 따라 전락물자에 해당되며, 본 물품의 해외수출시 대외무역법에 따라 전락물자 기술 수출인 통합고시에서 규정하는 허가기관의 장으로부터 수출허가를 득하시기 바랍니다.
 - 본 제품을 제3차에서 앞도 또는 재판매할 경우 해당 제3차에게 상기에 언급한 의무사항들을 사전에 충분히 고지하시기 바랍니다.
- 상기 외 OS 및 SW 별도입니다.
 - 3년 무상 품질 보증 합니다.



㈜락플레이스
03129 서울시 용포구 용포 33길 15 (연지동 연강빌딩 5층) Tel : 02-6251-7788 Fax : 02-6499-1478
rockPLACE, Inc.
15, Jong-ro 33-gil, Jongno-gu, Seoul, Korea 03129 Tel : 82-2-6251-7788 Fax : 82-2-6499-1478



견 적 서

REF No. : 2023RP108-1350	TERMS AND CONDITION
DATE : 2023. 08. 17.	납 기 : 발주후 4주이내
COMPANY : TAA(한국정보통신기술협회)	유지보수 :
ATTN : 황 세 진 선임연구원 귀하 TEL : 010-5110-4883	결제조건 : 납품 검수 후 30일
Email : hsejin314@tta.or.kr	유효기간 : 견적일로부터 1개월
FROM : ㈜락플레이스 허 운 범 차장 TEL : 010-6605-2146	

下記와 같이見積합니다.

(주) 락플레이스
대표이사 김 재 준, 김 연 수



Part No.	Description	수량	소비자가	공급단가	공급합계
(VAT 별도, 단위 : 원)					
Subscription	Red Hat Enterprise Linux Server				
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	2	4,089,000	2,400,000	4,800,000
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	2	8,000,000	2,000,000	4,000,000
합 계 금액					8,800,000

Part No.	Description	수량	소비자가	공급단가	공급합계
(VAT 별도, 단위 : 원)					
Subscription	Red Hat JBoss Web Server				
MW00123F3	Red Hat JBoss Web Server, 4-Core Standard 3Year - 전화/웹 지원 : 월-금, 9 a.m. - 5 p.m. 4시간내 응답 - unlimited incidents,	9	2,138,000	1,260,000	11,340,000
RP-CPS(WAS)	rockPLACE Support Carepack - JBoss Standard (3년) per 4Core 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	9	12,000,000	6,000,000	54,000,000
소 계 금액					65,340,000
합 계					74,140,000
부가세					7,414,000
합 계(부가세포함)					81,554,000

Remarks
1. Red Hat 제품은 연간 Subscription 제품이며, 기간이 만료될 시 Renewal을 하여야합니다.
2. 발주 시에는 반드시 고객정보(엔드유저명, 담당자, 연락처, Email)가 있어야 합니다.
3. OnSite 방문지원이 필요하실 경우에는 케어팩을 구매하셔야 합니다.

구분	Description	가격정책
1. OS 백업	OS 백업백구 통합관리 솔루션_cider_v2.0 (CiderV2) 유닉스,리눅스,윈도우 OS백업복구 통합관리 툴 Web기반의 간편한 UI를 통한 OS백업, 복구, 복제, 조희 기능 제공 베어메탈(Bare Metal) 자동 복구 AIX,HP-UX,Solaris ,Linux,Windows 5개 플랫폼을OS백업 통합관리 지원하는 유일한 솔루션	perpetual subscription
2. OSS 개발 관리툴	Toad Data Point Base edition Toad Data Point Professional edition 전세계 시장 점유율 1위 Multi DB (MySQL, Oracle, MSSQL, DB2, MonggoDB, 등) 개발툴 간편한 데이터베이스 오브젝트 관리 데이터 프로파일링을 통해서 데이터 품질 향상 (Professional Edition 이상에서 지원)	perpetual
	Toad Edge MySQL, PostgreSQL, MariaDB, EDB 구매한 라이선스에 맞게 전용 제품으로 활성화 (각각 별도 라이선스) 효율적인 SQL 작성(SQL Worksheet) 스키마 비교와 동기화 스크립트 생성 JSON 데이터 탐색 및 편집 / 데이터베이스 Snapshot 기능	subscription

Network Switch

23. 8. 16. 오후 5:27

상품상세정보

상품상세정보

네트워크스위치



[확대보기](#)

수량: 대

업체명: **주식회사에스비정보기술** [중소기업]

[계약자/공급자 정보조회](#)

계약방법: 다수공급자계약

규격명: 네트워크스위치, 다산네트웍스, (CN)D2224GP, 24port POE (공급)

price → 가격: 1,900,000 원 [다량납품할인을 확인](#)

단위: 대

원산지: 중국

주요부품1[원산지]: CPU[중국]

주요부품2[원산지]: CASE[중국]

핵심부품[원산지]: 중국[중국]

제조사: (주)다산네트웍스

납품장소: 수요기관 지정장소

인도조건: 현장설치도

공급지역: 전지역

부가세여부: 부가가치세포함

계약기간: 2022/11/03 ~ 2023/11/02

납품기한: 60일 (납품요구일로부터)

조달수수료여부: 조달수수료 별도 [조달수수료 안내·계산](#)

첨부파일: 2022/10/01_00226227500-(계약예규)물품구매(제조)계약일반

조건(기획재정부계약예규제583호20211201).hwp

2022/10/02_00226227500-물품다수공급자계약특수조건

(2021.7.1.시행).hwp

2022/10/03_00226227500-물품구매계약품질관리특수조건

(220501).hwp

2022/10/04_00226227500-규격서.zip

대분류: 09 - 전자/정보/통신/영상

중분류: 07 - 음향장비 및 신호장치

물품분류번호: 43222612

세부품명번호: 4322261201

물품식별번호: 24567061

계약번호: 00226227500-13


징수구분: 후징수

[TOP](#)

감추기

구매의사결정을 위한 구매정보

우선(의무)구매대상	해당 없음
평균배송일/납품기한	관련정보없음 / 60일 (납품요구일로부터)
품질보증조달물품여부	해당없음
본사소재지	경기도 의왕시 상고개로53, 10층 A동 1016호(포일동, 에이스청계타워)

견 적 서 (QUOTATION)	 D.BtoE <small>Data begin to end</small>
----------------------	--

견 적 일 자	2023년 9월 21일 목요일	
수 신	한국정보통신기술협회	귀 중
참 조	황세진 선임연구원님 (010-5110-4883 / hsejin314@tta.or.kr 귀 하	
유 효 기 간	견적일로부터 4개월	
결 제 조 건	협의	
건 명	TPC-C Performance & Quality Authentication 件	

견 적 번 호	D.BtoE-LIC-2023-0921-01		
상 호	주식회사 디비투이	대표이사	김 명 훈 
사업장주소	서울시 서초구 방배중앙로 21길42 경복빌딩 2층		
담당자	이사 최정환	이메일	ihchoi@db2e.co.kr
Mobile	010-3182-7003	Office	Tel)070-4464-0100, Fax)02-595-0105

귀사에 아래와 같이 견적합니다.

합계금액	일금일십삼억사천육백팔십구만육천칠백육십원	₩1,346,896,760	(부가세포함)
------	-----------------------	----------------	---------

NO	상세내역	수량	단위	소비자단가	금액	공급가	비고
1	IBM Db2 Advanced Edition Cartridge for IBM Cloud Pak for Data Virtual Processor Core License + SW Subscription & Support 12 Months (Ver.11.5.8.0)	36	VPC	128,307,000	4,619,052,000	785,200,000	
IBM Db2 Product 소계 (부가세별도)							
2	IBM Db2 License Renewal (2 Years)	2	ea	189,625,800	379,251,600	379,251,600	
IBM Db2 Product 소계 (부가세별도)							
3	IBM Db2 Implementaion & Support (3 Years)	3	ea	20,000,000	60,000,000	60,000,000	
IBM Db2 Support 소계 (부가세별도)							

IBM Db2 Product + Support 견적금액 합계 (부가세별도)	1,224,451,600	
IBM Db2 Product + Support 견적금액 합계 (부가세포함)	1,346,896,760	

(참고사항)
 - For Technical Supports, it indicates 24 * 7 * 4 Hours of support