



Red Hat – Industry Standard Benchmarks

Leadership Virtualized TPC-C Benchmark™ Performance & Price/Performance using Red Hat® Enterprise Linux® / KVM (> 1.3 Million tpmC, < 0.51 \$ / tpmC)

TPC-C Benchmark

DB2 ESE 9.7 Database

Red Hat Enterprise Linux 6.4

IBM x3650 M4

Version 3.0

February 2013





**Leadership Virtualized
TPC-C Benchmark™ Performance
& Price/Performance using
Red Hat® Enterprise Linux®
(> 1.3 Million tpmC, < 0.51 \$ / tpmC)**

1801 Varsity Drive
Raleigh NC 27606-2072 USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701
PO Box 13588
Research Triangle Park NC 27709 USA

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
- Microsoft Windows 2008 server and COM+ are registered trademarks of Microsoft Corporation
- IBM System x, IBM, ServeRAID and DB2 are trademarks of International Business Machines Corporation in the United States and other countries
- Linux is a registered trademark of Linus Torvalds
- Red Hat, Red Hat Enterprise Linux and the Red Hat "Shadowman" logo are registered trademarks of Red Hat, Inc. in the United States and other countries
- Intel and Xeon are registered trademarks of Intel Corporation.

All other trademarks referenced herein are the property of their respective owners.

© 2013 by Red Hat, Inc. This material may be distributed only subject to the terms and conditions set forth in the Open Publication License, V1.0 or later (the latest version is presently available at <http://www.opencontent.org/openpub/>).

The information contained herein is subject to change without notice. Red Hat, Inc. and IBM Inc. shall not be liable for technical or editorial errors or omissions contained herein.

Distribution of this work or derivative of this work in any standard (paper) book form for commercial purposes is prohibited unless prior permission is obtained from Red Hat Inc. and IBM Inc.

The GPG fingerprint of the security@redhat.com key is:
CA 20 86 86 2B D6 9D FC 65 F6 EC C4 21 91 80 CD DB 42 A6 0E



Table of Contents

1. Executive Summary	5
1.1 Competitive Comparison of RHEL 6.2 and Windows Server 2008.....	6
2. Overview of the TPC Benchmark C: The Order-Entry Benchmark	9
2.1 The Benchmark Model	9
2.2 More OLTP Features and Complexity	11
2.3 A Measure of Business Throughput	11
2.4 A Yardstick	12
3. Hardware Configuration	14
4. Transaction Mix.....	15
5. Data Base Design	15
6. Performance Metrics and Response Time.....	17
6.1 Response Times	17
6.2 Think and Keying Times.....	17
6.3 Response Time Frequency Distribution	18
6.4 Performance Curve for Response Time versus Throughput	23
6.5 Think Time Frequency Distribution.....	24
6.6 Throughput versus Elapsed Time.....	25
6.7 Steady State Determination: Ramp-up & Run Times	25
7. References.....	26



BLANK PAGE



1. Executive Summary

The Transaction Processing Council's Online Transaction Processing (OLTP) system benchmark, TPC-C, simulates a complete environment where a population of terminal operators executes transactions against a database.

1. TPC-C is considered the most rigorous OLTP database benchmark in the industry
2. The latest TPC-C benchmark result released by IBM achieves > 1 million tpmC using Red Hat Enterprise Linux (RHEL) / Kernel-based Virtual Machine (KVM):
 - This is the 1st and only audited result ever published using virtualization technology
 - This is the highest TPC-C result ever published using virtualization technology
3. This outstanding result demonstrates to customers the full potential of running their entire IT infrastructure virtualized. The powerful combination of RHEL with KVM technology running on an x3650 M4 and DB2 delivers the best result in the industry.

This benchmark result is a demonstration of the close and continued cooperation between IBM Corp. and Red Hat Inc. to showcase the superior combined performance of Red Hat Enterprise Linux (RHEL) and DB2 running on IBM's Intel Xeon-based System x servers.

For more details refer to:

1. TPC-C Benchmark Description:
<http://www.tpc.org/tpcc/detail.asp>
2. TPC-C Result Highlights:
http://www.tpc.org/tpcc/results/tpcc_result_detail.asp?id=113022201
3. Benchmark Executive Summary by IBM:
http://www.tpc.org/results/individual_results/IBM/ibm-linux-db2-x3650m4-kvm_ES.pdf
4. Benchmark Full Disclosure Report by IBM:
http://www.tpc.org/results/fdr/tpcc/ibm-linux-db2-x3650m4-kvm_FDR.pdf



1.1 Comparison of Bare Metal and Virtualized TPC-C Results on Identical Hardware

As shown in Table 1.1, Figure 1-1 and Figure 1-2, this result delivers 88% of bare metal performance as evidenced by previously TPC-C result on similar platform and demonstrates an outstanding scalability of virtualized infrastructures built on top of RHEL and KVM running on industry-standard x86 hardware.

Date	HW Configuration	tpmC	\$/tpmC	OS	DB
Feb. 2013	IBM System x 3650 M4 - Intel Xeon E5-2690 2.90 GHz (2 chips / 16 cores)	1.32M	0.51	Red Hat Enterprise Linux 6.4 / KVM Hypervisor	DB2 ESE 9.7
Apr. 2012	IBM Flex System x240 - Intel Xeon Processor E5-2690 2.90GHz (2 chips / 16 cores)	1.5M	0.53	Red Hat Enterprise Linux 6.2	DB2 ESE 9.7

Table 1-1



Red Hat Enterprise Linux / KVM produces best virtualized TPC-C performance (as of February 2013)

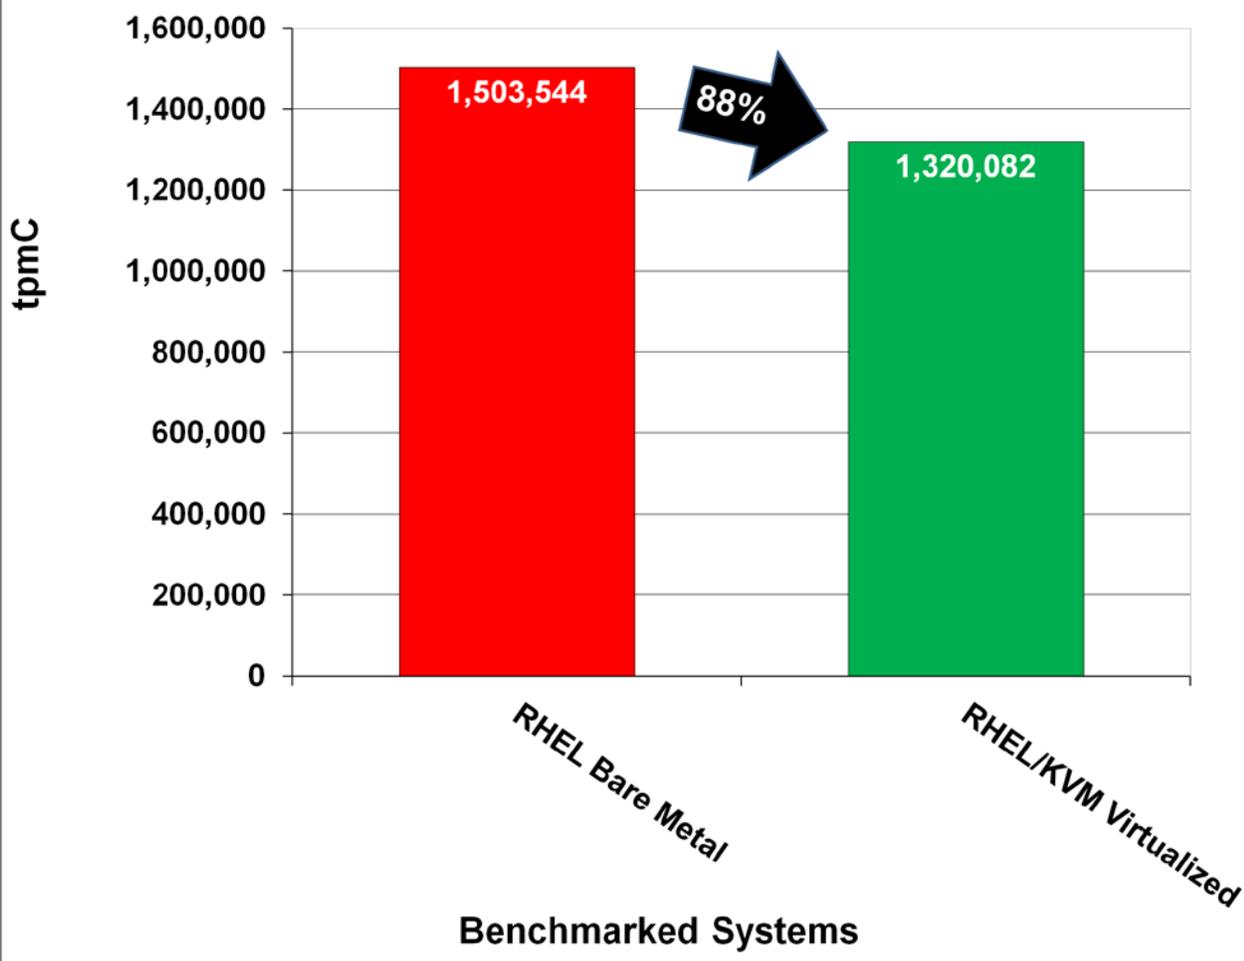


Figure 1-1



Red Hat Enterprise Linux / KVM produces best virtualized price/performance (as of February 2013)

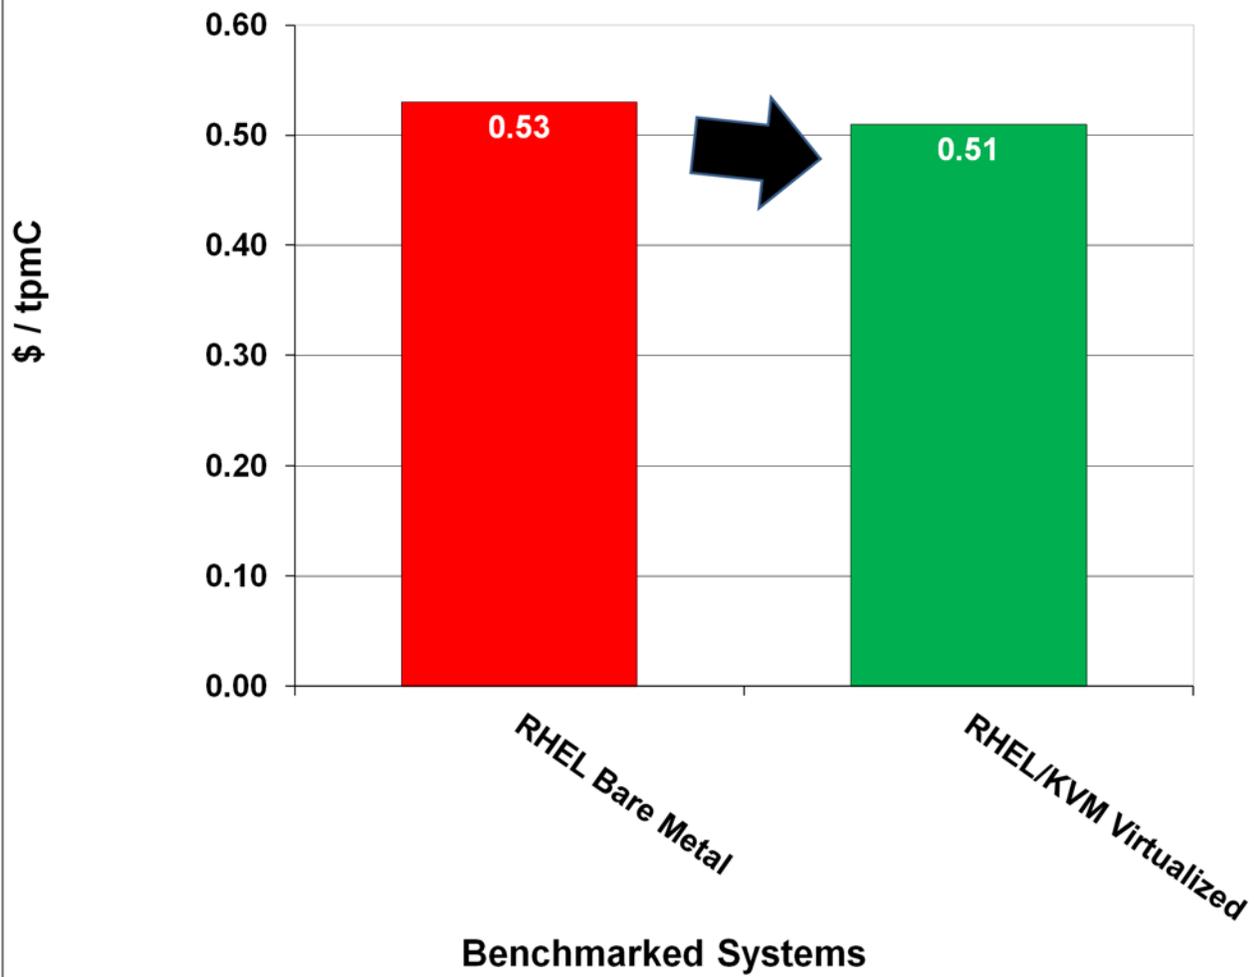


Figure 1-2



2. Overview of the TPC Benchmark C: The Order-Entry Benchmark

The goal of TPC benchmarks is to define a set of functional requirements that can be run on any transaction processing system, regardless of hardware or operating system. It is then up to the test sponsor to submit proof (in the form of a full disclosure report) that they have met all the requirements. This methodology allows any vendor, using "proprietary" or "open" systems, to implement the TPC benchmark and guarantees to end-users that they will see an apples-to-apples comparison. This is a dramatic departure from most other benchmarks where test sponsors are limited to comparing machines that run on just one operating system or benchmarks that execute the same set of software instructions.

TPC benchmarks also differ from other benchmarks in that TPC benchmarks are modeled after actual production applications and environments rather than stand-alone computer tests which may not evaluate key performance factors like user interface, communications, disk I/Os, data storage, and backup and recovery. The difficulty in designing TPC benchmarks lies in reducing the diversity of operations found in a production application, while retaining its essential performance characteristics, namely, the level of system utilization and the complexity of its operations. A large number of functions have to be performed to manage a production system. Since many of these functions are proportionally small in terms of system resource utilization or in terms of frequency of execution, they are not of primary interest for performance analysis. Although these functions are vital for a production system, within the context of a standard benchmark, they would merely create excessive diversity and expense and are, therefore, omitted.

2.1 The Benchmark Model

As an OLTP system benchmark, TPC-C simulates a complete environment where a population of terminal operators 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. However, it should be stressed that it is not the intent of TPC-C to specify how to best implement an Order-Entry system. 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.

In the TPC-C business model, a wholesale parts supplier (called the Company below) operates out of a number of warehouses and their associated sales districts. The TPC benchmark is designed to scale just as the Company expands and new warehouses are created. However, certain consistent requirements must be maintained as the benchmark is scaled. Each warehouse in the TPC-C model must supply ten sales districts, and each district serves three thousand customers. An operator from a sales district can select, at any time, one of the five operations or transactions offered by the Company's order-entry system. Like the transactions themselves, the frequency of the individual transactions are modeled after realistic scenarios.

The most frequent transaction consists of entering a new order which, on average, is comprised of ten different items. Each warehouse tries to maintain stock for the 100,000 items in the Company's catalog and fill orders from that stock. However, in reality, one warehouse will probably not have all the parts required to fill every order. Therefore, TPC-C requires that close to ten percent of all orders must be supplied by another warehouse of the Company. Another frequent transaction consists in recording a payment received from a customer. Less frequently, operators will request the status of a previously placed order, process a batch of ten orders for delivery, or query the system for potential supply shortages by examining the level of stock at the local warehouse. A total of five types of transactions, then, are used to model this business activity. The performance metric reported by TPC-C measures the number of orders that can be fully processed per minute and is expressed in tpm-C.

TPC-C was designed to carry over many of the characteristics of TPC-A, the TPC's standard version of DebitCredit. Therefore, TPC-C includes all the components of a basic OLTP benchmark. To make the benchmark applicable to systems of varying computing powers, TPC-C implementations must scale both the number of terminals and the size of the database proportionally to the computing power of the measured system. To test whether the measured system is a fully production-ready system with sufficient recovery capabilities, the database must provide what are defined as the ACID properties: atomicity, consistency, isolation, and durability. To facilitate independent verification of the benchmark results, the test sponsor must release, in a full disclosure report, all information necessary to reproduce the reported performance. This



performance, which measures the throughput of the system, must be reported along with the total system cost. The total system cost is a close approximation of the true cost of the vendor-supplied portion of the system to the end-user. It includes the cost of all hardware and software components; maintenance costs over 5 years; and sufficient storage capacity to hold the data generated over a period of 180 eight-hour days of operation at the reported throughput.

2.2 More OLTP Features and Complexity

TPC-C involves a mix of five concurrent transactions of different types and complexity either executed on-line or queued for deferred execution. This is one of the most substantial extensions of the basic OLTP benchmarking model as new components of the measured system are being stressed by having multiple transactions of different natures compete for system resources. Another substantial extension is the increased complexity of its database structure. The new database is comprised of nine types of records with a wide range of record and population sizes. As a result, there is greater diversity in the data manipulated by each of the five transactions. The data entered by operators in TPC-C include some of the basic characteristics of real-life data input. For example, operators may enter an invalid item number, forcing the transaction to be cancelled.

In moving toward modeling more realistic environments, TPC-C reduces the number of artificial limitations commonly found in other benchmarks. For example, to promote the use of fully-functional terminals or work-stations and screen management software, TPC-C requires all terminal inputs and displays to be usable by real-life operators. To that end, all screens must be formatted using labeled input and output fields, as specified, and must provide all the common screen manipulation features, including moving forward or backward through the input fields and entering numbers in right justified fields. In another area, any physical database design technique that can be used to improve the performance of a real-life application, such as partitioning or replication of data, is allowed in TPC-C. The use of database records by the transactions has been carefully defined to preclude test sponsors from gaining unrealistic advantages from any of these techniques.

2.3 A Measure of Business Throughput

The throughput of TPC-C is a direct result of the level of activity at the terminals. Each warehouse has ten terminals and all five transactions are available at each



terminal. A remote terminal emulator (RTE) is used to maintain the required mix of transactions over the performance measurement period. This mix represents the complete business processing of an order as it is entered, paid for, checked, and delivered. More specifically, the required mix is defined to produce an equal number of New-Order and Payment transactions and to produce one Delivery transaction, one Order-Status transaction, and one Stock-Level transaction for every ten New-Order transactions.

The tpm-C metric is the number of New-Order transactions executed per minute. Given the required mix and the wide range of complexity and types among the transactions, this metric more closely simulates a complete business activity, not just one or two transactions or computer operations. For this reason, the tpm-C metric is considered to be a measure of business throughput.

The RTE is also used to measure the response time of each transaction and to simulate keying times and think times. The keying time represents the time spent entering data at the terminal and the think time represents the time spent, by the operator, to read the result of the transaction at the terminal before requesting another transaction. Each transaction has a minimum keying time and a minimum think time. In addition, the response time of each transaction must be below a required threshold. These thresholds have been defined to give predominance to New-Order as the performance limiting transaction.

2.4 A Yardstick

Users of benchmark information and results, whether they be members of the press, market researchers, or commercial users, want to be assured that the benchmark results they see are valid measures of performance. To meet that demand, the TPC has designed its benchmarks to simulate and test systems with all the necessary production-oriented features, including backup and recovery features. In addition, the TPC requires complete documentation of the benchmark run (the full disclosure report). These reports are available to any user and pass through the TPC's own internal review process. All these requirements help to ensure that users of TPC benchmark results will see valid, objective measures of performance.

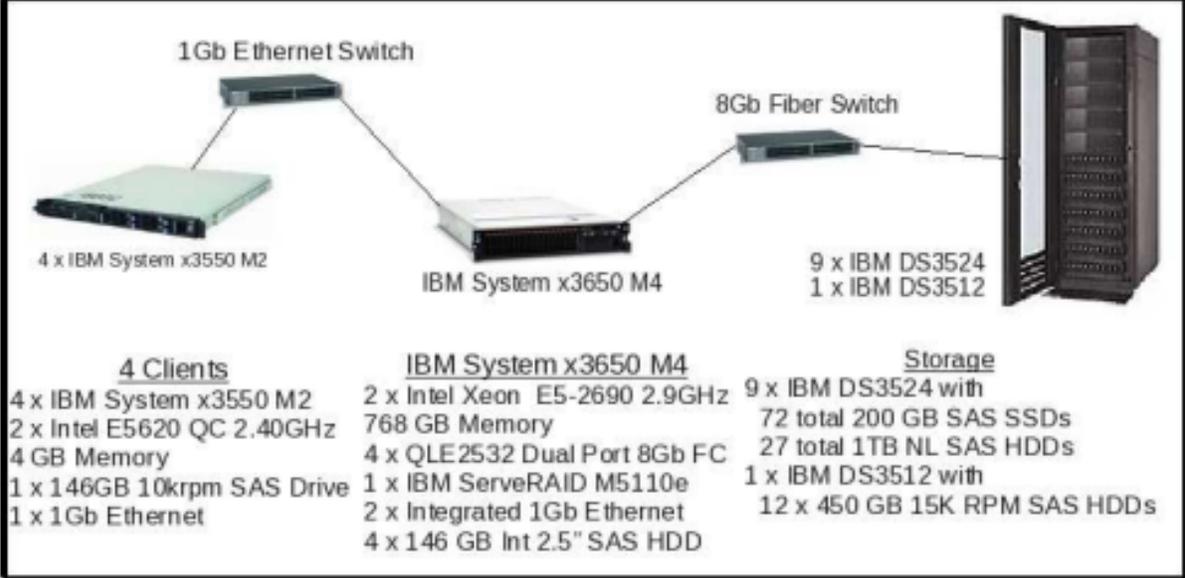
TPC-C follows the TPC's benchmarking philosophy and methodology in all the above respects, but it also includes new elements and more complex requirements. TPC-C's performance measurement metric, tpm-C, does not just measure a few basic computer or database transactions, but measures how



many complete business operations can be processed per minute. This new benchmark should give users a more extensive, more complex yardstick for measuring OLTP system performance.



3. Hardware Configuration

	IBM x3650 M4 with KVM DB2® 9.7		TPC-C Rev. 5.11				
			Report Date: February 25, 2013				
Total System Cost	TPC-C Throughput	Price/Performance		Availability Date			
\$667,882 USD	1,320,082	\$0.51 USD		February 25, 2013			
Database Processors/Cores/Threads	Database Manager	Operating System	Other Software	No. Users			
2/16/32 2.90 GHz Intel Xeon E5-2690	DB2 9.7	RHEL 6.4 with KVM	Microsoft Visual C++ Microsoft COM+	1,040,400			
							
<table style="width: 100%; border: none;"> <tr> <td style="width: 33%; border: none;"> <p style="text-align: center;"><u>4 Clients</u></p> 4 x IBM System x3550 M2 2 x Intel E5620 QC 2.40GHz 4 GB Memory 1 x 146GB 10krpm SAS Drive 1 x 1Gb Ethernet </td> <td style="width: 33%; border: none;"> <p style="text-align: center;"><u>IBM System x3650 M4</u></p> 2 x Intel Xeon E5-2690 2.9GHz 768 GB Memory 4 x QLE2532 Dual Port 8Gb FC 1 x IBM ServeRAID M5110e 2 x Integrated 1Gb Ethernet 4 x 146 GB Int 2.5" SAS HDD </td> <td style="width: 33%; border: none;"> <p style="text-align: center;"><u>Storage</u></p> 9 x IBM DS3524 with 72 total 200 GB SAS SSDs 27 total 1TB NL SAS HDDs 1 x IBM DS3512 with 12 x 450 GB 15K RPM SAS HDDs </td> </tr> </table>					<p style="text-align: center;"><u>4 Clients</u></p> 4 x IBM System x3550 M2 2 x Intel E5620 QC 2.40GHz 4 GB Memory 1 x 146GB 10krpm SAS Drive 1 x 1Gb Ethernet	<p style="text-align: center;"><u>IBM System x3650 M4</u></p> 2 x Intel Xeon E5-2690 2.9GHz 768 GB Memory 4 x QLE2532 Dual Port 8Gb FC 1 x IBM ServeRAID M5110e 2 x Integrated 1Gb Ethernet 4 x 146 GB Int 2.5" SAS HDD	<p style="text-align: center;"><u>Storage</u></p> 9 x IBM DS3524 with 72 total 200 GB SAS SSDs 27 total 1TB NL SAS HDDs 1 x IBM DS3512 with 12 x 450 GB 15K RPM SAS HDDs
<p style="text-align: center;"><u>4 Clients</u></p> 4 x IBM System x3550 M2 2 x Intel E5620 QC 2.40GHz 4 GB Memory 1 x 146GB 10krpm SAS Drive 1 x 1Gb Ethernet	<p style="text-align: center;"><u>IBM System x3650 M4</u></p> 2 x Intel Xeon E5-2690 2.9GHz 768 GB Memory 4 x QLE2532 Dual Port 8Gb FC 1 x IBM ServeRAID M5110e 2 x Integrated 1Gb Ethernet 4 x 146 GB Int 2.5" SAS HDD	<p style="text-align: center;"><u>Storage</u></p> 9 x IBM DS3524 with 72 total 200 GB SAS SSDs 27 total 1TB NL SAS HDDs 1 x IBM DS3512 with 12 x 450 GB 15K RPM SAS HDDs					
System Components	Each of the 4 Clients		Server				
	Quantity	Description	Quantity	Description			
Processors/Cores/ Threads	2/8/16	2.4GHz Intel Xeon E5620 4 x 256KB L2, 1 x 12MB L3	2/16/32	2.90 GHz Intel Xeon E5-2690 8 x 256KB L2, 1 x 20MB L3			
Memory	2	2 GB	24	32 GB			
Disk Controllers	1	SAS	4 1	QLogic 8Gb FC Dual-port HBA IBM ServeRAID M5110e			
Disk Drives	1	146 GB	27 12 72 4	1 TB 7200 RPM NL SAS 450 GB 15K RPM SAS 200GB SAS 2.5" SSD drives 146GB 15K RPM SAS HDD			
Total Storage		146 GB		40.33 TB			
Terminals	1	System Console	1	System Console			



4. Transaction Mix

Table 4-1 shows the mix of each of the transaction types executed by the System Under Test (SUT).

Transaction Mix	
New-Order	44.953%
Payment	43.018%
Order-Status	4.011%
Delivery	4.008%
Stock-Level	4.010%

Table 4-1: Numerical Quantities for Transaction and Terminal Profiles

5. Data Base Design

There were no restrictions on insert and/or delete operations to any of the tables. The space required for an additional five percent of the initial table cardinality was allocated to DB2 and priced as static space.

The insert and delete functions were verified by the auditor. In addition, the auditor verified that the primary key for each database table could be updated outside the range of its initial partition.

WAREHOUSE, DISTRICT, STOCK, CUSTOMER, HISTORY, ORDERS, ORDERLINE, and NEWORDER were horizontally partitioned into multiple tables.

For each partitioned table, a view was created over all table partitions to provide full transparency of data manipulation.

No tables were replicated.



Table 5-1 portrays the TPC Benchmark™ C defined tables and the number of rows for each table as they were built initially.

Table Name	Number of Rows
Warehouse	104,040
District	1,040,400
Customer	3,121,200,000
History	3,121,200,000
Orders	3,121,200,000
New Order	936,360,000
Order Line	31,211,013,719
Stock	10,404,000,000
Item	100,000

Table 5-1: Initial Cardinality of Tables

The database manager used for this testing was DB2 9.5. DB2 is a relational DBMS. DB2 remote stored procedures and embedded SQL statements were used. The DB2 stored procedures were invoked via SQL CALL statements. Both the client application and stored procedures were written in embedded C code.



6. Performance Metrics and Response Time

6.1 Response Times

Table 6-1 lists the response times and the ninetieth percentiles for each of the transaction types for the measured system.

<u>Response Times (in seconds)</u>	<u>90th %</u>	<u>Average</u>	<u>Maximum</u>
New Order	0.160	0.124	1.875
Payment	0.160	0.121	3.078
Order-Status	0.160	0.122	0.953
Delivery (interactive)	0.160	0.117	0.703
Delivery (deferred)	0.030	0.020	1.01
Stock-Level	0.160	0.123	1.110
Menu	0.141	0.117	1.046

Table 6-1: Response Times

6.2 Think and Keying Times

Table 6-2 lists the TPC-C keying and think times for the measured system.

<u>Keying/Think Times (in seconds)</u>	<u>Min.</u>	<u>Average</u>	<u>Max.</u>
New Order	18.000/0.00	18.000/12.035	18.062/120.343
Payment	3.000/0.00	3.000/12.037	3.062/120.344
Order-Status	2.000/0.00	2.000/10.033	2.047/100.329
Delivery	2.000/0.00	2.000/5.036	2.047/50.328
Stock-Level	2.000/0.00	2.000/5.035	2.047/50.328

Table 6-2: Think and Keying Times



6.3 Response Time Frequency Distribution

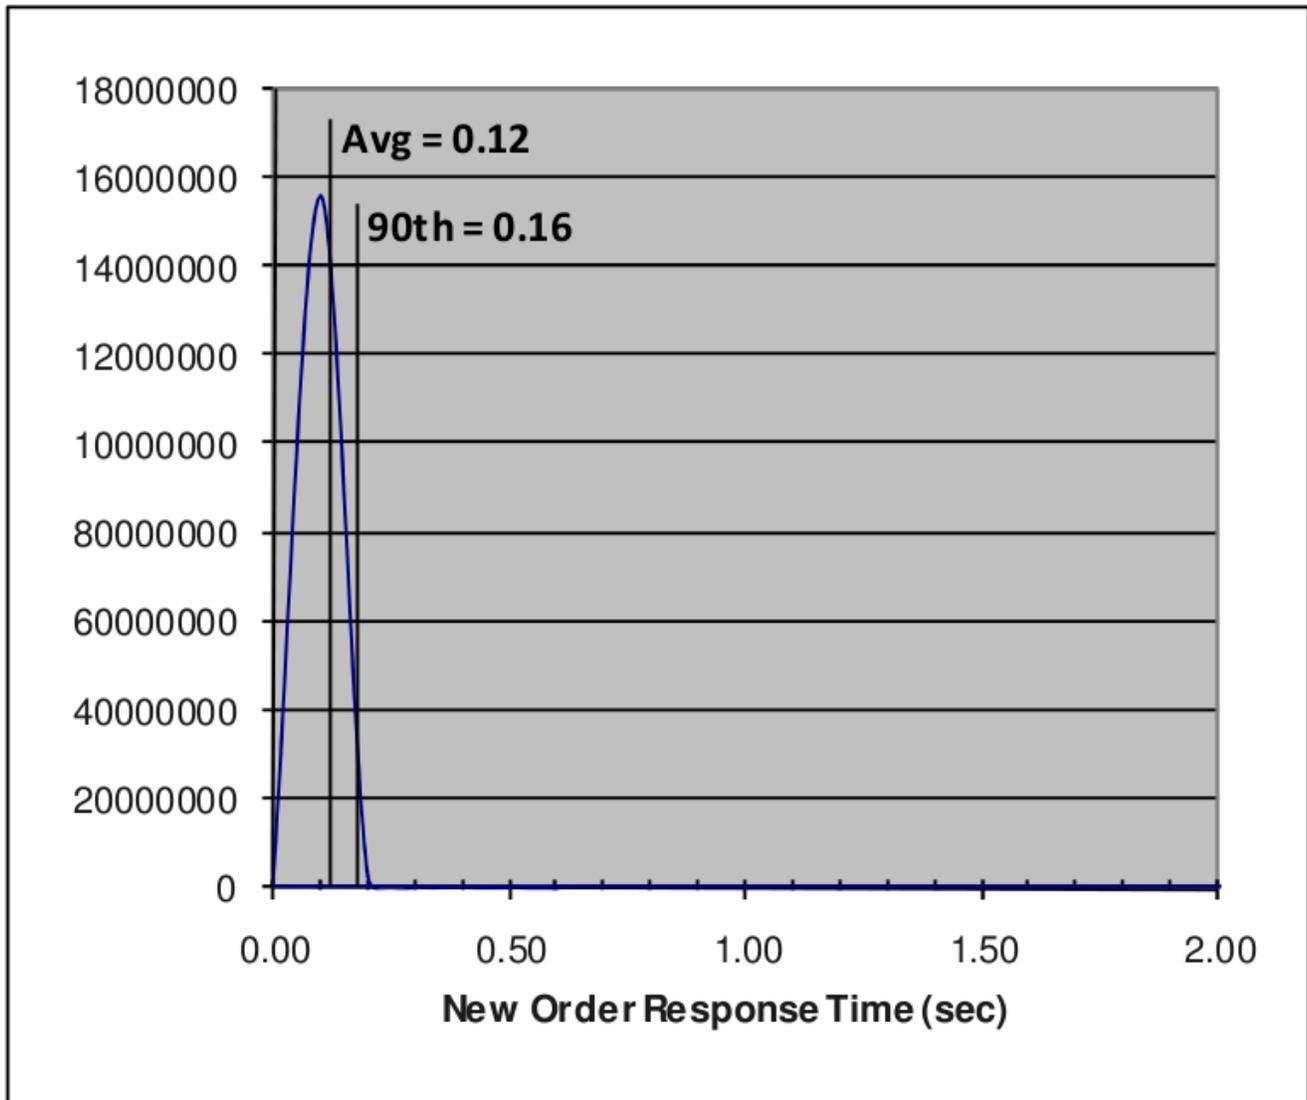


Figure 6-1: New-Order Response Time Distribution

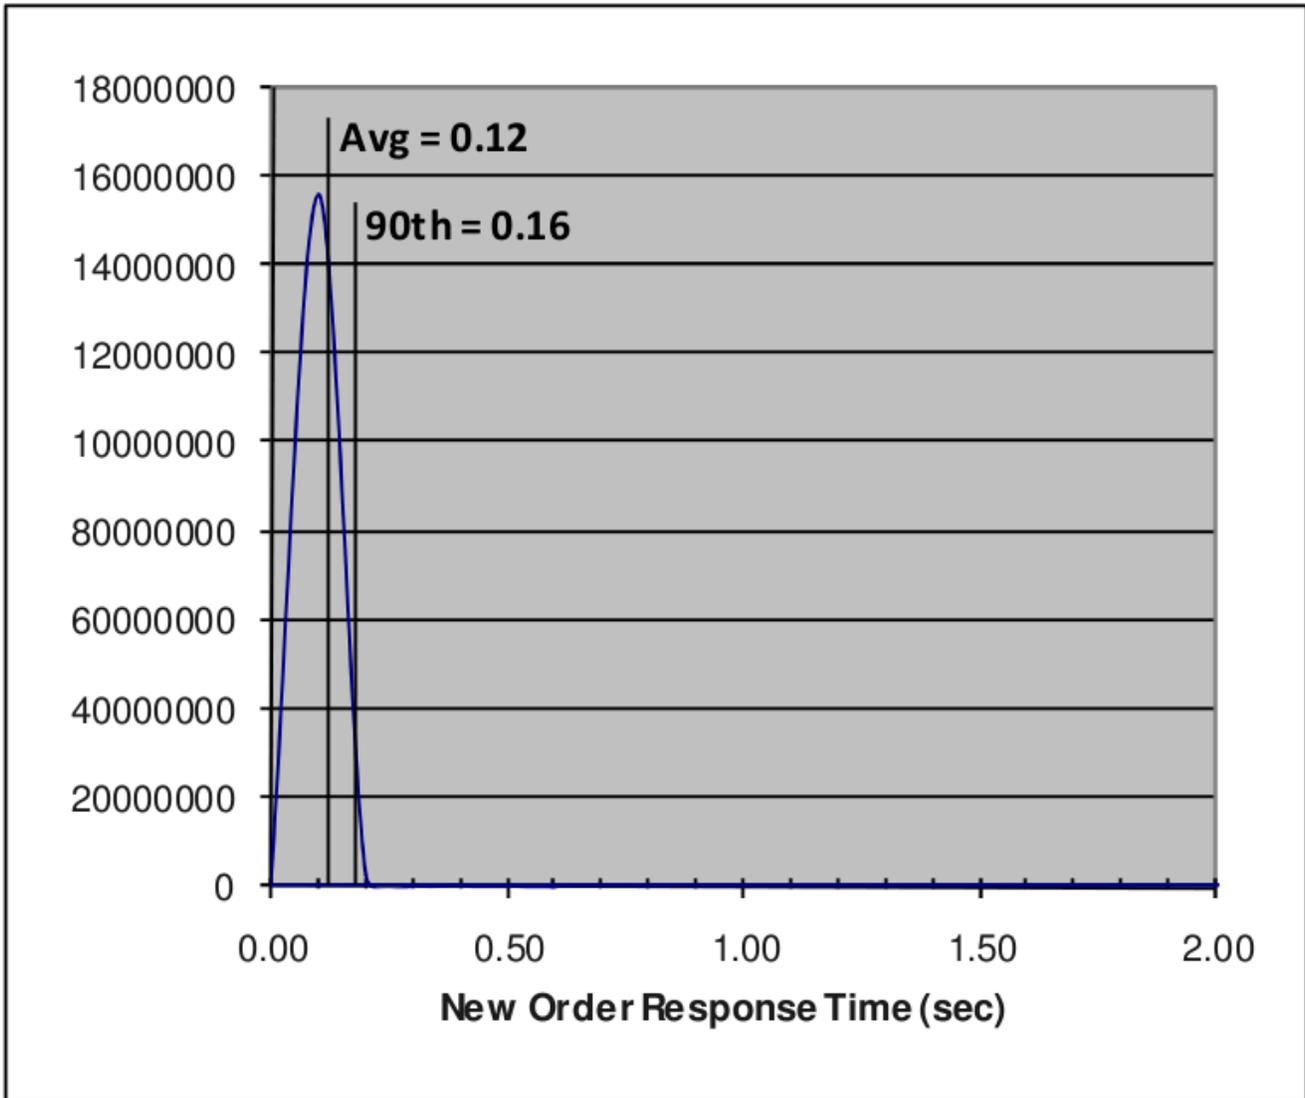


Figure 6-2: Payment Response Time Distribution

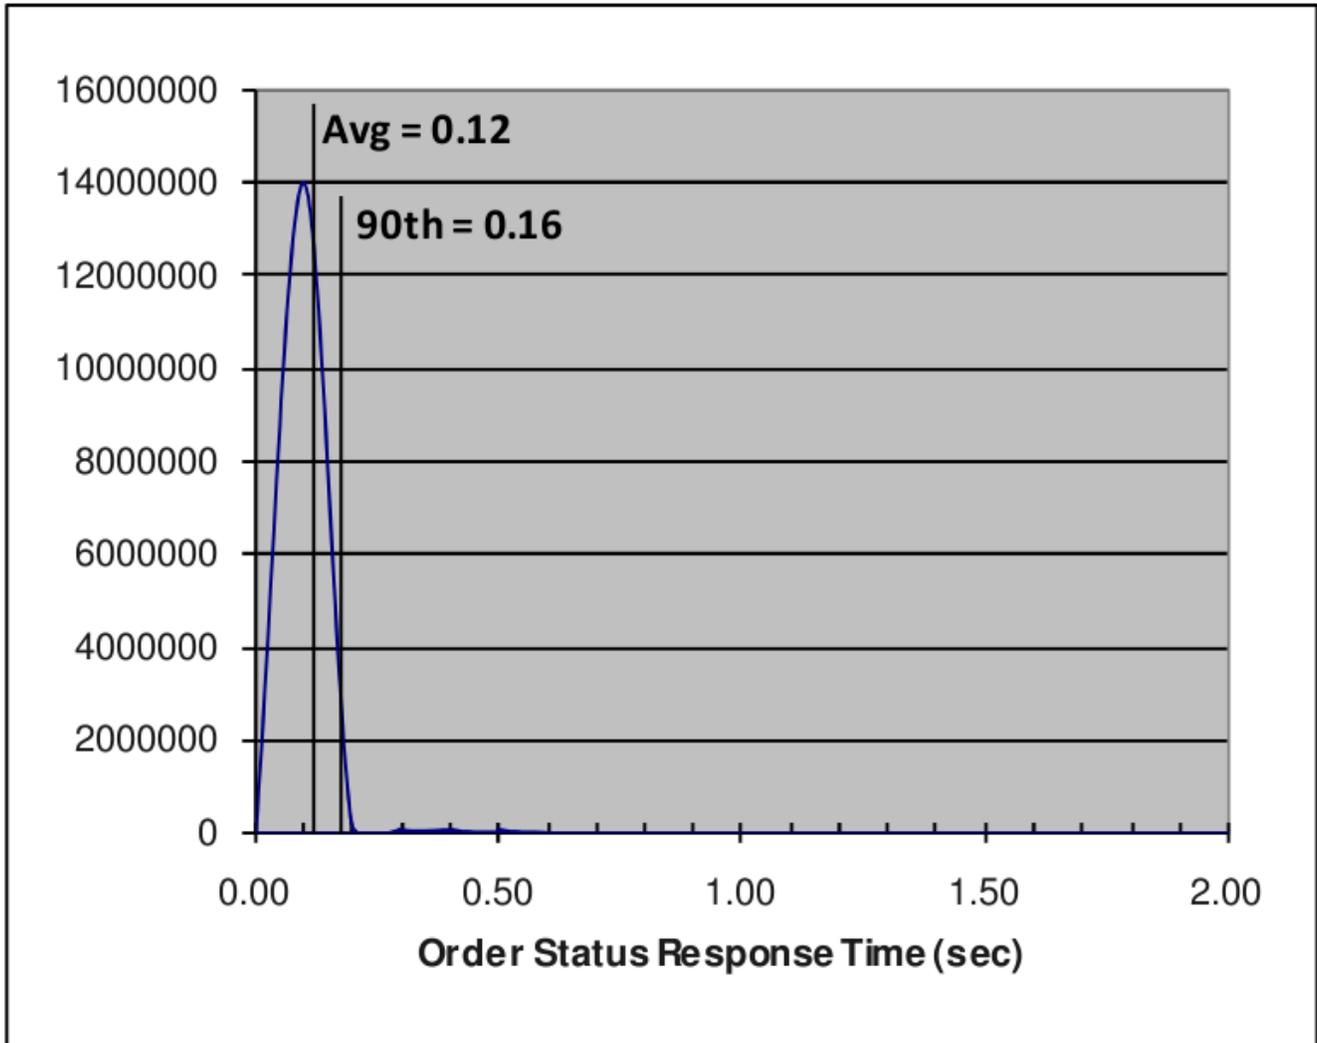


Figure 6-3: Order-Status Response Time Distribution

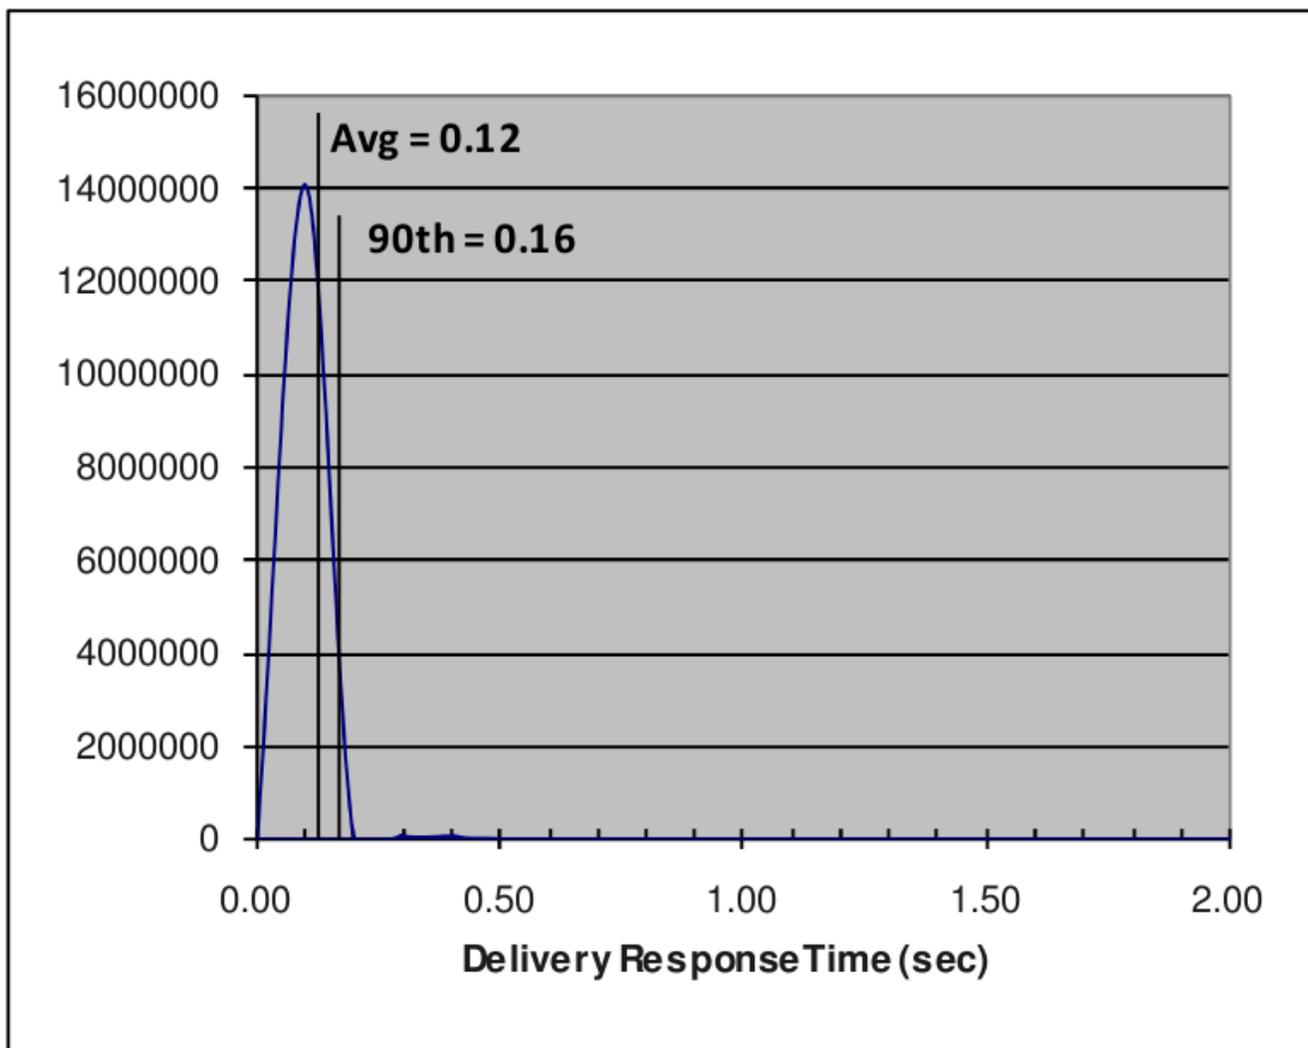


Figure 6-4: Delivery (Interactive) Response Time Distribution

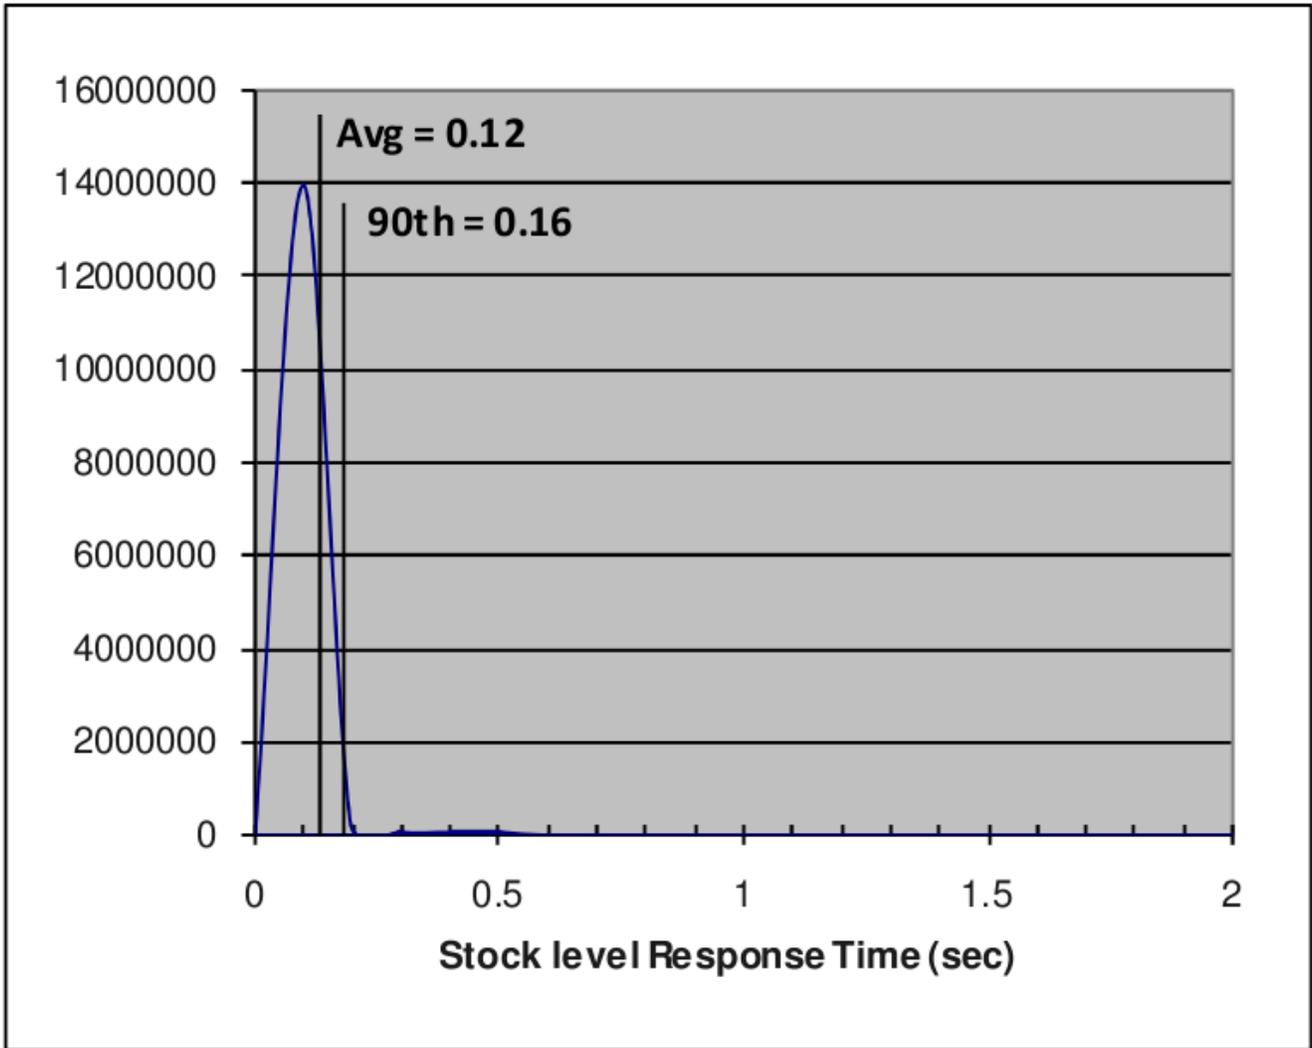


Figure 6-5: Stock Level Response Time Distribution



6.4 Performance Curve for Response Time versus Throughput

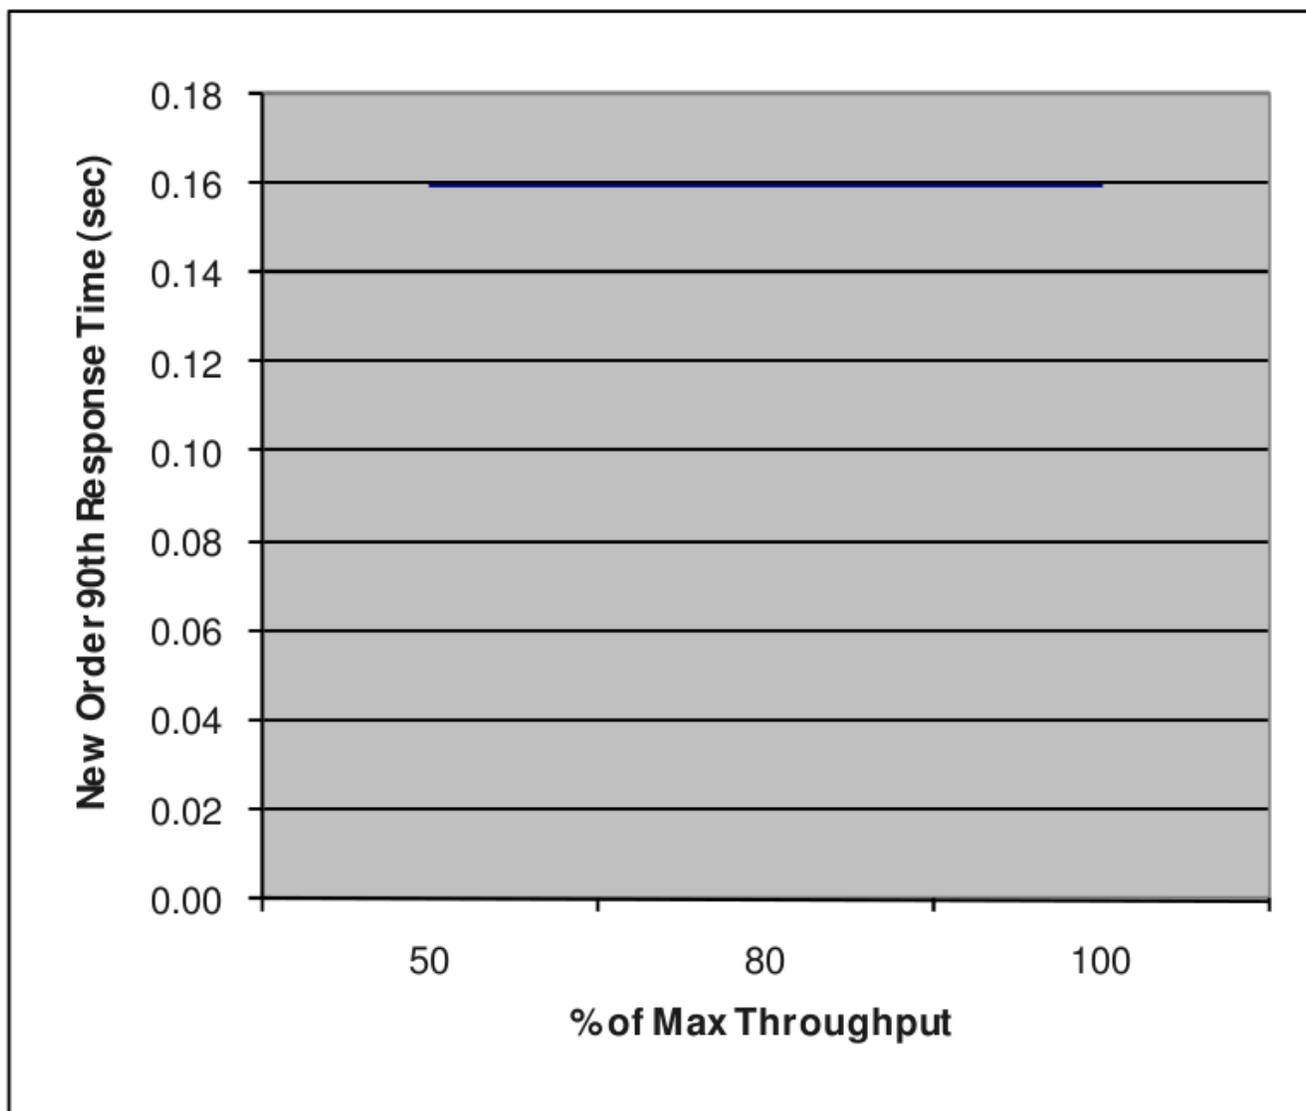


Figure 6-6: New-Order Response Time vs. Throughput



6.5 Think Time Frequency Distribution

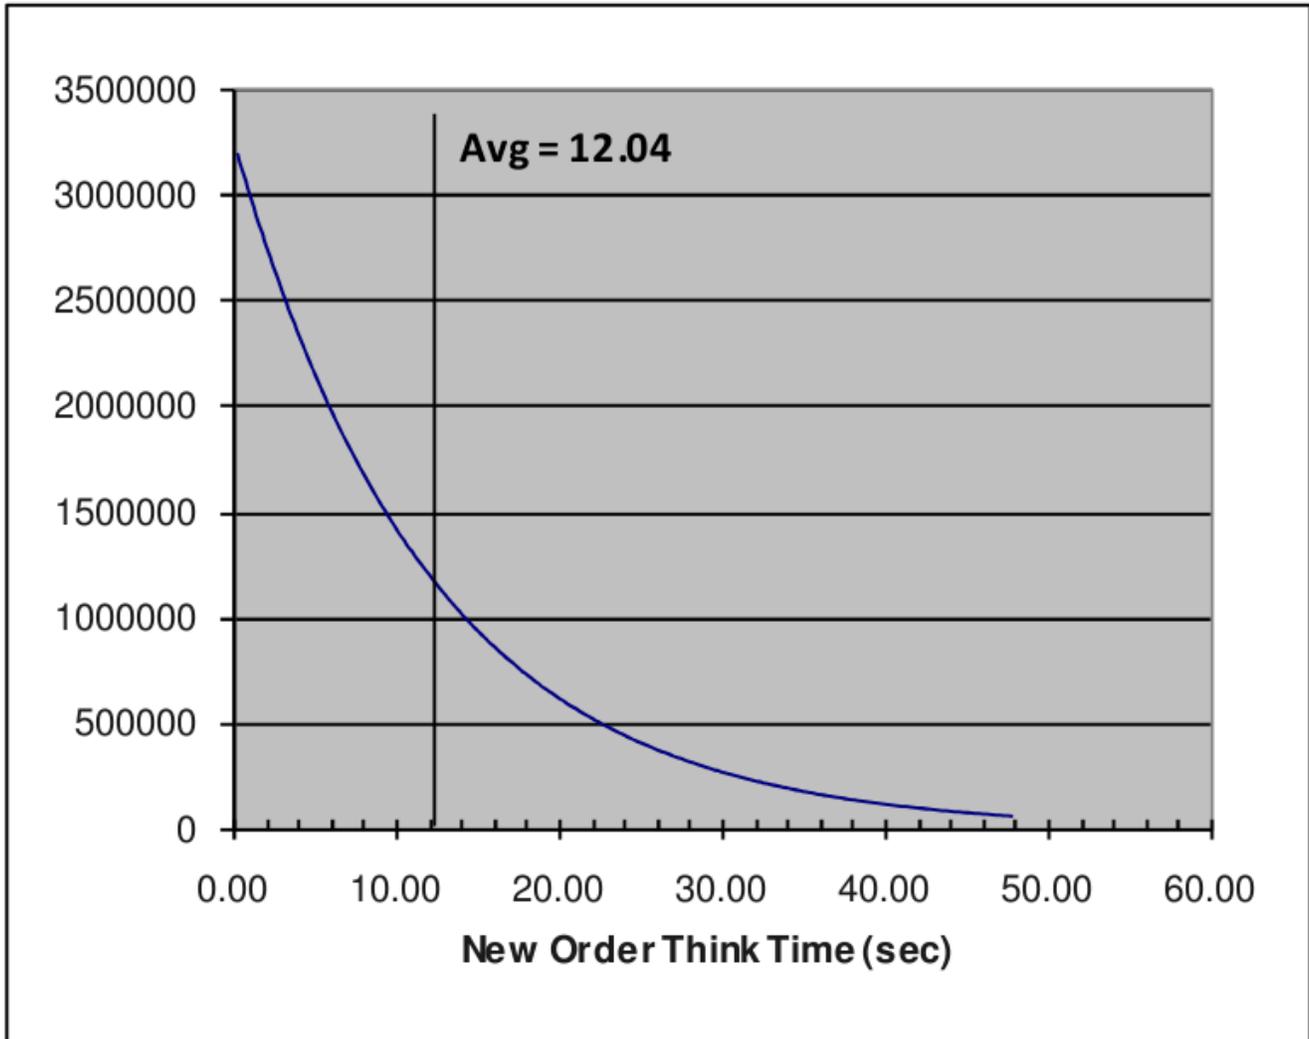


Figure 6-7: New-Order Think Time Distribution



6.6 Throughput versus Elapsed Time

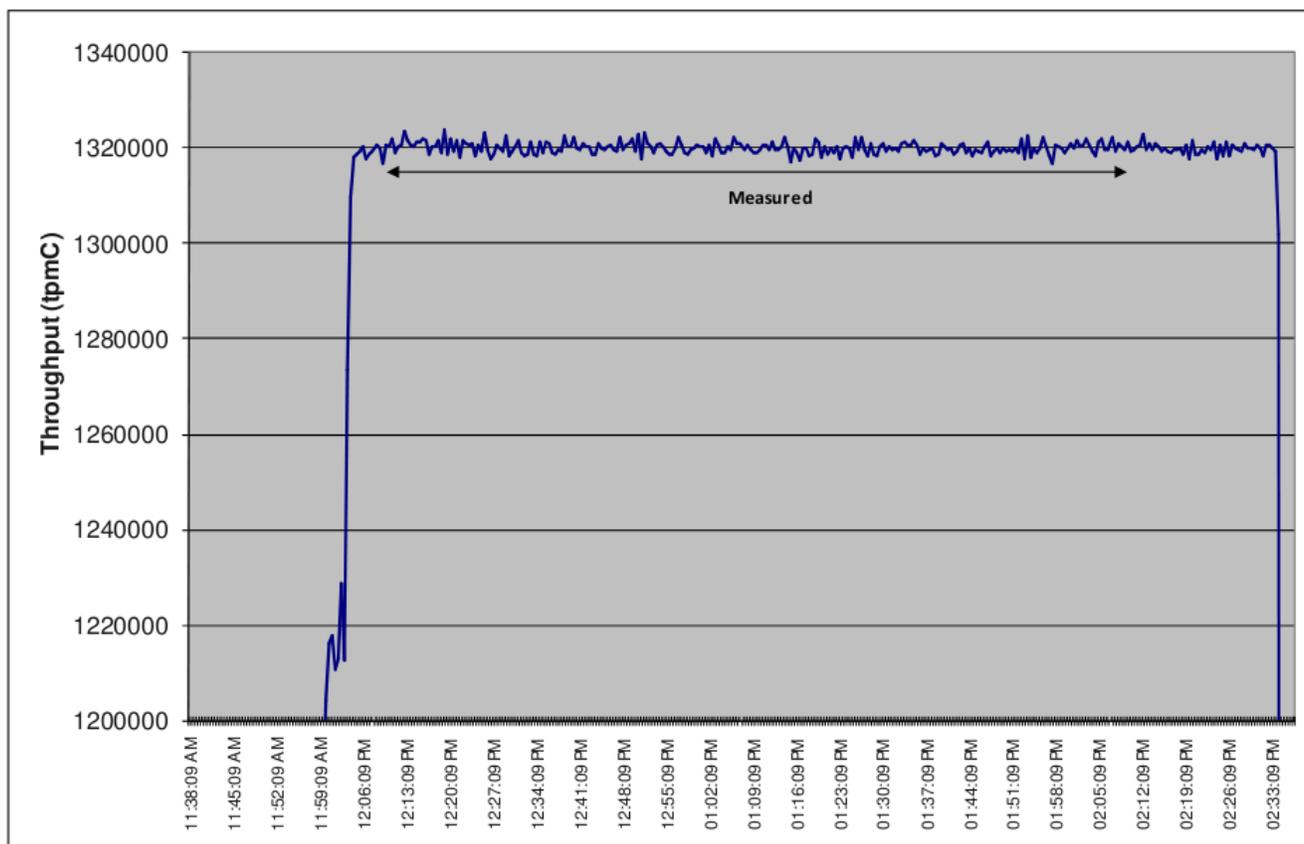


Figure 6-8: New-Order Throughput vs. Elapsed Time

6.7 Steady State Determination: Ramp-up & Run Times

All the emulated users were allowed to logon and do transactions. The user ramp-up phase is clearly visible on the graph above in Figure 6-8. New-Order throughput versus Elapsed Time graph shows that the system maintained a steady state during the measurement interval.

IBM used an internally developed Remote Terminal Emulator (RTE) for these tests. A total of 128,160 warehouses were configured; 119,616 were accessed during the runs. A ramp-up time of 73 minutes was specified, along with a run time of two hours.



7. References

1. Overview of the TPC Benchmark C: The Order-Entry Benchmark by Francois Raab, Walt Kohler, Amitabh Shah
<http://tpc.org/tpcc/detail.asp>
2. TPC-C Benchmark Description
<http://www.tpc.org/tpcc/detail.asp>
3. Benchmark Result Highlights
http://www.tpc.org/tpcc/results/tpcc_result_detail.asp?id=113022201
4. Benchmark Executive Summary by IBM
http://www.tpc.org/results/individual_results/IBM/ibm-linux-db2-x3650m4-kvm_ES.pdf
5. Benchmark Full Disclosure Report by IBM
http://www.tpc.org/results/fdr/tpcc/ibm-linux-db2-x3650m4-kvm_FDR.pdf