

TPC Benchmark™ C
Full Disclosure Report
for
IBM Netfinity 7600
using
Microsoft SQL Server 2000 Enterprise Edition
and
Windows 2000 Advanced Server

Submitted for Review
October 25, 2000



First Edition - October 2000

THE INFORMATION CONTAINED IN THIS DOCUMENT IS DISTRIBUTED ON AN AS IS BASIS WITHOUT ANY WARRANTY EITHER EXPRESSED OR IMPLIED. The use of this information or the implementation of any of these techniques is the customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item has been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environment do so at their own risk.

In this document, any references made to an IBM licensed program are not intended to state or imply that only IBM's licensed program may be used; any functionally equivalent program may be used.

This publication was produced in the United States. IBM may not offer the products, services, or features discussed in this document in other countries, and the information is subject to change without notice. Consult your local IBM representative for information on products and services available in your area.

© Copyright International Business Machines Corporation 2000. All rights reserved.

Permission is hereby granted to reproduce this document in whole or in part, provided the copyright notice as printed above is set forth in full text on the title page of each item reproduced.

U.S. Government Users - Documentation related to restricted rights: Use, duplication, or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Trademarks

IBM is a registered trademark and Netfinity is a trademark of International Business Machines Corporation.

The following terms used in this publication are trademarks of other companies as follows: TPC Benchmark, trademark of Transaction Processing Performance Council; Intel, Pentium and Xeon are registered trademarks of Intel Corporation; Microsoft, Windows and BenchCraft are trademarks or registered trademarks of Microsoft Corporation. Other company, product, or service names, which may be denoted by two asterisks (**), may be trademarks or service marks of others.

Notes

¹ MHz only measures microprocessor internal clock speed, not application performance. Many factors affect application performance.

² When referring to hard disk capacity, GB, or gigabyte, means one thousand million bytes. Total user-accessible capacity depends on operating environment.

Abstract

IBM Corporation conducted the TPC Benchmark™ C on the IBM Netfinity* 7600 configured as a client/server system. This report documents the full disclosure information required by the TPC Benchmark™ C Standard Specification, Revision 3.5, including the methodology used to achieve the reported results. All testing fully complied with this revision level.


The software used on the IBM Netfinity 7600 system includes Microsoft** Windows** 2000 Advanced Server operating system and Microsoft SQL Server 2000 Enterprise Edition database.

Two standard metrics, transactions per minute-C (tpmC) and price per tpmC (\$/tpmC), are reported as required by the TPC Benchmark C Standard Specification.

The benchmark results are summarized in the following table.

Hardware	Software	Total System Cost	tpmC	\$/tpmC	Total Solution Availability Date
IBM Netfinity 7600	Microsoft SQL Server 2000 Enterprise Edition Microsoft Windows 2000 Advanced Server	\$443,463	32,377.17	\$13.70	Oct. 25, 2000

The results of the benchmark and test methodology used were audited by Bradley J. Askins of InfoSizing, Inc. The auditor's attestation letter is contained in Section 9 of this report.

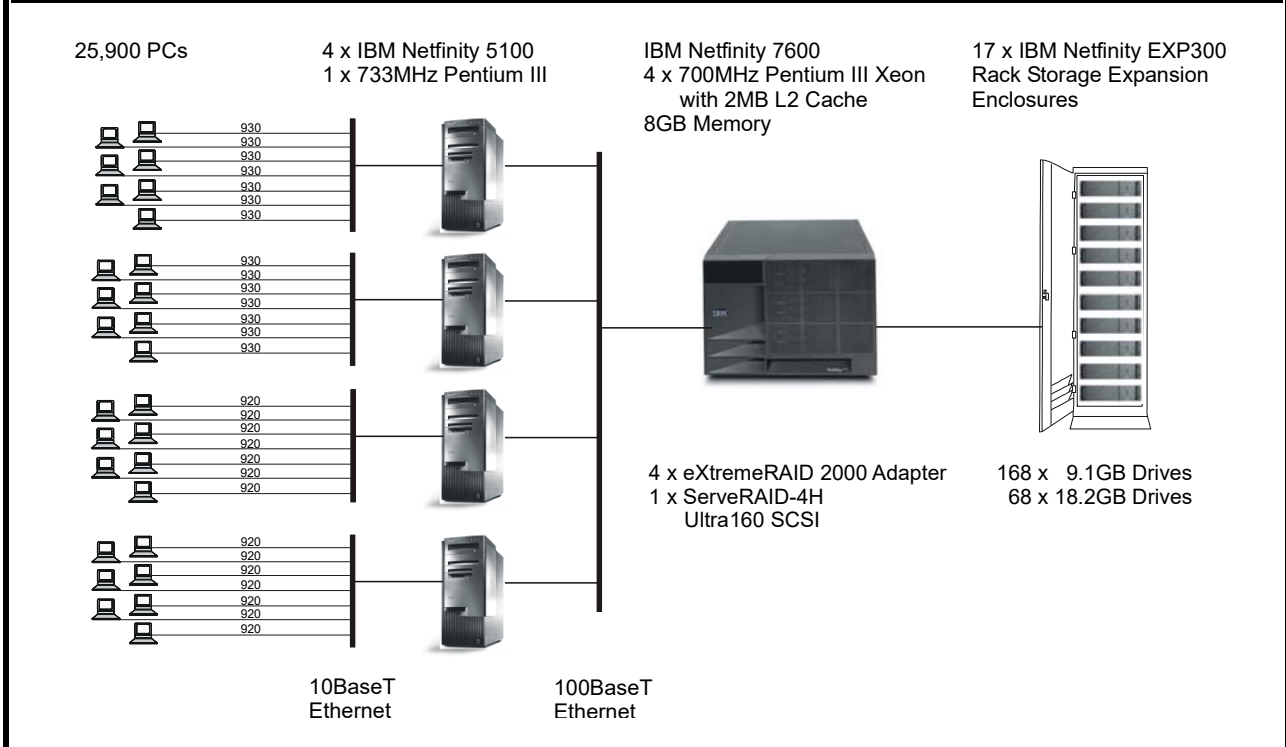


IBM Netfinity 7600 c/s

TPC-C Rev 3.5

Report Date: Oct. 25, 2000

Total System Cost		TPC-C Throughput	Price/Performance	Availability Date
\$443,463		32,377.17 tpmC	\$13.70 /tpmC	Oct. 25, 2000
Processors	Database Manager	Operating System	Other Software	Number of Users
4 Intel Pentium III Xeon 700MHz 2MB L2 Cache	Microsoft SQL Server 2000 Enterprise Edition	Microsoft Windows 2000 Advanced Server	Microsoft Visual C++ 6.0 Win32 Microsoft COM+	25,900



System Component	Qty	Server	Qty	Each of Four Clients
Processors	4	700MHz Pentium III Xeon w/2MB L2 Cache	1	733MHz Pentium III w/256KB L2 Cache
Cache				
Memory	16	512MB	2	128MB
Disk Controllers	4	Mylex eXtremeRAID 2000	1	Ultra160 SCSI Onboard
Disk Drives	1	ServeRAID-4H Ultra160	1	9.1GB Hard Disk
Total Storage	168	9.1GB (10000 rpm)		
Tape Drive	68	18.2GB (10000 rpm)		
	1	2766.4GB 10/20GB SCSI Tape Drive		

IBM Corporation		Netfinity 7600 c/s		TPC-C Revision 3.5			
				Report Date: Oct. 25, 2000			
Description	Order Number	Third-Party Brand	Pricing	Unit Price	Qty	Ext. Price	5-Yr. Maint.*
Server Hardware							
Netfinity 7600 700MHz/2MB Pentium III Xeon*	86655RY		1	\$9,660	1	\$9,660	\$7,995
700MHz/2MB L2 Cache Processor Upgrade	10K2332		1	3,757	3	11,271	0
512MB ECC SDRAM RDIMM Memory Kit	33L3117		1	1,940	15	29,100	0
Mylex eXtremeRAID 2000 Adapter***	E2000-4-32NB	Mylex	3	1,872	6	11,232	0
Netfinity ServeRAID-4H Ultra160 Adapter	37L6889		1	2,190	1	2,190	0
Netfinity 4.2M Ultra2 SCSI Cable	03K9311		1	113	18	2,034	0
Netfinity 10/100 Ethernet Adapter 2	34L1501		1	93	1	93	0
E54 14" (13.8" Viewable) Color Monitor*	6331B2N		1	179	1	179	270
10/20GB NS Internal SCSI Tape Drive	09N4042		1	398	1	398	0
Netfinity Rack*	9306900		1	1,560	2	3,120	2,520
Side Panel Kit	94G6669		1	185	1	185	0
Storage Hardware							
Netfinity EXP300 Rack Storage Enclosure*	35311RU		1	2,890	17	49,130	12,886
9.1GB 10K-4 Ultra160 SCSI Drive	37L7204		1	349	168	58,632	0
18.2GB 10K-4 Ultra160 SCSI Drive	37L7205		1	549	68	37,332	0
Subtotal						\$214,556	\$23,671
Server Software							
Microsoft SQL Server 2000 Enterprise Edition	810-00846	Microsoft	2	16,541	4	\$66,164	\$0
Microsoft Windows 2000 Advanced Server	C10-00475	Microsoft	2	2,399	1	2,399	0
Five-Year Maintenance for Software		Microsoft	2	10,475	1		10,475
Subtotal						\$68,563	\$10,475
Client Hardware							
IBM Netfinity 5100 / 733MHz/256KB Pentium III*	86582RY		1	2,205	4	\$8,820	\$22,780
9.1GB 10K Ultra160 SCSI Drive	37L7204		1	349	4	1,396	0
128MB 133MHz ECC SDRAM RDIMM	33L3123		1	279	4	1,116	0
Netfinity 10/100 Ethernet PCI Adapter	34L1501		1	93	4	372	0
E54 15" (13.8" Viewable) Color Monitor*	6331B2N		1	179	4	716	1,080
Subtotal						\$12,420	\$23,860
Client Software							
Microsoft Windows 2000 Server with COM+	C11-00821	Microsoft	2	738	4	2,952	0
Microsoft Visual C++ Professional 6.0 Win32	048-00317	Microsoft	2	549	1	549	0
Subtotal						\$3,501	\$0
User Connectivity							
8-Port 10Mbps Hub***	DEH2924	Generic	1	24	3,573	85,752	0
8-Port 10/100Mbps Nway Fast Ethernet Switch***	NX-DSS8	NETLUX	4	95	7	665	0
Subtotal						\$86,417	\$0
Total						\$385,457	\$58,006
Notes: * The standard 3-year warranty and the extended warranty on IBM hardware is for 7x24, 4-hour response time coverage. ** Five-year warranty. *** 10% or minimum 2 spares are added in place of on-site service (products have a 5-year return-to-vendor-warranty)				Five-Year Cost of Ownership:		\$443,463	
Pricing: 1 - Software House International; 2 - Microsoft Corp.; 3 - Mylex; 4 - NETLUX				tpmC Rating:		32,377.17	
Audited by Bradley J. Askins of InfoSizing, Inc.				\$ / tpmC:		\$13.70	
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 specification. If you find that stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.							

Numerical Quantities Summary			
MQTh, Computed Maximum Qualified Throughput: % throughput difference, reported and reproducibility runs:			32,377.17 tpmC 0.08%
Response Times (in seconds)	90%	Average	Maximum
New-Order	0.74	0.41	5.11
Payment	0.64	0.32	4.18
Order-Status	0.66	0.34	5.06
Delivery (Interactive)	0.12	0.12	1.20
Delivery (Deferred)	0.40	0.23	0.91
Stock-Level	1.95	1.36	4.68
Menu	0.12	0.12	1.56
Transaction Mix (in percent of total transactions)		Total Occurrences	Percent
New-Order		971,315	44.85
Payment		931,527	43.01
Order-Status		87,600	4.05
Delivery		87,496	4.04
Stock-Level		87,668	4.05
Emulation Delay (in seconds)	Response Time		Menu
New-Order	0.1		0.1
Payment	0.1		0.1
Order-Status	0.1		0.1
Delivery	0.1		0.1
Stock-Level	0.1		0.1
Keying/Think Times (in seconds)	Minimum	Average	Maximum
New-Order	18.00 / 0.00	18.01 / 12.04	18.03 / 120.50
Payment	3.00 / 0.00	3.02 / 12.03	3.03 / 120.50
Order Status	2.00 / 0.00	2.01 / 10.05	2.02 / 100.50
Delivery	2.00 / 0.00	2.01 / 5.06	2.02 / 50.50
Stock-Level	2.00 / 0.00	2.01 / 5.04	2.03 / 50.49
Test Duration			
Ramp-up time			9 minutes
Measurement interval			30 minutes
Number of transactions (all types) completed in measurement interval			2,253,103
Ramp-down time			38 minutes
Number of checkpoints in measurement interval			1
Checkpoint interval			30 minutes

Table of Contents

Abstract	3
Numerical Quantities Summary	6
Preface	11
General Items	12
Application Code Disclosure and Definition Statements	12
Benchmark Sponsor	12
Parameter Settings	12
Configuration Diagrams	12
<i>IBM Netfinity 7600 Measured Configuration</i>	13
<i>IBM Netfinity 7600 Priced Configuration</i>	14
<i>Drives and Storage Substitutions</i>	14
Clause 1: Logical Database Design Related Items	15
Table Definitions	15
Physical Organization of the Database	15
Insert and Delete Operations	15
Horizontal or Vertical Partitioning	15
Replication	15
Table Attributes	15
Clause 2: Transaction and Terminal Profiles Related Items	16
Random Number Generation	16
Screen Layout	16
Terminal Verification	16
Intelligent Terminals	16
Transaction Profiles	16
Deferred Delivery Mechanism	17
Clause 3: Transaction and System Properties Related Items	18
Atomicity Requirements	18
<i>Completed Transactions</i>	18
<i>Aborted Transactions</i>	18
Consistency Requirements	18
Isolation Requirements	19
Durability Requirements	19
<i>Loss of Data Test</i>	19
<i>Combined Loss of Log and Loss of System Test (Instantaneous Interruption and Loss of Memory)</i>	20
Clause 4: Scaling and Database Population Related Items	21
Cardinality of Tables	21
Distribution of Tables and Logs	22
Database Model Implemented	22
Partitions/Replications Mapping	22
180-Day Space Requirement	22
Clause 5: Performance Metrics and Response Time Related Items	23
Measured tpmC	23
Response Times	23
Keying/Think Times	23
Response Time Frequency Distribution Curves	24
Performance Curve for Response Time vs. Throughput	26
Throughput vs. Elapsed Time	27
Steady State Methodology	28
Work Performed during Steady State	28
<i>Transaction Flow</i>	28
Checkpoints	28
Reproducibility Methodology	28

Measurement Interval	29
Transaction Mix	29
Percentage of Total Mix	29
Number of Checkpoints	30
Clause 6: SUT, Driver and Communication Definition Related Items	31
Description of RTE	31
Emulated Components	31
Benchmarked and Targeted System Configuration Diagrams	31
Network Configuration	31
Network Bandwidth	31
Operator Intervention	31
Clause 7: Pricing Related Items	32
Hardware and Software Components	32
Availability Date	32
Measured tpmC	32
Country-Specific Pricing	32
Usage Pricing	32
System Pricing	33
Clause 9: Audit Related Items	34
Auditor	34
Availability of the Full Disclosure Report	34
<i>Attestation letter</i>	35
Appendix A: Source Code	37
<i>_delivery.h</i>	37
<i>_neworder.h</i>	38
<i>_orderstatus.h</i>	39
<i>_payment.h</i>	40
<i>_stocklevel.h</i>	41
<i>clientutils.c</i>	42
<i>clientutils.h</i>	44
<i>databuf.h</i>	45
<i>databuf.h.new</i>	49
<i>db_dblib_dll.dsp</i>	51
<i>delivery.h</i>	52
<i>dlldata.c</i>	53
<i>error.h</i>	54
<i>install.C</i>	56
<i>install.dsp</i>	66
<i>install.h</i>	67
<i>install.rc</i>	68
<i>install_com.cpp</i>	71
<i>isapi_dll_resource.h</i>	75
<i>isapi_dll.dsp</i>	75
<i>methods.h</i>	77
<i>mon_client.c</i>	79
<i>mon_client.h</i>	83
<i>neworder.h</i>	84
<i>orderstatus.h</i>	85
<i>payment.h</i>	86
<i>readregistry.cpp</i>	87
<i>readregistry.h</i>	88
<i>resource.h</i>	88
<i>stocklevel.h</i>	89
<i>rtetime.h</i>	90
<i>spinlock.h</i>	90

<i>tm_com_dll.dsp</i>	91
<i>tpcc.cpp</i>	93
<i>tpcc.def</i>	118
<i>tpcc.h</i>	119
<i>tpcc.rc</i>	121
<i>tpcc_com.cpp</i>	122
<i>tpcc_com.h</i>	124
<i>tpcc_com_all.cpp</i>	126
<i>tpcc_com_all.def</i>	131
<i>tpcc_com_all.dsp</i>	131
<i>tpcc_com_all.h</i>	133
<i>tpcc_com_all.idl</i>	135
<i>tpcc_com_all.rc</i>	136
<i>tpcc_com_all_resource.h</i>	137
<i>tpcc_type.h</i>	137
<i>tpcc_com_all.rgs</i>	138
<i>tpcc_com_all_i.c</i>	138
<i>tpcc_com_no.rgs</i>	140
<i>tpcc_com_os.rgs</i>	140
<i>tpcc_com_pay.rgs</i>	141
<i>tpcc_com_ps.def</i>	141
<i>tpcc_com_ps.dsp</i>	141
<i>tpcc_com_ps.h</i>	143
<i>tpcc_com_ps.idl</i>	146
<i>tpcc_com_ps_i.c</i>	147
<i>tpcc_com_ps_p.c</i>	148
<i>tpcc_com_sl.rgs</i>	159
<i>tpcc_dblib.cpp</i>	159
<i>tpcc_dblib.h</i>	170
<i>trans.h</i>	172
<i>txnlog.h</i>	175
<i>txn_base.h</i>	178
<i>webclnt.dsp</i>	178
<i>webclnt.dsw</i>	179
Stored Procedures	181
<i>neword.sql</i>	181
<i>payment.sql</i>	181
<i>ordstat.sql</i>	181
<i>delivery.sql</i>	181
<i>stocklev.sql</i>	181
<i>version.sql</i>	181
Appendix B: Database Design	183
<i>backup.sql</i>	183
<i>backupdev.sql</i>	183
<i>createdb.sql</i>	183
<i>tables.sql</i>	184
<i>idxcusnc.sql</i>	186
<i>idxcuscl.sql</i>	186
<i>idxdiscl.sql</i>	187
<i>idxitmcl.sql</i>	187
<i>idxnodcl.sql</i>	187
<i>idxodlcl.sql</i>	187
<i>idxordnc.sql</i>	188
<i>idxordcl.sql</i>	188
<i>idxstkcl.sql</i>	188

<i>idxwarcl.sql</i>	189
Loader Source Code	189
<i>getargs.c</i>	189
<i>strings.c</i>	193
<i>tpcc.h</i>	196
<i>tpcldr.c</i>	198
<i>time.c</i>	226
<i>random.c</i>	226
<i>tpcldr.dsp</i>	229
<i>tpcldr.dsw</i>	230
<i>tpcldr.mak</i>	230
Appendix C: Tunable Parameters	235
Microsoft Windows 2000 Advanced Server Configuration Parameters	235
<i>Microsoft SQL Server 2000 Startup Parameters</i>	235
<i>SQL Server 2000 Stack Size</i>	235
<i>Boot.ini</i>	235
<i>Microsoft SQL Server 2000 Configuration Parameters</i>	235
<i>Microsoft Windows 2000 Advanced Server Configuration Parameters</i>	236
Microsoft Windows 2000 Advanced Server Services	267
Microsoft Windows 2000 Advanced Server Registry Parameters	269
Disk Controller Configuration Parameters	270
<i>Mylex eXtremeRAID Adapter 0</i>	270
<i>Mylex eXtremeRAID Adapter 1</i>	274
<i>Mylex eXtremeRAID Adapter 2</i>	278
<i>Mylex eXtremeRAID Adapter 3</i>	283
<i>Netfinity ServeRAID-4H Ultra160 SCSI Adapter</i>	287
Client Configuration Parameters	292
<i>Microsoft Windows 2000 Server Configuration Parameters</i>	292
<i>COM+ Settings</i>	316
<i>TPCC Application Registry Parameters</i>	316
<i>Microsoft Internet Information Service Registry Parameters</i>	316
<i>Worldwide Web Service Registry Parameters</i>	317
RTE Input Parameters	317
Appendix D: 180-Day Space	326
Appendix E: Third-Party Quotations	327

Preface

The TPC Benchmark™ C was developed by the Transaction Processing Performance Council (TPC). The TPC was founded to define transaction processing benchmarks and to disseminate objective, verifiable performance data to the industry. This full disclosure report is based on the TPC Benchmark C Standard Specification Version 3.5, released October 25, 1999.

The TPC describes this benchmark in Clause 0.1 of the specification as follows:

TPC Benchmark C is an On Line Transaction Processing (OLTP) workload. It is 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 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 on 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.

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 any other environment 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 benchmarking when critical capacity planning and/or product evaluation decisions are contemplated.

General Items

Benchmark Sponsor

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

This benchmark was sponsored by International Business Machines Corporation.

Application Code Disclosure 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 and output functions.

Appendix A contains all source code implemented in this benchmark.

Parameter Settings

Settings must be provided for all customer-tunable parameters and options that have been changed from the defaults found in actual products, including but not limited to:

- ✓ Database tuning options*
- ✓ Recovery/commit options*
- ✓ Consistency/locking options*
- ✓ Operating system and application configuration parameters.*
- ✓ Compilation and linkage options and run-time optimizations used to create/install applications, OS, and/or databases.*

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

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

Configuration Diagrams

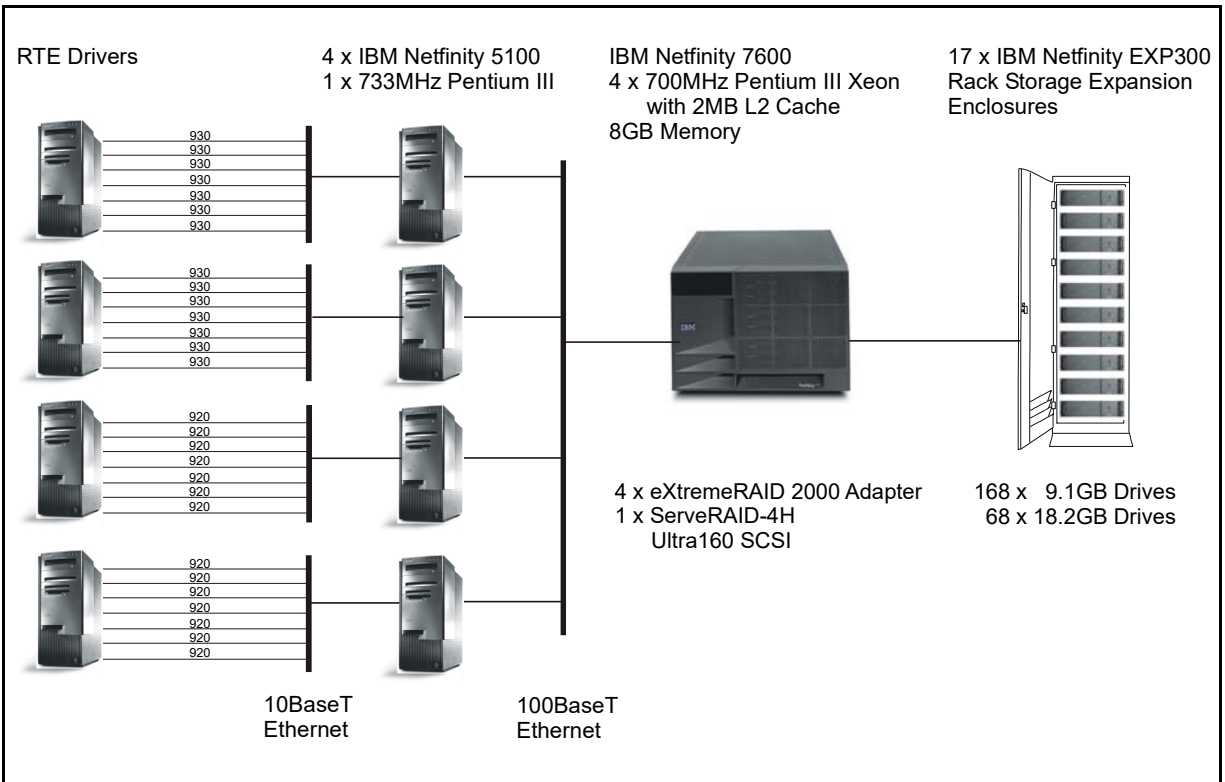
Diagrams of both measured and priced configurations must be provided, accompanied by a description of the differences.

The configuration diagrams for the tested and priced systems are provided on the following pages.

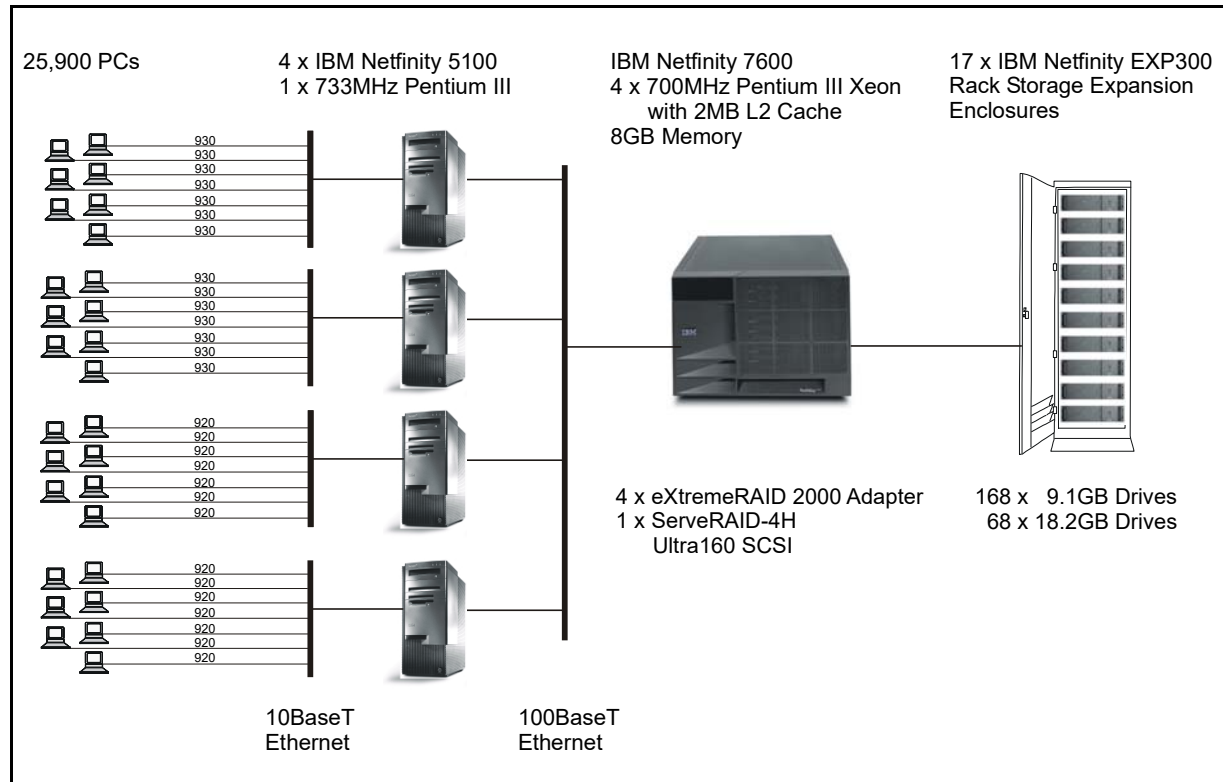
The Remote Terminal Emulator (RTE) used for these TPC Benchmark C tests is the Microsoft BenchCraft RTE. The components of the configuration being emulated by the RTE are the workstations and the Ethernet hubs. Appendix C contains a listing of the RTE scripts and inputs used in the benchmark testing.

The benchmarked configuration used IBM Netfinity 5100 systems as clients, which executed the terminal I/O and submitted transactions to COM+ servers, which are also running on the clients. These COM+ servers forwarded the transaction requests to the server, and returned the results to the RTE. Microsoft SQL Server 2000 Enterprise Edition is the DBMS executing on the server

IBM Netfinity 7600 Measured Configuration



IBM Netfinity 7600 Priced Configuration



Drives and Storage Substitutions

The priced and measured configurations were identical with one exception. The measured configuration used two different 9.1GB drives:

- v 9.1GB 10K-4 Ultra160 SCSI Drive
- v 9.1GB 10K Wide Ultra SCSI Drive

To verify equivalent performance of the drives, a regular TPC-C run was made with DISKPERF on to monitor the disk performance. The performance monitor data showed that the disk latency and disk queue depth for each array were equivalent.

Clause 1: Logical Database Design Related Items

Table Definitions

Listings must be provided for all table definition statements and all other statements used to set up the database. (8.1.2.1)

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

Physical Organization of the Database

The physical organization of tables and indexes within the database must be disclosed. (8.1.2.2)

Physical space was allocated to Microsoft SQL Server on the server disks as detailed in Figure 4-2.

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 restriction 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 maximum key value for these new rows. (8.1.2.3)

All insert and delete functions were fully operational during the running of the benchmark. The space required for an additional 5 percent of the initial table cardinality was allocated to Microsoft SQL Server 2000 and priced as static space.

Horizontal or Vertical Partitioning

While there are few restrictions placed upon horizontal or vertical partitioning of tables and rows in the TPC-C benchmark (see Clause 1.6), any such partitioning must be disclosed. (8.1.2.4)

Partitioning was not used in this benchmark.

Replication

Replication tables, if used, must be disclosed (see Clause 1.4.6). (8.1.2.5)

Replication was not used in this benchmark.

Table Attributes

Additional and/or duplicated attributes in any table must be disclosed, along with a statement on the impact on performance (see Clause 1.4.7). (8.1.2.6)

No additional attributes were used in this benchmark.

Clause 2: Transaction and Terminal Profiles Related Items

Random Number Generation

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

The seeds and offsets for the random number generator were collected and verified to be different for each driver. The auditor selected samples of the generated numbers from the database. The samples were verified to have no discernible patterns.

Screen Layout

The actual layouts of the terminal input/out screens must be disclosed.

All screen layouts followed the TPC Benchmark C Standard Specification.

Terminal Verification

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 must for the demonstration in 8.1.3.3 must be disclosed and commercially available (including supporting software and maintenance).

The auditor verified terminal features by direct experimentation. The benchmarked configuration uses Microsoft Internet Explorer 5.0 and HTML scripts as the terminal interface.

Intelligent Terminals

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

The terminals emulated in the priced configuration are IBM PC desktop computer systems. All processing of the input/output screens was handled by the IBM Netfinity 5100 clients. The screen input/output was managed via HTML strings that comply with the HTML Version 2.0 specification. A listing of the code used to implement the intelligent terminals is provided in Appendix A. All data manipulation was handled by the IBM Netfinity 7600.

Transaction Profiles

The percentage of home and remote order-lines in the New-Order transactions must be disclosed. (8.1.3.5)

The percentage of New-Order transactions that were rolled back as a result of an unused item number must be disclosed. (8.1.3.6)

The number of items per orders entered by New-Order transactions must be disclosed. (8.1.3.7)

The percentage of home and remote Payment transactions must be disclosed. (8.1.3.8)

The percentage of Payment and Order-Status transactions that used non-primary key (C_LAST) access to the database must be disclosed. (8.1.3.9)

The percentage of Delivery transactions that were skipped as a result of an insufficient number of rows in the NEW-ORDER table must be disclosed. (8.1.3.10)

The mix (i.e., percentages) of transaction types seen by the SUT must be disclosed. (8.1.3.11)

Table 2-1. Transaction Statistics

New Order	Value (%)
Home warehouse order lines	99.00
Remote warehouse order lines	1.00
Rolled back transactions	1.03
Average number of items per order	10.00
Payment	
Home warehouse payment transactions	85.01
Remote warehouse payment transactions	14.99
Non-Primary Key Access	
Payment transactions using C_LAST	60.07
Order-Status transactions using C_LAST	60.18
Delivery	
Delivery transactions skipped	0
Transaction Mix	
New-Order	44.85
Payment	43.01
Order-Status	4.05
Stock-Level	4.05
Delivery	4.04

Deferred Delivery Mechanism

The queuing mechanism used to defer execution of the Delivery transaction must be disclosed. (8.1.3.12)

The deferred delivery operation is queued by making an entry in an array within the application process (tpcc.dll) running on the client. Background threads within the application asynchronously process the queued delivery transactions.

The source code is listed in Appendix A.

Clause 3: Transaction and System Properties Related Items

The results of the ACID test 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. (8.1.4.1)

Atomicity Requirements

The system under test must guarantee that 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.

All ACID tests were conducted according to specification.

Completed Transactions

The following steps were performed to verify the Atomicity of completed transactions.

1. The balance was retrieved from the CUSTOMER table for a random Customer, District and Warehouse, giving BALANCE_1.
2. The Payment transaction was executed for the Customer, District and Warehouse used in step 1.
3. The balance was retrieved again for the Customer used in step 1 and step 2, giving BALANCE_2. It was verified that BALANCE_1 was greater than BALANCE_2 by AMT.

Aborted Transactions

The following steps were performed to verify the Atomicity of the aborted Payment transaction:

1. The Payment application code was changed to execute a rollback of the transaction instead of performing the commit.
2. Using the balance, BALANCE_2, from the CUSTOMER table retrieved for the completed transaction, the Payment transaction was executed for the Customer, District and Warehouse used in step 1 of section 3.1.1. The transaction rolled back due to the change in the application code from step 1.
3. The balance was retrieved again for the Customer used for step 2, giving BALANCE_3. It was verified that BALANCE_2 was equal to BALANCE_3.

Consistency Requirements

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

Consistency conditions one through four were tested using a shell script to issue queries to the database. The results of the queries demonstrated that the database was consistent for all four tests.

Isolation Requirements

Sufficient conditions must be enabled at either the system or the application level to ensure that the required isolation defined in Clause 3.4.1 is obtained.

Isolation tests one through seven were run using shell scripts to issue queries to the database. Each script included timestamps to demonstrate the concurrency of operations. The results of the queries were captured and placed in files. The auditor reviewed the results and verified that the isolation requirements had been met.

In addition, the phantom tests and the stock-level tests were run and verified.

Case A was followed for Isolation test seven.

Durability Requirements

The tested system must guarantee durability: the ability to preserve the effects of committed transactions and ensure database 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)*

Loss of Data Test

The following steps were successfully performed to pass the Durability test of failure of a disk unit with database tables:

1. The contents of the database were backed up to several database dump devices during the initial database load. There were no dump devices on the disk array from which a drive was removed as part of this test.
2. The current count of the total number of orders was determined by the sum of D_NEXT_O_ID for all rows in the district table giving SUM1.
3. A test was started with 150 users submitting transactions.
4. A disk containing a portion of each of the tables in the tpcc database was removed causing SQL Server to report errors accessing that device.
5. The run was aborted and SQL Server was restarted. Upon restart, the database tpcc reported numerous errors relating to the failed database device.
6. The transaction log was dumped to disk and the failed disk was replaced with a spare disk and was recovered.
7. The database was recovered and restored from the backup dump devices. Afterwards, the transaction log was applied to the database.
8. Step 2 was repeated to obtain the current count of the total number of orders giving SUM2.
9. It was verified that the sum of D_NEXT_O_ID after the database is recovered is greater than or equal to the sum of D_NEXT_O_ID before the run, plus all new order transactions completed during the run minus any rollback transactions.
10. Consistency Condition 3 was verified.

Combined Loss of Log and Loss of System Test (Instantaneous Interruption and Loss of Memory)

1. The current count of the total number of orders was determined by the sum of D_NEXT_O_ID for all rows in the district table giving SUM1.
2. A test was started under full load with all users submitting transactions.
3. One disk from the log array was removed. Since the disk was RAID-1 mirrored, SQL Server continued to process transactions without interruption.
4. The test continued under full load with all users submitting transactions. A checkpoint was issued, and the system continued to run for another 5 minutes.
5. The server under test was powered off, which removed power from the system and the memory.
6. The server was powered on again.
7. SQL Server was started to initiate automatic recovery from its log.
8. Step 1 was repeated to obtain the current count of the total number of orders giving SUM2.
9. It was verified that the sum of D_NEXT_O_ID after the database is recovered is greater than or equal to the sum of D_NEXT_O_ID before the run, plus all new order transactions completed during the run minus any rollback transactions.

Clause 4: Scaling and Database Population Related Items

Cardinality of Tables

The cardinality (e.g., the number of rows) of each table, as it existed at the start of the benchmark run (see Clause 4.2), must be disclosed. If the database was over-scaled and inactive rows of the WAREHOUSE table were deleted (see Clause 4.2.2), the cardinality of the WAREHOUSE table as initially configured and the number of rows deleted must be disclosed. (8.1.5.1)

The database was originally built with 2,590 warehouses, and the audited run used all 2,590 warehouses.

Table 4-1. Initial Cardinality of Tables

Table Name	Rows
Warehouse	2,590
District	25,900
Customer	100,000
History	23,310,000
Orders	77,700,000
New Order	77,700,000
Order Line	77,700,000
Stock	776,996,205
Item	259,000,000
Inactive Warehouses	0

Distribution of Tables and Logs

The distribution of tables and logs across all media must be explicitly depicted for the tested and priced systems. (8.1.5.2)

Figure 4-2 depicts the database configuration of the tested system to meet the 8-hour steady state requirement.

Figure 4-2. Data Distribution for the Benchmarked Configuration

Controller	Drives	Partition	Size	Use
1	56 - 18.2GB	E: F:	37530MB 21550MB	Customer and Stock Misc.
2	56 - 9.1GB	G: H:	37530MB 21550MB	Customer and Stock Misc.
3	56 - 9.1GB	I: J, Y:	37530MB 21550MB, 120000MB (NTFS)	Customer and Stock Misc., backup1
4	56 - 9.1GB	K: L, Z:	37530MB 21550MB, 120000MB (NTFS)	Customer and Stock Misc., backup2
5	12 - 8.2GB	O:	90000MB	Log file

Database Model Implemented

A statement must be provided that describes:

- 1. The database model implemented by the DBMS used (e.g., relational, network, hierarchical)*
- 2. The database interface (e.g., embedded, call level) and access language (e.g., SQL, DL/1, COBOL, read/write) used to implement the TPC-C transactions. 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.*

Microsoft SQL Server 2000 Enterprise Edition is a relational database. The interface used was Microsoft SQL Server stored procedures accessed with Remote Procedure Calls embedded in C code using the Microsoft DBLIB interface.

Partitions/Replications Mapping

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

The database was neither partitioned nor replicated.

180-Day Space Requirement

Details of the 180-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 (see Clause 4.2.3). (8.1.5.5)

See Appendix D for details about how the 180-day space requirements were calculated.

Clause 5: Performance Metrics and Response Time Related Items

Measured tpmC

Measured tpmC must be reported. (8.1.6.1)

Measured tpmC: 32,377.17 tpmC

Price per tpmC: \$13.70 per tpmC

Response Times

Ninetieth percentile, maximum and average response times must be reported for all transaction types as well as for the Menu response time. (8.1.6.2)

The TPC-C requirements for the average response time and the 90th percentile were met. Table 5-1 provides the response times for each of the transaction types and the menu for the measured system.

Table 5-1. Response Times in Seconds

Transaction Type	90%	Average	Maximum
New-Order	0.74	0.41	5.11
Payment	0.64	0.32	4.18
Order-Status	0.66	0.34	5.06
Delivery (interactive)	0.12	0.12	1.20
Delivery (deferred)	0.40	0.23	0.91
Stock-Level	1.95	1.36	4.68
Menu	0.12	0.12	1.56

Keying/Think Times

The minimum, the average, and the maximum keying and think times must be reported for each transaction type. (8.1.6.3)

Table 5-2 lists the keying/think times for the measured system.

Table 5-2. Keying/Think Times

Transaction Type	Minimum	Average	Maximum
New-Order	18.00 / 0.00	18.01 / 12.04	18.03 / 120.50
Payment	3.00 / 0.00	3.02 / 12.03	3.03 / 120.50
Order-Status	2.00 / 0.00	2.01 / 10.05	2.02 / 100.50
Delivery	2.00 / 0.00	2.01 / 5.06	2.02 / 50.50
Stock-Level	2.00 / 0.00	2.01 / 5.04	2.03 / 50.49

Response Time Frequency Distribution Curves

Response time frequency distribution curves (see Clause 5.6.1) must be reported for each transaction type. (8.1.6.4)

Figure 5-1. New-Order Transaction - Response Time Frequency Distribution




Figure 5-2. Payment Transaction - Response Time Frequency Distribution




Figure 5-3. Order-Status Transaction - Response Time Frequency Distribution



Figure 5-4. Delivery Transaction - Response Time Frequency Distribution


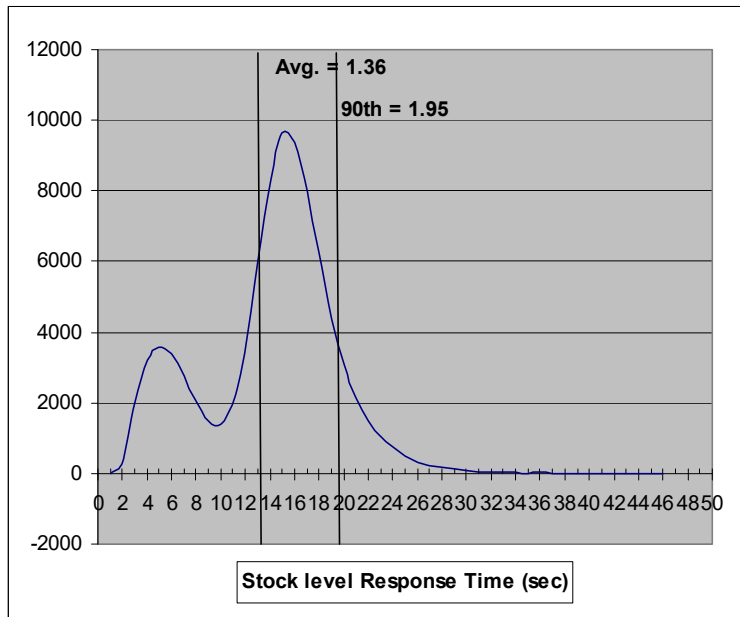


Figure 5-5. Stock-Level Transaction - Response Time Frequency Distribution



Performance Curve for Response Time vs. Throughput

The performance curve for response time vs. throughput (see Clause 5.6.2) must be reported for the New-Order transaction. (8.1.6.5)

Figure 5-6. New-Order Response Time vs. Throughput






Figure 5-7. New-Order Think Time Distribution



Throughput vs. Elapsed Time

A graph of throughput vs. elapsed time (see Clause 5.6.5) must be reported for the New-Order transaction.

Figure 5-8. New-Order Throughput vs. Elapsed Time



Steady State Methodology

The method used to determine that the SUT had reached a steady state prior to commencing the measurement interval (see Clause 5.5) must be described. (8.1.6.9)

Figure 5-8 shows that the system was in steady state at the beginning of the measurement interval.

Work Performed during Steady State

A description of how the work normally performed during a sustained test (e.g., checkpointing, writing redo/undo log records) actually occurred during the measurement interval must be reported. (8.1.6.10)

Transaction Flow

The RTE generated the required input data to choose a transaction from the menu. This data was time-stamped. The response for the requested transaction was verified and time-stamped in the RTE log files.

The RTE generated the required input data for the chosen transaction. It waited to complete the minimum required key time before transmitting the input screen. The transmission was time-stamped. The return of the screen with the required response data was time-stamped. The difference between these two time-stamps was the response time for that transaction and was logged in the RTE log.

The RTE then waited the required think time interval before repeating the process starting at selecting another transaction from the menu.

The RTE transmissions were sent to application processes running on the client machines through Ethernet LANs. These client application processes handled all screen I/O as well as all requests to the database on the server. The applications communicated with the database server over another Ethernet LAN using Microsoft SQL Server DBLIB library and RPC calls.

Checkpoints

Checkpoints were executed on the server during the ramp-up phase and at 30-minute intervals. Each measured run contained one checkpoint. SQL Server was started with trace flag 3502, which caused it to log the occurrence of the checkpoint. This information was used to verify that the checkpoints occurred at the appropriate times during the test run.

During a checkpoint, SQL Server flushes all dirty pages from its cache to disk. It places a record in the database transaction log indicating that the checkpoint has completed and that all transactions, which were committed prior to the checkpoint have been written to disk.

Reproducibility Methodology

A description of the method used to determine the reproducibility of the measurement results must be reported. (8.1.6.11)

A repeatability measurement was taken on the IBM Netfinity 7600 for the same length of time as the measured run. The repeatability measurement was 32,351.47 tpmC.

Measurement Interval

A statement of the duration of the measurement interval for the reported Maximum Qualified Throughput (tpmC) must be included. (8.1.6.12)

The measurement interval was 30 minutes.

Transaction Mix

The method of regulation of the transaction mix (e.g., card decks or weighted random distribution) must be described. If weighted distribution is used and the RTE adjusts the weights associated with each transaction type, the maximum adjustments to the weight from the initial value must be disclosed. (8.1.6.13)

See Table 5-3.

The RTE was given a weighted random distribution, which was not adjusted during the run.

Percentage of Total Mix

The percentage of the total mix for each transaction type must be disclosed.

See Table 5-3.

Table 5-3. Transaction Statistics and Transaction Mix

New Order	Value (%)
Home warehouse order lines	99.00
Remote warehouse order lines	1.00
Rolled back transactions	1.03
Average number of items per order	10.00
Payment	
Home warehouse payment transactions	85.01
Remote warehouse payment transactions	14.99
Non-Primary Key Access	
Payment transactions using C_LAST	60.07
Order-Status transactions using C_LAST	60.18
Delivery	
Delivery transactions skipped	0
Transaction Mix	
New-Order	44.85
Payment	43.01
Order-Status	4.05
Delivery	4.04
Stock Level	4.05

Number of Checkpoints

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

Checkpoints were performed during the ramp-up period and during each measured run interval. The measured interval checkpoint started 11 minutes and 5 seconds after the start of the measurement interval. The checkpoint in the measured interval lasted 10 minutes.

The checkpoints were verified to be clear of the protected zones around the beginning and end of the measurement intervals. The checkpoint interval was 30 minutes.

Clause 6: SUT, Driver and Communication Definition Related Items

Description of RTE

The RTE input parameters, code fragments, functions, etc., used to generate each transaction input field must be disclosed. (8.1.7.1)

The RTE used was Microsoft BenchCraft RTE. Benchcraft is a proprietary tool provided by Microsoft and is not commercially available. The RTE input is listed in Appendix C.

Emulated Components

It must be demonstrated that the functionality and performance of the components being emulated in the Driver System are equivalent to that of the priced system. The results of the test described in Clause 6.6.3.4 must be disclosed. (8.1.7.2)

No components were emulated.

Benchmarked and Targeted System Configuration Diagrams

A complete functional diagram of both the benchmarked configuration and the configuration of the proposed (target) system must be disclosed. A detailed list of all software and hardware functionality being performed on the Driver System, and its interface to the SUT must be disclosed (see Clause 6.6.3.6).

The driver RTE generated the transaction input data and transmitted it to the client in HTML format. The driver RTE received the output from the System under Test, time-stamped it, and forwarded it to the Master RTE for post-test processing. No other functionality was included on the driver RTE.

Detailed diagrams of the benchmarked and priced configurations are provided in the section called “General Items” at the beginning of this document.

Network Configuration

The network configurations of both the tested services and the 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 (see Clause 6.6.4). (8.1.7.4)

See the measured and priced configuration diagrams (pages 13 and 14) for details about the network configuration.

Network Bandwidth

The bandwidth of the network(s) used in the tested/priced configuration must be disclosed. (8.1.7.5)

The Ethernet used in the LAN complies with the IEEE.802.3 standard. The LANs that connected the driver RTEs to the clients had a bandwidth of 10Mbps. The LAN that connected the clients to the server had a bandwidth of 100Mbps.

Operator Intervention

If the configuration requires operator intervention (see Clause 6.6.6), the mechanism and the frequency of this intervention must be disclosed. (8.1.7.6)

The configuration did not require any operator intervention to sustain the reported throughput during the eight-hour period.

Clause 7: Pricing Related Items

Hardware and Software Components

A detailed list of the hardware and software used in the priced system must be reported. Each separately orderable item must have a 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(s) and effective date(s) must also be reported. (8.1.8.1)

The total 5-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. (8.1.8.2)

A detailed list of all hardware and software, including the 5-year price, is provided in the Executive Summary at the front of this report. All third-party quotations are included in Appendix E at the end of this document.

Availability Date

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 for the priced system must be the date at which all components are committed to be available. (8.1.8.3)

All hardware and software used in this benchmark are currently available.

Measured tpmC

A statement of the measured tpmC, as well as the respective calculations for the 5-year pricing, price/performance (price/tpmC) and the availability date must be included. (8.1.8.4)

{ Maximum Qualified Throughput:	32,377.17 tpmC
{ Price per tpmC:	\$13.70 per tpmC
{ Five-year cost of ownership:	\$443,463

Country-Specific Pricing

Additional Clause 7 related items may be included in the Full Disclosure Report for each country-specific priced configuration. Country-specific pricing is subject to Clause 7.1.7.

The configuration is priced for the United States of America.

Usage Pricing

For any usage pricing, the sponsor must disclose:

- v Usage level at which the component was priced.*
- v A statement of the company policy allowing such pricing. (8.1.8.6)*

The component pricing based on usage is shown below:

- v 4 Microsoft Windows 2000 Server*
- v 1 Microsoft Windows 2000 Advanced Server*
- v 4 Microsoft SQL Server 2000 Enterprise Edition (based on per-processor price)*

- v 5-year support for hardware components (except for components for which 10 percent spares are provided)

System Pricing

System pricing should include subtotals for the following components: Server Hardware, Server Software, Client Hardware, Client Software, and Network Components used for terminal connection (see Clause 7.2.2.3). (8.1.8.7)
System pricing must include line item indication where non-sponsoring companies' brands are used. System pricing must also include line item indication of third-party pricing. (8.1.8.8)

A detailed list of all hardware and software, including the 5-year price, is provided in the Executive Summary at the front of this report. All third-party quotations are included in Appendix E at the end of this document.

Clause 9: Audit Related Items

Auditor

The auditor's name, address, phone number, and a copy of the auditor's attestation letter indicating compliance must be included in the Full Disclosure Report. (8.1.9.1)

This implementation of the TPC-C benchmark was audited by Bradley J. Askins of InfoSizing, Inc. The auditor's attestation letter is provided in this section.

Availability of the Full Disclosure Report

The Full Disclosure Report must be readily available to the public at a reasonable charge, similar to the charges for similar documents by the test sponsor. The report must be made available when results are made public. In order to use the phrase "TPC BenchmarkTMC," the Full Disclosure Report must have been submitted to the TPC Administrator as well as written permission obtained to distribute same. (8.2)

Requests for the TPC Benchmark C Full Disclosure Report should be sent to:

TPC
650 N. Winchester Blvd.
Suite 1
San Jose, CA 95128

