

TPC Benchmark™ D
Full Disclosure Report

Data General Corporation

AViiON™ Model AV6600
With Oracle 8.0.3
and Windows NT 4.0

First Edition
August 1997

TPC Benchmark™ D Full Disclosure Report
Data General AV6600 Server With Oracle 8.0.3 and Windows NT 4.0
Edition 1

First Printing August 1997

All rights reserved. Permission is hereby granted to reproduce this document in whole or in part provided the copyright notice is included on the title page of each item reproduced.

Printed in U.S.A.

Data General Corporation (DGC) believes that the technical, pricing and discounting information in this document is accurate as of its publication date. The performance information in this document is for guidance only. System performance is highly dependent on many factors including system hardware, system and user software, and user-application characteristics. Customer applications must be carefully evaluated before estimating performance. DGC does not warrant or represent that a user can or will achieve similar performance as expressed in this document.

THE TERMS AND CONDITIONS GOVERNING THE SALE OF DGC HARDWARE PRODUCTS AND THE LICENSING OF DGC SOFTWARE CONSIST SOLELY OF THOSE SET FORTH IN THE WRITTEN CONTRACTS BETWEEN DGC AND ITS CUSTOMERS. NO REPRESENTATION OR OTHER AFFIRMATION OF FACT CONTAINED IN THIS DOCUMENT INCLUDING BUT NOT LIMITED TO STATEMENTS REGARDING PRICE, CAPACITY, RESPONSE-TIME PERFORMANCE, SUITABILITY FOR USE, OR PERFORMANCE OF PRODUCTS DESCRIBED HEREIN SHALL BE DEEMED TO BE A WARRANTY BY DGC FOR ANY PURPOSE, OR GIVE RISE TO ANY LIABILITY OF DGC WHATSOEVER.

DGC assumes no responsibility for any errors that may appear in this document. DGC reserves the right to make changes in specifications and other information contained in this document without prior notice, and the reader should in all cases consult DGC to determine whether any such changes have been made.

AViiON, and CLARiiON are U.S. registered trademarks of the Data General Corporation.

ORACLE, SQL*Loader, Oracle8, and Pro*C are registered trademarks of Oracle Corporation.

Windows NT is a trademark of the Microsoft Corporation.

TPC Benchmark is a trademark of the Transaction Processing Performance Council.

Additional Copies

Please use the toll-free number below to request additional copies of this report.

Telephone: 1-800-DATAGEN

Document Title: TPC Benchmark™ D Full Disclosure Report
Data General AV6600 Server With Oracle 8.0.3 and Windows NT 4.0
Edition 1

Abstract

This report documents the methodology and results of the TPC Benchmark D test conducted on the AV6600 using Oracle 8.0.3, in conformance with the requirements of the TPC-D Benchmark Specification. The operating system used for the benchmark was Windows NT 4.0 Enterprise Edition with Service Pack 3. The application was written in C and compiled using Microsoft Visual C++ for NT.

The TPC-D Benchmark 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.

The benchmark configuration, environment and methodology used to produce and validate the test results, and the pricing model used to calculate the price/performance were audited by Francois Raab, Information Paradigm, Inc., to verify compliance with the relevant TPC specifications. The auditor's letter of attestation is attached as Appendix G.

Data General Corporation		AViiON AV6600 with Oracle 8.03		TPC-D REV 1.2.3																																																													
				Report Date: 6-Aug-97																																																													
TOTAL SYSTEM COST		TPC-D POWER		TPC-D THROUGHPUT																																																													
627,123		528.6 QppD @ 100GB		213.9 QthD @ 100GB																																																													
				PRICE/PERFORMANCE																																																													
				1,865 \$ per QphD @ 100GB																																																													
DATABASE SIZE	DATABASE MANAGER	OPERATING SYSTEM	OTHER SOFTWARE	AVAILABILITY DATE																																																													
100GB	Oracle Version 8.0.3	Microsoft Windows NT 4.0 SP3	MKS Toolkit NT Resource Kit Visual C++	30-Nov-97 * see notes section																																																													
<p>Legend: □ Power Test, - - - - Geometric Mean, — Arithmetic Mean</p> <table border="1"> <caption>Approximate Query Execution Times (Seconds)</caption> <thead> <tr> <th>Query</th> <th>Arithmetic Mean (s)</th> <th>Geometric Mean (s)</th> </tr> </thead> <tbody> <tr><td>UF2</td><td>1200</td><td>1556</td></tr> <tr><td>UF1</td><td>1000</td><td>1556</td></tr> <tr><td>Q1</td><td>4800</td><td>1556</td></tr> <tr><td>Q2a</td><td>600</td><td>1556</td></tr> <tr><td>Q3</td><td>1200</td><td>1556</td></tr> <tr><td>Q4b</td><td>1800</td><td>1556</td></tr> <tr><td>Q5</td><td>3000</td><td>1556</td></tr> <tr><td>Q6</td><td>200</td><td>1556</td></tr> <tr><td>Q7</td><td>1200</td><td>1556</td></tr> <tr><td>Q8b</td><td>2200</td><td>1556</td></tr> <tr><td>Q9</td><td>4500</td><td>1556</td></tr> <tr><td>Q10</td><td>1400</td><td>1556</td></tr> <tr><td>Q11a</td><td>200</td><td>1556</td></tr> <tr><td>Q12b</td><td>4200</td><td>1556</td></tr> <tr><td>Q13</td><td>200</td><td>1556</td></tr> <tr><td>Q14c</td><td>200</td><td>1556</td></tr> <tr><td>Q15b</td><td>400</td><td>1556</td></tr> <tr><td>Q16</td><td>1000</td><td>1556</td></tr> <tr><td>Q17</td><td>100</td><td>1556</td></tr> </tbody> </table>						Query	Arithmetic Mean (s)	Geometric Mean (s)	UF2	1200	1556	UF1	1000	1556	Q1	4800	1556	Q2a	600	1556	Q3	1200	1556	Q4b	1800	1556	Q5	3000	1556	Q6	200	1556	Q7	1200	1556	Q8b	2200	1556	Q9	4500	1556	Q10	1400	1556	Q11a	200	1556	Q12b	4200	1556	Q13	200	1556	Q14c	200	1556	Q15b	400	1556	Q16	1000	1556	Q17	100	1556
Query	Arithmetic Mean (s)	Geometric Mean (s)																																																															
UF2	1200	1556																																																															
UF1	1000	1556																																																															
Q1	4800	1556																																																															
Q2a	600	1556																																																															
Q3	1200	1556																																																															
Q4b	1800	1556																																																															
Q5	3000	1556																																																															
Q6	200	1556																																																															
Q7	1200	1556																																																															
Q8b	2200	1556																																																															
Q9	4500	1556																																																															
Q10	1400	1556																																																															
Q11a	200	1556																																																															
Q12b	4200	1556																																																															
Q13	200	1556																																																															
Q14c	200	1556																																																															
Q15b	400	1556																																																															
Q16	1000	1556																																																															
Q17	100	1556																																																															
DATABASE LOAD TIME = 35:01:00		DISK SIZE/DATABASE SIZE = 3.59		RAID: N																																																													
SYSTEM COMPONENTS			DESCRIPTION																																																														
NUMBER OF NODES			Data General AViiON AV6600																																																														
PROCESSORS			6 x 200MHz PentiumPro, 512K Cache																																																														
MEMORY			2 GB																																																														
DISK DRIVES			83 x 4.3 GB and 1 x 2.1 GB																																																														
TOTAL GB OF STORAGE			359 GB																																																														

Data General Corporation	AViiON AV6600 with Oracle 8.03		TPC-D REV 1.2.3		
				Report Date: 6-Aug-97	
Description	Part Number	Unit Price	Qty	Extended Price	5 yr. Maint. Price
Server Hardware					
AV6600,3 CPU, 256MB, LAN, SCSI, CD	70706-AE	30,547	1	30,547	9,196
ADD-ON 512MB MEM (4X128MB DIMM)	7105	15,360	3	46,080	4,860
REPL 256MB WITH 512MB (4X128MB DIMM)	R7374	10,816	1	10,816	1,426
ADD-ON TRI-CPU BOARD (200MHZ/512K)	7100	12,997	1	12,997	1,497
D1600I 14" TERMINAL, WHITE, ERGONOMIC	6945W	293	1	293	357
D1200I/D1600I 101-KEYBOARD, PWR CORD	G6001A-A	56	1	56	107
60" DEEPRACK CABINET	14001-G7	1,772	1	1,772	0
2 channel SCSI controller	3944-AUWD	715	4	2,860	2,048
Subtotal:				105,421	19,491
Server Software					
ORACLE8 FOR NT	ORACLE NT	60,192	1	60,192	105,840
WIN NT EE PRE-INSTALL LMD	Q139AG21CD	3,999	1	3,999	10,560
Windows NT Resource Kit	NT RESRC KIT	499	1	499	0
MKS Toolkit for NT	MKS NT	499	1	499	0
Microsoft Visual C++ v. 4.0	VISUAL C++ 4.0	499	1	499	2,852
Subtotal:				65,688	119,252
Storage Devices					
4GB 1" HOT SWAP DISK FOR AV3600R/6600	61024-SJC	1,102	1	1,102	571
2GB 1" HOT SWAP DISK FOR AV3600R/6600	61024-SJC	777	1	777	571
HIGH PERFORMANCE ARRAY PACKAGE	79301R-E	28,574	4	114,296	26,876
NT ATTACHMENT KIT FOR CLARiiON	7945	488	1	488	0
REPLACE 8MB CACHE WITH 64MB	R7341	4,004	4	16,016	8,447
4GB 7200 RPM Disk	79012HC	1,313	82	107,666	40,461
Subtotal:				240,345	76,926
Total:				411,454	215,669
Five Year Cost of Ownership:					627,123
Audited By: Francois Raab of Information Paradigm, Inc.					
<p>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 not permitted. All discounts reflect standard pricing policies for the listed components. For complete details, see the pricing sections of the TPC benchmark specifications.</p> <p>If you find that the stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.</p> <p>AVAILABILITY NOTE: All components are available immediately, except NT 4.0 Enterprise Edition, available 30-Nov-97.</p>					

Data General Corporation	AViiON AV6600 with Oracle 8.03	TPC-D REV 1.2.3								
		Report Date:		6-Aug-97						
Numerical Quantities										
Measurement Results:										
Database Scaling (SF/Size)	=	100								
Total Data Storage / Database Size	=	3.59								
Database Load Time	=	35 hours 1 minutes 0 seconds								
Query Streams for Throughput Test	=	0								
TPC-D Power Metric (QppD@100GB)	=	528.6								
TPC-D Throughput Metric (QthD@100GB)	=	213.9								
Total System Price Over 5 Years	=	\$627,123								
TPC-D Price Performance Metric (\$/QphD@100GB)	=	\$1,865								
Measurement Intervals										
Measurement interval in Throughput Test (TS)	=	28609 seconds								
Duration of Stream Execution:										
Stream ID	Seed	Start Date	Start Time	End Date	End Time	Total Time				
UF1		7/31/97	11:07:21	7/31/97	11:22:34	0:15:13				
Stream 00	857038079	7/31/97	11:07:21	7/31/97	19:04:15	7:56:54				
UF2		7/31/97	18:43:37	7/31/97	19:04:15	0:20:38				
Timing Intervals (in Seconds)										
Query	Q1	Q2a	Q3	Q4b	Q5	Q6	Q7	Q8b	Q9	
Stream 00	4977.8	612.5	1210.1	1769.5	2950.5	198.6	1269.9	2185.9	4506.5	
Query	Q10	Q11a	Q12b	Q13	Q14c	Q15b	Q16	Q17	UF1	UF2
Stream 00	1398.5	186.4	4086.0	7.0	162.0	293.7	585.6	60.5	908.1	1237.5

Table of Contents

General Items	1
Test Sponsor.....	1
Parameter Settings	1
Configuration Items.....	1
Clause 1: Logical Database Design	3
Table Definitions	3
Physical Organization of Database.....	3
Horizontal Partitioning.....	3
Replication	3
Clause 2: Queries and Update Functions	4
Query Language	4
Random Number Generation.....	4
Substitution Parameters Generation.....	4
Query Text and Output Data from Database	4
Query Substitution Parameters and Seeds Used	5
Update Function Source Code.....	5
Database Maintenance Option.....	5
Clause 3: Database System Properties	6
Atomicity.....	6
Completed Transaction	6
Aborted Transaction	6
Consistency.....	6
Consistency Test.....	6
Isolation.....	6
Read-Write Conflict with Commit	7
Read-Write Conflict with Rollback	7
Write-Write Conflict with Commit	7
Write-Write Conflict with Rollback	7
Durability	7
Failure of a Durable Medium	7
System Crash.....	8
Memory Failure.....	8
Clause 4: Scaling and Database Population	9
Initial Cardinality of Tables.....	9
Distribution of Tables and Logs Across Media.....	9
Partitioning and Replication	10
DBGEN Version and Modifications	10
Database Content of the First Ten Rows.....	10
Database Load time	11
Data Storage Ratio.....	11

Database Load Mechanism Details and Illustration.....	11
Clause 5: Performance Metrics and Execution Rules	13
Steps in the Power Test.....	13
Timing Intervals.....	13
Number of Streams for The Throughput Test.....	13
Start/Finish Time of Each Query Stream	13
Total Elapsed Time	13
Start/Finish Time for Update Function	13
Timing Intervals for Each Query and Each Update.....	13
Performance Metrics	14
Reproducibility Method.....	14
Clause 6: SUT and Driver Implementation	15
Driver.....	15
Implementation Specific Layer (ISL)	15
Clause 7: Pricing.....	16
Hardware and Software Used in the Priced System.....	16
Total Five year Price	16
Availability Date.....	16
Clause 8: Auditor-Related Items.....	17
Auditor’s Report	17
Appendix A: Parameter Settings	A1
System Configuration Settings.....	1
Oracle Parameter Settings.....	1
Appendix B: Programs and Scripts	B1
Appendix C: Query Text and Query Output.....	C1
Appendix D: Seed and Query Substitution Parameters	D1
Appendix E: Implementation Specific Layer/Source Code	E1
qexecpl.c.....	
qexecpl.h.....	
gettime.c	
uf1.pc	
uf2_del.pc.....	
runuf1.sh	
runuf2.sh	
runpower1.sh	
Appendix F: Initial Ten Rows.....	F1
Appendix G: Auditor’s Attestation Letter	G1

General Items

Test Sponsor

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

Data General Corporation is the sponsor of this TPC Benchmark™ D.

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 but not limited to:

- *Database Tuning Options*
- *Optimizer/Query execution options*
- *Query processing tool/language configuration parameters*
- *Recovery/commit options*
- *Consistency/locking options*
- *Operating system and configuration parameter*
- *Configuration parameters and options for any other software in the pricing structure*
- *Compiler optimization options*

This requirement can be satisfied by providing a full list of all parameters and options, as long as all those which have been modified from their default values have been clearly identified and these parameters and options are only set once.

Details of system and database configurations and parameters are provided in Appendix A.

Configuration Items

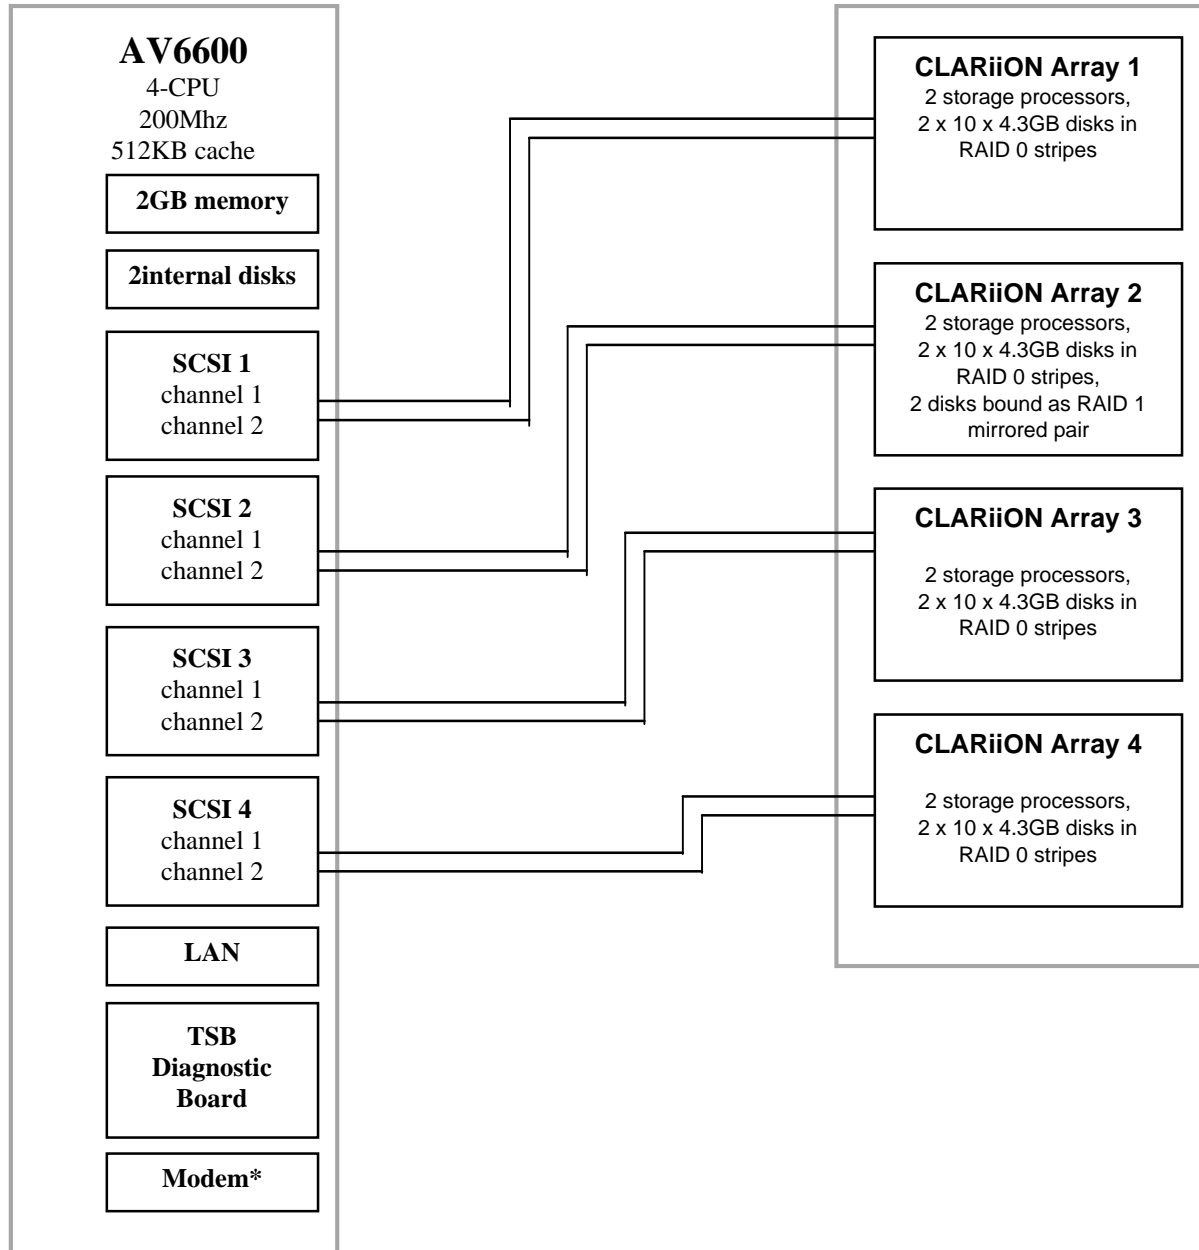
Diagrams of both measured and priced configurations must be provided, accompanied by a description of the differences. This includes, but is not limited to:

- *Number and type of processors*
- *Size of allocated memory, and any specific mapping/partitioning of memory in the test.*
- *Number and type of disk units (and controllers, if applicable).*
- *Number of channels or bus connections to disk units, including their protocol type.*
- *Number of LAN (e.g. Ethernet) Connections, including routers, workstations, terminals, etc., that were physically used in the test or are incorporated into the pricing structure.*
- *Type and the run-time execution location of software components (e.g., DBMS, query processing tools/languages, middle-ware components, software drivers, etc.).*

The server System Under Test (SUT), a Data General AViiON AV6600, depicted in the next diagram, consists of:

- (6) 200MHz PentiumPro® Processors, each with 512KB of L2 cache.
- 2GB of main memory.
- 4 SCSI dual controllers (plus an internal SCSI adapter).
- 1 Ethernet controller.
- 2 internal disk drives.
- 4 CLARiiON Disk Arrays, each with two storage processors; 3 arrays with 20 disk drives and one with 22 disk drives.

The system's disk I/O configuration consists of 9 SCSI buses. The internal SCSI adapter controls the 2 internal disks. Each of the other 8 SCSI channels (2 on each of 4 adapters) are connected to 1 Storage Processor controlling a 10 disk RAID 0 stripe; except for 1 channel, which in addition controls a pair of mirrored RAID 1 disks.



Measured and Priced Configuration

* The only difference between the measured and priced systems is that the measured system did not have a modem installed.

Clause 1: Logical Database Design

Table Definitions

Listings must be provided for all table definition statements and all other statements used to set up the test and qualification databases.

Appendix B contains the scripts that create and analyze the tables and indexes for the TPC-D database.

Physical Organization of Database

The physical organization of tables and indices, within the test and qualification databases, must be disclosed. If the column ordering of any table is different from that specified in Clause 1.4, it must be noted.

No record clustering or index clustering was used for this benchmark.

Horizontal Partitioning

Horizontal partitioning of tables and rows in the test and qualification databases (see Clause 1.5.4) must be disclosed.

The database layout is described in the tables in Clause 4.

Replication

Any replication of physical objects must be disclosed and must conform to the requirements of Clause 1.5.6.

Replication was not used for this benchmark.

Clause 2: Queries and Update Functions

Query Language

The query language used to implement the queries must be identified.

SQL was the query language used to implement all queries.

Random Number Generation

The method of verification for the random number generation must be described unless the supplied DBGEN and QGEN were used.

The 1.2.0 version of DBGEN and version 1.2.0 of QGEN were used to generate the random numbers for this TPC-D benchmark.

Substitution Parameters Generation

The method used to generate values for substitution parameters must be disclosed. If QGEN is not used for this purpose, then the source code of any non-commercial tool used must be disclosed. If QGEN is used, the version number, release number, modification number and patch level of QGEN must be disclosed.

The 1.2.0 version of QGEN was used to generate the substitution parameters.

Query Text and Output Data from Database

The executable query text used for query validation must be disclosed along with the corresponding output data generated during the execution of the query text against the qualification database. If minor modifications (see Clause 2.2.3) have been applied to any functional query definitions or approved variants in order to obtain executable query text, these modifications must be disclosed and justified. The justification for a particular minor query modification can apply collectively to all queries for which it has been used. The output data for the power and throughput tests must be made available electronically upon request.

Appendix C contains the query text and query output. The minor query modifications used in this implementation include the following:

- Wherever there is a date given in the WHERE clause of a select the TO_DATE function is used to convert the character input to the proper date format.
- Wherever there is a date returned by the query the TO_CHAR function is used to convert the date format to the proper output format.
- The Order table defined in the specification has been named the orders table due to conflict with a reserved word.
- Whenever there is a time modification (i.e. add 3 months to a date) the ADD_MONTHS function is used to calculate and adjust the time delta and keep it in the date/time format.
- CREATE TABLE statements have specific storage parameters in them. Namely, what tablespace to create them on, the sizes of the extents, and how full the extents should be before allocating a new one. The degree of parallelism was not specified for any temporary tables created.

Query Substitution Parameters and Seeds Used

All the query substitution parameters used during the performance test must be disclosed in tabular format, along with the seeds used to generate these parameters.

Appendix D contains the seed and query substitution parameters.

Update Function Source Code

The details of how the update functions were implemented must be disclosed (including source code of any non-commercial program used).

The update function is part of the implementation-specific driver code included in Appendix E.

Database Maintenance Option

The details of the database maintenance option selected (i.e., reset or evolve) must be disclosed (including source code of any non-commercial program used).

This implementation of the TPC-D benchmark uses the reset option.

Clause 3: Database System Properties

Atomicity

The results of the ACID tests must be disclosed along with a description of how the ACID requirements were met. This includes disclosing the code written to implement the ACID Transaction and Query.

Completed Transaction

Perform the ACID Transaction for a randomly selected set of input data and verify that the appropriate rows have been changed in the ORDER, LINEITEM, and HISTORY tables.

1. The total prices from the ORDER table and the extended price from the LINEITEM table were retrieved for a randomly selected order key.
2. The ACID Transaction was performed using the order key from step 1.
3. The ACID Transaction was committed.
4. The total price from the ORDER table and the extended price from the LINEITEM table were retrieved for the same order key. It was verified that the appropriate rows had been changed.

Aborted Transaction

Perform the ACID transaction for a randomly selected set of input data, substituting a ROLLBACK of the transaction for the COMMIT of the transaction. Verify that the appropriate rows have not been changed in the ORDER, LINEITEM, and HISTORY tables.

1. The total price from the ORDER table and the extended price from the LINEITEM table were retrieved for a randomly selected order key.
2. The ACID Transaction was performed using the order key from step 1. The transaction was stopped prior to the commit.
3. The ACID Transaction was rolled back.
4. The total price from the ORDER table and the extended price from the LINEITEM table were retrieved for the same order key, and were verified to have not been changed.

Consistency

Consistency is the property of the application that requires any execution of transactions to take the database from one consistent state to another.

Consistency Test

Verify that ORDER and LINEITEM tables are initially consistent, submit the prescribed number of ACID Transactions with randomly selected input parameters, and re-verify the consistency of the ORDER and LINEITEM tables.

1. The consistency of the ORDER and LINEITEM tables was verified based on a sample of O_ORDERKEYs.
2. 100 ACID transactions were submitted from each of 2 execution streams.
3. The consistency of the ORDER and LINEITEM tables was verified a second time with the same O_ORDERKEYs.

Isolation

Operations of concurrent transactions must yield results which are indistinguishable from the results which would be obtained by forcing each transaction to be serially executed to completion in some order.

Read-Write Conflict with Commit

Demonstrate isolation for the read-write conflict of a read-write transaction and a read-only transaction when the read-write transaction is committed.

1. An ACID transaction was started for a randomly selected O_KEY, L_KEY, and DELTA. The ACID transaction was suspended prior to commit.
2. An ACID query was started for the same O_KEY used in step 1. The ACID query completed and did not see the uncommitted changes made by the ACID transaction.
3. The ACID transaction was committed.

Read-Write Conflict with Rollback

Demonstrate isolation for the read-write conflict of a read-write transaction and a read-only transaction when the read-write transaction is rolled back.

1. An ACID transaction was started for a randomly selected O_KEY, L_KEY, and DELTA. The ACID transaction was suspended prior to rollback.
2. An ACID query was started for the same O_KEY used in step 1. The ACID query completed and did not see the uncommitted changes made by the ACID transaction.
3. The ACID transaction was rolled back.

Write-Write Conflict with Commit

Demonstrate isolation for the write-write conflict of two update transactions when the first transaction is committed.

1. An ACID transaction, T1, was started for a randomly selected O_KEY, L_KEY, and DELTA. The ACID transaction was suspended prior to commit.
2. A second ACID transaction, T2, was started using the same O_KEY and L_KEY and a different randomly selected DELTA.
3. T2 waited.
4. T1 was allowed to commit and then T2 completed.
5. It was verified that T2.L_EXTENDPRICE was calculated correctly.

Write-Write Conflict with Rollback

Demonstrate isolation for the write-write conflict of two update transactions when the first transaction is rolled back.

1. An ACID transaction, T1, was started for a randomly selected O_KEY, L_KEY, and DELTA. The ACID transaction was suspended prior to rollback.
2. A second ACID transaction, T2, was started using the same O_KEY and L_KEY and a different randomly selected DELTA.
3. T2 waited.
4. T1 was allowed to rollback and then T2 completed.
5. It was verified that T2.L_EXTENDPRICE was calculated correctly.

Durability

The tested system must guarantee durability: the ability to preserve the effects of committed transactions and insure database consistency after recovery from any one of the failures listed in Clause 3.5.2.

Failure of a Durable Medium

Guarantee the database and committed updates are preserved across a permanent irrecoverable failure of any single durable medium containing TPC-D database tables or recovery log tables.

The database logs were stored on a hardware mirrored pair (2 disks running RAID 1).

The tables for the database were stored on RAID-0 stripes.

1. The datafiles were backed up to an alternate disk media.
2. Two streams of ACID transactions were started.
3. While the test was running one side of the mirrored set of logs was disabled.
4. After it was determined that the test would still run with the loss of a log disk, a data disk was disabled.
5. The two streams of ACID transactions failed and recorded their numbers of committed transactions in success files.
6. The database was brought down.
7. The datafiles were restored to their state prior to the ACID transaction streams.
8. The database ran through its recovery mode.
9. The counts in the success files and the HISTORY table count were compared. The counts matched.

System Crash

Guarantee the database and committed updates are preserved across an instantaneous interruption (system crash/system hang) in processing which requires the system to reboot to recover.

1. Two streams of ACID transactions were started.
2. While the streams of ACID transactions were running the system was powered off.
3. When power was restored the system rebooted and the database was restarted.
4. The database went through a recovery period.
5. The success file and the HISTORY table counts were compared, and they matched.

Memory Failure

Guarantee the database and committed updates are preserved across failure of all or part of memory (loss of contents).

The system crash test and the memory failure test were combined. See the previous section.

Clause 4: Scaling and Database Population

Initial Cardinality of Tables

The cardinality (e.g., the number of rows) of each table of the test database, as it existed at the completion of the database load (see clause 4.2.5) must be disclosed.

Initial number of rows

Table	Occurrences
Orders	150,000,000
Lineitem	600,037,902
Customer	15,000,000
Part	20,000,000
Supplier	1,000,000
Partsupp	80,000,000
Nation	25
Region	5

Distribution of Tables and Logs Across Media

The distribution of tables and logs across all media must be explicitly described.

The database tables were distributed across 4 CLARiiON disk arrays. Each array contained two storage processors. The arrays were configured as follows:

- Database files were stored on 8 RAID 0 stripes, each consisting of (10) 4.3GB disks.
- Log information was stored on a mirrored pair, using (2) 4.3GB disks.
- Flat files were stored on (25) 4.3GB disks and (10) 2.1GB disks, which were physically removed before measurement.
- Database binaries were stored on one 4.3GB disk.
- The operating system was stored on one 2.1GB internal disk internal.

In the following tables, each row in the Controller column refers to one of two storage processors in each of the 4 CLARiiON disk arrays.

Distribution of Storage

Controller	Disk Drive	Storage Content
1	1-10	1 RAID 0 stripe for the database datafiles
2	11-20	1 RAID 0 stripe for the database datafiles
3	1-10	1 RAID 0 stripe for the database datafiles
4	11-20	1 RAID 0 stripe for the database datafiles
4	21,22	Mirrored pair containing one log
5	1-10	1 RAID 0 stripe for the database datafiles

6	11-20	1 RAID 0 stripe for the database datafiles
7	1-10	1 RAID 0 stripe for the database datafiles
8	11-20	1 RAID 0 stripe for the database datafiles

Partitioning and Replication

The mapping of database partitions/replications must be explicitly described. Implementations may use some form of RAID to ensure high availability. If used for data, auxiliary storage (e.g. indexes) or temporary space, the level of RAID used must be disclosed for each device.

No replication was used for this implementation.

Distribution of Partitioned Tables

Controller	Configuration	Usage	Fraction
1	RAID 0	lineitem, orders, parts, partsupp, supplier, nation, and region	1/8 th of each of these tables except nation and region 100%
2	RAID 0	lineitem, orders, parts, partsupp, supplier	1/8 th of each of these tables
3	RAID 0	lineitem, orders, parts, partsupp, supplier	1/8 th of each of these tables
4	RAID 0	lineitem, orders, parts, partsupp, supplier	1/8 th of each of these tables
4	Mirrored pair	Log	100%
5	RAID 0	lineitem, orders, parts, partsupp, supplier	1/8 th of each of these tables
6	RAID 0	lineitem, orders, parts, partsupp, supplier	1/8 th of each of these tables
7	RAID 0	lineitem, orders, parts, partsupp, supplier	1/8 th of each of these tables
8	RAID 0	lineitem, orders, parts, partsupp, supplier	1/8 th of each of these tables

DBGEN Version and Modifications

The version number, release number, modification number, and patch level of DBGEN must be disclosed. Any modifications to the DBGEN (see Clause 4.2.1) source code....must be disclosed. In the event that a program other than DBGEN was used to populate the database, it must be disclosed in its entirety.

The supplied DBGEN Version 1.2.0 was used for database population.

Database Content of the First Ten Rows

The content of the first ten rows of each table in the test database must be disclosed.

Appendix F contains the first ten rows of each table in the test database.

Database Load time

The database load time for the test database (see clause 4.3) must be disclosed.

Database load time was 35 hours 1 minute 0 seconds.

Data Storage Ratio

The data storage ratio must be disclosed. It is computed by dividing the total data storage of the priced configuration (expressed in GB) by the size chosen for the test database as defined in 4.1.3.1. The ratio must be reported to the nearest 1/100, rounded up.

Data Storage Ratio

Disk Type	Number of Disks	Space per Disk	Subtotal Disk Space
4.3GB	83	4.3GB	356.9GB
2.1GB	1	2.1GB	2.1GB

Total disk storage: 359.0GB

Data storage ratio: 3.59

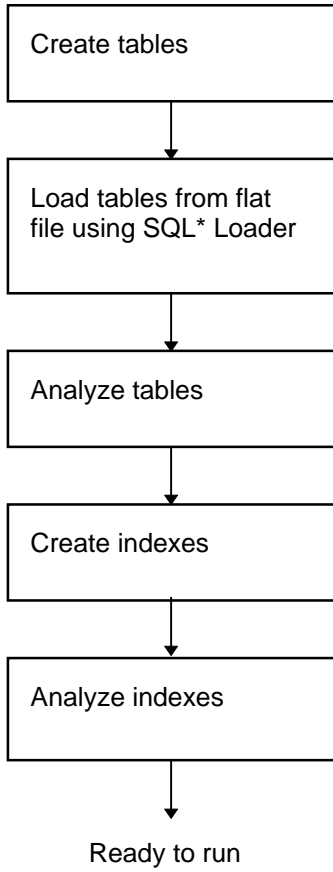
Database Load Mechanism Details and Illustration

The details of the database load must be disclosed, including a block diagram illustrating the overall process. Disclosure of the load procedure includes all steps, scripts, input and configuration files required to completely reproduce the test and qualification databases.

DBGEN was used to create flat files which were then loaded into the tablespaces using SQL*Loader. The flat files were split into as many as 160 files per table (as documented in the load scripts in the appendix), except the REGION and NATION tables which each had only one file.

PERL was used to perform the creation of the tablespaces, tables, and indexes of the database. It was also used to run the analyze scripts. The controller cache was set to 0MB write cache and 60 MB read cache during the database load and the power runs. This was done manually, through running of the CLARiiON Dass Manager.

After the load phase was completed, the flat file disks were physically removed from the system.



Database Load Process

Clause 5: Performance Metrics and Execution Rules

Steps in the Power Test

The details of the steps followed to implement the power test (e.g., system boot, database restart, etc.) must be disclosed.

The following steps were used to implement the power test:

1. Database restart.
2. UF1 update transaction.
3. Stream 00 execution.
4. UF2 update transaction.

Timing Intervals

The timing intervals (see Clause 5.3.6) for each query of the measured set and for both update functions must be reported for the power test.

The power test timing intervals are disclosed in the Numerical Quantities Summary earlier in this document.

Number of Streams for The Throughput Test

The number of execution streams used for the throughput test must be disclosed.

No separate throughput test was conducted. Therefore, the values of power test were used for the throughput test.

Start/Finish Time of Each Query Stream

The start time and finish time for each query execution stream must be reported for the throughput test.

The throughput test start time and finish time for each stream are disclosed in the Numerical Quantities Summary earlier in this document.

Total Elapsed Time

The total elapsed time of the measurement interval must be reported for the throughput test.

The total elapsed time of the throughput test was 28609 seconds.

Start/Finish Time for Update Function

Start and finish time for each update function in the update stream must be reported for the throughput test.

The start and finish time for each update function in the update stream are disclosed in the Numerical Quantities Summary earlier in this document.

Timing Intervals for Each Query and Each Update

The timing intervals (see Clause 5.3.6) for each query of each stream and for each update function must be reported for the throughput test.

The timing intervals for each query and each update function are contained in the Numerical Quantities Summary disclosed earlier in this document. Since no separate throughput test was conducted, the values for the power test are used for the throughput test.

Performance Metrics

The computed performance metrics, related numerical quantities and the price performance metric must be reported.

The performance metrics, and the numbers on which they are based, are contained in the Numerical Quantities section of the Executive Summary.

Reproducibility Method

A description of the method used to determine the reproducibility of the measurement results must be reported. This must include the performance metrics (QppD and QthD) from the reproducibility runs.

Performance results from the first two executions of the TPC-D benchmark indicated the following percent differences for the metrics:

Percentage Differences in Benchmark Executions

Run	QppD@100GB	QthD@100GB	QphD@100GB
1	528.6	213.9	336.3
2	531.0	213.5	336.7
Percent Difference	0.45%	0.19%	0.12%

Clause 6: SUT and Driver Implementation

Driver

A detailed description of how the driver performs its functions must be supplied, including any related source code or scripts. This description should allow an independent reconstruction of the driver.

QGEN is first called with a stream ID of 0 to generate the QET for the power test. The power test is performed by a shell script called **runpower1.sh**. This shell script calls **runuf1.sh**, which performs the update function. Queries are then executed with the **qexec.ott** ISL program. The last part of the power test is run through another shell script called **runuf2.sh**. Both wall-clock and high-resolution times are collected for the measurement intervals.

For the update functions multiple sqlldr processes loaded data into temporary tables. ORACLE8's parallel insert and delete functionality was used to perform the Update Functions, selecting data from the temporary tables.

Implementation Specific Layer (ISL)

If an implementation-specific layer is used, then a detailed description of how it performs its functions must be supplied, including any related source code or scripts. This description should allow an independent reconstruction of the implementation-specific layer.

Query execution text generated by QGEN is picked up by the ISL program, **qexec.ott**, which submits the query to the SUT. The ISL program utilize the Oracle Call Interface (OCI) to communicate with the Oracle database on the SUT. QETs directly generated by QGEN are read and submitted to the SUT via the ISL program as dynamic SQL statements. The ISL program then fetches the query execution output and reports it to the user. Timings are taken at intervals specified by Clause 5.3.6.2.

Clause 7: Pricing

Hardware and Software Used in the Priced System

A detailed list of hardware and software used in the priced system must be reported. Each 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, contents of the package must be disclosed. Pricing source(s) and effective date(s) of price(s) must also be reported.

A detailed list of the hardware and software used in the priced system is included in the executive summary at the beginning of this document.

Total Five year Price

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.

A detailed price sheet of all the hardware and software used in this configuration, including the 5-year maintenance cost, and total price, is included in the executive summary at the beginning of this document. All maintenance has been uplifted from 8x40 to 7x24. All AViiON hardware comes with a 1 year 48-hour response warrantee, which has been uplifted to 7x24 4-hour response. Software comes with a 90 day telephone warantee, which has been uplifted to 7x24 support with updates. Oracle is discounted 40%, and the maintenance 10% off the base prices.

Availability Date

The committed delivery date for general availability of products used in the price calculations must be reported. When the priced system includes products with different availability dates, the availability date reported on the executive summary must be the date by which all components are committed to being available. The full disclosure report must report availability dates individually for at least each of the categories for which a pricing subtotal must be provided.

All components are available as of the date reported in the executive summary.

Clause 8: Auditor-Related Items

Auditor's Report

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.

The auditor's letter of attestation is included in this full disclosure report as Appendix G.

This implementation was audited by: Francois Raab, Information Paradigm, TPC Certified D Auditor.

Further information regarding the audit process may be obtained from:

Information Paradigm Inc.
1373 North Franklin Street
Colorado Springs, CO 80903
719/473-7555

Requests for this TPC Benchmark D Full Disclosure Report should be sent to:
Transaction Processing Performance Council
c/o Shanley Public Relations
777 North First Street, Suite 6000
San Jose, CA 95112-6311
408/295-8894

Appendix A: Parameter Settings

System Configuration Settings

The following NT services were started:

- Alerter
- EventLog
- License Logging Service
- Messenger
- Net Logon
- OracleServiceTPCD
- Plug and Play
- Remote Procedure Call (RPC) Locator
- Remote Procedure Call (RPC) Service
- Schedule
- Server
- Simple TCP/IP Services
- Spooler
- TCP/IP NetBIOS Helper
- Workstation

The default NT registry settings were used, with the exception of tuning 1 parameter. This parameter deals with the SCSI disk controller concurrency. It was changed from the default of 16 to 64 (0x40 hexadecimal). This can be changed by regedit under HKEY_LOCAL_MACHINE/SYSTEM/CurrentControlSet/Services/aic78xx/Parameters/Device, setting the REG_DWORD 'NumberOfRequests' to 0x40 hexadecimal.

Oracle Parameter Settings

Appendix A

```
P_BUILD.ORA
# p_build.ora
enqueue_resources      = 15000
db_files               = 10000
sort_direct_writes    = true
# for creation and population
sort_area_size        = 5000000
sort_area_retained_size = 5000000
hash_area_size        = 45000000
always_anti_join      = hash
compatible             = 8.0.3
db_block_lru_latches  = 64
hash_multiblock_io_count= 8
hash_join_enabled     = true
sort_read_fac         = 15
log_checkpoint_interval = 2000000000
log_buffer            = 5242880
log_checkpoints_to_alert= true
parallel_min_servers  = 80
parallel_max_servers  = 80
optimizer_search_limit = 6
_advanced_dss_features = true
cleanup_rollback_entries= 100
max_dump_file_size    = 4000000
parallel_min_servers  = 80
parallel_max_servers  = 80
processes             = 160
remote_login_passwordfile = shared
dml_locks             = 2000
shared_pool_size      = 100000000
timed_statistics      = FALSE
transactions_per_rollback_segment = 2
background_dump_dest  = d:\tpcd\log\hundgig
user_dump_dest        = d:\tpcd\log\hundgig

P_INDEX.ORA
# p_index.ora
enqueue_resources      = 15000
db_files               = 10000
sort_direct_writes    = true
# for execution
sort_area_size        = 5000000
sort_area_retained_size = 5000000
hash_area_size        = 50000000
always_anti_join      = hash
compatible             = 8.0.3
db_block_lru_latches  = 64
hash_multiblock_io_count= 8
hash_join_enabled     = true
sort_read_fac         = 15
log_checkpoint_interval = 2000000000
log_buffer            = 5242880
log_checkpoints_to_alert= true
parallel_min_servers  = 80
parallel_max_servers  = 80
optimizer_search_limit = 6
_advanced_dss_features = true
cleanup_rollback_entries= 100
max_dump_file_size    = 4000000
db_file_multiblock_read_count = 32
dml_locks             = 2000
processes             = 160
remote_login_passwordfile = shared
rollback_segments=(r1,r2,r3,r4,r5,r6,r7,r8,r9,r10,\
r11,r12,r13,r14,r15,r16,r17,r18,r19,r20,\
r21,r22,r23,r24,r25,r26,r27,r28,r29,r30,\
r31,r32,r33,r34,r35,r36,r37,r38,r39,r40,\
r41,r42,r43,r44,r45,r46,r47,r48,r49,r50,\
r51,r52,r53,r54,r55,r56,r57,r58,r59,r60,\
r61,r62,r63,r64,r65,r66,r67,r68,r69,r70,\
r71,r72,r73,r74,r75,r76,r77,r78,r79,r80,\
r81,r82)
shared_pool_size      = 100000000
timed_statistics      = FALSE
transactions_per_rollback_segment = 2
background_dump_dest  = d:\tpcd\log\hundgig
user_dump_dest        = d:\tpcd\log\hundgig

INIT1_TPCD.ORA
#init1_tpcd.ora
```

Appendix A

```
# Init.ora for instance 1 created by bumpx.pl
thread = 1
ifile = d:\tpcd\dfs\hundgig\init_tpcd.ora

INIT_TPCD.ORA
# init_tpcd.ora
# Init.ora include file created by bumpx.pl
db_name = tpcd
control_files = d:\tpcd\dfs\hundgig\cntrltpcd.dbf
db_block_buffers = 1000
shared_pool_size = 35000000
parallel_max_servers = 144
parallel_min_servers = 0
max_dump_file_size = 5000
audit_trail = FALSE
global_names = FALSE
db_block_size = 16384
db_file_multiblock_read_count = 8
processes = 256
sessions = 256
transactions = 512
transactions_per_rollback_segment = 20
max_rollback_segments = 256
distributed_transactions = 0
nls_date_format = YYYY-MM-DD
db_files = 1023
open_cursors = 1024
sort_direct_writes = auto
sort_write_buffer_size = 65536
sort_write_buffers = 8
optimizer_mode = CHOOSE
optimizer_percent_parallel = 100
ifile = d:\tpcd\dfs\hundgig\tkinit.ora

TKINIT.ORA
enqueue_resources = 15000
db_block_buffers = 6000
db_files = 10000
sort_direct_writes = true
always_anti_join = hash
compatible = 8.0.3
db_block_lru_latches = 64
hash_multiblock_io_count = 8
hash_join_enabled = true
sort_read_fac = 15
log_checkpoint_interval = 2000000000
log_buffer = 5242880
log_checkpoints_to_alert = true
parallel_min_servers = 80
parallel_max_servers = 80
optimizer_search_limit = 6
_advanced_dss_features = true
_db_block_write_batch = 1024
cleanup_rollback_entries = 100
max_dump_file_size = 4000000
db_file_multiblock_read_count = 32
dml_locks = 2000
processes = 160
remote_login_passwordfile = shared
rollback_segments=(r1,r2,r3,r4,r5,r6,r7,r8,r9,r10,\
r11,r12,r13,r14,r15,r16,r17,r18,r19,r20,\
r21,r22,r23,r24,r25,r26,r27,r28,r29,r30,\
r31,r32,r33,r34,r35,r36,r37,r38,r39,r40,\
r41,r42,r43,r44,r45,r46,r47,r48,r49,r50,\
r51,r52,r53,r54,r55,r56,r57,r58,r59,r60,\
r61,r62,r63,r64,r65,r66,r67,r68,r69,r70,\
r71,r72,r73,r74,r75,r76,r77,r78,r79,r80,\
r81,r82)
hash_area_size = 15768240
sort_area_size = 1048576
sort_area_retained_size = 1048576
shared_pool_size = 100000000
timed_statistics = FALSE
transactions_per_rollback_segment = 2
background_dump_dest = d:\tpcd\log\hundgig
user_dump_dest = d:\tpcd\log\hundgig
```

Appendix B: Programs and Scripts

[NOTE: Some of the scripts are very long, and, in order to save space, are printed in very small font. The electronic version of the FDR has the complete listings. But all of the scripts may not be readable in the paper version.]

TPCD100_8_DG.CONF

```
# Configuration file for bumpx.p1

*matchon
#####
# preprocessing-like directives

%b-preproc

*sql
\svrmgr30 <<!
\set echo on;
\set termout on;
\connect internal;
\select to_char(sysdate, 'MM-DD-YYYY HH24:MI:SS') now from dual;
\}
\select to_char(sysdate, 'MM-DD-YYYY HH24:MI:SS') now from dual;
\exit;
\!

*load
\sqlldr80 {}

*mknod
\mknod {}

*dbgen
\dbgen {}

*sh
\{}

%-preproc
*matchoff
#####

# basic parameters

scale_factor = 100
dd_sql_area = d:\orant\rdbms80\admin\
dbs_area = d:\tpcd\dbs\hundgig\
max_bg = 64
sccre_max_bg = 40
dpop_max_bg = 40
dbgen_max_bg = 20
ixcre_max_bg = 1
anlyz_max_bg = 40
user=tpcd
passwd=tpcd
compatible=8.0

io_ifile=d:\tpcd\dbs\hundgig\tkinit.ora
io_control_files=d:\tpcd\dbs\hundgig\cntrl\tpcd.dbf
#skip_mk_initoras=include
#skip_mk_ldctlf=linetemp,orders,customer,parts,partsupp,supplier,nati
on,region

# create database parameters

db_maxdatafiles=10000

# Basic tablespace related parameters

skip_ts=ts_data

# Need to change the following base subdirectory for ecsun4
ts_def_area = g:\tpcd\dbs\hundgig\
ts_def_#files = 1

ts_sys_first_size=1000m
ts_sys_size=1000m
ts_sys_datafiles=sys.dbf
ts_sys_area=d:\tpcd\dbs\hundgig\

ts_log_first_size=1000m
ts_log_size=1000m
#ts_log_datafiles=log1.dbf,log2.dbf
#ts_log_area=f:\tpcd\dbs\hundgig\
ts_log_area=none
ts_log_datafiles=\\.\0:;\\.\P:

## Changed from original 1000m to make
## parallel DMLs (update functions) work.
##
ts_und0_#files=9
ts_und0_first_size=1000m
ts_und0_size=300m
ts_und0_datafiles=ts_und0_#
ts_und0_area=d:\tpcd\dbs\hundgig\
ts_und0_areas=g:\tpcd\dbs\hundgig\,h:\tpcd\dbs\hundgig\,&
i:\tpcd\dbs\hundgig\,j:\tpcd\dbs\hundgig\,&
k:\tpcd\dbs\hundgig\,l:\tpcd\dbs\hundgig\,&
m:\tpcd\dbs\hundgig\,n:\tpcd\dbs\hundgig\
ts_und0_#rs=83
```

```
ts_und0_rs_storage=(initial 10m next 2048k maxextents unlimited)

#ts_data_names=ts_ps,ts_p,ts_s
ts_data_names=ts_ps,ts_s
#ts_data_groups=ts_1,ts_o,ts_c
ts_data_groups=ts_1,ts_o,ts_c,ts_p

#ts_index_names=ts_ind1,ts_ind2,ts_ind3,ts_ind4,ts_ind5
ts_index_names=ts_ind1,ts_ind3,ts_ind4,ts_ind5

ts_1_group_#ts=80
ts_1_group_#files=2
ts_1_group_first_size=461m
ts_1_group_size=461m
ts_1_group_storage=(initial 64k next 460m pctincrease 0)
ts_1_area=g:\tpcd\dbs\hundgig\
ts_12_area=h:\tpcd\dbs\hundgig\
ts_13_area=i:\tpcd\dbs\hundgig\
ts_14_area=j:\tpcd\dbs\hundgig\
ts_15_area=k:\tpcd\dbs\hundgig\
ts_16_area=l:\tpcd\dbs\hundgig\
ts_17_area=m:\tpcd\dbs\hundgig\
ts_18_area=n:\tpcd\dbs\hundgig\
ts_19_area=g:\tpcd\dbs\hundgig\
ts_10_area=h:\tpcd\dbs\hundgig\
ts_11_area=i:\tpcd\dbs\hundgig\
ts_12_area=j:\tpcd\dbs\hundgig\
ts_13_area=k:\tpcd\dbs\hundgig\
ts_14_area=l:\tpcd\dbs\hundgig\
ts_15_area=m:\tpcd\dbs\hundgig\
ts_16_area=n:\tpcd\dbs\hundgig\
ts_17_area=g:\tpcd\dbs\hundgig\
ts_18_area=h:\tpcd\dbs\hundgig\
ts_19_area=i:\tpcd\dbs\hundgig\
ts_20_area=j:\tpcd\dbs\hundgig\
ts_21_area=k:\tpcd\dbs\hundgig\
ts_22_area=l:\tpcd\dbs\hundgig\
ts_23_area=m:\tpcd\dbs\hundgig\
ts_24_area=n:\tpcd\dbs\hundgig\
ts_25_area=g:\tpcd\dbs\hundgig\
ts_26_area=h:\tpcd\dbs\hundgig\
ts_27_area=i:\tpcd\dbs\hundgig\
ts_28_area=j:\tpcd\dbs\hundgig\
ts_29_area=k:\tpcd\dbs\hundgig\
ts_30_area=l:\tpcd\dbs\hundgig\
ts_31_area=m:\tpcd\dbs\hundgig\
ts_32_area=n:\tpcd\dbs\hundgig\
ts_33_area=g:\tpcd\dbs\hundgig\
ts_34_area=h:\tpcd\dbs\hundgig\
ts_35_area=i:\tpcd\dbs\hundgig\
ts_36_area=j:\tpcd\dbs\hundgig\
ts_37_area=k:\tpcd\dbs\hundgig\
ts_38_area=l:\tpcd\dbs\hundgig\
ts_39_area=m:\tpcd\dbs\hundgig\
ts_40_area=n:\tpcd\dbs\hundgig\
ts_41_area=g:\tpcd\dbs\hundgig\
ts_42_area=h:\tpcd\dbs\hundgig\
ts_43_area=i:\tpcd\dbs\hundgig\
ts_44_area=j:\tpcd\dbs\hundgig\
ts_45_area=k:\tpcd\dbs\hundgig\
ts_46_area=l:\tpcd\dbs\hundgig\
ts_47_area=m:\tpcd\dbs\hundgig\
ts_48_area=n:\tpcd\dbs\hundgig\
ts_49_area=g:\tpcd\dbs\hundgig\
ts_50_area=h:\tpcd\dbs\hundgig\
ts_51_area=i:\tpcd\dbs\hundgig\
ts_52_area=j:\tpcd\dbs\hundgig\
ts_53_area=k:\tpcd\dbs\hundgig\
ts_54_area=l:\tpcd\dbs\hundgig\
ts_55_area=m:\tpcd\dbs\hundgig\
ts_56_area=n:\tpcd\dbs\hundgig\
ts_57_area=g:\tpcd\dbs\hundgig\
ts_58_area=h:\tpcd\dbs\hundgig\
ts_59_area=i:\tpcd\dbs\hundgig\
ts_60_area=j:\tpcd\dbs\hundgig\
ts_61_area=k:\tpcd\dbs\hundgig\
ts_62_area=l:\tpcd\dbs\hundgig\
ts_63_area=m:\tpcd\dbs\hundgig\
ts_64_area=n:\tpcd\dbs\hundgig\
ts_65_area=g:\tpcd\dbs\hundgig\
ts_66_area=h:\tpcd\dbs\hundgig\
ts_67_area=i:\tpcd\dbs\hundgig\
ts_68_area=j:\tpcd\dbs\hundgig\
ts_69_area=k:\tpcd\dbs\hundgig\
ts_70_area=l:\tpcd\dbs\hundgig\
ts_71_area=m:\tpcd\dbs\hundgig\
ts_72_area=n:\tpcd\dbs\hundgig\
ts_73_area=g:\tpcd\dbs\hundgig\
ts_74_area=h:\tpcd\dbs\hundgig\
ts_75_area=i:\tpcd\dbs\hundgig\
ts_76_area=j:\tpcd\dbs\hundgig\
ts_77_area=k:\tpcd\dbs\hundgig\
ts_78_area=l:\tpcd\dbs\hundgig\
ts_79_area=m:\tpcd\dbs\hundgig\
ts_180_area=n:\tpcd\dbs\hundgig\

ts_o_group_#ts=80
ts_o_group_#files=2
ts_o_group_first_size=115m
ts_o_group_size=115m
ts_o_group_storage=(initial 64k next 112m pctincrease 0)
ts_o1_area=g:\tpcd\dbs\hundgig\
ts_o2_area=h:\tpcd\dbs\hundgig\
ts_o3_area=i:\tpcd\dbs\hundgig\
ts_o4_area=j:\tpcd\dbs\hundgig\
ts_o5_area=k:\tpcd\dbs\hundgig\
ts_o6_area=l:\tpcd\dbs\hundgig\
ts_o7_area=m:\tpcd\dbs\hundgig\
ts_o8_area=n:\tpcd\dbs\hundgig\
ts_o9_area=g:\tpcd\dbs\hundgig\
ts_o10_area=h:\tpcd\dbs\hundgig\
ts_o11_area=i:\tpcd\dbs\hundgig\
ts_o12_area=j:\tpcd\dbs\hundgig\
ts_o13_area=k:\tpcd\dbs\hundgig\
ts_o14_area=l:\tpcd\dbs\hundgig\
ts_o15_area=m:\tpcd\dbs\hundgig\
ts_o16_area=n:\tpcd\dbs\hundgig\
ts_o17_area=g:\tpcd\dbs\hundgig\
ts_o18_area=h:\tpcd\dbs\hundgig\
ts_o19_area=i:\tpcd\dbs\hundgig\
ts_o20_area=j:\tpcd\dbs\hundgig\
ts_o21_area=k:\tpcd\dbs\hundgig\
ts_o22_area=l:\tpcd\dbs\hundgig\
ts_o23_area=m:\tpcd\dbs\hundgig\
ts_o24_area=n:\tpcd\dbs\hundgig\
ts_o25_area=g:\tpcd\dbs\hundgig\
ts_o26_area=h:\tpcd\dbs\hundgig\
ts_o27_area=i:\tpcd\dbs\hundgig\
ts_o28_area=j:\tpcd\dbs\hundgig\
ts_o29_area=k:\tpcd\dbs\hundgig\
ts_o30_area=l:\tpcd\dbs\hundgig\
ts_o31_area=m:\tpcd\dbs\hundgig\
ts_o32_area=n:\tpcd\dbs\hundgig\
ts_o33_area=g:\tpcd\dbs\hundgig\
ts_o34_area=h:\tpcd\dbs\hundgig\
ts_o35_area=i:\tpcd\dbs\hundgig\
ts_o36_area=j:\tpcd\dbs\hundgig\
ts_o37_area=k:\tpcd\dbs\hundgig\
ts_o38_area=l:\tpcd\dbs\hundgig\
```


Appendix B

```
#anl_objects =
l_ored,o_clokod,l_sdqepso,l_rssco,l_pqesod,ps_pksksc,o_op,ps_spsa,p_cb
p,p_sptm,&
#p_tp,s_skey,s_sn,lineitem,orders,partsupp,parts,customer,supplier,nat
ion,region
anl_objects =
l_pqesod,p_sptm,l_rssco,l_ored,s_sn,s_skey,o_op,o_clokod,&
lineitem,orders,parts,customer,supplier,nation,region,ps_spsa,ps_pksksc,
&
ps_pksk,partsupp

anl_l_ored_type = index
#anl_l_ored_estimate = sample 200000 rows
anl_l_ored_estimate = sample 10 percent
anl_l_sdqepso_type = index
anl_l_sdqepso_estimate = sample 10 percent
anl_l_rssco_type = index
anl_l_rssco_estimate = sample 10 percent
anl_l_rssco_type = index
anl_l_pqesod_type = index
anl_l_pqesod_estimate = sample 2500 rows
anl_l_pqesod_estimate = sample 10 percent
anl_l_test_type = index
anl_l_test_estimate = sample 10 percent

anl_o_clokod_type = index
#anl_o_clokod_estimate = sample 2500 rows
anl_o_clokod_estimate = sample 10 percent
anl_o_op_type = index
#anl_o_op_estimate = sample 200000 rows
anl_o_op_estimate = sample 10 percent

anl_ps_skey_type = index
#anl_ps_skey_estimate = sample 200000 rows
anl_ps_skey_estimate = sample 10 percent
anl_ps_pksk_type = index
#anl_ps_pksk_estimate = sample 200000 rows
anl_ps_pksk_estimate = sample 10 percent
anl_ps_pksksc_type = index
#anl_ps_pksksc_estimate = sample 200000 rows
anl_ps_pksksc_estimate = sample 10 percent
anl_ps_spsa_type = index
#anl_ps_spsa_estimate = sample 200000 rows
anl_ps_spsa_estimate = sample 10 percent

anl_p_stpm_type = index
anl_p_stpm_estimate = sample 10 percent
anl_p_key_type = index
anl_p_key_estimate = sample 10 percent
anl_p_cbp_type = index
anl_p_cbp_estimate = sample 10 percent
anl_p_sptm_type = index
anl_p_sptm_estimate = sample 10 percent
anl_p_tp_type = index
anl_p_tp_estimate = sample 10 percent
anl_p_conbr_type = index
#anl_p_conbr_estimate = sample 5000 rows
anl_p_conbr_estimate = sample 10 percent

anl_s_skey_type = index
#anl_s_skey_estimate = sample 200000 rows
anl_s_skey_estimate = sample 10 percent
anl_s_sn_type = index
anl_s_sn_estimate = sample 10 percent

anl_r_rn_type = index
#anl_r_rn_estimate = sample 200000 rows
anl_r_rn_estimate = sample 10 percent

anl_lineitem_type = table
#anl_lineitem_estimate = sample 2500 rows
anl_lineitem_estimate = sample 10 percent
anl_orders_type = table
#anl_orders_estimate = sample 2500 rows
anl_orders_estimate = sample 10 percent
anl_partsupp_type = table
#anl_partsupp_estimate = sample 200000 rows
anl_partsupp_estimate = sample 10 percent
#anl_partsupp_estimate = sample 2 percent
anl_parts_type = table
#anl_parts_estimate = sample 4000 rows
anl_parts_estimate = sample 10 percent
anl_customer_type = table
#anl_customer_estimate = sample 40000 rows
anl_customer_estimate = sample 10 percent
anl_supplier_type = table
#anl_supplier_estimate = sample 200000 rows
anl_supplier_estimate = sample 10 percent
anl_nation_type = table
#anl_nation_estimate = sample 200000 rows
anl_nation_estimate = sample 10 percent
anl_region_type = table
#anl_region_estimate = sample 200000 rows
anl_region_estimate = sample 10 percent

#####
### Section for testing and temporary fixes
#####
#load_tables = lineitem
#tab_tables = lineitem
#ts_data_groups=
#s_data_names =
#ts_index_names = ts_ind5
#skip_ts = ts_data,ts_temp

#skip_create_user = true
#kip_create_tables = true

#load_tables = lineitem
#load_tables = orders,parts,partsupp,customer,supplier,nation,region
#load_tables =
lineitem,orders,parts,partsupp,customer,supplier,nation,region
#ind_indices = o_clokod
#con_constraints =

#anl_objects = l_pqesod,o_clokod,l_sdqepso,l_rssco,s_sn,p_stpm
#anlyz_max_bg = 1

#ind_indices=l_rssco,l_pqesod
#anl_objects=l_pqesod,l_sdqepso
#ind_indices=l_pqesod,l_sdqepso
#con_constraints=
```

```
BUILD.BAT
copy \tpcd\dfs\hundgig\p_build.ora \tpcd\dfs\hundgig\tkinit.ora
date/T > ldeventlog.v1
time/T >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
echo " database creation phase " >>
ldeventlog.v1
echo "*****" >>
ldeventlog.v1
call run_dbcre.bat >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
echo " table space creation phase " >>
ldeventlog.v1
echo "*****" >>
ldeventlog.v1
call run_sctso.bat >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
echo " user and table creation phase " >>
ldeventlog.v1
echo " database build/load STARTED " >>
ldeventlog.v1
date/T >> ldeventlog.v1
time/T >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
call run_scuto.bat >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
echo " database population phase " >>
ldeventlog.v1
echo "*****" >>
ldeventlog.v1
call run_dapop.bat >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
echo " change Oracle initialization parameters " >>
ldeventlog.v1
echo "*****" >>
ldeventlog.v1
copy \tpcd\dfs\hundgig\p_index.ora \tpcd\dfs\hundgig\tkinit.ora
echo "*****" >>
ldeventlog.v1
echo " index creation phase " >>
ldeventlog.v1
echo "*****" >>
ldeventlog.v1
call run_ixcre.bat >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
echo " analyze stats phase " >>
ldeventlog.v1
echo "*****" >>
ldeventlog.v1
call run_anlyz.bat >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
echo " database build/load COMPLETED " >>
ldeventlog.v1
date/T >> ldeventlog.v1
time/T >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1

date/T > ldeventsrc.v1
time/T >> ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
echo " database creation phase " >>
ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
cat 100g_dg_dbcre.dat >> ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
echo " table space creation phase " >>
ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
cat 100g_dg_sctso.dat >> ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
echo " user and table creation phase " >>
ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
cat 100g_dg_scuto.dat >> ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
echo " database population phase " >>
ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
cat 100g_dg_dapop.dat >> ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
echo " index creation phase " >>
ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
cat 100g_dg_ixcre.dat >> ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
echo " analyze stats phase " >>
ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
cat 100g_dg_anlyz.dat >> ldeventsrc.v1

LOAD.BAT
copy \tpcd\dfs\hundgig\p_build.ora \tpcd\dfs\hundgig\tkinit.ora
date/T > ldeventlog.v1
time/T >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
echo " user and table creation phase " >>
ldeventlog.v1
echo " database build/load STARTED " >>
ldeventlog.v1
date/T >> ldeventlog.v1
time/T >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
```

```

call run_scuto.bat >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
echo "          database population phase          " >>
ldeventlog.v1
echo "*****" >>
ldeventlog.v1
call run_dapop.bat >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
echo "          change oracle initialization parameters          " >>
ldeventlog.v1
echo "*****" >>
ldeventlog.v1
copy \tpcd\dfs\hundgig\p_index.ora \tpcd\dfs\hundgig\tkinit.ora
echo "*****" >>
ldeventlog.v1
echo "          index creation phase          " >>
ldeventlog.v1
echo "*****" >>
ldeventlog.v1
call run_ixcre.bat >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
echo "          analyze stats phase          " >>
ldeventlog.v1
echo "*****" >>
ldeventlog.v1
call run_anlyz.bat >> ldeventlog.v1
echo "*****" >>
ldeventlog.v1
echo "          database build/load COMPLETED          " >>
ldeventlog.v1
date/T >> ldeventlog.v1
time/T >> ldeventlog.v1
echo "*****" >>
ldeventsrc.v1
echo "          user and table creation phase          " >>
ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
cat 100g_dg_scuto.dat >> ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
echo "          database population phase          " >>
ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
cat 100g_dg_dapop.dat >> ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
echo "          index creation phase          " >>
ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
cat 100g_dg_ixcre.dat >> ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
echo "          analyze stats phase          " >>
ldeventsrc.v1
echo "*****" >>
ldeventsrc.v1
cat 100g_dg_anlyz.dat >> ldeventsrc.v1

100G_DG_SCUTO.DAT
#####
##
## preprocessing-like directives
%b-preproc
*sql
\svrmgr30 <<!
\set echo on;
\set termout on;
\connect internal;
\select to_char(sysdate, 'MM-DD-YYYY HH24:MI:SS') now from dual;
\}
\select to_char(sysdate, 'MM-DD-YYYY HH24:MI:SS') now from dual;
\exit;
\!
*load
\sqlldr80 {}
*mknod
\mknod {}
*dbgen
\dbgen {}
*sh
\{}
%e-preproc
%b-scuto
*bgon=64
#####
## Schema Creation Phase - User and Tables ONLY (no datafiles)
*time=Begin creating user and tables (no datafiles)
# creating tpcd user
*sql
{
drop user tpcd cascade;
grant resource,unlimited tablespace,connect
to tpcd identified by tpcd;
}
*wait
{
*sql
{
connect tpcd/tpcd;
}
*wait
{
*sql
{
connect tpcd/tpcd;
drop table lineitem;
create table lineitem (
_l_shipdate          date ,
_l_orderkey          number NOT NULL,
_l_discount          number ,
_l_extendedprice     number ,
_l_suppkey           number NOT NULL,
_l_quantity          number ,
_l_returnflag        char(1),
_l_partkey           number NOT NULL,
_l_linestatus        char(1),
_l_tax               number NOT NULL,
_l_commitdate        date ,
_l_receiptdate       date ,
_l_shipmode          varchar(10),
_l_linenumbr         number NOT NULL,
_l_shipinstruct      varchar(25) ,
_l_comment           varchar(44)
)
pctfree 1
pctused 99
initrans 10
tablespace ts_11
storage (initial 64k next 460m freelists 16 freelist groups 5
pctincrease 0)
parallel (degree 60 instances 1)
partition by range (l_shipdate)
(
partition item1 values less than (to_date('1992-03-24','YYYY-MM-DD'))
tablespace ts_11
partition item2 values less than (to_date('1992-05-01','YYYY-MM-DD'))
tablespace ts_12
partition item3 values less than (to_date('1992-06-01','YYYY-MM-DD'))
tablespace ts_13
partition item4 values less than (to_date('1992-07-01','YYYY-MM-DD'))
tablespace ts_14
partition item5 values less than (to_date('1992-08-01','YYYY-MM-DD'))
tablespace ts_15
partition item6 values less than (to_date('1992-09-01','YYYY-MM-DD'))
tablespace ts_16
partition item7 values less than (to_date('1992-10-01','YYYY-MM-DD'))
tablespace ts_17
partition item8 values less than (to_date('1992-11-01','YYYY-MM-DD'))
tablespace ts_18
partition item9 values less than (to_date('1992-12-01','YYYY-MM-DD'))
tablespace ts_19
partition item10 values less than (to_date('1993-01-01','YYYY-MM-DD'))
tablespace ts_110
partition item11 values less than (to_date('1993-02-01','YYYY-MM-DD'))
tablespace ts_111
partition item12 values less than (to_date('1993-03-01','YYYY-MM-DD'))
tablespace ts_112
partition item13 values less than (to_date('1993-04-01','YYYY-MM-DD'))
tablespace ts_113
partition item14 values less than (to_date('1993-05-01','YYYY-MM-DD'))
tablespace ts_114
partition item15 values less than (to_date('1993-06-01','YYYY-MM-DD'))
tablespace ts_115
partition item16 values less than (to_date('1993-07-01','YYYY-MM-DD'))
tablespace ts_116
partition item17 values less than (to_date('1993-08-01','YYYY-MM-DD'))
tablespace ts_117
partition item18 values less than (to_date('1993-09-01','YYYY-MM-DD'))
tablespace ts_118
partition item19 values less than (to_date('1993-10-01','YYYY-MM-DD'))
tablespace ts_119
partition item20 values less than (to_date('1993-11-01','YYYY-MM-DD'))
tablespace ts_120
partition item21 values less than (to_date('1993-12-01','YYYY-MM-DD'))
tablespace ts_121
partition item22 values less than (to_date('1994-01-01','YYYY-MM-DD'))
tablespace ts_122
partition item23 values less than (to_date('1994-02-01','YYYY-MM-DD'))
tablespace ts_123
partition item24 values less than (to_date('1994-03-01','YYYY-MM-DD'))
tablespace ts_124
partition item25 values less than (to_date('1994-04-01','YYYY-MM-DD'))
tablespace ts_125
partition item26 values less than (to_date('1994-05-01','YYYY-MM-DD'))
tablespace ts_126
partition item27 values less than (to_date('1994-06-01','YYYY-MM-DD'))
tablespace ts_127
partition item28 values less than (to_date('1994-07-01','YYYY-MM-DD'))
tablespace ts_128
partition item29 values less than (to_date('1994-08-01','YYYY-MM-DD'))
tablespace ts_129
partition item30 values less than (to_date('1994-09-01','YYYY-MM-DD'))
tablespace ts_130
partition item31 values less than (to_date('1994-10-01','YYYY-MM-DD'))
tablespace ts_131
partition item32 values less than (to_date('1994-11-01','YYYY-MM-DD'))
tablespace ts_132
partition item33 values less than (to_date('1994-12-01','YYYY-MM-DD'))
tablespace ts_133
partition item34 values less than (to_date('1995-01-01','YYYY-MM-DD'))
tablespace ts_134
partition item35 values less than (to_date('1995-02-01','YYYY-MM-DD'))
tablespace ts_135
partition item36 values less than (to_date('1995-03-01','YYYY-MM-DD'))
tablespace ts_136
partition item37 values less than (to_date('1995-04-01','YYYY-MM-DD'))
tablespace ts_137
partition item38 values less than (to_date('1995-05-01','YYYY-MM-DD'))
tablespace ts_138
partition item39 values less than (to_date('1995-06-01','YYYY-MM-DD'))
tablespace ts_139
partition item40 values less than (to_date('1995-07-01','YYYY-MM-DD'))
tablespace ts_140
partition item41 values less than (to_date('1995-08-01','YYYY-MM-DD'))
tablespace ts_141
partition item42 values less than (to_date('1995-09-01','YYYY-MM-DD'))
tablespace ts_142
partition item43 values less than (to_date('1995-10-01','YYYY-MM-DD'))
tablespace ts_143
partition item44 values less than (to_date('1995-11-01','YYYY-MM-DD'))
tablespace ts_144

```



```

partition ord62 values less than (to_date('1997-03-01','YYYY-MM-DD'))
tablespace ts_o62
partition ord63 values less than (to_date('1997-04-01','YYYY-MM-DD'))
tablespace ts_o63
partition ord64 values less than (to_date('1997-05-01','YYYY-MM-DD'))
tablespace ts_o64
partition ord65 values less than (to_date('1997-06-01','YYYY-MM-DD'))
tablespace ts_o65
partition ord66 values less than (to_date('1997-07-01','YYYY-MM-DD'))
tablespace ts_o66
partition ord67 values less than (to_date('1997-08-01','YYYY-MM-DD'))
tablespace ts_o67
partition ord68 values less than (to_date('1997-09-01','YYYY-MM-DD'))
tablespace ts_o68
partition ord69 values less than (to_date('1997-10-01','YYYY-MM-DD'))
tablespace ts_o69
partition ord70 values less than (to_date('1997-11-01','YYYY-MM-DD'))
tablespace ts_o70
partition ord71 values less than (to_date('1997-12-01','YYYY-MM-DD'))
tablespace ts_o71
partition ord72 values less than (to_date('1998-01-01','YYYY-MM-DD'))
tablespace ts_o72
partition ord73 values less than (to_date('1998-02-01','YYYY-MM-DD'))
tablespace ts_o73
partition ord74 values less than (to_date('1998-03-01','YYYY-MM-DD'))
tablespace ts_o74
partition ord75 values less than (to_date('1998-04-01','YYYY-MM-DD'))
tablespace ts_o75
partition ord76 values less than (to_date('1998-05-01','YYYY-MM-DD'))
tablespace ts_o76
partition ord77 values less than (to_date('1998-06-01','YYYY-MM-DD'))
tablespace ts_o77
partition ord78 values less than (to_date('1998-07-01','YYYY-MM-DD'))
tablespace ts_o78
partition ord79 values less than (to_date('1998-08-01','YYYY-MM-DD'))
tablespace ts_o79
partition ord80 values less than (to_date('1998-09-01','YYYY-MM-DD'))
tablespace ts_o80
)
drop table partsupp;
create table partsupp (
  ps_partkey          number NOT NULL,
  ps_suppkey         number NOT NULL,
  ps_supplycost      number NOT NULL,
  ps_availqty        number,
  ps_comment         varchar(199)
)
pctfree 0
pctused 99
tablespace ts_ps
storage (initial 64k next 83m pctincrease 0)
;
drop table parts;
create table parts (
  p_partkey          number NOT NULL,
  p_type            varchar(25),
  p_size            number,
  p_brand           varchar(10),
  p_name            varchar(55),
  p_container       varchar(10),
  p_mfg             varchar(25),
  p_retailprice     number,
  p_comment         varchar(23)
)
pctfree 0
pctused 99
tablespace ts_p1
storage (initial 64k next 2m pctincrease 0)
parallel (degree 60 instances 1)
partition by range (p_size)
(
  partition part1 values less than (2)
  tablespace ts_p1

  partition part2 values less than (3)
  tablespace ts_p2

  partition part3 values less than (4)
  tablespace ts_p3

  partition part4 values less than (5)
  tablespace ts_p4

  partition part5 values less than (6)
  tablespace ts_p5

  partition part6 values less than (7)
  tablespace ts_p6

  partition part7 values less than (8)
  tablespace ts_p7

  partition part8 values less than (9)
  tablespace ts_p8

  partition part9 values less than (10)
  tablespace ts_p9

  partition part10 values less than (11)
  tablespace ts_p10

  partition part11 values less than (12)
  tablespace ts_p11

  partition part12 values less than (13)
  tablespace ts_p12

  partition part13 values less than (14)
  tablespace ts_p13

  partition part14 values less than (15)
  tablespace ts_p14

  partition part15 values less than (16)
  tablespace ts_p15

  partition part16 values less than (17)
  tablespace ts_p16

  partition part17 values less than (18)
  tablespace ts_p17

  partition part18 values less than (19)
  tablespace ts_p18

  partition part19 values less than (20)
  tablespace ts_p19

  partition part20 values less than (21)
  tablespace ts_p20

  partition part21 values less than (22)
  tablespace ts_p21

  partition part22 values less than (23)
  tablespace ts_p22

  partition part23 values less than (24)
  tablespace ts_p23

  partition part24 values less than (25)
  tablespace ts_p24

  partition part25 values less than (26)
  tablespace ts_p25

  partition part26 values less than (27)
  tablespace ts_p26

  partition part27 values less than (28)
  tablespace ts_p27

  partition part28 values less than (29)
  tablespace ts_p28

  partition part29 values less than (30)
  tablespace ts_p29

  partition part30 values less than (31)
  tablespace ts_p30

  partition part31 values less than (32)
  tablespace ts_p31

  partition part32 values less than (33)
  tablespace ts_p32

  partition part33 values less than (34)
  tablespace ts_p33

  partition part34 values less than (35)
  tablespace ts_p34

  partition part35 values less than (36)
  tablespace ts_p35

  partition part36 values less than (37)
  tablespace ts_p36

  partition part37 values less than (38)
  tablespace ts_p37

  partition part38 values less than (39)
  tablespace ts_p38

  partition part39 values less than (40)
  tablespace ts_p39

  partition part40 values less than (41)
  tablespace ts_p40

  partition part41 values less than (42)
  tablespace ts_p41

  partition part42 values less than (43)
  tablespace ts_p42

  partition part43 values less than (44)
  tablespace ts_p43

  partition part44 values less than (45)
  tablespace ts_p44

  partition part45 values less than (46)
  tablespace ts_p45

  partition part46 values less than (47)
  tablespace ts_p46

  partition part47 values less than (48)
  tablespace ts_p47

  partition part48 values less than (49)
  tablespace ts_p48

  partition part49 values less than (50)
  tablespace ts_p49

  partition part50 values less than (51)
  tablespace ts_p50
)
;
drop table customer;
create table customer (
  c_custkey          number NOT NULL,
  c_mktsegment       varchar(10),
  c_nationkey        number,
  c_name             varchar(25),
  c_address          varchar(40),
  c_phone            varchar(15),
  c_acctbal         number,
  c_comment          varchar(117)
)
pctfree 0
pctused 99
tablespace ts_c1
storage (initial 64k next 19m pctincrease 0)
parallel (degree 60 instances 1)
partition by range (c_mktsegment)
(
  partition cust1 values less than ('AUTOMOBILF')
  tablespace ts_c1

  partition cust2 values less than ('BUILDINH')
  tablespace ts_c2

  partition cust3 values less than ('FURNITURF')
  tablespace ts_c3

  partition cust4 values less than ('HOUSEHOLE')
  tablespace ts_c4

  partition cust5 values less than ('MACHINERZ')
  tablespace ts_c5
)
;
drop table supplier;
create table supplier (
  s_suppkey          number NOT NULL,
  s_nationkey        number,
  s_comment          varchar(101),
  s_name             varchar(25),
  s_address          varchar(40),
  s_phone            varchar(15),
  s_acctbal         number
)
pctfree 0
pctused 99
tablespace ts_s
storage (initial 64k next 10m pctincrease 0)
parallel (degree 60 instances 1)
;
drop table nation;
create table nation (
  n_nationkey        number NOT NULL,
  n_name             varchar(25),
  n_regionkey        number,
  n_comment          varchar(152)
)

```


Appendix B

```
{
shutdown
}
wait
*time=Done data population
wait
*bgofff
%e-dapop

100G_DG_IXCRE.DAT
#####
##
# preprocessing-like directives

%b-preproc

*sql
\svrmgr30 <<!
\set echo on;
\set termout on;
\connect internal;
\select to_char(sysdate, 'MM-DD-YYYY HH24:MI:SS') now from dual;
{}
\select to_char(sysdate, 'MM-DD-YYYY HH24:MI:SS') now from dual;
\exit;
\!

*load
\sqlldr80 {}

*mknod
\mknod {}

*dbgen
\dbgen {}

*sh
\{}

%e-preproc
%b-ixcre
*bgon=1
#####
#####
# Index Creation Phase
*time=Begin index creation
*wait
*sql
{
shutdown
startup pfile=d:\tpcd\dbshundgig\init1_tpcd.ora
}
*wait
*time=Begin creating index ps_pksk
*sql
{
connect tpcd/tpcd;
drop index ps_pksk;
create unique index ps_pksk
on partsupp (ps_partkey,ps_supkey)
pctfree 1
tablespace ts_ind4
storage (initial 1m next 1m maxextents unlimited pctincrease 0)
parallel (degree 60 instances 1)
}
*wait
*time=Done creating index ps_pksk
*time=Begin creating index ps_pksksc
*sql
{
connect tpcd/tpcd;
drop index ps_pksksc;
create unique index ps_pksksc
on partsupp (ps_partkey,ps_supkey,ps_supplycost)
pctfree 1
tablespace ts_ind4
storage (initial 1m next 1m maxextents unlimited pctincrease 0)
parallel (degree 60 instances 1)
}
*wait
*time=Done creating index ps_pksksc
*time=Begin creating index ps_spsa
*sql
{
connect tpcd/tpcd;
drop index ps_spsa;
create unique index ps_spsa
on partsupp (ps_supkey,ps_partkey,ps_supplycost,ps_availqty)
pctfree 1
tablespace ts_ind4
storage (initial 1m next 1m maxextents unlimited pctincrease 0)
parallel (degree 60 instances 1)
}
*wait
*time=Done creating index ps_spsa
*time=Begin creating index l_ored
*sql
{
connect tpcd/tpcd;
drop index l_ored;
create index l_ored
on lineitem (l_orderkey,l_returnflag,l_extendedprice,l_discount)
pctfree 2
initrans 10
tablespace ts_ind1
storage (initial 5m next 5m freelists 16 freelist groups 5 maxextents
unlimited pctincrease 0)
parallel (degree 60 instances 1)
}
*wait
*time=Done creating index l_ored
*time=Begin creating index o_clokod
*sql
{
connect tpcd/tpcd;
drop index o_clokod;
create index o_clokod
on orders (o_clerk,o_orderkey,o_orderdate)
pctfree 2
initrans 10
tablespace ts_ind3
storage (initial 5m next 5m freelists 16 freelist groups 5 maxextents
unlimited pctincrease 0)
local (
partition oclokod1,
partition oclokod2,
partition oclokod3,
partition oclokod4,
partition oclokod5,
partition oclokod6,
partition oclokod7,
partition oclokod8,
partition oclokod9,
partition oclokod10,
partition oclokod11,
partition oclokod12,
partition oclokod13,
partition oclokod14,
partition oclokod15,
partition oclokod16,
partition oclokod17,
partition oclokod18,
partition oclokod19,
partition oclokod20,
partition oclokod21,
partition oclokod22,
partition oclokod23,
partition oclokod24,
partition oclokod25,
partition oclokod26,
partition oclokod27,
partition oclokod28,
partition oclokod29,
partition oclokod30,
partition oclokod31,
partition oclokod32,
partition oclokod33,
partition oclokod34,
partition oclokod35,
partition oclokod36,
partition oclokod37,
partition oclokod38,
partition oclokod39,
partition oclokod40,
partition oclokod41,
partition oclokod42,
partition oclokod43,
partition oclokod44,
partition oclokod45,
partition oclokod46,
partition oclokod47,
partition oclokod48,
partition oclokod49,
partition oclokod50,
partition oclokod51,
partition oclokod52,
partition oclokod53,
partition oclokod54,
partition oclokod55,
partition oclokod56,
partition oclokod57,
partition oclokod58,
partition oclokod59,
partition oclokod60,
partition oclokod61,
partition oclokod62,
partition oclokod63,
partition oclokod64,
partition oclokod65,
partition oclokod66,
partition oclokod67,
partition oclokod68,
partition oclokod69,
partition oclokod70,
partition oclokod71,
partition oclokod72,
partition oclokod73,
partition oclokod74,
partition oclokod75,
partition oclokod76,
partition oclokod77,
partition oclokod78,
partition oclokod79,
partition oclokod80
}
parallel (degree 60 instances 1)
}
*wait
*time=Done creating index o_clokod
*time=Begin creating index l_pqesod
*sql
{
connect tpcd/tpcd;
drop index l_pqesod;
create index l_pqesod
on lineitem (l_partkey,l_quantity,l_extendedprice,l_supkey,l_orderkey,l_discount)
pctfree 2
initrans 10
tablespace ts_ind1
storage (initial 5m next 5m freelists 16 freelist groups 5 maxextents
unlimited pctincrease 0)
parallel (degree 60 instances 1)
}
*wait
*time=Done creating index l_pqesod
*time=Begin creating index l_rossc
*sql
{
connect tpcd/tpcd;
drop index l_rossc;
create index l_rossc
on lineitem (l_receiptdate,l_orderkey,l_shipmode,l_shipdate,l_commitdate)
pctfree 2
initrans 10
tablespace ts_ind5
storage (initial 5m next 5m freelists 16 freelist groups 5 maxextents
unlimited pctincrease 0)
global partition by range (l_receiptdate)
(
partition lrossc1 values less than (to_date('1992-02-01','YYYY-MM-DD')),
partition lrossc2 values less than (to_date('1992-03-01','YYYY-MM-DD')),
partition lrossc3 values less than (to_date('1992-04-01','YYYY-MM-DD')),
partition lrossc4 values less than (to_date('1992-05-01','YYYY-MM-DD')),
partition lrossc5 values less than (to_date('1992-06-01','YYYY-MM-DD')),
partition lrossc6 values less than (to_date('1992-07-01','YYYY-MM-DD')),
partition lrossc7 values less than (to_date('1992-08-01','YYYY-MM-DD')),
partition lrossc8 values less than (to_date('1992-09-01','YYYY-MM-DD')),
partition lrossc9 values less than (to_date('1992-10-01','YYYY-MM-DD')),
partition lrossc10 values less than (to_date('1992-11-01','YYYY-MM-DD')),
partition lrossc11 values less than (to_date('1992-12-01','YYYY-MM-DD')),
partition lrossc12 values less than (to_date('1993-01-01','YYYY-MM-DD')),
partition lrossc13 values less than (to_date('1993-02-01','YYYY-MM-DD')),
partition lrossc14 values less than (to_date('1993-03-01','YYYY-MM-DD')),
partition lrossc15 values less than (to_date('1993-04-01','YYYY-MM-DD')),
partition lrossc16 values less than (to_date('1993-05-01','YYYY-MM-DD')),
partition lrossc17 values less than (to_date('1993-06-01','YYYY-MM-DD')),
partition lrossc18 values less than (to_date('1993-07-01','YYYY-MM-DD')),
partition lrossc19 values less than (to_date('1993-08-01','YYYY-MM-DD')),
partition lrossc20 values less than (to_date('1993-09-01','YYYY-MM-DD')),
partition lrossc21 values less than (to_date('1993-10-01','YYYY-MM-DD')),
partition lrossc22 values less than (to_date('1993-11-01','YYYY-MM-DD')),
partition lrossc23 values less than (to_date('1993-12-01','YYYY-MM-DD')),
partition lrossc24 values less than (to_date('1994-01-01','YYYY-MM-DD')),
partition lrossc25 values less than (to_date('1994-02-01','YYYY-MM-DD')),
partition lrossc26 values less than (to_date('1994-03-01','YYYY-MM-DD')),
partition lrossc27 values less than (to_date('1994-04-01','YYYY-MM-DD')),
partition lrossc28 values less than (to_date('1994-05-01','YYYY-MM-DD')),
partition lrossc29 values less than (to_date('1994-06-01','YYYY-MM-DD')),
partition lrossc30 values less than (to_date('1994-07-01','YYYY-MM-DD')),
partition lrossc31 values less than (to_date('1994-08-01','YYYY-MM-DD')),

```



```

append
fields terminated by '|'
(
  p_partkey      ,
  p_type        ,
  p_size        ,
  p_brand       ,
  p_name        ,
  p_container   ,
  p_mfggr       ,
  p_retailprice ,
  p_comment     ,
)
)

---
options
(
storage = (initial 2m next 2m)
)
unrecoverable
load
-- This is where INFILE should go.
into table parts
partition (PART2)
-- This is where PARTITION is specified.
append
fields terminated by '|'
(
  p_partkey      ,
  p_type        ,
  p_size        ,
  p_brand       ,
  p_name        ,
  p_container   ,
  p_mfggr       ,
  p_retailprice ,
  p_comment     ,
)
)

3 ... 99
---
options
(
storage = (initial 2m next 2m)
)
unrecoverable
load
-- This is where INFILE should go.
into table parts
partition (PART50)
-- This is where PARTITION is specified.
append
fields terminated by '|'
(
  p_partkey      ,
  p_type        ,
  p_size        ,
  p_brand       ,
  p_name        ,
  p_container   ,
  p_mfggr       ,
  p_retailprice ,
  p_comment     ,
)
)

PARTSUPP.CTL
---
--- partsuppctl for delimited records
---

options
(
storage = (initial 83m next 83m)
)
unrecoverable
load
into table partsupp
append
fields terminated by '|'
(
  ps_partkey     ,
  ps_suppkey     ,
  ps_supplycost  ,
  ps_availqty    ,
  ps_comment     ,
)
)

REGION.CTL
---
--- regionctl for delimited records
---

options
(
storage = (initial 64k next 64k)
)
unrecoverable
load
into table region
append
fields terminated by '|'
(
  r_regionkey    ,
  r_name         ,
  r_comment      ,
)
)

SUPPLIER.CTL
---
--- supplierctl for delimited records
---

options
(
storage = (initial 10m next 10m)
)
unrecoverable
load
into table supplier
append
fields terminated by '|'
(
  s_suppkey      ,
  s_nationkey    ,
  s_comment      ,
  s_name         ,
  s_address      ,
  s_phone        ,
  s_acctbal      ,
)
)

```


Appendix C: Query Text and Query Output

-- Query 1 (Original) Q1

```
SELECT L_RETURNFLAG,
L_LINESTATUS,
SUM(L_QUANTITY) AS SUM_QTY,
SUM(L_EXTENDEDPRI) AS SUM_BASE_PRICE,
SUM(L_EXTENDEDPRI * (1 - L_DISCOUNT)) AS
SUM_DISC_PRICE,
SUM(L_EXTENDEDPRI * (1 - L_DISCOUNT) * (1
+ L_TAX)) AS SUM_CHARGE,
AVG(L_QUANTITY) AS AVG_QTY,
AVG(L_EXTENDEDPRI) AS AVG_PRICE,
AVG(L_DISCOUNT) AS AVG_DISC,
COUNT(*) AS COUNT_ORDER
FROM LINEITEM
WHERE L_SHIPDATE <= TO_DATE('1998-12-
01', 'YYYY-MM-DD') - 86
GROUP BY L_RETURNFLAG, L_LINESTATUS
ORDER BY L_RETURNFLAG, L_LINESTATUS
```

L_RETURNFLAG	L_LINESTATUS	SUM_QTY	SUM_BASE_PRICE	SUM_DISC_PRICE	SUM_CHARGE	AVG_QTY	AVG_PRICE	AVG_DISC	COUNT_ORDER
A	F		3779444456.00						
5667241537767.56		5383869735321.49							
5599230751499.22		25.50		38237.70					
0.05		148210840.00							
N	F		98685681.00						
147989627721.00		140585223091.30							
146210064832.84		25.50		38244.11					
0.05		3869606.00							
N	O		7462395999.00						
11189890693167.18		10630397840056.28							
11055624575500.17		25.50		38238.27					
0.05		292635902.00							
R	F		3779061567.00						
5666619103069.79		5383304295411.96							
5598630538175.75		25.50		38235.27					
0.05		148203986.00							

4 rows processed.
Statement Processed in 4977.75 seconds.

-- Query 4 (Variant B) Q4 B

```
SELECT
O_ORDERPRIORITY,
COUNT(*) AS ORDER_COUNT
FROM ORDERS
WHERE O_ORDERKEY IN (
SELECT DISTINCT O_ORDERKEY
FROM LINEITEM, ORDERS
WHERE L_ORDERKEY = O_ORDERKEY
AND O_ORDERDATE >= TO_DATE('1995-02-
01', 'YYYY-MM-DD')
AND O_ORDERDATE < ADD_MONTHS(TO_DATE('1995-
02-01', 'YYYY-MM-DD'), 3)
AND L_COMMITDATE < L_RECEIPTDATE)
GROUP BY O_ORDERPRIORITY
ORDER BY O_ORDERPRIORITY
```

O_ORDERPRIORITY	ORDER_COUNT
1-URGENT	1017809.00
2-HIGH	1020535.00
3-MEDIUM	1018442.00
4-NOT SPECIFIED	1018575.00
5-LOW	1019620.00

5 rows processed.

-- Query 15 (Variant B) Q15 B

```
CREATE TABLE REVENUE0 (SUPPLIER_NO INTEGER,
TOTAL_REVENUE NUMBER(20, 2))
tablespace TS_S
pctfree 0
pctused 99
storage (initial 5m next 5m pctincrease 0)
```

```
INSERT INTO REVENUE0
SELECT L_SUPPKEY, SUM(L_EXTENDEDPRI*(1-
L_DISCOUNT))
FROM LINEITEM
WHERE L_SHIPDATE >= TO_DATE('1995-02-
01', 'YYYY-MM-DD')
AND L_SHIPDATE < ADD_MONTHS(TO_DATE('1995-
02-01', 'YYYY-MM-DD'), 3)
GROUP BY L_SUPPKEY
```

```
SELECT S_SUPPKEY, S_NAME, S_ADDRESS,
S_PHONE, TOTAL_REVENUE
FROM SUPPLIER, REVENUE0
WHERE S_SUPPKEY = SUPPLIER_NO
AND TOTAL_REVENUE = (SELECT
MAX(TOTAL_REVENUE) FROM REVENUE0)
ORDER BY S_SUPPKEY
```

S_SUPPKEY	S_NAME	S_ADDRESS	S_PHONE	TOTAL_REVENUE
62932.00	Supplier#000062932	0160hML21xB jPg404L4kmLL	882-650-3419	2415387.29

1 row processed.

```
DROP TABLE REVENUE0
Setting the number of rows to fetch to: 20
```

-- Query 10 (Original) Q10

```
SELECT
C_CUSTKEY,
C_NAME,
SUM(L_EXTENDEDPRI * (1-L_DISCOUNT)) AS
REVENUE,
C_ACCTBAL,
N_NAME,
C_ADDRESS,
C_PHONE,
C_COMMENT
FROM CUSTOMER, ORDERS, LINEITEM, NATION
WHERE C_CUSTKEY = O_CUSTKEY
AND L_ORDERKEY = O_ORDERKEY
AND O_ORDERDATE >= TO_DATE( '1993-12-01'
, 'YYYY-MM-DD')
AND O_ORDERDATE < ADD_MONTHS(TO_DATE( '1993-
12-01', 'YYYY-MM-DD'), 3)
AND L_RETURNFLAG = 'R'
AND C_NATIONKEY = N_NATIONKEY
GROUP BY C_CUSTKEY, C_NAME, C_ACCTBAL,
C_PHONE, N_NAME, C_ADDRESS, C_COMMENT
ORDER BY REVENUE DESC
```

C_CUSTKEY	C_NAME	REVENUE	C_ACCTBAL	N_NAME	C_ADDRESS	C_PHONE	C_COMMENT
12567388.00	Customer#012567388	974656.39	8448.78				

Appendix C

INDIA	1CzhCnh	ziyinPR04Nhn7jlxz0yggjs0xjl3Lx2XRAR203jk z516
nzQmn4grNk0 m322	18-238-696-	5x3ww7L71gn7 1 3z03x2Q4 oxQR03
7451		0Bkc3Cnx6LOR2wLhjk1c42w
g7NiLoi5m1NjP1k4jmkAQQP4L3xj5k2h2Q1yszm73mSQ		6369166.00 Customer#006369166
y2Cm65QR0k71h1A5Lnm1w150Am71BAC		807950.33 -104.56
8897564.00	Customer#008897564	SAUDI ARABIA nSza312ji2kkim0ks
884337.26	7453.06	30-673-771-1191
ALGERIA		SCyRiBxznB1NnjPgN x
jRg3lCnxQ0LygCCnM12 w	10-	41h61QN5BNw61MkQwlgxNMRiozPy0LA1hL7Q
781-743-5269		8713679.00 Customer#008713679
LOO1w26gm03Riio112P113wm06R7g74yskgxjxLL6AR7		805449.32 2503.02
33A625nSyNC3BANogj6kjcj5owh2OzNAh75zP6hiwsB3		IRAQ
z5RQ3QRh2xLihP21		1w06Qg1k7N1xxLS026Nnc0P5k3AQQC 5 21-
6725432.00	Customer#006725432	202-977-9481
861454.67	9780.31	RzAhzP0k3Q4LQjir4nQA654wg53z4A0g0BBY1A
UNITED KINGDOM	0S71x7N0zO23	Cn21SRPPMmPk1xn2wxRWA2BA3C2BPNSy6CRG7
OLmWLP5m2M6i06R	33-487-484-6205	17y5i65g25x1z076n0zjk5xk15L01mn
Bi3		984320.00 Customer#000984320
1S6KN73kCnBggmBLASiSjLCgMRCg4ijAmhQilZ2RPNQ5		804324.58 4844.65
LPxgQ310Ri4ZxLA6mCQMOQ7N0R12og1xj5k77gh Ln		IRAN yz1
5610056.00	Customer#005610056	2gP51zk7P31j6R3RCjk152Bk10zM4Sk6Rik 20-217-
860575.78	8506.50	654-6796
KENYA	6B7iNmnoQ1m5kNL 2Q	y00hwnxPn mmwPgSn00i jhi xn0
NPLzhrj	24-263-725-9874	72hMQQ061oz6i3R4Q OA MQy nx2gmk6
BLAw4P5lj2wxc7nLj3hynNPOzkPjz0y3		
7060918.00	Customer#007060918	20 rows processed.
856815.64	1436.46	
UNITED STATES		-- Query 11 (Variant A) Q11 A
kBS162gh1Cy7x64QOM34MhOBP21nkBOSOQ2	34-	
120-393-8506		
zz 4Lnlj0jNORL355m60i5ggw57gg0ms2gSMPxB		
R0w11n2mz00wmy2275		
2549213.00	Customer#002549213	CREATE TABLE PART_VALUE0
855577.74	741.59	(PARTNO NUMBER, VALUE NUMBER(20, 2))
CANADA	00wjCj7zyygn60	tablespace TS_S
13-945-194-9872		pctfree 0
i2hr7ggwNn1i4N2xm5QLn x7Qkk3yMi370Ay		pctused 99
S0n3Qkkw35LxQAMgxkPC71C6i2B5LOLzCMBPYNmAjCzB		storage (initial 5m next 5m pctincrease 0)
Ln3161hy4ikjmmLR		
12857572.00	Customer#012857572	INSERT INTO PART_VALUE0
844811.00	-571.13	SELECT
INDIA		PS_PARTKEY, SUM(PS_SUPPLYCOST*PS_AVAILQTY)
Ax6BL5hPoiLzCnRN6iR	18-	FROM PARTSUPP, SUPPLIER, NATION
175-215-2827		WHERE PS_SUPPKEY = S_SUPPKEY
j64jA3zkCSLkyig1kw124wABhxjQB1gs26j5wgigixAg		AND S_NATIONKEY = N_NATIONKEY
LOnl1x50xLzy onC5nj117n60Qn7		AND N_NAME = 'IRAN'
10058912.00	Customer#010058912	GROUP BY PS_PARTKEY
825007.68	1303.77	
RUSSIA	4CBQ73lgjz61c12x14	CREATE TABLE SUM_PART_VALUE0
32-675-100-7410		(TOTAL_VALUE NUMBER(20, 2))
wmClAMzSx4Pj2m7Q1lLgnQ3xlihr3zmn4mwOLk3		tablespace TS_S
0w6iimip1l		pctfree 0
1403759.00	Customer#001403759	pctused 99
818215.44	4510.10	storage (initial 5m next 5m pctincrease 0)
IRAN		INSERT INTO SUM_PART_VALUE0
05zmmkz5w6BLLRxA5wRwMP2NxBwkiy7MRP	20-	SELECT
825-378-5608		SUM(PS_SUPPLYCOST*PS_AVAILQTY) *
1PzP0MMRm3 0 3ky0CA01P1QNRhhkh6h6jL		0.0000010000
zw70MS541CRkiimz20SR17jMARBhw10mnQ33AgmM		FROM PARTSUPP, SUPPLIER, NATION
wMjzQ1xBL		WHERE PS_SUPPKEY = S_SUPPKEY
8921275.00	Customer#008921275	AND S_NATIONKEY = N_NATIONKEY
817259.07	-399.99	AND N_NAME = 'IRAN'
PERU	4xkQBPAajlxzq2	SELECT
0B2nzg	27-890-784-4314	PARTNO, VALUE
3gQ1z5563CL05A204nP 1Nx4j 0kORwgQ		FROM PART_VALUE0
2mShS6CNL54iy6xN10yww 0gwg51Q		WHERE VALUE >
5123nS0gPhmQy4AMQ00ynigx2		(SELECT SUM(TOTAL_VALUE) FROM
13614604.00	Customer#013614604	SUM_PART_VALUE0)
816130.68	3013.96	ORDER BY VALUE DESC
RUSSIA	0M5i5k4110	
5js713mAS3k5iA2mox g	32-539-305-	
2013		
4mB6nQA30RNMLzxB zPnck2S3P7RgLh213MQkiN		PARTNO
11553353.00	Customer#011553353	4140994.00
809286.32	9949.88	11920877.00
UNITED KINGDOM	k	23123216.10
h705NSz4y1r21h743j6iMnNrnjMzh0Azk	33-	8559734.00
225-301-4142		21920759.58
		19252399.00
		21093285.43
		6070561.00
		21068807.28

```
19932203.00      20946797.32
16018003.00      20757391.94
10479169.00      20730777.28
18118493.00      20681188.81
15027179.00      20648624.34
```

94553 rows processed.

DROP TABLE PART_VALUE0

DROP TABLE SUM_PART_VALUE0

-- Query 6 (Original) Q6

```
SELECT
SUM(L_EXTENDEDPRI * L_DISCOUNT) AS REVENUE
FROM LINEITEM
WHERE L_SHIPDATE >= TO_DATE('1995-01-01', 'YYYY-MM-DD')
AND L_SHIPDATE < ADD_MONTHS(TO_DATE('1995-01-01', 'YYYY-MM-DD'), 12)
AND L_DISCOUNT BETWEEN 0.06 - 0.01 AND 0.06 + 0.01
AND L_QUANTITY < 25
```

```
REVENUE
13412131094.15
```

1 row processed.
Setting the number of rows to fetch to: 100

-- Query 2 (Variant A) Q2 A

```
SELECT
S_ACCTBAL,
S_NAME,
N_NAME,
P_PARTKEY,
P_MFGR,
S_ADDRESS,
S_PHONE,
S_COMMENT
FROM PARTS, SUPPLIER, PARTSUPP, NATION,
REGION
WHERE P_PARTKEY = PS_PARTKEY
AND S_SUPPKEY = PS_SUPPKEY
AND P_SIZE = 22
AND P_TYPE LIKE '%NICKEL'
AND S_NATIONKEY = N_NATIONKEY
AND N_REGIONKEY = R_REGIONKEY
AND R_NAME = 'MIDDLE EAST'
AND (P_PARTKEY, PS_SUPPLYCOST) IN
(SELECT PS_PARTKEY, MIN(PS_SUPPLYCOST)
FROM PARTSUPP, SUPPLIER, NATION, REGION
WHERE S_SUPPKEY = PS_SUPPKEY
AND S_NATIONKEY = N_NATIONKEY
AND N_REGIONKEY = R_REGIONKEY
AND R_NAME = 'MIDDLE EAST'
GROUP BY PS_PARTKEY)
ORDER BY S_ACCTBAL DESC, N_NAME, S_NAME,
P_PARTKEY
```

```
S_ACCTBAL      S_NAME
N_NAME
P_PARTKEY      P_MFGR
S_ADDRESS
S_PHONE
S_COMMENT
9999.58        Supplier#000081596
JORDAN
11331562.00    Manufacturer#3
1jhx60NwRSR2LwiQi R nhxmR04Poj 2311 23-
802-625-7039
1Mgy0AMj3mN1
jgOgw2kLwA7m6Pi1300011B4nOx5j5A7i4nwPyOhnPjs
h550z s
9999.18        Supplier#000697011
IRAQ
```

```
3447007.00      Manufacturer#3
jgm3i17ykksim21ML5nRmLQS 21-
961-197-9510
PLk5RMgPwlgYRP1kAgis1iP0knszxiA1Phzwhj1Qj156
06j6305 314iyy
9998.95        Supplier#000148225
```

```
IRAN
10898214.00     Manufacturer#1
5Si4Sz L5RzkB67i7kBzi2R4zmO 20-
782-280-8482
zngNmMwkwn2j6z1Qn 2PB3z2 Pz P5AC26m
khCAPL6Bwik10
9998.55        Supplier#000498783
```

```
EGYPT
6748764.00      Manufacturer#2
5AS1h3MByB1ChN1SsizihMwMyxLm1 14-
434-458-3525
524230zimwPQQjy1R5gQAgh124B3P06zzMS5xy4gBQ7A
g3yn3mMA0N6jBMLzP257Sm00j5Alj
9998.55        Supplier#000498783
```

```
EGYPT
12748746.00     Manufacturer#1
5AS1h3MByB1ChN1SsizihMwMyxLm1 14-
434-458-3525
524230zimwPQQjy1R5gQAgh124B3P06zzMS5xy4gBQ7A
g3yn3mMA0N6jBMLzP257Sm00j5Alj
9998.41        Supplier#000110680
```

```
JORDAN
12610655.00     Manufacturer#1
CQ3x4 4gCNRog67A51wQiyzBi6ORQS2kQ5R2Q 23-
556-119-9861
1Ayo10SiNmLh1 2Azg3Bjx
ANnB4iwn1C6nPzBQRR6MMz 1w Lsy1Rz1g
9998.40        Supplier#000380860
```

```
JORDAN
8130851.00      Manufacturer#2
x1P3RhQMSL1hLgQkhs CLM 23-
814-345-7369
B44NRQ0g311LBiQ16RmmLni4yh60xxBgjNNO4
ho700hrjn7A31
02Qzy7MLNyz0knmB3P63g1Rj57RNRyy3BOnSNM7Nyzkh
Anh
9998.29        Supplier#000525187
```

```
IRAN
5775171.00      Manufacturer#2
6yjMQN5hxnR247h5gL73i4M 20-
293-265-8784
zKp4R7nPwwgWRm2ig3Ln7RNS0Bgxni721kSCCBPxSx32
NL2R2j1P00wA BO k1PzM 7jC462gz5wyxw m3R1CQj
9998.14        Supplier#000963957
```

```
IRAN
4713952.00      Manufacturer#5
5C35i6P1cywn6N4 wkj6Scxz 20-
755-335-4707
QQ R5gPQi5CAyCCC 15mihkQ11qiShyskBL0nzNwR0
ZOzgnOg4PR3PkOk06LYa6 ASwi3nBnAnNh5 M6BQ
1zi
9998.05        Supplier#000545616
```

```
EGYPT
7545615.00      Manufacturer#3
Rzc4yR3A11zgCgh 14-
510-874-4706
gPlgzC4ny5wRM5g6z1B1ynQnnN43zP77h11ck114A2Ry
wkszOS51Q BM3iPNA2w0xO 7 wS4P63z
```

100 rows processed.

-- Query 16 (Original) Q16

```
SELECT
P_BRAND,
P_TYPE,
P_SIZE,
COUNT(DISTINCT PS_SUPPKEY) AS SUPPLIER_CNT
FROM PARTSUPP, PARTS
WHERE P_PARTKEY = PS_PARTKEY
AND P_BRAND <> 'Brand#33'
AND P_TYPE NOT LIKE 'PROMO ANODIZED%'
AND P_SIZE IN (37, 50, 40, 36, 48, 10, 8, 43)
```

Appendix C

```
AND PS_SUPPKEY NOT IN (SELECT S_SUPPKEY
FROM SUPPLIER
WHERE S_COMMENT LIKE '%Better Business
Bureau%Complaints%')
GROUP BY P_BRAND, P_TYPE, P_SIZE
ORDER BY SUPPLIER_CNT DESC, P_BRAND, P_TYPE,
P_SIZE
```

P_BRAND	P_TYPE	P_SIZE
SUPPLIER_CNT		
Brand#11	ECONOMY PLATED STEEL	37.00
600.00		
Brand#44	SMALL POLISHED NICKEL	48.00
584.00		
Brand#54	LARGE BURNISHED STEEL	36.00
583.00		
Brand#11	PROMO BURNISHED STEEL	8.00
580.00		
Brand#15	ECONOMY ANODIZED NICKEL	48.00
580.00		
Brand#21	PROMO POLISHED TIN	36.00
579.00		
Brand#32	STANDARD BURNISHED STEEL	36.00
576.00		
Brand#55	LARGE BURNISHED STEEL	10.00
576.00		
Brand#25	ECONOMY POLISHED BRASS	36.00
574.00		
Brand#15	SMALL ANODIZED STEEL	36.00
572.00		

27840 rows processed.

```
-- Query 14 (Variant C)          Q14      C
```

```
CREATE TABLE ALL_SALES0 (TYPE VARCHAR2(25),
AMOUNT NUMBER(20, 2))
tablespace TS_S
pctfree 0
pctused 99
storage (initial 5m next 5m pctincrease 0)
```

```
CREATE TABLE SUM_PROMO_SALES0 (PROMO_AMOUNT
NUMBER(20, 2))
tablespace TS_S
pctfree 0
pctused 99
storage (initial 5m next 5m pctincrease 0)
```

```
CREATE TABLE SUM_ALL_SALES0 (ALL_AMOUNT
NUMBER(20, 2))
tablespace TS_S
pctfree 0
pctused 99
storage (initial 5m next 5m pctincrease 0)
```

```
INSERT INTO ALL_SALES0
SELECT P_TYPE, SUM(L_EXTENDEDPRI * (1-
L_DISCOUNT))
FROM LINEITEM, PARTS
WHERE L_PARTKEY = P_PARTKEY
AND L_SHIPDATE >= TO_DATE('1995-03-
01', 'YYYY-MM-DD')
AND L_SHIPDATE < ADD_MONTHS(TO_DATE('1995-
03-01', 'YYYY-MM-DD'), 1)
GROUP BY P_TYPE
```

```
INSERT INTO SUM_PROMO_SALES0
SELECT SUM(AMOUNT)
FROM ALL_SALES0
WHERE TYPE LIKE 'PROMO%'
```

```
INSERT INTO SUM_ALL_SALES0
SELECT SUM(AMOUNT)
FROM ALL_SALES0
```

```
SELECT 100.00 * PROMO_AMOUNT / ALL_AMOUNT AS
PROMO_REVENUE
FROM SUM_PROMO_SALES0, SUM_ALL_SALES0
```

```
PROMO_REVENUE
16.64
```

1 row processed.

```
DROP TABLE ALL_SALES0
```

```
DROP TABLE SUM_PROMO_SALES0
```

```
DROP TABLE SUM_ALL_SALES0
```

```
-- Query 8 (Variant B) Q8      B
```

```
SELECT
YEAR,
SUM(DECODE(NATION, 'IRAN', VOLUME, 0)) /
SUM(VOLUME) AS MKT_SHARE
FROM
(SELECT
TO_CHAR(O_ORDERDATE, 'YYYY') AS YEAR,
L_EXTENDEDPRI * (1 - L_DISCOUNT) AS VOLUME,
N2.N_NAME AS NATION
FROM PARTS, SUPPLIER, LINEITEM, ORDERS,
CUSTOMER, NATION N1, NATION N2,
REGION
WHERE P_PARTKEY = L_PARTKEY
AND S_SUPPKEY = L_SUPPKEY
AND L_ORDERKEY = O_ORDERKEY
AND O_CUSTKEY = C_CUSTKEY
AND C_NATIONKEY = N1.N_NATIONKEY
AND N1.N_REGIONKEY = R_REGIONKEY
AND R_NAME = 'MIDDLE EAST'
AND S_NATIONKEY = N2.N_NATIONKEY
AND O_ORDERDATE BETWEEN TO_DATE ('1995-01-
01', 'YYYY-MM-DD') AND
TO_DATE('1996-12-31', 'YYYY-MM-DD')
AND P_TYPE = 'LARGE ANODIZED NICKEL'
) ALL_NATIONS
GROUP BY YEAR
ORDER BY YEAR
```

```
YEAR MKT_SHARE
1995 0.04
1996 0.04
```

2 rows processed.

```
-- Query 12 (Variant B)          Q12      B
```

```
SELECT
L_SHIPMODE,
SUM(DECODE(O_ORDERPRIORITY, '1-URGENT', 1, '2-
HIGH', 1, 0)) AS HIGH_LINE_COUNT,
SUM(DECODE(O_ORDERPRIORITY, '1-URGENT', 0, '2-
HIGH', 0, 1)) AS LOW_LINE_COUNT
FROM
ORDERS,
LINEITEM
WHERE O_ORDERKEY = L_ORDERKEY
AND L_SHIPMODE IN ('TRUCK', 'FOB')
AND L_COMMITDATE < L_RECEIPTDATE
AND L_SHIPDATE < L_COMMITDATE
AND L_RECEIPTDATE >= TO_DATE('1996-01-
01', 'YYYY-MM-DD')
AND L_RECEIPTDATE <
ADD_MONTHS(TO_DATE('1996-01-01', 'YYYY-MM-
DD'), 12)
GROUP BY L_SHIPMODE
ORDER BY L_SHIPMODE
```

L_SHIPMODE	HIGH_LINE_COUNT	LOW_LINE_COUNT
FOB	626085.00	938294.00
TRUCK	623792.00	939515.00

2 rows processed.

-- Query 17 (Original) Q17

```
SELECT
SUM(L_EXTENDEDPRIICE)/7.0 AS AVG_YEARLY
FROM LINEITEM, PARTS
WHERE P_PARTKEY = L_PARTKEY
AND P_BRAND = 'Brand#33'
AND P_CONTAINER = 'WRAP BOX'
AND L_QUANTITY < (SELECT
0.2* AVG(L1.L_QUANTITY)
FROM LINEITEM L1
WHERE L1.L_PARTKEY = P_PARTKEY)

AVG_YEARLY
31939330.38
```

1 row processed.
Setting the number of rows to fetch to: 10

-- Query 3 (Original) Q3

```
SELECT
L_ORDERKEY,
SUM(L_EXTENDEDPRIICE*(1-L_DISCOUNT)) AS
REVENUE,
TO_CHAR(O_ORDERDATE, 'YYYY-MM-DD'),
O_SHIPPRIORITY
FROM CUSTOMER, ORDERS, LINEITEM
WHERE C_MKTSEGMENT = 'FURNITURE'
AND C_CUSTKEY = O_CUSTKEY
AND L_ORDERKEY = O_ORDERKEY
AND O_ORDERDATE < TO_DATE('1995-03-
16', 'YYYY-MM-DD')
AND L_SHIPDATE > TO_DATE('1995-03-16', 'YYYY-
MM-DD')
GROUP BY L_ORDERKEY, O_ORDERDATE,
O_SHIPPRIORITY
ORDER BY REVENUE DESC, O_ORDERDATE
```

L_ORDERKEY	REVENUE	O_SHIPPRIORITY	O_ORDERDATE
33655462.00	523031.55		1995-03-03
0.00			
231686086.00	487759.42		1995-03-05
0.00			
486605252.00	480690.71		1995-03-15
0.00			
391442626.00	480573.26		1995-02-28
0.00			
57909345.00	476835.45		1995-03-07
0.00			
420735010.00	475052.46		1995-02-04
0.00			
484101378.00	474720.47		1995-02-20
0.00			
554876450.00	465084.05		1995-03-10
0.00			
188730631.00	463203.23		1995-03-04
0.00			
203285025.00	461981.49		1995-02-17
0.00			

10 rows processed.
Statement Processed in 1210.14 seconds.

-- Query 5 (Original) Q5

```
SELECT
N_NAME,
SUM(L_EXTENDEDPRIICE*(1-L_DISCOUNT)) AS
REVENUE
FROM CUSTOMER, ORDERS, LINEITEM, SUPPLIER,
NATION, REGION
WHERE C_CUSTKEY = O_CUSTKEY
AND O_ORDERKEY = L_ORDERKEY
AND L_SUPPKEY = S_SUPPKEY
```

```
AND C_NATIONKEY = S_NATIONKEY
AND S_NATIONKEY = N_NATIONKEY
AND N_REGIONKEY = R_REGIONKEY
AND R_NAME = 'ASIA'
AND O_ORDERDATE >= TO_DATE('1995-01-
01', 'YYYY-MM-DD')
AND O_ORDERDATE < ADD_MONTHS(TO_DATE('1995-
01-01', 'YYYY-MM-DD'), 12)
GROUP BY N_NAME
ORDER BY REVENUE DESC
```

N_NAME	REVENUE
VIETNAM	5362466500.49
CHINA	5304508516.19
INDIA	5279722359.25
INDONESIA	5279067456.37
JAPAN	5259455185.53

5 rows processed.

-- Query 13 (Original) Q13

```
SELECT
YEAR,
SUM(REVENUE) AS REVENUE
FROM
(SELECT
TO_CHAR(O_ORDERDATE, 'YYYY') AS YEAR,
L_EXTENDEDPRIICE * (1-L_DISCOUNT) AS REVENUE
FROM LINEITEM, ORDERS
WHERE O_ORDERKEY = L_ORDERKEY
AND O_CLERK = 'Clerk#00000431'
AND L_RETURNFLAG = 'R'
) PERFORMANCE
GROUP BY YEAR
ORDER BY YEAR
```

YEAR	REVENUE
1992	17002547.10
1993	17534799.20
1994	14301025.42
1995	4195656.05

4 rows processed.

-- Query 7 (Original) Q7

```
SELECT
SUPP_NATION,
CUST_NATION,
YEAR,
SUM(VOLUME) AS REVENUE
FROM
(SELECT
N1.N_NAME AS SUPP_NATION,
N2.N_NAME AS CUST_NATION,
TO_CHAR(L_SHIPDATE, 'YYYY') AS YEAR,
L_EXTENDEDPRIICE * (1-L_DISCOUNT) AS VOLUME
FROM SUPPLIER, LINEITEM, ORDERS, CUSTOMER,
NATION N1, NATION N2
WHERE S_SUPPKEY = L_SUPPKEY
AND O_ORDERKEY = L_ORDERKEY
AND C_CUSTKEY = O_CUSTKEY
AND S_NATIONKEY = N1.N_NATIONKEY
AND C_NATIONKEY = N2.N_NATIONKEY
AND ((N1.N_NAME = 'IRAN' AND N2.N_NAME =
'JAPAN') OR
(N1.N_NAME = 'JAPAN' AND N2.N_NAME =
'IRAN'))
AND L_SHIPDATE BETWEEN TO_DATE('1995-01-
01', 'YYYY-MM-DD') AND
TO_DATE('1996-12-31', 'YYYY-MM-DD')
) SHIPPING
GROUP BY SUPP_NATION, CUST_NATION, YEAR
ORDER BY SUPP_NATION, CUST_NATION, YEAR
```

SUPP_NATION	CUST_NATION	YEAR	REVENUE

Appendix C

```
IRAN          JAPAN
1995 5241591238.61
IRAN          JAPAN
1996 5282135495.90
JAPAN        IRAN
1995 5243825176.18
JAPAN        IRAN
1996 5274453992.46
```

4 rows processed.

-- Query 9 (Original) Q9

```
SELECT
NATION,
YEAR,
SUM(AMOUNT) AS SUM_PROFIT
FROM
(SELECT
N_NAME AS NATION,
TO_CHAR(O_ORDERDATE, 'YYYY') AS YEAR,
L_EXTENDEDPRI * (1-L_DISCOUNT) -
PS_SUPPLYCOST * L_QUANTITY AS AMOUNT
FROM PARTS, SUPPLIER, LINEITEM, PARTSUPP,
ORDERS, NATION
WHERE S_SUPPKEY = L_SUPPKEY
```

```
AND PS_SUPPKEY = L_SUPPKEY
AND PS_PARTKEY = L_PARTKEY
AND P_PARTKEY = L_PARTKEY
AND O_ORDERKEY = L_ORDERKEY
AND S_NATIONKEY = N_NATIONKEY
AND P_NAME LIKE '%khaki%'
) PROFIT
GROUP BY NATION, YEAR
ORDER BY NATION, YEAR DESC
```

NATION	YEAR	SUM_PROFIT
ALGERIA	1998	2663786363.03
ALGERIA	1997	4568198128.11
ALGERIA	1996	4571901633.09
ALGERIA	1995	4574597299.56
ALGERIA	1994	4538956616.09
ALGERIA	1993	4555486757.11
ALGERIA	1992	4567189611.84
ARGENTINA	1998	2664412237.49
ARGENTINA	1997	4539530906.13
ARGENTINA	1996	4576185093.72

175 rows processed.

Queries processed: 17

Appendix D: Seed and Query Substitution Parameters

-- using 857038079 as a seed to the RNG

-- Query 1 Original

DELTA=86

-- Query 2 (Variant A)

SIZE=22
TYPE=NICKEL
REGION=MIDDLE EAST

-- Query 3

SEGMENT=FURNITURE
DATE=1995-03-16

-- Query 4 (Variant B)

DATE=1995-02-01

-- Query 5

REGION=ASIA
DATE=1995-01-01

-- Query 6

SHIPDATE=1995-01-01
DISCOUNT=0.01
QUANTITY=25

-- Query 7 (Original)

NATION1=IRAN
NATION2=JAPAN

-- Query 8 (Variant B)

NATION=IRAN
REGION=MIDDLE EAST
TYPE=LARGE ANODIZED NICKEL

-- Query 9 (Original)

COLOR=khaki

-- Query 10

DATE=1993-12-01

-- Query 11 (Variant A)

NATION=IRAN
FRACTION=0.000001

-- Query 12 (Variant B)

SHIPMODE1=TRUCK
SHIPMODE2=FOB
DATE=1996-01-01

-- Query 13 (Original)

CLERK=Clerk#000000431

-- Query 14 (Variant C)

DATE=1995-03-01

-- Query 15 (Variant B)

DATE=1995-02-01

-- Query 16

BRAND=Brand#33
TYPE=PROMO ANODIZED
SIZE IN (37, 50, 40, 36, 48, 10, 8, 43)

-- Query 17 (Original)

BRAND=Brand#33
CONTAINER=WRAP BOX

Appendix E: Implementation Specific Layer/Source Code

The following table lists the source code files used in the benchmark runs.

Implementation-Specific Driver Files

File Name	Description
qexecpl.c	SQL execution engine, OCI version
qexecpl.h	SQL statement execution front-end header file
gettime.c	Get wall clock and CPU time
uf1.pc	TPC-D benchmark UF1 driver, PRO*C version. Also used to restore UF2.
uf2_del.pc	TPC-D benchmark uf2 driver, PRO*C version. Also used to restore UF1.
runuf1.sh	Executes UF1
runuf2.sh	Executes UF2
runpower1.sh	Single-stream TPC-D power test

qexecpl.c

```

=====
copyright (c) 1996 Oracle Corp, Redwood Shores, CA
OPEN SYSTEMS PERFORMANCE GROUP
All Rights Reserved
=====
FILENAME
qexecpl.c
DESCRIPTION
SQL Execution Engine, OCI version
MODIFIED
pswong 04/02/96 - more polishing
pswong 03/25/96 - polish up
pswong 03/06/96 - created
=====*/

#include <stdio.h>
#include <string.h>
#include <setjmp.h>
/*#include <sys/param.h>*/
#include <stdlib.h>
#include <errno.h>
#include <math.h>
#include <string.h>
#include <sys/types.h>
#include <time.h>

#include "qexecpl.h"

/* Function Prototypes */
extern double gettime();

/* function prototypes from gen.c */
int get_statement();

/* Declare error handling functions */
void sql_error();
void sql_error2();

/* other prototypes */
int define_output_variables();
void process_select_list();
void usage();
void sqlinit();
void SQLexec();
void SQLexit();
void *memalloc();
void print_header();
void print_rows();
int OFEN();
void remove_newline();

char logname[UNAME_LEN]; /* username/passwd combo */

double tr_start = 0.0; /* query start time */
double tr_end = 0.0; /* query end time */

double s_tr_start = 0.0; /* statement start time */
double s_tr_end = 0.0; /* statement end time */

/* For our purpose of timing, we will treat comments as delimiters
/* for queries. Thus, we will collect query timings whenever we
/* encounter a comment (of course not for the first comment in a
/* file).

int end_flag = 0; /* flag to indicate that we have reached */
/* the end of a query */

int stmt_cnt = 0; /* Number of statements processed. */
int qry_cnt = 0; /* Number of query processed. */

double product = 1.0; /* cumulative product of query times */
int rows_ret = 0; /* the number of rows fetched */
int num_sel_list = 0; /* the number of select list item */

long num_to_fetch = -1; /* Number of rows to fetch. -1 means fetch all */

slisttype slist[MAX_SEL_LIST]; /* Array for describing Select List */
dlisttype *dlist[MAX_SEL_LIST]; /* Array of ptrs for Defining Select List */

char stmt[SQL_LEN]; /* The SQL statement or comment line. */
char cmt[81]; /* Buffer to save the comment. */

FILE *qtmp = stdin; /* fd for query template */
FILE *logfile = stdout; /* log and report files */
FILE *rep = stdout;

void *defbuf; /* Buffer pointer for ODEFIN */
int deflen = 0; /* size of data type for ODEFIN */
int deftype = 1; /* Oracle type number for ODEFIN */

time_t tim; /* To get wall clock time */

#define tpclda;
#define curq;
unsigned long tpchda[256];

/* usage: prints the usage of the program */
void usage() {
    fprintf(stderr, "\nusage: qexec.ott username/password [q<path name for
    query template file>]\n");
    fprintf(stderr, " [l<path name for log>] [r<path name
    for reports>]\n");
    fprintf(stderr, "options:\n");
    fprintf(stderr, "q<path for query> : full path name for the
    query template file.\n");
    fprintf(stderr, " (default is stdin)\n");
    fprintf(stderr, "l<path name for log> : full path name for log
    files.\n");
    fprintf(stderr, " (default is stdout)\n");
    fprintf(stderr, "r<path name for reports> : full path name for
    reports.\n");
    fprintf(stderr, " (default is stdout)\n");
    exit(-1);
}

void sql_error(lda, cur)
    lda;
    cur;
{
    char msg[2048];
    if (cur->rc) {

```

Appendix E

```

    oerhms(lda, cur->rc, (text *) msg, 2048);
    fprintf(stderr, "Error: SQL Error encountered in Oracle!\n");
    fputs(msg, stderr);
}

/* Rollback just in case */
oro1(lda);

fprintf(stderr, "Exiting Oracle...\n");
fflush(stderr);

SQLexit();
}
exit(1);

void sql_error2(lda, cur)
    ldadef *lda;
    csrdef *cur;
{
    char msg[2048];

    if (cur->rc) {
        oerhms(lda, cur->rc, (text *) msg, 2048);
        fprintf(stderr, "Error: SQL Error encountered in Oracle!\n");
        fputs(msg, stderr);
    }

    fflush(stderr);
}

void main(argc, argv)
    int argc;
    char *argv[];
{
    int retcode; /* Return code for get_statement */
    /* Initialize some variables */
    if ((argc > 5) || (argc < 2)) {
        usage();
    }

    /* argv[1] -- User and Password for Database */
    strcpy(logname, argv[1]);

    /* Process optional parameters */
    argc -= 1;
    argv += 1;

    while(--argc) {
        ++argv;
        switch(argv[0][0]) {
            case 'q':
                if ((qtemp = fopen(++(argv[0]), "r")) == NULL) {
                    fprintf(stderr, "Unable to open file '%s'\n", argv[0]);
                    fprintf(stderr, "%s: %s\n", argv[0], strerror(errno));
                    exit(-1);
                }
                break;
            case 'r':
                if ((rep = fopen(++(argv[0]), "a")) == NULL) {
                    fprintf(stderr, "Unable to open file '%s'\n", argv[0]);
                    fprintf(stderr, "%s: %s\n", argv[0], strerror(errno));
                    exit(-1);
                }
                break;
            case 'l':
                if ((logfile = fopen(++(argv[0]), "a")) == NULL) {
                    fprintf(stderr, "Unable to open file '%s'\n", argv[0]);
                    fprintf(stderr, "%s: %s\n", argv[0], strerror(errno));
                    exit(-1);
                }
                break;
            default:
                fprintf(stderr, "Invalid option: %c\n", argv[0][0]);
                usage();
                break;
        }
    }

    /* Do some initialization and establish connection with the database */

    SQLinit();

    /* May want to add some triggering mechanism here */

    time(&tim);
    fprintf(logfile, "Begin Execution at %s\n\n", ctime(&tim));

    /* Get the next statement and start processing it */

    while ((retcode = get_statement()) > 0) {
        switch (retcode) {
            /* If this is a comment, skips it */
            case COMMENT:
                if (!end_flag) {
                    end_flag = 0; /* reset query end flag */
                    /* save the comment so that we can print it out later on */
                    strcpy(cmnt, stmt);
                    break;
                }
                fprintf(logfile, "%s", stmt);
                fprintf(rep, "%s", stmt);
                break;
            /* if this is a set_row_fetch command */
            case SET_FETCHROW:
                fprintf(logfile, "Setting the number of rows to fetch to: %ld\n",
                    num_to_fetch);
                break;
            /* if this is a SQL statement */
            case SQL_STMT:
                /* Executes the query */
                SQLexec();
        }

        s_tr_end = gettimeofday();
        stmt_cnt++;

        /*
         * fprintf(logfile, "\nStatement Started at %.2f\n", s_tr_start);
         * fprintf(logfile, "Statement Ended at %.2f\n", s_tr_end);
         * fprintf(logfile, "Statement Processed in %.2f seconds.\n",
         *     (s_tr_end - s_tr_start));
         */
        break;

        /* should never reach here */
        default:
            fprintf(stderr, "Invalid statement type!!\n");
            SQLexit();
            break;
    }
}

/* Get Timing for the last query */
tr_end = gettimeofday();

time(&tim);
fprintf(logfile, "\nEnded Executing this Query at %s\n", ctime(&tim));
fprintf(logfile, "\nQuery Started at %.2f\n", tr_start);
fprintf(logfile, "Query Ended at %.2f\n", tr_end);
fprintf(logfile, "Query Processed in %.2f seconds\n",
    (tr_end - tr_start));

fprintf(rep, "%.2f\n", (tr_end - tr_start));

fprintf(logfile, "\nSQL statements processed: %d\n", stmt_cnt);
fprintf(logfile, "Queries processed: %d\n", qry_cnt);

fflush(rep);
fflush(logfile);

/* Close the query template file */
fclose(qtemp);

/* Disconnect from ORACLE. */
SQLexit();
exit(0);
}

/*
 * SQLinit(): Perform initialization tasks.
 *
 * Logs on to Oracle, opens some files and open a cursor for
 * later use.
 */

void SQLinit() {
    int i;

    /* preallocate MAX_PREALLOC members of the dlist array
     * initializes others to NULL so that we can determine who to free
     * later */
    for (i=0; i<MAX_SEL_LIST; i++) {
        if (i < MAX_PREALLOC)
            dlist[i] = (dlist *) memalloc(sizeof(dlist));
        else
            dlist[i] = NULL;
    }

    /* Connect to ORACLE. Program will call sql_error()
     * if an error occurs in connecting to the default database. */
    if (orlon(&tpclda, (ub1 *)tpchda, (text *)logname, -1, (text *)0, -1,
        0)) {
        if (olog(&tpclda, (ub1 *)tpchda, (text *)logname, -1, (text *)0, -1,
            (text *)0, -1, (ub4)OCI_LM_NBL)) {
            fprintf(stderr, "Error: Failed to log on\n");
            sql_error(&tpclda, &tpclda);
            exit(-1);
        }
    }

    printf("\nConnected to ORACLE as user: %s\n\n", logname);

    /* open a cursor for us to use */
    OOPEN(&tpclda, &curq);
}

/*
 * SQLexec() Executes the SQL statement.
 *
 * Parse the SQL statement.
 *
 * If DDL or DML statements, execute right away.
 *
 * Else describe and define select list outputs,
 *
 * execute and fetch results.
 */

void SQLexec()
{
    int i;

    if (!end_flag) {
        /* Clause 5.3.6.2: QI(i,s) is the time between the first character
         * of this query text is submitted and the first
         * character of the next query text is submitted.
         */
        tr_end = gettimeofday();

        if (qry_cnt) {
            time(&tim);
        }
    }
}

```

```

        fprintf(logfile, "\nEnded Executing this Query at %s\n",
ctime(&tim));
        fprintf(logfile, "\nQuery Started at %.2f\n", tr_start);
        fprintf(logfile, "Query Ended at %.2f\n", tr_end);
        fprintf(logfile, "Query Processed in %.2f seconds.\n\n",
(tr_end - tr_start));
    }
    fprintf(logfile, "-----\n\n");
    /* print comments for this query that we have saved */
    fprintf(logfile, "%s\n", cmnt);
    fprintf(rep, "%.2f\n", (tr_end - tr_start));
    fprintf(rep, "%s", cmnt);
    fprintf(logfile, "\nBegan Executing this Query at %s\n",
ctime(&tim));
    /* Let's fflush stuff so that we can see what's going on */
    fflush(logfile);
    fflush(rep);
}
tr_start = tr_end;
qry_cnt++;
end_flag = 1;
}
s_tr_start = gettime();
/* parse the query, use no defer option to determine the statement
type */
/* and catch any syntax error before we progress further.
*/
OPARSE(&tpclda, &curq, stmt, NA, FALSE, VER7);
/* Prints the query text to the logfile */
fprintf(logfile, "\n%s\n", stmt);
/* if this is a DDL or DML statement, execute it right away */
/* only worries about SELECT statements right now, cannot */
/* execute a stored PL/SQL procedure in this version */
if (curq.ft != 4) {
    OEXEC(&tpclda, &curq);
    return;
}
/* otherwise, this is a select statement */
/* Describe and define output variables */
num_sel_list = define_output_variables();
/* Executes the query and fetches the rows */
(void) process_select_list(num_sel_list);
/* Need to get the number of rows fetched first */
/* since the following statements will screw it up */
rows_ret = curq.rpc;
/* To control memory usage, let's free up the extra dlist entries */
/* that we have allocated. */
i=MAX_PREALLOC;
while(dlist[i] != NULL) {
    free(dlist[i]);
    dlist[i++] = NULL;
}
/* reset set_fetchrows */
num_to_fetch = -1;
}
void SQLexit() {
    int i;
    oclose(&curq);
    ologof(&tpclda);
    /* free all memory */
    for (i=0; i<MAX_SEL_LIST; i++) {
        if (dlist[i] != NULL) {
            free(dlist[i]);
            dlist[i] = NULL;
        }
    }
    /* Flush all output */
    fflush(rep);
    fflush(logfile);
}
/* define_output_variables(): Describe and define select-list items for
a query statement.
Returns the number of select-list items
for this query.
*/
int define_output_variables()
{
    int i;
    for (i=0; i<MAX_SEL_LIST; i++) {
        slist[i].buflen = MAX_COLNAME_SIZE;
        if (odescr(&curq, POS(i), (sb4 *) &slist[i].dbsize,
            (sb2 *) &slist[i].dbtype, (sb1 *) &slist[i].buf[0],
            (sb4 *) &slist[i].buflen, (sb4 *) &slist[i].dsz,
            (sb2 *) &slist[i].precision, (sb2 *) &slist[i].scale,
            (sb2 *) &slist[i].nullok)) {
            if (curq.rc == END_OF_LIST)
                break;
            else {
                sql_error(&tpclda, &curq);
                return -1;
            }
        }
        /* For formatting purpose, remove trailing blanks in select-list
name. */
        if (slist[i].buflen < MAX_COLNAME_SIZE)
            (slist[i].buf)[slist[i].buflen] = '\0';
        /* Well, we need to allocate for entries for dlist */
        if (i >= MAX_PREALLOC)
            dlist[i] = (dtype *) memalloc(sizeof(dtype));
        /* Let's check the sizes and types for this select list item */
        switch (slist[i].dbtype) {
        case NUM_TYPE:
            /* see if it is an integer or a floating point number */
            slist[i].dbsize = NUMWIDTH;
            /* The odescr will not give a good estimate to the scale if */
            /* no scale was given in the Oracle table definition. */
#ifdef HAVE_SCALE
            if (slist[i].scale != 0) {
                defbuf = (double *) dlist[i]->fbuf;
                deflen = FLT;
                deftype = FLT_TYPE;
                slist[i].dbtype = FLT_TYPE;
            } else {
                defbuf = (int *) dlist[i]->ibuf;
                deflen = INT;
                deftype = INT_TYPE;
                slist[i].dbtype = INT_TYPE;
            }
#else
            defbuf = (double *) dlist[i]->fbuf;
            deflen = FLT;
            deftype = FLT_TYPE;
            slist[i].dbtype = FLT_TYPE;
#endif
            /* default is character string */
            break;
        default:
            /* default is character string */
            defbuf = (char **) dlist[i]->sbuf;
            deflen = MAX_STR_LEN;
            deftype = STR_TYPE;
            break;
        }
        /* Define the column */
        ODEFIN(&tpclda, &curq, POS(i), defbuf, deflen, deftype, NA,
            (ub2 *)0, (text *)0, NA, NA, (ub2 *)dlist[i]->r1en, (ub2
*)0);
    }
    return i;
}
/* process_select_list(): Fetch rows from a query. */
void process_select_list(num)
int num; /* number of select list items */
{
    int ntf;
    /* Print the headers for the query execution result */
    print_header(num);
    /* See if we need to limit the rows to fetch */
    ntf = (num_to_fetch >= 0) ? num_to_fetch : MAX_ARRAY;
    /* Fetch the rows and print them out */
    if ((ntf > MAX_ARRAY) || (num_to_fetch == -1)) {
        OXFET(&tpclda, &curq, MAX_ARRAY, 0, 0);
        rows_ret = curq.rpc;
        print_rows(num, rows_ret);
        /* To avoid 1022 from OFEN */
        /* More rows to fetch... */
        if (curq.rc != 1403) {
            if (num_to_fetch == -1) {
                while (OFEN(&tpclda, &curq, MAX_ARRAY) != -1) {
                    print_rows(num, (curq.rpc-rows_ret));
                    rows_ret = curq.rpc;
                }
                /* print the final rows */
                print_rows(num, (curq.rpc-rows_ret));
                rows_ret = curq.rpc;
            } else {
                ntf -= MAX_ARRAY;
                while (OFEN(&tpclda, &curq, ((ntf>MAX_ARRAY) ?
MAX_ARRAY:ntf)) != -1) {
                    ntf -= MAX_ARRAY;
                    print_rows(num, (curq.rpc-rows_ret));
                    rows_ret = curq.rpc;
                    if (ntf <= 0) break;
                }
                print_rows(num, (curq.rpc-rows_ret));
                rows_ret = curq.rpc;
            }
        }
    }
}

```

Appendix E

```

    }
  } else {
    OEXFET(&tpclda,&curq,ntf,0,0);
    rows_ret = curq.rpc;
    print_rows(num,rows_ret);
  }
}
fprintf(logfile, "\n\n%d row%c processed.\n", rows_ret,
        rows_ret == 1 ? '\0' : 's');
}

int OFEN(lda,cur,nrows)
{
  ldadef lda;
  csrdef *cur;
  int nrows;
  {
    if(ofen(cur, nrows))
      if (cur->rc != 1403) {
        sql_error(lda,cur);
        return 1;
      } else
        return -1;
    }
  return 1;
}

int get_statement()
{
  char line[128];
  char *pos, *str;

  /* Reset statement buffer */
  stmt[0] = '\0';

  while (fgets(line, 127, qtemp) != NULL) {

    /* skip blank lines */
    if (line[0] == '\n')
      continue;

    /* remove blanks */
    str = line;

    while (*str == ' ') str++;

    /* Let's get the line together first */
    strcat(stmt, str);

    /* if this is a comment line */
    if ((str[0] == '-') && (str[1] == '-'))
      return COMMENT;

    /* see if this is a set_fetchrows line */
    if (strncmp(str, "set_fetchrows", 13) == 0) {
      pos = strchr(str, ';');
      *pos = '\0';
      pos = strchr(str, '=');
      num_to_fetch = atoi(++pos);
      return SET_FETCHROW;
    }

    /* if this is the end of the current statement */
    if ((pos = strchr(stmt, ';')) != NULL) {
      *pos = '\0';
      return SQL_STMT;
    }
  }
  return END_OF_FILE;
}

/* memalloc(): Allocates memory, exit program if we have a problem. */
void *memalloc(size)
{
  int size;

  void *tmp;

  if ((tmp = (void *) malloc(size)) == NULL) {
    fprintf(stderr, "Error in malloc\n");
    SQLexit(0);
    return NULL; /* should never reach here */
  } else {
    return tmp;
  }
}

void print_header(nsel)

```

qexecpl.h

```

/*-----
   Copyright (c) 1995 Oracle Corp, Redwood Shores, CA
   OPEN SYSTEMS PERFORMANCE GROUP
   All Rights Reserved
-----*/

FILENAME
qstream.h
DESCRIPTION
SQL statement execution front-end header file.
MODIFIED
pswong 03/07/96 - created
-----*/

#ifndef QSTREAMPL_H
#define QSTREAMPL_H

#include <stdio.h>
#include <string.h>
/*#include <sys/param.h>*/
#include <sys/types.h>
#include <time.h>

```

```

int nsel; /* Number of select list items */
{
  int i;
  int len = 0; /* Running column length */
  int cwid = 0;

  fprintf(logfile, "\n");
  for (i=0; i<nsel; i++) {
    /* format the output a little */
    cwid = MAX(slist[i].dbsize, slist[i].buflen);
    /* do a little bit of formatting */
    if (cwid > 80) {
      fprintf(logfile, "\n");
      len = 0;
    } else if ((len += cwid) > 80) {
      fprintf(logfile, "\n");
      len = cwid;
    }
    #ifndef FORMAT1
    if ((slist[i].dbtype == INT_TYPE) || (slist[i].dbtype == FLT_TYPE))
      fprintf(logfile, "%s ", cwid, slist[i].buf);
    else /* string type */
      fprintf(logfile, "%s ", -cwid, slist[i].buf);
    #else
    fprintf(logfile, "%s ", -cwid, slist[i].buf);
    #endif /* FORMAT1 */
  }
  fprintf(logfile, "\n");
}

void print_rows(ncol, nrow)
{
  int ncol;
  int nrow;
  {
    int i,j;
    int len;
    int cwid;

    for (i=0; i<nrow; i++) {
      len = 0;
      for (j=0; j<ncol; j++) {
        cwid = MAX(slist[j].dbsize, slist[j].buflen);
        /* do a little bit of formatting */
        if (cwid > 80) {
          fprintf(logfile, "\n");
          len = 0;
        } else if ((len += cwid) > 80) {
          fprintf(logfile, "\n");
          len = cwid;
        }

        switch(slist[j].dbtype) {
          case INT_TYPE:
            #ifdef HAVE_SCALE
            fprintf(logfile, "%*1d", cwid, (dlist[j]->ibuf)[i]);
            break;
            #endif /* HAVE_SCALE */
          case FLT_TYPE:
            #ifdef FORMAT1
            fprintf(logfile, "%*.2f ", cwid, (dlist[j]->fbuf)[i]);
            #else
            fprintf(logfile, "%*.2f ", -cwid, (dlist[j]->fbuf)[i]);
            #endif /* FORMAT1 */
            break;
          default:
            fprintf(logfile, "%s ", -(cwid), (dlist[j]->sbuf)[i]);
            break;
        }
      }
      fprintf(logfile, "\n");
    }
  }

  /* remove_newline(): Remove newline character from str. */
  void remove_newline(str)
  char *str;
  {
    char *p;

    while ((p = strchr(str, '\n')) != NULL)
      *p = ',';
  }
}

```

Appendix E

```

#endif /* TRUE */

#ifndef FALSE
#define FALSE 1
#endif /* FALSE */

#define MAX(x,y) ((x >= y) ? x : y)
#define MIN(x,y) ((x <= y) ? x : y)

/* defines and typedefs for parsing */
#define CRT_TBL 1
#define INS_STMT 3
#define SEL_STMT 4
#define UPD_STMT 5
#define DRP_VIEW 7
#define DRP_TBL 8
#define DEL_STMT 9
#define CRT_VIEW 10

/* defines and typedefs for query description */
#define MAX_COLNAME_SIZE 16 /* Maximum length of Column name */
#define MAX_SEL_LIST 16 /* Maximum items on a select list */
#define END_OF_LIST 1007 /* Error code when we reach the end of the
*/ /* select list.

/* types for describe */
#define CHAR_TYPE 1
#define NUM_TYPE 2
#define INT_TYPE 3
#define FLT_TYPE 4
#define STR_TYPE 5
#define DATE_TYPE 12

#define NUMWIDTH 16 /* width of the numeric fields */
#define POS(i) (i+1) /* The position is 1...n instead
#define IND(i) (i-1) /* of 0..n-1 as in an array.

typedef struct des
{
    sb4 dbsize;
    sb2 dbtype;
    sb1 buf[MAX_COLNAME_SIZE];
    sb4 buflen;
    sb2 dsize;
    sb2 precision;
    sb2 scale;
    sb2 nulltok;
} sltype;

/* defines and typedefs for query select list definition */
#define MAX_ARRAY 3 /* Maximum array size for array fetch */
#define MAX_STR_LEN 256 /* Maximum size for string variables
*/
#define MAX_PREALLOC 8 /* Maximum number of preallocated select list
*/ /* definitions.

/* Note: well, I don't want to waste too much memory at the start.
*/
/* Thus, I defined a MAX_PREALLOC so that we limit the amount
*/
/* of memory used for the dlist array.
*/
/* The program will dynamically allocate the remaining members
*/
/* of the dlist array if necessary.
*/

/* Actually, I can go one step further by making everything
*/
/* dynamic. But I don't want to call tons of mallocs and frees
*/
/* during the query execution. Another one of those memory-CPU
*/
/* trade-offs.....

#define INT sizeof(long)
#define STR sizeof(char)
#define FLT sizeof(double)

#define FLTP (double *)
#define INTP (long *)

#define STRP (char **)
typedef struct def
{
    long ibuf[MAX_ARRAY];
    double fbuf[MAX_ARRAY];
    char sbuf[MAX_ARRAY][MAX_STR_LEN];
    ub2 rlen[MAX_ARRAY]; /* return length */
} dlist;

extern int errno;

#define SQL_LEN 2048

#ifndef NULL
#define NULL 0
#endif

#ifndef NULLP
#define NULLP (void *)NULL
#endif /* NULLP */

#ifndef DISCARD
#define DISCARD (void)
#endif

#ifndef sword
#define sword int
#endif

#ifndef ub1
#define ub1 unsigned char
#endif

#define NA -1 /* ANSI SQL NULL */
#define VER7 2
#define NOT_SERIALIZABLE 8177 /* ORA-08177: transaction not
serializable */

#define ADR(object) ((ub1 *)&(object))
#define SZ(object) ((sword)sizeof(object))
#define SID(sid) ((sid == -1) ? 0 : sid)

/* For get_statement */
#define END_OF_FILE -1
#define COMMENT 1
#define SQL_STMT 2
#define SET_FETCHROW 3

#define OOPEN(lda, cursor)
    if (oopen((cursor), (lda), (text*)0, NA, NA, (text*)0, NA))\
        {sql_error(lda, cursor);} \
    else \
        DISCARD 0

#define DDEFIN(lda, cursor, pos, buf, buflen, ftype, scale, indp, fmt, ffmt1, ffmtt, rlen, rcode)
    if (odefin((cursor), (pos), (ub1*)(buf), (bufl), (ftype), (scale), \
        (indp), (text*)(fmt), (fmt1), (fmtt), (rlen), (rcode))\
        {sql_error(lda, cursor);} \
    else \
        DISCARD 0

#define OEXFET(lda, cursor, nrows, cancel, exact)
    if (oexfet((cursor), (nrows), (cancel), (exact)))\
        {if ((cursor)->rc == 1403) DISCARD 0; \
        else {sql_error(lda, cursor);} } \
    else \
        DISCARD 0

#define OPARSE(lda, cursor, sqlstm, sql1, defflg, lngflg)
    if (oparse((cursor), (text
*) (sqlstm), (sb4)(sql1), (defflg), (ub4)(lngflg)))\
        {sql_error(lda, cursor);} \
    else \
        DISCARD 0

#define OEXEC(lda, cursor)
    if (oexec((cursor)))\
        {sql_error(lda, cursor);} \
    else \
        DISCARD 0

#define ISOTXT "alter session set isolation_level = serializable"
#endif /* QSTREAMPL_H */

```

Appendix E

gettime.c

```
/*-----
   Copyright (c) 1996 Oracle Corp., Redwood Shores, CA
   OPEN SYSTEMS PERFORMANCE GROUP
   All Rights Reserved
-----*/

FILENAME
gttime.c
DESCRIPTION
wrapper for gettime
-----*/

#include <stdio.h>
#include <stdlib.h>

double gettime();

main() {
    printf("%f", gettime());
    exit(0);
}
```

uf1.pc

```
/*-----
   Copyright (c) 1995 Oracle Corp., Redwood Shores, CA
   OPEN SYSTEMS PERFORMANCE GROUP
   All Rights Reserved
-----*/

FILENAME
uf1.pc
DESCRIPTION
TPC-D benchmark uf1 driver, PRO*C version
MODIFIED
pswong 03/25/96 - more polish
pswong 02/28/96 - clean it up a little
pswong 04/21/95 - update to Spec 9.1, Use QGEN
-----*/

#include <stdio.h>
#include <string.h>
#include <setjmp.h>
/*#include <sys/param.h>*/
#include <errno.h>
#include <math.h>
#include <string.h>
/*#include <unistd.h>*/
#include <signal.h>
#include <stdlib.h>

#include "shared.h"

#define MAX_FILE_PATH_LEN 256
#define LINESIZE 256
#define LSIZE (LINESIZE-1)
#define UNAME_LEN 32

/* Include the SQL Communications Area. */
#include <sqlca.h>

/* Include the Oracle Communication Area */
EXEC SQL INCLUDE oraca;
EXEC ORACLE OPTION (ORACA=YES);

extern double gettime();

/* Declare error handling functions */
void sql_error();
void usage();
int get_ord_line();
int get_item_line();
char *nextpos();
void wakeup();

#define MAX_ITEMS 40
#define MAX_VNAME_LEN 30
#define MAX_INAME_LEN 30

/* batch size for inserts */
#define MAX_ORD_INSERTS 100
#define MAX_LINE_INSERTS (7 * MAX_ORD_INSERTS)

#ifdef NULL
#define NULL 0
#endif

/* structures to store query results */
/* typedefs in query.h

EXEC SQL BEGIN DECLARE SECTION;
VARCHAR lname[32]; /* UNAME_LEN username/passwd combo */

int linecnt = 0; /* used to specify the number of rows for array
inserts */
int ordcnt = 0; /* used to specify the number of rows for array
inserts */

/* declarations for ORDER */

long o_oke[100]; /*MAX_ORD_INSERTS;*/
long o_cke[100]; /*MAX_ORD_INSERTS;*/
varchar o_ostat[100][2]; /*MAX_ORD_INSERTS][2];*/
float o_tprice[100]; /*MAX_ORD_INSERTS];*/
varchar o_odate[100][13]; /*MAX_ORD_INSERTS][DATE_LEN];*/
varchar o_opri[100][15+1]; /*MAX_ORD_INSERTS][O_OPRIO_LEN+1];*/
varchar o_clk[100][15+1]; /*MAX_ORD_INSERTS][O_CLRK_LEN+1];*/
int o_spri[100]; /*MAX_ORD_INSERTS];*/
varchar o_cmnt[100][79+1]; /*MAX_ORD_INSERTS][O_CMNT_MAX+1];*/

/* declarations for LINEITEM */

long l_oke[700]; /*MAX_LINE_INSERTS;*/
long l_pke[700]; /*MAX_LINE_INSERTS;*/
long l_skey[700]; /*MAX_LINE_INSERTS;*/
*/
```

```
int l_lnum[700]; /*MAX_LINE_INSERTS;*/
float l_quan[700]; /*MAX_LINE_INSERTS;*/
float l_eprice[700]; /*MAX_LINE_INSERTS;*/
float l_disc[700]; /*MAX_LINE_INSERTS;*/
float l_tax[700]; /*MAX_LINE_INSERTS;*/
varchar l_rflag[700][2]; /*MAX_LINE_INSERTS][2];*/
varchar l_lstat[700][2]; /*MAX_LINE_INSERTS][2];*/
varchar l_sdate[700][13]; /*MAX_LINE_INSERTS][DATE_LEN];*/
varchar l_cdate[700][13]; /*MAX_LINE_INSERTS][DATE_LEN];*/
varchar l_rdate[700][13]; /*MAX_LINE_INSERTS][DATE_LEN];*/
varchar l_sinst[700][25+1]; /*MAX_LINE_INSERTS][L_INST_LEN+1];*/
varchar l_smode[700][10+1]; /*MAX_LINE_INSERTS][L_SMODE_LEN+1];*/
varchar l_cmnt[700][44+1]; /*MAX_LINE_INSERTS][L_CMNT_MAX+1];*/

EXEC SQL END DECLARE SECTION;

int set_id; /* set_id, global within the driver */
int run_id; /* run_id, global */
int proc_no; /* process number, global */
double sf = 1.0; /* scale factor, global */
double tr_end = 0.0; /* query end time */
double tr_start = 0.0; /* query start time */
double product = 1.0; /* cumulative product of query times */
int recover = 0; /* recover from previous UF2? */

FILE *logfile; /* log and report files */
FILE *ordfile, *itemfile; /* input data files */

char itemline[LINESIZE]; /* temp storage for input LINEITEM row */
/* usage: prints the usage of the program */
void usage() {
    fprintf(stderr, "\nusage: uf1.o[st]t <set_id> <run_id> <proc_no>
<scale factor> [f: <path name for reports> u:uid/p:passwd] \n\n");
    fprintf(stderr, " set_id :the set id for this
update set\n");
    fprintf(stderr, " run_id :the run id for this
TPCD run\n");
    fprintf(stderr, " proc_no :the process number
within this update stream\n");
    fprintf(stderr, " scale factor :scale factor for the
run\n");
    fprintf(stderr, " r :indicates that this
is a recovery run\n");
    fprintf(stderr, " f: <path name for log> :full path name for
reports\n");
    fprintf(stderr, " u:uid/p:passwd :Username/Password
string - default is tcpd/tpcd\n");
    exit(-1);
}

void sql_error()
{
    /* ORACLE error handler */
    fprintf(stderr, "Error: SQL Error encountered in oracle!\n");
    fprintf(stderr, "\n\n%.70s\n", sqlca.sqlerrm.sqlerrmc);
    fflush(stderr);

    EXEC SQL WHENEVER SQLERROR CONTINUE;
    EXEC SQL ROLLBACK WORK;
    exit(1);
}

void main(argc, argv)
int argc;
char *argv[];
{
    int orows = 0;
    int lrows = 0;
    int notdone = 1;
    int err = 0;
    char ordpath[MAX_FILE_PATH_LEN];
    char itempath[MAX_FILE_PATH_LEN];
    char logpath[MAX_FILE_PATH_LEN];

    printf("begin uf1\n");
    /* Initialize some variables */

    strcpy((char *) lname, arr, "tpcd/tpcd");
    lname.len = strlen((char *) lname, arr);

    /* get the signal handler for the pause() */
    /* signal(SIGUSR1, wakeup); */
}
```



```

pos2 = nextpos(pos1);
strcpy((char *) o_odate[i].arr, pos1);
o_odate[i].len = strlen((char *) o_odate[i].arr);
pos1 = pos2;

pos2 = nextpos(pos1);
strcpy((char *) o_opri[i].arr, pos1);
o_opri[i].len = strlen((char *) o_opri[i].arr);
pos1 = pos2;

pos2 = nextpos(pos1);
strcpy((char *) o_clk[i].arr, pos1);
o_clk[i].len = strlen((char *) o_clk[i].arr);
pos1 = pos2;

pos2 = nextpos(pos1);
o_spri[i] = atoi(pos1);
pos1 = pos2;

pos2 = nextpos(pos1);
strcpy((char *) o_cmnt[i].arr, pos1);
o_cmnt[i].len = strlen((char *) o_cmnt[i].arr);
pos1 = pos2;
}
return (MAX_ORD_INSERTS);
}

```

```

int get_item_line (rowcnt)
int rowcnt;
{
int i=0, j; /* at least one row will be scanned */
char *pos1, *pos2;
int notyet = 1;
int okey;

/* for each orderkey, extract the corresponding rows in LINEITEM */
for (j=0; j<rowcnt; j++) {

/* the first row should have been scanned here by now */
/* either from the main program or from the previous */
/* invocation of this function */
notyet = 1;

while (notyet) {

/* extract columns from the line */
pos1 = nextpos(itemline);
okey = atoi(itemline);

/* Compare the ORDERKEY values */
if (okey == o_okey[j]) {

l_okey[i] = okey;

pos2 = nextpos(pos1);
l_pkey[i] = atoi(pos1);
pos1 = pos2;

pos2 = nextpos(pos1);
l_skey[i] = atoi(pos1);
pos1 = pos2;

pos2 = nextpos(pos1);
l_lnum[i] = atoi(pos1);
pos1 = pos2;

pos2 = nextpos(pos1);
l_quan[i] = atof(pos1);
pos1 = pos2;

pos2 = nextpos(pos1);
l_eprice[i] = atof(pos1);
pos1 = pos2;

pos2 = nextpos(pos1);
l_disc[i] = atof(pos1);
}
}
}

```

```

pos1 = pos2;

pos2 = nextpos(pos1);
l_tax[i] = atof(pos1);
pos1 = pos2;

pos2 = nextpos(pos1);
strcpy((char *) l_rflag[i].arr, pos1);
l_rflag[i].len = strlen((char *) l_rflag[i].arr);
pos1 = pos2;

pos2 = nextpos(pos1);
strcpy((char *) l_lstat[i].arr, pos1);
l_lstat[i].len = strlen((char *) l_lstat[i].arr);
pos1 = pos2;

pos2 = nextpos(pos1);
strcpy((char *) l_sdate[i].arr, pos1);
l_sdate[i].len = strlen((char *) l_sdate[i].arr);
pos1 = pos2;

pos2 = nextpos(pos1);
strcpy((char *) l_cdate[i].arr, pos1);
l_cdate[i].len = strlen((char *) l_cdate[i].arr);
pos1 = pos2;

pos2 = nextpos(pos1);
strcpy((char *) l_rdate[i].arr, pos1);
l_rdate[i].len = strlen((char *) l_rdate[i].arr);
pos1 = pos2;

pos2 = nextpos(pos1);
strcpy((char *) l_sinst[i].arr, pos1);
l_sinst[i].len = strlen((char *) l_sinst[i].arr);
pos1 = pos2;

pos2 = nextpos(pos1);
strcpy((char *) l_smode[i].arr, pos1);
l_smode[i].len = strlen((char *) l_smode[i].arr);
pos1 = pos2;

pos2 = nextpos(pos1);
strcpy((char *) l_cmnt[i].arr, pos1);
l_cmnt[i].len = strlen((char *) l_cmnt[i].arr);
pos1 = pos2;
} else {

/* reset so that we won't run into seg faults */
itemline[strlen(itemline)] = '\0';
break;
}

/* increments array index */
i++;

/* get next line, if failed, return */
if (fgets(itemline, LSIZE, itemfile) == NULL) {
return(i);
}
}
}
return (i);
}

```

```

char *nextpos(start)
char *start;
{
char *mark = strchr(start, '|');
*mark = '\0';
mark++;
return (mark);
}

```

```

void wakeup() {
return;
}

```

uf2_del.pc

```

-----
copyright (c) 1995 Oracle Corp. Redwood Shores, CA
OPEN SYSTEMS PERFORMANCE GROUP
All Rights Reserved
-----
FILENAME
uf1_del.pc
DESCRIPTION
TPC-D benchmark uf2 driver, PRO*C version
Also used to restore uf1.
MODIFIED
pswong 03/30/96 - polished version
pswong 06/13/95 - version 2
pswong 06/07/95 - created
-----*/

#include <stdio.h>
#include <string.h>
#include <setjmp.h>
/*#include <sys/param.h>*/
#include <errno.h>
#include <math.h>
#include <string.h>
#include <signal.h>
/*#include <unistd.h>*/
#include <stdlib.h>

#include "shared.h"

#define MAX_FILE_PATH_LEN 256
#define LINESIZE 256
#define LSIZE (LINESIZE-1)
#define UNAME_LEN 32

/* max length of SQL statement */
#define SQL_LEN 4096

```

```

/* Didn't use bind variables because oracle will perform FTS with the
*/
/* DELETE statement as it cannot estimate the key range of the BETWEEN
*/
/* statement.
*/

/* We need so many different statement types to avoid cases like
*/
/* _LORDERKEY BETWEEN x AND x which was known to have problems.
*/

#define DELO_TXT1 "DELETE FROM ORDERS WHERE O_ORDERKEY BETWEEN %ld AND
%ld "
#define DELL_TXT1 "DELETE FROM LINEITEM WHERE L_ORDERKEY BETWEEN %ld
AND %ld "
#define DELO_TXT2 "DELETE FROM ORDERS WHERE O_ORDERKEY = %ld "
#define DELL_TXT2 "DELETE FROM LINEITEM WHERE L_ORDERKEY = %ld "
#define EQO_TXT "OR O_ORDERKEY = %ld "
#define EQL_TXT "OR L_ORDERKEY = %ld "
#define DEL_LEN 80 /* a little larger than it's necessary for safety
sake */
#define BTWO_TXT "OR O_ORDERKEY BETWEEN %ld AND %ld "
#define BTWL_TXT "OR L_ORDERKEY BETWEEN %ld AND %ld "
#define BTW_LEN 50 /* also a little large than necessary */

/* batch size for inserts
*/
/* This is the number of BETWEEN statements that can fit in the SQL
*/
/* statement. The -1 in the end is just for safety.
*/

#define MAX_INS ((SQL_LEN - DEL_LEN) / BTW_LEN - 1)
#define MAX_INSERTS ((getenv("MAX_DELETE") == NULL) ? 2 :
((atoi(getenv("MAX_DELETE")) > MAX_INS) ? MAX_INS :
atoi(getenv("MAX_DELETE"))))

/* The frequency of commits. Default is to commit every 4 times. */
#define COMMIT_FREQ 4

```

Appendix E

```

/* Include the SQL Communications Area. */
#include <sqlca.h>
/* Include the Oracle Communication Area */
EXEC SQL INCLUDE oraca;
EXEC ORACLE OPTION (ORACA=YES);

extern double gettimeofday();
/* Declare error handling functions */

void sql_error();
void usage();
int get_ord_line();
void wakeup();

#ifdef NULL
#define NULL 0
#endif

/* structures to store query results */
/* typedefs in query.h */

EXEC SQL BEGIN DECLARE SECTION;

varchar sqlstmt[4096]; /*SQL_LEN;*/ /* the delete statement
for orders */
varchar sqlstmt2[4096]; /*SQL_LEN;*/ /* the delete statement
for lineitem */
VARCHAR lname[32]; /*UNAME_LEN;*/ /* username/passwd combo */

int rowcnt = 0; /* used to specify the number of rows for array
inserts */

EXEC SQL END DECLARE SECTION;

int set_id; /* set_id, global within the driver */
int run_id; /* run_id, global */
int proc_no; /* process number, global */
int u_flag=2; /* usage flag, global */
double sf = 1.0; /* scale_factor, global */

FILE *ordfile; /* flatfile for ORDERS */

/* usage: prints the usage of the program */
void usage() {
    fprintf(stderr, "\nUsage: uf2[st]t <set_id> <run_id> <proc_no> <scale
factor> <usage flag> [u<uid>/passwd] \n\n");
    fprintf(stderr, "      set_id      :the set id for this
update set\n");
    fprintf(stderr, "      run_id      :the run id for this
TPCD run\n");
    fprintf(stderr, "      proc_no     :the process number
within this update stream\n");
    fprintf(stderr, "      scale factor  :scale factor for the
run\n");
    fprintf(stderr, "      usage flag   :enter 1 for uf2 and 2
for restoring ufl\n");
    fprintf(stderr, "      u<uid>/passwd :Username/Password
string - default is tcpd/tpcd\n");
    exit(-1);
}

void sql_error()
{
    /* ORACLE error handler */
    fprintf(stderr, "Error: error encountered in oracle!\n");
    fprintf(stderr, "\n\n%.70s\n", sqlca.sqlerrm.sqlerrmc);
    fflush(stderr);

    EXEC SQL WHENEVER SQLERROR CONTINUE;
    EXEC SQL ROLLBACK WORK;
    exit(1);
}

void main(argc,argv)
int argc;
char *argv[];
{
    int orows = 0;
    int lrows = 0;
    int o, l;
    int notdone = 1;
    int num_exec = 0;
    /* int err = 0; */
    char ordpath[MAX_FILE_PATH_LEN];

    /* Initialize some variables */
    printf("Begin uf2\n");

    strcpy((char *) lname, "tpcd/tpcd");
    lname.len = strlen((char *)lname.arry);

    /* signal(SIGUSR1, wakeup); */

    if ((argc > 7) || (argc < 6)) {
        usage();
    }

    /* argv[1] -- Set ID */
    if ((set_id = atoi(argv[1])) < 0) {
        usage();
        exit(-1);
    }

    /* argv[2] -- Run ID */
    run_id = atoi(argv[2]);

    /* argv[3] -- Process Number */
    proc_no = atoi(argv[3]);

    /* argv[4] -- Scale Factor */
    sf = atof(argv[4]);

    /* argv[5] -- Filename prefix for orderkeys */
    /* The input file by default should be located */
    /* in $TPCDB/update/data. */
    u_flag = atoi(argv[5]);

    switch(u_flag) {
    case 1:
        sprintf(ordpath, "%s/delete.%d.%d",
            getenv("UPDATE_DIR"), set_id, proc_no);
        break;
    case 2:
        sprintf(ordpath, "%s/okey.u%d.%d",
            getenv("UPDATE_DIR"), set_id, proc_no);
        break;
    default:
        fprintf(stderr, "Illegal usage flag!\n");
        exit(-1);
    }
    printf("ordpath is %s\n",ordpath);

    /* Process optional parameters */

    argc -= 5;
    argv += 5;

    while(--argc) {
        ++argv;
        switch(argv[0][0]) {
        case 'u':
            strncpy((char *) lname.arry, ++argv[0], UNAME_LEN);
            lname.len = strlen((char *)lname.arry);
            if (strcmp((char *) lname.arry, "/") == NULL) {
                fprintf(stderr, "Login name must be in the format of
userid/passwd\n");
                usage();
                exit(-1);
            }
            break;
        default:
            fprintf(stderr, "Unknown argument %s\n", argv[0]);
            usage();
            break;
        }
    }

    /* open the files */
    if ( (ordfile = fopen(ordpath, "r")) == NULL ) {
        fprintf(stderr, "Unable to open file %s\n", ordpath);
        fprintf(stderr, "%s: %s\n", ordpath, strerror(errno));
        exit(-1);
    }

    /* Establish sql_error() as the error handler. */
    EXEC SQL WHENEVER SQLERROR DO sql_error();

    /* Connect to ORACLE. Program will call sql_error()
if an error occurs in connecting to the default database. */
    EXEC SQL CONNECT :lname;

#ifdef DEBUG
    printf("\nConnected to ORACLE as user: %s\n\n", lname.arry);
#endif /* DEBUG */

    /* pause(); */

    /* start the restore from ufl */

    /* delete from ORDERS, hold the commit */

    notdone = 1; /* just in case */

    while (notdone) {
        rowcnt = 0;
        if ((rowcnt = get_ord_line()) < MAX_INSERTS) {
            notdone = 0;
        }

        sqlstmt.len = strlen((char *)sqlstmt.arry);
        sqlstmt2.len = strlen((char *)sqlstmt2.arry);

        /* Execute the deletes, hold commit */

        if (rowcnt != 0) {
            printf("exec sql statement %s\n",sqlstmt.arry);
            EXEC SQL EXECUTE IMMEDIATE :sqlstmt;
            o = sqlca.sqlerrd[2];
            printf("exec sql statement %s\n",sqlstmt2.arry);
            EXEC SQL EXECUTE IMMEDIATE :sqlstmt2;
            l = sqlca.sqlerrd[2];
        }

        /* Commit once in a while */

        if (!(++num_exec)%COMMIT_FREQ) {
            EXEC SQL COMMIT WORK;
        }

        orows += o;
        lrows += l;
    }

    /* COMMIT all deletes */

    EXEC SQL COMMIT WORK;
    EXEC SQL WHENEVER SQLERROR DO sql_error();

    /* Print the number of rows processed when we are recovering */
    if (u_flag == 2) {
        fprintf(stdout, "%d row%ds deleted from ORDERS\n", orows,
            (orows == 1 ? "\0" : "s"));
        fprintf(stdout, "%d row%ds deleted from LINEITEM\n", lrows,
            (lrows == 1 ? "\0" : "s"));
    }

#ifdef DEBUG
    fprintf(stdout, "%d row%ds deleted from ORDERS\n", orows,
        (orows == 1 ? "\0" : "s"));
    fprintf(stdout, "%d row%ds deleted from LINEITEM\n", lrows,
        (lrows == 1 ? "\0" : "s"));
#endif /* DEBUG */
}

```

```

printf("%d rows deleted from ORDERS\n",orows,(orows == 1 ? "\0" :
"s"));
printf("%d rows deleted from LINEITEM\n",lrows,(orows == 1 ? "\0" :
"s"));
/* end of the restore */
EXEC SQL WHENEVER SQLERROR CONTINUE;
/* Disconnect from ORACLE. */
EXEC SQL COMMIT WORK RELEASE;
EXEC SQL WHENEVER SQLERROR DO sql_error();
fclose(ordfile);
exit(0);
}

int get_ord_line (void) {
int i;
int rcnt = 0;
long ord1, ord2;
char *pos1;
char line[LINESIZE];
char o_stmt[BTW_LEN];
char l_stmt[BTW_LEN];

printf("begin get_ord_line\n");
fflush(stdout);
for (i=0; i<MAX_INSERTS; i++) {
if (fgets(line, LSIZE, ordfile) == NULL) {
return (i);
}
}

printf("line is %s\n",line);
fflush(stdout);
/* extract columns from the line */
pos1 = strchr((char *) line, '|');
*pos1 = '\0';
pos1++;
ord1 = atoi(line);
*(strchr((char *) pos1, '|')) = '\0';

```

```

ord2 = atoi(pos1);
printf("i = %d i is either 0 or not\n",i);
fflush(stdout);
if (i == 0) {
/* prepare statement for ORDERS and LINEITEM */
printf("ord1 is %d ord2 is %d\n",ord1,ord2);
fflush(stdout);
if (ord1 == ord2) {
sprintf((char *) sqlstmt.arr, DELO_TXT2, ord1);
sprintf((char *) sqlstmt2.arr, DELL_TXT2, ord1);
} else {
sprintf((char *) sqlstmt.arr, DELO_TXT1, ord1, ord2);
sprintf((char *) sqlstmt2.arr, DELL_TXT1, ord1, ord2);
}
} else {
/* add next BETWEEN statement for ORDERS and LINEITEM */
printf("ord1 is %d ord2 is %d\n",ord1,ord2);
fflush(stdout);
if (ord1 == ord2) {
sprintf(o_stmt, EQO_TXT, ord1, ord2);
strcat((char *) sqlstmt.arr, o_stmt);
sprintf(l_stmt, EQL_TXT, ord1, ord2);
strcat((char *) sqlstmt2.arr, l_stmt);
} else {
sprintf(o_stmt, BTWO_TXT, ord1, ord2);
strcat((char *) sqlstmt.arr, o_stmt);
sprintf(l_stmt, BTWL_TXT, ord1, ord2);
strcat((char *) sqlstmt2.arr, l_stmt);
}
}
}

printf("end get_ord_line\n");
fflush(stdout);
return (MAX_INSERTS);
}

void wakeup() {
return;
}

```

runuf1.sh

```

#!/bin/sh
#
#=====  

# Copyright (c) 1995 Oracle Corp, Redwood Shores, CA  

# OPEN SYSTEM PERFORMANCE GROUP  

# All Rights Reserved  

#=====  

# FILENAME  

# runuf1.sh  

# DESCRIPTION  

# runuf1.sh -l [<path name for reports>] -u [<uid/passwd>]  

# -p [<program>] <run_id> <scale factor> <pair number>  

# <parallelism>  

# USAGE  

# To execute UF1.  

#=====  

O=${TPCD_UPDATE_DATA}  

TPCD_DIR=${O}/tpcd  

SCRIPT_DIR=${TPCD_DIR}/update/scripts  

GTIME_DIR=${TPCD_DIR}/source  

GTIME=${GTIME_DIR}/gtime

usage() {
echo ""
echo "runuf1.sh -l [<path name for reports>] -u [<uid/passwd>]"
echo " -p [<run_id> <scale factor> <pair number>]"
echo " <parallelism>"
echo ""
echo "run_id : Run_ID of this update run."
echo "scale factor : Scale Factor of the database."
echo "set number : The update pair number that this update function  

belongs to"
echo "parallelism : The parallelism of the updates."
echo ""
echo "-l : Path name for reports on the update function.  

Debug use for UF1 only"
echo "-u : Userid/Password for Oracle. Default is  

tpcd/tpcd."
echo "-h : To Display this message."
echo ""
}

LOGPATH=.  

PASSWD="tpcd/tpcd"

set -- `getopt "l:u:h" "$@"` || usage

while :
do
case "$1" in
-u) shift; PASSWD=$1;;
-l) shift; LOGPATH=$1;;
-h) usage; exit 0;;
--) shift; break;;
esac
shift;
done

if [ $# -lt 4 ]
then
usage
exit 1
fi

RUN_ID=$1  

SF=$2  

SETNUM=$3  

PAR=$4

if [ ${SF} -eq 0 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/tenthgig  

LOG_DIR=${TPCD_DIR}/update/log/tenthgig

```

```

DBF_DIR=${TPCD_DIR}/DBS/TENTHGIG
fi

if [ ${SF} -eq 1 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/onegig  

LOG_DIR=${TPCD_DIR}/update/log/onegig  

DBF_DIR=${TPCD_DIR}/DBS/ONEGIG
fi

if [ ${SF} -eq 3 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/threegig  

LOG_DIR=${TPCD_DIR}/update/log/threegig  

DBF_DIR=${TPCD_DIR}/DBS/THREEGIG
fi

if [ ${SF} -eq 10 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/tengig  

LOG_DIR=${TPCD_DIR}/update/log/tengig  

DBF_DIR=${TPCD_DIR}/DBS/TENGIG
fi

if [ ${SF} -eq 30 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/thirtygig  

LOG_DIR=${TPCD_DIR}/update/log/thirtygig  

DBF_DIR=${TPCD_DIR}/DBS/THIRTYGIG
fi

if [ ${SF} -eq 100 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/hundgig  

LOG_DIR=${TPCD_DIR}/update/log/hundgig  

DBF_DIR=${TPCD_DIR}/DBS/HUNDGIG
fi

if [ ${SF} -eq 300 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/threehundgig  

LOG_DIR=${TPCD_DIR}/update/log/threehundgig  

DBF_DIR=${TPCD_DIR}/DBS/THREEHUNDGIG
fi

if [ ${SF} -eq 1000 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/onetbyte  

LOG_DIR=${TPCD_DIR}/update/log/onetbyte  

DBF_DIR=${TPCD_DIR}/DBS/ONETBYTE
fi

i=1  

PID=""

# Let's figure out the extent sizes in Mb  

LINEEXT=`echo $SF * 720 * 0.001 / '$PAR + 1 | bc`  

LINEEXT=`echo $LINEEXT | sed '\.[0123456789]*/s//`'  

ORDEXT=`echo $SF * 200 * 0.001 / '$PAR + 1 | bc`  

ORDEXT=`echo $ORDEXT | sed '\.[0123456789]*/s//`'

# perform the update function 1

START=`$GTIME`

# first create the temp tables

$SQLPLUS $PASSWD <<!  

drop table temp_item;  

create table temp_item (
l_shipdate date ,
l_orderkey number ,
l_discount number ,
l_extendedprice number ,

```

Appendix E

```

)
  _suppkey          number,
  _quantity         number,
  _returnflag       char(1),
  _partkey          number,
  _linestatus       char(1),
  _tax              number,
  _commitdate       date,
  _receiptdate      date,
  _shipmode         varchar(10),
  _linenumber       number,
  _shipinstruct     varchar(25),
  _comment          varchar(44)
)
pctfree 1
pctused 99
tablespace ts_s
storage (initial 1m next ${LINEEXT}m pctincrease 0)
parallel (degree 40 instances 1)
nologging;

drop table temp_ord;
create table temp_ord (
  o_orderdate       date,
  o_orderkey        number,
  o_custkey         number,
  o_orderpriority   varchar(15),
  o_shippriority    number,
  o_clerk           varchar(15),
  o_orderstatus     char(1),
  o_totalprice      number,
  o_comment         varchar(79)
)
pctfree 1
pctused 99
tablespace ts_s
storage (initial 1m next ${ORDEXT}m pctincrease 0)
parallel (degree 40 instances 1)
nologging;

exit;
!

while [ $i -le $PAR ]
do
# Kick off sqlldr to load the data into the temporary staging tables.
  FN=`expr $i - 1`
  FN=`expr $FN % 64`
  FN=`expr $FN + 1`
  DRVN=`awk -v a="$i" -v b="STS_S_DRVS" 'BEGIN {printf
"%s",substr(b,a,1)}'`
  $SQLLDR userid=$PASSWD control=${SCRIPT_DIR}/tempitem.ct1 \
  log=${LOG_DIR}/ti${i}_uf1.log \
  data=${UPDATE_DIR}/lineitem.tbl.us${SETNUM}.${i} \
  direct=true parallel=true \
  FILE=${DRVN}:${DBF_DIR}\TS_S_${FN}.DBF &

  $SQLLDR userid=$PASSWD control=${SCRIPT_DIR}/tempord.ct1 \
  log=${LOG_DIR}/to${i}_uf1.log \
  data=${UPDATE_DIR}/order.tbl.us${SETNUM}.${i} \
  direct=true parallel=true \
  FILE=${DRVN}:${DBF_DIR}\TS_S_${FN}.DBF &

  i=`expr $i + 1`
done
# All sqlldr processes have started, now wait for all them to finish.
wait
# now do the insert and then drop the staging tables
$SQLPLUS $PASSWD <<!
spool runuf1.log
set timing on;
alter session force parallel dml;

analyze table temp_ord estimate statistics sample 1000 rows;
analyze table temp_item estimate statistics sample 1000 rows;

insert /*+ PARALLEL(lineitem,20) */
into lineitem (select /*+ PARALLEL(temp_item,2) */ * from temp_item);

insert /*+ PARALLEL(orders,20) */
into orders (select /*+ PARALLEL(temp_ord,2) */ * from temp_ord);

commit;

drop table temp_item;
drop table temp_ord;
spool off
exit;
!
END= ` $GTIME `
# Done
echo ""
echo "Update Function 1 Set $SETNUM done!"
echo "Elapsed Time is `echo $END - $START | bc`"
echo ""

```

runuf2.sh

```

#!/bin/sh
#
#=====  

# Copyright (c) 1995 Oracle Corp. Redwood Shores, CA  

# OPEN SYSTEM PERFORMANCE GROUP  

# All Rights Reserved  

#=====  

# FILENAME  

# runuf2.sh  

# DESCRIPTION  

# runuf2.sh [-u <uid/passwd to login>] [-p <program>] <run_id>  

# <scale factor> <pair number> <parallelism>  

# USAGE  

# To execute UF2.  

#=====  

O=${TPCD_UPDATE_DATA}  

TPCD_DIR=${O}/tpcd  

SCRIPT_DIR=${TPCD_DIR}/update/scripts  

GTIME_DIR=${TPCD_DIR}/source  

GTIME=${GTIME_DIR}/gtime

usage() {
echo "runuf2.sh [-u <user/passwd>] <run_id> <scale factor> <pair number> <parallelism>"
echo ""
echo "run_id : Run_ID of this update run."  

echo "scale factor : Scale Factor of the database."  

echo "set number : The update pair number that this update function  

belongs to."  

echo "parallelism : The parallelism of the updates."  

echo ""
echo "-u : Userid/Password for Oracle. Default is tpcd/tpcd"  

echo "-h : To Display this message."  

echo ""
}

PASSWD="tpcd/tpcd"

set -- `getopt "u:h" "$@"` || usage

while :
do
case "$1" in
-u) shift; PASSWD=$1;;
-h) usage; exit 0;;
-) shift; break;;
esac
shift;
done

if [ $# -lt 4 ]
then
usage
exit 1
fi

RUN_ID=$1
SF=$2
SETNUM=$3
PAR=$4

if [ ${SF} -eq 0 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/tenthgig
LOG_DIR=${TPCD_DIR}/update/log/tenthgig
DBF_DIR=${TPCD}\DBS\TENTHGIG
fi

if [ ${SF} -eq 1 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/onegig
LOG_DIR=${TPCD_DIR}/update/log/onegig
DBF_DIR=${TPCD}\DBS\ONEGIG
fi

if [ ${SF} -eq 3 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/threegig
LOG_DIR=${TPCD_DIR}/update/log/threegig
DBF_DIR=${TPCD}\DBS\THREEGIG
fi

if [ ${SF} -eq 10 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/tengig
LOG_DIR=${TPCD_DIR}/update/log/tengig
DBF_DIR=${TPCD}\DBS\TENGIG
fi

if [ ${SF} -eq 30 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/thirtygig
LOG_DIR=${TPCD_DIR}/update/log/thirtygig
DBF_DIR=${TPCD}\DBS\THIRTYGIG
fi

if [ ${SF} -eq 100 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/hundgig
LOG_DIR=${TPCD_DIR}/update/log/hundgig
DBF_DIR=${TPCD}\DBS\HUNDGIG
fi

if [ ${SF} -eq 300 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/threehundgig
LOG_DIR=${TPCD_DIR}/update/log/threehundgig
DBF_DIR=${TPCD}\DBS\THREEHUNDGIG
fi

if [ ${SF} -eq 1000 ]
then
UPDATE_DIR=${TPCD_DIR}/update/data/onetbyte
LOG_DIR=${TPCD_DIR}/update/log/onetbyte
DBF_DIR=${TPCD}\DBS\ONETBYTE
fi

i=1
PID=""
START= ` $GTIME `
# first create the temp tables
$SQLPLUS $PASSWD <<!

```

Appendix E

```
drop table temp_okey;
create table temp_okey (
  t_orderkey          number
)
pctfree 1
pctused 99
tablespace ts_s
storage (initial 64k next 1m pctincrease 0)
parallel (degree 40 instances 1)
nologging;
exit;
!
# All processes have started, now wait for all the update processes to
finish.

wait

while [ $i -le $PAR ]
do
  FN=`expr $i - 1`
  FN=`expr $FN % 64`
  FN=`expr $FN + 1`
  DRVN=`awk -v a="$i" -v b="$TS_S_DRVN" 'BEGIN {printf
"%s", substr(b,a,1)}'`
  $SQLldr userid=$PASSWD control=${SCRIPT_DIR}/tempokey.ct1 \
  log=${LOG_DIR}/ti${i}.uf2.log \
  data=${UPDATE_DIR}/delete.${SETNUM}.${i} \
  direct=true parallel=true \
  FILE=${DRVN}:${DBF_DIR}\TS_S_${FN}.DBF &
  i=`expr $i + 1`
done
```

```
done
wait
$SQLPLUS $PASSWD <<!
spool runuf2.log;
set timing on;
analyze table temp_okey estimate statistics sample 1000 rows;
alter session force parallel dml;
delete /*+ PARALLEL(orders,20) INDEX(ORDERS, O_OP) */
from orders where o_orderkey in
(select /*+ PARALLEL(temp_okey,20) */ * from temp_okey);
delete /*+ PARALLEL(lineitem,20) INDEX(LINEITEM, L_ORED) */
from lineitem where l_orderkey in
(select /*+ PARALLEL(Temp_okey,20) */ * from temp_okey);
commit;
drop table temp_okey;
spool off;
exit;
!
END=`$GTIME`
# Done
echo ""
echo "Update Function 2 Set $SETNUM done!"
echo "Elapsed Time is `echo $END - $START | bc`"
echo ""
```

runpower1.sh

```
#!/bin/sh
# =====
# Copyright (c) 1996 Oracle Corp, Redwood Shores, CA
# OPEN SYSTEMS PERFORMANCE DIVISION
# All Rights Reserved
# =====
# FILENAME
# runpower1.sh
# DESCRIPTION
# Usage: runpower1.sh [-p <program for query stream>]
# [-u1 <program for UF1>] [-u2 <program for
UF2>]
# [-o] [-s] [-h] [-u <user/password>]
# <scale factor> <update parallelism>
#
# Single stream TPC-D Power Test.
# =====

ORACLE_HOME=d:/orant
TPCD_HOME=/tpcd
SCRIPT_DIR=${TPCD_HOME}/scripts
SQL_DIR=${TPCD_HOME}/sql
UPD_DIR=${TPCD_HOME}/update
SRC_DIR=${TPCD_HOME}/source

RUN_ID_FILE=${TPCD_HOME}/r_id

UPD_SQL=${UPD_DIR}/sql
UPD_SPT=${TPCD_HOME}/update/scripts
UPD_SRC=${TPCD_HOME}/update/source
UPD_DAT=${UPD_DIR}/data

TPCD_BIN=${TPCD_HOME}/bin
TPCD_LOG=${TPCD_HOME}/log
TPCD_RPT=${TPCD_HOME}/rpt

OUT=${TPCD_HOME}/out

GTIME=${SRC_DIR}/gtime

DF=/dev/null
HID=1
INTERVAL=60
COUNT=1200

# The defaults
USER="tpcd/tpcd"
QPROG=${SRC_DIR}/qexec.exe

U1PROG=${UPD_SRC}/uf1.exe
U2PROG=${UPD_SRC}/uf2.exe

usage () {
echo " "
echo "Usage: $0 [-p <program for query stream>] [-u1 <program for
UF1>] [-u2 <program for UF2>] [-o] [-s] [-h] [-u
<user/password>]
<scale factor> <update parallelism>"
echo ""
echo "scale factor : The scale factor of the run."
echo "update ||ism : The parallelism to use for the UFs."
echo ""
echo "-p <program> : Program for Query Stream."
echo " : Default is $QPROG."
echo "-u1 <program> : Program for UF1."
echo " : Default is $U1PROG."
echo "-u2 <program> : Program for UF2."
echo " : Default is $U2PROG."
echo "-o : Collect Oracle statistics."
echo "-s : Collect System statistics."
echo "-u <user/passwd> : User/Password. Default is tpcd/tpcd."
echo "-h : Displays this message."
}
set -- `getopt "p:u1:u2:u:osh" "$@"` || usage

while :
do
```

```
case "$1" in
-u1) shift; U1PROG=$1;;
-u2) shift; U2PROG=$1;;
-p) shift; QPROG=$1;;
-o) shift; OSTAT=1;;
-s) shift; SSTAT=1;;
-h) usage; exit 0;;
-) shift; break;;
esac
#shift;
done

if [ "$#" -ne "2" ]
then
usage
exit 1
fi

SF=$1
PARAM=$2
THROUGHPUT=1

if [ $SF -eq 1 ]
then
QRY_DIR=${TPCD_HOME}/queries/q_one_0.sql
INIT_DIR=${TPCD_HOME}/dbs/onegig
elif [ $SF -eq 3 ]
then
QRY_DIR=${TPCD_HOME}/queries/q_three_0.sql
INIT_DIR=${TPCD_HOME}/dbs/three
elif [ $SF -eq 10 ]
then
QRY_DIR=${TPCD_HOME}/queries/q_ten_0.sql
INIT_DIR=${TPCD_HOME}/dbs/tengig
elif [ $SF -eq 100 ]
then
QRY_DIR=${TPCD_HOME}/queries/q_hund_0.sql
INIT_DIR=${TPCD_HOME}/dbs/hundgig
elif [ $SF -eq 300 ]
then
QRY_DIR=${TPCD_HOME}/queries/q_threehund_0.sql
INIT_DIR=${TPCD_HOME}/dbs/threehund
elif [ $SF -eq 1000 ]
then
QRY_DIR=${TPCD_HOME}/queries/q_tbyte_0.sql
INIT_DIR=${TPCD_HOME}/dbs/tbyte
fi

export SF PARA THROUGHPUT

# ${UPD_SPT}/genuf1.sh 1 ${PARAM} ${SF}
# ${UPD_SPT}/genuf1.sh -r 1 ${PARAM} ${SF}
# ${UPD_SPT}/genuf2.sh 1 ${PARAM} ${SF}

if [ ! -f $RUN_ID_FILE ]
then
echo "0" > $RUN_ID_FILE
fi

RUN_ID=`cat $RUN_ID_FILE`
RUN_ID=`expr $RUN_ID + 1`
echo $RUN_ID > $RUN_ID_FILE

echo "TPC-D Power Test.Run `date`"
echo "RUNID is $RUN_ID"
echo ""

if [ $OSTAT ]
then
svrmgr30 @$ORACLE_HOME/rdbms0/admin/utl1bstat
fi

if [ $SSTAT ]
then
${TPCD_BIN}/start_ntstat.sh $RUN_ID
fi

START=`$GTIME`
echo "TPC-D Power Test Execution starts at $START"
echo ""
```

Appendix E

```
# Execute UF1
echo "Start UF1 at `date`"
sh ${UPD_SPT}/runuf1.sh -u ${USER} ${RUN_ID} ${SF} 1 ${PARAM} \
  > ${OUT}/uf1.${RUN_ID}.${HID} 2>&1
# Execute Query Stream
echo "End UF1. Start Query Stream at `date`"
${QPROG} ${USER} q${QRY_DIR} r${TPCD_RPT}/rpt.${RUN_ID}.${HID} \
  | ${OUT}/qs.${RUN_ID}.${HID} 2>&1
# Execute UF2
echo "End Query Stream. Start UF2 at `date`"
sh ${UPD_SPT}/runuf2.sh -u ${USER} ${RUN_ID} ${SF} 1 ${PARAM} \
  > ${OUT}/uf2.${RUN_ID}.${HID} 2>&1
echo "END UF2 `date`"
END=`$GTIME`
echo "TPC-D Power Test Execution ends at $END"
echo "Measurement Interval is `echo $END - $START | bc`"
echo ""
echo "-----"
echo ""
```

```
if [ $OSTAT ]
then
  svrmgr30 @$ORACLE_HOME/rdbms80/admin/utlstat
  mv report.txt ora_stat_${RUN_ID}.txt
fi
if [ $SSTAT ]
then
  ${TPCD_BIN}/kill_ntstat.sh ${RUN_ID}
  mv report_os.txt sys_stat_${RUN_ID}.txt
fi
echo "-- update function 1" >> ${TPCD_RPT}/rpt.${RUN_ID}.${HID}
awk -f utime.awk ${OUT}/uf1.${RUN_ID}.${HID} >>
  ${TPCD_RPT}/rpt.${RUN_ID}.${HID}
echo "-- update function 2" >> ${TPCD_RPT}/rpt.${RUN_ID}.${HID}
awk -f utime.awk ${OUT}/uf2.${RUN_ID}.${HID} >>
  ${TPCD_RPT}/rpt.${RUN_ID}.${HID}
cat ${TPCD_RPT}/rpt.${RUN_ID}.${HID} | awk
  'BEGIN{getline;print}{printf"%s\t", $0;getline;print $0;}' >
  ${TPCD_RPT}/mpinter.${RUN_ID}
cat ${TPCD_RPT}/rpt.${RUN_ID}.${HID} >
  ${TPCD_RPT}/pt_rpt.${RUN_ID}.${HID}
${TPCD_BIN}/metric ${SF} < ${TPCD_RPT}/rpt.${RUN_ID}.${HID} >>
  ${TPCD_RPT}/pt_rpt.${RUN_ID}.${HID}
${TPCD_BIN}/metric ${SF} < ${TPCD_RPT}/rpt.${RUN_ID}.${HID} >>
  ${TPCD_RPT}/mpinter.${RUN_ID}
#rm -f ${TPCD_RPT}/rpt.${RUN_ID}.${HID}
cat ${INIT_DIR}/tkinit.ora >> ${TPCD_RPT}/pt_rpt.${RUN_ID}.${HID}
```

Appendix F

Appendix F: Initial Ten Rows

O_ORDERDA	O_ORDERKEY	O_CUSTKEY	O_ORDERPRIORITY	O_SHIPPRIO	O_CLERK	O	O_TOTALPRI	O_COMMENT	
16-MAR-94	600070	3065861	4-NOT SPECIFIED		0 Clerk#000054498	F	128827.92	B1PwLhCjki44g0 025gnRzAhs AQ MBO	
04-MAR-94	600198	2241085	2-HIGH		0 Clerk#000072411	F	149792.12	13z0703m1N5N0w MjXRjQm41S4	
km3zz07NQ	xLChi M 4P5ryA7wmn231SAkr	ANYk6hyBk6PLjLR							
03-MAR-94	600386	2214397	3-MEDIUM		0 Clerk#000084084	F	143577.43	R7AMNMC 3N621o15 2BoJQBzS	
47w0g4o									
25-MAR-94	600608	2075792	1-URGENT		0 Clerk#000084636	F	299721.73	lin 71issm4ib	
Rx56yQL4555wglgBS3qn3R5i0km771	7hwmGwB2R51j								
14-MAR-94	600679	3473629	2-HIGH		0 Clerk#000054337	F	56068.1	4mnQk2kr 6AwikCnzSqh13zRbi0 ik	
1M54101kwnSjh42P2h3C1n									
31-MAR-94	600838	5953672	5-LOW		0 Clerk#000024247	F	72338.76	hhmwg6zCAS7146Ak	
xj1S7NPB1N34giLcn75zR047LkALNmA	6xwjCx1Ri5632A75AMCRR1z4								
22-MAR-94	601190	10142651	2-HIGH		0 Clerk#000075145	F	193829.68	0hzyLAPiMi3k 57QnzLCh	
11-MAR-94	601286	11880113	1-URGENT		0 Clerk#000082486	F	153549.17	hzOB3Q6jMhBxhmSyNz4CjQ00P7	
06-MAR-94	601633	13083533	1-URGENT		0 Clerk#000015877	F	57313.24	16mk N C3Mxkzmz6ZLgJAhY7	
ORR2g2BAMzh1S									
02-MAR-94	602017	1312997	2-HIGH		0 Clerk#000007383	F	334707.8	COq15nSg67zP 4kNAk2R1jx6k	
10 rows selected.									
L_SHIPDAT	L_ORDERKEY	L_DISCOUNT	L_EXTENDED	L_SUPPKEY	L_QUANTITY	L_L_PARTKEY	L_L_TAX	L_COMMITD	L_RECEIPT
L_SHIPMODE	L_LINENUMB	L_SHIPINSTRUCT		L_COMMENT					
06-JUL-98	86657188	.07	29544.48	129932	16 N	7629917	O	.04 20-MAY-98 17-JUL-98 AIR	
4 TAKE BACK RETURN		nA3P3z jNgmCB7 whi3yPxh							
27-JUL-98	86657318	.02	54610.71	56628	33 N	14556599	O	.05 17-JUN-98 30-JUL-98 REG AIR	
5 TAKE BACK RETURN		A430Nw4Cmngk5BChAjzkz65h3y5Czj0x42A4NLZSMmSN							
09-JUL-98	86657318	.08	52731.36	495590	33 N	12745553	O	.07 18-JUN-98 16-JUL-98 TRUCK	
6 NONE		m3kyAgQ66 34R4							
12-JUL-98	86657408	.02	3344.49	240029	3 N	3990025	O	.08 22-AUG-98 06-AUG-98 FOB	
2 NONE		y2RzPQ0BAMkQwx7M4ymNCLBCOOQ7 ialiw1NQc							
25-JUL-98	86657408	.03	65963	162308	50 N	912307	O	.07 09-AUG-98 21-AUG-98 SHIP	
4 TAKE BACK RETURN		5y561nz4AijSCON04i17P							
30-JUL-98	86657408	0	72444.8	887814	40 N	13137774	O	.08 25-JUN-98 23-AUG-98 MAIL	
5 COLLECT COD		57sJc7xB0Bjz3m7mgPk257x7x2NR3Q1A							
18-JUL-98	86657602	.05	2977.54	358432	2 N	13358431	O	.01 03-JUL-98 28-JUL-98 TRUCK	
2 TAKE BACK RETURN		573RB1w24PwjgNmKMA7kjkxCPVh6mZQQ							
09-JUL-98	86657603	.06	8077.8	794409	6 N	2044402	O	.02 21-JUN-98 17-JUL-98 SHIP	
4 COLLECT COD		1wL7AmCjCM4yxg1BP3SiRin1i							
22-JUL-98	86657638	.07	86421.15	36962	45 N	9286934	O	.01 04-AUG-98 02-AUG-98 RAIL	
2 DELIVER IN PERSON		gQw3h7AZLSN							
12-JUL-98	86657767	.07	17661.44	656057	16 N	4156048	O	.07 29-AUG-98 03-AUG-98 RAIL	
2 DELIVER IN PERSON		RkgQmlhx7k4g4Cjy yBLhyP6CRO6							
10 rows selected.									
C_CUSTKEY	C_MKTSEGME	C_NATIONKE	C_NAME		C_ADDRESS			C_PHONE	
C_ACCTBAL	C_COMMENT								
2343757	BUILDING	0	Customer#002343757		Omi0hz2n13nAM			10-288-494-8922	
8061.92	60 hx5AQPCKkNnQzSknMknk2i4 5C								
2343773	BUILDING	13	Customer#002343773		N0in0N1lhomhax			23-783-956-2263	
6442.11	iAlMhhwQw1Rgng gNA5GjNmyjL64N61lgy0nP4P6MORy7l y								
2343774	BUILDING	22	Customer#002343774		7Lw6zSsshLzqizlgnnz			32-644-259-4728	
8504.63	QngMkB33g NPMw06PhPnM6nnC04xxQ6037mj0x4M1 2								
2343779	BUILDING	10	Customer#002343779		i7hzg4 gL3z40Cz53047nBhR4CA			20-352-990-5668	
8656.41	7100LImwP4hB22C01AaxL1RP0MOQm72imi51q2N2By5337 QznBn11l								
2343792	BUILDING	2	Customer#002343792		xS3x1qjkkrc hp57LR3g2mMnhkP Onm			12-670-459-1107	
4884.19	1N7ALZmhs ojsQA0Q 45z3 4QONS5yzBBnxNS 7mPlzgQnmyAB0w53BlN0NPOx036MB7Rm06iS216QzL								
2343794	BUILDING	10	Customer#002343794		y1BN4Qmxi iPBAL4AR2N			20-548-928-6601	
4734.53	31j2 13h053Mni7Bgjas273in47m7isiL kg5ARP								
2343803	BUILDING	16	Customer#002343803		7N51zo2Qz4PA n2n1SyPm 3nzy			26-735-958-3573	
7233.44	45m46zjP7QRyzh5Bnnzmj 43NQmzN1Pz1jo4g0LwW77wSN4xA6gy3hBRy6 xBmRn1QSPCS								
2343806	BUILDING	23	Customer#002343806		2w1Rj0wsOP jmywBi			33-550-986-7094	
5073.36	Sh1Ax61zjMmysB1QLGZA3iROWgxS01k2 65jNycjh5w5Sxw0 3RgmRCRmmM1 7Q3B13SR								
2343807	BUILDING	14	Customer#002343807		1517LLQ1igwj			24-984-652-9565	
4977.94	AA7mg16ki 5M1x1i 5iR0Mbg5C5OLC674mh5h16m6PMk3kBlwyQj55SLg60j04Q OQmi 6LgAPmzkOP7wm								
2343813	BUILDING	11	Customer#002343813		OR Qn L1SnnjwBzQ 73 Qn7ga3			21-172-489-9055	
9358.31	m37jOP11m0ZM nx2ACnyjP OLwXLRz2nwQ4g2gj35OkMcmPR7Qynji7S3nL3								
10 rows selected.									
P_PARTKEY	P_TYPE	P_SIZE	P_BRAND	P_NAME					
P_CONTAINE	P_MFGR	P_RETAILPR	P_COMMENT						
13790289	STANDARD BRUSHED STEEL	11	Brand#55	lavender peru midnight salmon firebrick				WRAP	
PACK	Manufacturer#5	1378.6	Q750gw11B1CwX1PN						
13790313	PROMO BURNISHED NICKEL	11	Brand#53	orchid brown firebrick burlywood drab				JUMBO	
JAR	Manufacturer#5	1402.63	6hj7z60z10j5A6						
13790399	LARGE POLISHED COPPER	11	Brand#14	rose olive sienna tan				WRAP	
JAR	Manufacturer#1	1488.71	kB56gLMQ411 iB						

Appendix F

13790543	STANDARD POLISHED TIN	11	Brand#43	drab sky floral frosted chocolate	LG
BAG	Manufacturer#4	1632.86	JMPjz12yk4k7m		
13790706	PROMO BURNISHED STEEL	11	Brand#51	midnight brown magenta blush orchid	JUMBO
BOX	Manufacturer#5	1796.02	P71mBMLnQOR		
13790708	LARGE POLISHED BRASS	11	Brand#35	sky chartreuse saddle dim tomato	SM
JAR	Manufacturer#3	1798.02	SCNzh5		
13790780	LARGE BURNISHED TIN	11	Brand#53	cream lawn rose firebrick burnished	SM
PACK	Manufacturer#5	1870.1	6g5OySw2w40iR4CAP3 n		
13790798	STANDARD POLISHED STEEL	11	Brand#33	sandy salmon khaki olive honeydew	MED
PACK	Manufacturer#3	1888.11	A6RBxLah25		
13790916	PROMO PLATED BRASS	11	Brand#51	burnished turquoise deep purple sky	LG
DRUM	Manufacturer#5	2006.23	h50gokmhBC1 Bz		
13790990	SMALL PLATED TIN	11	Brand#33	floral medium burnished tan	WRAP
BAG	Manufacturer#3	2080.31	jw6yiBg2iMORXP4m		

S_SUPPKEY	S_NATIONKE	S_COMMENT	S_PHONE	S_ACCTBAL	S_NAME
256033	9	QRnOM0AlzSmikg6nx010hkBLAQ3xPmxOL1M26j	2Myx7zB0i7B4LSxjL43hhgM7nknPQMLNxo4xi4S53		
Supplier#000256033		CwgOx4BMxx3	19-693-495-3538	3964.65	
256034	21	L34xmzigMj0 53wQwC2 Qqn6Q4iw	KSRLLCCw4BAjw4yLPNCihz		
Supplier#000256034		wPR635207R7iPN14QQP	31-956-256-8259	4900.02	
256035	21	isk5Bm00yAha31S76017CAhAA			
Supplier#000256035		Qn5kxML06h1100S1A2RS	31-895-218-1031	6635.15	
256036	7	Ak5P4gPqhrn2hhln06nSR0C1zbx25mjm	iyM100CQ0RB0P510Ny6x72ML5nMgOmLPP45		
Supplier#000256036		7123 ok2CNiiz3m3nko03w5	17-695-362-9989	8034.77	
256037	23	00L56Lh6A1mSM Bn0jM0Cjhgwni223Bn1j127n5w1jm	Pj57LMk2RMk7Sh5Mk w14PR3glBM50yi1		
Supplier#000256037		z4P5N6S1A40hk0n230116	33-392-388-4607	2515.64	
256038	7	pZgn 54ygn1Bqh0g0BSx5BPjcmQLg0Ai0q7zMWLlmgc7	N4N13kc h5BP		
Supplier#000256038		min2Icj321yy6yr o	17-440-792-6246	5295.65	
256039	11	GziAOS2N20xih4yy4wBA 65LxMhQ3HPmPyjja3ggN3n1wx01jr45mSRRsg1m0AA			
Supplier#000256039		216mPnMS4w w62w12PO5nNhsXp1ly55m0j4	21-668-360-2610	1064.35	
256040	11	nmiNkw622niBQA3A2Smn mLNLh6xB44zLL3y1PNMR			
Supplier#000256040		Loj6y12jPj2zk1 wMCj4C B15	21-911-163-1284	442.51	
256041	17	PzQPCM371h0hh6w0711C zM7z52kk7gi4N Bhz MShPwBmxBmnQjmin6w70gPMnymmNwwMn			
Supplier#000256041		3Lz0n06zSNR1hNCyn0 iLhg1	27-910-931-2221	-775.46	
256042	6	PjnzCok1kz 0007yQ41A12nm501AC1M255Bhjoky6BS473iCgC76y1BnPPgzgP5nR			
Supplier#000256042		4yAAXORoz1mw7wCRO32M53zRjnhw7ignNm	16-503-876-4075	-301.02	

PS_PARTKEY	PS_SUPPKEY	PS_SUPPLYC	PS_AVAILQT	PS_COMMENT
4Q7s1Qmzx2jos3	1	2	771.64	3325 00PL56QKQRskg2z7MANNj4i1h2zLQQLiQnAlML1S6 k4hg3P5hk3ywMLwy 7gjR3
LxMLMCOB	1	250002	993.49	8076 ns07M1n4N7L1xgAyMzMNn07k0N1hjyShgcy30A 27QML0SQ77CPpgkCQAQCwz5M3MmSSAQ
gAN3	1	500002	337.09	3956 6215k jLCizn10b162nP41LQy431k0zyzn2M6L3h731C1jh1x3x5ghj1oyL76A0hzPk2Cs2jKxn
2MP1x0wC23ik2OmK4Nnxzm	1	750002	357.84	4069 2340CA5ghw0P0gs3n2jcs35yam 3L5C7iB k7 wL R52LLOACQ6i6060B
nQhBM5R0i5N7A5w4B452k15060zGhg6SRB7nPlhQCjgjr17SBA77	2	3	378.49	8895 MMNOM3BnMM6NBzjB 2mg i jALB
00NSjwPw	2	250003	915.27	4969 6S66zNlykhii26wwaxz1PRMxggAy446yylPBSwP R16ggNkyikkyhymMShNgQmBim1N60
x6kn3BR61kmhMzjQk6S343LM61z2Qh512MSi4nQ5Bghx1h401yyn463mx0h7Bxgx1	2	500003	438.37	8539 BPOgj3k MgQR2
L1gSxyp1	2	750003	306.39	3025 y5Bny3Aw02nyMxgzP5BS14gg7McN1LSkih56gOMoy4Qynj5P3im jomkrQhRR3hy1Cn4jn
1AwBj0gSnm1z1zAMA6417zgs154nLCC0Q6BC1lgxyB6Bkoj6QC6n4mw2w7jgCNP	3	4	920.92	4651 p7 437MnnM0Pik
PyRmlw076ko3igxhS64h5x6PBLM2Pxko0j3NMRgzP6S1ghhw4Nnn04my031zcyQQ4M3gs2ko2iOmxPz1	3	250004	498.13	4093

N_NATIONKE	N_NAME	N_REGIONKE	N_COMMENT
0	ALGERIA	1	2Cxh17 L1iwk6hmh300izngn32CPwCikyLk6khmzSRA
1	ARGENTINA	1	zQn3Okwz1wl.n7PLS30hCgn56kP5PyRikgi1B71L
2	BRAZIL	1	gLm50nACAmnBCj2k1ki7RCPngPxnCojNg4k OiAg57COSOm1NwCnoyLx40R SC
3	CANADA	1	4yMO AhnQ5Lh wzQAM662Aw1Byc17CxmzRwNR5na104 x
4	EGYPT	4	11im5126 cxj NMQmLxoiKn102j2m3Ah4yNR1QQiL507j2Qs1YN
5	ETHIOPIA	0	NS7n LSOP Oz5n1A1B2S02nN01Mh4SBXP iRhBO 047R26 2B1M
6	FRANCE	3	3mjmiZ1 s 3L3k2hNNhN1P4w370xRxyN15wn
7	GERMANY	3	z nOP4Rkwo CmzBB 516mAg 1Byw40M3QYNPA
8	INDIA	2	MN1R5RC1RMj1111w7Myn M11y1N1MmBQ17PL4C
9	INDONESIA	2	SjPmQO71Lj 7Abj6Mx1Aqk3nLwi73BPxzCwjzmn4z1Lzgg6nNZ0j0w zxC66gP6ykrPmG

R_REGIONKE	R_NAME	R_COMMENT
0	AFRICA	xSx31zz31C11z40Anmm05AjiOxC3AMMNOC0kACgwnng3g1P7LLLwy1Qy7R
1	AMERICA	kgyh3LsnC72k6z1Az0LP3k2L4QB1QL106730j01SPj0ng07C0100SbgmgRQ41gPCmk21A425ik1yAR4y
2	ASIA	NSg6x1M1A11zm6m0ROA0jx nhrA77NgRxBwL1M6Py RjySB3RLwkyPkwMM2R1BQ xazkOgkjm110gAgh1
3	EUROPE	z1SL7Qwg12hMBL51h1z0M45QkjsHwSy104MLoh7wn1ARLQPYAyAi15761Li7A1nR1S RQ4SLny7B
4	MIDDLE EAST	R11xmPLZ3cy2mN1g4QmBNNASM Acki MPki70i

Appendix G: Auditor's Attestation Letter

Test Sponsor: Richard Clifton
 VP, Database Engineering
 Data General Corporation
 4400 Computer Drive
 Westboro, MA 01886

5 August 1997

I verified the TPC Benchmark™ D performance of the following configuration:

Platform: AViiON Model AV6600
 DataBase Manager: Oracle Version 8.03
 Operating System: Microsoft Windows NT 4.0 SP3
 The results were:

CPU's Speed	Memory	Disks	QppD@100GB	QthD@100GB
AViiON Model AV6600				
6 x Pentium Pro (200 MHz)	512KB L2 2GB Main	1 x 2.1GB 83 x 4.3 GB	528.6	213.9

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 TIME table was not used
- The input variables were generated by QGEN
- The database was populated using DBGEN
- The database was maintained by the "Reset" method
- The throughput metric was computed using the results from the power test
- The ratio between the longest and the shortest query was such that no adjustment was necessary
- A compliant implementation specific layer was used
- The query text was produced using compliant variants and minor modifications
- The database records were defined with the proper layout and size

Appendix G

- The database was properly scaled to 100GB and populated accordingly
- The database load time was correctly measured and reported
- The ACID Properties were verified and met
- The reported execution times were correctly measured and reported
- Measurement repeatability was verified
- At least 8 hours of database log was configured
- The system pricing was verified for major components and maintenance
- The major pages from the FDR were verified for accuracy

Additional Audit Notes:

None.

Respectfully Yours,

A handwritten signature in black ink, appearing to read "François Raab", with a long horizontal flourish extending to the right.

François Raab
President

AViiON Model AV6600