5 Application of DBT-1 to a Commercial RDBMS

5.1 Overview

It is easily anticipated that users who are intended to introduce an opensource RDBMS will raise this question: what differences are there in functionality, performance, extensibility, and reliability between the opensource RDBMS and the commercial RDBMS that has been familiar to them. In this project, we ported the OSDL-developed DBT-1 benchmark to a commercial RDBMS so that comparisons can be made between opensource RDBMSs and the commercial RDBMS in terms of their performance of web transaction applications. We set Oracle 9i, which is representative in terms of share, as the target of porting though more than one commercial RDBMS is available as of December 2004. The source we selected is the DBT-1 for PostgreSQL 7.4.x which is the sole RDBMS that is presently supported by DBT-1.

This chapter describes the procedures for using DBT-1 in practical situations and other practical hints and instructions for making measurements.

5.2 Environment Definition

5.2.1 Installation

5.2.1.1 Installing Linux・Oracle

Oracle9i which is discussed in this chapter has several subversions for different OS distributions and OS versions. The Oracle installation procedure for the OS may differ greatly depending on the OS distribution and the Oracle subversion you selected. For this reason, we will not provide a detailed description of the procedure for installing Oracle9i. Install it while referring to the manual associated with the target version of Oracle and Oracle's support information such as Metalink.

The procedure explained below assumes Oracle 9i Database Release2 (9.2.0.4.0). The OSs that are selected as the measurement environments are RedHat EL3, Miracle Asian Linux v3.0, and SuSE ES9.

5.2.1.2 Installing DBT-1

In this section, we install dbt1-v2.1-ora-1.0 which we developed for this time's project.

5.2.1.2.1 Prerequisite conditions

dbt1-v2.1-ora-1.0 presumes the following conditions:

- The Linux account dbt1si is created as the benchmark user and its home directory is set to /home/dbt1si/.
- dbt1-v2.1-ora-1.0.tar.gz is extracted in the directory immediately below

/home/dbt1si/.

- An Oracle database with SID=dbt1 is created using Oracle Database Configuration Assistant or Oracle Net8 Configuration Assistant and TNS Listener is configured for the database. The database for this time's project was created in the dedicated database mode.

5.2.1.2.2 Creating a Linux user as the benchmark user

Create a Linux user following the procedure shown below. Run the following command as the root:

```
# useradd dbt1si -d /home/dbt1si –g [Oracle's administrator group]
```

Specify the group that is established when installing Oracle as [Oracle's administrator group]. Edit /home/dbt1si/.bash_profile as follows:

```
<Sample settings>

ORACLE_BASE=/home/oracle
ORACLE_HOME=$ORACLE_BASE/product/9.2.0
ORACLE_SID=dbt1
NLS_LANG=Japanese_Japan.JA16EUCTILDE
ORA_NLS33=$ORACLE_HOME/ocommon/nls/admin/data
PATH=$PATH:$ORACLE_HOME/bin
ORACLE_DOC=$ORACLE_HOME/doc
CLASSPATH=$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/jlib
CLASSPATH=$CLASSPATH:$ORACLE_HOME/network/jlib
CLASSPATH=$CLASSPATH:$ORACLE_HOME/jdbc/lib/classes12.zip
CLASSPATH=$CLASSPATH:$ORACLE_HOME/jdbc/lib/nls_charaset12.zip
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$ORACLE_HOME/lib
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$ORACLE_HOME/ctx/lib:$ORACLE_HOME/jdbc/lib

export PATH ORACLE_BASE ORACLE_HOME ORACLE_SID NLS_LANG ORA_NLS33
export ORACLE_DOC CLASSPATH LD_LIBRARY_PATH
```

5.2.1.2.3 Creating an Oracle user for running the benchmark

Before installing DBT-1, create an Oracle user for running the benchmark and give it the dba privileges. Start sqlplus and type:

```
sqlplus> connect sys/[sys password] as sysdba
sqlplus> startup;
sqlplus> create user dbt1 identified by dbt1 default tablespace [default data table area name];
sqlplus> grant dba to dbt1;
```

5.2.1.2.4 Installing the DBT-1 package for Oracle

Firstly, expand the Oracle version of DBT-1 tarball dbt1-v2.1-ora-1.0.tar.gz into /home/dbt1si/ as the dbt1si user. The tarball is assumed to be located in the /tmp directory.

```
$ cd ~
$ tar xfz /tmp/dbt1-v2.1-ora-1.0.tar.gz
```

Then, edit the make environment configuration file (/home/dbt1si/dbt1-v2.1/make.common). By default, $ORACLE_HOME= /home/oracle/product/9.2.0 is assumed. If $ORACLE_HOME disagrees with the above settings, edit the part "/home/oracle/product/9.2.0" according to the actual $ORACLE_HOME value.

Finally, run make and make install. Do not use configure as described in the PostgreSQL version of the DBT-1 user manual. Run the make command in the directory /home/dbt1si/dbt1-v2.1 (not using configure).

```
$ cd ~/dbt1si
$ make
$ make install
```

This completes the installation of DBT-1.

5.2.2 Parameter Configuration

5.2.2.1 Oracle

Configure the DBT1 database that has been generated under Oracle as shown below. Leave the other parameters unchanged. In this chapter, the settings shown below were handled as default parameters that are used before tuning was carried out.

Although these values are allowed to be altered through sqlplus since Oracle 9i, we follow the same old procedure as that for Oracle8i and earlier; that is, editing $ORACLE_HOME/dbs/spfiledbt1.ora. The Oracle parameters that have been changed

are listed in Table 5.2-1.

Table 5.2-1 Oracle Parameters

Parameter name	Value
max session	335
db_cache_size	25165824
sort_area_size	524288
db_cache_advice	ON
db_writer_processes	1
log_buffer	524288
log_parallelism	1
open_cursors	300
processes	150
sga_max_size	135337420
shared_pool_size	83886080
pga_aggregate_target	25165824
object_cache_optimal_size	102400

Restart the DBMS to have the changed parameters be reflected. Type from within sqlplus:

```
sqlplus> connect sys/[sys password] as sysdba
sqlplus> shutdown
sqlplus> startup
```

5.2.2.2 OS

We set the OS parameters as shown below for the measurement projected described in this chapter. Note that the maximum number of file descriptors (-n) must not be set to 4096 and that measurements may not be possible due to errors if the number of simultaneously connected users, which is explained later, is set to a value as great as several thousands. Table 5.2-2 shows a sample output of "ulimit –a."

Table 5.2-2 ulimit –a Result Output

RedHat EL AS3	SuSE ES9	Miracle Asian Linux
core file size (blocks, -c) 0	core file size (blocks, -c) 0	core file size (blocks, -c) 0
data seg size (kbytes, -d) unlimited	data seg size (kbytes, -d) unlimited	data seg size (kbytes, -d) unlimited
file size (blocks, -f) unlimited	file size (blocks, -f) unlimited	file size (blocks, -f) unlimited
max locked memory (kbytes, -l) 4	max locked memory (kbytes, -l) unlimited	max locked memory (kbytes, -l) 4
max memory size (kbytes, -m) unlimited	max memory size (kbytes, -m) unlimited	max memory size (kbytes, -m) unlimited
open files (-n) 4096	open files (-n) 4096	open files (-n) 4096
pipe size (512 bytes, -p) 8	pipe size (512 bytes, -p) 8	pipe size (512 bytes, -p) 8
stack size (kbytes, -s) 10240	stack size (kbytes, -s) unlimited	stack size (kbytes, -s) 10240
cpu time (seconds, -t) unlimited	cpu time (seconds, -t) unlimited	cpu time (seconds, -t) unlimited
max user processes (-u) 7168	max user processes (-u) 30207	max user processes (-u) 7168
virtual memory (kbytes, -v) unlimited	virtual memory (kbytes, -v) unlimited	virtual memory (kbytes, -v) unlimited

5.3 Evaluation Procedure

5.3.1 Creating the Database

5.3.1.1 Determining the size

DBT-1 uses the ITEM (product) count and CUSTOMER count as the factor to determine the size of data. Here, it is assumed to be 10,000 and 1,000, respectively, for illustrative purposes (reference: these settings yield a total text-equivalent data size of approximately 3 GB). The test can determine these values according to his or her interest (e.g., want to simulate an application that has many users but few items).

5.3.1.2 Generating load data

Generate the load data according to the procedure described in Section 3.1. If the environment variable $DBT1_RAWDATA is set (for example to the directory /work/medium that gives the dbt1si user the read and write permissions) and designated as the location where the entity of the loading text data is to exist, type as follows as the dbt1si user:

```
$ cd ~/dbt1-v2.1/datagen
$ ./datagen –d ORACLE –i 10000 –u 1000 –p $DBT1_RAWDATA
```

It may take several hours to generate 3GB of data in some environments. A symbolic link is made from /tmp to the actually generated data file. It must be noted that the path /tmp is used unconditionally when loading data.

Make sure that the following sql file has been generated by datagen by checking the files' time stamp:

```
$ ls –l /home/dbt1si/dbt1-v2.1/scripts/oracle/create_sequence.sql
```

5.3.1.3 Generating the database

Generate the contents of the DBT-1 database, i.e., schema. The data generation script internally makes use of "/home/dbt1si/dbt1-v2.1/scripts/oracle/set_run_env.sh." This file need not be edited if the above-mentioned steps are observed. Let us review the contents of this file for confirmation. Given below is an excerpt of the file.

```
export ORA_USER=dbt1
export ORA_PASS=dbt1
export ORA_SID=dbt1
export ORA_LOAD_LOGPATH=/tmp
export ORA_LOAD_ERRPATH=/tmp
export ORA_LOAD_DATAPATH=/tmp
export DBT1_PERL_MODULE=/home/dbt1si/dbt1-v2.1/perlmodules
```

To customize the location of the table areas for storing data and indexes, in particular, customize (change the Storage clause) the SQL statements

- create_tables.sql
- create_indexes.sql

in ~/dbt1-v2.1/scripts/oracle/. No changes are required if the default table areas are all right.

Finally, create a schema. Run the following command as the dbt1si user. Things are OK if the command terminates with no error.

```
$ cd /home/dbt1si/dbt1-v2.1/scripts/oracle
$ ./create_dbt1_schema.sh
```

An SQL（delete_dbt1_schema.sql) for deleting a schema is made available in case schema generation fails. Start sqlplus from the same directory and run @delete_dbt1_schema.

```
$ sqlplus dbt1/dbt1@dbt1

SQL*Plus: Release 9.2.0.4.0 - Production on month Dec 27 19:57:10 2004

Copyright (c) 1982, 2002, Oracle Corporation.  All rights reserved.

Oracle9i Enterprise Edition Release 9.2.0.4.0 - Production
With the Partitioning, OLAP and Oracle Data Mining options
Connected to
JServer Release 9.2.0.4.0 - Production
SQL> @delete_dbt1_schema.sql
```

5.3.2 Configuring DBT-1 Parameters

It is assumed here that DBT-1 is run using a startup script which is supplied in advance in the DBT-1 package (the script will then collect statistics about system operations in the background). In this case, edit $DBT1_HOME/data_collect/oracle/run.config. The meanings of the parameters in the file are identical to those for the PostgreSQL version of DBT-1. See Chapter 4. The parameters "bdbname," "username," and "password" must be set to the Oracle's SID, user name, and password, respectively (no changes are required if the above-mentioned configuration steps are followed).

Table 5.3-1 DBT-1 Measurement Parameters

#	Item name	Value
1	# database config	
2	items	10000
3	gcustomers	2,880,000
4	dbhost	localhost
5	bdbname	DBT1
6	username	pgsql
7	password	pgsql
8	# cache config	
9	cache	0
10	xcache_host	localhost
11	lcache_port	9999
12	mconnection	10
13	# appServer config	
14	appserver	1
15	server_host0	localhost
16	nserver_port0	9992
17	q_size0	1000
18	a_size0	1000
19	rconnection0	80
20	# dbdriver config	
21	ydriver_host0	localhost
22	vrate0	100
23	eus0	100
24	zduration0	1200
25	think_time0	7.2
26	jdriver_server_host0	localhost
27	kdriver_server_port0	9992
28	out_dir	/tmp
29	db_param	-c tcpip_socket=on
30	redirect_tmp	1
31	redirect_xlog	1

Although DBT-1 can be run as either web 2 or 3 layer model as explained in Chapter 8, it assumes the web 3 layer model which is commonly adopted in the present world. In other words, DBT-1 is configured as dbdriver (Web client emulator)<->AppServer (with

connection pool)<->DBMS. The standard settings for this time's measurement project in this configuration (some parameters have been changed appropriately through tuning, though) are listed in Table 5.3-1.

Since any shortage of free transaction work area will cause wait conditions in locations that are irrelevant to the DB, it is desirable, if possible, that q_size0 and z_size0 be set to values greater than the eu settings, though these also depend on the value of Think_time.

The DBT-1's run-time settings such as the number of simultaneously connected users and the number of appServer's DB connections can be given as start-time options to dbdriver and appServer if they are to be started directly. DBT-1 runs with much the same options as those for the PostgreSQL version of DBT-1 except that the "dbname" is set to the Oracle's SID. For details, therefore, refer to the PostgreSQL edition of DBT-1 manuals. If the parameters "dbname, "username," and "password" are omitted, they are all assumed to be "dbt1."

<Example>

$ **./appCache --host localhost --dbname dbt1 --username dbt1 --password dbt1 --port 9999 --db_connection 10 --item_count 1000**

$ **./appServer --host localhost --dbname dbt1 --user dbt1 -password dbt1 --server_port 9992 --db_connection 20 --txn_q_size 100 --txn_array_size 100 --item_count 1000 --access_cache --cache_host localhost --cache_port 9999**

$ **./dbdriver --server_name localhost --port 9992 --item_count 1000 --customer_count 8640 --emulated_users 100 --rampup_rate 60 --think_time 1.6 --duration 900**

5.3.3 Running DBT-1

5.3.3.1 Starting the database server

It is necessary to run the DBT-1 database for Oracle first. To clear Oracle's internal cache, which is not required if Oracle is already active, it is necessary to restart Oracle. Start sqlplus as the dbt1si user and type as follows:

sqlplus> **connect sys/**[sys password] **as sysdba**
sqlplus> **startup;**

If no TNS listener is up, start one following the step shown below before starting the database.

$ **lsnrctl start**

5.3.3.2 Starting the DBT-1 application servers

Run the following as the dbt1si user:

```
$ cd ~/dbt1-v2.1/data_collect
$ ./dbt1_slave.pl
```

5.3.3.3 Starting the DBT-1 client emulator

Start the DBT-1 workbench. Run the following as the dbt1si user:

```
$ cd ~/dbt1-v2.1/data_collect/oracle
$ ./dbt1_master.pl –f run.config
```

It is possible to specify the configuration file in the –f option. When running benchmarks for two or more cases, it is advisable that the configuration data be saved in separate files and specified in the –f option as required.

5.3.4 Collecting the Execution Results

When the benchmark is finished, its execution results are directed to out_dir specified in the run.config file as with the PostgreSQL version of DBT-1. Unlike the PostgreSQL version of DBT-1, however, neither "db_stat" nor "ipcs" directories are created in /tmp since this version of DBT-1 does not snap the DB state.

The major items to be referenced are the ratios and average responses (in seconds) for each transaction type and BT/seconds (bogo transactions/sec = number of requests via web/sec), which are found in the file called BT. Some items of information are obtainable in addition to these. Typical execution results are summarized in Table 5.3-2.

Table 5.3-2 Files Recording Execution Results

File	Description	Remarks
BT	Ratios and average responses (in seconds) for each transaction type and BT/seconds	Throughputs in short
config.txt	Run-time OS parameters, CPU, memory size, and other system environment information, record of DBT-1 configuration	
param.out	List of PostgreSQL's parameters	
run.meminfo0.out	Memory-related information at the start of DBT-1	
run.meminfo1.out	Memory-related information at the end of DBT-1	
indexes.out	Utilization statistics about PostgreSQL's user indexes at the start of DBT-1	
run.iostat.out	Information about system I/O	Equivalent to iostat -d 60. See man iostat.
run.vmstat.out	Statistics about system memory, I/O, and CPU	Equivalent vmstat 60. See man vmstat.
result.mix.log	Transaction type, request time, and response time for each transaction	BT contains edited forms of these items.
ips.csv	Progress of average transactions at 30 second intervals	The input is result.mix.log

5.4 Concluding Remarks

We have verified that DBT-1 and the procedure described in this chapter which we developed for the current project can successfully be used to measure the performance of Oracle transactions. It is conceivable, however, that the Oracle's performance limit may exceed the load capacity of DBT-1's dbdriver proper depending on proper database tuning or on the specifications of the hardware used (a situation in which the performance limit cannot be reached by increasing the value of eu). In such a case, it is necessary to take appropriate measures such as using two or more occurrences of dbdriver and appServer or setting the think time to a smaller value. In conclusion, we want to add that, during the course of the development stage, we could run DBT-1 up to an eu value of 2,000 in this environment with RedHat and that we could make measurements that exceeded the performance limit of the database we used in the current project.