

## Looking for Real Exam Questions for IT Certification Exams!

We guarantee you can pass any IT certification exam at your first attempt with just 10-12 hours study of our guides.

Our study guides contain actual exam questions, you will get word to word same on your actual test; accurate answers with detailed explanation verified by experts and all graphics and drag-n-drop exhibits shown just as on the real test.

To test the quality of our guides, you can download the one-fourth portion of any guide from <http://www.certificationking.com> absolutely free.

Besides, we also offer complete version of following exams absolutely free. You can start your certification from these free guides and if you are satisfied you can buy the rest

- ◆ Microsoft: 70-270, 70-305 ◆ Cisco: 642-901 ◆ Oracle: 1Z0-007, 200 ◆ CompTIA: 220-601
- ◆ SUN: 310-011, 310-043 ◆ Citrix: 1Y0-A01, 1Y0-256 ◆ CIW: 1D0-420 ◆ Novell: 50-686
- ◆ Adobe: 9A0-029 ◆ Apple: 9L0-005, 9L0-505 ◆ Avaya: 132-S-100 ◆ Cognos: COG-105
- ◆ CWNP: PW0-100 ◆ EMC: E20-001 ◆ Hyperion: 4H0-002 ◆ HP: HP0-771, HP0-J24
- ◆ IBM: 000-253, 000-700 ◆ Juniper: JN0-100, JN0-201 ◆ Lotus: LOT-737
- ◆ Nortel: 920-803 ◆ SAS: A00-201 ◆ SNIA: S10-100 ◆ Sybase: 510-015
- ◆ Symantec: 250-101 ◆ TeraData: NR0-011

For pricing and placing order, please visit <http://certificationking.com/order.html>

We accept all major credit cards through [www.paypal.com](http://www.paypal.com)

For other payment options and any further query, feel free to mail us at [info@certificationking.com](mailto:info@certificationking.com)

## 1z0-054

### QUESTION NO: 1

After running SQL Performance Analyzer (SPA), you observe a few regressed SQL statements in the SPA output. Identify the two actions that you would suggest for these regressed SQL statements. (Choose two.)

- A. Running SQL Access Advisor
- B. Adding them to SQL Plan Baseline
- C. Submitting them to SQL Tuning Advisor
- D. Running Automatic Database Diagnostic Monitor (ADDM)

**Answer:** B, C

### QUESTION NO: 2

View the Exhibit exhibit1 to examine the series of SQL commands and parameter settings.

```
SQL> SHOW PARAMETER OPTIMIZER
```

| NAME                                 | TYPE    | VALUE    |
|--------------------------------------|---------|----------|
| optimizer_capture_sql_plan_baselines | boolean | TRUE     |
| optimizer_dynamic_sampling           | integer | 2        |
| optimizer_features_enable            | string  | 11.1.0.6 |
| optimizer_index_caching              | integer | 0        |
| optimizer_index_cost_adj             | integer | 100      |
| optimizer_mode                       | string  | ALL_ROWS |
| optimizer_secure_view_merging        | boolean | TRUE     |
| optimizer_use_invisible_indexes      | boolean | FALSE    |
| optimizer_use_pending_statistics     | boolean | FALSE    |
| optimizer_use_sql_plan_baselines     | boolean | TRUE     |

```
SQL> SELECT * FROM sh.sales WHERE quantity_sold > 10 ORDER BY prod_id;
SQL> SELECT * FROM sh.sales WHERE quantity_sold > 40 ORDER BY prod_id;
SQL> ALTER SESSION SET OPTIMIZER_MODE=FIRST_ROWS;
SQL> SELECT * FROM sh.sales WHERE quantity_sold > 40 ORDER BY prod_id;
```

View the Exhibit exhibit2 to examine the plans available in the SQL plan baseline.

## 1z0-054

| Select                   | Name                          | SQL Text   | Enabled | Accepted | Fixed | Auto Purge | Created                 | Last Modified           |
|--------------------------|-------------------------------|--|---------|----------|-------|------------|-------------------------|-------------------------|
| <input type="checkbox"/> | SYS_SQL_PLAN_89447021cf314e9e | select * from hr.employees where job_id='CLERK'              | YES     | YES      | NO    | YES        | Jul 20, 2008 7:02:30 PM | Jul 20, 2008 7:16:48 PM |
| <input type="checkbox"/> | SYS_SQL_PLAN_891470210572d2e8 | select * from hr.employees where job_id='CLERK'              | YES     | NO       | NO    | YES        | Jul 20, 2008 7:20:45 PM | Jul 20, 2008 7:20:45 PM |
| <input type="checkbox"/> | SYS_SQL_PLAN_7ed8568135b3cdca | SELECT NAME<br>NAME_COL_PLUS_SHOW_PARAM,DECODE<br>(TYPE,1... | YES     | YES      | NO    | YES        | Jul 21, 2008 2:40:44 PM | Jul 21, 2008 2:40:44 PM |
| <input type="checkbox"/> | SYS_SQL_PLAN_4698b35ddf463620 | select * from table(dbms_xplan.display (null,null,'...       | YES     | YES      | NO    | YES        | Jul 20, 2008 7:04:22 PM | Jul 20, 2008 7:04:22 PM |
| <input type="checkbox"/> | SYS_SQL_PLAN_467a776254bc8843 | select * from sh.sales where quantity_sold > 40 or...        | YES     | YES      | NO    | YES        | Jul 21, 2008 2:25:42 PM | Jul 21, 2008 2:25:42 PM |
| <input type="checkbox"/> | SYS_SQL_PLAN_467a776211df68d0 | select * from sh.sales where quantity_sold > 40 or...        | YES     | NO       | YES   | YES        | Jul 21, 2008 2:41:22 PM | Jul 21, 2008 2:41:56 PM |

The first plan (in red) is created when OPTIMIZER\_MODE is set to ALL\_ROWS and the second plan (in blue) is created when OPTIMIZER\_MODE is set to FIRST\_ROWS. Which SQL plan baseline would be used if the SQL query in exhibit1 is executed again when the value of OPTIMIZER\_MODE is set to FIRST\_ROWS?

- A. the second plan, because it is a fixed plan
- B. the first plan, because it is an accepted plan
- C. the second plan, because it is the latest generated plan in FIRST\_ROW mode
- D. A new plan, because the second plan in FIRST\_ROW mode is not an accepted plan

**Answer: B**

### QUESTION NO: 3

You work as a DBA for a company and you have the responsibility of managing one of its online transaction processing (OLTP) systems. The database encountered performance-related problems and you generated an Automatic Workload Repository (AWR) report to investigate it further. View the Exhibits and examine the AWR report.

#### Top 5 Timed Foreground Events

| Event                   | Waits  | Time(s) | Avg wait (ms) | % DB time | Wait Class  |
|-------------------------|--------|---------|---------------|-----------|-------------|
| DB CPU                  |        | 584     |               | 29.08     |             |
| library cache: mutex X  | 14,721 | 71      | 5             | 3.53      | Concurrency |
| latch: shared pool      | 1,158  | 55      | 48            | 2.76      | Concurrency |
| cursor: pin S wait on X | 3,777  | 50      | 13            | 2.50      | Concurrency |
| log file sync           | 672    | 17      | 25            | 0.83      | Commit      |

# 1z0-054

## Time Model Statistics

- Total time in database user-calls (DB Time): 2008.5s
- Statistics including the word "background" measure background process time, and so do not contribute to the DB time statistic
- Ordered by % of DB time desc, Statistic name

| Statistic Name                             | Time (s) | % of DB Time |
|--|----------|--------------|
| sql execute elapsed time                   | 1,731.94 | 86.23        |
| DB CPU                                     | 584.11   | 29.08        |
| parse time elapsed                         | 533.72   | 26.57        |
| hard parse elapsed time                    | 416.43   | 20.73        |
| connection management call elapsed time    | 33.26    | 1.66         |
| PL/SQL compilation elapsed time            | 10.58    | 0.53         |
| Java execution elapsed time                | 8.01     | 0.40         |
| failed parse elapsed time                  | 5.20     | 0.26         |
| PL/SQL execution elapsed time              | 3.66     | 0.18         |
| hard parse (sharing criteria) elapsed time | 1.94     | 0.10         |
| hard parse (bind mismatch) elapsed time    | 1.33     | 0.07         |
| sequence load elapsed time                 | 0.41     | 0.02         |
| repeated bind elapsed time                 | 0.05     | 0.00         |
| DB time                                    | 2,008.48 |              |
| background elapsed time                    | 32.06    |              |
| background cpu time                        | 4.79     |              |

## Load Profile

|                   | Per Second | Per Transaction | Per Exec | Per Call |
|-------------------|------------|-----------------|----------|----------|
| DB Time(s):       | 3.8        | 12.6            | 0.01     | 0.00     |
| DB CPU(s):        | 1.1        | 3.7             | 0.00     | 0.00     |
| Redo size:        | 6,062.3    | 20,190.1        |          |          |
| Logical reads:    | 5,982.5    | 19,924.3        |          |          |
| Block changes:    | 25.5       | 84.9            |          |          |
| Physical reads:   | 2,778.2    | 9,252.7         |          |          |
| Physical writes:  | 2.9        | 9.7             |          |          |
| User calls:       | 1,263.4    | 4,207.7         |          |          |
| Parses:           | 506.6      | 1,687.3         |          |          |
| Hard parses:      | 53.3       | 177.5           |          |          |
| W/A MB processed: | 726,646.9  | 2,420,040.5     |          |          |
| Logons:           | 1.1        | 3.5             |          |          |
| Executes:         | 513.1      | 1,708.9         |          |          |
| Rollbacks:        | 0.1        | 0.3             |          |          |
| Transactions:     | 0.3        |                 |          |          |

**A Composite Solution With Just One Click – Certification Guaranteed**

# 1z0-054

## Dictionary Cache Stats

- "Pct Misses" should be very low (< 2% in most cases)
- "Final Usage" is the number of cache entries being used

| Cache                | Get Requests | Pct Miss | Scan Reqs | Pct Miss | Mod Reqs | Final Usage |
|----------------------|--------------|----------|-----------|----------|----------|-------------|
| do_awr_control       | 13           | 69.23    | 0         |          | 2        | 1           |
| do_database_links    | 1,074        | 0.56     | 0         |          | 0        | 0           |
| do_global_oids       | 15,419       | 2.67     | 0         |          | 0        | 13          |
| do_histogram_data    | 77,565       | 21.21    | 0         |          | 0        | 571         |
| do_histogram_defs    | 168,045      | 23.16    | 0         |          | 0        | 1,014       |
| do_object_grants     | 44,042       | 4.17     | 0         |          | 0        | 59          |
| do_objects           | 358,789      | 3.30     | 0         |          | 0        | 396         |
| do_profiles          | 548          | 2.19     | 0         |          | 0        | 1           |
| do_rollback_segments | 230          | 0.00     | 0         |          | 0        | 38          |
| do_segments          | 99,805       | 15.72    | 0         |          | 5        | 279         |
| do_sequences         | 25           | 100.00   | 0         |          | 25       | 0           |
| do_tablespace        | 85,868       | 0.04     | 0         |          | 0        | 5           |
| do_users             | 179,387      | 0.35     | 0         |          | 0        | 20          |
| global database name | 927          | 0.11     | 0         |          | 0        | 1           |
| kqsubheap_object     | 197          | 30.46    | 0         |          | 0        | 0           |
| outstanding_alerts   | 19           | 94.74    | 0         |          | 0        | 1           |

[Back to Top](#)

## Library Cache Activity

- "Pct Misses" should be very low

| Namespace       | Get Requests | Pct Miss | Pin Requests | Pct Miss | Reloads | Invali- dations |
|-----------------|--------------|----------|--------------|----------|---------|-----------------|
| BODY            | 1,832        | 1.36     | 3,673        | 1.55     | 23      | 0               |
| CLUSTER         | 2,761        | 1.81     | 1,590        | 3.14     | 0       | 0               |
| INDEX           | 947          | 35.59    | 947          | 35.00    | 1       | 0               |
| JAVA DATA       | 4            | 75.00    | 873          | 0.89     | 0       | 0               |
| SQL AREA        | 340,330      | 23.79    | 602,683      | 12.78    | 22,142  | 5,231           |
| TABLE/PROCEDURE | 145,469      | 2.49     | 191,059      | 8.55     | 5,812   | 0               |
| TRIGGER         | 5,539        | 0.23     | 5,539        | 0.29     | 0       | 0               |

What could be the problem in this database?

- A. Java pool is not configured.
- B. The CPU in the system is slow.
- C. The shared pool size is inadequate.
- D. The database buffer cache is inadequate.
- E. The OPEN\_CURSORS parameter is set to a small value.

**Answer: C**

**A Composite Solution With Just One Click – Certification Guaranteed**

# 1z0-054

## QUESTION NO: 4

You are working on an online transaction processing (OLTP) system. You detected free buffer waits events for your database instance running in a machine that has multiple CPUs. You increased the database buffer cache size as the first step. After a few hours of work on the database, further investigation shows that the same event is being recorded. What would be your next step to avoid this event in future?

- A. Decrease the value of the DBWR\_IO\_SLAVES parameter.
- B. Set the USE\_INDIRECT\_DATA\_BUFFERS parameter to TRUE.
- C. Increase the value of the DB\_WRITER\_PROCESSES parameter.
- D. Increase the value of the DB\_FILE\_MULTIBLOCK\_READ\_COUNT parameter.

**Answer: C**

## QUESTION NO: 5

You are working on a development database that was upgraded to Oracle Database 11g from Oracle Database 9i. An ADDM finding in this database says that the shared pool is inadequately sized, as shown in the Exhibit.



You diagnosed that this is due to different kinds of workloads and this occurs only during peak hours. You tried to resize this by shrinking the database buffer cache but that caused inadequate buffer cache problems. The following are the related parameter settings:

```
SQL> show parameter sga
```

```
NAME TYPE VALUE
```

```
-----
```

```
lock_sga boolean FALSE
```

```
pre_page_sga boolean FALSE
```

```
sga_max_size big integer 300M
```

```
sga_target big integer 0
```

```
SQL> show parameter target
```

```
NAME TYPE VALUE
```

```
-----
```

```
.....
```

```
fast_start_mttr_target integer 0
```

```
memory_max_target big integer 0
```

**A Composite Solution With Just One Click – Certification Guaranteed**

## 1z0-054

memory\_target big integer 0  
pga\_aggregate\_target big integer 100M  
sga\_target big integer 0

You want to balance the memory between the System Global Area (SGA) components within SGA without affecting the size of the Program Global Area (PGA). Which action would solve this problem?

- A. Set the SGA\_TARGET parameter to 300M.
- B. Set the SGA\_MAX\_SIZE parameter to 400M.
- C. Set the MEMORY\_TARGET parameter to 100M.
- D. Set the MEMORY\_MAX\_TARGET parameter to 300M.

**Answer: A**

### QUESTION NO: 6

You observed that some of the queries are performing poorly on the SALES\_RECORDS table in your database.

On further investigation, you find that at the end of each day the contents of the SALES\_RECORDS table are transferred to the SALES table and deleted from the SALES\_RECORDS table. The deleted operations cause the table to be sparsely populated. You decided to use the ALTER TABLE...SHRINK SPACE COMPACT command to shrink the table. Why would you choose this method? (Choose all that apply.)

- A. because it can be used during peak hours
- B. because it avoids unnecessary cursor invalidation
- C. because it adjusts the high-water mark (HWM) immediately
- D. because you have long-running queries that might span the shrink operation
- E. because it does not allow any data manipulation language (DML) operations, thereby making the shrink operation faster

**Answer: A, B, D**

### QUESTION NO: 7

View the Exhibit and examine a portion of the output obtained from the following query:

## 1z0-054

| STAT_ID    | STAT_NAME                               | VALUE    |
|------------|---|----------|
| 3649082374 | DB time                                 | 61021783 |
| 2748282437 | DB CPU                                  | 3890625  |
| 4157170894 | background elapsed time                 | 42472524 |
| 2451517896 | background cpu time                     | 2796875  |
| 4127043053 | sequence load elapsed time              | 0        |
| 1431595225 | parse time elapsed                      | 10983653 |
| 372226525  | hard parse elapsed time                 | 10480831 |
| 2821698184 | sql execute elapsed time                | 50353110 |
| 1990024365 | connection management call elapsed time | 855906   |

SQL> SELECT \* FROM v\$sys\_time\_model;

Select three correct interpretations of the time model statistics. (Choose three.)

- A. DB time includes the wait time of all the nonidle and idle user sessions.
- B. SQL execute elapsed time includes the time spent in performing fetches of query results.
- C. DB CPU includes the CPU time spent on database user-level calls and background CPU time.
- D. SQL execute elapsed time includes components of the hard parse elapsed time like bind elapsed time.
- E. DB time includes the connection management call elapsed time excluding the background process time.

**Answer:** B, D, E

### QUESTION NO: 8

A batch workload that historically completed in the maintenance window between 10:00 PM and midnight is currently showing poor performance and completing at 2 AM. To help in the diagnosis of the performance degradation, the senior DBA in your organization asks you to execute the awrddrpt.sql script to produce a Compare Periods report. Which two statements are true about the report produced by this script? (Choose two.)

- A. It is refreshed automatically based on the moving window baseline.
- B. It compares details between any two selected time periods of the same duration.
- C. It normalizes the statistics by the amount of time spent on the database for each time period.
- D. It compares details between two consecutive time periods of the same or different durations and is refreshed every 60 minutes.

**Answer:** B, C

### QUESTION NO: 9

You plan to use adaptive thresholds as part of the performance tuning activity. You decide to

**A Composite Solution With Just One Click – Certification Guaranteed**



## 1z0-054

increase the window size of the default moving window baseline for all metric observations and comparisons in your database. The following error occurs when you try to increase the window size through Enterprise Manager:

Failed to commit: ORA-13541: system moving window baseline size (1296000) greater than retention (1036800) ORA-06512: at "SYS.DBMS\_WORKLOAD\_REPOSITORY", line 601  
ORA-06512: at line 2 Which action would allow you to perform the preceding task successfully?

- A. increasing the flashback retention period
- B. increasing the retention period for SQL Management Base
- C. increasing the undo retention period for the database instance
- D. increasing the Automatic Workload Repository (AWR) retention period

**Answer: D**

### QUESTION NO: 10

Identify two correct statements about the Active Session History (ASH) data. (Choose two.)

- A. A part of SGA memory is used to store ASH data as rolling buffer.
- B. The ASH data can be analyzed between any two small time intervals.
- C. All ASH data in memory are flushed to disk by MMON in every 3 seconds.
- D. All ASH data in memory are flushed to disk by MMNL process whenever the buffer is full.

**Answer: A, B**

### QUESTION NO: 11

A user in a session executed the following SQL statement to set the optimizer mode:

```
ALTER SESSION SET OPTIMIZER_MODE = ALL_ROWS
```

What impact would it have on the goal of the optimizer for that session? (Choose all that apply.)

- A. Statement level OPTIMIZER\_MODE hints take precedence over the session-level setting.
- B. The OPTIMIZER\_MODE parameter set at instance level takes precedence over the session-level value.
- C. The optimizer uses a cost-based approach, regardless of the presence of statistics; it optimizes with a goal of best response time.
- D. The optimizer uses a cost-based approach for all SQL statements in the session, regardless of the presence of statistics; it optimizes with a goal of best throughput.

**Answer: A, D**

### QUESTION NO: 12

You are working on an online transaction processing (OLTP) system. By day most of the

**A Composite Solution With Just One Click – Certification Guaranteed**

## 1z0-054

application users perform queries accessing the most recently added or modified rows. The applications have most of the queries based on multiple tables. But at night, some batch processing is also done. Which two actions would you recommend to choose a goal for the optimizer based on the needs of your application? (Choose two.)

- A. setting the OPTIMIZER\_MODE parameter to ALL\_ROWS at the instance level
- B. setting the OPTIMIZER\_MODE parameter to FIRST\_ROWS\_n at the instance level
- C. asking the developer to add a hint /\*ALL\_ROWS\*/ in the long-running batch processing queries
- D. asking the developer to add a hint /\*FIRST\_ROWS\_n\*/ in the long-running batch processing queries

**Answer:** B, C

### QUESTION NO: 13

Which three factors influence the optimizer's behavior while choosing an optimization approach and goal for a SQL statement? (Choose three.)

- A. parsing of a SQL statement
- B. operating system (OS) statistics
- C. object statistics in the data dictionary
- D. the OPTIMIZER\_MODE initialization parameter
- E. optimizer SQL hints for changing the query optimizer goal

**Answer:** C, D, E

### QUESTION NO: 14

Examine the initialization parameter values for the instance given below:

NAME TYPE VALUE

```
-----  
optimizer_capture_sql_plan_baselines boolean FALSE  
optimizer_dynamic_sampling integer 2  
optimizer_features_enable string 11.1.0.6  
optimizer_index_caching integer 0  
optimizer_index_cost_adj integer 100  
optimizer_mode string ALL_ROWS  
db_file_multiblock_read_count integer 64
```

The index created on the column used in the WHERE clause of the query. You notice that the query is not using the index. Instead of an index scan, a full table scan is used.

View the Exhibit and examine the autotrace output for a query.

**A Composite Solution With Just One Click – Certification Guaranteed**

# 1z0-054

```
select * from employees where employee_id=107;
```

Execution Plan

-----  
Plan hash value: 1601196873

| Id  | Operation         | Name | Rows | Bytes | Cost (%CPU) | Time     |
|-----|-------------------|------|------|-------|-------------|----------|
| 0   | SELECT STATEMENT  |      | 1    | 71    | 3 (0)       | 00:00:01 |
| * 1 | TABLE ACCESS FULL | T    | 1    | 71    | 3 (0)       | 00:00:01 |

Predicate Information (identified by operation id):

-----  
1 - filter("EMPLOYEE\_ID"=107)

What could be the reason for it? (Choose all that apply.)

- A. The OPTIMIZER\_INDEX\_COST\_ADJ initialization parameter has a low value.
- B. The DB\_FILE\_MULTIBLOCK\_READ\_COUNT initialization parameter has a low value.
- C. The statistics for the table and all the indexes associated with the table are not current.
- D. The table has less than DB\_FILE\_MULTIBLOCK\_READ\_COUNT blocks under the high-water mark.

**Answer:** C, D

## QUESTION NO: 15

Examine the initialization parameter values for the instance given below:

NAME TYPE VALUE

-----  
optimizer\_capture\_sql\_plan\_baselines boolean FALSE  
optimizer\_dynamic\_sampling integer 2  
optimizer\_features\_enable string 11.1.0.6  
optimizer\_index\_caching integer 0  
optimizer\_index\_cost\_adj integer 100  
optimizer\_mode string ALL\_ROWS  
db\_file\_multiblock\_read\_count integer 64

You notice that the one of the queries is using a full table scan (view Exhibit1) instead of index unique scan (view Exhibit2). The index is present on the column that is accessed in the WHERE clause of the query. The cost for a full table scan is more than that for an index unique scan.

# 1z0-054

```
select * from employees where employee_id=107;
```

Execution Plan

Plan hash value: 1601196873

| Id  | Operation         | Name | Rows | Bytes | Cost (%CPU) | Time     |
|-----|-------------------|------|------|-------|-------------|----------|
| 0   | SELECT STATEMENT  |      | 1    | 71    | 3 (0)       | 00:00:01 |
| * 1 | TABLE ACCESS FULL | T    | 1    | 71    | 3 (0)       | 00:00:01 |

Predicate Information (identified by operation id):

1 - filter("EMPLOYEE\_ID"=107)

Plan hash value: 1076294677

| Id  | Operation                   | Name   | Rows | Bytes | Cost (%CPU) | Time     |
|-----|-----------------------------|--------|------|-------|-------------|----------|
| 0   | SELECT STATEMENT            |        | 1    | 71    | 1 (0)       | 00:00:01 |
| 1   | TABLE ACCESS BY INDEX ROWID | T      | 1    | 71    | 1 (0)       | 00:00:01 |
| * 2 | INDEX UNIQUE SCAN           | EMP_PK | 1    |       | 1 (0)       | 00:00:01 |

Predicate Information (identified by operation id):

2 - access("EMPLOYEE\_ID"=107)

Why would the optimizer choose full table scan over index unique scan? (Choose all that apply.)

- A. The OPTIMIZER\_INDEX\_COST\_ADJ initialization parameter is set to a low value.
- B. The OPTIMIZER\_INDEX\_COST\_ADJ initialization parameter is set to a high value.
- C. The DB\_FILE\_MULTIBLOCK\_READ\_COUNT initialization parameter is set to a low value.
- D. The statistics for the table and all the indexes associated with the table are not current.

**Answer:** B, D

## QUESTION NO: 16

The columns CUST\_CITY, CUST\_STATE\_PROVINCE, and COUNTRY\_ID are frequently used together in the WHERE clause of the queries. The CUSTOMERS table is a big table with 20 GB of data. You notice that the selectivity for these three columns varies from the selectivity that the optimizer calculates. What would you recommend to influence the selectivity calculated by the optimizer?

- A. creating function-based indexes by concatenating all the columns
- B. updating histogram statistics for these columns by using the DBMS\_STATS.GATHER\_TABLE\_STATS procedure

**A Composite Solution With Just One Click – Certification Guaranteed**

## 1z0-054

C. using the DBMS\_STATS.CREATE\_EXTENDED\_STATS function to create a virtual column and create index on the virtual column

D. using the DBMS\_STATS.CREATE\_EXTENDED\_STATS function to create a virtual column and the DBMS\_STATS.GATHER\_TABLE\_STATS procedure to collect statistics on the virtual column

**Answer: D**

### QUESTION NO: 17

View Exhibit1 to examine the description of the CUSTOMERS table.

```
SQL> DESC customers
Name                               Null?    Type
-----
CUST_ID                             NOT NULL NUMBER
CUST_FIRST_NAME                     NOT NULL VARCHAR2 (20)
CUST_LAST_NAME                      NOT NULL VARCHAR2 (40)
CUST_GENDER                         NOT NULL CHAR (1)
CUST_YEAR_OF_BIRTH                 NOT NULL NUMBER (4)
CUST_MARITAL_STATUS                VARCHAR2 (20)
CUST_STREET_ADDRESS                NOT NULL VARCHAR2 (40)
CUST_POSTAL_CODE                   NOT NULL VARCHAR2 (10)
CUST_CITY                          NOT NULL VARCHAR2 (30)
CUST_CITY_ID                       NOT NULL NUMBER
CUST_STATE_PROVINCE                NOT NULL VARCHAR2 (40)
CUST_STATE_PROVINCE_ID             NOT NULL NUMBER
COUNTRY_ID                         NOT NULL NUMBER
CUST_MAIN_PHONE_NUMBER             NOT NULL VARCHAR2 (25)
CUST_INCOME_LEVEL                  VARCHAR2 (30)
CUST_CREDIT_LIMIT                  NUMBER
CUST_EMAIL                         VARCHAR2 (30)
CUST_TOTAL                        NOT NULL VARCHAR2 (14)
CUST_TOTAL_ID                     NOT NULL NUMBER
CUST_SRC_ID                        NUMBER
CUST_EFF_FROM                      DATE
CUST_EFF_TO                        DATE
CUST_VALID                         VARCHAR2 (1)
```

You observed that optimizer selectivity is not accurate when the CUST\_STATE\_PROVINCE and COUNTRY\_ID columns are used together in the WHERE clause of a query. View Exhibit2 to examine the query execution plan and the commands executed to gather the statistics.

## 1z0-054

```
SQL> SELECT COUNT(*) FROM customers WHERE country_id=52790 and cust_state_province='CA';

COUNT(*)
-----
      3341

SQL> exec dbms_stats.gather_table_stats(null,'CUSTOMERS',method_opt=>'for all columns size 1');

PL/SQL procedure successfully completed.

SQL> explain plan for select count(*) from customers where COUNTRY_ID=52790 and CUST_STATE_PROVINCE='CA';

Explained.

SQL> select * from table(dbms_xplan.display())
2 ;

PLAN_TABLE_OUTPUT
-----
Plan hash value: 296924608

-----
| Id | Operation          | Name          | Rows  | Bytes | Cost (%CPU)| Time     |
-----|-----|-----|-----|-----|-----|-----|-----|
|  0 | SELECT STATEMENT   |               |      1 |    16 |    406 (1)| 00:00:05 |
|  1 |   SORT AGGREGATE   |               |      1 |    16 |           |          |
|* 2 |    TABLE ACCESS FULL| CUSTOMERS    |     20 |   320 |    406 (1)| 00:00:05 |
-----

Predicate Information (identified by operation id):

PLAN_TABLE_OUTPUT
-----

      2 - filter("CUST_STATE_PROVINCE"='CA' AND "COUNTRY_ID"=52790)

14 rows selected.
```

The optimizer predicts that 20 rows will be processed rather than the 3,341 rows, which is the actual number of rows returned from the table. What can you do to make the optimizer detect the actual number of rows?

- A. Set the STATISTICS\_LEVEL parameter to ALL.
- B. Set the OPTIMIZER\_USE\_PENDING\_STATISTICS parameter to FALSE.
- C. Create extended statistics for the CUST\_STATE\_PROVINCE and COUNTRY\_ID columns.
- D. Increase the STALE\_PERCENT value for the CUSTOMERS table by using the DBMS\_STATS.SET\_TABLE\_PREFS procedure.

**Answer: C**

### QUESTION NO: 18

View Exhibit1 to examine the description of the CUSTOMERS table.

# 1z0-054

```
SQL> DESC customers
Name                                         Null?    Type
-----
CUST_ID                                     NOT NULL NUMBER
CUST_FIRST_NAME                             NOT NULL VARCHAR2 (20)
CUST_LAST_NAME                              NOT NULL VARCHAR2 (40)
CUST_GENDER                                 NOT NULL CHAR (1)
CUST_YEAR_OF_BIRTH                          NOT NULL NUMBER (4)
CUST_MARITAL_STATUS                         VARCHAR2 (20)
CUST_STREET_ADDRESS                         NOT NULL VARCHAR2 (40)
CUST_POSTAL_CODE                            NOT NULL VARCHAR2 (10)
CUST_CITY                                   NOT NULL VARCHAR2 (30)
CUST_CITY_ID                               NOT NULL NUMBER
CUST_STATE_PROVINCE                         NOT NULL VARCHAR2 (40)
CUST_STATE_PROVINCE_ID                     NOT NULL NUMBER
COUNTRY_ID                                  NOT NULL NUMBER
CUST_MAIN_PHONE_NUMBER                     NOT NULL VARCHAR2 (25)
CUST_INCOME_LEVEL                           VARCHAR2 (30)
CUST_CREDIT_LIMIT                           NUMBER
CUST_EMAIL                                  VARCHAR2 (30)
CUST_TOTAL                                  NOT NULL VARCHAR2 (14)
CUST_TOTAL_ID                              NOT NULL NUMBER
CUST_SRC_ID                                 NUMBER
CUST_EFFECTIVE_FROM                         DATE
CUST_EFFECTIVE_TO                           DATE
CUST_VALID                                  VARCHAR2 (1)
```

The CUSTOMERS table has been updated heavily today. In a frequently used SQL statement, you notice that estimated rows and the actual number of rows fetched differ greatly. The COUNTRY\_ID column has an index.

View Exhibit2 and examine the query execution plan.

```
SQL> SELECT cust_id, cust_last_name, cust_total
2 FROM customers
3 WHERE country_id = 52790

296320 rows selected.

Execution Plan
-----
Plan hash value: 2008213504

-----
| Id | Operation          | Name          | Rows  | Bytes | Cost (%CPU)| Time     |
-----
|  0 | SELECT STATEMENT   |               |       |       |              |          |
|*  1 | TABLE ACCESS FULL| CUSTOMERS    | 73966 | 2383K | 798 (1)    | 00:00:10 |
-----

Predicate Information (identified by operation id):
-----

   1 - filter("COUNTRY_ID"=52790)

Statistics
-----
   1 recursive calls
   0 db block gets
 27844 consistent gets
  8597 physical reads
   0 redo size
6373693 bytes sent via SQL*Net to client
217714 bytes received via SQL*Net from client
 19756 SQL*Net roundtrips to/from client
   0 sorts (memory)
   0 sorts (disk)
296320 rows processed
```

What would you recommend to improve the optimizer's estimation?

## 1z0-054

- A. setting the STATISTICS\_LEVEL parameter to ALL
- B. setting the OPTIMIZER\_USE\_PENDING\_STATISTICS parameter to FALSE
- C. creating extended statistics for the CUST\_LAST\_NAME, CUST\_ID, and CUST\_TOTAL columns
- D. updating the statistics for the CUSTOMERS table by using the DBMS\_STATS.GATHER\_TABLE\_STATS procedure

**Answer:** D

### **QUESTION NO:** 19

You are working on a decision support system (DSS). The index is available on the COUNTRY\_ID column of the CUSTOMERS table. View the Exhibit and examine the parameter settings and the query execution plan.



## 1z0-054

```
NAME                                TYPE                                VALUE
-----                                -----                                -----
db_file_multiblock_read_count       integer                             49
```

```
SQL> SELECT BLOCKS, EMPTY_BLOCKS FROM ALL_TABLES
WHERE TABLE_NAME = 'CUST';
```

```
      BLOCKS  EMPTY_BLOCKS
-----  -
      2902         0
```

```
SQL> EXEC DBMS_STATS.GATHER_TABLE_STATS('SH','CUST');
```

PL/SQL procedure successfully completed.

```
SQL> select cust_id, cust_last_name from sh.cust where country_id=52790;
```

74080 rows selected.

### Execution Plan

Plan hash value: 260468903

```
-----
| Id | Operation          | Name | Rows  | Bytes | Cost (%CPU)| Time     |
-----
|  0 | SELECT STATEMENT   |      | 74150 | 1303K | 793  (1)| 00:00:10 |
|*  1 | TABLE ACCESS FULL| CUST | 74150 | 1303K | 793  (1)| 00:00:10 |
-----
```

Predicate Information (identified by operation id):

```
-----
1 - filter("COUNTRY_ID"=52790)
```

### Statistics

```
-----
0 recursive calls
0 db block gets
7659 consistent gets
2848 physical reads
0 redo size
1584703 bytes sent via SQL*Net to client
54738 bytes received via SQL*Net from client
4940 SQL*Net roundtrips to/from client
```

Why is the query using a full table scan instead of an index scan?

- A. because the histogram statistics for the COUNTRY\_ID column are not updated
- B. because the index statistics for the index on the COUNTRY\_ID column are not current
- C. because the DB\_FILE\_MULTIBLOCK\_READ\_COUNT initialization parameter is set to a higher value
- D. because the optimizer predicts that most of the blocks in the table are accessed. Therefore, it uses a full table scan, even though indexes are available.

**Answer: D**

**QUESTION NO: 20**

View Exhibit1 and examine the indexes on the CUSTOMERS table.

**A Composite Solution With Just One Click – Certification Guaranteed**

## 1z0-054

```
SQL> SELECT i.table_name, column_name, o.object_name , o.object_id, i.index_type
FROM   user_objects o,
       user_indexes i,
       user_ind_columns c
WHERE  o.object_type= 'INDEX'
AND    i.table_name LIKE '%CUST%'
AND    i.index_name = o.object_name
AND    i.index_name = c.index_name
```

| TABLE_NAME | COLUMN_NAME         | OBJECT_NAME           | OBJECT_ID | INDEX_TYPE |
|------------|---------------------|-----------------------|-----------|------------|
| CUSTOMERS  | CUST_GENDER         | CUSTOMERS_GENDER_BIX  | 70685     | BITMAP     |
| CUSTOMERS  | CUST_MARITAL_STATUS | CUSTOMERS_MARITAL_BIX | 70686     | BITMAP     |
| CUSTOMERS  | CUST_ID             | CUSTOMERS_PK          | 70473     | NORMAL     |
| CUSTOMERS  | CUST_YEAR_OF_BIRTH  | CUSTOMERS_YOB_BIX     | 70687     | BITMAP     |
| CUSTOMERS  | COUNTRY_ID          | CUST_COUNTRY          | 79047     | NORMAL     |

The statistics for the CUSTOMERS table have been updated recently by using the following command:

```
SQL> EXEC
DBMS_STATS.GATHER_TABLE_STATS('SH','CUSTOMERS',method_opt=>'FOR ALL
INDEXED COLUMNS SIZE AUTO');
```

View Exhibit2 to examine a query plan. Even though the index is present on the COUNTRY\_ID and CUST\_GENDER columns, the query uses a full table scan. What could be the reason?

```
SQL> SELECT i.table_name, column_name, o.object_name , o.object_id, i.index_type
FROM   user_objects o,
       user_indexes i,
       user_ind_columns c
WHERE  o.object_type= 'INDEX'
AND    i.table_name LIKE '%CUST%'
AND    i.index_name = o.object_name
AND    i.index_name = c.index_name
```

| TABLE_NAME | COLUMN_NAME         | OBJECT_NAME           | OBJECT_ID | INDEX_TYPE |
|------------|---------------------|-----------------------|-----------|------------|
| CUSTOMERS  | CUST_GENDER         | CUSTOMERS_GENDER_BIX  | 70685     | BITMAP     |
| CUSTOMERS  | CUST_MARITAL_STATUS | CUSTOMERS_MARITAL_BIX | 70686     | BITMAP     |
| CUSTOMERS  | CUST_ID             | CUSTOMERS_PK          | 70473     | NORMAL     |
| CUSTOMERS  | CUST_YEAR_OF_BIRTH  | CUSTOMERS_YOB_BIX     | 70687     | BITMAP     |
| CUSTOMERS  | COUNTRY_ID          | CUST_COUNTRY          | 79047     | NORMAL     |

- A. because the histogram statistics for the COUNTRY\_ID column are not updated
- B. because the DB\_FILE\_MULTIBLOCK\_READ\_COUNT initialization parameter is set to a high value
- C. because the optimizer calculates the cost of accessing blocks by using a full table scan to be less as compared to index scans, even though indexes are available
- D. because indexes on CUST\_GENDER and COUNTRY\_ID columns are of different types, the index on the CUST\_GENDER column is bitmap index, and on COUNTRY\_ID columns is btree index.

**Answer: C**

# 1z0-054

## QUESTION NO: 21

You are working on an online transaction processing (OLTP) system. The CUST table in the SH schema was populated by direct load and after that it has gone through a lot of updates and deletions. The statistics for the CUST and SALES tables were updated recently. View the Exhibit and examine the query plan. The query is performing a lot of I/O for a query that fetches only 168 rows. To investigate further, you queried the ALL\_TABLES view to find out PCTUSED, PCTFREE, and the number of rows in the CUST table, as given below:

```
SQL> SELECT table_name ,blocks, pct_used, pct_free, num_rows  
2 FROM all_tables
```

```
3* WHERE table_name = 'CUST';
```

```
TABLE_NAME BLOCKS PCT_USED PCT_FREE NUM_ROWS
```

```
-----
```

```
CUST 13768 80 111060
```

What would you recommend to reduce the I/O?

```
SQL> SELECT time_id, quantity_sold, amount_sold  
FROM sales s, cust c  
WHERE c.cust_id = s.cust_id AND cust_first_name = 'Jimmy'  
ORDER BY time_id
```

```
168 rows selected.
```

```
Execution Plan
```

```
Plan hash value: 980033780
```

| Id  | Operation           | Name  | Rows  | Bytes | TempSpc | Cost (%CPU) | Time     | Pstart | Pstop |
|-----|---------------------|-------|-------|-------|---------|-------------|----------|--------|-------|
| 0   | SELECT STATEMENT    |       | 11120 | 358K  |         | 4344 (1)    | 00:00:53 |        |       |
| 1   | SORT ORDER BY       |       | 11120 | 358K  | 1064K   | 4344 (1)    | 00:00:53 |        |       |
| * 2 | HASH JOIN           |       | 11120 | 358K  |         | 4243 (1)    | 00:00:51 |        |       |
| * 3 | TABLE ACCESS FULL   | CUST  | 85    | 1020  |         | 3738 (1)    | 00:00:45 |        |       |
| 4   | PARTITION RANGE ALL |       | 918K  | 18M   |         | 498 (4)     | 00:00:06 | 1      | 28    |
| 5   | TABLE ACCESS FULL   | SALES | 918K  | 18M   |         | 498 (4)     | 00:00:06 | 1      | 28    |

```
Predicate Information (identified by operation id):
```

```
2 - access("C"."CUST_ID"="S"."CUST_ID")  
3 - filter("CUST_FIRST_NAME"='Jimmy')
```

```
Statistics
```

```
14863 recursive calls  
0 db block gets  
18448 consistent gets  
13688 physical reads  
0 redo size  
4716 bytes sent via SQL*Net to client  
541 bytes received via SQL*Net from client  
13 SQL*Net roundtrips to/from client  
137 sorts (memory)  
0 sorts (disk)  
168 rows processed
```

- A. reorganizing the table to use fewer blocks
- B. rebuilding the index on the CUST\_FIRST\_NAME column
- C. increasing the value for the PCTFREE attribute for the CUST table
- D. increasing the value for the PGA\_AGGREGATE\_TARGET initialization parameter

**Answer: A**

## 1z0-054

### QUESTION NO: 22

Your system has been upgraded from Oracle Database 10g to Oracle Database 11g. You imported SQL Tuning Sets (STS) from the previous version that contains important SQL statements. You are not sure about the elapsed time that the SQL statements will have in the new system due to the version change of the optimizer. Which tool provides a comparative report with respect to the elapsed time in both versions of the optimizer?

- A. SQL Access Advisor
- B. SQL Tuning Advisor (STA)
- C. SQL Performance Analyzer (SPA)
- D. Automatic Workload Repository (AWR) compare period report

**Answer: C**

### QUESTION NO: 23

Which three statements are true about performance analysis by SQL Performance Analyzer? (Choose three.)

- A. It detects changes in SQL execution plans.
- B. It shows only the overall impact on workload.
- C. It produces results that can be used to create the SQL plan baseline.
- D. It generates recommendations to run SQL Tuning Advisor to tune regressed SQL statements.
- E. The importance of SQL statements to be analyzed is based on the size of the objects accessed.

**Answer: A, C, D**

### QUESTION NO: 24

The database application developers are planning to make some major schema changes such as creating new indexes and materialized views. They want to check the net impact of these changes on the workload performance. This activity has to be performed in the production database, so they want only the query part of the data manipulation language (DML) statements to be executed so that the side effects to the database or user data can be prevented. What should they use to achieve this?

- A. Database Replay
- B. SQL Tuning Advisor
- C. SQL Access Advisor
- D. SQL Performance Analyzer

**Answer: D**

**A Composite Solution With Just One Click – Certification Guaranteed**

## 1z0-054

### QUESTION NO: 25

You upgraded database from the Oracle Database 10g to Oracle Database 11g. To test the performance SQL on the upgraded database. You want to build the before-change performance data by using SQL Performance Analyzer .Which method would allow the task to execute quickly?

- A. the EXPLAIN PLAN method
- B. the TEST EXECUTE method
- C. the COMPARE PERFORMANCE method
- D. the OPTIMIZER\_USE\_SQL\_PLAN\_BASELINE parameter set to TRUE

**Answer:** A

### QUESTION NO: 26

Examine the following block that executes a SQL Performance Analyzer task:

```
begin dbms_sqlpa.execute_analysis_task(task_name => 'SPA_SKD4', execution_type => 'TEST
EXECUTE',
execution_name => 'before');
end;
```

Which two statements are true about the execution? (Choose two.)

- A. It generates execution plans and execution statistics.
- B. It executes only the query part of the DML statements.
- C. The test plans are generated and added to plan baseline.
- D. The test plans are executed, but not generated if they are available in the plan baseline.

**Answer:** A, B

### QUESTION NO: 27

You plan to collect statistics for an important table in your database to help improved query performance. You are uncertain about the optimizer behavior after the statistics have been collected because currently there are no statistics for indexes or partitions on the table. You wish to test the performance effect of the newly collected statistics. How would you prevent the optimizer from using the newly collected statistics for other sessions immediately after statistics collection?

- A. Set the PUBLISH preference to FALSE for the table.
- B. Set the NO\_VALIDATE preference to TRUE for the table.
- C. Set the OPTIMIZER\_DYNAMIC\_SAMPLING parameter to 0.

**A Composite Solution With Just One Click – Certification Guaranteed**

## 1z0-054

- D. Set the OPTIMIZER\_USE\_SQL\_PLAN\_BASELINES parameter to FALSE for the session.
- E. Set the OPTIMIZER\_USE\_PENDING\_STATISTICS parameter to FALSE for the session.

**Answer: A**

### QUESTION NO: 28

You have the automatic optimizer statistics gathering enabled for a 24x7 database as part of the automatic maintenance task. You observed that a few tables are being updated regularly, but not being queried often. However statistics are being gathered for these tables automatically and regularly, thereby affecting the database performance.

What would you do to increase the interval of statistics collection for these tables so that the statistics collection jobs spend less time during the automatic maintenance task?

- A. Set the OPTIMIZER\_USE\_PENDING parameter statistics to TRUE.
- B. Increase the value of the STALE\_PERCENT preference for these tables.
- C. Mark those associated indexes that are less frequently used as invisible.
- D. Increase the value of the ESTIMATE\_PERCENT preference for these tables

**Answer: B**

### QUESTION NO: 29

You have the automatic optimizer statistics gathering enabled for a 24x7 database as part of the automatic maintenance task. Because the automatic optimizer statistics collection runs during maintenance window(in the night), the statistics on tables that are significantly modified throughout the day becomes stale. This leads to poor query performance on these tables and you want to minimize the overhead in statistics collection. Identify a solution to this problem.

- A. Change the STALE\_PERCENT preference to zero for the tables.
- B. Set the OPTIMIZER\_USE\_PENDING parameter statistics to TRUE.
- C. Set the OPTIMIZER\_DYNAMIC\_SAMPLING parameter to zero during day time.
- D. Gather statistics on the table when the tables have a representative number of rows, and then lock the statistics.

**Answer: D**

### QUESTION NO: 30

You observe that suboptimal execution plans for the queries are being generated on a table that previously used less resources. You have collected statistics on these tables two days ago. The optimizer statistics retention period is set to 31 days. You are able to find the timestamp information about statistics update from the DBA\_TAB\_STATS\_HISTORY view. Because it is a

**A Composite Solution With Just One Click – Certification Guaranteed**

## 1z0-054

frequently queried table, you would like the optimizer to generate better plans.

Which action would enable you to use the previous set of statistics on the objects that may lead to better execution plans?

- A. restoring statistics from statistics history up to the desired time
- B. deleting all AWR snapshots collected after the time of desired statistics collection
- C. applying the flashback table technique until the time of desired statistics collection
- D. setting the `OPTIMIZER_PENDING_STATISTICS` parameter to `TRUE` to use the previous version of statistics

**Answer: A**

### QUESTION NO: 31

The `SALES` table in your database undergoes frequent changes. While observing the statistics for the table in the middle of the day, you find a change since the last observation. But the information about the statistics collection is not recorded in the `DBA_TAB_STATS_HISTORY`. Identify the reason for this.

- A. The AWR retention time was set to zero.
- B. The `STATISTICS_LEVEL` parameter was set to `BASIC`.
- C. Some of the statistics were collected using the `ANALYZE` command.
- D. There were user-defined statistics that were set by using the `DBMS_STATS.SET_TABLE_STATS` procedure.

**Answer: C**

### QUESTION NO: 32

You have a range-partitioned table in your database. Each partition in the table contains the sales data for a quarter. The partition related to the current quarter is modified frequently, whereas other partitions undergo fewer data manipulations. The preferences for the table are set to their default values. You collect statistics for the table using the following command at regular intervals:

```
SQL> EXECUTE  
DBMS_STATS.GATHER_TABLE_STATS('SH','SALES',GRANULARITY=>'AUTO');
```

You need statistics to be collected more quickly. How do you achieve this?

- A. Set the `PUBLISH` preference to `TRUE` for the partition table.
- B. Set the `NO_VALIDATE` preference to `TRUE` for the partition table.
- C. Set the `INCREMENTAL` preference to `TRUE` for the partition table.
- D. Increase the value of the `STALE_PERCENT` preference for the partition table.

**Answer: C**

**A Composite Solution With Just One Click – Certification Guaranteed**

## 1z0-054

### QUESTION NO: 33

Which is the correct description of SQL profiling while using SQL Tuning Advisor?

- A. It is a set of recommendations by the optimizer to create new indexes.
- B. It is auxiliary information collected by the optimizer for a SQL statement to eliminate estimation error.
- C. It is a set of recommendations by the optimizer to refresh stale statistics to avoid bad execution plan.
- D. It is a set of recommendations by the optimizer to restructure a SQL statement to avoid bad execution plan.

**Answer: B**

### QUESTION NO: 34

You used SQL Tuning Advisor for a long-running SQL statement that suggested a SQL profile which can be used by the query subsequently for a better execution plan. After implementing the profile, you altered the profile to include it in a category as follows:

```
SQL> EXEC DBMS_SQLTUNE.ALTER_SQL_PROFILE ( name=>
'SYS_SQLPROF_0146ced728214000, attribute_name => 'CATEGORY', value => 'DEV2').
```

But the performance has not improved for the user session. The profile is enabled and on investigating the query plan, you find that the query from the user session does not use the SQL profile. View the Exhibit and examine the parameters for the user session. What could be the reason for this?

| NAME                                 | TYPE    | VALUE                          |
|--------------------------------------|---------|--------------------------------|
| optimizer_capture_sql_plan_baselines | boolean | FALSE                          |
| optimizer_use_sql_plan_baselines     | boolean | TRUE                           |
| sql_trace                            | boolean | FALSE                          |
| sql_version                          | string  | NATIVE                         |
| sqltune_category                     | string  | DEFAULT                        |
| optimizer_mode                       | string  | ALL_ROWS                       |
| control_management_pack_access       | string  | DIAGNOSTIC+TUNING              |
| resource_manager_cpu_allocation      | integer | 2                              |
| resource_manager_plan                | string  | SCHEDULER[0x2C0F]:DEFAULT_MAIN |

- A. The CONTROL\_MANAGEMENT\_PACK\_ACCESS parameter is not set to DEV2.
- B. The OPTIMIZER\_MODE parameter is set to ALL\_ROWS for the user session.
- C. The SQLTUNE\_CATEGORY parameter is set to DEFAULT for the user session.
- D. The OPTIMIZER\_USE\_PLAN\_BASELINES parameter is set to TRUE for the user session.

**Answer: C**

### QUESTION NO: 35

You used SQL Tuning Advisor for a long-running SQL statement that suggested a SQL profile which can be used by the query subsequently for better execution plan. View the Exhibit. But you

**A Composite Solution With Just One Click – Certification Guaranteed**



# 1z0-054

want certain user sessions not to use this SQL profile for their queries. How would you implement this?

| Select Name   | SQL Text  | Category | Status  | Created                 | Last Modified           |
|---|---|----------|---------|-------------------------|-------------------------|
| <input type="checkbox"/> SYS_SQLPROF_0146ced728214000 | SELECT /*+ ORDERED USE NL(C) FULL(C) FULL(S) */ CO... | DEFAULT  | ENABLED | Jul 19, 2008 7:47:50 PM | Jul 20, 2008 1:07:39 PM |

TIP The table will display maximum of 2000 rows. Use search criteria to get the desired results.

- A. Alter the SQL profile to change the category of the SQL profile.
- B. Set the OPTIMIZER\_USE\_PENDING\_STATISTICS to TRUE the desired sessions.
- C. Use database resource manager to prevent the use of the SQL profile by these user sessions.
- D. Use database resource manager to preUse the resource management feature in profiles of these users to prevent the use of the SQL profile.

Answer: A

## QUESTION NO: 36


View the Exhibit and examine the result of Automatic SQL Tuning for a period. None of the recommended SQL profiles are implemented. What would you do to allow the Automatic SQL Tuning implement the SQL profiles automatically?

Begin Date **Jun 19, 2008 1:33:49 PM (UTC+07:00)** End Date **Jul 20, 2008 1:33:49 PM (UTC+07:00)**


### Overall Task Statistics

Executions **28** Candidate SQL **134** Distinct SQL Examined **20**

#### SQL Examined Status



#### Breakdown by Finding Type



### Profile Effect Statistics

#### Tuned SQL DB Time Benefit (seconds per week)

Implemented (sec) **0** Potential (sec) **9**



- A. Set the OPTIMIZER\_MODE parameter to ALL\_ROWS for the database instance.

**A Composite Solution With Just One Click – Certification Guaranteed**

## 1z0-054

- B. Set the SQLTUNE\_CATEGORY parameter to DEFAULT for the database instance.
- C. Set the OPTIMIZER\_CAPTURE\_SQL\_PLAN\_BASELINES parameter to TRUE for the database instance.
- D. Use the DBMS\_SQLTUNE.SET\_TUNING\_TASK\_PARAMETERS procedure to set ACCEPT\_SQL\_PROFILES to TRUE.

**Answer: D**

### QUESTION NO: 37

You receive a recommendation for a SQL statement through the automatic SQL tuning process that suggests implementing a SQL profile. You implement the SQL profile. The tables associated with the SQL statement grow, and indexes are created and dropped on these tables. What would happen to the SQL profile that is created for the SQL statement?

- A. It becomes outdated immediately after the tables change.
- B. It is maintained automatically along with the changes in the tables.
- C. The profile will be used for the changed plans of the SQL statement.
- D. The plan for the SQL statement is frozen until it is regenerated by the SQL Tuning Advisor.

**Answer: C**

### QUESTION NO: 38

Identify the type of recommendation that can automatically be implemented by the Automatic Tuning Advisor as part of automatic SQL tuning task?

- A. recommendation for SQL profiles
- B. recommendation about effective indexing
- C. recommendation about stale or no statistics
- D. recommendation about the structure of SQL statements

**Answer: A**

### QUESTION NO: 39

A user session executes a query and the plan for the query is shown in the Exhibit as Plan-1. After verifying with the SQL Access advisor, an index is created on the table in the JOB\_ID column and the query is executed again. A new plan is generated against the second query shown in the Exhibit as Plan-2. When explaining the plan for the second query, you observe that the optimizer uses the first plan instead of the second. The following parameters are set for the user session:

```
OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES=TRUE  
OPTIMIZER_USE_SQL_PLAN_BASELINES=TRUE  
SQLTUNE_CATEGORY=DEFAULT
```

**A Composite Solution With Just One Click – Certification Guaranteed**

# 1z0-054

Why does the optimizer not use the second plan?

## Plan-1

Plan hash value: 1445457117

| Id  | Operation         | Name      | Rows | Bytes | Cost (%CPU) | Time     |
|-----|-------------------|-----------|------|-------|-------------|----------|
| 0   | SELECT STATEMENT  |           | 6    | 372   | 3 (0)       | 00:00:01 |
| * 1 | TABLE ACCESS FULL | EMPLOYEES | 6    | 372   | 3 (0)       | 00:00:01 |

Predicate Information (identified by operation id):

1 - filter("JOB\_ID"='CLERK')

## Plan-2

Plan hash value: 1139333089

| Id  | Operation                   | Name      | Rows | Bytes | Cost (%CPU) | Time     |
|-----|-----------------------------|-----------|------|-------|-------------|----------|
| 0   | SELECT STATEMENT            |           | 6    | 372   | 2 (0)       | 00:00:01 |
| 1   | TABLE ACCESS BY INDEX ROWID | EMPLOYEES | 6    | 372   | 2 (0)       | 00:00:01 |
| * 2 | INDEX RANGE SCAN            | JBIND     | 6    |       | 1 (0)       | 00:00:01 |

Predicate Information (identified by operation id):

2 - access("JOB\_ID"='CLERK')

- A. because the second plan is not verified
- B. because the second plan is not part of the SQL plan baseline
- C. because the SQLTUNE\_CATEGORY parameter is set to DEFAULT
- D. because the explain plan option in the user session does not allow the new plan because SQLTUNE\_CATEGORY parameter is set to DEFAULT

**Answer: A**

**QUESTION NO: 40**

View the Exhibit and examine the plans in the SQL plan baseline. You observe that the last two SELECT statements (outlined in red and blue) are identical, but have two different plans. Also, the second plan (in blue) is not being used by the optimizer. Which two steps do you perform to ensure that the optimizer uses the second plan (in blue) if the plan is more efficient? (Choose two.)

**A Composite Solution With Just One Click – Certification Guaranteed**

## 1z0-054

Select All | Select None

| Select                   | Name                          | SQL Text  | Enabled | Accepted | Fixed | Auto Purge | Created                 | Last Modified           |
|--------------------------|-------------------------------|---|---------|----------|-------|------------|-------------------------|-------------------------|
| <input type="checkbox"/> | SYS_SQL_PLAN_ba5e158c52d2775d | DELETE FROM PLAN_TABLE WHERE STATEMENT_ID=:1          | YES     | YES      | NO    | YES        | Jul 20, 2008 5:12:07 PM | Jul 20, 2008 5:12:07 PM |
| <input type="checkbox"/> | SYS_SQL_PLAN_331fa51fd463620  | SELECT PLAN_TABLE_OUTPUT FROM TABLE(DBMS_XPLAN.DIS... | YES     | YES      | NO    | YES        | Jul 20, 2008 5:12:06 PM | Jul 20, 2008 5:12:06 PM |
| <input type="checkbox"/> | SYS_SQL_PLAN_184921b7cf314e9e | select * from employees where job_id='CLERK'          | YES     | YES      | NO    | YES        | Jul 20, 2008 5:15:31 PM | Jul 20, 2008 5:15:31 PM |
| <input type="checkbox"/> | SYS_SQL_PLAN_184921b70572d2e8 | select * from employees where job_id='CLERK'          | YES     | NO       | NO    | YES        | Jul 20, 2008 5:17:33 PM | Jul 20, 2008 5:17:33 PM |

**TIP** The table will display maximum of 2000 rows. Use search criteria to get the desired results.

- A. Mark the second plan (in blue) as fixed.
- B. Use SQL Tuning Advisor to tune the plan.
- C. Use SQL Access Advisor to tune the plan.
- D. Use the DBMS\_SPM.EVOLVE\_SQL\_PLAN\_BASELINE function to evolve the plan.

**Answer:** B, D

### QUESTION NO: 41

Identify the two direct sources from where SQL plans can be loaded into the SQL plan baselines. (Choose two.)

- A. Cursor cache
- B. Stored outline
- C. SQL Tuning Set
- D. Automatic Workload Repository (AWR) snapshots

**Answer:** A, C

### QUESTION NO: 42

Your system is upgraded from Oracle Database 10g to Oracle Database 11g. You imported SQL Tuning Sets (STS) from the previous version. After running the SQL Performance Analyzer for optimizer version change, you observe performance regression for a few SQL statements. What would you suggest to avoid performance regression for these SQL statements?

- A. Include the old plans in the SQL plan baseline.
- B. Use DBMS\_SPM.EVOLVE\_SQL\_PLAN\_BASELINE to evolve plans.
- C. Set the OPTIMIZER\_USE\_PENDING\_STATISTICS parameter to TRUE.
- D. Increase the value for the OPTIMIZER\_DYNAMIC\_SAMPLING parameter.

**Answer:** A

### QUESTION NO: 43

The OPTIMIZER\_CAPTURE\_SQL\_PLAN\_BASELINES parameter is set to TRUE for a 24x7 database instance. The space consumption has significantly increased by the automatically

**A Composite Solution With Just One Click – Certification Guaranteed**

## 1z0-054

captured plan baselines. You are suggested to run the following block:

```
BEGIN  
dbms_spm.configure('SPACE_BUDGET_PERCENT', 30);  
END;
```

After a few weeks of work, you observe warnings on exceeding space usage by SQL Management Base (SMB) in alert log file. Which two actions can be solutions for this? (Choose two.)

- A. increasing the AWR retention period
- B. purging the older SQL management objects
- C. increasing the size of the SYSTEM tablespace
- D. increasing the size of the SYSAUX tablespace

**Answer:** B, D

### QUESTION NO: 44

Your company wants to migrate a single instance Oracle 11g Release 1 database to a RAC environment. The database supports middle-tier applications using connection pooling for connecting to database.

Being a part of the performance management team, you want to test the performance of the SQL statements and the peak workload of the applications in the new environment before upgrading the production system.

Given below are some of the steps in random order to accomplish the task:

- 1) Capture the real workload on the existing system.
- 2) Create a SQL Tuning Set (STS) for the statements executed by the application.
- 3) Transport STS on the test environment and use SQL Tuning Advisor to compare the before and after performance.
- 4) Replay the workload on the test environment.
- 5) Analyze the divergence after replay.
- 6) Use Automatic Workload Repository (AWR) to capture the information for top SQL on the existing system.
- 7) Transport STS on the test environment and use SQL Performance Analyzer to compare the before and after performance.

Identify the correct steps required to perform the test.

- A. 2 and 3 only
- B. 6 and 7 only
- C. 1, 4, and 5 only
- D. 2, 3, and 5 only
- E. 2, 3, 4, and 5 only

**Answer:** C

### QUESTION NO: 45

**A Composite Solution With Just One Click – Certification Guaranteed**

## 1z0-054

You are working on an online transaction processing (OLTP) system. The middle-tier applications use connection pooling to connect to the database. Presently, you have a single-node database. The company plans to migrate the database to a RAC environment. Before you move to a RAC environment, you want to test the performance of the SQL statements and the peak workload on the new environment. To accomplish the Database Capture for replay, you identified the peak workload period on the existing system and started the Database Capture. Which client requests to the database can be captured as part of the workload capture? (Choose all that apply.)

- A. flashback query
- B. distributed transactions
- C. logging in and logging out of sessions
- D. all DDL statements having bind variables
- E. direct path load of data from external files

**Answer:** C, D

### QUESTION NO: 46

You are a DBA in ABC Corp. You are working on a DSS system.. The applications directly connect to the database to perform transactions. The application users also perform direct loads. The company wants to upgrade from Oracle Database 10g version 10.2.0.1 to Oracle Database 11g. You want to test the performance of the SQL statements in the new environment before upgrading the production database. To accomplish this, you set up the test database with Oracle Database 11g, with the same data and schemas as in production.

What would you recommend to accomplish the task?

- A. gathering all the SQL statements and using SQL Tuning Advisor on the test database for performance analysis
- B. capturing the workload on the production database, replaying it on the test machine, and comparing the performance
- C. creating a SQL Tuning Set by capturing the SQL workload on the production database and using SQL Performance Analyzer for performance analysis
- D. configuring AWR to capture the maximum number of SQL statements, transporting it to the test system, and using SQL Access Advisor for performance analysis

**Answer:** C

### QUESTION NO: 47

You are a DBA in ABC Corp. You are working on an online transaction processing (OLTP) system. The applications running on the database use connection pooling to connect to the database to perform transactions.

The company wants to upgrade the CPU and the memory for the machine on which the database is

**A Composite Solution With Just One Click – Certification Guaranteed**

## 1z0-054

running. To test the realistic workload on the new machine with increased CPU and memory, you identified the peak workload time for the database and captured the workload on the production system.

You set up the test machine with the upgraded CPU and memory and installed the database with the same version as production.

Which three actions would you perform to accomplish the realistic replay of the peak workload on the test machine? (Choose three.)

- A. Preprocess the captured workload for the specified target database version.
- B. Set up as many replay clients as there are actual clients to the production system.
- C. Set up all the replay clients on the same machine on which the test database is running.
- D. Resolve all external references such as database links, external tables, and directory objects.
- E. Restore the replay database to match the capture database at the start of the workload capture.

**Answer:** A, D, E

### QUESTION NO: 48

You are a DBA in ABC Corp. You are working on an online transaction processing (OLTP) system. The applications running on the database use connection pooling to connect to the database to perform transactions.

The company wants to upgrade the CPU and the memory for the machine on which the production database is running. To test the realistic workload on the new machine with increased CPU and memory, you identified the peak workload time for the database and captured the workload on the production system. You set up the test machine with the upgraded CPU and memory, and performed the following operations for replay:

- 1) You preprocessed the captured workload.
- 2) You restored the replay database to match the capture database at the start of the workload capture.
- 3) You resolved all external references.
- 4) You set up replay clients.

You have the following replay considerations:

- 1) load testing to be performed
- 2) all the sessions to connect immediately
- 3) user waits between issuing calls to be kept to a minimum

What replay options would you choose to accomplish the replay?

- A. SYNCHRONIZATION = FALSE, THINK\_TIME\_SCALE = 0, CONNECT\_TIME\_SCALE = 0 and THINK\_TIME\_AUTO\_CORRECT = FALSE
- B. SYNCHRONIZATION = TRUE, THINK\_TIME\_SCALE = 100, CONNECT\_TIME\_SCALE = 0 and THINK\_TIME\_AUTO\_CORRECT = FALSE
- C. SYNCHRONIZATION = FALSE, THINK\_TIME\_SCALE = 100, CONNECT\_TIME\_SCALE = 100 and THINK\_TIME\_AUTO\_CORRECT = TRUE

**A Composite Solution With Just One Click – Certification Guaranteed**