1Z0-054^{Q&As}

Oracle Database 11g: Performance Tuning

Pass Oracle 1Z0-054 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4lead.com/1Z0-054.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle
Official Exam Center

⚙ **Instant Download** After Purchase

⚙ **100% Money Back** Guarantee

⚙ **365 Days** Free Update

⚙ **800,000+** Satisfied Customers

QUESTION 1

You are managing an online transaction processing (OLTP) system. Many users complain about the slow response time. You investigated and found that in the Top 5 events, direct path read accounts for 73.2% of the total wait time.Which two are the possible reasons for the high direct path read time? (Choose two.)

A. The I/O subsystem is slow.

B. The SQL statements are using high CPU.

C. The SQL statement has a long parsing time.

D. The SQL statements are performing a large number of physical and logical reads.

Correct Answer: AD

QUESTION 2

Your database has the ACCTG service configured for an accounting application running on a middle tier. These service is used by the middle-tier applications to connect to the database by using connection pools. The application has three modules LEDGER, ACCOUNTS, and BILLING. You asked the developers to invoke the DBMS_APPLICATION_INFO.SET_MODULE procedure to set the module name and action in the code for the application. You enabled tracing at the service level by executing the following command:

SQL> exec DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE(\\'ACCTG\\');

You want to consolidate the trace output from the trace files created based on module. Which is the appropriate solution for this?

A. Using trace files as input for the tkprof utility to consolidate the tracing of sessions for a module

B. Using the trace file as input for the trcsess utility to consolidate and process the output for a module

C. Using the trcsess utility to consolidate all trace files into a single output file, which can then be processed by the tkprof utility

D. Using the trace file as input for the trcsess utility and creating output files that can directly be used by the users for performance or debugging purposes

Correct Answer: C

QUESTION 3

For which two systems would you recommend configuring DBWR I/O slaves instead of multiple database writer processes? (Choose two.)

A. for systems where you have a single CPU

B. for systems where asynchronous I/O is not available

C. for systems with multiple CPUs or multiple processor groups

D. for systems where asynchronous I/O is available and configured

E. for systems where throughput is an important performance consideration

Correct Answer: AB

QUESTION 4

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)
```

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  |           | 73966 | 2383K |   798   (1)| 00:00:10 |
|*  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?

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

Correct Answer: D

QUESTION 5

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

Correct Answer: D

Latest 1Z0-054 Dumps   1Z0-054 PDF Dumps   1Z0-054 Practice Test

To Read the Whole Q&As, please purchase the Complete Version from Our website.

Try our product !

100% Guaranteed Success
100% Money Back Guarantee
365 Days Free Update
Instant Download After Purchase
24x7 Customer Support
Average 99.9% Success Rate
More than 800,000 Satisfied Customers Worldwide
Multi-Platform capabilities - Windows, Mac, Android, iPhone, iPod, iPad, Kindle

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications.
You can view Vendor list of All Certification Exams offered:

https://www.pass4lead.com/allproducts

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket: