

1Z0-053 Dumps

Oracle Database 11g: Administration II

<https://www.certleader.com/1Z0-053-dumps.html>



NEW QUESTION 1

Which of the following NLS_SORT parameter values would result in case-insensitive and accent-insensitive binary sorts?

- A. NLS_SORT = BINARY
- B. NLS_SORT = BINARY_AI
- C. NLS_SORT = BINARY_CI
- D. NLS_SORT = BINARY_AI_CI
- E. Binary sorts are case insensitive and accent insensitive by default.

Answer: B

NEW QUESTION 2

Given two different character sets (A and B), which of the following must be true for A to be considered a strict superset of B? (Choose all that apply.)

- A. A must contain all of the characters defined in B.
- B. A must be Unicode.
- C. The encoded values in A must match the encoded values in B for all characters defined in B.
- D. A must be a multi-byte character set.
- E. The encoded values in A must match the encoded values in B for all numeric and alphabetic characters in B.

Answer: AC

NEW QUESTION 3

The NLS_TERRITORY parameter specifies the default conventions to be used for which of the following globalization elements? (Choose all that apply.)

- A. Date format
- B. Decimal character
- C. Group separator
- D. First day of the month
- E. None of the above
- F. All of the above

Answer: ABC

Explanation: Parameter type String

Syntax NLS_TERRITORY = territory Default value Operating system-dependent Modifiable ALTER SESSION

Range of values Any valid territory name Basic Yes

NLS_TERRITORY specifies the name of the territory whose conventions are to be followed for day and week numbering.

This parameter also establishes the default date format, the default decimal character and group separator, and the default ISO and local currency symbols.

For information on these settings, see "NLS_DATE_FORMAT", "NLS_NUMERIC_CHARACTERS", "NLS_CURRENCY", and "NLS_ISO_CURRENCY".

NEW QUESTION 4

Automatic data conversion will occur if which of the following happens?

- A. The client and server have different NLS_LANGUAGE settings.
- B. The client and server character sets are not the same, and the database character set is not a strict superset of the client character set.
- C. The client and server are in different time zones.
- D. The client requests automatic data conversion.
- E. The AUTO_CONVERT initialization parameter is set to TRUE.

Answer: B

NEW QUESTION 5

Case-insensitive sorts are always accent insensitive by default.

- A. True
- B. False

Answer: B

NEW QUESTION 6

Which NLS parameter directly governs linguistic searches?

- A. NLS_SEARCH_L
- B. NLS_SORT
- C. NLS_SEARCH
- D. NLS_SORT_L
- E. None of the above

Answer: B

NEW QUESTION 7

Which is not a valid locale definition file type?

- A. Language
- B. Linguistic sort
- C. Calendar
- D. Territory
- E. Character set

Answer: C

NEW QUESTION 8

What can you determine about the following linguistic sorts based only on their names?

- 1. GERMAN
- 2. FRENCH_M

- A. 1 is a monolingual sort.
- B. 2 is a monolingual sort.
- C. 1 is case insensitive.
- D. Both 1 and 2 are case insensitive.
- E. Case sensitivity is unknown.

Answer: A

NEW QUESTION 9

NLS parameters can be set using the five methods listed. Put the methods in order from highest to lowest according to Oracles order of precedence:

- a: Default setting
- b: Client environment variable
- c: Explicit ALTER SESSION statement
- d: Inside SQL function
- e: Server initialization parameter

- A. b, d, e, a, c
- B. e, a, b, c, d
- C. d, c, b, e, a
- D. a, b, d, c, e
- E. d, c, b, a, e

Answer: C

Explanation: *Table 3-1 Methods of Setting NLS Parameters and Their Priorities*

Priority	Method
1 (highest)	Explicitly set in SQL functions
2	Set by an ALTER SESSION statement
3	Set as an environment variable
4	Specified in the initialization parameter file
5	Default

C:\Users\albo\Desktop\1-1.jpg

NEW QUESTION 10

Which of the following datatypes store time-zone information in the database?

- A. TIMESTAMP
- B. DATE
- C. TIMESTAMP WITH TIME ZONE
- D. TIMESTAMP WITH LOCAL TIME ZONE
- E. DATETIME

Answer: C

NEW QUESTION 10

In a database with the database character set of US7ASCII and a national character set of UTF-8, which datatypes would be capable of storing Unicode data by default?

- A. VARCHAR2
- B. CHAR
- C. NVARCHAR2
- D. CLOB
- E. LONG

Answer: C

NEW QUESTION 14

Which of the following would be affected by setting NLS_LENGTH_SEMANTICS=CHAR?

- A. All objects in the database

- B. Tables owned by SYS and SYSTEM
- C. Data dictionary tables
- D. NCHAR columns
- E. CHAR columns

Answer: E

NEW QUESTION 17

The NLS_LANGUAGE parameter specifies the default conventions to be used for which of the following globalization elements?

- A. Languages for server messages
- B. Day and month names and abbreviations
- C. Symbols to represent a.m., p.m., AD, and BC
- D. Affirmative and negative response strings (YES, NO)
- E. None of the above
- F. All of the above

Answer: F

Explanation: Parameter type String

Syntax NLS_LANGUAGE = language

Default value Operating system-dependent, derived from the NLS_LANG environment variable

Modifiable ALTER SESSION

Range of values Any valid language name Basic Yes

NLS_LANGUAGE specifies the default language of the database. This language is used for messages, day and month names, symbols for AD, BC, a.m., and p.m., and the default sorting mechanism. This parameter also determines the default values of the parameters NLS_DATE_LANGUAGE and NLS_SORT.

NEW QUESTION 18

What is the name of the file that identifies the set of available locale definitions?

- A. locale.def
- B. lxdef.ora
- C. lx1boot.nlb
- D. lx1boot.ora
- E. lang.def

Answer: C

NEW QUESTION 22

Which of the following are valid program types for a lightweight job? (Choose all that apply.)

- A. PLSQL_BLOCK
- B. EXECUTABLE
- C. JAVA_STORED_PROCEDURE
- D. STORED_PROCEDURE
- E. EXTERNAL

Answer: AD

Explanation: Job_type

Job action type ('PLSQL_BLOCK', 'STORED_PROCEDURE', 'EXECUTABLE', or 'CHAIN')

NEW QUESTION 24

To set the history retention period for either window logging or job logging individually, which parameters of the SET_SCHEDULER_ATTRIBUTE procedure need to be used? (Choose all that apply.)

- A. LOG_HISTORY
- B. JOB_LOG_RETENTION
- C. WINDOW_LOG_RETENTION
- D. WHICH_LOG
- E. LOG_NAME

Answer: AD

Explanation: -----

Attributes list of SET_SCHEDULE_ATTRIBUTE:

'default_timezone': Repeating jobs and windows that use the calendaring syntax retrieve the time zone from this attribute when start_date is not specified. See "Calendaring Syntax" for more information.

'email_server': The SMTP server address that the Scheduler uses to send e-mail notifications for job state events. E-mail notifications cannot be sent if this attribute is NULL. 'email_sender': The default e-mail address of the sender of job state e-mail notifications. 'email_server_credential': The schema and name of an existing credential object that SYS has execute object privileges on. Default is NULL. The username and password stored in this credential are used to authenticate with the e-mail server when sending e-mail notifications. This functionality is available with Oracle Database 11g Release 2 (11.2.0.2).

'email_server_encryption': This attribute indicates whether or not encryption is enabled for this email server connection, and if so, at what point encryption starts, and with which protocol. This functionality is available starting with Oracle Database 11g Release 2 (11.2.0.2). Values are:

NONE: the default, indicating no encryption used

SSL_TLS: indicating that either SSL or TLS are used, from the beginning of the connection STARTTLS: indicating that the connection starts unencrypted, but the command STARTTLS is sent to the e-mail server and starts encryption

'event_expiry_time': The time, in seconds, before a job state event generated by the Scheduler expires from the Scheduler event queue. If NULL, job state events expire after 24 hours.

'log_history': The number of days that log entries for both the job log and the window log are retained. Default is 30 and the range of valid values is 0 through 1000000. 'max_job_slave_processes': This Scheduler attribute is not used.

The PURGE_LOG Procedure uses "WHICH_LOG" and "LOG_NAME" attributes. Syntax

```
DBMS_SCHEDULER.PURGE_LOG (
  log_history          IN PLS_INTEGER  DEFAULT 0,
  which_log           IN VARCHAR2     DEFAULT 'JOB_AND_WINDOW_LOG',
  job_name            IN VARCHAR2     DEFAULT NULL);
```

C:\Users\albo\Desktop\1-1.jpg Parameters

log_history

This specifies how much history (in days) to keep. The valid range is 0 - 1000000. If set to 0, no history is kept.

which_log

This specifies the log type. Valid values are: job_log, window_log, and job_and_window_log.

job_name

This specifies which job-specific entries must be purged from the job log. This can be a comma-delimited list of job names and job classes. Whenever job_name has a value other than NULL, the which_log argument implicitly includes the job log.

NEW QUESTION 28

Which three statements are true about persistent configuration? (Choose three.)

- A. A user cannot set privileges on the persistent lightweight jobs
- B. Persistent lightweight jobs generate a large amount of metadata
- C. It is possible to create fully self-contained persistent lightweight jobs
- D. The use of a template is mandatory to create persistent lightweight jobs
- E. Persistent lightweight jobs are useful when users need to create a large number of jobs in a short time

Answer: ADE

NEW QUESTION 30

View the Exhibits exhibit1 and exhibit2.

Both the processes use PROG_1 as the job template that is already available. The time taken by the jobs are recorded in the TEST_LOG table. While comparing the time taken to create the jobs, the process in exhibit1 takes less time than the process in exhibit2.

What is the reason for this? lightweight-job-1 (exhibit):

```
DECLARE
  newjob sys.job;
  newjobarr sys.job_array;
  newjobname VARCHAR2(30);
BEGIN
  newjobarr := sys.job_array();
  newjobarr.extend(10);
  FOR j IN 1..41 LOOP
    FOR i IN 1..10 LOOP
      newjob := sys.job(job_name => 'LWTJK' || to_char(i) || '_' || to_char(j),
                       job_style => 'LIGHTWEIGHT',
                       job_template => 'PROG_1',
                       enabled => TRUE
                      );
      newjobarr(i) := newjob;
    END LOOP;
  END LOOP;
  INSERT INTO TEST_LOG
  VALUES ('LWTJ', 'START', SYSTIMESTAMP);
  DBMS_SCHEDULER.CREATE_JOBS(newjobarr, 'TRANSACTIONAL');
  INSERT INTO TEST_LOG
  VALUES ('LWTJ', 'END', SYSTIMESTAMP);
END LOOP;
END;
```

lightweight-job-2 (exhibit):

```

DECLARE
newjob sys.job;
newjobarr sys.job_array;
newjobname VARCHAR2(30);
BEGIN
newjobarr := sys.job_array();
newjobarr.extend(10);
FOR j in 1..61 LOOP
FOR i IN 1..10 LOOP
newjob := sys.job(job_name => 'REGJK' || to_char(i) || '_' || to_char(j),
job_style => 'REGULAR',
job_template => 'PROG_1',
enabled => TRUE
);
newjobarr(i) := newjob;
END LOOP;
END LOOP;
INSERT INTO TEST_LOG
VALUES ('REGJOB', 'START', SYSTIMESTAMP);
DBMS_SCHEDULER.CREATE_JOBS(newjobarr, 'TRANSACTIONAL');
INSERT INTO TEST_LOG
VALUES ('REGJOB', 'END', SYSTIMESTAMP);
END LOOP;
END;

```

- A. It updates several tables in the SYSTEM tablespace instead of creating new tables.
- B. It creates jobs temporarily in memory only.
- C. It creates less metadata for the jobs.
- D. It writes the job metadata to disk in compressed format.

Answer: C

Explanation: The difference between two blocks of code is the "JOB_STYLE". The lightweight job creates less metadata for the jobs.

NEW QUESTION 33

Which of the following calendaring syntax expressions would evaluate to the last day of every month?

- A. FREQ = MONTHLY; BYMONTHDAY = 31
- B. FREQ = MONTHLY; BYMONTHDAY = -1
- C. FREQ = DAILY; BYDAY = -1
- D. FREQ = MONTHLY; BYDAY = 31
- E. FREQ = DAILY; BYMONTHDAY = LAST_DAY

Answer: B

Explanation: FREQ

This specifies the type of recurrence. It must be specified. The possible predefined frequency values are YEARLY, MONTHLY, WEEKLY, DAILY, HOURLY, MINUTELY, and SECONDLY. Alternatively, specifies an existing schedule to use as a user-defined frequency.

BYMONTHDAY

This specifies the day of the month as a number. Valid values are 1 to 31. An example is 10, which means the 10th day of the selected month. You can use the minus sign (-) to count backward from the last day, so, for example, BYMONTHDAY=-1 means the last day of the month and BYMONTHDAY=-2 means the next to last day of the month.

NEW QUESTION 35

Which of the following is true about job chains?

- A. They consist of one or more Scheduler programs.
- B. They are used to implement dependency scheduling.
- C. They are used to implement time-based scheduling.
- D. They are used to implement event-based scheduling.
- E. None of the above.

Answer: B

Explanation: Creating and Managing Job Chains

A job chain ("chain") is a named series of tasks that are linked together for a combined objective. Chains are the means by which you can implement dependency based scheduling, in which jobs are started depending on the outcomes of one or more previous jobs.

NEW QUESTION 38

You create two resource plans, one for data warehouse loading jobs at night and the other for application jobs at day time. You want the resource plans to activate automatically so that the resource allocation is optimum as desired by the activity. How would you achieve this?

- A. Implement job classes

- B. Implement Scheduler windows
- C. Implement the mapping rule for the consumer groups
- D. Set the SWITCH_TIME resource plan directive for both the resource plans

Answer: B

NEW QUESTION 39

Which of the following is not a valid calendaring syntax element?

- A. FREQ
- B. BYHOUR
- C. RUNDATE
- D. INTERVAL
- E. BYMINUTE

Answer: C

Explanation: [Calendaring Syntax \(link\)](#)

Table 129-10 Values for repeat_interval FREQ

This specifies the type of recurrence. It must be specified. The possible predefined frequency values are YEARLY, MONTHLY, WEEKLY, DAILY, HOURLY, MINUTELY, and SECONDLY. Alternatively, specifies an existing schedule to use as a user-defined frequency.

INTERVAL

This specifies a positive integer representing how often the recurrence repeats. The default is 1, which means every second for secondly, every day for daily, and so on. The maximum value is 99.

BYMONTH

This specifies which month or months you want the job to execute in. You can use numbers such as 1 for January and 3 for March, as well as three-letter abbreviations such as FEB for February and JUL for July.

BYWEEKNO

This specifies the week of the year as a number. It follows ISO-8601, which defines the week as starting with Monday and ending with Sunday; and the first week of a year as the first week, which is mostly within the Gregorian year. The first week is equivalent to the following two variants: the week that contains the first Thursday of the Gregorian year; and the week containing January 4th. The ISO-8601 week numbers are integers from 1 to 52 or 53; parts of week 1 may be in the previous calendar year; parts of week 52 may be in the following calendar year; and if a year has a week 53, parts of it must be in the following calendar year.

As an example, in the year 1998, the ISO week 1 began on Monday December 29th, 1997; and the last ISO week (week 53) ended on Sunday January 3rd, 1999. So December 29th, 1997, is in the ISO week 1998-01, and January 1st, 1999, is in the ISO week 1998-53. byweekno is only valid for YEARLY. Examples of invalid specifications are "FREQ=YEARLY; BYWEEKNO=1; BYMONTH=12" and "FREQ=YEARLY; BYWEEKNO=53; BYMONTH=1".

BYYEARDAY

This specifies the day of the year as a number. Valid values are 1 to 366. An example is 69, which is March 10 (31 for January, 28 for February, and 10 for March). 69 evaluates to March 10 for non-leap years and March 9 in leap years. -2 will always evaluate to December 30th independent of whether it is a leap year.

BYDATE

This specifies a list of dates, where each date is of the form [YYYY] MMDD. A list of consecutive dates can be generated by using the SPAN modifier, and a date can be adjusted with the OFFSET modifier. An example of a simple BYDATE clause follows: BYDATE=0115,0315,0615,0915,1215,20060115 The following SPAN example is equivalent to BYDATE=0110,0111,0112,0113,0114, which is a span of 5 days starting at 1/10: BYDATE=0110+SPAN:5D

The plus sign in front of the SPAN keyword indicates a span starting at the supplied date. The minus sign indicates a span ending at the supplied date, and the "^" sign indicates a span of n days or weeks centered around the supplied date. If n is an even number, it is adjusted up to the next odd number.

Offsets adjust the supplied date by adding or subtracting n days or weeks. BYDATE=0205- OFFSET: 2W is equivalent to BYDATE=0205-14D (the OFFSET: keyword is optional), which is also equivalent to BYDATE=0122.

BYMONTHDAY

This specifies the day of the month as a number. Valid values are 1 to 31. An example is 10, which means the 10th day of the selected month. You can use the minus sign (-) to count backward from the last day, so, for example, BYMONTHDAY=-1 means the last day of the month and BYMONTHDAY=-2 means the next to last day of the month.

BYDAY

This specifies the day of the week from Monday to Sunday in the form MON, TUE, and so on. Using numbers, you can specify the 26th Friday of the year, if using a YEARLY frequency, or the 4th THU of the month, using a MONTHLY frequency. Using the minus sign, you can say the second to last Friday of the month. For example, -1 FRI is the last Friday of the month.

BYHOUR

This specifies the hour on which the job is to run. Valid values are 0 to 23. As an example, 10 means 10 a.m.

BYMINUTE

This specifies the minute on which the job is to run. Valid values are 0 to 59. As an example, 45 means 45 minutes past the chosen hour.

BYSECOND

This specifies the second on which the job is to run. Valid values are 0 to 59. As an example, 30 means 30 seconds past the chosen minute.

BYSETPOS

This selects one or more items, by position, in the list of timestamps that result after the whole calendaring expression is evaluated. It is useful for requirements such as running a job on the last workday of the month. Rather than attempting to express this with the other BY clauses, you can code the calendaring expression to evaluate to a list of every workday of the month, and then add the BYSETPOS clause to select only the last item of that list. Assuming that workdays are Monday through Friday, the syntax would then be: FREQ=MONTHLY; BYDAY=MON,TUE,WED,THU,FRI; BYSETPOS=-1

Valid values are 1 through 9999. A negative number selects an item from the end of the list (-1 is the last item, -2 is the next to last item, and so on) and a positive number selects from the front of the list. The BYSETPOS clause is always evaluated last. BYSETPOS is only supported with the MONTHLY and YEARLY frequencies. The BYSETPOS clause is applied to the list of timestamps once per frequency period. For example, when the frequency is defined as MONTHLY, the Scheduler determines all valid timestamps for the month, orders that list, and then applies the BYSETPOS clause. The Scheduler then moves on to the next month and repeats the procedure. Assuming a start date of Jun 10, 2004, the example evaluates to: Jun 30, Jul 30, Aug 31, Sep 30, Oct 29, and so on.

INCLUDE This includes one or more named schedules in the calendaring expression. That is, the set of timestamps defined by each included named schedule is added to the results of the calendaring expression. If an identical timestamp is contributed by both an included schedule and the calendaring expression, it is included in the resulting set of timestamps only once. The named schedules must have been defined with the CREATE_SCHEDULE procedure.

EXCLUDE

This excludes one or more named schedules from the calendaring expression. That is, the set of timestamps defined by each excluded named schedule is removed from the results

of the calendaring expression. The named schedules must have been defined with the CREATE_SCHEDULE procedure.

INTERSECT

This specifies an intersection between the calendaring expression results and the set of timestamps defined by one or more named schedules. Only the timestamps that appear both in the calendaring expression and in one of the named schedules are included in the resulting set of timestamps. For example, assume that the named schedule last_sat indicates the last Saturday in every month, and that for the year 2005, the only months where the last day of the month is also a Saturday are April and December.

Assume also that the named schedule end_qtr indicates the last day of each quarter in 2005:

3/31/2005, 6/30/2005, 9/30/2005, 12/31/2005 These calendaring expressions result in the dates that follow:

3/31/2005, 4/30/2005, 6/30/2005, 9/30/2005, 12/31/2005

FREQ=MONTHLY; BYMONTHDAY=-1; INTERSECT=last_sat,end_qtr

In this example, the terms FREQ=MONTHLY; BYMONTHDAY=-1 indicate the last day of each month.

PERIODS

This identifies the number of periods that together form one cycle of a user-defined frequency. It is used in the repeat_interval expression of the schedule that defines the user-defined frequency. It is mandatory when the repeat_interval expression in the main schedule contains a BYPERIOD clause. The following example defines the quarters of a fiscal year.

FREQ=YEARLY; BYDATE=0301,0601,0901,1201; PERIODS=4 BYPERIOD

This selects periods from a user-defined frequency. For example, if a main schedule names a user-defined frequency schedule that defines the fiscal quarters shown in the previous example, the clause BYPERIOD=2,4 in the main schedule selects the 2nd and 4th fiscal quarters.

Topic 20, Globalization

NEW QUESTION 41

Which DBMS_SCHEDULER procedure(s) can be used to alter an existing job? (Choose all that apply.)

- A. SET_ATTRIBUTE_NULL
- B. ALTER_JOB
- C. ALTER_JOB_PARAMETERS
- D. ALTER
- E. SET_ATTRIBUTE

Answer: AE

Explanation: SET_ATTRIBUTE_NULL Procedure, Changes an attribute of an object to NULL SET_ATTRIBUTE Procedure, Changes an attribute of a job, schedule, or other Scheduler object

There is NO such "ALTER_JOB", "ALTER_JOB_PARAMETERS" and "ALTER" procedure available.

NEW QUESTION 42

Which three statements are true about windows? (Choose three.)

- A. Only one window can be open at any given time
- B. Consumer groups are associated with windows
- C. Windows work with job classes to control resource allocation
- D. The database service name must be provided during windows creation
- E. Windows can automatically start job or change resource allocation among jobs for various time periods.

Answer: ACE

NEW QUESTION 47

How many different calendars does Oracle 11g support?

- A. 22
- B. 7
- C. 6
- D. 15
- E. 2

Answer: B

NEW QUESTION 51

Which two statements are true with respect to the maintenance window? (Choose two.)

- A. A DBA can enable or disable an individual task in all maintenance windows.
- B. A DBA cannot change the duration of the maintenance window after it is created.
- C. In case of a long maintenance window, all Automated Maintenance Tasks are restarted every four hours.
- D. A DBA can control the percentage of the resource allocated to the Automated Maintenance Tasks in each window.

Answer: AD

NEW QUESTION 54

When a job exceeds the date specified in its END_DATE attribute, which of the following will happen? (Choose all that apply.)

- A. The job will be dropped automatically if the value of the AUTO_DROP attribute is TRUE.
- B. The job will only be disabled if the value of the AUTO_DROP attribute is FALSE.
- C. The STATE attribute of the job will be set to COMPLETED if the value of the AUTO_DROP attribute is FALSE.
- D. All objects referenced by the job will be dropped if the value of the AUTO_DROP attribute is TRUE and the value of the CASCADE attribute is TRUE.
- E. The STATE column of the job table will be set to COMPLETED for the job.

Answer: ABC

Explanation: auto_drop, If TRUE (the default), indicates that the job should be dropped once completed. end_date, This attribute specifies the date and time after which the job expires and is no longer run. After the end_date, if auto_drop is TRUE, the job is dropped. If auto_drop is FALSE, the job is disabled and the STATE of the job is set to COMPLETED.

If no value for end_date is specified, the job repeats forever unless max_runs or max_failures is set, in which case the job stops when either value is reached. The

value for end_date must be after the value for start_date. If it is not, an error is generated when the job is enabled.

NEW QUESTION 57

Which two statements are true about a job chain? (Choose two.)

- A. A job chain can contain a nested chain of jobs.
- B. The jobs in a job chain cannot have more than one dependency.
- C. A job of the CHAIN type can be run using event-based or time-based schedules.
- D. The jobs in a job chain can be executed only by using the events generated by the Scheduler

Answer: AC

Explanation: Defining Chain Steps

After creating a chain object, you define one or more chain steps. Each step can point to one of the following:

- ? A Scheduler program object (program)
- ? Another chain (a nested chain)
- ? An event schedule, inline event, or file watcher

You define a step that points to a program or nested chain by using the DEFINE_CHAIN_STEP procedure. An example is the following, which adds two steps to my_chain1:

```
BEGIN
  DBMS_SCHEDULER.DEFINE_CHAIN_STEP (
    chain_name      => 'my_chain1',
    step_name       => 'my_step1',
    program_name    => 'my_program1');
  DBMS_SCHEDULER.DEFINE_CHAIN_STEP (
    chain_name      => 'my_chain1',
    step_name       => 'my_step2',
    program_name    => 'my_chain2');
END;
```

C:\Users\albo\Desktop\1-1.jpg

NEW QUESTION 58

When setting arguments for a job, which procedure do you use for types that cannot be implicitly converted to and from a VARCHAR2 datatype?

- A. SET_JOB_ARGUMENT_VALUE
- B. SET_JOB_VALUE_ANYDATA
- C. SET_JOB_ANYDATA_VALUE
- D. SET_SPECIAL_JOB_VALUE
- E. SET_JOB_ANYTYPE_VALUE

Answer: C

Explanation: Topic 19, Administering the Scheduler

NEW QUESTION 60

Which of the following are not disabled by default?

- A. Jobs
- B. Chains
- C. Windows
- D. Window groups
- E. Schedule

Answer: E

NEW QUESTION 61

Which three elements can a job chain process involve? (Choose three)

- A. an event
- B. a schedule
- C. a program
- D. another chain
- E. a lightweight job

Answer: BDE

Explanation: Don't really understand the answer so far. Defining Chain Steps

After creating a chain object, you define one or more chain steps. Each step can point to one of the following:

- ? A Scheduler program object (program)
- ? Another chain (a nested chain)
- ? An event schedule, inline event, or file watcher

NEW QUESTION 63

A schedule defined entirely within the confines of a Scheduler job object is known as a (n) .

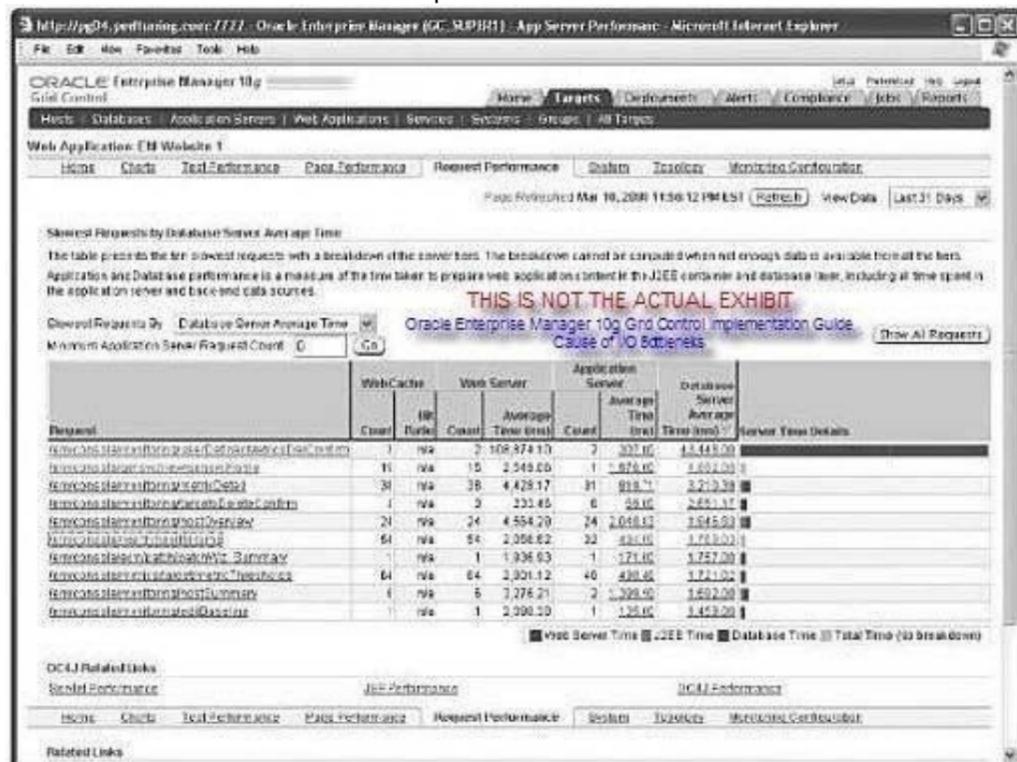
- A. Fixed schedule
- B. Inline schedule
- C. Stored schedule
- D. Hard-coded schedule
- E. None of the above

Answer: B

Explanation: "A schedule defined within a job object is know as an inline schedule, where as an independent schedule object is referred to as a stored schedule. Inline schedules cannot be reference by any other objects." http://www.datadisk.co.uk/html_docs/oracle/scheduler.htm

NEW QUESTION 64

View the exhibit and examine the output. Which statement can be an inference from the output? Exhibit:



- A. The FRA disk group has an asynchronous I/O bottleneck
- B. The least number of I/Os are performed on the last data file in the list
- C. The number of times that the backup or restore process directed the OS to wait until an I/O was complete is the highest for the last data file in the list
- D. The number of times that the backup or restore process made an OS call to poll for I/O completion in Nonblocking mode is the least for the FRA disk group

Answer: A

Explanation: Topic 18, Automating Tasks with the Scheduler

NEW QUESTION 67

What is the danger associated with stopping a running job by using the STOP_JOB procedure?

- A. The job will need to be reenabled before it will execute again.
- B. The job may hold locks on objects referenced within it.
- C. All jobs within the job group will also be stopped.
- D. The job may leave data in an inconsistent state.
- E. There is no danger in using the STOP_JOB procedure.

Answer: D

NEW QUESTION 71

What is the default value for the ENABLED attribute of a job or program when it is created?

- A. TRUE
- B. FALSE
- C. There is no default
- D. It must be defined at creation time.
- E. PENDING
- F. NULL

Answer: B

NEW QUESTION 73

Which statement is true about a Scheduler-generated event?

- A. It can be generated when a file arrives on the file system.
- B. it indicates state changes that occur within the Scheduler itself.
- C. it is raised by an application that is consumed by the Scheduler to start a job.
- D. it requires jobs to be defined with an event condition and a queue specification.

Answer: D

Explanation: Consuming Job State Events with your Application

To consume job state events, your application must subscribe to the Scheduler event queue

SYS.SCHEDULER\$_EVENT_QUEUE. This queue is a secure queue and is owned by SYS. To create a subscription to this queue for a user, do the following:

1. Log in to the database as the SYS user or as a user with the MANAGE ANY QUEUE privilege.
 2. Subscribe to the queue using a new or existing agent.
 3. Run the package procedure DBMS_AQADM.ENABLE_DB_ACCESS as follows: DBMS_AQADM.ENABLE_DB_ACCESS(agent_name, db_username); where agent_name references the agent that you used to subscribe to the events queue, and db_username is the user for whom you want to create a subscription.
- There is no need to grant dequeue privileges to the user. The dequeue privilege is granted on the Scheduler event queue to PUBLIC.

NEW QUESTION 74

For which two situations would you use functionality provided by the Resource Manager? (Choose two.)

- A. setting idle timeout limits on resource plans
- B. saving storage space by using compressed backup sets
- C. creating jobs that will run automatically at a scheduled time
- D. assigning priorities to jobs to manage access to system resources
- E. creating alerts to perform notification when tablespaces are low on available space resources

Answer: AD

NEW QUESTION 76

To control the execution of a server process when it is receiving bad packets from a potentially malicious client, you set the

SEC_PROTOCOL_ERROR_FURTHER_ACTION initialization parameter as follows:

```
SQL> ALTER SYSTEM SET SEC_PROTOCOL_ERROR_FURTHER_ACTION = Drop,10;
```

What is the significance of this setting?

- A. It terminates the client connection after 10 bad packets and the client cannot reconnect to the same instance.
- B. It terminates the client connection after 10 bad packets but the client can still reconnect, and attempt the same operation again.
- C. It terminates the client connection 10 seconds after receiving a bad packet and the client cannot reconnect to the same instance.
- D. It terminates the client connection after receiving a bad packet and the client can reconnect to the same instance after 10 minutes.

Answer: B

NEW QUESTION 79

Within a resource-plan definition, what differentiates a top-level plan from a subplan?

- A. A subplan has the PLAN_SUB parameter value set to SUB.
- B. A top-level plan has the GROUP_OR_PLAN parameter set to the name of the subplan in the resource-plan definition.
- C. There is no difference in the resource-plan definition.
- D. A subplan always has the CPU_MTH parameter value set to RATIO.
- E. The string TOP_LEVEL is appended to the name of top-level resource plans.

Answer: C

NEW QUESTION 83

You observed the following output for a user session:

```
SQL > SELECT sid, event, seconds_in_wait FROM v$session_wait WHERE sid = 18;

SID EVENT                                SECONDS_IN_WAIT
-----
18  statement suspended, wait error to be cleared 648
```

What do you infer from the preceding output?

- A. Resumable set for session with sid 18
- B. The user session has entered into a deadlock
- C. The database instance is enabled to use asynchronous commit
- D. The threshold warning limit is exceeded for the tablespace that is used by the user session

Answer: A

NEW QUESTION 84

Every resource plan must contain an allocation to which consumer group?

- A. LOW_GROUP
- B. SYS_GROUP
- C. DEFAULT_GROUP

- D. BASE_GROUP
- E. OTHER_GROUPS

Answer: E

NEW QUESTION 85

Which statement is true about a running session that belongs to the online transaction processing (OLTP) group?

```
BEGIN
DBMS_RESOURCE_MANAGER.CREATE_PLAN_DIRECTIVE (
  PLAN                => 'PRIUSERS',
  GROUP_OR_SUBPLAN    => 'OLTP',
  COMMENT              => 'OLTP GROUP',
  MGMT_P1              => 75,
  SWITCH_GROUP        => 'LOW_GROUP',
  SWITCH_IO_REQS      => 10000,
  SWITCH_IO_MEGABYTES => 2500,
  SWITCH_FOR_CALL     => TRUE);
END;
/
```

- A. It permanently switches to the low_group consumer group if the session exceeds 10,000 I/O requests or 2,500 MB of data transfer.
- B. It performs the first 10000 I/O requests or 2,500 MB of data transfer in the LOW-GROUP consumer group, and then switches to the original group.
- C. It switches to the LOW_GROUP consumer group if the session exceeds 10000 I/O requests or 2500 MB of data transfer and returns to the original group after the operation.
- D. It switches to the LOW_GROUP consumer group if the session exceeds 10000 I/O requests or 2500 MB of data transfer for queries, but not for data manipulation language (DML) operations.

Answer: C

Explanation: Refer to here.

Example 2

The following PL/SQL block creates a resource plan directive for the OLTP group that temporarily switches any session in that group to the LOW_GROUP consumer group if the session exceeds 10,000 I/O requests or exceeds 2,500 Megabytes of data transferred. The session is returned to its original group after the offending top call is complete.

```
BEGIN
  DBMS_RESOURCE_MANAGER.CREATE_PLAN_DIRECTIVE (
    PLAN                => 'DAYTIME',
    GROUP_OR_SUBPLAN    => 'OLTP',
    COMMENT              => 'OLTP group',
    MGMT_P1              => 75,
    SWITCH_GROUP        => 'LOW_GROUP',
    SWITCH_IO_REQS      => 10000,
    SWITCH_IO_MEGABYTES => 2500,
    SWITCH_FOR_CALL     => TRUE);
END;
/
```

C:\Users\albo\Desktop\1-1.jpg

NEW QUESTION 87

View the Exhibit and examine the resource consumption details for the current plan in use by the database instance. Which two statements are true based on the output? (Choose two.) Exhibit:

```
SELECT name, active_sessions, queue_length,
       consumed_cpu_time, cpu_waits, cpu_wait_time
FROM v$rsrc_consumer_group;
```

NAME	ACTIVE_SESSIONS	QUEUE_LENGTH	CONSUMED_CPU_TIME	CPU_WAITS	CPU_WAIT_TIME
OLTP_ORDER_ENTRY	1	0	29690	467	6709
OTHER_GROUPS	0	0	5982366	4089	60425
SYS_GROUP	1	0	2420704	914	19540
DSS_QUERIES	4	2	4594660	3004	55700

- A. An attempt to start a new session by the user belonging to DSS_QUERIES fails with an error
- B. A user belonging to DSS_QUERIES can log in to a new session but the session will be queued
- C. The CPU_WAIT_TIME column indicates the total time that sessions in the consumer group waited for the CPU due to resource management
- D. The CPU_WAIT_TIME column indicates the total time that sessions in the consumer group waited for the CPU due to resource management, I/O waits, and

latch or enqueue contention

Answer: BC

Explanation: V\$RSRC_CONSUMER_GROUP Use the V\$RSRC_CONSUMER_GROUP view to monitor resources consumed, including CPU, I/O, and parallel servers. It can also be used to monitor statistics related to CPU resource management, runaway query management, parallel statement queuing, and so on. All of the statistics are cumulative from the time when the plan was activated.

SELECT name, active_sessions, queue_length, consumed_cpu_time, cpu_waits, cpu_wait_time FROM v\$rsrc_consumer_group;

NAME	ACTIVE_SESSIONS	QUEUE_LENGTH	CONSUMED_CPU_TIME	CPU_WAITS
CPU_WAIT_TIME				
-----	-----	-----	-----	-----
OLTP_ORDER_ENTRY	1	0	29690	467
6709				
OTHER_GROUPS	0	0	5982366	4089
60425				
SYS_GROUP	1	0	2420704	914
19540				
DSS_QUERIES	4	2	4594660	3004
55700				

C:\Users\albo\Desktop\1-1.jpg

In the preceding query results, the DSS_QUERIES consumer group has four sessions in its active session pool and two more sessions queued for activation. A key measure in this view is CPU_WAIT_TIME. This indicates the total time that sessions in the consumer group waited for CPU because of resource management. Not included in this measure are waits due to latch or enqueue contention, I/O waits, and so on.

NEW QUESTION 92

Which of these represent the main components of Database Resource Manager? (Choose all that apply.)

- A. Resource consumer groups
- B. Resource plans
- C. Resource-plan groups
- D. Resource-plan directives
- E. All of the above

Answer: ABD

NEW QUESTION 97

Which of the following objects can be directly referenced by a window object? (Choose all that apply.)

- A. Schedule object
- B. Program object
- C. Job object
- D. Resource plan
- E. Resource consumer group

Answer: AD

NEW QUESTION 100

You need to configure fine-grained access control to external network resources from within your database. You create an access control list (ACL) using the DBMS_NETWORK_ACL_ADMIN package.

Which statement is true regarding the ACL created?

- A. It is a list of remote database links stored in the XML file that are available to the users of the database.
- B. It is a list of users and network privileges stored in the XML file according to which a group of users can connect to one or more hosts.
- C. It is a list of users and network privileges stored in the data dictionary according to which a group of users can connect to one or more hosts.
- D. It is the list of the host names or the IP addresses stored in the data dictionary that can connect to your database through PL/SQL network utility packages such as UTL_TCP.

Answer: B

Explanation: The DBMS_NETWORK_ACL_ADMIN package provides the interface to administer the network Access Control List (ACL).

Refer to here for About Fine-Grained Access to External Network Services Configuring fine-grained access control for users and roles that need to access external network services from the database. This way, specific groups of users can connect to one or more host computers, based on privileges that you grant them.

Typically, you use this feature to control access to applications that run on specific host addresses.

To configure fine-grained access control to external network services, you create an access control list (ACL), which is stored in Oracle XML DB. You can create the access control list by using Oracle XML DB itself, or by using the DBMS_NETWORK_ACL_ADMIN and DBMS_NETWORK_ACL_UTILITY PL/SQL packages.

This guide explains how to use these packages to create and manage the access control list. To create an access control list by using Oracle XML DB and for general conceptual information about access control lists, see Oracle XML DB Developer's Guide.

NEW QUESTION 103

The database users regularly complain about the difficulty in performing transactions. On investigation, you find that some users perform long-running transactions that consume huge amounts of space in the undo tablespace, which caused the problem. You want to control the usage of the undo tablespace only for these user sessions and you do not want these sessions to perform long-running operations.

How would you achieve this?

- A. Implement a profile for the users.
- B. Implement external roles for the users.

- C. Set the threshold for the undo tablespace.
- D. Implement a Database Resource Manager plan.

Answer: B

NEW QUESTION 105

View the Exhibit and examine the steps that you executed to create a database resource plan.

Subsequently, you execute the following procedure which results in an error: SQL> EXECUTE dbms_resources_manager.validate_pending_area ();
What could be the reason?

```
SQL> EXECUTE dbms_resource_manager.create_pending_area();

PL/SQL procedure successfully completed.

SQL> EXECUTE dbms_resource_manager.create_consumer_group(consumer_group => 'OLTP',-
> comment => 'Online users');

PL/SQL procedure successfully completed.

SQL> EXECUTE dbms_resource_manager.create_plan (plan => 'PRIUSERS',-
> comment => 'DSS/Batch priority, ...' );

PL/SQL procedure successfully completed.

SQL> EXECUTE dbms_resource_manager.create_plan_directive (plan => 'PRIUSERS',-
> group_or_subplan => 'OLTP',comment => 'Online Group',CPU_P1 => 60);

PL/SQL procedure successfully completed.
```

- A. The pending area is automatically submitted after the plan creation
- B. The procedure must be executed before creating the resources plan directive
- C. The SYS_GROUP resource consumer group is not included in the resource plan directive
- D. The OTHER_GROUPS resources consumer group is not included in the resource plan directive

Answer: D

NEW QUESTION 106

The Database Resource Manager is automatically enabled in the maintenance window that runs the Automated Maintenance Task. What is the reason for this?

- A. to prevent the creation of an excessive number of scheduler job classes
- B. to allow the Automated Maintenance Tasks to use system resources without any restriction
- C. to allow resource sharing only among the Automated Maintenance Tasks in the maintenance window
- D. to prevent the Automated Maintenance Tasks from consuming excessive amounts of system resources

Answer: D

Explanation: Oracle Database Resource Manager (the Resource Manager) is an infrastructure that provides granular control of database resources allocated to users, applications, and services. The Resource Manager solves many resource allocation problems that an operating system does not manage well, including:

- ? Excessive overhead
- ? Inefficient scheduling
- ? Inappropriate allocation of resources
- ? Inability to manage database-specific resources

The Resource Manager helps overcome these problems by giving the database more control over allocation of hardware resources and enabling you to prioritize work within the database. You can classify sessions into groups based on session attributes, and then allocate resources to these groups to optimize hardware utilization.

NEW QUESTION 111

You plan to control the sessions performing a huge number of I/O operations. Your requirement is to kill the session when it exceeds a specified number of I/Os. Which statement describes a solution to the above?

- A. Set a threshold for the default system-defined moving window baseline.
- B. Add directives to the Automatic Database Diagnostic Monitor (ADDM).
- C. Modify the profile for the targeted users for which control needs to be imposed.
- D. Implement the database resource manager to add the SWITCH_IO_REQS and SWITCH_GROUP directives.

Answer: D

NEW QUESTION 114

View the exhibit and examine the TRANS table's storage information. After a massive delete operation, you executed the following statement to shrink the TRANS table:

```
SQL> ALTER TABLE trans SHRINK SPACE CASCADE;
```

Which statement describes the outcome of the command? Exhibit:

Select	Name	Type	Extent Management	Segment Management	Status	Size (MB)	Used (MB)
<input checked="" type="radio"/>	SYSAUX	PERMANENT	LOCAL	AUTO	ONLINE	330.000	322.000
<input checked="" type="radio"/>	SYSTEM	PERMANENT	LOCAL	MANUAL	ONLINE	470.000	462.625
<input checked="" type="radio"/>	TEMP	TEMPORARY	LOCAL	MANUAL	ONLINE	20.000	6.000
<input checked="" type="radio"/>	TT	PERMANENT	LOCAL	MANUAL	ONLINE	0.102	.102
<input checked="" type="radio"/>	UNDOTBS1	UNDO	LOCAL	MANUAL	ONLINE	90.000	15.813
<input checked="" type="radio"/>	TRANS	PERMANENT	LOCAL	AUTO	ONLINE	127.500	1.625

- A. An error is produced.
- B. The table and all related objects are compacted and the position of the high-water mark (HWM) for the table is adjusted
- C. The table and related indexes are compacted but the position of the high-water mark (HWM) for the table remains unchanged
- D. The unused space in the table is reclaimed and returned to the tablespace and the data manipulation language (DML) triggers on the table are fired during the shrinking process

Answer: B

NEW QUESTION 118

View the Exhibit for some of the current parameter settings. A user logs in to the HR schema and issues the following commands:

SQL> CREATE TABLE emp (empno NUMBER(3), ename VARCHAR2(20), sal NUMBER(8,2));

SQL> INSERT INTO emp(empno,ename) VALUES(1,'JAMES');

At this moment, a second user also logs in to the HR schema and issues the following command:

SQL> ALTER TABLE emp MODIFY sal NUMBER(10,2);

What happens in the above scenario? Exhibit:

NAME	TYPE	VALUE
db_file_multiblock_read_count	integer	107
ddl_lock_timeout	integer	60
distributed_lock_timeout	integer	60
dml_locks	integer	748
lock_sga	boolean	FALSE
enable_ddl_logging	boolean	FALSE
resumable_timeout	integer	0

- A. The second user's session immediately produces the resource busy error.
- B. The second user's command executes successfully.
- C. The second user's session waits for a time period before producing the resource busy error.
- D. A deadlock is created.

Answer: C

NEW QUESTION 123

In which cases is reference partitioning effective in enhancing performance?

- A. It is effective only in partition pruning.
- B. It is effective only in partitionwise joins provided that the query predicates are different from the partitioning key.
- C. It is effective in both partition pruning as well as partitionwise joins provided that the query predicates are identical to the partitioning key.
- D. It is effective in both partition pruning as well as partitionwise joins irrespective of whether the query predicates are different from or identical to the partitioning key.

Answer: D

NEW QUESTION 127

You notice that a long-running transaction is suspended due to a space constraint, and there is no AFTER SUSPEND triggered event addressing the issue. You also note that the critical transaction is just about to reach the RESUMABLE_TIMEOUT value.

Which of these actions is appropriate?

- A. Abort the session, fix the space problem, then resubmit the transaction.
- B. Use the DBMS_RESUMABLE.SET_SESSION_TIMEOUT procedure to extend the time-out for the session while you fix the problem.
- C. Do nothing, let the transaction fail, then fix the problem.
- D. Use Segment Shrink to clean up the table.
- E. Use the DBMS_RESUMABLE.SET_TIMEOUT procedure to extend the time-out for the session while you fix the problem.

Answer: B

Explanation: Topic 17, Managing Resources

NEW QUESTION 128

View the Exhibit to examine the parameters set for your database instance.

You execute the following command to perform I/O calibration after the declaration of bind variables in the session that are used in the command:

```
SQL> EXECUTE dbms_resource_manager.calibrate_io( num_physical_disks=>1,
max_latency=>50, max_iops=>:max_iops, max_mbps=>:max_mbps, actual_latency=>:actual_latency);
```

Which statement describes the consequence? Exhibit:

NAME	TYPE	VALUE
filesystemio_options	string	ASYNCH
backup_tape_io_slaves	boolean	FALSE
dbwr_io_slaves	integer	0
disk_asynch_io	boolean	TRUE
tape_asynch_io	boolean	TRUE
optimizer_use_pending_statistics	boolean	FALSE
statistics_level	string	TYPICAL
timed_os_statistics	integer	0
timed_statistics	boolean	FALSE
aq_tm_processes	integer	0
db_writer_processes	integer	1
gcs_server_processes	integer	0
global_txn_processes	integer	1
job_queue_processes	integer	1000
log_archive_max_processes	integer	4
processes	integer	150

- A. The command produces an error.
- B. The calibration process runs successfully and populates all the bind variables.
- C. The calibration process runs successfully but the latency time is not computed.
- D. The calibration process runs successfully but only the latency time is computed.

Answer: A

Explanation: Requisition of Calibrate I/O (link)

Before running I/O calibration, ensure that the following requirements are met:

? The user must be granted the SYSDBA privilege

? timed_statistics must be set to TRUE

? Asynchronous I/O must be enabled

? When using file systems, asynchronous I/O can be enabled by setting the FILESYSTEMIO_OPTIONS initialization parameter to SETALL.

? Ensure that asynchronous I/O is enabled for data files by running the following query:

```
COL NAME FORMAT A50
```

```
SELECT NAME,ASYNCH_IO FROM V$DATAFILE F,V$IOSTAT_FILE I WHERE F.FILE#=I.FILE_NO
AND FILETYPE_NAME='Data File';
```

Additionally, only one calibration can be performed on a database instance at a time.

NEW QUESTION 130

You plan to control idle sessions that are blocking other sessions from performing transactions. Your requirement is to automatically terminate these blocking sessions when they remain idle for a specified amount of time.

How would you accomplish this task?

- A. Set metric threshold
- B. Implement Database Resource Manager
- C. Enable resumable timeout for user sessions
- D. Add directives to Automatic Database Diagnostic Monitor (ADDM)

Answer: B

NEW QUESTION 131

Evaluate the following block of code:

```
BEGIN
  DBMS_NETWORK_ACL_ADMIN.CREATE_ACL (
    acl =>'mycompany-com-permissions.xml',
    principal => 'ACCT_MGR',
    is_grant => TRUE,
    privilege => 'connect');
  DBMS_NETWORK_ACL_ADMIN.ASSIGN_ACL (
    acl => 'mycompany-com-permissions.xml',
    host => '*.mycompany.com');
END;
```

What is the outcome of the above code?

- A. It produces an error because a fully qualified host name needs to be specified.
- B. It produces an error because the range of ports associated with the hosts has not been specified.
- C. It creates an access control list (ACL) with the user ACCT_MGR who gets the CONNECT and RESOLVE privileges.
- D. It creates an access control list (ACL) with the user ACCT_MGR who gets the CONNECT privilege but not the RESOLVE privilege.

Answer: C

NEW QUESTION 136

View the Exhibit to examine the error obtained during the I/O calibration process. There are no data files on raw devices. What is the reason for this error?

```
SQL> EXECUTE dbms_resource_manager.calibrate_io( -
> num_physical_disks=>1, -
> max_latency=>50, -
> max_iops=>:max_iops, -
> max_mbps=>:max_mbps, -
> actual_latency=>:actual_latency);
BEGIN dbms_resource_manager.calibrate_io( num_physical_disks=>1, max_latency=>50, max_iops=>:max_iops,
max_mbps=>:max_mbps, actual_latency=>:actual_latency); END;
*
ERROR at line 1:
ORA-56708: Could not find any datafiles with asynchronous i/o capability
ORA-06512: at "SYS.DBMS_RMIN", line 453
ORA-06512: at "SYS.DBMS_RESOURCE_MANAGER", line 1153
ORA-06512: at line 1
```

- A. The DISK_ASYNCH_IO parameter is set to TRUE.
- B. The FILESYSTEMIO_OPTIONS parameter is set to NONE.
- C. Another session runs the I/O calibration process concurrently.
- D. The pending area has not been created before running the I/O calibration process.

Answer: B

Explanation: 9.1.1.2 FILESYSTEMIO_OPTIONS Initialization Parameter

You can use the FILESYSTEMIO_OPTIONS initialization parameter to enable or disable asynchronous I/O or direct I/O on file system files. This parameter is platform-specific and has a default value that is best for a particular platform.

FILESYSTEMIO_OPTIONS can be set to one of the following values:

ASYNCH: enable asynchronous I/O on file system files, which has no timing requirement for transmission.

DIRECTIO: enable direct I/O on file system files, which bypasses the buffer cache. SETALL: enable both asynchronous and direct I/O on file system files.

NONE: disable both asynchronous and direct I/O on file system files.

To enable asynch I/O, set two values in the init.ora file (or spfile) and recycle the instances:

disk_asynch_io = true filesystemio_options = asynch

The first one is default; so you may not have to change it. The default of the second one is NONE; so you will probably have to change it. After setting these values and recycling the instance,

NEW QUESTION 141

Which tasks can be accomplished using the DBMS_LOB.SETOPTIONS procedure?

- A. only encryption and compression settings for all SecureFile LOBs
- B. only encryption and deduplication settings for only SecureFile CLOBs
- C. deduplication, encryption, and compression settings for all SecureFile LOBs
- D. deduplication, encryption, and compression settings only for SecureFile CLOBs

Answer: C

Explanation: Refer to here.

DBMS_LOB.SETOPTIONS()

This procedure sets compression, deduplication and encryption features. It enables the features to be set on a per-LOB basis, overriding the default LOB settings.

This call incurs a round trip to the server to make the changes persistent.

NEW QUESTION 143

Which statement describes the information returned by the DBMS_SPACE.SPACE_USAGE procedure for LOB space usage?

- A. It returns space usage of only BasicFile LOB chunks.
- B. It returns space usage of only SecureFile LOB chunks.
- C. It returns both BasicFile and SecureFile LOB space usage for only nonpartitioned tables.
- D. It returns both BasicFile and SecureFile LOB space usage for both partitioned and nonpartitioned tables.

Answer: B

Explanation: SPACE_USAGE Procedures

The first form of the procedure shows the space usage of data blocks under the segment High Water Mark. You can calculate usage for LOBs, LOB PARTITIONS and LOB SUBPARTITIONS. This procedure can only be used on tablespaces that are created with auto segment space management. The bitmap blocks, segment header, and extent map blocks are not accounted for by this procedure. Note that this overload cannot be used on SECUREFILE LOBs.

```
DBMS_SPACE.SPACE_USAGE (
  segment_owner      IN  VARCHAR2,
  segment_name       IN  VARCHAR2,
  segment_type       IN  VARCHAR2,
  unformatted_blocks OUT NUMBER,
  unformatted_bytes  OUT NUMBER,
  fs1_blocks         OUT NUMBER,
  fs1_bytes          OUT NUMBER,
  fs2_blocks         OUT NUMBER,
  fs2_bytes          OUT NUMBER,
  fs3_blocks         OUT NUMBER,
  fs3_bytes          OUT NUMBER,
  fs4_blocks         OUT NUMBER,
  fs4_bytes          OUT NUMBER,
  full_blocks        OUT NUMBER,
  full_bytes         OUT NUMBER,
  partition_name     IN  VARCHAR2 DEFAULT NULL);
```

C:\Users\albo\Desktop\1-1.jpg

The second form of the procedure returns information about SECUREFILE LOB space usage. It will return the amount of space in blocks being used by all the SECUREFILE LOBs in the LOB segment. The procedure displays the space actively used by the LOB column, freed space that has retention expired, and freed space that has retention unexpired. Note that this overload can be used only on SECUREFILE LOBs.

```
DBMS_SPACE.SPACE_USAGE (
  segment_owner      IN  VARCHAR2,
  segment_name       IN  VARCHAR2,
  segment_type       IN  VARCHAR2,
  segment_size_blocks OUT NUMBER,
  segment_size_bytes  OUT NUMBER,
  used_blocks        OUT NUMBER,
  used_bytes         OUT NUMBER,
  expired_blocks     OUT NUMBER,
  expired_bytes      OUT NUMBER,
  unexpired_blocks   OUT NUMBER,
  unexpired_bytes    OUT NUMBER,
  partition_name     IN  VARCHAR2 DEFAULT NULL);
```

C:\Users\albo\Desktop\1-1.jpg

NEW QUESTION 146

You issued the following command on the temporary tablespace LMTEMP in your database:
SQL>ALTER TABLESPACE ltemp SHRINK SPACE KEEP 20M;
Which requirement must be fulfilled for this command to succeed?

- A. The tablespace must be locally managed.
- B. The tablespace must have only one temp file.
- C. The tablespace must be made nondefault and offline.
- D. The tablespace can remain as the default but must have no active sort operations.

Answer: A

NEW QUESTION 151

Evaluate the following command and its output: SQL>SELECT * FROM dba_temp_free_space;
TABLESPACE_NAME TABLESPACE_SIZE ALLOCATED_SPACE FREE_SPACE
===== LMTEMP 250609664 101048576 149561088
Which two statements correctly interpret the output? (Choose two.)

- A. FREE_SPACE indicates only the space that is currently unallocated.
- B. ALLOCATED_SPACE indicates only the space currently allocated and in use.
- C. FREE_SPACE indicates only the space that is currently allocated and available for reuse.
- D. ALLOCATED_SPACE indicates both the space currently allocated and used, and the space that is available for reuse.
- E. FREE_SPACE indicates both the space that is currently allocated and available for reuse, and the space that is currently unallocated.

Answer: DE

NEW QUESTION 154

For which of the following can you use Segment Shrink? (Choose all that apply.)

- A. Heap tables
- B. Tables with function-based indexes
- C. Indexes
- D. Partitions and subpartitions
- E. None of the above

Answer: ACD

NEW QUESTION 159

When shrinking a table segment, you choose to shrink all the indexes for that table using the SHRINK SPACE command. Which clause should you use?

- A. INCLUDING DEPENDENCIES
- B. INCLUDING DEPENDENCIES CASCADE
- C. COMPACT
- D. CASCADE
- E. None of the above

Answer: D

NEW QUESTION 163

ENCRYPT_TS is an encrypted tablespace that contains tables with data.

Which statement is true regarding the effect of queries and data manipulation language (DML) statements on the encrypted data in the tables?

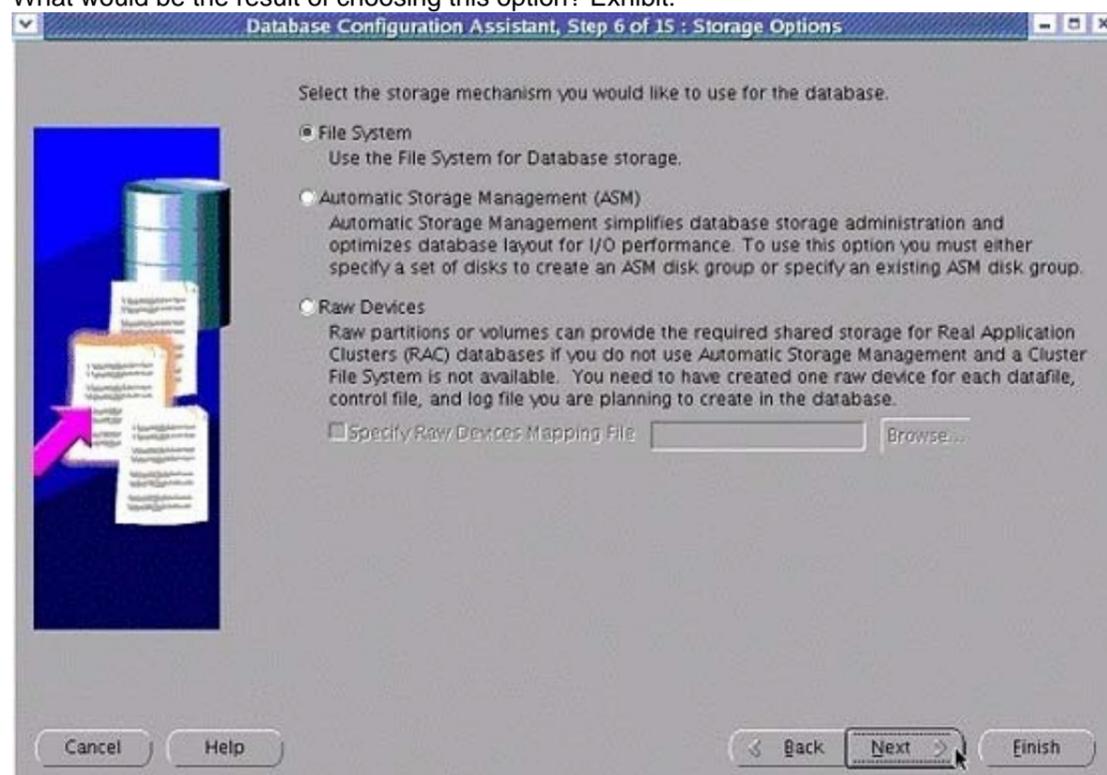
- A. The data is decrypted during SORT and JOIN operations.
- B. The data remains encrypted when it is read into memory.
- C. The data remains encrypted when it is stored in the redo logs.
- D. The data remains encrypted in the UNDO tablespace provided that the UNDO tablespace was created with the encryption option enabled.

Answer: C

NEW QUESTION 164

View the Exhibit.

You are creating a database by using Database Configuration Assistant (DBCA). You have chosen the File System option as the storage mechanism. What would be the result of choosing this option? Exhibit:



- A. Disk mirroring and striping would be done automatically
- B. The database files would be managed by the operating system's file system
- C. DBCA would not save the database files by using Optimal File Architecture (OFA)
- D. The data files are automatically spread across all available storage devices to optimize performance and resource utilization

Answer: B

NEW QUESTION 167

You decided to use Direct NFS configuration in a non-RAC Oracle installation and created the oranfstab file in / etc. Which two statements are true regarding this oranfstab file? (Choose two.)

- A. Its entries are specific to a single database.
- B. It contains file systems that have been mounted by Direct NFS.
- C. It is globally available to all Oracle 11g databases on the machine.
- D. It contains file systems that have been mounted by the kernel NFS system.

Answer: CD

NEW QUESTION 172

Which two statements are true about the compressed backups in RMAN? (Choose two.)

- A. Compressed backups can only be taken on the tape drives.
- B. The binary compression creates some performance overhead during backup operation.
- C. The ZLIB compression algorithm can be used only if the COMPATIBLE initialization parameter is set to 11.1.0.

D. The media manager compression for the tape drive should be enabled for taking compressed backups on the tape.

Answer: BD

NEW QUESTION 177

Which statements are true regarding table compression? (Choose all that apply.)

- A. It saves disk space and reduces memory usage.
- B. It saves disk space but has no effect on memory usage.
- C. It incurs extra CPU overhead during DML as well as direct loading operations.
- D. It incurs extra CPU overhead during DML but not direct loading operations.
- E. It requires uncompress operation during I/O.

Answer: AC

NEW QUESTION 180

What is the proper procedure to recover a lost tempfile?

- A. Restore the backup copy of the tempfile from the backup media.
- B. Re-create the tempfile with the create tempfile command.
- C. Copy an existing tempfile from another database.
- D. Re-create the tempfile with the create tablespace command.
- E. Re-create the tempfile with the alter tablespace command.

Answer: E

Explanation: ALTER TABLESPACE <tablespace_name> ADD TEMPFILE <new file name>; ALTER TABLESPACE <tablespace_name> TEMPFILE <old_file_name> OFFLINE; ALTER TABLESPACE <tablespace_name> TEMPFILE <> ONLINE;

NEW QUESTION 183

You need to create a partitioned table to store historical data and you issued the following command:

```
CREATE TABLE purchase_interval
PARTITION BY RANGE (time_id) INTERVAL (NUMTOYMINTERVAL(1,'month'))
STORE IN (tbs1,tbs2,tbs3) (
PARTITION p1 VALUES LESS THAN(TO_DATE('1-1-2005', 'dd-mm-yyyy')), PARTITION p2 VALUES LESS THAN(TO_DATE('1-1- 2007', 'dd-mm-yyyy'))
) AS
SELECT * FROM purchases
WHERE time_id < TO_DATE('1-1-2007','dd-mm-yyyy');
```

What is the outcome of the above command?

- A. It returns an error because the range partitions P1 and P2 should be of the same range.
- B. It creates two range partitions (P1, P2). Within each range partition, it creates monthwise subpartitions.
- C. It creates two range partitions of varying range.
- D. For data beyond '1-1-2007,' it creates partitions with a width of one month each.
- E. It returns an error because the number of tablespaces (TBS1,TBS2,TBS3) specified does not match the number of range partitions (P1,P2) specified.

Answer: C

NEW QUESTION 186

Following is the list of locations in random order where orafstab can be placed.

1./etc/mtab 2.\$ORACLE_HOME/dbs/orafstab 3./etc/orafstab

What is the sequence in which Direct NFS will search the locations?

- A. 1, 2, 3
- B. 3, 2, 1
- C. 2, 3, 1
- D. 1, 3, 2

Answer: C

NEW QUESTION 188

Which statements are true regarding system-partitioned tables? (Choose all that apply.)

- A. Only a single partitioning key column can be specified.
- B. All DML statements must use partition-extended syntax.
- C. The same physical attributes must be specified for each partition.
- D. Unique local indexes cannot be created on a system-partitioned table.
- E. Traditional partition pruning and partitionwise joins are not supported on these tables.

Answer: DE

NEW QUESTION 192

Which dependent object will get invalidated even if it is not affected by the table redefinition?

- A. packages
- B. triggers
- C. synonyms

D. views

Answer: B

Explanation: Results of the Redefinition Process

The following are the end results of the redefinition process:

The original table is redefined with the columns, indexes, constraints, grants, triggers, and statistics of the interim table.

Dependent objects that were registered, either explicitly using REGISTER_DEPENDENT_OBJECT or implicitly using COPY_TABLE_DEPENDENTS, are renamed automatically so that dependent object names on the redefined table are the same as before redefinition.

Note:

If no registration is done or no automatic copying is done, then you must manually rename the dependent objects.

The referential constraints involving the interim table now involve the redefined table and are enabled.

? Any indexes, triggers, materialized view logs, grants, and constraints defined on the original table (prior to redefinition) are transferred to the interim table and are dropped when the user drops the interim table. Any referential constraints involving the original table before the redefinition now involve the interim table and are disabled.

? Some PL/SQL objects, views, synonyms, and other table-dependent objects may become invalidated. Only those objects that depend on elements of the table that were changed are invalidated. For example, if a PL/SQL procedure queries only columns of the redefined table that were unchanged by the redefinition, the procedure remains valid. See "Managing Object Dependencies" for more information about schema object dependencies.

NEW QUESTION 194

You want to disable resumable space allocation for all sessions. Which value should be assigned to the RESUMABLE_TIMEOUT parameter to disable resumable space allocation for all sessions?

- A. 10
- B. 100
- C. NULL

Answer: A

NEW QUESTION 199

Which of these components correctly identify the unique value of the NAME column in the DBA_RESUMABLE view?

- A. Username, instance number, session ID
- B. Instance number, username, session ID
- C. Instance number, session ID, username
- D. Username, session ID, instance number
- E. None of the above

Answer: D

NEW QUESTION 200

You have three temporary tablespace groups named G1, G2, and G3 in your database. You are creating a new temporary tablespace as follows:

```
CREATE TEMPORARY TABLESPACE TEMP1 TEMPFILE 'u1/data/temp1.dbf' SIZE 10M TABLESPACE GROUP '';
```

Which statement regarding the above command is correct?

- A. It will create the tablespace TEMP1 in group G1.
- B. It will create the tablespace TEMP1 in group G3.
- C. It will not add the tablespace TEMP1 to any group.
- D. It will create the tablespace TEMP1 in the default group.

Answer: C

NEW QUESTION 204

The ADMIN_EMP table has columns EMPNO, ENAME, DEPTNO, and SAL. It has a materialized view EMP_MV with a materialized log and an ENAME_IDX index on the ENAME column. You need to perform an online table redefinition on the ADMIN_EMP table to move it from the TBS1 tablespace to the TBS2 tablespace in the same schema.

What action is required for the dependent objects when you perform online redefinition on the table?

- A. The materialized view should have a complete refresh performed after the online table redefinition is completed.
- B. The materialized view should have a fast refresh performed after the online table redefinition is completed.
- C. The materialized view, materialized log, and the index should be dropped and re-created after the online table redefinition is complete.
- D. The materialized view and materialized log should be dropped and all constraints disabled and re-created after the online table redefinition is complete.

Answer: A

Explanation: Restrictions for Online Redefinition of Tables

? After redefining a table that has a materialized view log, the subsequent refresh of any dependent materialized view must be a complete refresh.

NEW QUESTION 208

To enable resumable space allocation for the instance, which of the following initialization parameters should you set to a nonzero value?

- A. RESUMABLE_SPACE_TIME
- B. RESUMABLE_SPACE
- C. RESUMABLE_TIME
- D. RESUMABLE_TIMEOUT

E. TIME_RESUMABLE

Answer: D

NEW QUESTION 212

Evaluate the following command:

```
SQL>ALTER SYSTEM SET db_securefile = 'IGNORE';
```

What is the impact of this setting on the usage of SecureFiles?

- A. It forces BasicFiles to be created even if the SECUREFILE option is specified to create the LOB.
- B. It forces SecureFiles to be created even if the BASICFILE option is specified to create the LOB.
- C. It does not allow the creation of SecureFiles and generates an error if the SECUREFILE option is specified to create the LOB.
- D. It ignores the SECUREFILE option only if a Manual Segment Space Management tablespace is used and creates a BasicFile.

Answer: A

Explanation: Parameter type String

Syntax DB_SECUREFILE = { NEVER | PERMITTED | ALWAYS | IGNORE }

Default value PERMITTED

Modifiable ALTER SESSION, ALTER SYSTEM

Basic No

DB_SECUREFILE specifies whether or not to treat LOB files as SecureFiles. Values:

NEVER

Any LOBs that are specified as SecureFiles are created as BasicFile LOBs. All SecureFile- specific storage options and features (for example, compress, encrypt, deduplicate) will cause an exception. The BasicFile LOB defaults will be used for storage options not specified.

PERMITTED

LOBs are allowed to be created as SecureFiles. ALWAYS

Attempts to create all LOBs as SecureFile LOBs but creates any LOBs not in an Automatic Segment Space Managed tablespace as BasicFile LOBs, unless SECUREFILE is explicitly specified. Any BasicFile LOB storage options that are specified will be ignored and the SecureFile defaults will be used for all storage options not specified.

IGNORE

The SECUREFILE keyword and all SecureFile options are ignored.

If the COMPATIBLE initialization parameter is not set to 11.1.0 or higher, then LOBs are not treated as SecureFiles.

NEW QUESTION 216

You create a new Automatic Database Diagnostic Monitor (ADDM) task: instance_analysis_mode_task. To view the ADDM report, you use the following command:

```
SQL> SELECT dbms_addm.get_report('my_instance_analysis_mode_task') FROM dual;
```

You want to suppress ADDM output relating to Segment Advisor actions on user SCOTT's segments.

What would you do to achieve this?

- A. Add a finding directive for the ADDM task.
- B. Add a segment directive for the ADDM task.
- C. Add a parameter directive for the ADDM task.
- D. Disable the Segment Advisor from the Automatic Maintenance Task.

Answer: B

NEW QUESTION 217

Observe the following PL/SQL block: BEGIN

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

Which statement is correct regarding the above PL/SQL block?

- A. It automatically purges the SQL management objects when SMB occupies more than 30% of the SYSAUX tablespace.
- B. It reserves 30% of the space in the SYSAUX tablespace for SQL Management Base (SMB).
- C. It reserves 30% of the space in the SYSTEM tablespace for SMB.
- D. It generates a weekly warning in the alert log file when SMB occupies more than 30% of the SYSAUX tablespace.

Answer: D

NEW QUESTION 220

You need to perform an online table redefinition of an existing SALES table to partition it into two tablespaces TBS1 and TBS2. The SALES table has a materialized view, materialized log, indexes, referential integrity constraint, and triggers with the PRECEDES clause existing on it.

What action is required for dependent objects when you perform online table redefinition?

- A. The dependent materialized view should have a complete refresh performed after the online table redefinition process.
- B. Triggers with the PRECEDES clause should be disabled before the online table redefinition process.
- C. Referential integrity constraints must be manually enabled after the online table redefinition process.
- D. The materialized log should be dropped before the online table redefinition process.

Answer: A

Explanation: When performing the online table redefinition, you will:

Copy dependent objects (such as triggers, indexes, materialized view logs, grants, and constraints) and statistics from the table being redefined to the interim table, using one of the following two methods. Method 1 is the preferred method because it is more automatic, but there may be times that you would choose to use method 2. Method 1 also enables you to copy table statistics to the interim table.

Results of the Redefinition Process ([link](#))

The following are the end results of the redefinition process:

? The original table is redefined with the columns, indexes, constraints, grants, triggers, and statistics of the interim table.
? Dependent objects that were registered, either explicitly using REGISTER_DEPENDENT_OBJECT or implicitly using COPY_TABLE_DEPENDENTS, are renamed automatically so that dependent object names on the redefined table are the same as before redefinition.

Note:

If no registration is done or no automatic copying is done, then you must manually rename the dependent objects.

The referential constraints involving the interim table now involve the redefined table and are enabled. Any indexes, triggers, materialized view logs, grants, and constraints defined on the original table (prior to redefinition) are transferred to the interim table and are dropped when the user drops the interim table. Any referential constraints involving the original table before the redefinition now involve the interim table and are disabled.

Some PL/SQL objects, views, synonyms, and other table-dependent objects may become invalidated. Only those objects that depend on elements of the table that were changed are invalidated. For example, if a PL/SQL procedure queries only columns of the redefined table that were unchanged by the redefinition, the procedure remains valid. See "Managing Object Dependencies" for more information about schema object dependencies. Restrictions for Online Redefinition of Tables ([link](#))

After redefining a table that has a materialized view log, the subsequent refresh of any dependent materialized view must be a complete refresh.

NEW QUESTION 222

Evaluate the following SQL statement used to create the PRODUCTS table:

```
CREATE TABLE products (
product_id NUMBER(3) PRIMARY KEY, product_desc VARCHAR2(25),
qty NUMBER(8,2), rate NUMBER(10,2),
total_value AS ( qty * rate)) PARTITION BY RANGE (total_value) (
PARTITION p1 VALUES LESS THAN (100000), PARTITION p2 VALUES LESS THAN (150000), PARTITION p3 VALUES LESS THAN (MAXVALUE))
COMPRESS FOR ALL OPERATIONS;
```

Which statement is true regarding this command?

- A. It executes successfully but partition pruning cannot happen for this partition key.
- B. It produces an error because the TOTAL_VALUE column cannot be used as a partition key.
- C. It produces an error because compression cannot be used for the TOTAL_VALUE partition key.
- D. It executes successfully but the values in the TOTAL_VALUE column would not be physically stored in the partitions.

Answer: D

NEW QUESTION 225

Which statements are true regarding SecureFile LOBs? (Choose all that apply.)

- A. The amount of undo retained is user controlled.
- B. SecureFile LOBs can be used only for nonpartitioned tables.
- C. Fragmentation is minimized by using variable-sized chunks dynamically.
- D. SecureFile encryption allows for random reads and writes of the encrypted data.
- E. It automatically detects duplicate LOB data and conserves space by storing only one copy.

Answer: CDE

Explanation: One of the new feature of Oracle 11g is that Oracle SecureFiles use variable chunk sizes, which can be as large as 64 MB. By storing these chunks next to one another, Oracle also minimizes fragmentation.

About Deduplication

SecureFiles Intelligent Deduplication, available with the Oracle Advanced Compression Option, enables Oracle Database to automatically detect duplicate LOB data within a LOB column or partition, and conserve space by storing only one copy of the data.

Note that you must have a license for the Oracle Advanced Compression Option before implementing SecureFiles Intelligent Deduplication. See Oracle Database Licensing Information for more information.

Note also that Oracle Streams does not support SecureFiles LOBs that are deduplicated. About Encryption

SecureFiles Intelligent Encryption, available with the Oracle Advanced Security Option, introduces a new encryption facility for LOBs. The data is encrypted using Transparent Data Encryption (TDE), which allows the data to be stored securely, and still allows for random read and write access.

NEW QUESTION 227

A user receives the following error while performing a large volume of inserts into a table:

ERROR at line 1:

ORA-01536: space quota exceeded for tablespace 'USERS'

The issue is resolved by increasing the space quota on the USERS tablespace for the user. But the user may perform such transaction in the future. You want to ensure that the command waits rather than produce an error when such an event occurs the next time.

What can you do to achieve this before running the command in the future?

- A. Set RESUMABLE_TIMEOUT for the instance.
- B. Set the RESOURCE_LIMIT Parameter to TRUE.
- C. Enable the database instance to use asynchronous commit.
- D. Set the LOG_CHECKPOINT_TIMEOUT parameter to a nonzero value for the database instance.

Answer: A

Explanation: How Resumable Space Allocation Works

The following is an overview of how resumable space allocation works. Details are contained in later sections.

? A statement executes in resumable mode only if its session has been enabled for resumable space allocation by one of the following actions:

The ALTER SESSION ENABLE RESUMABLE statement is issued in the session before the statement executes when the RESUMABLE_TIMEOUT initialization parameter is set to a nonzero value. The ALTER SESSION ENABLE RESUMABLE TIMEOUT timeout_value statement is issued in the session before the statement executes, and the timeout_value is a nonzero value.

? A resumable statement is suspended when one of the following conditions occur

(these conditions result in corresponding errors being signalled for non-resumable statements):

Out of space condition Maximum extents reached condition Space quota exceeded condition.

? When the execution of a resumable statement is suspended, there are

mechanisms to perform user supplied operations, log errors, and query the status of the statement execution. When a resumable statement is suspended the following actions are taken:

The error is reported in the alert log.

The system issues the Resumable Session Suspended alert.

If the user registered a trigger on the AFTER SUSPEND system event, the user trigger is executed. A user supplied PL/SQL procedure can access the error message data using the DBMS_RESUMABLE package and the DBA_ or USER_RESUMABLE view.

? Suspending a statement automatically results in suspending the transaction. Thus all transactional resources are held through a statement suspend and resume.

When the error condition is resolved (for example, as a result of user intervention or perhaps sort space released by other queries), the suspended statement automatically resumes execution and the Resumable Session Suspended alert is cleared.

? A suspended statement can be forced to throw the exception using the

DBMS_RESUMABLE.ABORT() procedure. This procedure can be called by a DBA, or by the user who issued the statement.

? A suspension time out interval, specified by the RESUMABLE_TIMEOUT

initialization parameter or by the timeout value in the ALTER SESSION ENABLE RESUMABLETIMEOUT statement, is associated with resumable statements. A

resumable statement that is suspended for the timeout interval wakes up and returns the exception to the user if the error condition is not resolved within the timeout interval.

? A resumable statement can be suspended and resumed multiple times during execution.

NEW QUESTION 229

What happens when you run the SQL Tuning Advisor with limited scope?

- A. Access path analysis is not performed for SQL statements.
- B. SQL structure analysis is not performed for SQL statements.
- C. SQL Profile recommendations are not generated for SQL statements.
- D. Staleness and absence of statistics are not checked for the objects in the SQL Tuning Advisor.

Answer: B

Explanation: Topic 16, Space Management

NEW QUESTION 233

Which steps are mandatory to enable Direct NFS?

1. Mount all required file systems using the kernel NFS driver.
2. Create an orafstab file containing the attributes for each NFS server to be accessed using Direct NFS.
3. Replace the ODM library libodm11.so_stub with libodm11.so.

- A. 2 and 3
- B. 1 and 3
- C. 1 and 2
- D. 1, 2 and 3

Answer: B

NEW QUESTION 236

You work with a newly created database. Presently, there is no application load on the database instance. You want to create a baseline for tuning the application, so you decide to collect recommendations that can be implemented to improve application performance.

What action must you take to achieve this?

- A. Run Segment Advisor
- B. Run the SQL Tuning Advisor (STA)
- C. Run the Automatic Workload Repository (AWR) report
- D. Run the SQL Access Advisor with a hypothetical workload

Answer: D

NEW QUESTION 240

The SQL Tuning Advisor configuration has default settings in your database instance.

Which recommendation is automatically implemented after the SQL Tuning Advisor is run as part of the automatic maintenance task?

- A. statistics recommendations
- B. SQL Profile recommendations
- C. Index-related recommendations
- D. restructuring of SQL recommendations

Answer: B

NEW QUESTION 243

Which is the source used by Automatic SQL Tuning that runs as part of the AUTOTASK framework?

- A. SQL statements that are part of the AWR baseline only
- B. SQL statements based on the AWR top SQL identification
- C. SQL statements that are part of the available SQL Tuning Set (STS) only
- D. SQL statements that are available in the cursor cache and executed by a user other than SYS

Answer: B

NEW QUESTION 244

Which of the following identifies and creates an index to minimize the DB time for a particular SQL statement?

- A. The SGA Tuning Advisor
- B. The SQL Access Advisor
- C. The SQL Tuning Advisor
- D. The Memory Advisor

Answer: C

Explanation: The SQL Access Advisor allows the DBA to gather global recommendations for a workload. The SQL Tuning advisor is more granular, tuning a single statement.

The main functions of the SQL Access advisor is to recommend missing indexes and materialized views, but a comprehensive task analysis will also create SQL Profiles that can be used within the SQL Tuning advisor. The DBA defines the SQL used in the SQL Access Advisor task, and can choose current SQL, a user-defined set of SQL, a historical workload, or a hypothetical workload.

A hypothetical workload is very useful because the DBA need-only specify the tables that participate in the queries, and the SQL Access Advisor gathers the appropriate SQL statements to create the workload. The SQL Tuning Advisor (STA) is primarily designed to replace the manual tuning of SQL statements and speed up the overall SQL tuning process. The SQL Tuning Advisor studies poorly executing SQL statements and evaluates resource consumption in terms of CPU, I/O, and temporary space.

The SQL Tuning Advisor (STA) works with the Automatic Tuning Optimizer (ATO) to analyze historical SQL workload using data from the AWR, and it generates recommendations for new indexes and materialized views that will reduce the disk I/O associated with troublesome SQL statements.

http://www.dba-oracle.com/oracle10g_tuning/t_sql_tuning_advisor.htm

NEW QUESTION 248

You have a very large table that your users access frequently. Which of the following advisors will recommend any indexes to improve the performance of queries against this table?

- A. The Automatic Memory Manager (AMM)
- B. The SQL Tuning Advisor
- C. The Segment Advisor
- D. The SQL Access Advisor

Answer: D

Explanation: The SQL Access Advisor analyzes all SQL running during a given time period and recommends indexes and materialized views to improve the overall performance of the database.

A is incorrect because there is no such advisor as the Automatic Memory Manager.

B is incorrect because the SQL Tuning Advisor looks only at a single SQL statement and provides recommendations.

C is incorrect because the Segment Advisor recommends segment shrink when table and index segments are heavily fragmented.

NEW QUESTION 252

Your company wants to upgrade the production database to a Real Application Clusters (RAC) environment.

You set up the best RAC database and want to replay a recorded workload captured from the production database on the test machine.

The following steps may be used to replay the database workload:

- 1) Preprocess the captured workload
- 2) Restart the database in RESTRICTED mode.
- 3) Set up Replay Clients.
- 4) Restore the test database to the point when the capture started.
- 5) Remap connections.

Which is the correct sequence of the required steps?

- A. 1, 4, 3, 5
- B. 1, 4, 5, 3
- C. 2, 1, 5, 3, 4
- D. 1, 2, 4, 5, 3

Answer: B

Explanation: 11.1 Steps for Replaying a Database Workload

Proper planning of the workload replay and preparation of the replay system ensures that the replay will be accurate. Before replaying a database workload, review and complete the

following steps as appropriate:

- ? Setting Up the Replay Directory
- ? Restoring the Database
- ? Resolving References to External Systems
- ? Remapping Connections
- ? User Remapping
- ? Specifying Replay Options
- ? Using Filters with Workload Replay
- ? Setting Up Replay Clients

NEW QUESTION 256

Which tuning tool recommends how to optimize materialized views so that these views can take advantage of the general query rewrite feature?

- A. Segment Advisor
- B. SQL Access Advisor

- C. Undo Advisor
- D. SQL Tuning Advisor

Answer: B

NEW QUESTION 258

The SQL Tuning Advisor has been configured with default configurations in your database instance. Which recommendation is automatically implemented without the DBA's intervention after the SQL Tuning Advisor is run as part of the AUTOTASK framework?

- A. statistics recommendations
- B. SQL profile recommendations
- C. index-related recommendations
- D. restructuring of SQL recommendations

Answer: B

NEW QUESTION 259

An index called ORD_CUSTNAME_IX has been created on the CUSTNAME column in the ORDERS table using the following command:

```
SQL>CREATE INDEX ord_custname_ix ON orders (custname);
```

The ORDERS table is frequently queried using the CUSTNAME column in the WHERE clause. You want to check the impact on the performance of the queries if the index is not available. You do not want the index to be dropped or rebuilt to perform this test.

Which is the most efficient method of performing this task?

- A. disabling the index
- B. making the index invisible
- C. making the index unusable
- D. using the MONITORING USAGE clause for the index

Answer: B

Explanation: Invisible Indexes

An invisible index is maintained by Oracle Database for every data manipulation language (DML) statement, but is ignored by the optimizer unless you explicitly set the parameter OPTIMIZER_USE_INVISIBLE_INDEXES to TRUE on a session or system level.

Making an index invisible is an alternative to making it unusable or dropping it. Using invisible indexes, you can:

? Test the removal of an index before dropping it

? Create invisible indexes temporarily for specialized, nonstandard operations, such as online application upgrades, without affecting the behavior of existing applications

For more information, see Oracle Database Administrator's Guide.

NEW QUESTION 264

Which is true concerning Database Replay in an Oracle Real Application Cluster (RAC) database?

- A. Workload capture is per instance.
- B. You only need to restart one instance to begin workload capture.
- C. Specifically in RAC, you shut down all instances, restart them individually, and begin workload capture with the last instance started.
- D. RAC does not support workload capture, but it does support workload replay.
- E. None of the above.

Answer: E

NEW QUESTION 268

Which two statements regarding a SQL profile are true? (Choose two.)

- A. It is built by Automatic Tuning Optimizer.
- B. It cannot be stored persistently in the data dictionary.
- C. It can be used by the query optimizer automatically.
- D. It can be created manually by using the CREATE PROFILE command.

Answer: AC

NEW QUESTION 270

What recommendations does the SQL Access Advisor provide for optimizing SQL queries? (Choose all that apply.)

- A. selection of SQL plan baselines
- B. partitioning of tables and indexes
- C. creation of index-organized tables
- D. creation of bitmap, function-based, and B-tree indexes
- E. optimization of materialized views for maximum query usage and fast refresh

Answer: BDE

Explanation: Refer to here.

SQL Access Advisor helps you achieve your performance goals by recommending the proper set of materialized views, materialized view logs, partitions, and indexes for a given workload.

SQL Access Advisor index recommendations include bitmap, function-based, and B-tree indexes. SQL Access Advisor materialized view recommendations include fast refreshable and full refreshable MVs, for either general rewrite or exact text match rewrite.

NEW QUESTION 273

Which of the following advisors is run in every maintenance window by the auto-task system?

- A. The Memory Advisor
- B. The SQL Tuning Advisor
- C. The Undo Advisor
- D. The SQL Access Advisor

Answer: B

Explanation: About Automated Maintenance Tasks (link)

Oracle Database has three predefined automated maintenance tasks:

? Automatic Optimizer Statistics Collection

? Automatic Segment Advisor

? Automatic SQL Tuning Advisor

By default, all three automated maintenance tasks are configured to run in all maintenance windows.

NEW QUESTION 278

What two statements are true regarding the recommendations received from the SQL Access Advisor? (Choose two.)

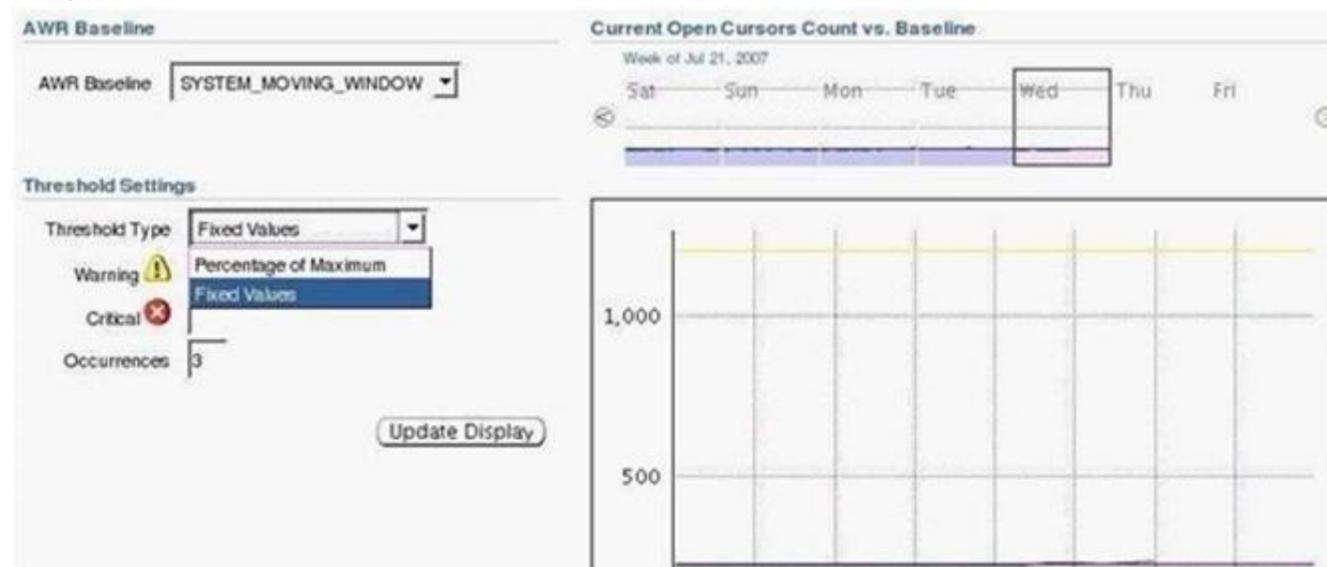
- A. It cannot generate recommendations that support multiple workload queries.
- B. It can recommend partitioning on tables provided that the workloads have some predicates and joins on the columns of the NUMBER or DATE type.
- C. It can recommend partitioning only on tables that have at least 10,000 rows.
- D. It can recommend only B-tree indexes and not bitmap or function-based indexes.

Answer: BC

NEW QUESTION 279

View the Exhibit that sets the threshold for the Current Open Cursors Count metric. Why is the Significance Level threshold type not available in the threshold setting?

Exhibit:



- A. because AWR baseline is not enabled
- B. because Current Open Cursors Count is not a basic metric
- C. because the STATISTICS_LEVEL parameter is set to BASIC
- D. because the AWR baseline is a system-defined moving window baseline

Answer: B

Explanation:

ORACLE Enterprise Manager 11g Database Control

Database Instance: iocp > Baseline Metric Thresholds >

Edit Thresholds: Current Open Cursors Count

Cancel
Last

Current Open Cursors Count is not a basic metric.
Metrics that are not basic do not support Significance Level thresholds.
Note: Non-basic metrics display average values over AWR snapshots, not per-minute metric values.

AWR Baseline

Name: SYSTEM_MOVING_WINDOW

Threshold Settings

Threshold Type: Fixed Values

Critical:

Warning:

Occurrences: 2

Preview

Current Open Cursors Count vs. Baseline

Week of Mar 22, 2014

C:\Users\albo\Desktop\1-1.jpg

NEW QUESTION 280

View the Exhibit for some of the parameter settings. You start a session and issue the following command:
SQL>CREATE INDEX emp_ename ON emp(ename) TABLESPACE users INVISIBLE;
What is the outcome of the above command? Exhibit:

NAME	TYPE	VALUE
optimizer_dynamic_sampling	integer	2
optimizer_features_enable	string	11.1.0.6
optimizer_use_invisible_indexes	boolean	FALSE
optimizer_index_caching	integer	0
optimizer_index_cost_adj	integer	100
skip_unusable_indexes	boolean	TRUE

- A. The index is not used by the optimizer but is maintained during DML operations.
- B. The index is not used by the optimizer and is not maintained during DML operations.
- C. The index is used by the optimizer only if a hint is specified in the query statement and is maintained during DML operations.
- D. The index is used by the optimizer only if a hint is specified in the query statement but is not maintained during DML operations.

Answer: A

NEW QUESTION 284

Which of the following advisors within the Oracle advisory framework will analyze a single SQL statement and make recommendations for performance improvement?

- A. SQL Repair Advisor
- B. SQL Optimizer
- C. SQL Access Advisor
- D. SQL Tuning Advisor

Answer: D

NEW QUESTION 288

Which three are the valid statements in relation to SQL plan baselines? (Choose three.)

- A. The plans can be manually loaded to the SQL plan baseline.
- B. The plans in the SQL plan baseline are verified and accepted plans.
- C. The plans generated for every SQL statement are stored in the SQL plan baseline by default.
- D. The plan baselines are stored temporarily in the memory as long as the database instance is running.
- E. For the SQL plan baselines to be accessible to the optimizer, the SYSAUX tablespace must be online.

Answer: ABE

Explanation: The SQL management base (SMB), which is part of the data dictionary, stores the SQL plan baselines and plan history in the SYSAUX tablespace.

The SMB also contains SQL profiles. The SMB uses automatic space management.

Capturing SQL Plan Baselines (refer to here).

You can configure the SQL Plan Baseline Capture phase for automatic capture of plan history and SQL plan baselines for repeatable SQL statements.

Alternatively, you can manually load plans as SQL plan baselines.

Not all the new generated SQL plan can be accepted, so that all SQL plan stored in the plan history, only after evolving the plans in the plan history, it becomes in the plan baseline.

You can evolve an existing SQL plan baseline by manually loading plans from the shared SQL area or from a SQL tuning set. When you manually load plans into a SQL plan baseline, the database adds these loaded plans as accepted plans.

NEW QUESTION 292

Your company wants to upgrade the current production database to the RAC environment. To perform testing before migrating to the RAC environment, you performed the workload capture on the production database to record the peak workload. You set up the test RAC database and want to replay the recorded workload on the test machine. Note the following steps that you may require to replay the database workload:

- 1) Preprocess the captured workload.
- 2) Restart the database in RESTRICTED mode.
- 3) Set up the Replay Clients.
- 4) Restore the test database to the point when the capture started.
- 5) Remap connections.

Arrange the steps required in the correct sequence to accomplish this task on the test machine.

- A. 1, 4, 5, 3 (2 is not required.)
- B. 1, 4, 3, 5 (2 is not required.)
- C. 1, 2, 4, 5 (3 is not required.)
- D. 2, 1, 5, 3, 4
- E. 1, 2, 4, 5, 3

Answer: A

Explanation: 11.1 Steps for Replaying a Database Workload

Proper planning of the workload replay and preparation of the replay system ensures that the replay will be accurate. Before replaying a database workload, review and complete the following steps as appropriate:

- ? Setting Up the Replay Directory
- ? Restoring the Database
- ? Resolving References to External Systems
- ? Remapping Connections
- ? User Remapping
- ? Specifying Replay Options
- ? Using Filters with Workload Replay
- ? Setting Up Replay Clients

NEW QUESTION 297

Which two statements about Oracle Direct Network File System (NFS) are true? (Choose two.)

- A. It bypasses the OS file system cache.
- B. A separate NFS interface is required for use across Linux, UNIX, and Windows platforms.
- C. It uses the operating system kernel NFS layer for user tasks and network communication modules.
- D. File systems need not be mounted by the kernel NFS system when being served through Direct NFS.
- E. Oracle Disk Manager can manage NFS on its own, without using the operating system kernel NFS driver.

Answer: AE

Explanation: Direct NFS provides faster performance than what can be provided by the operating system's NFS driver as Oracle bypasses the operating system and generates exactly the requests it needs (no user configuration or tuning required). Data is cached just once in user space, which saves memory (no second copy in kernel space). Performance is further improved by load balancing across multiple network interfaces (if available).

NEW QUESTION 301

Which of the following is a potential performance tuning recommendation from the SQL Access Advisor?

- A. Create new indexes.
- B. Modify existing indexes.
- C. Implement partitioning on a nonpartitioned table.
- D. Create materialized views.
- E. All of the above

Answer: E

Explanation: Overview of SQL Access Advisor

Materialized views, partitions, and indexes are essential when tuning a database to achieve optimum performance for complex, data-intensive queries. SQL Access Advisor helps you achieve your performance goals by recommending the proper set of materialized views, materialized view logs, partitions, and indexes for a given workload. Understanding and using these structures is essential when optimizing SQL as they can result in significant performance improvements in data retrieval. The advantages, however, do not come without a cost. Creation and maintenance of these objects can be time consuming, and space requirements can be significant. In particular, partitioning of an unpartitioned base table is a complex operation that must be planned carefully.

NEW QUESTION 303

You plan to use SQL Performance Analyzer to analyze the SQL workload. You created a SQL Tuning Set as a part of the workload capturing. What information is

captured as part of this process? (Choose all that apply.)

- A. the SQL text
- B. the execution plan
- C. the execution context
- D. the execution frequency
- E. the system change number (SCN)

Answer: ACD

Explanation: <http://www.oracle.com/technetwork/cn/articles/o69ocp-099954.html>

NEW QUESTION 308

View the Exhibit1 to examine the series of SQL commands. View the Exhibit2 to examine the plans available in the SQL plan baseline. The baseline in the first row of the Exhibit is created when OPTIMIZER_MODE was set to FIRST_ROWS. Which statement is true if the SQL query in exhibit1 is executed again when the value of OPTIMIZER_MODE is set to FIRST_ROWS?

```
SQL> SELECT signature, sql_handle, plan_name, origin, enabled,
accepted, fixed, autopurge
FROM dba_sql_plan_baselines;
```

SIGNATURE	SQL_HANDLE	PLAN_NAME	ORIGIN	ENABLED	ACCEPTED	FIXED
8.062E+18	SYS_SQL_6fa2	SYS_SQL_PLAN_1ea	AUTO-CAPTURE	YES	NO	NO
8.062E+18	SYS_SQL_6fe2	SYS_SQL_PLAN_4be	AUTO-CAPTURE	YES	YES	NO
...						
...						
...						

parameter-optimizer (exhibit):

```
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 > 40 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;
```

- A. The optimizer uses a new plan because none of the plans in the exhibit2 are fixed plans.
- B. The optimizer uses the plan in the second row of the exhibit2 because it is an accepted plan.
- C. The optimizer uses the plan in the first row of the exhibit2 because it is the latest generated plan.
- D. The optimizer uses the plan in the first row of the exhibit2 because OPTIMIZER_MODE was set to FIRST_ROW during its creation.

Answer: B

Explanation: Setting the OPTIMIZER_MODE Initialization Parameter(Link)

The OPTIMIZER_MODE initialization parameter establishes the default behavior for choosing an optimization approach for the instance.

OPTIMIZER_MODE Initialization Parameter Values

? ALL_ROWS, The optimizer uses a cost-based approach for all SQL statements in the session regardless of the presence of statistics and optimizes with a goal of best throughput (minimum resource use to complete the entire statement). This is the default value.

? FIRST_ROWS_n, The optimizer uses a cost-based approach, regardless of the presence of statistics, and optimizes with a goal of best response time to return the first n number of rows, where n equals 1, 10, 100, or 1000.

? FIRST_ROWS, The optimizer uses a mix of cost and heuristics to find a best plan for fast delivery of the first few rows.

Note that using heuristics sometimes leads the optimizer to generate a plan with a cost that is significantly larger than the cost of a plan without applying the heuristic. FIRST_ROWS is available for backward compatibility and plan stability; use FIRST_ROWS_n instead.

NEW QUESTION 312

Your system has been upgraded from Oracle Database 10g to Oracle Database 11g. You imported SQL Tuning Sets (STS) from the previous version. After changing the OPTIMIZER_FEATURES_ENABLE parameter to 10.2.0.4 and running the SQL Performance Analyzer, you observed performance regression for a few SQL statements. What would you do with these SQL statements?

- A. Set OPTIMIZER_USE_PLAN_BASELINES to FALSE to prevent the use of regressed plans.
- B. Capture the plans from the previous version using STS and then load them into the stored outline.

- C. Capture the plans from the previous version using STS and then load them into SQL Management Base (SMB).
- D. Set OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES to FALSE to prevent the plans from being loaded to the SQL plan baseline.

Answer: C

Explanation: The SQL management base (SMB) is a part of the data dictionary that resides in the SYSAUX tablespace. It stores statement logs, plan histories, SQL plan baselines, and SQL profiles.

Parameters Relating to Stored Outline Migration:

? OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES, Enables or disables the automatic recognition of repeatable SQL statement and the generation of SQL plan baselines for these statements.

? OPTIMIZER_USE_SQL_PLAN_BASELINES, Enables or disables the use of SQL plan baselines stored in SQL Management Base.

NEW QUESTION 316

Performance divergence indicated in the Workload Replay report is most likely due to what?

- A. DML and SQL statement results that do not match between the capture and replay systems
- B. When errors that occur in the capture system don't occur in the replay system
- C. Top SQL statements
- D. Infrastructure or system-configuration differences
- E. Time-of-day differences between capture and replay systems

Answer: D

NEW QUESTION 318

Which two changes and their effect on the system can be tested by using the Database Replay feature? (Choose two.)

- A. multiplexing of the control file
- B. adding the redo log member to the database
- C. database and operating system upgrades
- D. changing the database storage to ASM-managed storage

Answer: CD

NEW QUESTION 321

To generate recommendations to improve the performance of a set of SQL queries in an application, you execute the following blocks of code:

```
BEGIN
  dbms_advisor.create_task(dbms_advisor.sqlaccess_advisor, 'TASK1');
END;
/
BEGIN
  dbms_advisor.set_task_parameter('TASK1', 'ANALYSIS_SCOPE', 'ALL');
  dbms_advisor.set_task_parameter('TASK1', 'MODE', 'COMPREHENSIVE');
END;
/
BEGIN
  dbms_advisor.execute_task('TASK1');
  dbms_output.put_line(dbms_advisor.get_task_script('TASK1'));
END;
/
```

The blocks of code execute successfully; however, you do not get the required outcome. What could be the reason?

- A. A template needs to be associated with the task.
- B. A workload needs to be associated with the task.
- C. The partial or complete workload scope needs to be associated with the task.
- D. The type of structures (indexes, materialized views, or partitions) to be recommended need to be specified for the task.

Answer: B

NEW QUESTION 322

Examine the following PL/SQL block:

```
SQL> SET SERVEROUTPUT ON;
SQL> SET LONG 10000;
SQL> DECLARE report clob;
BEGIN
  report := DBMS_SPM.EVOLVE_SQL_PLAN_BASELINE();
  DBMS_OUTPUT.PUT_LINE(report);
END;
/
```

Which statement describes the effect of the execution of the above PL/SQL block?

- A. The plan baselines are verified with the SQL profiles.

- B. All fixed plan baselines are converted into nonfixed plan baselines.
- C. All the nonaccepted SQL profiles are accepted into the plan baseline.
- D. The nonaccepted plans in the SQL Management Base are verified with the existing plan baselines.

Answer: D

NEW QUESTION 323

Which of these appropriately describes the results of a manual SQL Tuning Advisor task?

- A. A list of SQL statements and recommendations for tuning
- B. A list of SQL statements that have been tuned by the Advisor, with before and after metrics
- C. Graphs showing the actual performance improvement made by the Advisor after it implemented the recommended changes
- D. All of the above

Answer: A

NEW QUESTION 324

When creating a SQL tuning set, which of the following steps allows the DBA to reduce the size of the SQL set by selecting specific operators and values?

- A. Filter versions
- B. Filter loads
- C. Filter tasks
- D. Filter options

Answer: D

NEW QUESTION 325

View the Exhibit to examine the Automatic SQL Tuning result details. Which action would you suggest for the selected SQL statement in the Exhibit?

Only profiles that significantly improve SQL performance were implemented.

Select	SQL Text	Parsing Schema	SQL ID	Statistics	SQL Profile	Index	Restructure SQL	Miscellaneous	Error Date
<input type="radio"/>	SELECT NULL AS table_cat, Lower...	SYSMAN	361qjn3w9vfh	✓	(99.9%) ✓				7/12/07
<input checked="" type="radio"/>	SELECT EXECUTION_ID, STATUS, STATUS_DET...	SYSMAN	lyk8b9986nkk7		(60%) ✓	(97.9%) ✓			7/12/07
<input type="radio"/>	SELECT /*+ INDEX(sqlobj\$ (signature cate...	SYS	8b75qwpna202v					✓	7/12/07
<input type="radio"/>	select OBJOID, CLSOID, RUNTIME, PRI, JO...	SYS	8vf1dtwgk1xy5					✓	7/12/07
<input type="radio"/>	select smontabv.cnt, smontab.time_mp, ...	SYS	4q8mr2bvy6gr					✓	7/12/07
<input type="radio"/>	select t.ts#,t.file#,t.block#,nv(t.boj...	SYS	1qu8t96d0bdmu					✓	7/12/07
<input type="radio"/>	select obj#, dataobj#, part#, hibounden...	SYS	130dvvr5e8bqn					✓	7/12/07
<input type="radio"/>	select privilege#,level from sysauth\$ co...	SYS	0n6b2sojwb74n					✓	7/12/07
<input type="radio"/>	select value(p\$) from "XDB"."XDB\$RESOURC...	SYS	23y48d28wkg2r					✓	7/12/07
<input type="radio"/>	SELECT obj_type, plan_id, name, flags, L...	SYS	On1nassmccz0c					✓	7/12/07

- A. Accept the recommended SQL profile.
- B. Collect statistics for the related objects.
- C. Run the Access Advisor for the SQL statement.
- D. Run the Segment Advisor for recommendations.

Answer: C

NEW QUESTION 326

The OPTIMIZER_USE_PLAN_BASELINES parameter is set to TRUE. The optimizer generates a plan for a SQL statement but does not find a matching plan in the SQL plan baseline.

Which two operations are performed by the optimizer in this scenario? (Choose two.)

- A. The optimizer adds the new plan to the plan history.
- B. The optimizer selects the new plan for the execution of the SQL statement.
- C. The optimizer adds the new plan to the SQL plan baseline as an accepted plan.
- D. The optimizer adds the new plan to the SQL plan baseline but not in the ENABLED state.
- E. The optimizer costs each of the accepted plans in the SQL plan baseline and picks the one with the lowest cost.

Answer: AE

Explanation: 15.2.2 Selecting SQL Plan Baselines

During the SQL plan baseline selection phase, Oracle Database detects plan changes based on the stored plan history, and selects plans to avoid potential

performance regressions for a set of SQL statements.

Each time the database compiles a SQL statement, the optimizer does the following:

1. Uses a cost-based search method to build a best-cost plan
2. Tries to find a matching plan in the SQL plan baseline
3. Does either of the following depending on whether a match is found:

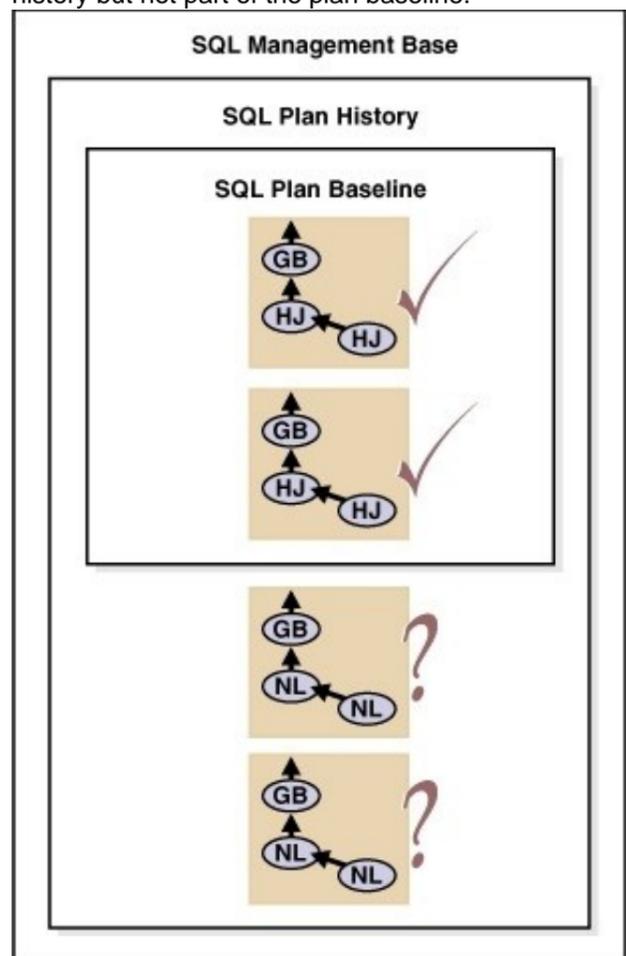
? If found, then the optimizer proceeds using the matched plan

? If not found, then the optimizer evaluates the cost of each accepted plan in the SQL plan baseline and selects the plan with the lowest cost

The best-cost plan found by the optimizer that does not match any plans in the plan history for the SQL statement represents a new plan. The database adds this plan as a nonaccepted plan to the plan history. The database does not use the new plan until it is verified to not cause a performance regression. However, if a change in the system (such as a dropped index) causes all accepted plans to become non-reproducible, then the optimizer selects the best-cost plan. Thus, the presence of a SQL plan baseline causes the optimizer to use conservative plan selection strategy for the SQL statement.

To enable the use of SQL plan baselines, set the OPTIMIZER_USE_SQL_PLAN_BASELINES initialization parameter to TRUE (default). A SQL plan baseline contains one or more accepted plans.

The plan history is the set of plans, both accepted and not accepted, that the optimizer generates for a SQL statement over time, the plans in the baseline form a subset of the plan history. For example, after the optimizer generates the first acceptable plan for a SQL plan baseline, subsequent plans are part of the plan history but not part of the plan baseline.



C:\Users\albo\Desktop\1-1.jpg

NEW QUESTION 331

You want the Automatic SQL Tuning process to stop accepting and implementing the recommended SQL profiles automatically. Which action would you perform to achieve this?

- A. Edit the automatic maintenance window group configuration.
- B. Set the CURSOR_SHARING parameter to EXACT for the database instance.
- C. Use the DBMS_SQLTUNE.SET_TUNING_TASK_PARAMETERS procedure to set ACCEPT_SQL_PROFILES to FALSE.
- D. Set the SQLTUNE_CATEGORY parameter to DEFAULT for the database instance.

Answer: C

NEW QUESTION 336

You want to analyze a SQL Tuning Set (STS) using SQL Performance Analyzer in a test database.

Which two statements are true regarding the activities performed during the test execution of SQLs in a SQL Tuning Set? (Choose two.)

- A. Every SQL statement in the STS is considered only once for execution.
- B. The SQL statements in the STS are executed concurrently to produce the execution plan and execution statistics.
- C. The execution plan and execution statistics are computed for each SQL statement in the STS.
- D. The effects of DDL and DML are considered to produce the execution plan and execution statistics.

Answer: AC

NEW QUESTION 337

Which client requests to the database can be captured as a 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: CD

NEW QUESTION 342

Which tasks are run automatically as part of the Automated Maintenance Task by default? (Choose all that apply.)

- A. Segment Advisor
- B. SQL Access Advisor
- C. Optimizer statistics gathering
- D. Automatic SQL Tuning Advisor
- E. Automatic Database Diagnostics Monitor

Answer: ACD

Explanation: About Automated Maintenance Tasks (link)

Oracle Database has three predefined automated maintenance tasks:

? Automatic Optimizer Statistics Collection—Collects optimizer statistics for all schema objects in the database for which there are no statistics or only stale statistics. The statistics gathered by this task are used by the SQL query optimizer to improve the performance of SQL execution.

See Also:

Oracle Database Performance Tuning Guide for more information on automatic statistics collection

? Automatic Segment Advisor—Identifies segments that have space available for reclamation, and makes recommendations on how to defragment those segments.

You can also run the Segment Advisor manually to obtain more up-to-the-minute recommendations or to obtain recommendations on segments that the Automatic Segment Advisor did not examine for possible space reclamation.

See Also:

"Using the Segment Advisor" for more information.

? Automatic SQL Tuning Advisor—Examines the performance of high-load SQL statements, and makes recommendations on how to tune those statements. You can configure this advisor to automatically implement SQL profile recommendations.

See Also:

Oracle Database Performance Tuning Guide for more information on SQL Tuning Advisor By default, all three automated maintenance tasks are configured to run in all maintenance windows.

NEW QUESTION 346

In Oracle 11g, which recommendations does the SQL Access Advisor generate? (Choose all that apply.)

- A. partitioning recommendations
- B. statistics collection recommendations
- C. index creation recommendations
- D. materialized view recommendations
- E. materialized view log recommendations

Answer: ACDE

Explanation: Overview of SQL Access Advisor

Materialized views, partitions, and indexes are essential when tuning a database to achieve optimum performance for complex, data-intensive queries. SQL Access Advisor helps you achieve your performance goals by recommending the proper set of materialized views, materialized view logs, partitions, and indexes for a given workload. Understanding and using these structures is essential when optimizing SQL as they can result in significant performance improvements in data retrieval. The advantages, however, do not come without a cost. Creation and maintenance of these objects can be time consuming, and space requirements can be significant. In particular, partitioning of an unpartitioned base table is a complex operation that must be planned carefully.

NEW QUESTION 351

While deploying a new application module, the software vendor ships the application software along with appropriate SQL plan baselines for the new SQLs being introduced. Which two statements describe the consequences? (Choose two.)

- A. The plan baselines can be evolved over time to produce better performance.
- B. The newly generated plans are directly placed into the SQL plan baseline without being verified.
- C. The new SQL statements initially run with the plans that are known to produce good performance under standard test configuration.
- D. The optimizer does not generate new plans for the SQL statements for which the SQL plan baseline has been imported.

Answer: AC

Explanation: Purpose of SQL Plan Baselines

Common scenarios where SQL plan management can improve or preserve SQL performance include:

A database upgrade that installs a new optimizer version usually results in plan changes for a small percentage of SQL statements. Most of these plan changes result in either no performance change or improvement. However, some plan changes may cause performance regressions. SQL plan baselines significantly minimize potential regressions resulting from an upgrade.

Ongoing system and data changes can impact plans for some SQL statements, potentially causing performance regressions. SQL plan baselines help minimize performance regressions and stabilize SQL performance.

Deployment of new application modules means introducing new SQL statements into the database. The application software may use appropriate SQL execution plans developed in a standard test configuration for the new statements. If the system configuration is significantly different from the test configuration, then the database can evolve SQL plan baselines over time to produce better performance.

NEW QUESTION 354

You installed Oracle Database 11g and are performing a manual upgrade of the Oracle9i database. As a part of the upgrade process, you execute the following script:

```
SQL> @utlu11i.sql
```

Which statement about the execution of this script is true?

- A. It must be executed from the Oracle Database 11g environment.
- B. It must be executed only after the SYSAUX tablespace has been created.
- C. It must be executed from the environment of the database that is being upgraded.
- D. It must be executed only after AUTOEXTEND is set to ON for all existing tablespaces.
- E. It must be executed from both the Oracle Database 11g and Oracle Database 9i environments.

Answer: C

Explanation: Script Name: utlu111i.SQL

Used to: Pre-Upgrade Information

Analyzes the database to be upgraded, detailing requirements and issues for the upgrade to Oracle Database 11g Release 1 (11.1)

NEW QUESTION 356

Which two prerequisites are needed for performing workload capture and replay? (Choose two.)

- A. Close all sessions performing queries using database links.
- B. running the database in shared server mode
- C. The database on which the workload is replayed has to be a restore of the original database to a specific SCN.
- D. setting up the directory to capture the workload

Answer: CD

NEW QUESTION 357

Which two statements about the SQL Management Base (SMB) are true? (Choose two.)

- A. It contains only SQL profiles generated by SQL Tuning Advisor.
- B. It stores plans generated by the optimizer using a stored outline.
- C. It is part of the data dictionary and stored in the SYSAUX tablespace.
- D. It is part of the data dictionary and stored in the SYSTEM tablespace.
- E. It contains the statement log, the plan history, plan baselines, and SQL profiles.

Answer: CE

NEW QUESTION 362

Identify the activities performed as part of the Automatic SQL Tuning process in the maintenance window? (Choose all that apply.)

- A. generating the SQL profile
- B. testing and accepting the SQL profile
- C. generating a list of candidate SQLs for tuning
- D. adding tuned SQL plans into the SQL plan baseline
- E. tuning each SQL statement in the order of importance
- F. generating baselines that include candidate SQLs for tuning

Answer: ABCE

Explanation: 17.2.1 How Automatic SQL Tuning Works (link)

Oracle Database automatically runs SQL Tuning Advisor on selected high-load SQL statements from the Automatic Workload Repository (AWR) that qualify as tuning candidates. This task, called Automatic SQL Tuning, runs in the default maintenance windows on a nightly basis. By default, automatic SQL tuning runs for at most one hour. You can customize attributes of the maintenance windows, including start and end time, frequency, and days of the week.

After automatic SQL tuning begins, the database performs the following steps:

1. Identifies SQL candidates in the AWR for tuning

Oracle Database analyzes statistics in AWR and generates a list of potential SQL statements that are eligible for tuning. These statements include repeating high-load statements that have a significant impact on the database.

The database tunes only SQL statements that have an execution plan with a high potential for improvement. The database ignores recursive SQL and statements that have been tuned recently (in the last month), parallel queries, DML, DDL, and SQL statements with performance problems caused by concurrency issues. The database orders the SQL statements that are selected as candidates based on their performance impact. The database calculates the impact by summing the CPU time and the I/O times in AWR for the selected statement in the past week.

2. Tunes each SQL statement individually by calling SQL Tuning Advisor

During the tuning process, the database considers and reports all recommendation types, but it can implement only SQL profiles automatically.

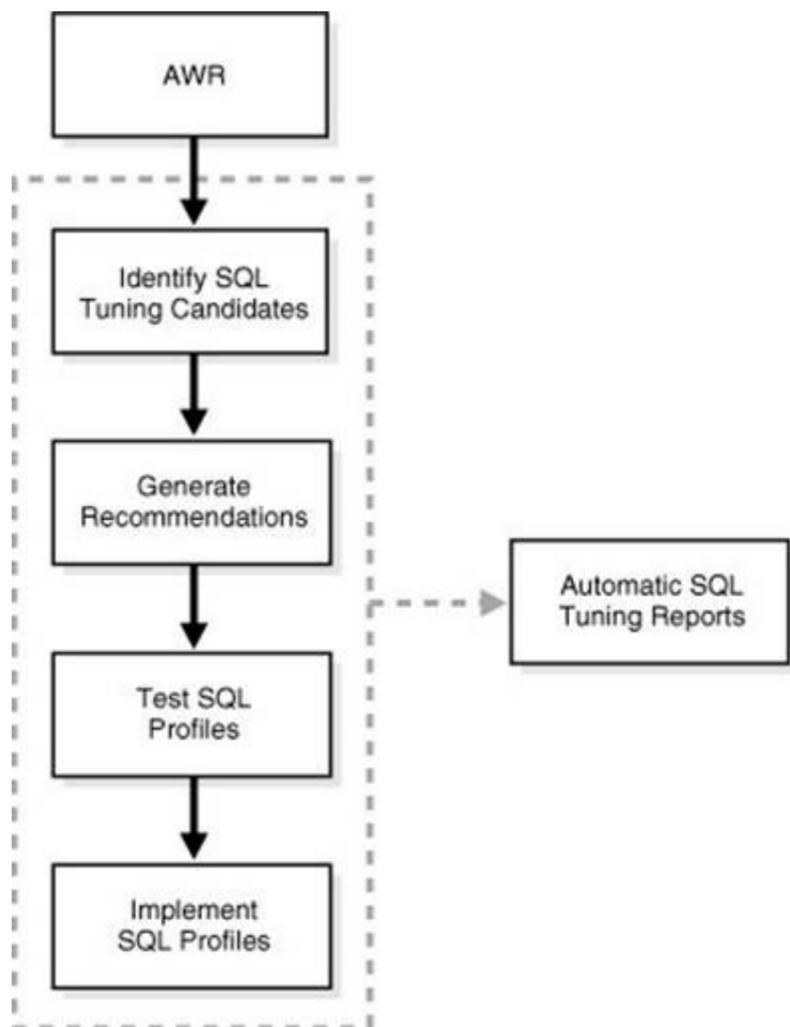
3. Tests SQL profiles by executing the SQL statement

If a SQL profile is recommended, the database tests the new profile by executing the SQL statement both with and without the profile. If the performance improvement improves at least threefold, then the database accepts the SQL profile, but only if the ACCEPT_SQL_PROFILES task parameter is set to TRUE. Otherwise, the automatic SQL tuning reports merely report the recommendation to create a SQL profile.

4. Optionally, implements the SQL profiles provided they meet the criteria of threefold performance improvement

The database considers other factors when deciding whether to implement the SQL profile. For example, the database does not implement a profile when the objects referenced in the statement have stale optimizer statistics. SQL profiles that have been implemented automatically show type is AUTO in the DBA_SQL_PROFILES view.

If the database uses SQL plan management, and if a SQL plan baseline exists for the SQL statement, then the database adds a new plan baseline when creating the SQL profile. As a result, the optimizer uses the new plan immediately after profile creation. See Chapter 15, "Using SQL Plan Management". At any time during or after the automatic SQL tuning process, you can view the results using the automatic SQL tuning report. This report describes in detail all the SQL statements that were analyzed, the recommendations generated, and the SQL profiles that were automatically implemented.



C:\Users\albo\Desktop\1-1.jpg

NEW QUESTION 366

When executing a SQL workload, you choose to generate execution plans only, without collecting execution statistics. Which two statements describe the implications of this? (Choose two.)

- A. It produces less accurate results of the comparison analysis.
- B. It automatically calls the SQL Tuning Advisor for recommendations.
- C. It shortens the time of execution and reduces the impact on system resources.
- D. Only the changes in the execution plan, and not performance regression, are detected.

Answer: AC

NEW QUESTION 370

Which package provides API's for the SQL Tuning Advisor?

- A. DBMS_MONITOR
- B. DBMS_STATS
- C. DBMS_SQLTUNE
- D. DBMS_ADVISOR

Answer: C

Explanation: View the Exhibit 1.

In the CUSTOMERS_OBE table, when the value of CUST_STATE_PROVINCE is "CA", the value of COUNTRY_ID is "US."

View the Exhibit exhibit2 to examine the commands and query plans. The optimizer can sense 8 rows instead of 29 rows, which is the actual number of rows in the table.

What can you do to make the optimizer detect the actual selectivity? desc-customers_obe (exhibit):

```
SQL> DESCRIBE CUSTOMERS_OBE
```

Name	Null?	Type
CUST_ID		NUMBER
CUST_FIRST_NAME	NOT NULL	VARCHAR2 (20)
CUST_LAST_NAME	NOT NULL	VARCHAR2 (40)
CUST_GENDER		CHAR (1)
CUST_YEAR_OF_BIRTH		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_STATE_PROVINCE		VARCHAR2 (40)
COUNTRY_ID	NOT NULL	CHAR (2)
CUST_MAIN_PHONE_NUMBER		VARCHAR2 (25)
CUST_INCOME_LEVEL		VARCHAR2 (30)
CUST_CREDIT_LIMIT		NUMBER
CUST_EMAIL		VARCHAR2 (30)

C:\Users\albo\Desktop\1-1.jpg

sql-select-customers_obe (exhibit):

```
SQL> SELECT COUNT(*) FROM customers_obe WHERE country_id = 'US' AND cust_state_province = 'CA';
```

COUNT(*)
29

```
SQL> EXECUTE dbms_stats.gather_table_stats(null,'customers_obe', method_opt => 'for all columns size 1');
```

```
SQL> EXPLAIN PLAN FOR SELECT * FROM customers_obe WHERE country_id = 'US' AND cust_state_province = 'CA';
```

Explained.

```
SQL> SELECT plan_table_output FROM TABLE(dbms_xplan.display('plan_table',null,'BASIC ROWS'));
```

PLAN_TABLE_OUTPUT			
Plan hash value: 520139036			
Id	Operation	Name	Rows
0	SELECT STATEMENT		8
1	TABLE ACCESS FULL	CUSTOMERS_OBE	8

8 rows selected.

C:\Users\albo\Desktop\1-1.jpg

- A. Change the STALE_PERCENT value for the CUSTOMERS_OBE table.
- B. Set the STATISTICS_LEVEL parameter to TYPICAL.
- C. Create extended statistics for the CUST_STATE_PROVINCE and CUSTOMERS_OBE columns.
- D. Set the OPTIMIZER_USE_PENDING_STATISTICS parameter to FALSE. Answer: C

Real example refer to here. Determining Single Column Statistics

A good example of correlated or related columns are the country_id and cust_state_province columns in the CUSTOMERS_OBE table. When the value of cust_state_province is 'CA', the value of country_id is 'US'. There is a skew in the data in these two columns, which means the majority of rows in the table have the values 'CA' and 'US'. Both the relationship between the columns and the skew in the data make it difficult for the optimizer to calculate the selectivity or cardinality of these columns correctly when they are used together in a query. Extended statistics should help in this situation.

NEW QUESTION 371

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 1Z0-053 Exam with Our Prep Materials Via below:

<https://www.certleader.com/1Z0-053-dumps.html>