



Oracle

Exam Questions 1Z0-053

Oracle Database 11g: Administration II

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

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 3

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 4

Case-insensitive sorts are always accent insensitive by default.

- A. True
- B. False

Answer: B

NEW QUESTION 5

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 6

What elements of globalization can be explicitly defined using the NLS_LANG environment variable? (Choose all that apply.)

- A. NLS_LANGUAGE
- B. NLS_SORT
- C. NLS_CALENDAR
- D. NLS_CHARACTERSET
- E. NLS_TERRITORY

Answer: ADE

NEW QUESTION 7

The NLS_SORT parameter sets the default sort method for which of the following operations? (Choose all that apply.)

- A. WHERE clause
- B. ORDER BY clause
- C. BETWEEN clause
- D. NLSSORT function
- E. NLS_SORT function

Answer: AD

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

Globalization support is implemented through the text- and character-processing functions provided by which Oracle feature?

- A. RSTLNE
- B. NLSRTL
- C. LISTENER
- D. NLSSORT
- E. Linguistic sorts

Answer: B

Explanation: NLSRTL: NLS Runtime Type Library.

NEW QUESTION 10

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 13

Which of the following are valid settings for the NLS_COMP parameter? (Choose all that apply.)

- A. ASCII
- B. ANSI
- C. BINARY
- D. MONOLINGUAL
- E. MULTILINGUAL

Answer: BC

Explanation:

Parameter type	String
Syntax	NLS_COMP = { BINARY LINGUISTIC ANSI }
Default value	BINARY
Modifiable	ALTER SESSION
Basic	No

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

NLS_COMP specifies the collation behavior of the database session. Values:

BINARY

Normally, comparisons in the WHERE clause and in PL/SQL blocks is binary unless you specify the NLSSORT function.

LINGUISTIC

Comparisons for all SQL operations in the WHERE clause and in PL/SQL blocks should use the linguistic sort specified in the NLS_SORT parameter. To improve the performance, you can also define a linguistic index on the column for which you want linguistic comparisons.

ANSI

A setting of ANSI is for backwards compatibility; in general, you should set NLS_COMP to LINGUISTIC

NEW QUESTION 18

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 21

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 25

Which view shows all valid values for the NLS_LANGUAGE, NLS_SORT, NLS_TERRITORY, and NLS_CHARACTERSET parameters?

- A. V\$VALID_NLS_VALUES
- B. NLS_VALID_VALUES
- C. NLS_VALUE_OPTIONS
- D. V\$NLS_VALUE_OPTIONS
- E. V\$NLS_VALID_VALUES

Answer: E

NEW QUESTION 26

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 27

Which NLS parameter can be used to change the default Oracle sort method from binary to linguistic for the SQL SELECT statement?

- A. NLS_LANG
- B. NLS_COMP
- C. NLS_SORT
- D. None of the above

Answer: D

Explanation: Neither NLS_SORT nor NLS_COMP cannot change sorting from BINARY to LINGUISTIC. Because the NLS_SORT depends on NLS_COMP, if they are different, the sorting method will always be BINARY.

(Refer to NLS_SORT) The exact operators and query clauses that obey the NLS_SORT parameter depend on the value of the NLS_COMP parameter. If an operator or clause does not obey the NLS_SORT value, as determined by NLS_COMP, the collation used is BINARY.

The BINARY comparison is faster and uses less resources than any linguistic comparison but for text in a natural language, it does not provide ordering expected by users.

The value of NLS_SORT affects execution plans of queries. Because a standard index cannot be used as a source of values sorted in a linguistic order, an explicit sort operation must usually be performed instead of an index range scan. A functional index on the NLSSORT function may be defined to provide values sorted in a linguistic order and reintroduce the index range scan to the execution plan.

NLS_COMP specifies the collation behavior of the database session.

NEW QUESTION 31

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 33

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 jog 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 38

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;
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..41 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;
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 43

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 48

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 49

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 52

Consider the following code snippet:

```
BEGIN
DBMS_SCHEDULER.SET_ATTRIBUTE (
name => 'lne_job1',
attribute => 'job_priority',
value => 1);
END;
/
```

If this code were executed, which of the following statements would be true?

- A. The priority of the `lne_job1` job would be set to 1.
- B. The `lne_job1` job would be executed synchronously.
- C. The `lne_job1` job would run immediately in the users current session.
- D. The `lne_job1` job would retain its current priority.
- E. The job will immediately take priority over all running jobs.

Answer: A

Explanation: `job_priority`

This attribute specifies the priority of this job relative to other jobs in the same class as this job. If multiple jobs within a class are scheduled to be executed at the same time, the job priority determines the order in which jobs from that class are picked up for execution by the job coordinator. It can be a value from 1 through 5, with 1 being the first to be picked up for job execution.

If no job priority is specified when creating a job, the default priority of 3 is assigned to it.

NEW QUESTION 56

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 59

How many different calendars does Oracle 11g support?

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

Answer: B

NEW QUESTION 61

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 65

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 66

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 69

Which Scheduler view(s) can be queried to see which jobs are currently executing? (Choose all that apply.)

- A. DBA_SCHEDULER_JOB_RUN_DETAILS
- B. DBA_SCHEDULER_RUNNING_JOBS
- C. DBA_SCHEDULER_CURRENT_JOBS
- D. DBA_SCHEDULER_JOBS
- E. DBA_SCHEDULER_EXECUTING_JOBS

Answer: BD

Explanation: DBA_SCHEDULER_JOB_RUN_DETAILS displays log run details for all Scheduler jobs in the database. DBA_SCHEDULER_RUNNING_JOBS displays information about all running Scheduler jobs in the database.

DBA_SCHEDULER_JOBS displays information about all Scheduler jobs in the database. DBA_SCHEDULER_CURRENT_JOBS is INVALID. DBA_SCHEDULER_EXECUTING_JOBS is INVALID.

NEW QUESTION 70

Which of the following Scheduler elements encourage object reuse? (Choose all that apply.)

- A. Schedule objects
- B. Program arguments
- C. Job classes
- D. Job arguments
- E. All of the above

Answer: ABD

Explanation: Don't really get the exact source of answer, just keep it in memory with common sense.

NEW QUESTION 74

The user SYS creates a job by using the following command:

```
BEGIN DBMS_SCHEDULER.CREATE_JOB (  
  job_name => 'update_sales',  
  job_type => 'STORED-PROCEDURE',  
  job_action => 'OPS.SALES_PKG.UPDATE_SALES_SUMMARY',  
  start_date => '28-DEC-07 07.00.00 PM Australia/Sydney', repeat_interval => 'FREQ=DAILY; INTERVAL=2', end_date => '20-JAN-08 07.00.00 PM  
Australia/Sydney', comments => 'New sales job');  
END;  
/
```

Which two statements are true about the job that was created by the preceding command?
(Choose two.)

- A. The job is enabled by default after creation
- B. The job is automatically dropped after the end date
- C. The job executes with the privileges of the user SYS
- D. The globalization environment that exists at the time of the job creation prevails at the job runs

Answer: BC

Explanation: By default, jobs are created with auto_drop set to TRUE. 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 75

Which two statements correctly describe the relationship among the Scheduler components: job, program, and schedule? (Choose two)

- A. A job is specified as part of a program definition
- B. A program can be used in the definition of multiple jobs
- C. A program and job can be specified as part of a schedule definition
- D. A program and schedule can be specified as part of a job definition

Answer: BD

Explanation: Programs

A program object (program) describes what is to be run by the Scheduler. A program includes:

? An action: For example, the name of a stored procedure, the name of an executable found in the operating system file system (an "external executable"), or the text of a PL/SQL anonymous block.

? A type: STORED_PROCEDURE, PLSQL_BLOCK, or EXTERNAL, where EXTERNAL indicates an external executable.

? Number of arguments: The number of arguments that the stored procedure or external executable accepts.

A program is a separate entity from a job. A job runs at a certain time or because a certain event occurred, and invokes a certain program. You can create jobs that point to existing program objects, which means that different jobs can use the same program and run the program at different times and with different settings. With the right privileges, different users can use the same program without having to redefine it. Therefore, you can create program libraries, where users can select from a list of existing programs.

If a stored procedure or external executable referenced by the program accepts arguments, you define these arguments in a separate step after creating the program. You can optionally define a default value for each argument.

Oracle Scheduler

Program objects (programs) contain metadata about the command that the Scheduler will run, including default values for any arguments. Schedule objects (schedules) contain information about run date and time and recurrence patterns. Job objects (jobs) associate a program with a schedule. To define what is executed and when, you assign relationships among programs, schedules, and jobs.

A job is the combination of a schedule and a program, along with any additional arguments required by the program.

NEW QUESTION 77

Which of the following are not disabled by default?

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

Answer: E

NEW QUESTION 82

Which of the following tasks is not performed by the job coordinator?

- A. Update job log when a job completes
- B. Spawn and remove job slaves
- C. Write/read job info to/from memory cache
- D. Query job table
- E. Pass job information to job slaves

Answer: A

NEW QUESTION 87

Which three components does the Scheduler use for managing tasks within the Oracle environment? (Choose three.)

- A. a job
- B. a program
- C. a schedule
- D. a PL/SQL procedure

Answer: ABC

Explanation: The Scheduler objects include:

- ? Programs
- ? Schedules
- ? Jobs
- ? Destinations
- ? Chains
- ? File Watchers
- ? Credentials
- ? Job Classes
- ? Windows
- ? Groups

NEW QUESTION 91

Which three statements are true regarding the functioning of the Autotask Background Process (ABP)? (Choose three.)

- A. It creates jobs without considering the priorities associated with them.
- B. It translates tasks into jobs for execution by the scheduler.
- C. It determines the list of jobs that must be created for each maintenance window.
- D. It is spawned by the MMON background process at the start of the maintenance window.
- E. It maintains a repository in the SYSTEM tablespace to store the history of the execution of all tasks.

Answer: BCD

Explanation: AutoTask Overview

Oracle 11g added a new component to the Oracle automated task infrastructure, that began in Oracle 10g, called AutoTask. AutoTask provides a central component that is responsible for management of scheduled maintenance tasks.

AutoTask Architecture

AutoTask builds on much of the existing architecture of Oracle 10g. The AutoTask architecture consists of the following components:

AutoTask Background Process (ABP) - APB is spawned by MMON and is responsible for managing the AutoTask features. It coordinates the AutoTask clients and the scheduler. It also maintains AutoTask related history that can be seen in the DBA_AUTOTASK_TASK view.

AutoTask Clients - Automated maintenance tasks that are scheduled by AutoTask. Oracle 11g includes clients for statistics gathering, the Segment Advisor, and the Automatic SQL Tuning Advisor.

The AutoTask Maintenance Windows - Individual maintenance windows exist for different days of the week.

These windows can be modified as required.

Resource Manager - A resource plan is enabled that is designed to limit the amount of resources that the AutoTask tasks can consume. This resource plan can be modified as required.

OEM - You can manage the start time and duration of the various AutoTask maintenance windows and add or remove maintenance tasks with OEM. You can also use OEM to enable or disable maintenance tasks.

The Scheduler - AutoTask operations rely heavily on the Scheduler. The AutoTask architecture uses scheduler windows and the scheduler infrastructure to execute AutoTask tasks.

NEW QUESTION 96

You notice that a job in a chain has not completed on a nonconstrained RAC database. Which of these are valid reasons why that might occur?

- A. The job priority is 1 and the resource consumer group CPU emphasis allocation is a low percentage.
- B. The job affinity is to a service and one node in that service is unavailable.
- C. The job affinity is to an instance and that instance is unavailable.
- D. There is no service affinity.
- E. None of the above.

Answer: C

NEW QUESTION 99

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 104

Which DBMS_SCHEDULER procedures can be used to enable a program? (Choose all that apply.)

- A. ENABLE
- B. ENABLE_PROGRAM
- C. VALIDATE_PROGRAM

- D. SET_ATTRIBUTE
- E. SET_ENABLED

Answer: AD

Explanation: Refer to here.

ENABLE Procedure, Enables a program, job, chain, window, or window group SET_ATTRIBUTE Procedure, Changes an attribute of a job, schedule, or other Scheduler object

NEW QUESTION 109

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 110

Which DBMS_RESOURCE_MANAGER procedure prioritizes consumer-group mappings?

- A. CREATE_MAPPING_PRIORITY
- B. SET_MAPPING_PRIORITY
- C. SET_MAPPING_ORDER
- D. PRIORITIZE_MAPPING_ORDER
- E. This functionality is not available through the DBMS_RESOURCE_MANAGER package.

Answer: B

NEW QUESTION 112

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 116

You upgraded Oracle Database 10g to Oracle Database 11g. How would this affect the existing users' passwords?

- A. All passwords automatically become case-sensitive.
- B. All passwords remain non-case-sensitive till they are changed.
- C. All passwords remain non-case-sensitive and cannot be changed.
- D. All passwords remain non-case-sensitive until their password attribute in the profile is altered.

Answer: B

NEW QUESTION 119

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 122

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 126

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 128

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 129

Which statements describe the capabilities of the DBMS_NETWORK_ACL_ADMIN package? (Choose all that apply.)

- A. It can be used to allow the access privilege settings for users but not roles.
- B. It can be used to allow the access privilege settings for users as well as roles.
- C. It can be used to control the time interval for which the access privilege is available to a user.
- D. It can be used to selectively restrict the access for each user in a database to different host computers.
- E. It can be used to selectively restrict a user's access to different applications in a specific host computer.

Answer: BCD

NEW QUESTION 132

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

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 136

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 139

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 140

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 141

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 143

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 147

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 type="radio"/>	SYSAUX	PERMANENT	LOCAL	AUTO	ONLINE	330.000	322.000
<input type="radio"/>	SYSTEM	PERMANENT	LOCAL	MANUAL	ONLINE	470.000	462.625
<input type="radio"/>	TEMP	TEMPORARY	LOCAL	MANUAL	ONLINE	20.000	6.000
<input type="radio"/>	TT	PERMANENT	LOCAL	MANUAL	ONLINE	0.102	.102
<input type="radio"/>	UNDOTBS1	UNDO	LOCAL	MANUAL	ONLINE	90.000	15.813
<input 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 149

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 154

Your database initialization parameter file has the following entry: SEC_MAX_FAILED_LOGIN_ATTEMPTS=3
Which statement is true regarding this setting?

- A. It drops the connection after the specified number of login attempts fail for any user.
- B. It is enforced only if the password profile is enabled for the user.
- C. It locks the user account after the specified number of attempts.
- D. It drops the connection after the specified number of login attempts fail only for users who have the SYSDBA privilege.

Answer: A

NEW QUESTION 158

Which of the following describes how a distributed resumable transaction behaves?

- A. The resumable setting on the initiating session determines the resumable conditions for the entire distributed transaction.
- B. The resumable setting for the initiating instance determines the resumable conditions for the entire distributed transaction.
- C. The resumable setting on the initiating session controls only that part of the transaction that occurs within the local instance; remote resumable settings determine the behavior of the distributed parts of the transaction.
- D. None of the above.

Answer: C

NEW QUESTION 160

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 164

You installed Oracle Database 11g at fresh. Which statements are true regarding the default audit settings in this database? (Choose all that apply.)

- A. The audit trail is stored in an operating system file.
- B. Auditing is disabled for all privileges.
- C. The audit trail is stored in the database.
- D. Auditing is enabled for all privileges.
- E. Auditing is enabled for certain privileges related to database security.

Answer: CE

NEW QUESTION 165

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 166

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 168

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 171

Evaluate the following command and its output: SQL>SELECT * FROM dba_temp_free_space;

```
TABLESPACE_NAME TABLESPACE_SIZE ALLOCATED_SPACE 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 173

Which statement is true regarding online redefinition for the migration of BasicFile LOBs to SecureFile LOBs?

- A. It cannot be done in parallel.
- B. It can be done at the table level or partition level.
- C. It does not require additional storage because the operation is done online.
- D. Local and global indexes are maintained automatically during the operation.

Answer: B

Explanation: Refer to here

Compatibility and Upgrading

All features described in this document are enabled with compatibility set to 11.2.0.0.0 or higher. There is no downgrade capability after 11.2.0.0.0 is set.

If you want to upgrade BasicFiles LOBs to SecureFiles LOBs, you must use typical methods for upgrading data (CTAS/ITAS, online redefinition, export/import, column to column copy, or using a view and a new column). Most of these solutions require twice the disk space used by the data in the input LOB column.

However, partitioning and taking these actions on a partition-by-partition basis lowers the disk space requirements. Migrating Columns from BasicFiles LOBs to SecureFiles LOBs

The method of migrating LOBs columns is presented in this section. Preventing Generation of REDO Space when Migrating to SecureFiles LOBs

Generation of redo space can cause performance problems during the process of migrating BasicFiles LOB columns. Redo changes for the table are logged during the migration process only if the table has LOGGING set.

Redo changes for the column being converted from BasicFiles LOB to SecureFiles LOB are logged only if the storage characteristics of the SecureFiles LOB column indicate LOGGING. The logging setting (LOGGING or NOLOGGING) for the LOB column is inherited from the tablespace in which the LOB is created.

To prevent generation of redo space during migration make sure that you specify the NOLOGGING storage parameter for the new SecureFiles LOB column(s).

You can turn LOGGING on once your migration is complete.

Online Redefinition for BasicFiles LOBs

Online redefinition is the only recommended method for migration of BasicFiles LOBs to SecureFiles LOBs. It can be done at the table or partition level.

Online Redefinition Advantages

? No requirement to take the table or partition offline

? Can be done in parallel

Online Redefinition Disadvantages

? Additional storage equal to the entire table or partition and all LOB segments must be available

? Global indexes must be rebuilt

NEW QUESTION 176

Evaluate the following command:

```
SQL> CREATE TABLE design_data (id NUMBER, doc CLOB)
```

```
LOB(doc) STORE AS SECUREFILE(DEDUPLICATE);
```

Which statement is true regarding the above command?

- A. The LOB values are automatically compressed.
- B. The LOB values are cached by default in the buffer cache.
- C. The LOB values are automatically stored in encrypted mode.
- D. All LOB data that is identical in two or more rows in a LOB column share the same data blocks.

Answer: D

Explanation: DEDUPLICATE or KEEP_DUPLICATES (Link)

The option DEDUPLICATE enables you to specify that LOB data which is identical in two or more rows in a LOB column should share the same data blocks. The database combines LOBs with identical content into a single copy, reducing storage and simplifying storage management. The opposite of this option is KEEP_DUPLICATES.

NEW QUESTION 178

Which of the following options describes Segment Shrink?

- A. Reclaims space above and below the high-water mark without using additional space
- B. Moves rows to a new physical location, resetting the high-water mark, but uses additional space during the operation
- C. Deallocates space above the high-water mark that is currently not in use
- D. None of the above

Answer: A

NEW QUESTION 179

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 183

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.

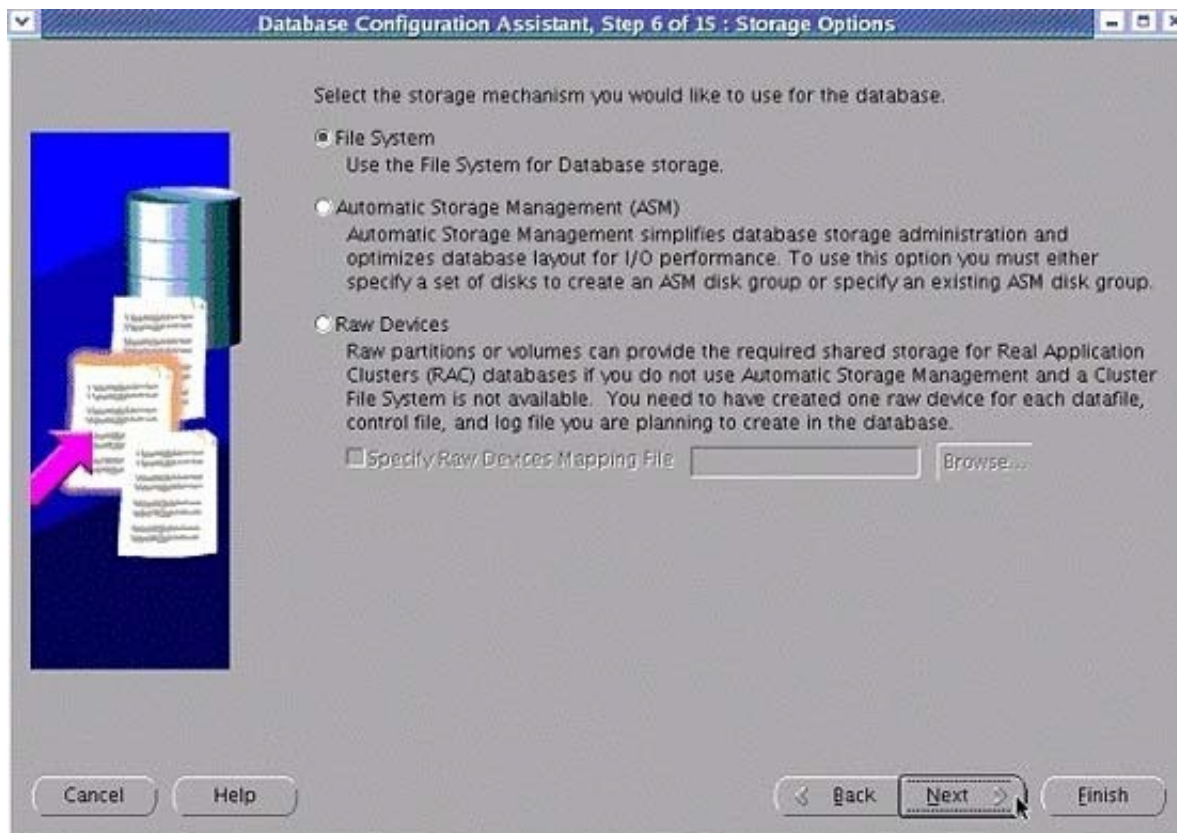
Answer: C

NEW QUESTION 188

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 191

Sales details are being stored on a daily basis in the SALES_2007 table. A large amount of data is added to the table daily. To save disk space, you issued the following command:

```
ALTER TABLE sales_2007 COMPRESS FOR ALL OPERATIONS;
```

What would be the outcome of this command?

- A. It produces an error because data already exists in the table.
- B. It produces an error because compression can be enabled at table creation only.
- C. It compresses all data added or modified henceforth but the existing data in the table is not compressed immediately.
- D. It immediately compresses all existing data as well as new data, resulting from either fresh additions or modifications to existing data.

Answer: C

NEW QUESTION 196

You opened the encryption wallet and then issued the following command:

```
SQL>CREATE TABLESPACE securespace DATAFILE '/home/user/oradata/secure01.dbf' SIZE 150M  
ENCRYPTION USING '3DES168' DEFAULT STORAGE(ENCRYPT);
```

Then you closed the wallet. Later, you issued the following command to create the EMPLOYEES table in the SECURESPACE tablespace and you use the NO SALT option for the EMPID column.

What is the outcome?

- A. It creates the table and encrypts the data in it.
- B. It generates an error because the wallet is closed.
- C. It creates the table but does not encrypt the data in it.
- D. It generates an error because the NO SALT option cannot be used with the ENCRYPT option.

Answer: B

NEW QUESTION 198

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 201

You executed the following SQL statement to shrink the EMPLOYEES table segment stored in the EXAMPLE tablespace:

```
ALTER TABLE employees SHRINK SPACE CASCADE;
```

Which statement is correct in this scenario?

- A. The EMPLOYEES table will be changed to read-only mode during the shrink operation.
- B. The indexes created on the EMPLOYEES table will need to be rebuilt after the shrink operation is over.

- C. The shrink behavior will be cascaded to all dependent segments of the table that support a shrink operation.
- D. Data manipulation language (DML) operations will not be possible on the EMPLOYEES table during the COMPACTION phase of the shrink operation.

Answer: C

NEW QUESTION 204

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 206

Which of the following supplied functions is used to identify external tables, directories, and BFILES?

- A. DBMS_TDB.CHECK_DIRECTORIES
- B. DBMS_TDB.CHECK_EXTERNAL
- C. DBMS_TDB.CHECK_BFILE
- D. DBMS_TDB.CHECK_EXT

Answer: B

NEW QUESTION 207

Which statement describes the effect of table redefinition on the triggers attached to the table?

- A. All triggers on the table are invalidated and are automatically revalidated with the next DML execution on the table.
- B. All triggers on the table are invalidated and must be manually recompiled before the next DML execution on the table.
- C. All triggers on the table remain valid.
- D. Only triggers that are affected by the changes to the structure of the table are invalidated and automatically revalidated with the next DML execution on the table.

Answer: A

NEW QUESTION 212

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 213

You want to enable resumable space allocation at the instance level.

Which two actions would enable resumable space allocation at the instance level? (Choose two.)

- A. issuing the ALTER SYSTEM ENABLE RESUMABLE; statement
- B. issuing the ALTER SESSION ENABLE RESUMABLE; statement
- C. modifying the RESUMABLE_TIMEOUT initialization parameter to a nonzero value
- D. issuing the ALTER SYSTEM SET RESUMABLE_TIMEOUT=<nonzero value>; statement

Answer: CD

NEW QUESTION 218

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 223

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 226

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 228

Which statement about the enabling of table compression in Oracle Database 11g is true?

- A. Compression can be enabled at the table, tablespace, or partition level for direct loads only.
- B. Compression can be enabled only at the table level for both direct loads and conventional DML.
- C. Compression can be enabled at the table, tablespace, or partition level for conventional DML only.
- D. Compression can be enabled at the table, tablespace, or partition level for both direct loads and conventional DML.

Answer: D

NEW QUESTION 232

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 234

The BOOKINGS table contains online booking information. When a booking is confirmed, the details are transferred to an archival table BOOKINGS_HIST and deleted from the

BOOKINGS table. There is no fixed time interval between each online booking and its confirmation. Because sufficient space is not always available from the delete operations the high-water mark (HWM) is moved up and many rows are inserted below the HWM of the table. The BOOKINGS table has Automatic Segment Space Management (ASSM) and row movement enabled. The table is accessible in 24x7 mode.

What is the most efficient method to reclaim the space released by the delete operations in the BOOKINGS table?

- A. Perform EXPORT, DROP, and IMPORT operations on the BOOKINGS table sequentially
- B. Shrink the BOOKINGS table by using the ALTER TABLE... SHRINK SPACE command
- C. Move the BOOKINGS table to a different location by using the ALTER TABLE... MOVE command
- D. Deallocate the space in the BOOKINGS table by using the ALTER TABLE ... DEALLOCATE UNUSED command

Answer: B

NEW QUESTION 235

Examine the following command:

```
SQL> ALTER TABLE booking SHRINK SPACE COMPACT;
```

Which activity is performed when the preceding command is executed?

- A. The shrink operation touches every block in the BOOKING table
- B. The high-water mark (HWM) for the BOOKING table is shifted from its original position
- C. The progress of the shrink operation is saved in the bitmap blocks of the BOOKING table
- D. The data manipulation language (DML) triggers on the BOOKING table are executed because the shrink operation is internally handled by the INSERT/DELETE operation

Answer: C

NEW QUESTION 240

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 241

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 246

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 250

A PL/SQL procedure queries only those columns of a redefined table that were unchanged by the online table redefinition.

What happens to the PL/SQL procedure after the online table redefinition?

- A. It remains valid.
- B. It becomes invalid for all options of online table redefinition but automatically gets revalidated the next time it is used.
- C. It becomes invalid for all options of online table redefinition and is automatically recompiled during online redefinition of the table.
- D. It becomes invalid only if the storage parameters have been modified and it automatically gets revalidated the next time it is used.

Answer: A

NEW QUESTION 254

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 256

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 RESUMABLE TIMEOUT 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 258

Examine the following PL/SQL block:

DECLARE

my_plans PLS_INTEGER; BEGIN

my_plans := DBMS_SPM.LOAD_PLANS_FROM_CURSOR_CACHE(sql_id => '99twu5t2dn5xd');

END;

/

Which statement is true about the plan being loaded into the SQL plan baseline by the above command?

- A. It is loaded with the FIXED status.
- B. It is loaded with the ACCEPTED status.
- C. It is not loaded with the ENABLED status.
- D. It is not loaded with the ACCEPTED status.

Answer: B

Explanation: You can create SQL plan baselines by manually loading existing plans for a set of SQL statements as plan baselines. The database does not verify manually loaded plans for performance, but adds them as accepted plans to existing or new SQL plan baselines. (Refer to here 15.2.1.2 Creating Baselines from Existing Plans)

The DBMS_SPM package supports the SQL plan management feature by providing an interface for the DBA or other user to perform controlled manipulation of plan history and SQL plan baselines maintained for various SQL statements. `LOAD_PLANS_FROM_CURSOR_CACHE` Functions, This function loads one or more plans present in the cursor cache for a SQL statement, or a set of SQL statements. It has four overloads: using SQL statement text, using SQL handle, using SQL ID, or using attribute_name and attribute_value pair.

By default, the parameter of `LOAD_PLANS_FROM_CURSOR_CACHE` has:

1. fixed = 'NO', means loaded with non-fixed plans.
2. enabled = 'YES', means that it is enabled for use by the optimizer. Refer to here for details.

NEW QUESTION 263

You work in a data warehouse environment that involves the execution of complex queries. The current content of the SQL cache contains the ideal workload for analysis. You want to analyze only a few most resource-intensive statements.

What must you do to receive recommendations on efficient use of indexes and materialized views to improve query performance?

- A. Run the SQL Access Advisor.
- B. Run the SQL Tuning Advisor (STA).
- C. Run the Automatic Workload Repository (AWR) report.
- D. Run the Automatic Database Diagnostic Monitor (ADDM).

Answer: A

Explanation: In addition, SQL Access Advisor can recommend partitioning on an existing unpartitioned base table to improve performance. Furthermore, it may recommend new indexes and materialized views that are themselves partitioned. While creating new partitioned indexes and materialized view is no different from the unpartitioned case, partitioning existing base tables should be executed with care. This is especially true when indexes, views, constraints, or triggers are defined on the table. See "Special Considerations when Script Includes Partitioning Recommendations" for a list of issues involving base table partitioning for performing this task online.

NEW QUESTION 266

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 268

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 269

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 274

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 276

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 278

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 280

Which of these recommendations should be followed before capturing a workload? (Choose all that apply.)

- A. Make sure your replay database has the same structure as the capture database, except without data.
- B. Make sure the replay and capture databases are similar in data content.
- C. Perform a clean shutdown and restart of the capture database before beginning a workload capture.
- D. Start the capture database in UNRESTRICTED mode, then start the capture.
- E. Define inclusion and exclusion filters.

Answer: BC

NEW QUESTION 283

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 285

View the Exhibit.

Examine the following command that is executed for the TRANSPORT table in the SH schema:

```
SQL> SELECT DBMS_STATS.CREATE_EXTENDED_STATS('sh', 'customers_obe',  
'(country_id, cust_state_province)') FROM dual;
```

Which statement describes the significance of this command? 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)
CUST_RY_ID	NOT NULL	CHAR(2)
CUST_MAIN_PHONE_NUMBER		VARCHAR2(25)
CUST_INCOME_LEVEL		VARCHAR2(30)
CUST_CREDIT_LIMIT		NUMBER
CUST_EMAIL		VARCHAR2(30)

- A. It collects statistics into the pending area in the data dictionary.
- B. It creates a virtual hidden column in the CUSTOMERS_OBE table.
- C. It collects statistics with AUTO_SAMPLE_SIZE for ESTIMATE_PERCENT.
- D. It creates a histogram to hold skewed information about the data in the columns.

Answer: B

Explanation: DBMS_STATS.CREATE_EXTENDED_STATS Function

Creates a virtual column for a user specified column group or an expression in a table This function creates a column statistics entry in the system for a user specified column group or an expression in a table. Statistics for this extension will be gathered when user or auto statistics gathering job gathers statistics for the table. We call statistics for such an extension, "extended statistics". This function returns the name of this newly created entry for the extension.

NEW QUESTION 286

Which three functions are performed by the SQL Tuning Advisor? (Choose three.)

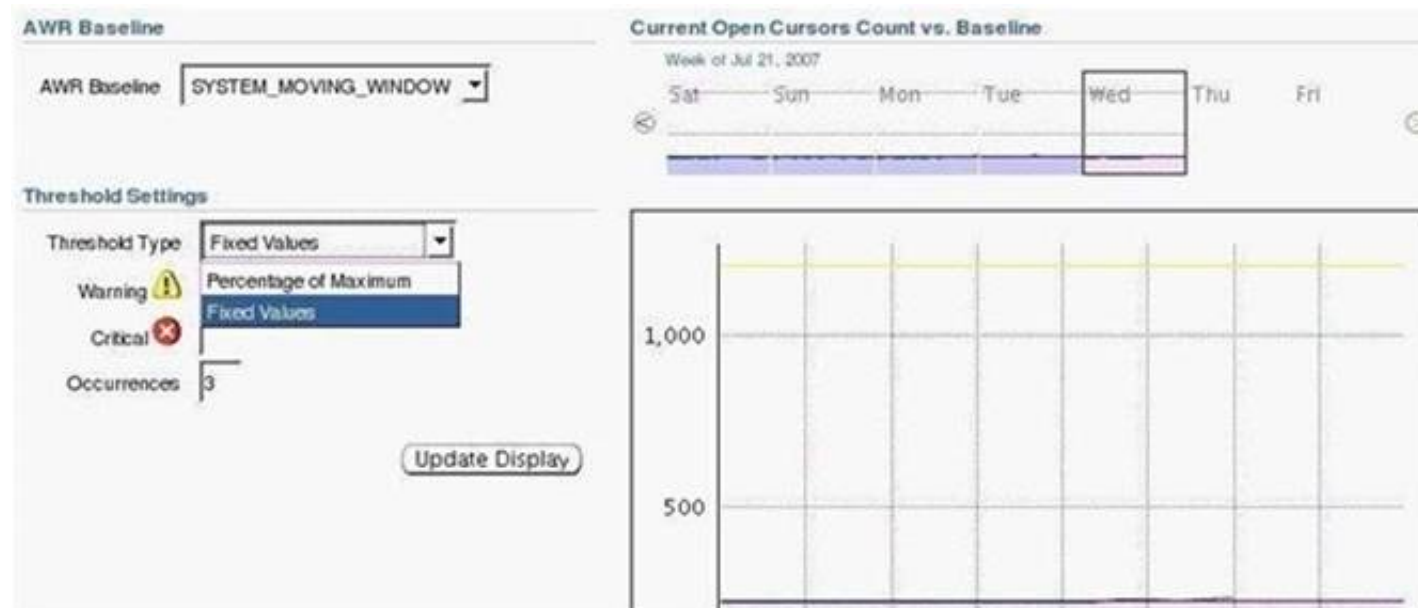
- A. Building the SQL profile
- B. Recommending optimization of materialized views
- C. Checking query objects for missing and stale statistics
- D. Recommending bitmap, function-based, and B-tree indexes
- E. Recommending restructuring SQL queries that are using bad plans

Answer: ACE

NEW QUESTION 288

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: locp > 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
 Sat Sun Mon Tue Wed

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

NEW QUESTION 290

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 293

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 294

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 296

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 298

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 299

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 302

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_6fa2	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 304

View the Exhibit to examine the replay settings for replay parameters.
What is the implication for setting the values for replay parameters? (Choose all that apply.) Exhibit:

ORACLE Enterprise Manager 11g Database Control

Set Up Workload Replay: Customize Options

Database: orcl

Capture Name:

Connections | Replay Parameters

Some replay parameters can be modified to change the behavior of the replay.

Name	Value
synchronization	TRUE
connect_time_scale	100
think_time_scale	100
think_time_auto_correct	TRUE

- A. The COMMIT order in the captured workload is preserved during replay.
- B. The value 100 in the THINK_TIME_SCALE parameter attempts to make the replay client shorten the think time between calls.
- C. The value 100 in the CONNECT_TIME_SCALE parameter attempts to connect all sessions as captured.
- D. The value 100 in the THINK_TIME_SCALE parameter attempts to match the captured user think time while replaying.
- E. The value 100 in the CONNECT_TIME_SCALE parameter attempts to connect all sessions immediately as soon as the replay begins.

Answer: ACD

NEW QUESTION 306

Which two client requests are captured during database replay Capture? (Choose two)

- A. Flashback queries
- B. Shared server requests
- C. Login and logoff activities of sessions
- D. Direct path load of data from external files by using utilities such as SQL *loader
- E. Data definition language (DDL) and data manipulation language (DML) operations

Answer: CE

Explanation: 9.3 Workload Capture Restrictions

Certain types of user sessions and client requests may sometimes be captured in a workload, but they are not supported by Database Replay. Capturing these session and request types in a workload may result in errors during workload replay.

The following types of user sessions and client requests are not supported by Database Replay:

- ? Direct path load of data from external files using utilities such as SQL*Loader
 - ? Non-PL/SQL based Advanced Queuing (AQ)
 - ? Flashback queries
 - ? Oracle Call Interface (OCI) based object navigations
 - ? Non SQL-based object access
 - ? Distributed transactions, Any distributed transactions that are captured will be replayed as local transactions.
 - ? XA transactions, XA transactions are not captured or replayed. All local transactions are captured.
 - ? JAVA_XA transactions, If the workload uses the JAVA_XA package, JAVA_XA function and procedure calls are captured as normal PL/SQL workload. To avoid problems during workload replay, consider dropping the JAVA_XA package on the replay system to enable the replay to complete successfully.
 - ? Database Resident Connection Pooling (DRCP)
 - ? Workloads using OUT binds
 - ? Multi-threaded Server (MTS) and shared server sessions with synchronization mode set to OBJECT_ID
 - ? Migrated sessions, The workload is captured for migrated sessions. However, user logins or session migration operations are not captured. Without a valid user login or session migration, the replay may cause errors because the workload may be replayed by a wrong user.
- Typically, Database Replay refrains from capturing these types of non-supported user sessions and client requests. Even when they are captured, Database Replay will not replay them. Therefore, it is usually not necessary to manually filter out non-supported user sessions and client requests. In cases where they are captured and found to cause errors during replay, consider using workload capture filters to exclude them from the workload.

NEW QUESTION 309

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 314

Which statement most accurately describes the implementation of a SQL Access Advisor recommendation?

- A. SQL Access Advisor recommendations are automatically implemented.
- B. Individual SQL Access Advisor recommendations can be scheduled for implementation.
- C. All SQL Access Advisor recommendations for a specific task must be implemented at the same time.
- D. SQL Access Advisor recommendations are automatically scheduled for implementation during the maintenance window.
- E. None of the above.

Answer: B

Explanation: Implementing the SQL Access Advisor Recommendations (link)

A SQL Access Advisor recommendation can range from a simple suggestion to a complex solution that requires partitioning a set of existing base tables and implementing a set of database objects such as indexes, materialized views, and materialized view logs. You can select the recommendations for implementation and schedule when the job should be executed.

NEW QUESTION 318

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 323

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

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

Answer: AC

Explanation: The AWR snapshots needs to be loaded to STS at first, then load to SQL plan.

----- (Link)

You can perform manual plan loading by:

Loading Plans from SQL Tuning Sets and AWR Snapshots

To load plans from a SQL tuning set, use the `LOAD_PLANS_FROM_SQLSET` function of the `DBMS_SPM` package.

The following example loads the plans stored in the SQL tuning set named `tset1`:

```
DECLARE
my_plans PLS_INTEGER; BEGIN
my_plans := DBMS_SPM.LOAD_PLANS_FROM_SQLSET( sqlset_name => 'tset1'); END;
/
```

To load plans from Automatic Workload Repository (AWR), load the plans stored in AWR snapshots into a SQL tuning set before using the `LOAD_PLANS_FROM_SQLSET` function as described in this section.

Loading Plans from the Shared SQL Area

To load plans from the shared SQL area, use the `LOAD_PLANS_FROM_CURSOR_CACHE` function of the `DBMS_SPM` package. In the following example, Oracle Database loads the plans located in the shared SQL area for the SQL statement identified by its `sql_id`:

```
DECLARE
my_plans PLS_INTEGER; BEGIN
my_plans := DBMS_SPM.LOAD_PLANS_FROM_CURSOR_CACHE( sql_id => '99twu5t2dn5xd');
END;
/
```

NEW QUESTION 325

You run the SQL Tuning Advisor (STA) to tune a SQL statement that is part of a fixed SQL plan baseline. The STA generates a SQL profile for the SQL statement, which recommends that you accept the profile.

Which statement is true when you accept the suggested SQL profile?

- A. The tuned plan is not added to the SQL plan baseline.
- B. The tuned plan is added to the fixed SQL plan baseline as a fixed plan.
- C. The tuned plan is added to the fixed SQL plan baseline as a nonfixed plan.
- D. The tuned plan is added to a new nonfixed SQL plan baseline as a nonfixed plan.

Answer: C

Explanation: 15.4 Using Fixed SQL Plan Baselines (Refer to here)

When you tune a SQL statement with a fixed SQL plan baseline using SQL Tuning Advisor, a SQL profile recommendation has special meaning. When the SQL profile is accepted, the database adds the tuned plan to the fixed SQL plan baseline as a non-fixed plan. However, as described above, the optimizer does not use the tuned plan when a reproducible fixed plan is present. Therefore, the benefit of SQL tuning may not be realized. To enable the use of the tuned plan, manually alter the tuned plan to a fixed plan by setting its `FIXED` attribute to `YES`.

NEW QUESTION 327

The Automatic Database Diagnostic Monitor (ADDM) analysis runs every 60 minutes on your database. Your database is facing a series of interrelated problems over a period of two hours. You need to ensure that the ADDM analysis is run over a time span of two hours in future. What would you do?

- A. Create two custom ADDM tasks.
- B. Modify the AWR snapshot time interval to two hours.
- C. Create a new scheduler window for a time period of two hours.
- D. Modify the Automatic Workload Repository (AWR) snapshot retention period to two hours.

Answer: B

NEW QUESTION 329

While tuning a SQL statement, the SQL Tuning Advisor finds an existing SQL profile for the statement that has stale statistics available. What would the optimizer do in this situation?

- A. It updates the existing SQL profiles with current statistics.
- B. It makes the statistics information available to GATHER_STATS_JOB.
- C. It initiates the statistics collection process by running GATHER_STATS_JOB.
- D. It logs a warning message in the alert log so that the DBA can perform statistics collection manually.

Answer: B

NEW QUESTION 331

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 333

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 334

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 337

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.

View Recommendations Previous 1-25 of 73 Next 25

Select	SQL Text	Parsing Schema	SQL ID	Statistics	SQL Profile	Index	Restructure SQL	Miscellaneous	Error	Date
<input type="radio"/>	SELECT NULL AS table_cat, t.owner...	SYSMAN	361qjn3w9uflh	✓	(99.9%) ✓					7/12/07
<input checked="" type="radio"/>	SELECT EXECUTION_ID, STATUS, STATUS_DETA...	SYSMAN	lyk8tb9986ark7		(69%) ✓	(97.9%) ✓				7/12/07
<input type="radio"/>	SELECT /*+ INDEX(sqlobb\$ (signature cate...	SYS	8b75qwpna202v					✓		7/12/07
<input type="radio"/>	select OBJOID, CLSOID, RUNTIME, PRI, JO...	SYS	8vf1dhwgk1xy5					✓		7/12/07
<input type="radio"/>	select smontabv.cnt, smontab.time_mp, ...	SYS	4q8mr2bvy6qr					✓		7/12/07
<input type="radio"/>	select t.ts#,t.file#,t.block#,rvl(t.bobj...	SYS	1qu8t96d0bdmu					✓		7/12/07
<input type="radio"/>	select obj#, dataobj#, part#, hiboundlen...	SYS	130dvvr5s8bqn					✓		7/12/07
<input type="radio"/>	select privilege#,level from sysauth\$ co...	SYS	0b6b2sqmb74n					✓		7/12/07
<input type="radio"/>	select value(p\$) from "XDB"."XDB\$RESOURC...	SYS	23y48d28wkq2r					✓		7/12/07
<input type="radio"/>	SELECT obj_type, plan_id, name, flags, L...	SYS	0n1napsmccz0c					✓		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 340

You plan to set up the Automatic Workload Repository (AWR) baseline metric thresholds for a moving window baseline. Which action would you take before performing this task?

- A. Compute the baseline statistics.
- B. Take an immediate AWR snapshot.
- C. Decrease the window size for the baseline.
- D. Decrease the expiration time for the baseline.

Answer: A

Explanation: A moving window baseline corresponds to all AWR data that exists within the AWR retention period. This is useful when using adaptive thresholds because the database can use AWR data in the entire AWR retention period to compute metric threshold values.

NEW QUESTION 343

In a Database Replay workload capture, what client request information is gathered? (Choose all that apply.)

- A. SQL text
- B. Shared server requests (Oracle MTS)
- C. Bind variable values
- D. Information about transactions
- E. Remote DESCRIBE and COMMIT operations

Answer: ACD

NEW QUESTION 345

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 346

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 347

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 350

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 351

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 355

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

1Z0-053 Practice Exam Features:

- * 1Z0-053 Questions and Answers Updated Frequently
- * 1Z0-053 Practice Questions Verified by Expert Senior Certified Staff
- * 1Z0-053 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * 1Z0-053 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The 1Z0-053 Practice Test Here](#)