

# 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

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 5

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 6

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 7**

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 8**

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 9**

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

- A. LOG\_HISTORY
- B. JOB\_LOG\_RETENTION
- C. WINDOW\_LOG\_RETENTION
- D. WHICH\_LOG
- E. LOG\_NAME

**Answer:** AD

**Explanation:** -----

Attributes list of SET\_SCHEDULE\_ATTRIBUTE:

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

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

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

NONE: the default, indicating no encryption used

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

'event\_expiry\_time': The time, in seconds, before a job state event generated by the Scheduler expires from the Scheduler event queue. If NULL, job state events

expire after 24 hours.

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

The PURGE\_LOG Procedure uses "WHICH\_LOG" and "LOG\_NAME" attributes. Syntax

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

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

log\_history

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

which\_log

This specifies the log type. Valid values are: job\_log, window\_log, and job\_and\_window\_log.

job\_name

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

#### NEW QUESTION 10

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

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

**Answer:** ADE

#### NEW QUESTION 10

Which three statements are true regarding persistent lightweight jobs? (Choose three.)

- A. Persistent lightweight jobs modify several tables in the data dictionary to generate a lot of redo.
- B. The user cannot set privileges on persistent lightweight jobs.
- C. Persistent lightweight jobs are useful when users need to create a large number of jobs in a short time.
- D. Persistent lightweight jobs are useful when users need to create a small number of jobs that run infrequently.
- E. The use of a template is mandatory to create persistent lightweight jobs.

**Answer:** BCE

#### NEW QUESTION 13

Which of the following is true about job chains?

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

**Answer:** B

**Explanation:** Creating and Managing Job Chains

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

#### NEW QUESTION 17

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 21

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

- A. SET\_ATTRIBUTE\_NULL
- B. ALTER\_JOB
- C. ALTER\_JOB\_PARAMETERS
- D. ALTER
- E. SET\_ATTRIBUTE

**Answer:** AE

**Explanation:** SET\_ATTRIBUTE\_NULL Procedure, Changes an attribute of an object to NULL SET\_ATTRIBUTE Procedure, Changes an attribute of a job, schedule, or other Scheduler object  
There is NO such "ALTER\_JOB", "ALTER\_JOB\_PARAMETERS" and "ALTER" procedure available.

#### NEW QUESTION 25

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 27

How many different calendars does Oracle 11g support?

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

**Answer:** B

#### NEW QUESTION 32

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 36

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 37

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

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

**Answer:** AC

**Explanation:** Defining Chain Steps

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

? A Scheduler program object (program)

? Another chain (a nested chain)

? An event schedule, inline event, or file watcher

You define a step that points to a program or nested chain by using the DEFINE\_CHAIN\_STEP procedure. An example is the following, which adds two steps to my\_chain1:

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

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

#### NEW QUESTION 42

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 46

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 51

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 56

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 57**

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 58**

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

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

**Answer:** BDE

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

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

? A Scheduler program object (program)

? Another chain (a nested chain)

? An event schedule, inline event, or file watcher

**NEW QUESTION 63**

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 65

You want to schedule a job to rebuild all indexes on the SALES table after the completion of a bulk load operation. The bulk load operation must also be a scheduled job that executes as soon as the first file that contains data arrives on the system. How would you create these jobs?

- A. Create both jobs by using events raised by the scheduler
- B. Create both jobs by using events raised by the application
- C. Create a job to rebuild indexes by using events raised by the application and then create another job to perform bulk load by using events raised by the scheduler
- D. Create a job to rebuild indexes by using events raised by the Scheduler and then create another job to perform bulk load by using events raised by the application

**Answer:** D

**Explanation:** The bulk load operation can be done by setting up a file watcher, then raise the event to inform the scheduler, it is kind of application based event.

After the bulk load job, you can use Scheduler raised event, e.g. JOB COMPLETE, to fire the index rebuild job.

There are two kinds of events consumed by the Scheduler:

? Events raised by your application

An application can raise an event to be consumed by the Scheduler. The Scheduler reacts to the event by starting a job. For example, when an inventory tracking system notices that the inventory has gone below a certain threshold, it can raise an event that starts an inventory replenishment job.

See "Starting Jobs with Events Raised by Your Application".

? File arrival events raised by a file watcher

You can create a file watcher—a Scheduler object introduced in Oracle Database 11g Release 2—to watch for the arrival of a file on a system. You can then configure a job to start when the file watcher detects the presence of the file. For example, a data warehouse for a chain of stores loads data from end-of-day revenue reports uploaded from the point- of-sale systems in the stores. The data warehouse load job starts each time a new end-of- day report arrives.

#### NEW QUESTION 66

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

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

**Answer:** B

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

#### NEW QUESTION 67

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 69

What is the danger associated with stopping a running job by using the STOP\_JOB procedure?

- A. The job will need to be reenabled before it will execute again.

- B. The job may hold locks on objects referenced within it.
- C. All jobs within the job group will also be stopped.
- D. The job may leave data in an inconsistent state.
- E. There is no danger in using the STOP\_JOB procedure.

**Answer:** D

#### NEW QUESTION 71

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

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

**Answer:** B

#### NEW QUESTION 75

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 79

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 83

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 88

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 91

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 96**

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

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

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

Answer: C

**Explanation:** Refer to here.

Example 2

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

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

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

**NEW QUESTION 99**

View the Exhibit and examine the resource consumption details for the current plan in use by the database instance.

Which two statements are true based on the output? (Choose two.) Exhibit:

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

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

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

Answer: BC

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

SELECT name, active\_sessions, queue\_length, consumed\_cpu\_time, cpu\_waits, cpu\_wait\_time FROM v\$rsrc\_consumer\_group;

NAME	ACTIVE_SESSIONS	QUEUE_LENGTH	CONSUMED_CPU_TIME	CPU_WAITS
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 101

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 103

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

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

**Answer:** AD

### NEW QUESTION 107

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 112

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

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

**Answer:** D

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

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

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

**NEW QUESTION 113**

USER\_DATA is a nonencrypted tablespace that contains a set of tables with data. You want to convert all existing data in the USER\_DATA tablespace and the new data into the encrypted format. Which methods would you use to achieve this? (Choose all that apply.)

- A. Use Data Pump to transfer the existing data to a new encrypted tablespace.
- B. Use ALTER TABLE MOVE to transfer the existing data to a new encrypted tablespace.
- C. Use CREATE TABLE AS SELECT to transfer the existing data to a new encrypted tablespace.
- D. Enable row movement for each table to be encrypted and then use ALTER TABLESPACE to encrypt the tablespace.
- E. Encrypt the USER\_DATA tablespace using the ALTER TABLESPACE statement so that all the data in the tablespace is automatically encrypted.

**Answer:** ABC

**NEW QUESTION 118**

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

SQL> ALTER TABLE trans SHRINK SPACE CASCADE;

Which statement describes the outcome of the command? Exhibit:

Select	Name	Type	Extent Management	Segment Management	Status	Size (MB)	Used (MB)
<input checked="" type="radio"/>	SYSAUX	PERMANENT	LOCAL	AUTO	ONLINE	330.000	322.000
<input 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 122**

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 126**

You notice that a long-running transaction is suspended due to a space constraint, and there is no AFTER SUSPEND triggered event addressing the issue. You also note that the critical transaction is just about to reach the RESUMABLE\_TIMEOUT value. Which of these actions is appropriate?

- A. Abort the session, fix the space problem, then resubmit the transaction.
- B. Use the DBMS\_RESUMABLE.SET\_SESSION\_TIMEOUT procedure to extend the time-out for the session while you fix the problem.
- C. Do nothing, let the transaction fail, then fix the problem.
- D. Use Segment Shrink to clean up the table.
- E. Use the DBMS\_RESUMABLE.SET\_TIMEOUT procedure to extend the time-out for the session while you fix the problem.

**Answer: B**

**Explanation:** Topic 17, Managing Resources

**NEW QUESTION 127**

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

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

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

Which statement describes the consequence? Exhibit:

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

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

**Answer: A**

**Explanation:** Requisition of Calibrate I/O (link)

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

- ? The user must be granted the SYSDBA privilege
- ? timed\_statistics must be set to TRUE
- ? Asynchronous I/O must be enabled
- ? When using file systems, asynchronous I/O can be enabled by setting the FILESYSTEMIO\_OPTIONS initialization parameter to SETALL.
- ? Ensure that asynchronous I/O is enabled for data files by running the following query:

```
COL NAME FORMAT A50
SELECT NAME,ASYNCH_IO FROM V$DATAFILE F,V$IOSTAT_FILE I WHERE F.FILE#=I.FILE_NO
AND FILETYPE_NAME='Data File';
```

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

**NEW QUESTION 131**

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 135**

Which three statements correctly describe the features of the I/O calibration process? (Choose three.)

- A. Only one I/O calibration process can run at a time.
- B. It automates the resource allocation for the Automated Maintenance Tasks.
- C. It improves the performance of the performance-critical sessions while running.
- D. It can be used to estimate the maximum number of I/Os and maximum latency time for the system.
- E. The latency time is computed only when the TIMED\_STATISTICS initialization parameter is set to TRUE.

**Answer:** ADE

**NEW QUESTION 138**

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 139**

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 143**

Evaluate the following block of code:

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

What is the outcome of the above code?

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

**Answer:** C

**NEW QUESTION 145**

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 148

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 153

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 156

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 159

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 164

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 167

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

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

**Answer: BD**

#### NEW QUESTION 170

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 175

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 179

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 180

You have 100 segments in the USERS tablespace. You realize that the USERS tablespace is running low on space. You can use Segment Advisor to .

- A. identify the segments that you should shrink
- B. modify the storage parameters for the tablespace
- C. automatically shrink the segments with unused space
- D. check the alerts generated for the tablespace that is running low on space

Answer: A

#### NEW QUESTION 184

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 rang
- 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 188

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 193

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

1./etc/mtab 2.\$ORACLE\_HOME/dbs/oranfstab 3./etc/oranfstab

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

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

Answer: C

#### NEW QUESTION 198

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 201

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 206

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 210

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 215

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 220

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 224

Observe the following PL/SQL block: BEGIN  
dbms\_spm.configure('SPACE\_BUDGET\_PERCENT', 30); END;  
Which statement is correct regarding the above PL/SQL block?

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

**Answer:** D

#### NEW QUESTION 229

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 233

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 235

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

ERROR at line 1:

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

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

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

- A. Set RESUMABLE\_TIMEOUT for the instance.
- B. Set the RESOURCE\_LIMIT Parameter to TRUE.
- C. Enable the database instance to use asynchronous commit.
- D. Set the LOG\_CHECKPOINT\_TIMEOUT parameter to a nonzero value for the database instance.

**Answer:** A

**Explanation:** How Resumable Space Allocation Works

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

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

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

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

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

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

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

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

The error is reported in the alert log.

The system issues the Resumable Session Suspended alert.

If the user registered a trigger on the AFTER SUSPEND system event, the user trigger is executed. A user supplied PL/SQL procedure can access the error message data using the DBMS\_RESUMABLE package and the DBA\_ or USER\_RESUMABLE view.

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

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

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

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

? A suspension time out interval, specified by the RESUMABLE\_TIMEOUT

initialization parameter or by the timeout value in the ALTER SESSION ENABLE RESUMABLETIMEOUT statement, is associated with resumable statements. A resumable statement that is suspended for the timeout interval wakes up and returns the exception to the user if the error condition is not resolved within the timeout interval.

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

#### NEW QUESTION 237

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 240

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

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

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

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

Which is the correct sequence of the required steps?

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

**Answer:** B

**Explanation:** 11.1 Steps for Replaying a Database Workload

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

following steps as appropriate:

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

#### NEW QUESTION 245

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

- A. Segment Advisor
- B. SQL Access Advisor
- C. Undo Advisor
- D. SQL Tuning Advisor

**Answer:** B

#### NEW QUESTION 246

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 247

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

- A. Workload capture is per instance.

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

**Answer: E**

**NEW QUESTION 250**

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 253**

Which statements are true regarding the system-defined moving window baseline in Oracle Database 11g? (Choose all that apply.)

- A. It does not allow you to change the moving window size.
- B. Adaptive threshold functionalities use it by default to compute statistics.
- C. It is created by default with the window size being equal to the AWR retention time.
- D. It is created when the first snapshot is collected by the Automatic Workload Repository (AWR).

**Answer: BC**

**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. Oracle Database automatically maintains a system-defined moving window baseline. The default window size for the system-defined moving window baseline is the current AWR retention period, which by default is 8 days. If you are planning to use adaptive thresholds, consider using a larger moving window—such as 30 days—to accurately compute threshold values. You can resize the moving window baseline by changing the number of days in the moving window to a value that is equal to or less than the number of days in the AWR retention period. Therefore, to increase the size of a moving window, you must first increase the AWR retention period accordingly.

**NEW QUESTION 258**

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

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

**Answer: BC**

**NEW QUESTION 263**

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 268**

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

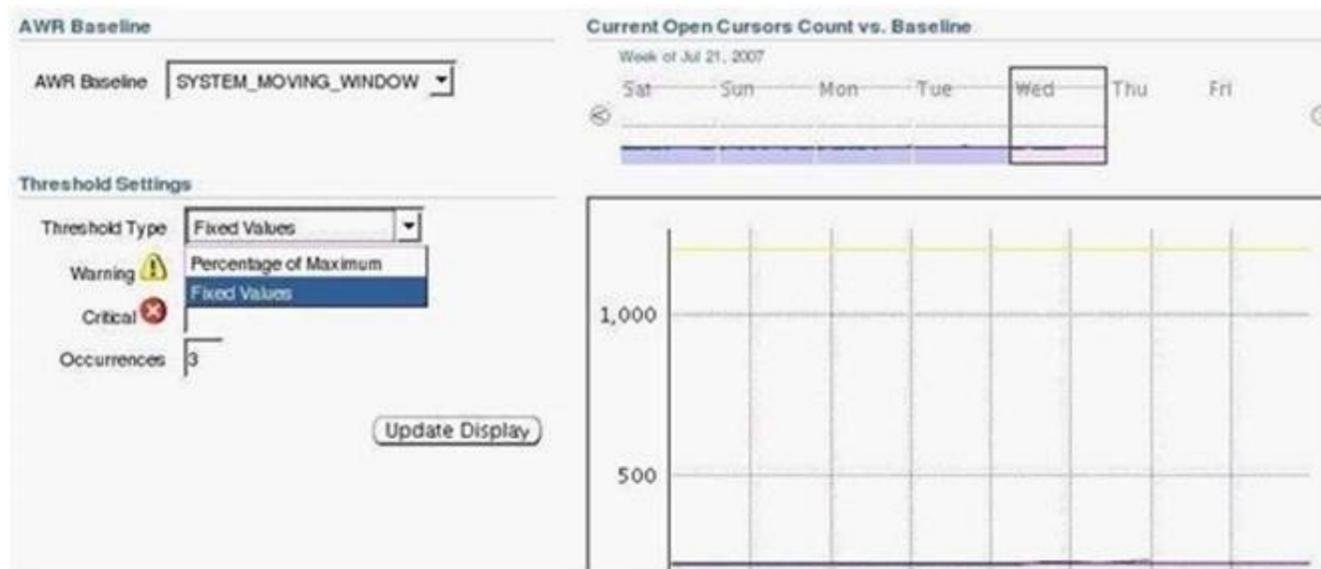
- A. Building the SQL profile
- B. Recommending optimization of materialized views
- C. Checking query objects for missing and stale statistics
- D. Recommending bitmap, function-based, and B-tree indexes
- E. Recommending restructuring SQL queries that are using bad plans

**Answer:** ACE

**NEW QUESTION 272**

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

Exhibit:



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

**Answer:** B

**Explanation:**

**ORACLE Enterprise Manager 11g**  
 Database Control

Database Instance: iocp > [Baseline Metric Thresholds](#) >  
**Edit Thresholds: Current Open Cursors Count**

**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 276

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 280

Which two statements are true regarding hot patching? (Choose two.)

- A. It requires relinking of the Oracle binary.
- B. It does not require database instance shutdown.
- C. It can detect conflicts between two online patches.
- D. It is available for installing all patches on all platforms.
- E. It works only in a single database instance environment.

**Answer:** BC

#### Explanation: Online Patching

Regular patches typically contain .o (object) files and/or .a (archive) libraries, and therefore require a relink of the RDBMS binary. Online patches, however, contain .so files, which are dynamic/shared libraries, and do not require a relink of the RDBMS binary. Consequently, since a relink is not needed, you can apply or roll back online patches while the RDBMS instance is running. This simplifies administration, because no downtime is needed, and also results in a much quicker turnaround time for installing or de-installing Online Patches.

A regular RDBMS patch can require many minutes to install, since it requires instance shutdown, a relink, and instance startup. On the other hand, you can install an online patch in just a few seconds.

Online patches are only applicable for Oracle RDBMS and not any other products. Online patches are currently supported on the following Windows and UNIX platforms for version 11.2.0.1.0 and later:

- ? Linux x86
- ? Linux x86\_64
- ? HP-UX Itanium
- ? Solaris SPARC 64-bit
- ? Solaris AMD 64-bit
- ? AIX (AIX 6.1 and later)

#### About Patch Conflicts:

All patches may not be compatible with one another. For example, if you apply a patch, all the bugs the patch fixes could reappear after you apply another patch. This is called a conflict situation. OPatch detects such situations and raises an error when it detects a conflict.

#### NEW QUESTION 284

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 286

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

- A. Create new indexes.
- B. Modify existing indexes.

- C. Implement partitioning on a nonpartitioned table.
- D. Create materialized views.
- E. All of the above

**Answer:** E

**Explanation:** Overview of SQL Access Advisor

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

**NEW QUESTION 287**

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

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

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

parameter-optimizer (exhibit):

```
SQL> SHOW PARAMETER OPTIMIZER
```

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

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

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

**Answer:** B

**Explanation:** Setting the OPTIMIZER\_MODE Initialization Parameter(Link)

The OPTIMIZER\_MODE initialization parameter establishes the default behavior for choosing an optimization approach for the instance.

OPTIMIZER\_MODE Initialization Parameter Values

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

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

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

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

**NEW QUESTION 290**

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 regressedplans.

- 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 292

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 294

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 297

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 299

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, Lower...	SYSMAN	361qjn3w9uflh	✓	(99.9%) ✓				7/12/07
<input checked="" type="radio"/>	SELECT EXECUTION_ID, STATUS, STATUS_DETA...	SYSMAN	lyk8t9986atk7		(69%) ✓	(97.9%) ✓			7/12/07
<input type="radio"/>	SELECT /*+ INDEX(sqbbj\$ (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	4q8mr2bvy6qf					✓	7/12/07
<input type="radio"/>	select l.ts#,l.file#,l.block#,nv{l.bobj...	SYS	1qu8t96d0bdmu					✓	7/12/07
<input type="radio"/>	select obj#, dataobj#, part#, hbounden...	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	23y48d28wkg2r					✓	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 304**

To view the results of the most recent Automatic SQL Tuning Advisor task, which sequence should you follow?

- A. EM Database home page, Software and Support, SQL Advisors, Automatic SQL Tuning Advisor
- B. EM Database home page, Software and Support, Advisor Central, SQL Advisors, Automatic SQL Tuning Advisor
- C. EM Database home page, Software and Support, Support Workbench, Advisor Central, SQL Advisors, Automatic SQL Tuning Advisor
- D. Either B or C
- E. All of the above

**Answer: D**

**NEW QUESTION 307**

Which of the following represents the correct sequence of events for Database Replay?

- A. Capture, analyze, preprocess, replay
- B. Capture, preprocess, analyze, replay
- C. Capture, preprocess, replay, analyze
- D. Analyze, capture, preprocess, replay
- E. None of the above

**Answer: C**

**Explanation: Database Replay**

Database Replay allows workloads to be captured from production systems and re-executed with high fidelity on test copies of production databases. This enables detailed analysis of how the proposed changes may affect production systems; for instance, patching or upgrading database software.

Task Name	Description	Go to Task
1 Capture Workload	Capture a workload from the production environment. This can be scheduled to accommodate a database restart if desired.	
2 Preprocess Workload	Preprocessing prepares a captured workload for replay. You must do this once for every captured workload. Preprocessing is best performed in the test database. The captured workload must be accessible from the test database.	
3 Replay Workload	Replay the preprocessed workload on a test copy of the production database.	

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

**NEW QUESTION 308**

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 311

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 313

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 315

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 317

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

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

**Answer:** CD

#### NEW QUESTION 319

Evaluate the following code:

```
SQL>VARIABLE task_name VARCHAR2(255);
SQL>VARIABLE sql_stmt VARCHAR2(4000);
SQL>BEGIN
    :sql_stmt :=| 'SELECT COUNT(*) FROM customers WHERE cust_state_province
= 'CA'';
    :task_name := 'MY_QUICKTUNE_TASK';
    DBMS_ADVISOR.QUICK_TUNE
(DBMS_ADVISOR.SQLACCESS_ADVISOR, :task_name, :sql_stmt);
END;
```

What is the outcome of this block of code?

- A. It creates a task and workload, and executes the task.
- B. It creates a task and workload but does not execute the task.
- C. It produces an error because a template has not been created.

D. It produces an error because the SQL Tuning Set has not been created.

**Answer:** A

#### NEW QUESTION 324

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

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

**Answer:** CE

#### NEW QUESTION 327

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

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

**Answer:** ABCE

#### Explanation: 17.2.1 How Automatic SQL Tuning Works (link)

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

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

1. Identifies SQL candidates in the AWR for tuning

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

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

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

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

3. Tests SQL profiles by executing the SQL statement

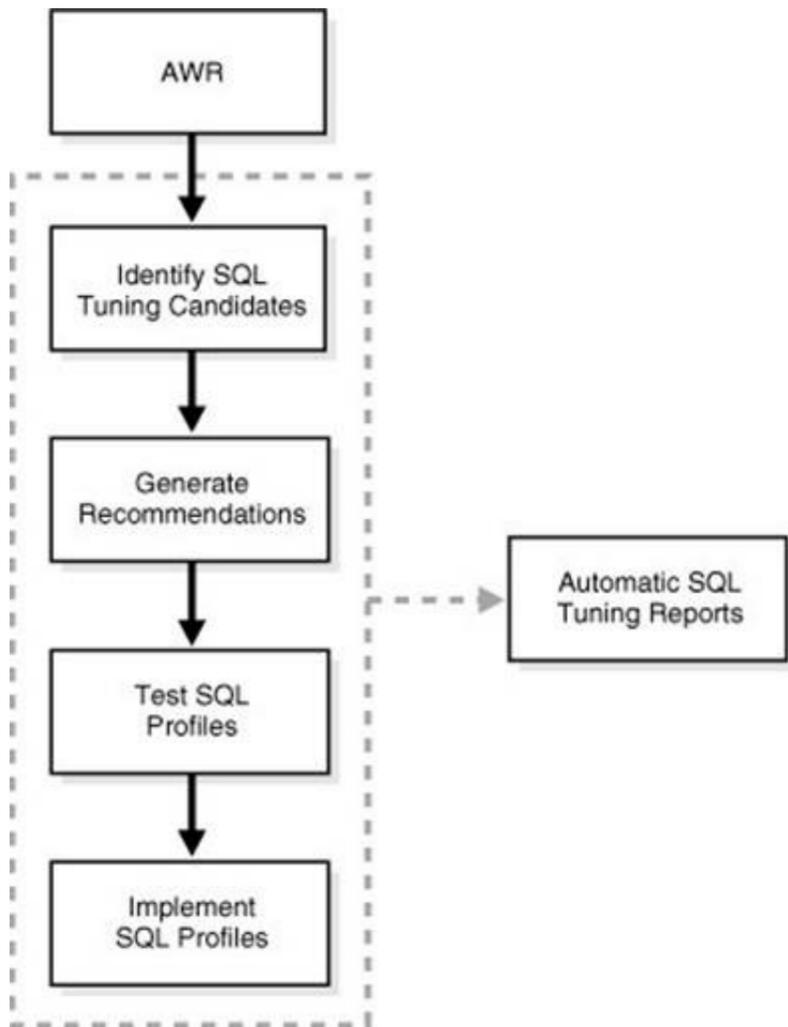
If a SQL profile is recommended, the database tests the new profile by executing the SQL statement both with and without the profile. If the performance improvement improves at least threefold, then the database accepts the SQL profile, but only if the ACCEPT\_SQL\_PROFILES task parameter is set to TRUE.

Otherwise, the automatic SQL tuning reports merely report the recommendation to create a SQL profile.

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

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

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



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

**NEW QUESTION 329**

View the Exhibit-1 to observe the maintenance window property. View the Exhibit-2 to examine the output of the query. Which two statements describe the conclusions? (Choose two.) EM-window (exhibit):

ORACLE Enterprise Manager 11g Database Control

Database Instance: orcl > Scheduler Windows > View Window: TUESDAY\_WINDOW

Name: TUESDAY\_WINDOW  
 Resource Plan: DEFAULT\_MAINTENANCE\_PLAN  
 Enabled: TRUE  
 Priority: LOW  
 Description: Tuesday window for maintenance tasks

Schedule: Repeat By Weeks  
 Interval (Weeks): 1  
 Days of Week: Tuesday  
 Repeat Time: Hour:10 Minute:00 Second:00 PM  
 Duration (min): 0 hour(s) 20 minute(s)  
 Available to Start: Not Available After

sql-dba\_autotask\_client (exhibit):

```

SQL> SELECT client_name,
2 mean_job_duration,
3 resource_percentage
4 FROM DBA_AUTOTASK_CLIENT;
  
```

CLIENT NAME	MEAN_JOB_DURATION	RESOURCE_PERCENTAGE
auto optimized stats collection	+0000000000 00:18:41.000000000	25
auto space advisor	+0000000000 00:02:36.666666666	25
sql tuning advisor	+0000000000 00:04:36.500000000	25

- A. RESOURCE\_PERCENTAGE should be increased.
- B. The repeat time for the window should be decreased.
- C. RESOURCE\_PERCENTAGE should be decreased.
- D. The window duration should be increased.

**Answer: AD**

#### NEW QUESTION 331

Which statement about Automatic Memory Management with Oracle 11g is true?

- A. You cannot specify MEMORY\_TARGET if you explicitly specify SGA\_TARGET or PGA\_AGGREGATE\_TARGET values that are greater than zero in your parameter file.
- B. Oracle can reallocate memory between the SGA and PGA automatically as needed.
- C. To use Automatic Memory Management, you must explicitly set both the MEMORY\_TARGET and MEMORY\_MAX\_TARGET parameters in your parameter file.
- D. You can set the MEMORY\_TARGET parameter to a maximum value of the current SGA size plus the current PGA size.

**Answer: B**

**Explanation:** Topic 15, Managing Database Performance

#### NEW QUESTION 335

Which two statements about workload capture and replay operations are true? (Choose two.)

- A. The clients must be created manually on the test machines to perform more realistic testing.
- B. Restart the database in RESTRICTED mode before beginning workload replay to enable a clean state for workload replay.
- C. Restart the database in RESTRICTED mode before beginning workload capture to enable a clean state for capturing the workload.
- D. The application state of the capture system must be identical to the application state of the replay system when the workload replay begins.

**Answer: CD**

#### NEW QUESTION 338

Note the following parameter settings: Which setting is NOT allowed?

- A. ALTER SYSTEM SET DB\_CACHE\_SIZE=50M;
- B. ALTER SYSTEM SET DB\_8K\_CACHE\_SIZE=10M;
- C. ALTER SYSTEM SET DB\_4K\_CACHE\_SIZE=10M;
- D. ALTER SYSTEM SET DB\_16K\_CACHE\_SIZE=10M;

**Answer: B**

#### NEW QUESTION 341

Using Oracle Enterprise Manager to set SGA pool values manually, for which of the following pools does Oracle EM offer advice to set the value appropriately? (Choose all that apply.)

- A. DB\_CACHE\_SIZE
- B. SHARED\_POOL\_SIZE
- C. LARGE\_POOL\_SIZE
- D. JAVA\_POOL\_SIZE
- E. SGA\_MAX\_SIZE
- F. SORT\_AREA\_SIZE

**Answer: AB**

#### NEW QUESTION 342

Which statements are true regarding the Query Result Cache? (Choose all that apply.)

- A. It can be set at the system, session, or table level.
- B. It is used only across statements in the same session.
- C. It can store the results from normal as well as flashback queries.
- D. It can store the results of queries based on normal, temporary, and dictionary tables.

**Answer: AC**

#### NEW QUESTION 345

For Oracle 11g, Oracle strongly recommends that you configure your database to use which of the following memory management features?

- A. Automatic PGA Memory Management
- B. Automatic SGA Memory Management
- C. Automatic Shared Memory Management
- D. Automatic Memory Management
- E. Manual SGA Memory Management
- F. None of the above

**Answer: D**

#### NEW QUESTION 350

Examine the parameter setting in your database:

NAME	TYPE	VALUE
archive_lag_target	integer	0
db_flashback_retention_target	integer	1440
fast_start_io_target	integer	0
fast_start_mttr_target	integer	0
memory_max_target	big integer	808M
memory_target	big integer	808M
pga_aggregate_target	big integer	0
sga_target	big integer	0

SQL> SHOW PARAMETER SGA\_MAX\_SIZE

NAME	TYPE	VALUE
sga_max_size	big integer	808M

Which statement is correct about the database?

- A. Automatic memory management is disabled because PGA\_AGGREGATE\_TARGET and SGA\_TARGET are not set
- B. The instance is started but the database will not be opened until PGA\_AGGREGATE\_TARGET and SGA\_TARGET are set
- C. The database is opened but users cannot perform transactions until PGA\_AGGREGATE\_TARGET and SGA\_TARGET are set
- D. Automatic memory management is enabled and, as per policy, 60% of the memory for System Global Area (SGA) and 40% of the memory for Program Global Area (PGA) will be distributed at startup

**Answer: D**

**NEW QUESTION 353**

Which statements about the MEMORY\_TARGET initialization parameter are true? (Choose all that apply.)

- A. MEMORY\_TARGET can be increased up to the value of MEMORY\_MAX\_TARGET, if MEMORY\_MAX\_TARGET is set to a value greater than zero
- B. MEMORY\_MAX\_TARGET defaults to a value of zero if MEMORY\_TARGET is not set
- C. MEMORY\_TARGET represents the total amount of memory that can be allocated to SGA and PGA memory structures.
- D. MEMORY\_TARGET is static and cannot be modified without shutting down the instance

**Answer: ABC**

**NEW QUESTION 356**

In Oracle 11g, by default which one of the following conditions implicitly enables Automatic PGA Memory Management?

- A. Setting a nonzero value for SGA\_TARGET
- B. Configuring Automatic Shared Memory Management
- C. Configuring Automatic Memory Management
- D. Setting a nonzero value for SGA\_MAX\_SIZE and PGA\_AGGREGATE\_TARGET
- E. None of the above

**Answer: C**

**NEW QUESTION 359**

In your database, the RESULT\_CACHE\_MODE parameter has been set to MANUAL in the initialization parameter file. You issued the following command:  
 SQL>SELECT /\*+ RESULT\_CACHE \*/ sale\_category, sum(sale\_amt) FROM sales GROUP BY sale\_category;  
 Where would the result of this query be stored?

- A. database buffer cache
- B. shared pool
- C. PGA
- D. large pool

**Answer: B**

**NEW QUESTION 360**

You have applications that have frequently executed queries, and produce small and static result sets. You configure the sqlnet.ora file in the client machine to set a nonzero value for the OCI\_RESULT\_CACHE\_MAX\_SIZE parameter.  
 What is the purpose of this configuration?

- A. to avoid round trips to the server by enabling caching of query results in client memory
- B. to improve performance by storing a copy of the data from the private SQL area of the PGA
- C. to enhance the query performance by creating a cache in the client memory for sorting operations
- D. to avoid the storing of query plans and results in the server by creating a cache in the client memory

**Answer: A**

**NEW QUESTION 361**

View the Exhibit to examine the output produced by the following query at three different times since the database instance started and has undergone workloads of different capacities:

SQL> SELECT substr(component, 0, 10) COMP, current\_size CS, user\_specified\_size US FROM v\$memory\_dynamic\_components

WHERE current\_size!=0; What do you infer from this? Exhibit:

```

First execution:
=====
COMP                CS                US
-----
shared pool         58720256          0
large pool          4194304           0
java pool           4194304           0
SGA Target          176160768         0
DEFAULT bu         100663296         0
Shared IO           8388608           8388608
PGA Target          117440512         0

6 rows selected.

Second execution:
=====
COMP                CS                US
-----
shared pool         58720256          0
large pool          4194304           0
java pool           4194304           0
SGA Target          192937984         0
DEFAULT bu         117440512         0
Shared IO           8388608           8388608
PGA Target          100663296         0

6 rows selected.

Third execution:
=====
COMP                CS                US
-----
shared pool         62914560          0
large pool          100663296         0
java pool           4194304           0
SGA Target          192937984         0
DEFAULT bu           8388608           0
Shared IO           8388608           8388608
PGA Target          100663296         0

6 rows selected.

```

- A. All sessions are connected to the database instance in dedicated mode, and no RMAN or parallel query operations have been performed.
- B. The database instance is running with manual shared memory management.
- C. The database instance is running with manual PGA management.
- D. The database instance has the MEMORY\_TARGET value set to a nonzero value.

Answer: D

**NEW QUESTION 363**

View the Exhibit to examine the parameter values. You are planning to set the value for the MEMORY\_TARGET parameter of your database instance. What value would you assign? Exhibit:

SQL> SHOW PARAMETER TARGET

NAME	TYPE	VALUE
archive_lag_target	integer	0
db_flashback_retention_target	integer	1440
fast_start_io_target	integer	0
fast_start_mtrr_target	integer	3600
memory_max_target	big integer	0
memory_target	big integer	0
pga_aggregate_target	big integer	90M
sga_target	big integer	272M

- A. 1440 MB
- B. 90 MB
- C. 362 MB
- D. 272 MB

Answer: C

**NEW QUESTION 365**

You set the following parameters in the parameter file and restarted the database:

MEMORY\_MAX\_TARGET=0 MEMORY\_TARGET=500M PGA\_AGGREGATE\_TARGET=90M  
 SGA\_TARGET=270M

Which two statements are true regarding these parameters after the database instance is restarted? (Choose two.)

- A. The MEMORY\_MAX\_TARGET parameter is automatically set to 500 MB.
- B. The value of the MEMORY\_MAX\_TARGET parameter remains zero till it is changed manually.
- C. The PGA\_AGGREGATE\_TARGET and SGA\_TARGET parameters are automatically set to zero.
- D. The lower bounds of PGA\_AGGREGATE\_TARGET and SGA\_TARGET parameters are set to 90 MB and 270 MB, respectively.

Answer: AD

#### NEW QUESTION 368

The environmental variable ORACLE\_BASE is set. You want to check the diagnostic files created as part of the Automatic Diagnostic Repository (ADR). View the Exhibit and note the various parameters set in your database.

What will be the location of the ADR base? Exhibit:

```
SQL> SELECT name, value FROM v$pparameter WHERE name LIKE '%dest';
```

NAME	VALUE
log_archive_dest	
log_archive_duplex_dest	
log_archive_min_succeed_dest	
standby_archive_dest	
db_create_file_dest	
db_recovery_file_dest	/u01/app/oracle/flash_recovery_area
background_dump_dest	
user_dump_dest	
core_dump_dest	
audit_file_dest	/u01/app/oracle/admin/orcl/adump
diagnostic_dest	

11 rows selected.

- A. It is set to ORACLE\_BASE.
- B. It is set to ORACLE\_HOME/log.
- C. It is set to /u01/app/oracle/admin/orcl/adump.
- D. It is set to /u01/app/oracle/flash\_recovery\_area.

Answer: A

**Explanation:** The ADR root directory is known as ADR base. Its location is set by the DIAGNOSTIC\_DEST initialization parameter. If this parameter is omitted or left null, the database sets DIAGNOSTIC\_DEST upon startup as follows:

? If environment variable ORACLE\_BASE is set, DIAGNOSTIC\_DEST is set to the directory designated by ORACLE\_BASE.

? If environment variable ORACLE\_BASE is not set, DIAGNOSTIC\_DEST is set to ORACLE\_HOME/log.

#### NEW QUESTION 371

Which two statements are true regarding the Automatic Diagnostic Repository (ADR) in Oracle Database 11g? (Choose two.)

- A. A single ADR can support multiple ADR homes for different database instances.
- B. The alert files are stored in XML file format in the TRACE directory of each ADR home.
- C. If the environmental variable ORACLE\_BASE is set, then DIAGNOSTIC\_DEST is set to \$ORACLE\_BASE.
- D. The BACKGROUND\_DUMP\_DEST initialization parameter overrides the DIAGNOSTIC\_DEST initialization parameter for the location of the alert log file.

Answer: AC

**Explanation:** The ADR root directory is known as ADR base. Its location is set by the DIAGNOSTIC\_DEST initialization parameter. If this parameter is omitted or left null, the database sets DIAGNOSTIC\_DEST upon startup as follows:

? If environment variable ORACLE\_BASE is set, DIAGNOSTIC\_DEST is set to the directory designated by ORACLE\_BASE.

? If environment variable ORACLE\_BASE is not set, DIAGNOSTIC\_DEST is set to

ORACLE\_HOME/log

Within ADR base, there can be multiple ADR homes, where each ADR home is the root directory for all diagnostic data—traces, dumps, the alert log, and so on—for a particular instance of a particular Oracle product or component. For example, in an Oracle Real Application Clusters environment with Oracle ASM, each database instance, Oracle ASM instance, and listener has an ADR home.

#### NEW QUESTION 374

Which statements are true regarding the concept of problems and incidents in the fault diagnosability infrastructure for Oracle Database 11g? (Choose all that apply.)

- A. Only the incident metadata and dumps are stored in the Automatic Diagnostic Repository (ADR).
- B. The problem key is the same as the incident number.
- C. The database sends an incident alert to the Oracle Enterprise Manager Database Home page.
- D. Every problem has a problem key, which is a text string that describes the problem.
- E. The database makes an entry into the alert log file when problems and incidents occur.

Answer: CDE

**Explanation:** Reference at here

Fault Diagnosability Infrastructure Overview

The fault diagnosability infrastructure aids in preventing, detecting, diagnosing, and resolving problems. The problems that are targeted in particular are critical errors such as those caused by code bugs, metadata corruption, and customer data corruption.

When a critical error occurs, it is assigned an incident number, and diagnostic data for the error (such as trace files) are immediately captured and tagged with this number. The data is then stored in the Automatic Diagnostic Repository (ADR)—a file-based repository outside the database—where it can later be retrieved by incident number and analyzed. About Incidents and Problems

A problem is a critical error in a database instance, Oracle Automatic Storage Management (Oracle ASM) instance, or other Oracle product or component. Critical errors manifest as internal errors, such as ORA-00600, or other severe errors, such as ORA-07445 (operating system exception) or ORA-04031 (out of memory in the shared pool). Problems are tracked in the ADR. Each problem has a problem key, which is a text string that describes the problem. It includes an error code (such as ORA 600) and in some cases, one or more error parameters.

An incident is a single occurrence of a problem. When a problem (critical error) occurs multiple times, an incident is created for each occurrence. Incidents are timestamped and tracked in the Automatic Diagnostic Repository (ADR). Each incident is identified by a numeric incident ID, which is unique within the ADR. When

an incident occurs, the database:

- ? Makes an entry in the alert log.
- ? Sends an incident alert to Oracle Enterprise Manager (Enterprise Manager).
- ? Gathers first-failure diagnostic data about the incident in the form of dump files (incident dumps).
- ? Tags the incident dumps with the incident ID.
- ? Stores the incident dumps in an ADR subdirectory created for that incident.

#### NEW QUESTION 375

Which two statements are true regarding the starting of the database instance using the following command? (Choose two.)

```
SQL>STARTUP UPGRADE
```

- A. It enables all system triggers.
- B. It allows only SYSDBA connections.
- C. It ensures that all job queues remain active during the upgrade process.
- D. It sets system initialization parameters to specific values that are required to enable database upgrade scripts to be run.

**Answer:** BD

#### NEW QUESTION 378

Which of the following are true concerning block media recovery? (Choose all that apply.)

- A. Any gap in archive logs ends the recovery.
- B. If a gap in archive logs is encountered, RMAN will search forward for newer versions of the blocks that are not corrupt.
- C. Uncorrupted blocks from the flashback logs may be used to speed recovery.
- D. The database can be in NOARCHIVELOG mode.
- E. None of the above.

**Answer:** BC

**Explanation:** Overview of Block Media Recovery ([link](#)) Basic Concepts of Block Media Recovery

Whenever block corruption has been automatically detected, you can perform block media recovery manually with the RECOVER ... BLOCK command. By default, RMAN first searches for good blocks in the real-time query physical standby database, then flashback logs and then blocks in full or level 0 incremental backups.

Prerequisites for Block Media Recovery ([link](#))

The following prerequisites apply to the RECOVER ... BLOCK command:

? The target database must run in ARCHIVELOG mode and be open or mounted with a current control file.

? If the target database is a standby database, then it must be in a consistent state, recovery cannot be in session, and the backup must be older than the corrupted file.

? The backups of the data files containing the corrupt blocks must be full or level 0 backups and not proxy copies.

If only proxy copy backups exist, then you can restore them to a nondefault location on disk, in which case RMAN considers them data file copies and searches them for blocks during block media recovery.

? RMAN can use only archived redo logs for the recovery.

RMAN cannot use level 1 incremental backups. Block media recovery cannot survive a missing or inaccessible archived redo log, although it can sometimes survive missing redo records.

? Flashback Database must be enabled on the target database for RMAN to search the flashback logs for good copies of corrupt blocks.

If flashback logging is enabled and contains older, uncorrupted versions of the corrupt blocks, then RMAN can use these blocks, possibly speeding up the recovery.

? The target database must be associated with a real-time query physical standby database for RMAN to search the database for good copies of corrupt blocks.

Restore Failover ([link](#))

RMAN automatically uses restore failover to skip corrupted or inaccessible backups and look for usable backups. When a backup is not found, or contains corrupt data, RMAN automatically looks for another backup from which to restore the desired files.

#### NEW QUESTION 383

Which two kinds of failures make the Data Recovery Advisor (DRA) generate a manual checklist? (Choose two.)

- A. failures because a data file is renamed by error
- B. failures when no standby database is configured
- C. failures that require no archive logs to be applied for recovery
- D. failures due to loss of connectivity-for example, an unplugged disk cable

**Answer:** AD

#### NEW QUESTION 386

Which statement is true regarding the retention policy for the incidents accumulated in the Automatic Diagnostic Repository (ADR)?

- A. The incident metadata is purged when the problem is resolved and the DBA closes the SR.
- B. The incident files and dumps are not retained in the ADR for the manually created incidents.
- C. The incident files are retained but the incident metadata is purged when the problem is resolved and the DBA closes the SR
- D. The default setting is for one year after which the incident metadata is purged from the ADR and the files are retained for one month.

**Answer:** D

**Explanation:**

**ORACLE Enterprise Manager 11g**  
 Database Control

Database Instance: iocp > Support Workbench >  
**View Incident Packaging Configuration**

---

**Incident Data Retention**

Incident Metadata Retention Period (day)	365
Incident Files Retention Period (day)	30

**Packaging Settings**

These settings are used in selecting incidents and files from a problem when the problem is added to a package.

Cutoff Age for Incident Inclusion (day)	90
Leading Incidents Count	3
Trailing Incidents Count	3
Correlation Time Proximity (min)	5
Time Window for Package Content (hour)	24

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

**NEW QUESTION 389**

What view might you use to try to determine how long a particular backup will take?

- A. V\$SESSION\_EVENT
- B. V\$SESSION
- C. V\$WAITS
- D. V\$WAITSTAT
- E. V\$SESSION\_LONGOPS

**Answer:** E

**NEW QUESTION 393**

Which of the following are NOT fundamental tasks of the Support Workbench? (Choose all that apply.)

- A. View long-running SQL workloads
- B. View problem details
- C. Gather additional diagnostic information
- D. Create a Service Request
- E. Clean up incident data after upload to Oracle Support

**Answer:** AE

**Explanation:** The Enterprise Manager Support Workbench (Support Workbench) is a facility that enables you to investigate, report, and in some cases, repair problems (critical errors), all with an easy-to-use graphical interface. The Support Workbench provides a self-service means for you to gather first-failure diagnostic data, obtain a support request number, and upload diagnostic data to Oracle Support with a minimum of effort and in a very short time, thereby reducing time-to-resolution for problems. The Support Workbench also recommends and provides easy access to Oracle advisors that help you repair SQL-related problems, data corruption problems, and more.

**NEW QUESTION 395**

Choose the correct order to package and upload data for an incident to Oracle Support.

- A. Schedule, create new package, view manifest, view contents
- B. Create new package, view manifest, view contents, schedule
- C. Schedule, create new package, view contents, view manifest
- D. Create new package, view contents, view manifest, schedule
- E. None of the above.

**Answer:** D

**Explanation:**

**ORACLE Enterprise Manager 11g**  
 Database Control

---

Create New Package   View Contents   View Manifest   Schedule

**Quick Packaging: Create New Package**

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

**NEW QUESTION 399**

You executed the following commands:

SQL> ALTER SESSION SET OPTIMIZER\_USE\_PENDING\_STATISTICS = false; SQL> EXECUTE DBMS\_STATS.SET\_TABLE\_PREFS('SH', 'CUSTOMERS',

'PUBLISH', 'false');

SQL> EXECUTE DBMS\_STATS.GATHER\_TABLE\_STATS('SH', 'CUSTOMERS');

Which statement is correct regarding the above statistics collection on the SH.CUSTOMERS table in the above session?

- A. The statistics are stored in the pending statistics table in the data dictionary.
- B. The statistics are treated as the current statistics by the optimizer for all sessions.
- C. The statistics are treated as the current statistics by the optimizer for the current sessions only.
- D. The statistics are temporary and used by the optimizer for all sessions until this session terminates.

**Answer: A**

**NEW QUESTION 401**

Which tasks can be accomplished using the Enterprise Manager Support Workbench in Oracle Database 11g? (Choose all that apply.)

- A. Generate reports on data failure such as data file failures.
- B. You can track the Service Request (SR) and implement repairs.
- C. You can package and upload diagnostic data to Oracle Support.
- D. You can manually run health checks to gather diagnostic data for a problem.

**Answer: BCD**

**NEW QUESTION 405**

View the Exhibit to examine the output for the V\$DIAG\_INFO view. Which statements are true regarding the location of diagnostic traces? (Choose all that apply.)

Exhibit:

```
SELECT * FROM V$DIAG_INFO;
```

INST_ID	NAME	VALUE
1	Diag Enabled	TRUE
1	ADR Base	/u01/oracle
1	ADR Home	/u01/oracle/diag/rdbms/orclbi/orclbi
1	Diag Trace	/u01/oracle/diag/rdbms/orclbi/orclbi/trace
1	Diag Alert	/u01/oracle/diag/rdbms/orclbi/orclbi/alert
1	Diag Incident	/u01/oracle/diag/rdbms/orclbi/orclbi/incident
1	Diag Cdump	/u01/oracle/diag/rdbms/orclbi/orclbi/cdump
1	Health Monitor	/u01/oracle/diag/rdbms/orclbi/orclbi/hm
1	Default Trace File	/u01/oracle/diag/rdbms/orclbi/orclbi/trace/orcl_ora_22769.trc
1	Active Problem Count	8

- A. The path to the location of the background as well as the foreground process trace files is /u01/oracle/diag/ rdbms/orclbi/orclbi/trace.
- B. The location of the text alert log file is /u01/oracle/diag/rdbms/orclbi/orclbi/alert.
- C. The location of the trace file for the current session is /u01/oracle/diag/rdbms/orclbi/orclbi/trace.
- D. The location of the XML-formatted version of the alert log is /u01/oracle/diag/rdbms/orclbi/orclbi/alert.

**Answer: ACD**

**NEW QUESTION 409**

Identify the three predefined server-generated alerts. (Choose three.)

- A. Drop User
- B. Tablespace Space Usage
- C. Resumable Session Suspended
- D. Recovery Area Low On Free Space
- E. SYSTEM Tablespace Size Increment

**Answer: BCD**

**NEW QUESTION 410**

View the Exhibit to examine the metrics with a threshold.

Which statement is true regarding the Number of Transactions (per second) metric?

Exhibit:

Select Metric Name	Alert Counts		Threshold Type	Adaptive?
	Critical	Warning AWR Baseline		
Workload Volume Metrics	0	3		
Cumulative Logons (per second)	0	0	Fixed Values	
Current Open Cursors Count	0	0	Fixed Values	
Number of Transactions (per second)	0	3 SYSTEM_MOVING_WINDOW	Significance Level	✓

- A. Oracle uses statistical relevance to determine when an adaptive threshold has been breached for the metric.
- B. The statistics for the metric values observed over the baseline time period are not examined to determine threshold values.
- C. Oracle determines when an adaptive threshold has been breached based on the maximum value captured by the baseline.
- D. The total concurrent number of threshold violations, which must occur before an alert is raised for the metric, has been set to zero.

**Answer:** A

**NEW QUESTION 413**

.....

## **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](#)**