Topic 1 - Single Topic

Question #1 *Topic 1*

Examine the parameters for your database instance:

NAME	TYPE	VALUE
undo_management	string	AUTO
undo_retention	integer	1200
undo_tablespace	string	UNDOTBS1

You execute the following command:

SQL> ALTER TABLESPACE undotbs1 RETENTION NOGUARANTEE;

Which statement is true in this scenario?

A. Undo data is written to flashback logs after 1200 seconds.

B. Inactive undo data is retained for 1200 seconds even if subsequent transactions fail due to lack of space in the undo tablespace.

C. You can perform a Flashback Database operation only within the duration of 1200 seconds.

D. An attempt is made to keep inactive undo for 1200 seconds but transactions may overwrite the undo before that time has elapsed.

Correct Answer: *D*

dancymonkey Highly Voted 3 years, 11 months ago

B -> if RETENTION GRANTEE, it will let lack of space in the undo tablespace
D -> NOGUARANTEE -> it won't have lack of space so it may overwrite
THe answer is D

upvoted 11 times

Nick_144 Most Recent 3 years, 9 months ago

answer is A

upvoted 1 times

Idelmo 4 years, 3 months ago

La respuesta es la B

upvoted 1 times

hggz 4 years ago

D is correct

upvoted 9 times

Question #2 *Topic 1*

A user establishes a connection to a database instance by using an Oracle Net connection. You want to ensure the following:
1. The user account must be locked after five unsuccessful login attempts.
2. Data read per session must be limited for the user.
3. The user cannot have more than three simultaneous sessions.
4. The user must have a maximum of 10 minutes session idle time before being logged off automatically.
How would you accomplish this?

A. by granting a secure application role to the user

B. by implementing Database Resource Manager

C. by using Oracle Label Security options

D. by assigning a profile to the user

Correct Answer: *D*

Kitevn Highly Voted 3 years, 3 months ago
The answer is D
upvoted 5 times

Nick_144 Most Recent 3 years, 9 months ago
D is the answer
upvoted 3 times

efa 4 years, 2 months ago
CPU and memory limits are implemented with profiles.

CPU and session-oriented resource limits are managed through profiles. Profiles let you set limits for several resources, including CPU time, memory, and the number of logical reads performed during a user session or database call.
upvoted 1 times

Phoenix22 4 years, 4 months ago
Why is not B?
upvoted 1 times

ogdru 3 years, 8 months ago
because resource manager focus on allocate resources but 1. failed login must use profile
upvoted 4 times

Question #3 *Topic 1*

As a user of the ORCL database, you establish a database link to the remote HQ database such that all users in the ORCL database may access tables only from the SCOTT schema in the HQ database. SCOTT's password is TIGER. The service mane "HQ" is used to connect to the remote HQ database.
Which command would you execute to create the database link?

A. CREATE DATABASE LINK HQ USING 'HQ';

B. CREATE DATABASE LINK HQ CONNECT TO CURRENT_USER USING 'HQ';

C. CREATE PUBLIC DATABASE LINK HQ CONNECT TO scott IDENTIFIED BY tiger USING 'HQ';

D. CREATE DATABASE LINK HQ CONNECT TO scott IDENTIFIED BY tiger USING 'HQ';

Correct Answer: *C*

Nick_144 Highly Voted 3 years, 9 months ago
answer is C
upvoted 6 times

Kitevn Most Recent 3 years, 3 months ago
For all users in ORCL database may access the database link => we have to create PUBLIC DATABASE LINK
upvoted 4 times

Question #4 *Topic 1*

What happens if a maintenance window closes before a job that collects optimizer statistics completes?

A. The job is terminated and the gathered statistics are not saved.

B. The job is terminated but the gathered statistics are not published.

C. The job continues to run until all statistics are gathered.

D. The job is terminated and statistics for the remaining objects are collected the next time the maintenance window opens.

Correct Answer: *D*
The stop_on_window_close attribute controls whether the GATHER_STATS_JOB continues when the maintenance window closes. The default setting for the stop_on_window_close attribute is TRUE, causing Scheduler to terminate GATHER_STATS_JOB when the maintenance window closes. The remaining objects are then processed in the next maintenance window.
References:
https://docs.oracle.com/cd/B19306_01/server.102/b14211/stats.htm#g49431

Impaler Highly Voted 3 years, 11 months ago
D is indeed correct

The stop_on_window_close attribute controls whether the GATHER_STATS_JOB continues when the maintenance window closes. The default settir for the stop_on_window_close attribute is TRUE, causing Scheduler to terminate GATHER_STATS_JOB when the maintenance window closes. The remaining objects are then processed in the next maintenance window.

https://docs.oracle.com/cd/B19306_01/server.102/b14211/stats.htm#g49431
upvoted 8 times

Nick_144 Most Recent 3 years, 9 months ago
answer is C
upvoted 3 times

katetel 3 years, 5 months ago
no answer is d
upvoted 4 times

Question #5 *Topic 1*

You plan to create a database by using the Database Configuration Assistant (DBCA), with the following specifications:
Applications will connect to the database via a middle tier.
The number of concurrent user connections will be high.
The database will have mixed workload, with the execution of complex BI queries scheduled at night.
Which DBCA option must you choose to create the database?

A. a General Purpose database template with default memory allocation

B. a Data Warehouse database template, with the dedicated server mode option and AMM enabled

C. a General Purpose database template, with the shared server mode option and Automatic Memory Management (AMM) enabled

D. a default database configuration

Correct Answer: *C*
References:

Community vote distribution

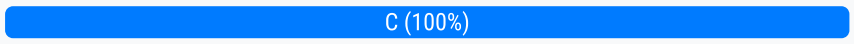

Earlg_ray 2 years, 1 month ago

Data warehouse : Users perform numerous, complex queries that process large volumes of data.
answer is b

https://docs.oracle.com/database/121/ADMQS/GUID-E59E0AE9-AF91-4DD2-B311-C92E3EFE9948.htm

upvoted 1 times

jackymak 2 years, 3 months ago

Selected Answer: C

I chose C

upvoted 1 times

jackymak 2 years, 3 months ago

Because AMM can distribute the memory to PGA in the working hours
and moving it back to SGA for the daily night batch jobs.

upvoted 1 times

Oracle2020 2 years, 12 months ago

The number of concurrent user connections will be high. It should be share served and AMM

upvoted 2 times

JMAN1 3 years, 5 months ago

Does anyone know real answer or why it is correct? thanks.

upvoted 1 times

Question #6 *Topic 1*

Which two statements are true about the logical storage structure of an Oracle database? (Choose two.)

A. An extent contains data blocks that are always physically contiguous on disk.

B. An extent can span multiple segments.

C. Each data block always corresponds to one operating system block.

D. It is possible to have tablespaces of different block sizes.

E. A data block is the smallest unit of I/O in data files.

Correct Answer: *DE*
References:

Cobra 2 years, 3 months ago
D&E is correct
upvoted 1 times

Grace1_2239478234 2 years, 5 months ago
D cannot be correct as you choose the block size per database not per tablespace - AE
You set the data block size for each Oracle database when you create the database.
upvoted 1 times

NorwayOracle 2 years, 4 months ago
D is correct. From the documentation:

Use the BLOCKSIZE clause of the CREATE TABLESPACE statement to create a tablespace with a block size different from the database standard block size. In order for the BLOCKSIZE clause to succeed, you must have already set the DB_CACHE_SIZE and at least one DB_nK_CACHE_SIZE initialization parameter.
upvoted 1 times

nathansoc 3 years, 5 months ago
entity-relationship diagram for physical and logical storage. The crow's foot notation represents a one-to-many relationship.
upvoted 1 times

nathansoc 3 years, 5 months ago
https://docs.oracle.com/cd/E11882_01/server.112/e40540/img/cncpt227.gif
upvoted 1 times

Desmon2175 3 years, 10 months ago
D,E correct
upvoted 4 times

JMAN1 3 years, 5 months ago
reference - https://docs.oracle.com/cd/A87860_01/doc/server.817/a76965/c02block.htm
upvoted 2 times

Question #7 *Topic 1*

Which two statements correctly describe the relationship between data files and logical database structures? (Choose two.)

A. A segment cannot span data files.

B. A data file can belong to only one tablespace.

C. An extent cannot span data files.

D. The size of an Oracle data block in a data file should be the same as the size of an OS block.

Correct Answer: *BC*
A single extent can never span data files.
https://docs.oracle.com/database/121/CNCPT/logical.htm#CNCPT1095

Cobra 2 years, 3 months ago
B&C is correct
upvoted 1 times

gio 2 years, 10 months ago
BC, C is true because of "should"
upvoted 2 times

gio 2 years, 10 months ago
BD, sorry! D is true because of "should"
upvoted 1 times

nathansoc 3 years, 5 months ago
https://docs.oracle.com/cd/E11882_01/server.112/e40540/logical.htm?sa=X&ved=2ahUKEwjV28eVqvvuAhUM_RQKHbW9DAAQ9QF6BAgJEAI#CNCPT3000
upvoted 1 times

Desmon2175 3 years, 10 months ago
A is faulse because asegment can contain informations from many datafile
D is faulse because a datablock can store on or more OS blocks.... B,C are correct
upvoted 3 times

Question #8 *Topic 1*

Which statement is true about the Log Writer process?

A. It writes when it receives a signal from the checkpoint process (CKPT).

B. It writes concurrently to all members of multiplexed redo log groups.

C. It writes after the Database Writer process writes dirty buffers to disk.

D. It writes when a user commits a transaction.

Correct Answer: *D*
References:
http://docs.oracle.com/cd/B19306_01/server.102/b14220/process.htm
(see log writer process (LGWR))

nabhan2001nabhan 6 months, 1 week ago
When a user process commits a transaction
– When an online redo log switch occurs
– When the redo log buffer is one-third full or contains 1 MB of buffered data
– Before a DBWn process writes modified buffers to disk
upvoted 1 times

Cobra 2 years, 3 months ago
only D is correct
upvoted 1 times

hggz 4 years ago
The LGWR process writes redo information from the redo log buffer to the online redo log files under a variety of circumstances:

■ When a user commits a transaction, even if this is the only transaction in the log buffer.
■ When the redo log buffer becomes one-third full.
■ When the buffer contains approximately 1MB of changed records. This total does not include deleted or inserted records.
■ When a database checkpoint is performed.
■ When 3 seconds have elapsed since the last redo log buffer write.
upvoted 2 times

nabhan2001nabhan 6 months, 1 week ago
not ckeckpoint .
upvoted 1 times

Nickhodges94 4 years, 6 months ago
I believe both B and D are true.

The LGWR background process does write to disk upon COMMIT:

"When a user issues a COMMIT statement, LGWR puts a commit record in the redo log buffer and writes it to disk immediately, along with the transaction's redo entries. The corresponding changes to data blocks are deferred until it is more efficient to write them. This is called a fast comr mechanism."

However, It also writes simultaneously to all multiplexed redo log files:

"LGWR writes synchronously to the active mirrored group of redo log files."

Reference :
https://docs.oracle.com/cd/B19306_01/server.102/b14220/process.htm#i7261
upvoted 2 times

hggz 4 years, 5 months ago
D is correct, B - no.

It writes concurrently to all multiplexed redo log groups, but not to all members in each group.
upvoted 7 times

Question #9 *Topic 1*

The ORCL database is configured to support shared server mode. You want to ensure that a user connecting remotely to the database instance has a one-to-one ratio between client and server processes.
Which connection method guarantees that this requirement is met?

A. connecting by using an external naming method

B. connecting by using the easy connect method

C. creating a service in the database by using the dbms_service.create_service procedure and using this service for creating a local naming service

D. connecting by using the local naming method with the server = dedicated parameter set in the tnsnames.ora file for the net service

E. connecting by using a directory naming method

Correct Answer: *D*

Impaler 3 years, 11 months ago

"To request a dedicated server connection when Oracle Database is configured for shared server, users must connect using a net service name tha is configured to use a dedicated server. Specifically, the net service name value should include the SERVER=DEDICATED clause in the connect descriptor."

So D is correct

upvoted 4 times

Question #10 *Topic 1*

Which two tasks can be performed on an external table? (Choose two.)

A. partitioning the table

B. creating an invisible index

C. updating the table by using an update statement

D. creating a public synonym

E. creating a view

Correct Answer: *DE*
http://docs.oracle.com/cd/B28359_01/server.111/b28310/tables013.htm#ADMIN01507
You can, for example select, join, or sort external table data. You can also create views and synonyms for external tables. However, no DML operations
(UPDATE, INSERT, or DELETE) are possible, and no indexes can be created, on external tables.

napura Highly Voted 3 years, 2 months ago
Validation:
This exam has been validated against Oracle Database 12.1.0.1.0.

https://education.oracle.com/oracle-database-12c-administration/pexam_1Z0-062

So correct answers are D and E
upvoted 6 times

LazyLH Most Recent 4 years, 5 months ago
Why would the answer not be ADE? I thought you could partition external tables.
https://docs.oracle.com/en/database/oracle/oracle-database/12.2/admin/managing-tables.html#GUID-2A801016-0399-4925-AD1B-A02683E81B
upvoted 1 times

TuxBingo 4 years ago
Hi.
The problem is they don't specified for wich version this options can apply, because in 12.2 you can create partitions in external tables.
upvoted 2 times

Question #11 *Topic 1*

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

A. It can contain a nested chain of jobs.

B. It can be used to implement dependency-based scheduling.

C. It cannot invoke the same program or nested chain in multiple steps in the chain.

D. It cannot have more than one dependency.

E. It can be executed using event-based or time-based schedules.

Correct Answer: *ABE*
References:

kanchaveli 2 years, 9 months ago
A,B,E
.....
upvoted 4 times

DANG257315 3 years, 4 months ago
https://docs.oracle.com/database/121/ADMIN/schedover.htm#ADMIN12357
upvoted 1 times

nathansoc 3 years, 5 months ago
https://docs.oracle.com/cd/B28359_01/server.111/b28310/schedover005.htm#ADMIN12362
upvoted 1 times

nathansoc 3 years, 5 months ago
https://docs.oracle.com/cd/B28359_01/server.111/b28310/scheduse009.htm#ADMIN10021
upvoted 1 times

Question #12 *Topic 1*

The HR user receives the following error while inserting data into the sales table:

```
ERROR at line 1:
ORA-01653: unable to extend table HR.SALES by 128 in tablespace USERS
```

On investigation, you find that the users tablespace uses Automatic Segment Space Management (ASSM). It is the default tablespace for the HR user with an unlimited quota on it.
Which two methods would you use to resolve this error? (Choose two.)

A. Altering the data file associated with the USERS tablespace to extend automatically

B. Adding a data file to the USERS tablespace

C. Changing segment space management for the USERS tablespace to manual

D. Creating a new tablespace with autoextend enabled and changing the default tablespace of the HR user to the new tablespace

E. Enabling resumable space allocation by setting the RESUMABLE_TIMEOUT parameter to a nonzero value

Correct Answer: *AB*

yosiw96816 3 years, 8 months ago
Agree - A, B
upvoted 1 times

Question #13 *Topic 1*

Which three factors influence the optimizer's choice of an execution plan? (Choose three.)

A. the optimizer_mode initialization parameter

B. operating system (OS) statistics

C. cardinality estimates

D. object statistics in the data dictionary

E. fixed baselines

Correct Answer: *ACD*
References:

hggz Highly Voted 3 years, 11 months ago
I think ADE

Init parameter (A)
Fixed Baseline (E) https://docs.oracle.com/database/121/TGSQL/tgsql_spm.htm#TGSQL618

C is not true. Remember, cardinality is a fancy word for “number of rows returned” by a SQL operation.
http://www.dba-oracle.com/t_sql_tuning_cardinality_plan.htm
upvoted 9 times

TuxBingo 3 years, 3 months ago
Agree with hggz, correct answer is: ADE
upvoted 3 times

yosiw96816 Highly Voted 3 years, 8 months ago
I agree with ACD
https://docs.oracle.com/cd/E18283_01/server.112/e16638/optimops.htm
upvoted 6 times

jokerking Most Recent 2 years, 9 months ago
I think ACE
upvoted 1 times

luizyto 2 years, 11 months ago
I think that C it's correct " Cardinality estimates must be as accurate as possible because they influence all aspects of the execution plan"
https://docs.oracle.com/database/121/TGSQL/tgsql_optcncpt.htm#GUID-5D925135-6CF3-40B8-8AEF-12159055DE1F
upvoted 1 times

luizyto 2 years, 11 months ago
Ok I think is D instead of C.
upvoted 2 times

luizyto 2 years, 11 months ago
I think that C it's correct " Cardinality estimates must be as accurate as possible because they influence all aspects of the execution plan"
https://docs.oracle.com/database/121/TGSQL/tgsql_optcncpt.htm#GUID-5D925135-6CF3-40B8-8AEF-12159055DE1F
upvoted 2 times

yukclam9 3 years, 5 months ago
C is incorrect. First we have plans then we have cardinality estimates.
upvoted 1 times

Question #14 *Topic 1*

Examine the resources consumed by a database instance whose current Resource Manager plan is displayed.

```
SQL> 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_WAITS   CPU_WAITS
CPU_WAIT_TIME
---------- ------------ ----------   ---------------   -------   -----
----
OLTP__ORDER__ENTRY         1            0              29690        467
6709
OTHES__GROUPS              0            0              5982366     4089
60425
SYS_GROUP                  1            0              2420704      914
19540
DSS_QUERIES                4            2              4594660     3004
55700
```

Which two statements are true? (Choose two.)

A. An attempt to start a new session by a user belonging to DSS_QUERIES fails with an error.

B. An attempt to start a new session by a user belonging to OTHER_GROUPS fails with an error.

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 I/O waits and latch or enqueue contention.

E. A user belonging to the DSS__QUERIES resource consumer group can create a new session but the session will be queued.

Correct Answer: *CE*

hggz Highly Voted 4 years, 5 months ago

Correct.

https://docs.oracle.com/cd/B28359_01/server.111/b28310/dbrm009.htm#CHDCJCGB

upvoted 7 times

napura Most Recent 3 years, 2 months ago

CPU_WAIT_TIME NUMBER=Cumulative amount of time that sessions waited for CPU because of resource management. This does not include wai due to latch or enqueue contention, I/O waits, and so on.

https://docs.oracle.com/cd/E18283_01/server.112/e17110/dynviews_2156.htm
Correct answers are C&E

upvoted 3 times

x38 4 years, 3 months ago

B should fail too because there is not "OTHER_GROUPS_ resource group (it's OTHES_GROUPS in the display)

upvoted 1 times

joker_li 3 years, 2 months ago

B is correct , because active session and seq is zero

upvoted 1 times

Question #15 *Topic 1*

Which action takes place when a file checkpoint occurs?

A. The checkpoint position is advanced in the checkpoint queue.

B. All buffers for a checkpointed file that were modified before a specific SCN are written to disk by DBWn and the SCN is stored in the control file.

C. The Database Writer process (DBWn) writes all dirty buffers in the buffer cache to data files.

D. The Log Writer process (LGWR) writes all redo entries in the log buffer to online redo log files.

Correct Answer: *B*

jackymak 2 years, 3 months ago
I will try C.
upvoted 2 times

Cobra 2 years, 3 months ago
B is correct
upvoted 2 times

Dragon67 1 year, 12 months ago
please can i get your email, got some question on this exam
upvoted 1 times

exampro46 3 years, 2 months ago
FILE LEVEL CHECKPOINT
File level checkpoint is writing dirty blocks associated with set of files belonging to a tablespace. This happens when we put a tablespace in begin backup mode or when we take a tablespace offline or when we make tablespace readonly. Oracle writes all dirty blocks associated with datafiles that tablespace to database before it changes the status of that tablespace. That means the answer is B. Not C because there written "all dirty buffers", but in B written "all buffers for a checkpointed file" which is correct.
upvoted 2 times

napura 3 years, 2 months ago
Checkpoint Process (CKPT)

The checkpoint process (CKPT) updates the control file and data file headers with checkpoint information and signals DBWn to write blocks to dis Checkpoint information includes the checkpoint position, SCN, location in online redo log to begin recovery, and so on.
https://docs.oracle.com/cd/E11882_01/server.112/e40540/process.htm#BABHFGFF

answer is B
upvoted 2 times

yosiw96816 3 years, 8 months ago
Agree with B - Tablespace and data file checkpoints:
The database writes to disk all buffers modified by redo before a specific target. A tablespace checkpoint is a set of data file checkpoints, one for each data file in the tablespace. These checkpoints occur in a variety of situations, including making a tablespace read-only or taking it offline normal, shrinking a data file, or executing ALTER TABLESPACE BEGIN BACKUP.
https://docs.oracle.com/database/121/CNCPT/startup.htm#CNCPT89045
upvoted 4 times

dad 3 years, 9 months ago
all dirty buffers are written to disk means a full checkpoint and there will never be a full checkpoint except in two circumstances: an orderly shutdown and at the DBA's request. So, the answer is B
upvoted 2 times

Gah 3 years, 10 months ago
B Or c? I Need help in this please
upvoted 1 times

santhoshpsk 4 years, 1 month ago
The question is about "file" checkpoint in which case DB writer won't write "all" dirty buffers but only writes dirty buffers belongs to that particula file. So the answer should be B.
upvoted 4 times

hggz 4 years ago
I think you are right: https://docs.oracle.com/database/121/BRADV/glossary.htm#GUID-22212AAE-6F3C-4E84-B940-5AB2261A7EDA
upvoted 2 times

joe_doe 4 years, 1 month ago

C confirmed DBWr will do nothing when user commit a transaction,but it'll work when check point occured
upvoted 1 times

hggz 4 years, 5 months ago
I agree with C

B is for incremental check point:

"During an incremental database checkpoint, the control file is updated with the checkpoint position; data files are not updated. Data file headers are updated with checkpoint information by the DBWn process when dirty buffers are written to the files."
upvoted 2 times

hggz 4 years ago
I think it's may be C: data file checkpoint

"A data structure that defines an SCN in the redo thread of a database for a particular data file. Every data file has a checkpoint SCN, which yo can view in V$DATAFILE.CHECKPOINT_CHANGE#. All changes with an SCN lower than this SCN are guaranteed to be in the data file."

https://docs.oracle.com/database/121/BRADV/glossary.htm#GUID-22212AAE-6F3C-4E84-B940-5AB2261A7EDA
upvoted 1 times

bkk 4 years, 5 months ago
I also say C, as B is for incremental check point.
upvoted 1 times

verratti 4 years, 5 months ago
I think The answer is C
upvoted 4 times

Question #16 *Topic 1*

Examine the structure of the sales table, which is stored in a locally managed tablespace with Automatic Segment Space Management (ASSM) enabled.

```
Name                 Null?                Type
-------------------  -------------------  -----------------
PROD_ID              NOT NULL             NUMBER
CUST_ID              NOT NULL             NUMBER
TIME_ID              NOT NULL             DATE
CHANNEL_ID           NOT NULL             NUMBER
PROMO_ID             NOT NULL             NUMBER
QUANTITY_SOLD        NOT NULL             NUMBER (10,2)
AMOUNT_SOLD          NOT NULL             NUMBER (10,2)
```

You want to perform online segment shrink to reclaim fragmented free space below the high water mark.
What should you ensure before the start of the operation?

A. Row movement is enabled.

B. Referential integrity constraints for the table are disabled.

C. No queries are running on this table.

D. Extra disk space equivalent to the size of the segment is available in the tablespace.

E. No pending transaction exists on the table.

Correct Answer: *A*

bigswole 2 years, 9 months ago
row movement must be enabled before shrink operation A is correct
upvoted 2 times

Question #17 *Topic 1*

Which task would you recommend before using the Database Upgrade Assistant (DBUA) to upgrade a single-instance Oracle 11g R2 database to Oracle
Database 12c?

A. shutting down the database instance that is being upgraded

B. executing the catctl.pl script to run the upgrade processes in parallel

C. running the Pre-Upgrade Information Tool

D. copying the listener.ora file to the new ORACLE_HOME

Correct Answer: *C*
References:
http://docs.oracle.com/cd/E11882_01/server.112/e23633/upgrade.htm#UPGRD12395

sela 1 year, 10 months ago
C
For all upgrades, before using DBUA to upgrade your system, Oracle strongly recommends that you run the Pre-Upgrade Information Tool manually. DBUA runs the Pre-Upgrade Information Tool as part of the prerequisite checks it performs before starting the upgrade. However, to reduce downtime, Oracle recommends that you run the Pre-Upgrade Information Tool as part of your upgrade planning, so that you can analyze the database, and take proactive steps before your planned upgrade date.
https://docs.oracle.com/en/database/oracle/oracle-database/12.2/upgrd/upgrading-oracle-database-upgrade-assistant-dbua.html#GUID-04161A85-6A67-4404-B7A7-69A3713C5F99
upvoted 1 times

Cobra 2 years, 3 months ago
why A isent correct ?
upvoted 1 times

jackymak 2 years, 6 months ago
Dose it requested to shut down the DB?
upvoted 1 times

Question #18 *Topic 1*

Your database is open and the listener LISTENER is up. You issue the command:

LSNRCTL> RELOAD -
What is the effect of reload on sessions that were originally established by listener?

A. Only sessions based on static listener registrations are disconnected.

B. Existing connections are not disconnected; however, they cannot perform any operations until the listener completes the re-registration of the database instance and service handlers.

C. The sessions are not affected and continue to function normally.

D. All the sessions are terminated and active transactions are rolled back.

Correct Answer: *C*

Cobra 2 years, 3 months ago
c is correct
upvoted 1 times

JMAN1 2 years, 7 months ago
I hope this will help you study :)
If the listener is running and modifications are made to the listener.ora file manually, with Oracle Net Manager or with Enterprise Manager, you must reload the listener to refresh the listener with the most current information. The reload command rereads the listener.ora file for the new definitions. As you can see, it is not necessary to stop and start the listener to reload it. Although stopping and restarting the listener can also accomplish a reload, using the reload command is better because the listener is not actually stopped, which makes this process more efficient. The following code shows an example of the reload command:
LSNRCTL> reload listener1
Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=btlnx63)(PORT=1522)))
The command completed successfully
LSNRCTL>
upvoted 1 times

Question #19 *Topic 1*

Which statement is true regarding the startup of a database instance?

A. The instance does not start up normally and requires manual media recovery after a shutdown using the abort option.

B. Uncommitted transactions are rolled back during the startup of the database instance after a shutdown using the immediate option.

C. There is no difference in the underlying mechanics of the startup whether the database is shut down by using the immediate option or the abort option.

D. Media recovery is required when the database is shut down by using either the immediate option or the abort option.

E. Instance recovery is not required if the database instance was shut down by using SHUTDOWN IMMEDIATE.

Correct Answer: *E*
References:
http://docs.oracle.com/cd/A87860_01/doc/server.817/a76956/start.htm

Cobra 2 years, 3 months ago
E is correct
upvoted 2 times

Question #20 *Topic 1*

Examine the memory-related parameters set in the SPFILE of an Oracle database:

```
memory_max_target=6G
memory_target=5G
pga_aggregate_target=500M
sga_max_size=0
sga_target=0
```

Which statement is true?

A. Only SGA components are sized automatically.

B. Memory is dynamically re-allocated between the SGA and PGA as needed.

C. The size of the PGA cannot grow automatically beyond 500 MB.

D. The value of the MEMORY_TARGET parameter cannot be changed dynamically.

Correct Answer: *B*

Cobra 2 years, 3 months ago

B is correct

upvoted 1 times

JMAN1 2 years, 7 months ago

B is correct as _gio_ wrote.

racle Database can manage the SGA memory and instance PGA memory completely automatically. You designate only the total memory size to b used by the instance, and Oracle Database dynamically exchanges memory between the SGA and the instance PGA as needed to meet processin demands.

upvoted 2 times

gio 2 years, 10 months ago

https://docs.oracle.com/database/121/ADMIN/memory.htm#ADMIN11201

upvoted 3 times

Question #21 *Topic 1*

Which two statements are true about extents? (Choose two.)

A. Blocks belonging to an extent can be spread across multiple data files.

B. Data blocks in an extent are logically contiguous but can be non-contiguous on disk.

C. The blocks of a newly allocated extent, although free, may have been used before.

D. Data blocks in an extent are automatically reclaimed for use by other objects in a tablespace when all the rows in a table are deleted.

Correct Answer: *BC*

khan Highly Voted 5 years, 3 months ago

Answer BC

https://docs.oracle.com/database/121/CNCPT/logical.htm#CNCPT1063
Deallocation of Extents:
In general, the extents of a user segment do not return to the tablespace unless you drop the object using a DROP statement.

For example, if you delete all rows in a table, then the database does not reclaim the data blocks for use by other objects in the tablespace. You c also drop the segment using the DBMS_SPACE_ADMIN package.

upvoted 8 times

Cobra Most Recent 2 years, 3 months ago

B & C is correct

upvoted 1 times

Question #22 *Topic 1*

You execute the commands:

```
SQL>CREATE USER sidney
    IDENTIFIED BY out_standing1
    DEFAULT TABLESPACE users
    QUOTA 10M ON users
    TEMPORARY TABLESPACE temp
    ACCOUNT UNLOCK;

SQL> GRANT CREATE SESSION TO sidney;
```

Which two statements are true? (Choose two.)

A. The create user command fails if any role with the name Sidney exists in the database.

B. The user Sidney can connect to the database instance but cannot perform sort operations because no space quota is specified for the temp tablespace.

C. The user Sidney is created but cannot connect to the database instance because no profile is default.

D. The user Sidney can connect to the database instance but requires relevant privileges to create objects in the users tablespace.

E. The user Sidney is created and authenticated by the operating system.

Correct Answer: *AD*

References:

napura Highly Voted 3 years, 2 months ago

Correct answers are A & D

upvoted 5 times

Cobra Most Recent 2 years, 3 months ago

A & D is correct

upvoted 1 times

Question #23 *Topic 1*

Examine the query and its output:

```
SQL> SELECT reason, metric_value FROM dba_outstanding_alerts;

REASON                                   METRIC_VALUE
---------------------------------------  ------------------
Tablespace [TEST] is                     28.125
[28 percent] full

Metrics "Current Logons Count"           29
is at 29

Metrics "Database Time Spent             99.0375405
Waiting (%)" is at 99.03754 for
event class "Application"

db_recovery_file_dest_size of            97
4294967296 bytes is 97.298 used
and has 116228096 remaining
bytes available.
```

After 30 minutes, you execute the same query:

```
SQL> SELECT reason,metric_value FROM dba_outstanding_alerets;

REASON                                   METRIC_VALUE
---------------------------------------  --------------
Tablespace [TEST] is                     28.125
[28 percent full
```

What might have caused three of the alerts to disappear?

A. The threshold alerts were cleared and transferred to DBA_ALERT_HISTORY.

B. An Automatic Workload Repository (AWR) snapshot was taken before the execution of the second query.

C. An Automatic Database Diagnostic Monitor (ADOM) report was generated before the execution of the second query.

D. The database instance was restarted before the execution of the second query.

Correct Answer: *D*

hggz Highly Voted 4 years, 5 months ago

A is correct

"Threshold alerts are written to DBA_OUTSTANDING_ALERTS. Nonthreshold alerts are written
only to DBA_ALERT_HISTORY. Entries from DBA_OUTSTANDING_ALERTS are cleared when the alert condition is cleared."

upvoted 9 times

SantiBZ_07032022_1744 Most Recent 2 years, 2 months ago

I think so A is correct

upvoted 1 times

luizyto 2 years, 10 months ago

Please A it's correct change it to A.

upvoted 1 times

jobao 3 years, 6 months ago

A is correct

upvoted 2 times

Gah 3 years, 10 months ago

A is the correct answer?

upvoted 2 times

santhoshpsk 4 years, 1 month ago

Yes. A is correct. Those alerts are state-full alerts.

upvoted 1 times

vineetarneja 4 years, 3 months ago

yeah. i also think A is correct because db_recovery_file_dest_size alert will still be there even after the instance restart.

upvoted 1 times

verratti 4 years, 5 months ago

A is correct

upvoted 3 times

Question #24 *Topic 1*

Which two statements are true? (Choose two.)

A. A role cannot be assigned external authentication.

B. A role can be granted to other roles.

C. A role can contain both system and object privileges.

D. The predefined resource role includes the unlimited_tablespace privilege.

E. All roles are owned by the sys user.

F. The predefined connect role is always automatically granted to all new users at the time of their creation.

Correct Answer: *BC*

References:

Cobra 2 years, 3 months ago

B & C is correct

upvoted 1 times

Nickhodges94 4 years, 6 months ago

This is new in Oracle 12c. The RESOURCE Role used to grant unlimited tablespace, however it stopped as of Oracle 12.1.0.2

Reference: https://docs.oracle.com/database/121/DBSEG/release_changes.htm#DBSEG000

upvoted 4 times

Question #25 *Topic 1*

Identify three valid options for adding a pluggable database (PDB) to an existing multitenant container database (CDB).

A. Use the CREATE PLUGGABLE DATABASE statement to create a PDB using the files from the SEED.

B. Use the CREATE DATABASE . . . ENABLE PLUGGABLE DATABASE statement to provision a PDB by copying file from the SEED.

C. Use the DBMS_PDB package to clone an existing PDB.

D. Use the DBMS_PDB package to plug an Oracle 12c non-CDB database into an existing CDB.

E. Use the DBMS_PDB package to plug an Oracle 11 g Release 2 (11.2.0.3.0) non-CDB database into an existing CDB.

Correct Answer: *ACD*
Use the CREATE PLUGGABLE DATABASE statement to create a pluggable database (PDB).
This statement enables you to perform the following tasks:
* (A) Create a PDB by using the seed as a template
Use the create_pdb_from_seed clause to create a PDB by using the seed in the multitenant container database (CDB) as a template. The files associated with the seed are copied to a new location and the copied files are then associated with the new PDB.
* (C) Create a PDB by cloning an existing PDB
Use the create_pdb_clone clause to create a PDB by copying an existing PDB (the source PDB) and then plugging the copy into the CDB. The files associated with the source PDB are copied to a new location and the copied files are associated with the new PDB. This operation is called cloning a PDB.
The source PDB can be plugged in or unplugged. If plugged in, then the source PDB can be in the same CDB or in a remote CDB. If the source PDB is in a remote CDB, then a database link is used to connect to the remote CDB and copy the files.
* Create a PDB by plugging an unplugged PDB or a non-CDB into a CDB
Use the create_pdb_from_xml clause to plug an unplugged PDB or a non-CDB into a CDB, using an XML metadata file.

Cobra 2 years, 3 months ago
A & C & D
upvoted 2 times

Question #26 *Topic 1*

Your database supports a DSS workload that involves the execution of complex queries: Currently, the library cache contains the ideal workload for analysis. You want to analyze some of the queries for an application that are cached in the library cache.
What must you do to receive recommendations about the efficient use of indexes and materialized views to improve query performance?

A. Create a SQL Tuning Set (STS) that contains the queries cached in the library cache and run the SQL Tuning Advisor (STA) on the workload captured in the STS.

B. Run the Automatic Workload Repository Monitor (ADDM).

C. Create an STS that contains the queries cached in the library cache and run the SQL Performance Analyzer (SPA) on the workload captured in the STS.

D. Create an STS that contains the queries cached in the library cache and run the SQL Access Advisor on the workload captured in the STS.

Correct Answer: *D*
* SQL Access Advisor is primarily responsible for making schema modification recommendations, such as adding or dropping indexes and materialized views.
SQL Tuning Advisor makes other types of recommendations, such as creating SQL profiles and restructuring SQL statements.
* The query optimizer can also help you tune SQL statements. By using SQL Tuning Advisor and SQL Access Advisor, you can invoke the query optimizer in advisory mode to examine a SQL statement or set of statements and determine how to improve their efficiency. SQL Tuning Advisor and SQL Access Advisor can make various recommendations, such as creating SQL profiles, restructuring SQL statements, creating additional indexes or materialized views, and refreshing optimizer statistics.
Note:
* Decision support system (DSS) workload
* The library cache is a shared pool memory structure that stores executable SQL and PL/SQL code. This cache contains the shared SQL and PL/SQL areas and control structures such as locks and library cache handles.
References:

bkk 1 year, 11 months ago
apologies, lesson 22 page 19.
upvoted 1 times

bkk 1 year, 11 months ago
Correct answer is A
Student guide lesson 22 page 22
SQL Tuning Advisor analyzes one or more SQL statements one at a time. It examines statistics, SQL profiles, indexes, materialized views, and restructured SQL
upvoted 1 times

Question #27 *Topic 1*

The following parameter are set for your Oracle 12c database instance:
OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES=FALSE
OPTIMIZER_USE_SQL_PLAN_BASELINES=TRUE
You want to manage the SQL plan evolution task manually. Examine the following steps:
1. Set the evolve task parameters.
2. Create the evolve task by using the DBMS_SPM.CREATE_EVOLVE_TASK function.
3. Implement the recommendations in the task by using the DBMS_SPM.IMPLEMENT_EVOLVE_TASK function.
4. Execute the evolve task by using the DBMS_SPM.EXECUTE_EVOLVE_TASK function.
5. Report the task outcome by using the DBMS_SPM.REPORT_EVOLVE_TASK function.
Identify the correct sequence of steps:

A. 2, 4, 5

B. 2, 1, 4, 3, 5

C. 1, 2, 3, 4, 5

D. 1, 2, 4, 5

Correct Answer: *B*
* Evolving SQL Plan Baselines

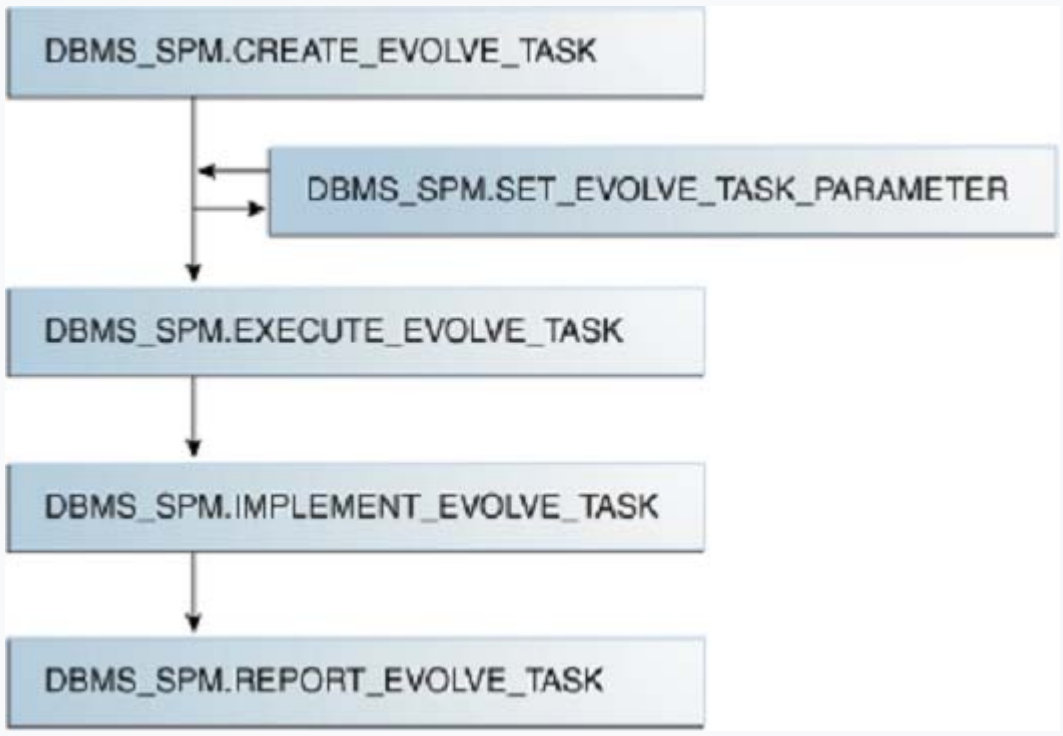

*
2. Create the evolve task by using the DBMS_SPM.CREATE_EVOLVE_TASK function.
This function creates an advisor task to prepare the plan evolution of one or more plans for a specified SQL statement. The input parameters can be a SQL handle, plan name or a list of plan names, time limit, task name, and description.
1. Set the evolve task parameters.

SET_EVOLVE_TASK_PARAMETER -
This function updates the value of an evolve task parameter. In this release, the only valid parameter is TIME_LIMIT.
4. Execute the evolve task by using the DBMS_SPM.EXECUTE_EVOLVE_TASK function.
This function executes an evolution task. The input parameters can be the task name, execution name, and execution description. If not specified, the advisor generates the name, which is returned by the function.
3: IMPLEMENT_EVOLVE_TASK
This function implements all recommendations for an evolve task. Essentially, this function is equivalent to using ACCEPT_SQL_PLAN_BASELINE for all recommended plans. Input parameters include task name, plan name, owner name, and execution name.
5. Report the task outcome by using the DBMS_SPM_EVOLVE_TASK function.
This function displays the results of an evolve task as a CLOB. Input parameters include the task name and section of the report to include.
References:

Question #28 *Topic 1*

In a recent Automatic Workload Repository (AWR) report for your database, you notice a high number of buffer busy waits. The database consists of locally managed tablespaces with free list managed segments.
On further investigation, you find that buffer busy waits is caused by contention on data blocks.
Which option would you consider first to decrease the wait event immediately?

A. Decreasing PCTUSED

B. Decreasing PCTFREE

C. Increasing the number of DBWN process

D. Using Automatic Segment Space Management (ASSM)

E. Increasing db_buffer_cache based on the V$DB_CACHE_ADVICE recommendation

Correct Answer: *D*
* Automatic segment space management (ASSM) is a simpler and more efficient way of managing space within a segment. It completely eliminates any need to specify and tune the pctused,freelists, and freelist groups storage parameters for schema objects created in the tablespace. If any of these attributes are specified, they are ignored.
* Oracle introduced Automatic Segment Storage Management (ASSM) as a replacement for traditional freelists management which used one-way linked-lists to manage free blocks with tables and indexes. ASSM is commonly called "bitmap freelists" because that is how Oracle implement the internal data structures for free block management.
Note:
* Buffer busy waits are most commonly associated with segment header contention onside the data buffer pool (db_cache_size, etc.).
* The most common remedies for high buffer busy waits include database writer (DBWR) contention tuning, adding freelists (or ASSM), and adding missing indexes.

Question #29 *Topic 1*

Examine this command:
SQL > exec DBMS_STATS.SET_TABLE_PREFS (SH, CUSTOMERS, PUBLISH, false);
Which three statements are true about the effect of this command? (Choose three.)

A. Statistics collection is not done for the CUSTOMERS table when schema stats are gathered.

B. Statistics collection is not done for the CUSTOMERS table when database stats are gathered.

C. Any existing statistics for the CUSTOMERS table are still available to the optimizer at parse time.

D. Statistics gathered on the CUSTOMERS table when schema stats are gathered are stored as pending statistics.

E. Statistics gathered on the CUSTOMERS table when database stats are gathered are stored as pending statistics.

Correct Answer: *CDE*
* SET_TABLE_PREFS Procedure
This procedure is used to set the statistics preferences of the specified table in the specified schema.
* Example:

Using Pending Statistics -
Assume many modifications have been made to the employees table since the last time statistics were gathered. To ensure that the cost-based optimizer is still picking the best plan, statistics should be gathered once again; however, the user is concerned that new statistics will cause the optimizer to choose bad plans when the current ones are acceptable. The user can do the following:
EXEC DBMS_STATS.SET_TABLE_PREFS('hr', 'employees', 'PUBLISH', 'false');
By setting the employees tables publish preference to FALSE, any statistics gather from now on will not be automatically published. The newly gathered statistics will be marked as pending.

Question #30 *Topic 1*

Examine the following impdp command to import a database over the network from a pre-12c Oracle database (source):

```
$>   impdp <user_name> full=Y network_link=hrdb_test transportable=always
     transport_datafiles=
          '/u01/app/oracle/oradata/hrdb/sales01.dbf',
          '/u01/app/oracle/oradata/hrdb/cust01.dbf',
          '/u01/app/oracle/oradata/hrdb/emp01.dbf',
     version=12 logfile=import.log
```

Which three are prerequisites for successful execution of the command? (Choose three.)

A. The import operation must be performed by a user on the target database by a user with the DATAPUMP_IMP_FULL_DATABASE role, and the database link must connect to a user with the DATAPUMP_EXP_FULL_DATABASE role on the source database.

B. All the user-defined tablespaces must be in read-only mode on the source database.

C. The export dump file must be created before starting the import on the target database.

D. The source and target database must be running on the same operating system (OS) with the same endianness.

E. The impdp operation must be performed by the same user that performed the expdp operation.

Correct Answer: *ABD*

In this case we have run the impdp without performing any conversion if endian format is different then we have to first perform conversion.

Oracle2020 2 years, 11 months ago

If you read the command, you will see that it does not import data from any dumpfile. It does it through a network link, and it does not convert from endian format either. Is correct A, B, D

upvoted 4 times

jackymak 2 years, 6 months ago

Does that mean I need to 'alter database open read only' the source database for the data transfer?

upvoted 1 times

boscoDBA 3 years, 2 months ago

can you import on target database without export dump file?

upvoted 1 times

jackymak 2 years, 6 months ago

It is extracting the data from data files rather than data dump files.

upvoted 1 times

Question #31 *Topic 1*

Which two are true concerning a multitenant container database with three pluggable database? (Choose two.)

A. All administration tasks must be done to a specific pluggable database.

B. The pluggable databases increase patching time.

C. The pluggable databases reduce administration effort.

D. The pluggable databases are patched together.

E. Pluggable databases are only used for database consolidation.

Correct Answer: *CD*

aihamza 1 month, 3 weeks ago

C & D are correct answers.

upvoted 1 times

sela 1 year, 10 months ago

A. All administration tasks must be done to a specific pluggable database.False
B. The pluggable databases increase patching time.False
C. The pluggable databases reduce administration effort.True
D. The pluggable databases are patched together.True
E. Pluggable databases are only used for database consolidation.False

upvoted 1 times

Question #32 *Topic 1*

Examine the current value for the following parameters in your database instance:

SGA_MAX_SIZE = 1024M -

SGA_TARGET = 700M -

DB_8K_CACHE_SIZE = 124M -

LOG_BUFFER = 200M -
You issue the following command to increase the value of DB_8K_CACHE_SIZE:
SQL> ALTER SYSTEM SET DB_8K_CACHE_SIZE=140M;
Which statement is true?

A. It fails because the DB_8K_CACHE_SIZE parameter cannot be changed dynamically.

B. It succeeds only if memory is available from the autotuned components if SGA.

C. It fails because an increase in DB_8K_CACHE_SIZE cannot be accommodated within SGA_TARGET.

D. It fails because an increase in DB_8K_CACHE_SIZE cannot be accommodated within SGA_MAX_SIZE.

Correct Answer: *D*
* The SGA_TARGET parameter can be dynamically increased up to the value specified for the SGA_MAX_SIZE parameter, and it can also be reduced.
* Example:
For example, suppose you have an environment with the following configuration:

SGA_MAX_SIZE = 1024M -

SGA_TARGET = 512M -

DB_8K_CACHE_SIZE = 128M -
In this example, the value of SGA_TARGET can be resized up to 1024M and can also be reduced until one or more of the automatically sized components reaches its minimum size. The exact value depends on environmental factors such as the number of CPUs on the system. However, the value of
DB_8K_CACHE_SIZE remains fixed at all times at 128M
* DB_8K_CACHE_SIZE

Size of cache for 8K buffers -
* For example, consider this configuration:

SGA_TARGET = 512M -

DB_8K_CACHE_SIZE = 128M -
In this example, increasing DB_8K_CACHE_SIZE by 16 M to 144M means that the 16M is taken away from the automatically sized components. Likewise, reducing DB_8K_CACHE_SIZE by 16M to 112M means that the 16M is given to the automatically sized components.

Modasser Highly Voted 3 years, 6 months ago
D is the correct answer
ref: https://docs.oracle.com/cd/B28359_01/server.111/b28310/memory004.htm#ADMIN11203
upvoted 5 times

Ray520 2 years, 3 months ago
Please can you explain why it's D? After reading the document and trying to understand it seems B is correct.
upvoted 1 times

IamGo0ke 3 years ago
Answer is B
ref: https://docs.oracle.com/cd/B28359_01/server.111/b28310/memory004.htm#ADMIN11218

Increasing DB_8K_CACHE_SIZE by 16 M to 144M means that the 16M is taken away from the automatically sized components. Likewise, reducing DB_8K_CACHE_SIZE by 16M to 112M means that the 16M is given to the automatically sized components.

upvoted 7 times

soniat 3 years ago

have taking the test?

upvoted 2 times

IamGo0ke 2 years, 12 months ago

Yes, I did today and I passed. Quite a number of the questions here came out with other few new questions. The set of questions here w a great help.

upvoted 2 times

luizyto Most Recent 1 year, 10 months ago

Answer is B, we always configure sga_target =sga_max_size and sometimes we configure cache_size to n_value without reducing the sga_target, s the static value for a components of sga is taken from sga_target. It does not make sense to sum sga_target+cache_size. It is sga_target-cache_size=memory available for automatic components of sga not defined.

upvoted 1 times

Cobra 2 years, 3 months ago

in my openion B is correct

upvoted 2 times

Cobra 2 years, 3 months ago

B is correct answer

upvoted 2 times

Ray520 2 years, 4 months ago

Log_buffer and db_nk_cache_size are manually sized sga parameters and take their value from sga_max_size and not sga_target. So log_buffer an db_8k_cache_size in this example cannot exceed 324M (sga_max_size - sga_target = 324M in this example). Hence the increase to 140M will fail, s D is indeed correct.
Reference:
https://docs.oracle.com/cd/B28359_01/server.111/b28310/memory004.htm#ADMIN11203

upvoted 2 times

Ray520 2 years, 4 months ago

Log_buffer and db_nk_cache_size are manually sized sga parameters and take their value from sga_max_size and not sga_target. So log_buffer an db_8k_cache_size in this example cannot exceed 340M (sga_max_size - sga_target = 340M in this example). Hence the increase to 140M will fail, s D is indeed correct.
Reference:
https://docs.oracle.com/cd/B28359_01/server.111/b28310/memory004.htm#ADMIN11203

upvoted 1 times

Aleksandar 2 years, 10 months ago

The answer is B

upvoted 4 times

Phoenix22 4 years, 4 months ago

Answer is surely B.

As D says ...
It fails because an increase in DB_8K_CACHE_SIZE cannot be accommodated within
SGA_MAX_SIZE.

while having value of 124M , and willing to increase it to 144M , which is 20MB extra . For this SGA size it's with granule of 4MB. So it should work fine under SGA_TARGET.

SGA_MAX_SIZE is set to 1024M while SGA_TARGET=700M, so talking about SMS is not relevant here.

upvoted 3 times

hggz 4 years, 5 months ago

B is correct.

Ref: https://docs.oracle.com/cd/B28359_01/server.111/b28310/memory004.htm#ADMIN11202

upvoted 3 times

verratti 4 years, 5 months ago

i think B is right

upvoted 3 times

Question #33 *Topic 1*

Which three statements are true concerning unplugging a pluggable database (PDB)? (Choose three.)

A. The PDB must be open in read only mode.

B. The PDB must be dosed.

C. The unplugged PDB becomes a non-CDB.

D. The unplugged PDB can be plugged into the same multitenant container database (CDB)

E. The unplugged PDB can be plugged into another CDB.

F. The PDB data files are automatically removed from disk.

Correct Answer: *BDE*
B, not A: The PDB must be closed before unplugging it.
D: An unplugged PDB contains data dictionary tables, and some of the columns in these encode information in an endianness-sensitive way. There is no supported way to handle the conversion of such columns automatically. This means, quite simply, that an unplugged PDB cannot be moved across an endianness difference.
E (not F): To exploit the new unplug/plug paradigm for patching the Oracle version most effectively, the source and destination CDBs should share a filesystem so that the PDBs datafiles can remain in place.
References:

aihamza 1 month, 3 weeks ago
B, D & E are correct answers.
upvoted 1 times

sela 1 year, 10 months ago
Answer: C,D,E
upvoted 1 times

jackymak 2 years, 6 months ago
B is NOT dosed BUT closed.
upvoted 1 times

Question #34 *Topic 1*

Examine the following command:
CREATE TABLE (prod_id number(4),
Prod_name varchar2 (20),
Category_id number(30),
Quantity_on_hand number (3) INVISIBLE);
Which three statements are true about using an invisible column in the PRODUCTS table? (Choose three.)

A. The %ROWTYPE attribute declarations in PL/SQL to access a row will not display the invisible column in the output.

B. The DESCRIBE commands in SQL *Plus will not display the invisible column in the output.

C. Referential integrity constraint cannot be set on the invisible column.

D. The invisible column cannot be made visible and can only be marked as unused.

E. A primary key constraint can be added on the invisible column.

Correct Answer: *ABE*
AB: You can make individual table columns invisible. Any generic access of a table does not show the invisible columns in the table. For example, the following operations do not display invisible columns in the output:
* SELECT * FROM statements in SQL
* DESCRIBE commands in SQL*Plus
* %ROWTYPE attribute declarations in PL/SQL
* Describes in Oracle Call Interface (OCI)
Incorrect:
Not D: You can make invisible columns visible.
You can make a column invisible during table creation or when you add a column to a table, and you can later alter the table to make the same column visible.
References:

sela 1 year, 10 months ago
a,b,e correction
12C New Feature : Invisible Columns (Doc ID 1566790.1)
https://oracle-base.com/articles/12c/invisible-columns-12cr1
upvoted 1 times

sela 1 year, 10 months ago
a,d,e
12C New Feature : Invisible Columns (Doc ID 1566790.1)
https://oracle-base.com/articles/12c/invisible-columns-12cr1
upvoted 1 times

jackymak 2 years, 6 months ago
If the column cannot be found by DESCRIBE, how do I know the invisible column is invisible but not 'not existed'
upvoted 1 times

jackymak 2 years, 6 months ago
SQL> SET COLINVISIBLE ON
SQL> DESC tab1
upvoted 1 times

gio 2 years, 10 months ago
I have tried to reference with foreign key constraint a invisible columns and it works
upvoted 1 times

Question #35 *Topic 1*

You wish to enable an audit policy for all database users, except SYS, SYSTEM, and SCOTT.
You issue the following statements:
SQL> AUDIT POLICY ORA_DATABASE_PARAMETER EXCEPT SYS;
SQL> AUDIT POLICY ORA_DATABASE_PARAMETER EXCEPT SYSTEM;
SQL> AUDIT POLICY ORA_DATABASE_PARAMETER EXCEPT SCOTT;
For which database users is the audit policy now active?

A. All users except SYS

B. All users except SCOTT

C. All users except sys and SCOTT

D. All users except sys, system, and SCOTT

Correct Answer: *B*
If you run multiple AUDIT statements on the same unified audit policy but specify different EXCEPT users, then Oracle Database uses the last exception user list, not any of the users from the preceding lists. This means the effect of the earlier AUDIT POLICY ... EXCEPT statements are overridden by the latest AUDIT
POLICY ... EXCEPT statement.
Note:
* The ORA_DATABASE_PARAMETER policy audits commonly used Oracle Database parameter settings. By default, this policy is not enabled.
* You can use the keyword ALL to audit all actions. The following example shows how to audit all actions on the HR.EMPLOYEES table, except actions by user pmulligan.
Example Auditing All Actions on a Table
CREATE AUDIT POLICY all_actions_on_hr_emp_pol
ACTIONS ALL ON HR.EMPLOYEES;
AUDIT POLICY all_actions_on_hr_emp_pol EXCEPT pmulligan;
References:

Question #36 *Topic 1*

On your Oracle 12c database, you invoked SQL *Loader to load data into the EMPLOYEES table in the HR schema by issuing the following command:
$> sqlldr hr/hr@pdb table=employees
Which two statements are true regarding the command? (Choose two.)

A. It succeeds with default settings if the EMPLOYEES table belonging to HR is already defined in the database.

B. It fails because no SQL *Loader data file location is specified.

C. It fails if the HR user does not have the CREATE ANY DIRECTORY privilege.

D. It fails because no SQL *Loader control file location is specified.

Correct Answer: *AC*
Note:
* SQL*Loader is invoked when you specify the sqlldr command and, optionally, parameters that establish session characteristics.

rdnkrkmz Highly Voted 5 years, 6 months ago
Unfortunately there is one correct answer which is A.
C is incorrect, I tested on 12.2
upvoted 6 times

NorwayOracle 2 years, 6 months ago
Agree, the create any directory-privilege is only relevant when using External_table=execute as a parameter.
upvoted 1 times

bkk Highly Voted 4 years, 5 months ago
But the exam is validated against 12.1
upvoted 5 times

Kitevn Most Recent 3 years, 2 months ago
A,C is correct.
When we grant CREATE ANY DIRECTORY to HR, sqlloader will create a temporary directory which is the location of executing sqlldr. If we put the controlfile, parameter file and input data present (name of these file must be the same as name of table), sqlldr will successfully load data into table.
upvoted 1 times

smalls87 3 years, 11 months ago
Nevermind, B & C are wrong, because the command won't fail if you omit the .dat & .ctl file. It will ask for it after the sqlldr command is executed
upvoted 1 times

smalls87 3 years, 11 months ago
I thought to load data using SQL Loader you need to specify a control file with mapping....
upvoted 1 times

Question #37 *Topic 1*

After implementing full Oracle Data Redaction, you change the default value for the NUMBER data type as follows:

```
SQL> SELECT NUMBER_VALUE FROM REDACTION_VALUES_FOR_TYPE_FULL;

NUMBER_VALUE
------------
           0

SQL> EXEC DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES(-1)

PL/SQL procedure successfully completed.

SQL> select number_value from redaction_values_for_type_full;

NUMBER_VALUE
------------
          -1
```

After changing the value, you notice that FULL redaction continues to redact numeric data with zero.

What must you do to activate the new default value for numeric full redaction?

A. Re-enable redaction policies that use FULL data redaction.

B. Re-create redaction policies that use FULL data redaction.

C. Re-connect the sessions that access objects with redaction policies defined on them.

D. Flush the shared pool.

E. Restart the database instance.

Correct Answer: *E*

About Altering the Default Full Data Redaction Value

You can alter the default displayed values for full Data Redaction polices. By default, 0 is the redacted value when Oracle Database performs full redaction

(DBMS_REDACT.FULL) on a column of the NUMBER data type. If you want to change it to another value (for example, 7), then you can run the DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES procedure to modify this value. The modification applies to all of the Data Redaction policies in the current database instance. After you modify a value, you must restart the database for it to take effect.

Note:

* The DBMS_REDACT package provides an interface to Oracle Data Redaction, which enables you to mask (redact) data that is returned from queries issued by low-privileged users or an application.

* UPDATE_FULL_REDACTION_VALUES Procedure

This procedure modifies the default displayed values for a Data Redaction policy for full redaction.

* After you create the Data Redaction policy, it is automatically enabled and ready to redact data.

* Oracle Data Redaction enables you to mask (redact) data that is returned from queries issued by low-privileged users or applications. You can redact column data by using one of the following methods:

/ Full redaction.

/ Partial redaction.

/ Regular expressions.

/ Random redaction.

/ No redaction.

References:

TuxBingo 4 years ago

E is the correct, check why:
About Altering the Default Full Data Redaction Value
You can alter the default displayed values for full Data Redaction polices.

By default, 0 is the redacted value when Oracle Database performs full redaction (DBMS_REDACT.FULL) on a column of the NUMBER data type. If you want to change it to another value (for example, 7), then you can run the DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES procedure to modify this value. The modification applies to all of the Data Redaction policies in the current database instance. After you modify a value, you must restart the database for it to take effect. You can find the current values by querying the REDACTION_VALUES_FOR_TYPE_FULL data dictiona view.
https://docs.oracle.com/database/121/ASOAG/configuring-oracle-data-redaction-policies.htm#ASOAG945

Regards.

upvoted 4 times

Question #38 *Topic 1*

You must track all transactions that modify certain tables in the sales schema for at least three years.

Automatic undo management is enabled for the database with a retention of one day.

Which two must you do to track the transactions? (Choose two.)

A. Enable supplemental logging for the database.

B. Specify undo retention guarantee for the database.

C. Create a Flashback Data Archive in the tablespace where the tables are stored.

D. Create a Flashback Data Archive in any suitable tablespace.

E. Enable Flashback Data Archiving for the tables that require tracking.

Correct Answer: *DE*
E: By default, flashback archiving is disabled for any table. You can enable flashback archiving for a table if you have the FLASHBACK ARCHIVE object privilege on the Flashback Data Archive that you want to use for that table.
D: Creating a Flashback Data Archive
/ Create a Flashback Data Archive with the CREATE FLASHBACK ARCHIVE statement, specifying the following:
Name of the Flashback Data Archive
Name of the first tablespace of the Flashback Data Archive
(Optional) Maximum amount of space that the Flashback Data Archive can use in the first tablespace
/ Create a Flashback Data Archive named fla2 that uses tablespace tbs2, whose data will be retained for two years:
CREATE FLASHBACK ARCHIVE fla2 TABLESPACE tbs2 RETENTION 2 YEAR;

sela 1 year, 10 months ago

D,E

16.9 Using Flashback Data Archive

A Flashback Data Archive provides the ability to track and store transactional changes to a table over its lifetime. A Flashback Data Archive is usef for compliance with record stage policies and audit reports.

A Flashback Data Archive consists of one or more tablespaces or parts thereof. You can have multiple Flashback Data Archives. If you are logged as SYSDBA, you can specify a default Flashback Data Archive for the system. A Flashback Data Archive is configured with retention time. Data archived in the Flashback Data Archive is retained for the retention time specified when the Flashback Data Archive was created.

By default, Flashback Data Archive is not enabled for any tables. You can enable Flashback Data Archive for a table if all of these conditions are tr

You have the FLASHBACK ARCHIVE object privilege on the Flashback Data Archive to use for that table.

When choosing a Flashback Data Archive for a specific table, consider the data retention requirements for the table and the retention times of th Flashback Data Archives on which you have the FLASHBACK ARCHIVE object privilege.
https://docs.oracle.com/database/121/ADFNS/adfns_flashback.htm#ADFNS01011

upvoted 1 times

Question #39 *Topic 1*

You are the DBA supporting an Oracle 11g Release 2 database and wish to move a table containing several DATE, CHAR, VARCHAR2, and NUMBER data types, and the tables indexes, to another tablespace.
The table does not have a primary key and is used by an OLTP application.
Which technique will move the table and indexes while maintaining the highest level of availability to the application?

A. Oracle Data Pump.

B. An ALTER TABLE MOVE to move the table and ALTER INDEX REBUILD to move the indexes.

C. An ALTER TABLE MOVE to move the table and ALTER INDEX REBUILD ONLINE to move the indexes.

D. Online Table Redefinition.

E. Edition-Based Table Redefinition.

Correct Answer: *D*
* Oracle Database provides a mechanism to make table structure modifications without significantly affecting the availability of the table. The mechanism is called online table redefinition. Redefining tables online provides a substantial increase in availability compared to traditional methods of redefining tables.
* To redefine a table online:
Choose the redefinition method: by key or by rowid
* By keySelect a primary key or pseudo-primary key to use for the redefinition. Pseudo-primary keys are unique keys with all component columns having NOT
NULL constraints. For this method, the versions of the tables before and after redefinition should have the same primary key columns. This is the preferred and default method of redefinition.
* By rowidUse this method if no key is available. In this method, a hidden column named M_ROW$$ is added to the post-redefined version of the table. It is recommended that this column be dropped or marked as unused after the redefinition is complete. If COMPATIBLE is set to 10.2.0 or higher, the final phase of redefinition automatically sets this column unused. You can then use the ALTER TABLE ... DROP UNUSED COLUMNS statement to drop it.
You cannot use this method on index-organized tables.
Note:
* When you rebuild an index, you use an existing index as the data source. Creating an index in this manner enables you to change storage characteristics or move to a new tablespace. Rebuilding an index based on an existing data source removes intra-block fragmentation. Compared to dropping the index and using the CREATE INDEX statement, re-creating an existing index offers better performance.
Incorrect:
Not E: Edition-based redefinition enables you to upgrade the database component of an application while it is in use, thereby minimizing or eliminating down time.

sela 1 year, 10 months ago
D. Online Table Redefinition.
upvoted 1 times

Question #40 *Topic 1*

To implement Automatic Management (AMM), you set the following parameters:

```
MEMORY_MAX_TARGET=600M
SGA_MAX_SIZE=500M
MEMORY_TARGET=600M
OPEN_CURSORS=300
SGA_TARGET=300M
PROCESSES=150
STATISTICS_LEVEL=BASIC
PGA_AGGREGATE_TARGET=0
```

When you try to start the database instance with these parameter settings, you receive the following error message:

SQL > startup -
ORA-00824: cannot set SGA_TARGET or MEMORY_TARGET due to existing internal settings, see alert log for more information.
Identify the reason the instance failed to start.

A. The PGA_AGGREGATE_TARGET parameter is set to zero.

B. The STATISTICS_LEVEL parameter is set to BASIC.

C. Both the SGA_TARGET and MEMORY_TARGET parameters are set.

D. The SGA_MAX_SIZE and SGA_TARGET parameter values are not equal.

Correct Answer: *B*
Example:

SQL> startup force -
ORA-00824: cannot set SGA_TARGET or MEMORY_TARGET due to existing internal settings
ORA-00848: STATISTICS_LEVEL cannot be set to BASIC with SGA_TARGET or MEMORY_TARGET

Cobra 2 years, 2 months ago
B IS CORRECT
upvoted 1 times

Cobra 2 years, 2 months ago
https://docs.oracle.com/en/database/oracle/oracle-database/12.2/refrn/STATISTICS_LEVEL.html#:~:text=STATISTICS_LEVEL%20specifies%20the%20level%20of%20collection%20for%20database,a20variety%20of%20purposes%2C%20including%20making%20self-management%20decisions.?msclkid=ba84e1c4ceaa11ec8f34e3541b7ab01
upvoted 1 times

Question #41 *Topic 1*

What are two benefits of installing Grid Infrastructure software for a stand-alone server before installing and creating an Oracle database?

A. Effectively implements role separation

B. Enables you to take advantage of Oracle Managed Files.

C. Automatically registers the database with Oracle Restart.

D. Helps you to easily upgrade the database from a prior release.

E. Enables the Installation of Grid Infrastructure files on block or raw devices.

Correct Answer: *AC*
C: To use Oracle ASM or Oracle Restart, you must first install Oracle Grid Infrastructure for a standalone server before you install and create the database. Otherwise, you must manually register the database with Oracle Restart.
Desupport of Block and Raw Devices
With the release of Oracle Database 11g release 2 (11.2) and Oracle RAC 11g release 2 (11.2), using Database Configuration Assistant or the installer to store
Oracle Clusterware or Oracle Database files directly on block or raw devices is not supported.
If you intend to upgrade an existing Oracle RAC database, or an Oracle RAC database with Oracle ASM instances, then you can use an existing raw or block device partition, and perform a rolling upgrade of your existing installation. Performing a new installation using block or raw devices is not allowed.
References:

vineetarneja Highly Voted 4 years, 3 months ago
No. I think AC are only correct. i checked for E option and I see that the RAW devices are no longer supported to store the Clusterware files. PFB t reference URL.

https://docs.oracle.com/database/121/CWLIN/storage.htm#CDEBGJAG

Regards
Vineet Arneja
vineetarneja1989@gmail.com
upvoted 6 times

jokerking Most Recent 2 years, 9 months ago
I think AC are only correct
upvoted 1 times

hggz 4 years, 5 months ago
I think C and E are correct.

ASM can work with raw and block devices

https://docs.oracle.com/cd/B28359_01/server.111/b31107/asmprepare.htm#OSTMG02150
upvoted 1 times

Bernardes 3 years, 4 months ago
AC
At 12.1 the support for raw and block devices was ended.
"Desupport of Block and Raw Devices
With the release of Oracle Database 11g release 2 (11.2) and Oracle RAC 11g release 2 (11.2), using Database Configuration Assistant or the installer to store Oracle Clusterware or Oracle Database files directly on block or raw devices is not supported."
upvoted 1 times

Question #42 *Topic 1*

Identify two correct statements about multitenant architectures.

A. Multitenant architecture can be deployed only in a Real Application Clusters (RAC) configuration.

B. Multiple pluggable databases (PDBs) share certain multitenant container database (CDB) resources.

C. Multiple CDBs share certain PDB resources.

D. Multiple non-RAC CDB instances can mount the same PDB as long as they are on the same server.

E. Patches are always applied at the CDB level.

F. A PDB can have a private undo tablespace.

Correct Answer: *BE*

B: Using 12c Resource manager you will be able control CPU, Exadata I/O, sessions and parallel servers. A new 12c CDB Resource Manager Plan will use so- called "Shares" (resource allocations) to specify how CPU is distributed between PDBs. A CDB Resource Manager Plan also can use "utilization limits" to limit the
CPU usage for a PDB. With a default directive, you do not need to modify the resource plan for each PDB plug and unplug.

E: New paradigms for rapid patching and upgrades.
The investment of time and effort to patch one multitenant container database results in patching all of its many pluggable databases. To patch a single pluggable database, you simply unplug/plug to a multitenant container database at a different Oracle Database software version.

Incorrect:

Not A:

* The Oracle RAC documentation describes special considerations for a CDB in an Oracle RAC environment.
* Oracle Multitenant is a new option for Oracle Database 12c Enterprise Edition that helps customers reduce IT costs by simplifying consolidation, provisioning, upgrades, and more.

It is supported by a new architecture that allows a container database to hold many pluggable databases. And it fully complements other options, including Oracle
Real Application Clusters and Oracle Active Data Guard. An existing database can be simply adopted, with no change, as a pluggable database; and no changes are needed in the other tiers of the application.

Not D: You can unplug a PDB from one CDB and plug it into a different CDB without altering your schemas or applications. A PDB can be plugged into only one
CDB at a time.

not F:

* UNDO tablespace can NOT be local and stays on the CDB level.
* Redo and undo go hand in hand, and so the CDB as a whole has a single undo tablespace per RAC instance.

sela 1 year, 10 months ago

B,E
Student Guide page :2-39 and 2-40

upvoted 1 times

rdnkrkmz 5 years, 6 months ago

But this Exam is for 12.1, so F is not correct.

upvoted 2 times

hggz 4 years, 5 months ago

You right, local undo is feature of release 12.2

upvoted 2 times

rdnkrkmz 5 years, 6 months ago

As of release 12.2

```
COLUMN property_name FORMAT A30
COLUMN property_value FORMAT A30

SELECT property_name, property_value
FROM database_properties
WHERE property_name = 'LOCAL_UNDO_ENABLED';

PROPERTY_NAME PROPERTY_VALUE
------------------------------ ------------------------------
LOCAL_UNDO_ENABLED TRUE
```

then, PDB can have local undo tablespace
upvoted 1 times

Question #43 *Topic 1*

You upgrade your Oracle database in a multiprocessor environment. As a recommended you execute the following script:

SQL > @utlrp.sql -
Which two actions does the script perform? (Choose two.)

A. Parallel compilation of only the stored PL/SQL code

B. Sequential recompilation of only the stored PL/SQL code

C. Parallel recompilation of any stored PL/SQL code

D. Sequential recompilation of any stored PL/SQL code

E. Parallel recompilation of Java code

F. Sequential recompilation of Java code

Correct Answer: *CE*
utlrp.sql and utlprp.sql
The utlrp.sql and utlprp.sql scripts are provided by Oracle to recompile all invalid objects in the database. They are typically run after major database changes such as upgrades or patches. They are located in the $ORACLE_HOME/rdbms/admin directory and provide a wrapper on the UTL_RECOMP package. The utlrp.sql script simply calls the utlprp.sql script with a command line parameter of "0". The utlprp.sql accepts a single integer parameter that indicates the level of parallelism as follows.
0 - The level of parallelism is derived based on the CPU_COUNT parameter.
1 - The recompilation is run serially, one object at a time.
N - The recompilation is run in parallel with "N" number of threads.
Both scripts must be run as the SYS user, or another user with SYSDBA, to work correctly.
References:

Question #44 *Topic 1*

Which two statements are true concerning dropping a pluggable database (PDB)? (Choose two.)

A. The PDB must be open in read-only mode.

B. The PDB must be in mount state.

C. The PDB must be unplugged.

D. The PDB data files are always removed from disk.

E. A dropped PDB can never be plugged back into a multitenant container database (CDB).

Correct Answer: *BC*
http://docs.oracle.com/database/121/ADMIN/cdb_plug.htm#ADMIN13658

jackymak 2 years, 5 months ago
I don't understand how to mount the PDB but unplugged it.
The other options seem to be not right though.
upvoted 1 times

Ray520 2 years ago
It's either and or: Either it should be unplugged or in MOUNT state.
upvoted 1 times

Question #45 *Topic 1*

You notice a high number of waits for the db file scattered read and db file sequential read events in the recent Automatic Database Diagnostic Monitor (ADDM) report. After further investigation, you find that queries are performing too many full table scans and indexes are not being used even though the filter columns are indexed.
Identify three possible reasons for this.

A. Missing or stale histogram statistics

B. Undersized shared pool

C. High clustering factor for the indexes

D. High value for the DB_FILE_MULTIBLOCK_READ_COUNT parameter

E. Oversized buffer cache

Correct Answer: *ACD*
D: DB_FILE_MULTIBLOCK_READ_COUNT is one of the parameters you can use to minimize I/O during table scans. It specifies the maximum number of blocks read in one I/O operation during a sequential scan. The total number of I/Os needed to perform a full table scan depends on such factors as the size of the table, the multiblock read count, and whether parallel execution is being utilized for the operation.

Question #46 *Topic 1*

Which three features work together, to allow a SQL statement to have different cursors for the same statement based on different selectivity ranges? (Choose three.)

A. Bind Variable Peeking

B. SQL Plan Baselines

C. Adaptive Cursor Sharing

D. Bind variable used in a SQL statement

E. Literals in a SQL statement

Correct Answer: *ACD*

verratti Highly Voted 4 years, 5 months ago
A C AND E
upvoted 5 times

luizyto Most Recent 1 year, 10 months ago
I think should be A,C,E
upvoted 1 times

TuxBingo 3 years, 3 months ago
Read the question, is asking what cause a SQL statement to have more than one.
upvoted 1 times

hggz 4 years, 4 months ago
B,C,D correct:

B. SQL Plan Baselines: https://docs.oracle.com/database/121/TGSQL/tgsql_spm.htm#TGSQL615

"For example, a SQL statement with bind variables whose values result in different selectivities may have several optimal plans. Creating a materialized view or an index or repartitioning a table may make current plans more expensive than other plans."

C. Adaptive Cursor Sharing: https://docs.oracle.com/database/121/TGSQL/tgsql_cursor.htm#TGSQL-GUID-42432DFE-D56A-4B9A-8FCF-94FB86012AFC

"The adaptive cursor sharing feature enables a single statement that contains bind variables to use multiple execution plans. Cursor sharing is "adaptive" because the cursor adapts its behavior so that the database does not always use the same plan for each execution or bind variable value."

D. Bind variable used in a SQL statement
upvoted 2 times

hggz 3 years, 11 months ago
After some time i think ACD
upvoted 2 times

5kong 4 years, 5 months ago
http://davetechnotes.blogspot.com/2009/02/literal-vs-bind-variables-bind-variable.html
upvoted 1 times

verratti 4 years, 5 months ago
appears to be correct.
upvoted 1 times

soniat 3 years, 4 months ago
ACE is the correct answer. d is wrong When using bind variables there will be only one cursor and one execution plan in the Shared Pool regardless of the value being send to the bind variable. If the bind variable is selecting from a column which is highly skewed then the executi plan may not be optimal.
upvoted 1 times

Question #47 *Topic 1*

You notice a performance change in your production Oracle 12c database. You want to know which change caused this performance difference. Which method or feature should you use?

A. Compare Period ADDM report

B. AWR Compare Period report

C. Active Session History (ASH) report

D. Taking a new snapshot and comparing it with a preserved snapshot

Correct Answer: *A*

TuxBingo Highly Voted 4 years ago

The correct answer is A, check why:
Comparing Current System Performance to a Baseline Period
You may have noticed a performance change on a production system and would like to know why, or you may have implemented a change to a production system and want to know the effect of the change, such as increased concurrency waits.

The Compare Period ADDM compares the performance of the database server in two time periods, and returns a report describing the performance changes and the root origin of the changes.
https://docs.oracle.com/database/121/TDPPT/tdppt_degrade.htm#TDPPT525

upvoted 6 times

vineetarneja Most Recent 4 years, 3 months ago

But ADDM only shows the recommendations. it does not show what is the source of the problem. it only focus on the solution part. here we want to know the root cause of the problem. so i also think that the answer should be B. if some one has a better justification about answer A, please explain.

Regards
Vineet Arneja
vineetarneja1989@gmail.com

upvoted 1 times

euler 4 years, 11 months ago

answer is B, not A

upvoted 3 times

5kong 4 years, 8 months ago

not B. AWR is just collector. but ADDM is analysis. it can check what is different.

upvoted 5 times

Havok 4 years, 6 months ago

correct. AWR takes the snapshots. ADDM uses the two snapshots and analyzes.

upvoted 2 times

Question #48 *Topic 1*

You want to capture column group usage and gather extended statistics for better cardinality estimates for the CUSTOMERS table in the SH schema.
Examine the following steps:
1. Issue the SELECT DBMS_STATS.CREATE_EXTENDED_STATS (SH, CUSTOMERS) FROM dual statement.
2. Execute the DBMS_STATS.SEED_COL_USAGE (null, SH, 500) procedure.
3. Execute the required queries on the CUSTOMERS table.
4. Issue the SELECT DBMS_STATS.REPORT_COL_USAGE (SH, CUSTOMERS) FROM dual statement.
Identify the correct sequence of steps.

A. 3, 2, 1, 4

B. 2, 3, 4, 1

C. 4, 1, 3, 2

D. 3, 2, 4, 1

Correct Answer: *B*
Step 1 (2). Seed column usage -
Oracle must observe a representative workload, in order to determine the appropriate column groups. Using the new procedure DBMS_STATS.SEED_COL_USAGE, you tell Oracle how long it should observe the workload.
Step 2: (3) You don't need to execute all of the queries in your work during this window. You can simply run explain plan for some of your longer running queries to ensure column group information is recorded for these queries.
Step 3. (1) Create the column groups
At this point you can get Oracle to automatically create the column groups for each of the tables based on the usage information captured during the monitoring window. You simply have to call the DBMS_STATS.CREATE_EXTENDED_STATS function for each table.This function requires just two arguments, the schema name and the table name. From then on, statistics will be maintained for each column group whenever statistics are gathered on the table.
Note:
* DBMS_STATS.REPORT_COL_USAGE reports column usage information and records all the SQL operations the database has processed for a given object.
* The Oracle SQL optimizer has always been ignorant of the implied relationships between data columns within the same table. While the optimizer has traditionally analyzed the distribution of values within a column, he does not collect value-based relationships between columns.
* Creating extended statisticsHere are the steps to create extended statistics for related table columns withdbms_stats.created_extended_stats:
1 - The first step is to create column histograms for the related columns.2 Next, we run dbms_stats.create_extended_stats to relate the columns together.
Unlike a traditional procedure that is invoked via an execute ("exec") statement, Oracle extended statistics are created via a select statement.

PashaAllahverdi 2 years, 3 months ago
https://oracle-base.com/articles/11g/extended-statistics-enhancements-11gr2
upvoted 1 times

Question #49 *Topic 1*

Which three statements are true about Automatic Workload Repository (AWR)? (Choose three.)

A. All AWR tables belong to the SYSTEM schema.

B. The AWR data is stored in memory and in the database.

C. The snapshots collected by AWR are used by the self-tuning components in the database

D. AWR computes time model statistics based on time usage for activities, which are displayed in the v$SYS time model and V$SESS_TIME_MODEL views.

E. AWR contains system wide tracing and logging information.

Correct Answer: *BCD*

khalilshahin01 Highly Voted 3 years, 11 months ago

Tricky one! ADR contains system wide tracing and logging information (NOT AWR!).
Therefore BCD

upvoted 8 times

Angeltje1992 Most Recent 4 years, 3 months ago

I Think D is right:
"Time model statistics based on time usage for activities, displayed in the V$SYS_TIME_MODEL and V$SESS_TIME_MODEL views"

upvoted 1 times

5kong 4 years, 2 months ago

#274 E. An AWR snapshot contains system-wide tracing and logging information.
is false

upvoted 1 times

euler 4 years, 11 months ago

BCE? D is confuse

upvoted 1 times

hggz 4 years, 4 months ago

Yes correct:
https://oracle-base.com/articles/10g/automatic-workload-repository-10g
"Time model statistics indicating the amount of DB time associated with a process from the V$SESS_TIME_MODEL and V$SYS_TIME_MODEL views."

upvoted 4 times

hggz 4 years, 4 months ago

Link for 12.1: https://docs.oracle.com/database/121/TGDBA/gather_stats.htm#TGDBA169

upvoted 2 times

yosiw96816 3 years, 8 months ago

In your provided link:
Time model statistics based on time usage for activities, displayed in the V$SYS_TIME_MODEL and V$SESS_TIME_MODEL views which makes BCD correct

upvoted 2 times

Question #50 *Topic 1*

You upgraded your database from pre-12c to a multitenant container database (CDB) containing pluggable databases (PDBs).
Examine the query and its output:

```
SQL> SELECT * FROM v$PWFILE_users;

USERNAME                       SYSDB SYSOP SYSAS SYSBA SYSDG SYSKM       CON_ID
------------------------------ ----- ----- ----- ----- ----- ----- ------------
SYS                            TRUE  TRUE  FALSE FALSE FALSE FALSE            0
```

Which two tasks must you perform to add users with SYSBACKUP, SYSDG, and SYSKM privilege to the password file? (Choose two.)

A. Assign the appropriate operating system groups to SYSBACKUP, SYSDG, SYSKM.

B. Grant SYSBACKUP, SYSDG, and SYSKM privileges to the intended users.

C. Re-create the password file with SYSBACKUP, SYSDG, and SYSKM privilege and the FORCE argument set to No.

D. Re-create the password file with SYSBACKUP, SYSDG, and SYSKM privilege, and FORCE arguments set to Yes.

E. Re-create the password file in the Oracle Database 12c format.

Correct Answer: *BD*
* orapwd
/ You can create a database password file using the password file creation utility, ORAPWD.
The syntax of the ORAPWD command is as follows:
orapwd FILE=filename [ENTRIES=numusers] [FORCE={y|n}] [ASM={y|n}]
[DBUNIQUENAME=dbname] [FORMAT={12|legacy}] [SYSBACKUP={y|n}] [SYSDG={y|n}]
[SYSKM={y|n}] [DELETE={y|n}] [INPUT_FILE=input-fname]
force - whether to overwrite existing file (optional),
* v$PWFILE_users
/ 12c: V$PWFILE_USERS lists all users in the password file, and indicates whether the user has been granted the SYSDBA, SYSOPER, SYSASM, SYSBACKUP,
SYSDG, and SYSKM privileges.
/ 10c: sts users who have been granted SYSDBA and SYSOPER privileges as derived from the password file.

ColumnDatatypeDescription -
USERNAMEVARCHAR2(30)The name of the user that is contained in the password file
SYSDBAVARCHAR2(5)If TRUE, the user can connect with SYSDBA privileges
SYSOPERVARCHAR2(5)If TRUE, the user can connect with SYSOPER privileges
Incorrect:
not E: The format of the v$PWFILE_users file is already in 12c format.

Tekdeel 3 years, 4 months ago
BASED and TRUE
upvoted 1 times

KiwiFace 3 years, 4 months ago
B is very wrong I think
upvoted 1 times

Question #51 *Topic 1*

An application accesses a small lookup table frequently. You notice that the required data blocks are getting aged out of the default buffer cache. How would you guarantee that the blocks for the table never age out?

A. Configure the KEEP buffer pool and alter the table with the corresponding storage clause.

B. Increase the database buffer cache size.

C. Configure the RECYCLE buffer pool and alter the table with the corresponding storage clause.

D. Configure Automata Shared Memory Management.

E. Configure Automatic Memory Management.

Correct Answer: *A*
Schema objects are referenced with varying usage patterns; therefore, their cache behavior may be quite different. Multiple buffer pools enable you to address these differences. You can use a KEEP buffer pool to maintain objects in the buffer cache and a RECYCLE buffer pool to prevent objects from consuming unnecessary space in the cache. When an object is allocated to a cache, all blocks from that object are placed in that cache. Oracle maintains a DEFAULT buffer pool for objects that have not been assigned to one of the buffer pools.

sela 1 year, 10 months ago

A
https://docs.oracle.com/database/121/TGDBA/tune_buffer_cache.htm#TGDBA551
Configuring the KEEP Pool
The purpose of the KEEP buffer pool is to retain objects in memory, thus avoiding I/O operations. Each object kept in memory results in a trade-c
It is more beneficial to keep frequently-accessed blocks in the cache. Avoid retaining infrequently-used blocks in the cache, as this results in less space for other, more active blocks

upvoted 1 times

Question #52 *Topic 1*

You conned using SQL Plus to the root container of a multitenant container database (CDB) with SYSDBA privilege.
The CDB has several pluggable databases (PDBs) open in the read/write mode.
There are ongoing transactions in both the CDB and PDBs.
What happens alter issuing the SHUTDOWN TRANSACTIONAL statement?

A. The shutdown proceeds immediately. The shutdown proceeds as soon as all transactions in the PDBs are either committed or rolled hack.

B. The shutdown proceeds as soon as all transactions in the CDB are either committed or rolled back.

C. The shutdown proceeds as soon as all transactions in both the CDB and PDBs are either committed or rolled back.

D. The statement results in an error because there are open PDBs.

Correct Answer: *B*
* SHUTDOWN [ABORT | IMMEDIATE | NORMAL | TRANSACTIONAL [LOCAL]]
Shuts down a currently running Oracle Database instance, optionally closing and dismounting a database. If the current database is a pluggable database, only the pluggable database is closed. The consolidated instance continues to run.
Shutdown commands that wait for current calls to complete or users to disconnect such as SHUTDOWN NORMAL and SHUTDOWN TRANSACTIONAL have a time limit that the SHUTDOWN command will wait. If all events blocking the shutdown have not occurred within the time limit, the shutdown command cancels with the following message:
ORA-01013: user requested cancel of current operation
* If logged into a CDB, shutdown closes the CDB instance.
To shutdown a CDB or non CDB, you must be connected to the CDB or non CDB instance that you want to close, and then enter

SHUTDOWN -
Database closed.
Database dismounted.
Oracle instance shut down.
To shutdown a PDB, you must log into the PDB to issue the SHUTDOWN command.

SHUTDOWN -
Pluggable Database closed.
Note:
* Prerequisites for PDB Shutdown
When the current container is a pluggable database (PDB), the SHUTDOWN command can only be used if:
The current user has SYSDBA, SYSOPER, SYSBACKUP, or SYSDG system privilege.
The privilege is either commonly granted or locally granted in the PDB.
The current user exercises the privilege using AS SYSDBA, AS SYSOPER, AS SYSBACKUP, or AS SYSDG at connect time.
To close a PDB, the PDB must be open.

jackymak 2 years, 5 months ago
I think it is C.
Otherwise, what is going to happen in PDBs?
upvoted 1 times

jackymak 2 years, 5 months ago
tested: PDB will shutdown immediate
upvoted 2 times

NorwayOracle 2 years, 4 months ago
Did the PDB's shutdown by cancel ongoing transactions or did it shut down using immediate or transaction option?
upvoted 2 times

Ray_gk 2 years, 2 months ago
I have tested this too. For me, the PDB shutdown immediately upon executing the shutdown transactional command from the root container despite the ongoing transaction in the PDB. I did not issue a commit or rollback in the PDB, nor did I use the shutdown immediate command in the root container.
upvoted 1 times

Ray_gk 2 years, 2 months ago

Based on my test I am going with A. I forgot to mention both the root container and PDB shutdown as if I used the shutdown immediate command. Actually, per Oracle documentation, shutdown immediate or abort are the only valid clauses both the root container and PDBs. Seems as though shutdown immediate is the default if you choose another clause aside from abort.

upvoted 1 times

Question #53 *Topic 1*

You are planning the creation of a new multitenant container database (CDB) and want to store the ROOT and SEED container data files in separate directories.
You plan to create the database using SQL statements.
Which three techniques can you use to achieve this? (Choose three.)

A. Use Oracle Managed Files (OMF).

B. Specify the SEED FILE_NAME_CONVERT clause.

C. Specify the PDB_FILE_NAME_CONVERT initialization parameter.

D. Specify the DB_FILE_NAMECONVERT initialization parameter.

E. Specify all files in the CREATE DATABASE statement without using Oracle managed Files (OMF).

Correct Answer: *ABC*

You must specify the names and locations of the seed's files in one of the following ways:
* (A) Oracle Managed Files
* (B) The SEED FILE_NAME_CONVERT Clause
* (C) The PDB_FILE_NAME_CONVERT Initialization Parameter

Question #54 *Topic 1*

You are about to plug a multi-terabyte non-CDB into an existing multitenant container database (CDB).
The characteristics of the non-CDB are as follows:
Version: Oracle Database 11g Release 2 (11.2.0.2.0) 64-bit
Character set: AL32UTF8
National character set: AL16UTF16
O/S: Oracle Linux 6 64-bit
The characteristics of the CDB are as follows:
Version: Oracle Database 12c Release 1 64-bit
Character Set: AL32UTF8
National character set: AL16UTF16
O/S: Oracle Linux 6 64-bit
Which technique should you use to minimize down time while plugging this non-CDB into the CDB?

A. Transportable database

B. Transportable tablespace

C. Data Pump full export/import

D. The DBMS_PDB package

E. RMAN

Correct Answer: *B*
* Overview, example:
- Log into ncdb12c as sys
- Get the database in a consistent state by shutting it down cleanly.
- Open the database in read only mode
- Run DBMS_PDB.DESCRIBE to create an XML file describing the database.
- Shut down ncdb12c
- Connect to target CDB (CDB2)
- Check whether non-cdb (NCDB12c) can be plugged into CDB(CDB2)
- Plug-in Non-CDB (NCDB12c) as PDB(NCDB12c) into target CDB(CDB2).
- Access the PDB and run the noncdb_to_pdb.sql script.
- Open the new PDB in read/write mode.
* You can easily plug an Oracle Database 12c non-CDB into a CDB. Just create a PDB manifest file for the non-CDB, and then use the manifest file to create a cloned PDB in the CDB.
* Note that to plug in a non-CDB database into a CDB, the non-CDB database needs to be of version 12c as well. So existing 11g databases will need to be upgraded to 12c before they can be part of a 12c CDB.

Community vote distribution
D (100%)

TuxBingo Highly Voted 4 years ago
I agree with the answer is B.
Why--
http://www.oracle.com/technetwork/database/upgrade/upgrading-oracle-database-wp-12c-1896123.pdf
TABLE 2 . DATABASE UPGRADE AND MIGRATION METHODS – PAGE 7
upvoted 7 times

soniat 3 years, 4 months ago
cis the correct answer
upvoted 1 times

soniat 3 years, 4 months ago
i meant d
upvoted 2 times

Ray520 Most Recent 2 years, 4 months ago
The key here is "minimize downtime". Transportable tablespaces (starts from db version 11.2.0.3) is the fastest way to upgrade.
upvoted 1 times

Kotek___3 2 years, 5 months ago

Selected Answer: D

I think that the answer is D as there is no need for transport and dbms_pdb describe allows to use datafiles of a non-CDB db, therefore less downtime

upvoted 1 times

Ray520 2 years ago

version has to be 12c for both to plug non-CDB into an CDB using DBMS_PDB package. Here versions are different

upvoted 1 times

Question #55 *Topic 1*

Your database supports an online transaction processing (OLTP) application. The application is undergoing some major schema changes, such as addition of new indexes and materialized views. You want to check the impact of these changes on workload performance.
What should you use to achieve this?

A. Database replay

B. SQL Tuning Advisor

C. SQL Access Advisor

D. SQL Performance Analyzer

E. Automatic Workload Repository compare reports

Correct Answer: *D*
You can use the SQL Performance Analyzer to analyze the SQL performance impact of any type of system change. Examples of common system changes include:
Database upgrades
Configuration changes to the operating system, hardware, or database
Database initialization parameter changes
Schema changes, such as adding new indexes or materialized views
Gathering optimizer statistics
SQL tuning actions, such as creating SQL profiles
http://docs.oracle.com/cd/B28359_01/server.111/b28318/intro.htm#CNCPT961

Question #56 *Topic 1*

An administrator account is granted the CREATE SESSION and SET CONTAINER system privileges.
A multitenant container database (CDB) instant has the following parameter set:

THREADED_EXECUTION = FALSE -
Which four statements are true about this administrator establishing connections to root in a CDB that has been opened in read only mode?
(Choose four.)

A. You can conned as a common user by using the connect statement.

B. You can connect as a local user by using the connect statement.

C. You can connect by using easy connect.

D. You can connect by using OS authentication.

E. You can connect by using a Net Service name.

F. You can connect as a local user by using the SET CONTAINER statement.

Correct Answer: *ACDE*

TuxBingo Highly Voted 4 years ago

I share this explanation:
The given answer of C,D,E,F is wrong. The question tells you an administrator user can login (CREATE SESSION) and can issue SET CONTAINER commands. It also tells you it is a CDB and that the multi-threaded architecture is disabled. That means OS authentication will work. So D is corre
As long as TCP/IP is configured, then using EZ Connect would work, so C is correct. Assuming there's a listener and a tnsnames.ora file correctly configured would make E correct. Since the connection is to the ROOT container, that immediately eliminates any answer referencing LOCAL user because ROOT cannot have any LOCAL users. SELECT DISTINCT(COMMON) from DBA_USERS returns YES in CDB$ROOT. Further evidence of COMMON only users in root is provided if you attempt to create a local user in CDB$ROOT (i.e. a user without the C## prefix). Oracle returns an error if you attempt it. So, that eliminates B and F, leaving B as the fourth choice.

Answer: ACDE

Link for reference: https://www.thegeekdiary.com/oracle-12c-new-feature-multi-threaded-architecture-of-processes/

upvoted 7 times

Question #57 *Topic 1*

Examine the following query output:

```
SQL> SELECT name, force_logging FROM v$database;

NAME      FORCE_LOGGING
--------- ---------------------------------------
PROD      NO
```

You issue the following command to import tables into the hr schema:
$ > impdp hr/hr directory = dumpdir dumpfile = hr_new.dmp schemas=hr TRANSFORM=DISABLE_ARCHIVE_LOGGING: Y
Which statement is true?

A. All database operations performed by the impdp command are logged.

B. Only CREATE INDEX and CREATE TABLE statements generated by the import are logged.

C. Only CREATE TABLE and ALTER TABLE statements generated by the import are logged.

D. None of the operations against the master table used by Oracle Data Pump to coordinate its activities are logged.

Correct Answer: *C*
Oracle Data Pump disable redo logging when loading data into tables and when creating indexes.
The new TRANSFORM option introduced in data pumps import provides the flexibility to turn off the redo generation for the objects during the course of import.
The Master Table is used to track the detailed progress information of a Data Pump job.
The Master Table is created in the schema of the current user running the Pump Dump export or import, and it keeps tracks of lots of detailed information.

vineetarneja Highly Voted 4 years, 3 months ago
Justification
Even with this parameter specified, there is still redo logging for other operations of Oracle Data Pump. This includes all CREATE and ALTER statements, except CREATE INDEX, and all operations against the master table used by Oracle Data Pump during the import."

http://docs.oracle.com/database/121/NEWFT/chapter12101.htm#NEWFT253
-
Vineet Arneja
vineetarneja1989@gmail.com
upvoted 9 times

Hacova Most Recent 2 years, 11 months ago
Answer C
Even with this parameter specified, there is still redo logging for other operations of Oracle Data Pump. This includes all CREATE and ALTER statements, except CREATE INDEX, and all operations against the master table used by Oracle Data Pump during the import.
upvoted 3 times

Abhirup 3 years, 11 months ago
correct option ; D
upvoted 1 times

dancymonkey 3 years, 11 months ago
D
DISABLE_ARCHIVE_LOGGING:[Y | N]

If set to Y, then the logging attributes for the specified object types (TABLE and/or INDEX) are disabled before the data is imported. If set to N (th default), then archive logging is not disabled during import. After the data has been loaded, the logging attributes for the objects are restored to their original settings. If no object type is specified, then the DISABLE_ARCHIVE_LOGGING behavior is applied to both TABLE and INDEX object types. This transform works for both file mode imports and network mode imports. It does not apply to transportable tablespace imports.

Note:

If the database is in FORCE LOGGING mode, then the DISABLE_ARCHIVE_LOGGING option will not disable logging when indexes and tables are created.

https://docs.oracle.com/database/121/SUTIL/GUID-64FB67BD-EB67-4F50-A4D2-5D34518E6BDB.htm#SUTIL939
upvoted 1 times

Question #58 *Topic 1*

You notice a performance change in your production Oracle database and you want to know which change has made this performance difference.
You generate the Compare Period Automatic Database Diagnostic Monitor (ADDM) report to further investigation.
Which three findings would you get from the report? (Choose three.)

A. It detects any configuration change that caused a performance difference in both time periods.

B. It identifies any workload change that caused a performance difference in both time periods.

C. It detects the top wait events causing performance degradation.

D. It shows the resource usage for CPU, memory, and I/O in both time periods.

E. It shows the difference in the size of memory pools in both time periods.

F. It gives information about statistics collection in both time periods.

Correct Answer: *ABD*
Keyword: shows the difference.
* Full ADDM analysis across two AWR snapshot periods
Detects causes, measure effects, then correlates them
Causes: workload changes, configuration changes
Effects: regressed SQL, reach resource limits (CPU, I/O, memory, interconnect)
Makes actionable recommendations along with quantified impact
* Identify what changed
/ Configuration changes, workload changes
* Performance degradation of the database occurs when your database was performing optimally in the past, such as 6 months ago, but has gradually degraded to a point where it becomes noticeable to the users. The Automatic Workload Repository (AWR) Compare Periods report enables you to compare database performance between two periods of time.
While an AWR report shows AWR data between two snapshots (or two points in time), the AWR Compare Periods report shows the difference (ABE) between two periods (or two AWR reports with a total of four snapshots). Using the AWR Compare Periods report helps you to identify detailed performance attributes and configuration settings that differ between two time periods.
References:

sela 1 year, 10 months ago
C,D,E
https://oracle-base.com/articles/10g/automatic-database-diagnostic-monitor-10g
upvoted 1 times

Question #59 *Topic 1*

Examine the parameter for your database instance:

```
NAME                                 TYPE        VALUE
------------------------------------ ----------- ---------
optimizer_adaptive_reporting_only    boolean     FALSE
optimizer_capture_sql_plan_baselines boolean     FALSE
optimizer_dynamic_sampling           integer     2
optimizer_features_enable            string      12.1.0.1
```

You generated the execution plan for the following query in the plan table and noticed that the nested loop join was done. After actual execution of the query, you notice that the hash join was done in the execution plan:

```
SQL> SELECT product_name
FROM   order_items o, product_information p
WHERE  o.unit_price = 15
AND    quantity > 1
AND    p.product_id = o.product_id;

30 rows selected.
```

Identify the reason why the optimizer chose different execution plans.

A. The optimizer used a dynamic plan for the query.

B. The optimizer chose different plans because automatic dynamic sampling was enabled.

C. The optimizer used re-optimization cardinality feedback for the query.

D. The optimizer chose different plan because extended statistics were created for the columns used.

Correct Answer: *A*

Ray520 2 years ago

it's A. Optimizer_dynamiz_sampling value is 2 (default) which means it uses dynamic statistics if at least one table has no statistics.

upvoted 1 times

Oracle2020 2 years, 5 months ago

I Think B
Explanation/Reference:
Explanation:
* optimizer_dynamic_sampling
OPTIMIZER_DYNAMIC_SAMPLING controls both when the database gathers dynamic statistics, and the size of the sample that the optimizer uses to gather the statistics.
Range of values0 to 11

upvoted 1 times

TuxBingo 4 years ago

The correct answer is A, Why?
Oracle Documentation: https://docs.oracle.com/database/121/TGSQL/tgsql_astat.htm#TGSQL451 -- Table 13-1 Dynamic Statistics Levels -- 2

Use dynamic statistics if at least one table in the statement has no statistics. This is the default value, so for that reason: The optimizer used a dynamic plan for the query.

upvoted 4 times

jackymak 2 years, 3 months ago

Thanks, It should be A.
If optimizer_dynamic_sampling values = 11
then C will be right.

upvoted 1 times

jackymak 2 years, 3 months ago

***I mean "then B will be right"

upvoted 1 times

Mamlouk_Med 5 years, 5 months ago

I Think B

upvoted 2 times

5kong 4 years, 8 months ago

why? automatic dynamic sampling was enabled ??
the answer is A

upvoted 6 times

Question #60 *Topic 1*

Which three statements are true about adaptive SQL plan management? (Choose three.)

A. It automatically performs verification or evolves non-accepted plans, in COMPREHENSIVE mode when they perform better than existing accepted plans.

B. The optimizer always uses the fixed plan, if the fixed plan exists in the plan baseline.

C. It adds new, bettor plans automatically as fixed plans to the baseline.

D. The non-accepted plans are automatically accepted and become usable by the optimizer if they perform better than the existing accepted plans.

E. The non-accepted plans in a SQL plan baseline are automatically evolved, in COMPREHENSIVE mode, during the nightly maintenance window and a

Correct Answer: *ADE*

With adaptive SQL plan management, DBAs no longer have to manually run the verification or evolve process for non-accepted plans. When automatic SQL tuning is in COMPREHENSIVE mode, it runs a verification or evolve process for all SQL statements that have non-accepted plans during the nightly maintenance window. If the non-accepted plan performs better than the existing accepted plan (or plans) in the SQL plan baseline, then the plan is automatically accepted and becomes usable by the optimizer. After the verification is complete, a persistent report is generated detailing how the non-accepted plan performs compared to the accepted plan performance. Because the evolve process is now an AUTOTASK, DBAs can also schedule their own evolve job at end time.

Note:

* The optimizer is able to adapt plans on the fly by predetermining multiple subplans for portions of the plan.
* Adaptive plans, introduced in Oracle Database 12c, enable the optimizer to defer the final plan decision for a statement until execution time. The optimizer instruments its chosen plan (the default plan) with statistics collectors so that it can detect at runtime, if its cardinality estimates differ greatly from the actual number of rows seen by the operations in the plan. If there is a significant difference, then the plan or a portion of it will be automatically adapted to avoid suboptimal performance on the first execution of a SQL statement.

References:

sela 1 year, 10 months ago

a,d,e
2.2.4.2 Adaptive SQL Plan Management

With adaptive SQL plan management, DBAs no longer have to manually run the verification or evolve process for non-accepted plans. When automatic SQL tuning is in COMPREHENSIVE mode, it runs a verification or evolve process for all SQL statements that have non-accepted plans during the nightly maintenance window. If the non-accepted plan performs better than the existing accepted plan (or plans) in the SQL plan baseline, then the plan is automatically accepted and becomes usable by the optimizer. After the verification is complete, a persistent report is generated detailing how the non-accepted plan performs compared to the accepted plan performance. Because the evolve process is now an AUTOTASK, DBAs can also schedule their own evolve job at end time.

Unaccepted plans in a SQL plan baseline are automatically evolved during the nightly maintenance window and a persistent verification report is generated which means a DBA no longer has to manual evolve plans and they can go back days or weeks later and review what plans were evolv during each of the nightly maintenance windows.

https://docs.oracle.com/database/121/NEWFT/chapter12101.htm#NEWFT205

upvoted 1 times

TuxBingo 4 years ago

Check the documentation: https://docs.oracle.com/database/121/NEWFT/chapter12101.htm#NEWFT205 Section: 2.2.4.2 Adaptive SQL Plan Management The key of this is: COMPREHENSIVE mode

upvoted 4 times

Question #61 *Topic 1*

You create a new pluggable database, HR_PDB, from the seed database. Which three tablespaces are created by default in HR_PDB? (Choose three.)

A. SYSTEM

B. SYSAUX

C. EXAMPLE

D. UNDO

E. TEMP

F. USERS

Correct Answer: *ABE*

* A PDB would have its SYSTEM, SYSAUX, TEMP tablespaces.It can also contains other user created tablespaces in it.

* Oracle Database creates both the SYSTEM and SYSAUX tablespaces as part of every database.

* tablespace_datafile_clauses

Use these clauses to specify attributes for all data files comprising the SYSTEM and SYSAUX tablespaces in the seed PDB.

Incorrect:

Not D: a PDB can not have an undo tablespace. Instead, it uses the undo tablespace belonging to the CDB.

Note:

* Example:

```
CONN pdb_admin@pdb1 -
SELECT tablespace_name FROM dba_tablespaces;

TABLESPACE_NAME -
------------------------------

SYSTEM -

SYSAUX -

TEMP -

USERS -
SQL>
```

Question #62 *Topic 1*

Which two statements are true about variable extent size support for large ASM files? (Choose two.)

A. The metadata used to track extents in SGA is reduced.

B. Rebalance operations are completed faster than with a fixed extent size

C. An ASM Instance automatically allocates an appropriate extent size.

D. Resync operations are completed faster when a disk comes online after being taken offline.

E. Performance improves in a stretch cluster configuration by reading from a local copy of an extent.

Correct Answer: *AC*
A: Variable size extents enable support for larger ASM datafiles, reduce SGA memory requirements for very large databases (A), and improve performance for file create and open operations.
C: You don't have to worry about the sizes; the ASM instance automatically allocates the appropriate extent size.
Note:
* The contents of ASM files are stored in a disk group as a set, or collection, of data extents that are stored on individual disks within disk groups. Each extent resides on an individual disk. Extents consist of one or more allocation units (AU). To accommodate increasingly larger files, ASM uses variable size extents.
* The size of the extent map that defines a file can be smaller by a factor of 8 and 64 depending on the file size. The initial extent size is equal to the allocation unit size and it increases by a factor of 8 and 64 at predefined thresholds. This feature is automatic for newly created and resized datafiles when the disk group compatibility attributes are set to Oracle Release 11 or higher.

TuxBingo 4 years ago
I agree with the answer, please refere to the link: https://docs.oracle.com/cd/B28359_01/server.111/b31107/asmcon.htm#OSTMG94058
upvoted 2 times

Question #63 *Topic 1*

You executed a DROP USER CASCADE on an Oracle 11g release 1 database and immediately realized that you forgot to copy the OCA.EXAM_RESULTS table to the OCP schema.
The RECYCLE_BIN enabled before the DROP USER was executed and the OCP user has been granted the FLASHBACK ANY TABLE system privilege.
What is the quickest way to recover the contents of the OCA.EXAM_RESULTS table to the OCP schema?

A. Execute FLASHBACK TABLE OCA.EXAM_RESULTS TO BEFORE DROP RENAME TO OCP.EXAM_RESULTS; connected as SYSTEM.

B. Recover the table using traditional Tablespace Point In Time Recovery.

C. Recover the table using Automated Tablespace Point In Time Recovery.

D. Recovery the table using Database Point In Time Recovery.

E. Execute FLASHBACK TABLE OCA.EXAM_RESULTS TO BEFORE DROP RENAME TO EXAM_RESULTS; connected as the OCP user.

Correct Answer: *C*
RMAN tablespace point-in-time recovery (TSPITR).
Recovery Manager (RMAN) TSPITR enables quick recovery of one or more tablespaces in a database to an earlier time without affecting the rest of the tablespaces and objects in the database.
Fully Automated (the default)
In this mode, RMAN manages the entire TSPITR process including the auxiliary instance. You specify the tablespaces of the recovery set, an auxiliary destination, the target time, and you allow RMAN to manage all other aspects of TSPITR.
The default mode is recommended unless you specifically need more control over the location of recovery set files after TSPITR, auxiliary set files during TSPITR, channel settings and parameters or some other aspect of your auxiliary instance.

TuxBingo 4 years ago
Agree with the answer C
Lik Oracle documentation:https://docs.oracle.com/cd/E11882_01/backup.112/e10642/rcmtspit.htm#BRADV89790
upvoted 2 times

Phoenix22 4 years, 4 months ago
Why is not E?
upvoted 1 times

TuxBingo 4 years ago
Hi Phoenix22
When you perform a drop user or drop user cascade, the recycle bin is purge by default. Reference:
https://docs.oracle.com/cd/B19306_01/server.102/b14200/statements_9008.htm "When you drop a user, Oracle Database also purges all of th user's schema objects from the recycle bin."

Regards.
upvoted 5 times

Blade69 2 years, 4 months ago
This is a version 11g database if it was version 12 then E is the correct answer.
upvoted 1 times

Question #64 *Topic 1*

In your multitenant container database (CDB) containing pluggable database (PDBs), the HR user executes the following commands to create and grant privileges on a procedure:

```
CREATE OR REPLACE PROCEDURE create_test_v (v_emp_id NUMBER, v_ename VARCHAR2, v_SALARY NUMBER, v_dept_id NUMBER)

BEGIN -
INSERT INTO hr.test VALUES (V_emp_id, V_ename, V_salary, V_dept_id);
END;
/
GRANT EXECUTE ON CREATE_TEST TO john, jim, smith, king;
```

How can you prevent users having the EXECUTE privilege on the CREATE_TEST procedure from inserting values into tables on which they do not have any privileges?

A. Create the CREATE_TEST procedure with definer's rights.

B. Grant the EXECUTE privilege to users with GRANT OPTION on the CREATE_TEST procedure.

C. Create the CREATE_TEST procedure with invoker's rights.

D. Create the CREATE_TEST procedure as part of a package and grant users the EXECUTE privilege the package.

Correct Answer: *C*

If a program unit does not need to be executed with the escalated privileges of the definer, you should specify that the program unit executes with the privileges of the caller, also known as the invoker. Invoker's rights can mitigate the risk of SQL injection.

Incorrect:

Not A: By default, stored procedures and SQL methods execute with the privileges of their owner, not their current user. Such definer-rights subprograms are bound to the schema in which they reside. not B: Using the GRANT option, a user can grant an Object privilege to another user or to PUBLIC.

Question #65 *Topic 1*

You created a new database using the "create database" statement without specifying the "ENABLE PLUGGABLE" clause.
What are two effects of not using the "ENABLE PLUGGABLE database" clause?

A. The database is created as a non-CDB and can never contain a PDB.

B. The database is treated as a PDB and must be plugged into an existing multitenant container database (CDB).

C. The database is created as a non-CDB and can never be plugged into a CDB.

D. The database is created as a non-CDB but can be plugged into an existing CDB.

E. The database is created as a non-CDB but will become a CDB whenever the first PDB is plugged in.

Correct Answer: *AD*
A (not B,not E): The CREATE DATABASE ... ENABLE PLUGGABLE DATABASE SQL statement creates a new CDB. If you do not specify the ENABLE
PLUGGABLE DATABASE clause, then the newly created database is a non-CDB and can never contain PDBs.
D: You can create a PDB by plugging in a Non-CDB as a PDB.
The following graphic depicts the options for creating a PDB:

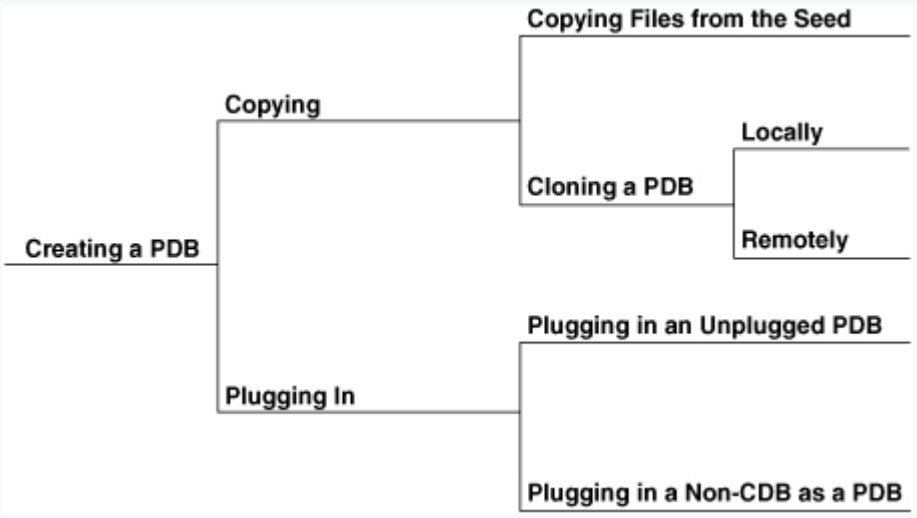

Incorrect:
Not E: For the duration of its existence, a database is either a CDB or a non-CDB. You cannot transform a non-CDB into a CDB or vice versa. You must define a database as a CDB at creation, and then create PDBs within this CDB.

Question #66 *Topic 1*

What is the effect of specifying the "ENABLE PLUGGABLE DATABASE" clause in a "CREATE DATABASE" statement?

A. It will create a multitenant container database (CDB) with only the root opened.

B. It will create a CDB with root opened and seed read only.

C. It will create a CDB with root and seed opened and one PDB mounted.

D. It will create a CDB that must be plugged into an existing CDB.

E. It will create a CDB with root opened and seed mounted.

Correct Answer: *B*

* The CREATE DATABASE ... ENABLE PLUGGABLE DATABASE SQL statement creates a new CDB. If you do not specify the ENABLE PLUGGABLE DATABASE clause, then the newly created database is a non-CDB and can never contain PDBs.

Along with the root (CDB$ROOT), Oracle Database automatically creates a seed PDB (PDB$SEED). The following graphic shows a newly created CDB:

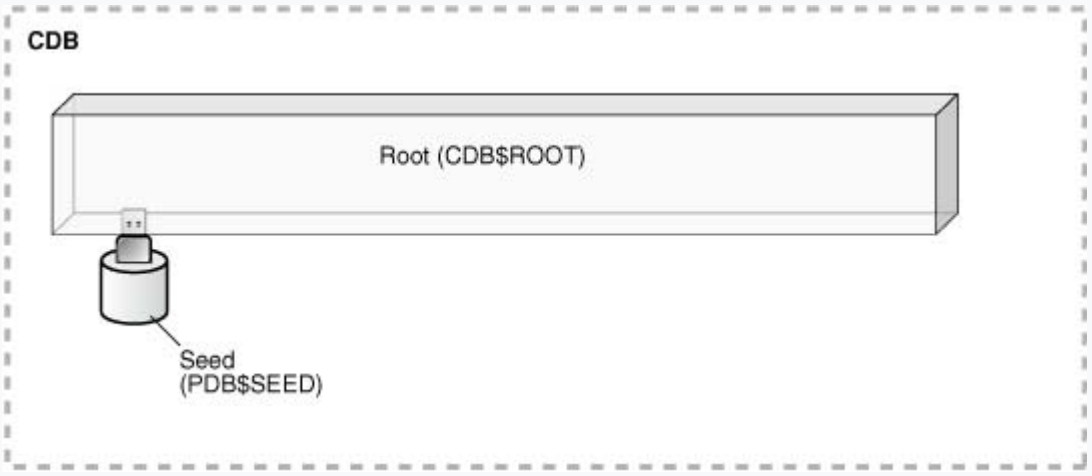

* Creating a PDB

Rather than constructing the data dictionary tables that define an empty PDB from scratch, and then populating its Obj$ and Dependency$ tables, the empty PDB is created when the CDB is created. (Here, we use empty to mean containing no customer-created artifacts.) It is referred to as the seed PDB and has the name

PDB$Seed. Every CDB non-negotiably contains a seed PDB; it is non-negotiably always open in read-only mode. This has no conceptual significance; rather, it is just an optimization device. The create PDB operation is implemented as a special case of the clone PDB operation.

jackymak 2 years, 5 months ago

That means create a CDB will not include a PDB by default?

upvoted 2 times

Question #67 *Topic 1*

You have installed two 64G flash devices to support the Database Smart Flash Cache feature on your database server that is running on Oracle Linux.
You have set the DB_SMART_FLASH_FILE parameter:
DB_FLASH_CACHE_FILE= /dev/flash_device_1 , /dev/flash_device_2
How should the DB_FLASH_CACHE_SIZE be configured to use both devices?

A. Set DB_FLASH_CACHE_ZISE = 64G.

B. Set DB_FLASH_CACHE_ZISE = 64G, 64G

C. Set DB_FLASH_CACHE_ZISE = 128G.

D. DB_FLASH_CACHE_SIZE is automatically configured by the instance at startup.

Correct Answer: *B*
* Smart Flash Cache concept is not new in Oracle 12C - DB Smart Flash Cache in Oracle 11g.
In this release Oracle has made changes related to both initialization parameters used by DB Smart Flash cache. Now you can define many files|devices and its sizes for "Database Smart Flash Cache" area. In previous releases only one file|device could be defined.
DB_FLASH_CACHE_FILE = /dev/sda, /dev/sdb, /dev/sdc
DB_FLASH_CACHE_SIZE = 32G, 32G, 64G
So above settings defines 3 devices which will be in use by "DB Smart Flash Cache"
/dev/sda size 32G
/dev/sdb size 32G
/dev/sdc size 64G
New view V$FLASHFILESTAT its used to determine the cumulative latency and read counts of each file|device and compute the average latency

Question #68 *Topic 1*

Examine the following parameters for a database instance:

MEMORY_MAX_TARGET=0 -

MEMORY_TARGET=0 -

SGA_TARGET=0 -

PGA_AGGREGATE_TARGET=500m -
Which three initialization parameters are not controlled by Automatic Shared Memory Management (ASMM)? (Choose three.)

A. LOG_BUFFER

B. SORT_AREA_SIZE

C. JAVA_POOL_SIZE

D. STREAMS_POOL_SIZE

E. DB_16K_CACHE_SZIE

F. DB_KEEP_CACHE_SIZE

Correct Answer: *AEF*
Manually Sized SGA Components that Use SGA_TARGET Space
SGA Component, Initialization Parameter
/ The log buffer

LOG_BUFFER -
/ The keep and recycle buffer caches

DB_KEEP_CACHE_SIZE -

DB_RECYCLE_CACHE_SIZE -
/ Nonstandard block size buffer caches

DB_nK_CACHE_SIZE -
Note:
* In addition to setting SGA_TARGET to a nonzero value, you must set to zero all initialization parameters listed in the table below to enable full automatic tuning of the automatically sized SGA components.
* Table, Automatically Sized SGA Components and Corresponding Parameters

SGA Component	Initialization Parameter
Fixed SGA and other internal allocations needed by the Oracle Database instance	N/A
The shared pool	SHARED_POOL_SIZE
The large pool	LARGE_POOL_SIZE
The Java pool	JAVA_POOL_SIZE
The buffer cache	DB_CACHE_SIZE
The Streams pool	STREAMS_POOL_SIZE

vineetarneja Highly Voted 4 years, 3 months ago
The option B SORT_AREA_SIZE is a part of pga , so ideally it is also not controlled by ASMM as ASMM only controls the components of SGA.

what do you guys think?

Regards
Vineet Arneja
vineetarneja1989@gmail.com
upvoted 7 times

jackymak Most Recent 2 years, 5 months ago

A,B,F is what I think.
upvoted 1 times

IamGo0ke 3 years, 3 months ago

A,E,F is correct. Under automatic memory control the following memory components will be managed automatically, you can override the ASMM by specifying the size of the component in the init.ora file or setting its value (spfile), remember that the values will be deducted from the SGA_TARGET value. The values also state the minimum value that will be allocated.

Buffer cache (DB_CACHE_SIZE)
shared pool (SHARED_POOL_SIZE)
large pool (LARGE_POOL_SIZE)
java pool (JAVA_POOL_SIZE)
streams pool (STREAMS_POOL_SIZE)
When using ASMM you still need to manually configure

log buffer (LOG_BUFFER)
db buffer cache keep pool (DB_KEEP_CACHE_SIZE)
db buffer recycle pool (DB_RECYCLE_CACHE_SIZE)
db buffer nK block size pools (DB_nK_CACHE_SIZE) https://sites.google.com/site/oracledb009/database-concepts/asmm
upvoted 2 times

fridaytar 3 years, 5 months ago

It should be - A, B, F ??
upvoted 2 times

Question #69 *Topic 1*

Examine the contents of SQL loader control file:

```
LOAD DATA
INFILE myfile1.dat
INFILE myfile2.dat
FIELD NAMES FIRST FILE
APPEND
INTO TABLE EMP
FIELDS CSV WITH EMBEDDED
DATE FORMAT "DD-Month_YYYY"
(empno,
 ename,
 job,
 mgr,
 hiredate DATE,
 sal,
 comm,
 deptno,
 entrydate DATE)
```

Which three statements are true regarding the SQL* Loader operation performed using the control file? (Choose three.)

A. An EMP table is created if a table does not exist. Otherwise, if the EMP table is appended with the loaded data.

B. The SQL* Loader data file myfile1.dat has the column names for the EMP table.

C. The SQL* Loader operation fails because no record terminators are specified.

D. Field names should be the first line in the both the SQL* Loader data files.

E. The SQL* Loader operation assumes that the file must be a stream record format file with the normal carriage return string as the record terminator.

Correct Answer: *ABE*

A: The APPEND keyword tells SQL*Loader to preserve any preexisting data in the table. Other options allow you to delete preexisting data, or to fail with an error if the table is not empty to begin with.

B (not D):

Note:

* SQL*Loader-00210: first data file is empty, cannot process the FIELD NAMES record

Cause: The data file listed in the next message was empty. Therefore, the FIELD NAMES FIRST FILE directive could not be processed.

Action: Check the listed data file and fix it. Then retry the operation

E:

* A comma-separated values (CSV) (also sometimes called character-separated values, because the separator character does not have to be a comma) file stores tabular data (numbers and text) in plain-text form. Plain text means that the file is a sequence of characters, with no data that has to be interpreted instead, as binary numbers. A CSV file consists of any number of records, separated by line breaks of some kind; each record consists of fields, separated by some other character or string, most commonly a literal comma or tab. Usually, all records have an identical sequence of fields.

* Fields with embedded commas must be quoted.

Example:

1997,Ford,E350,"Super, luxurious truck"

Note:

* SQL*Loader is a bulk loader utility used for moving data from external files into the Oracle database.

Community vote distribution

ABE (100%)

jackymak 2 years, 5 months ago

Selected Answer: ABE

A maybe not right, C,D must be wrong.
Seem to be the CREATE TABLE statement will be there while generating .ctl file,
but you need to manually execute the CREATE statement anyway.
https://docs.safe.com/fme/html/FME_Desktop_Documentation/FME_ReadersWriters/sqlldr/sqlldr_table_creation.htm

upvoted 1 times

hggz 4 years, 4 months ago

I'am confused with this, "A" can't be correct, but some guys tell what thay chosen this answer and got that point.

http://www.briefmenow.org/oracle/which-three-statements-are-true-regarding-the-sql-loader-operation-performed-using-the-control-file-2/
upvoted 3 times

Question #70 *Topic 1*

In your multitenant container database (CDB) containing pluggable database (PDBs), you granted the CREATE TABLE privilege to the common user C # #
A_ADMIN in root and all PDBs. You execute the following command from the root container:
SQL > REVOKE create table FROM C # # A_ADMIN;
What is the result?

A. It executes successfully and the CREATE TABLE privilege is revoked from C # # A_ADMIN in root only.

B. It fails and reports an error because the CONTAINER=ALL clause is not used.

C. It excludes successfully and the CREATE TABLE privilege is revoked from C # # A_ADMIN in root and all PDBs.

D. It fails and reports an error because the CONTAINER=CURRENT clause is not used.

E. It executes successfully and the CREATE TABLE privilege is revoked from C # # A_ADMIN in all PDBs.

Correct Answer: *A*
REVOKE ..FROM -
If the current container is the root:
/ Specify CONTAINER = CURRENT to revoke a locally granted system privilege, object privilege, or role from a common user or common role.
The privilege or role is revoked from the user or role only in the root. This clause does not revoke privileges granted with CONTAINER = ALL.
/ Specify CONTAINER = ALL to revoke a commonly granted system privilege, object privilege on a common object, or role from a common user or common role.
The privilege or role is revoked from the user or role across the entire CDB. This clause can revoke only a privilege or role granted with CONTAINER = ALL from the specified common user or common role. This clause does not revoke privileges granted locally with CONTAINER = CURRENT. However, any locally granted privileges that depend on the commonly granted privilege being revoked are also revoked.
If you omit this clause, then CONTAINER = CURRENT is the default.
References:

bigswole 2 years, 10 months ago
yeah this is a tricky one. when i read the portion "you granted the CREATE TABLE privilege to the common user C ## A_ADMIN in root and all PDBs". i incorrectly assumed that Container = All clause was implied in this scenario. I tested it by granting the system privilege separately and confirmed the answer is A
upvoted 3 times

exampro46 3 years, 2 months ago
I tested it when you grant with container=all the privilege adds to all CDB and PDBS and the answer is B, but when you grant with separate in CD and PDBS then answer is A, I hate this questions which not clear but tricky is there written "in root and all PDB". That means privilege given separate I think and answer is A.
upvoted 2 times

rcgenilo27 4 years, 11 months ago
B is the right answer
upvoted 3 times

5kong 4 years, 8 months ago
why? If you omit this clause, then CONTAINER = CURRENT is the default.
upvoted 5 times

katetel 2 years, 11 months ago
yes answer is a
upvoted 1 times

Question #71 *Topic 1*

Which two statements are true concerning the Resource Manager plans for individual pluggable databases (PDB plans) in a multitenant container database
(CDB)? (Choose two.)

A. If no PDB plan is enabled for a pluggable database, then all sessions for that PDB are treated to an equal degree of the resource share of that PDB.

B. In a PDB plan, subplans may be used with up to eight consumer groups.

C. If a PDB plan is enabled for a pluggable database, then resources are allocated to consumer groups across all PDBs in the CDB.

D. If no PDB plan is enabled for a pluggable database, then the PDB share in the CDB plan is dynamically calculated.

E. If a PDB plan is enabled for a pluggable database, then resources are allocated to consumer groups based on the shares provided to the PDB in the CDB plan

Correct Answer: *AE*
A: Setting a PDB resource plan is optional. If not specified, all sessions within the PDB are treated equally.
*
In a non-CDB database, workloads within a database are managed with resource plans.
In a PDB, workloads are also managed with resource plans, also called PDB resource plans.
The functionality is similar except for the following differences:
/ Non-CDB Database

Multi-level resource plans -

Up to 32 consumer groups -

Subplans -
/ PDB Database
Single-level resource plans only

Up to 8 consumer groups -
(not B) No subplans

Question #72 *Topic 1*

You use a recovery catalog for maintaining your database backups.
You execute the following command:
$rman TARGET / CATALOG rman / cat@catdb
RMAN > BACKUP VALIDATE DATABASE ARCHIVELOG ALL;
Which two statements are true? (Choose two.)

A. Corrupted blocks, if any, are repaired.

B. Checks are performed for physical corruptions.

C. Checks are performed for logical corruptions.

D. Checks are performed to confirm whether all database files exist in correct locations

E. Backup sets containing both data files and archive logs are created.

Correct Answer: *BD*
B (not C): You can validate that all database files and archived redo logs can be backed up by running a command as follows:
RMAN> BACKUP VALIDATE DATABASE ARCHIVELOG ALL;
This form of the command would check for physical corruption. To check for logical corruption,
RMAN> BACKUP VALIDATE CHECK LOGICAL DATABASE ARCHIVELOG ALL;
D: You can use the VALIDATE keyword of the BACKUP command to do the following:
Check datafiles for physical and logical corruption
Confirm that all database files exist and are in the correct locations.
Note:
You can use the VALIDATE option of the BACKUP command to verify that database files exist and are in the correct locations (D), and have no physical or logical corruptions that would prevent RMAN from creating backups of them. When performing a BACKUP...VALIDATE, RMAN reads the files to be backed up in their entirety, as it would during a real backup. It does not, however, actually produce any backup sets or image copies (Not A, not E).

Question #73 *Topic 1*

Which three statements are true concerning the multitenant architecture? (Choose three.)

A. Each pluggable database (PDB) has its own set of background processes.

B. A PDB can have a private temp tablespace.

C. PDBs can share the sysaux tablespace.

D. Log switches occur only at the multitenant container database (CDB) level.

E. Different PDBs can have different default block sizes.

F. PDBs share a common system tablespace.

G. Instance recovery is always performed at the CDB level.

Correct Answer: *BDG*

B:

* A PDB would have its SYSTEM, SYSAUX, TEMP tablespaces. It can also contains other user created tablespaces in it.

* There is one default temporary tablespace for the entire CDB. However, you can create additional temporary tablespaces in individual PDBs.

D:

* There is a single redo log and a single control file for an entire CDB

* A log switch is the point at which the database stops writing to one redo log file and begins writing to another. Normally, a log switch occurs when the current redo log file is completely filled and writing must continue to the next redo log file.

G: instance recovery -

The automatic application of redo log records to uncommitted data blocks when an database instance is restarted after a failure.

Incorrect:

Not A:

* There is one set of background processes shared by the root and all PDBs.

* High consolidation density. The many pluggable databases in a single container database share its memory and background processes, letting you operate many more pluggable databases on a particular platform than you can single databases that use the old architecture.

Not C: There is a separate SYSAUX tablespace for the root and for each PDB.

Not F: There is a separate SYSTEM tablespace for the root and for each PDB.

Question #74 *Topic 1*

You notice that the elapsed time for an important database scheduler Job is unacceptably long.
The job belongs to a scheduler job class and window.
Which two actions would reduce the job's elapsed time? (Choose two.)

A. Increasing the priority of the job class to which the job belongs

B. Increasing the job's relative priority within the Job class to which it belongs

C. Increasing the resource allocation for the consumer group mapped to the scheduler job's job class within the plan mapped to the scheduler window

D. Moving the job to an existing higher priority scheduler window with the same schedule and duration

E. Increasing the value of the JOB_QUEUE_PROCESSES parameter

F. Increasing the priority of the scheduler window to which the job belongs

Correct Answer: *BC*
B: Job priorities are used only to prioritize among jobs in the same class.
Note: Group jobs for prioritization
Within the same job class, you can assign priority values of 1-5 to individual jobs so that if two jobs in the class are scheduled to start at the same time, the one with the higher priority takes precedence. This ensures that you do not have a less important job preventing the timely completion of a more important one.
C: Set resource allocation for member jobs
Job classes provide the link between the Database Resource Manager and the Scheduler, because each job class can specify a resource consumer group as an attribute. Member jobs then belong to the specified consumer group and are assigned resources according to settings in the current resource plan.

TuxBingo 4 years ago
I agree B,C – is correct answer.

See Prioritizing Jobs
http://docs.oracle.com/database/121/ADMIN/scheduse.htm#CIAEGDJB

Setting Relative Job Priorities Within a Job Class
http://docs.oracle.com/database/121/ADMIN/scheduse.htm#i1016329bis
upvoted 1 times

Question #75 *Topic 1*

You plan to migrate your database from a File system to Automata Storage Management (ASM) on same platform.
Which two methods or commands would you use to accomplish this task? (Choose two.)

A. RMAN CONVERT command

B. Data Pump Export and import

C. Conventional Export and Import

D. The BACKUP AS COPY DATABASE . . . command of RMAN

E. DBMS_FILE_TRANSFER with transportable tablespace

Correct Answer: *AD*
A:
1. Get the list of all datafiles.
Note: RMAN Backup of ASM Storage
There is often a need to move the files from the file system to the ASM storage and vice versa. This may come in handy when one of the file systems is corrupted by some means and then the file may need to be moved to the other file system.
D: Migrating a Database into ASM
* To take advantage of Automatic Storage Management with an existing database you must migrate that database into ASM. This migration is performed using
Recovery Manager (RMAN) even if you are not using RMAN for your primary backup and recovery strategy.
* Example:
Back up your database files as copies to the ASM disk group.
BACKUP AS COPY INCREMENTAL LEVEL 0 DATABASE
FORMAT '+DISK' TAG 'ORA_ASM_MIGRATION';
References:

Question #76 *Topic 1*

You run a script that completes successfully using SQL*Plus that performs these actions:
1. Creates a multitenant container database (CDB)
2. Plugs in three pluggable databases (PDBs)
3. Shuts down the CDB instance
4. Starts up the CDB instance using STARTUP OPEN READ WRITE
Which two statements are true about the outcome after running the script? (Choose two.)

A. The seed will be in mount state.

B. The seed will be opened read-only.

C. The seed will be opened read/write.

D. The other PDBs will be in mount state.

E. The other PDBs will be opened read-only.

F. The PDBs will be opened read/write.

Correct Answer: *BD*
B: The seed is always read-only.
D: Pluggable databases can be started and stopped using SQL*Plus commands or the ALTER PLUGGABLE DATABASE command.

Question #77 *Topic 1*

You execute the following piece of code with appropriate privileges:

```
BEGIN
  DBMS_REDACT.ADD_POLICY(
    OBJECT_SCHEMA => 'SCOTT',
    OBJECT_NAME   => 'EMP',
    POLICY_NAME   => 'SCOTT_EMP',
    COLUMN_NAME   => 'SAL',
    FUNCTION_TYPE => DBMS_REDACT.FULL,
    EXPRESSION    => 'SYS_CONTEXT("SYS_SESSION_ROLES","MGR") = "FALSE"');
END;
/

CREATE VIEW SCOTT.EMP_V AS SELECT * FROM SCOTT.EMP;

BEGIN
  DBMS_REDACT.ADD_POLICY(
    OBJECT_SCHEMA => 'SCOTT',
    OBJECT_NAME   => 'EMP_V',
    POLICY_NAME   => 'SCOTT_EMP_V',
    COLUMN_NAME   => 'SAL',
    FUNCTION_TYPE => DBMS_REDACT.NONE,
    EXPRESSION    => 'SYS_CONTEXT("SYS_SESSION_ROLES","MGR") = "FALSE"');
END;
/
```

User SCOTT has been granted the CREATE SESSION privilege and the MGR role.

Which two statements are true when a session logged in as SCOTT queries the SAL column in the view and the table? (Choose two.)

A. Data is redacted for the EMP.SAL column only if the SCOTT session does not have the MGR role set.

B. Data is redacted for EMP.SAL column only if the SCOTT session has the MGR role set.

C. Data is never redacted for the EMP_V.SAL column.

D. Data is redacted for the EMP_V.SAL column only if the SCOTT session has the MGR role set.

E. Data is redacted for the EMP_V.SAL column only if the SCOTT session does not have the MGR role set.

Correct Answer: *AC*

Note:

* DBMS_REDACT.FULL completely redacts the column data.
* DBMS_REDACT.NONE applies no redaction on the column data. Use this function for development testing purposes. LOB columns are not supported.
* The DBMS_REDACT package provides an interface to Oracle Data Redaction, which enables you to mask (redact) data that is returned from queries issued by low-privileged users or an application.
* If you create a view chain (that is, a view based on another view), then the Data Redaction policy also applies throughout this view chain. The policies remain in effect all of the way up through this view chain, but if another policy is created for one of these views, then for the columns affected in the subsequent views, this new policy takes precedence.

Question #78 *Topic 1*

Your database is open and the LISTENER listener running. You stopped the wrong listener LISTENER by issuing the following command:
1snrctl > STOP
What happens to the sessions that are presently connected to the database Instance?

A. They are able to perform only queries.

B. They are not affected and continue to function normally.

C. They are terminated and the active transactions are rolled back.

D. They are not allowed to perform any operations until the listener LISTENER is started.

Correct Answer: *B*
The listener is used when the connection is established. The immediate impact of stopping the listener will be that no new session can be established from a remote host. Existing sessions are not compromised.

Question #79 *Topic 1*

Which three statements are true about using flashback database in a multitenant container database (CDB)? (Choose three.) (Choose three.)

A. The root container can be flashed back without flashing back the pluggable databases (PDBs).

B. To enable flashback database, the CDB must be mounted.

C. Individual PDBs can be flashed back without flashing back the entire CDB.

D. The DB_FLASHBACK RETENTION_TARGET parameter must be set to enable flashback of the CDB.

E. A CDB can be flashed back specifying the desired target point in time or an SCN, but not a restore point.

Correct Answer: *ABD*

Muhammadnagah Highly Voted 4 years, 9 months ago
ABD

CDB must be mounted to set the parameter DB_FLASHBACK RETENTION_TARGET and enable flashback database
upvoted 16 times

jackymak 2 years, 5 months ago
www.dba-oracle.com/p_db_flashback_recovery_target.htm
There is a default values for DB_FLASHBACK RETENTION_TARGET.
upvoted 1 times

abdullahalbyati Highly Voted 4 years, 4 months ago
can we add a downvote button? some of these comments are making the confusing tricky questions more confusing
upvoted 7 times

NorthKorean 4 years, 4 months ago
we should upvote the right one
upvoted 2 times

Ray520 Most Recent 2 years ago
I don't know why A is marked true (although C and E are definitely incorrect, so all that's left is A) because if we flahsback CDB, PDB also gets flashed back. What am I missing here?
upvoted 1 times

charan94 3 years, 4 months ago
i think we can flashback pdb leaving cdb. But i think this is possilbe later releases of 12.1
upvoted 1 times

TuxBingo 4 years ago
ABD is the correct answer, check this url: https://oracle-base.com/articles/12c/multitenant-flashback-of-container-database-12cr1
upvoted 2 times

rcgenilo27 4 years, 11 months ago
ACD is the correct answer
upvoted 1 times

Ray520 2 years, 3 months ago
you can do PITR for ndividual PDB, not flashback. Flashback command only works in CDB. So C is incorrect.
upvoted 1 times

Question #80 *Topic 1*

You execute the following PL/SQL:

```
BEGIN
DBMS_FGA.add_policy(
object_schema => 'JIM',
object_name => 'PRODUCTS',
policy_name => 'PROD_AUDIT',
audit_condition => 'PRICE > 10000',
audit_column => 'PRICE');
END;
/
```

Which two statements are true? (Choose two.)

A. Fine-Grained Auditing (FGA) is enabled for the PRICE column in the PRODUCTS table for SELECT statements only when a row with PRICE > 10000 is accessed.

B. FGA is enabled for the PRODUCTS.PRICE column and an audit record is written whenever a row with PRICE > 10000 is accessed.

C. FGA is enabled for all DML operations by JIM on the PRODUCTS.PRICE column.

D. FGA is enabled for the PRICE column of the PRODUCTS table and the SQL statements is captured in the FGA audit trial.

Correct Answer: *AB*

DBMS_FGA.add_policy -

* The DBMS_FGA package provides fine-grained security functions.

* ADD_POLICY Procedure

This procedure creates an audit policy using the supplied predicate as the audit condition.

Incorrect:

Not C: object_schema -

The schema of the object to be audited. (If NULL, the current log-on user schema is assumed.)

Community vote distribution

AD (100%)

yosiw96816 Highly Voted 3 years, 6 months ago

Not B - statement_types is not set (default is SELCT)
AD

upvoted 6 times

yosiw96816 3 years, 6 months ago

REF https://docs.oracle.com/database/121/ARPLS/d_fga.htm#ARPLS66351

upvoted 2 times

Ray520 Most Recent 2 years ago

Selected Answer: AD

STATEMENT_TYPE-> SELECT (default) since not mentioned, hence A.
D-SQL text and bind variables get captured in audit logs (audit_policy default is DB+EXTENDED which sends the audit trail to the SYS.FGA_LOG$ table in the database and includes SQL Text and SQL Bind)

upvoted 1 times

jackymak 2 years, 5 months ago

Selected Answer: AD

I think AD

upvoted 3 times

katetel 2 years, 11 months ago

the answer is a,b

upvoted 2 times

charan94 3 years, 3 months ago

I think Answer is BD

upvoted 1 times

NorthKorean 4 years, 4 months ago

What about D ?

upvoted 2 times

jackymak 2 years, 5 months ago

enable IN BOOLEAN DEFAULT TRUE,

upvoted 1 times

Question #81 *Topic 1*

You execute the following commands to audit database activities:
SQL > ALTER SYSTEM SET AUDIT_TRIAL=DB, EXTENDED SCOPE=SPFILE;
SQL > AUDIT SELECT TABLE, INSERT TABLE, DELETE TABLE BY JOHN By SESSION WHENEVER SUCCESSFUL;
Which statement is true about the audit record that generated when auditing after instance restarts?

A. One audit record is created for every successful execution of a SELECT, INSERT OR DELETE command on a table, and contains the SQL text for the SQL Statements.

B. One audit record is created for every successful execution of a SELECT, INSERT OR DELETE command, and contains the execution plan for the SQL statements.

C. One audit record is created for the whole session if john successfully executes a SELECT, INSERT, or DELETE command, and contains the execution plan for the SQL statements.

D. One audit record is created for the whole session if JOHN successfully executes a select command, and contains the SQL text and bind variables used.

E. One audit record is created for the whole session if john successfully executes a SELECT, INSERT, or DELETE command on a table, and contains the

Correct Answer: *A*
Note:
* BY SESSION
In earlier releases, BY SESSION caused the database to write a single record for all SQL statements or operations of the same type executed on the same schema objects in the same session. Beginning with this release (11g) of Oracle Database, both BY SESSION and BY ACCESS cause Oracle Database to write one audit record for each audited statement and operation.
* BY ACCESS
Specify BY ACCESS if you want Oracle Database to write one record for each audited statement and operation.
Note:
If you specify either a SQL statement shortcut or a system privilege that audits a data definition language (DDL) statement, then the database always audits by access. In all other cases, the database honors the BY SESSION or BY ACCESS specification.
* For each audited operation, Oracle Database produces an audit record containing this information:
/ The user performing the operation
/ The type of operation
/ The object involved in the operation
/ The date and time of the operation
References:

Community vote distribution
A (100%)

luizyto 1 year, 9 months ago

Selected Answer: A
https://docs.oracle.com/database/121/SQLRF/statements_4007.htm#SQLRF01107
upvoted 1 times

luizyto 1 year, 9 months ago

A, one audit for each statement.
https://docs.oracle.com/database/121/SQLRF/statements_4007.htm#SQLRF01107
upvoted 1 times

jackymak 2 years, 5 months ago

I think It is E, D is missing the INSERT and DELETE
upvoted 1 times

bigswole 2 years, 9 months ago

E is correct since AUDIT_TRAIL=db, extended was used

"db, extended

Performs all actions of AUDIT_TRAIL=db, and also populates the SQL bind and SQL text CLOB-type columns of the SYS.AUD$ table, when availabl
These two columns are populated only when this parameter is specified."

- https://docs.oracle.com/database/121/REFRN/GUID-BD86F593-B606-4367-9FB6-8DAB2E47E7FA.htm#REFRN10006

upvoted 2 times

katetel 2 years, 10 months ago

correct answer is : E
E. One audit record is created for the whole session if john successfully executes a SELECT, INSERT, or DELETE command on a table, and contains the execution plan, SQL text, and bind variables used.

upvoted 3 times

lollo1234 3 years, 2 months ago

I go for D. BY SESSION means that 1 Record for the whole session is beeing created. Also AUDIT_TRIAL=DB, EXTENDED means:
Performs all actions of AUDIT_TRAIL=db, and also populates the SQL bind and SQL text CLOB-type columns of the SYS.AUD$ table, when availabl
Reference: https://docs.oracle.com/cd/E11882_01/server.112/e40402/initparams017.htm#REFRN10006

upvoted 2 times

TuxBingo 3 years, 3 months ago

Correct Asnwer is: A

upvoted 1 times

charan94 3 years, 4 months ago

Answer is D. As question clear mentions user john should be audited..

upvoted 1 times

Modasser 3 years, 6 months ago

Answer is A

upvoted 2 times

dancymonkey 3 years, 11 months ago

E is not clear
I would go for D

upvoted 1 times

Question #82 *Topic 1*

You support Oracle Database 12c Oracle Database 11g, and Oracle Database log on the same server.
All databases of all versions use Automatic Storage Management (ASM).
Which three statements are true about the ASM disk group compatibility attributes that are set for a disk group? (Choose three.)

A. The ASM compatibility attribute controls the format of the disk group metadata.

B. RDBMS compatibility together with the database version determines whether a database Instance can mount the ASM disk group.

C. The RDBMS compatibility setting allows only databases set to the same version as the compatibility value, to mount the ASM disk group.

D. The ASM compatibility attribute determines some of the ASM features that may be used by the Oracle disk group.

E. The ADVM compatibility attribute determines the ACFS features that may be used by the Oracle 10 g database.

Correct Answer: *ABD*
AD: The value for the disk group COMPATIBLE.ASM attribute determines the minimum software version for an Oracle ASM instance that can use the disk group.
This setting also affects the format of the data structures for the Oracle ASM metadata on the disk.
B: The value for the disk group COMPATIBLE.RDBMS attribute determines the minimum COMPATIBLE database initialization parameter setting for any database instance that is allowed to use the disk group. Before advancing the COMPATIBLE.RDBMS attribute, ensure that the values for the COMPATIBLE initialization parameter for all of the databases that access the disk group are set to at least the value of the new setting for COMPATIBLE.RDBMS.
For example, if the COMPATIBLE initialization parameters of the databases are set to either 11.1 or 11.2, then COMPATIBLE.RDBMS can be set to any value between 10.1 and 11.1 inclusively.
Not E:
/The value for the disk group COMPATIBLE.ADVM attribute determines whether the disk group can contain Oracle ASM volumes. The value must be set to 11.2 or higher. Before setting this attribute, the COMPATIBLE.ASM value must be 11.2 or higher. Also, the Oracle ADVM volume drivers must be loaded in the supported environment.
/ You can create an Oracle ASM Dynamic Volume Manager (Oracle ADVM) volume in a disk group. The volume device associated with the dynamic volume can then be used to host an Oracle ACFS file system.
The compatibility parameters COMPATIBLE.ASM and COMPATIBLE.ADVM must be set to 11.2 or higher for the disk group.
Note:
* The disk group attributes that determine compatibility are COMPATIBLE.ASM, COMPATIBLE.RDBMS. and COMPATIBLE.ADVM. The COMPATIBLE.ASM and
COMPATIBLE.RDBMS attribute settings determine the minimum Oracle Database software version numbers that a system can use for Oracle ASM and the database instance types respectively. For example, if the Oracle ASM compatibility setting is 11.2, and RDBMS compatibility is set to 11.1, then the Oracle ASM software version must be at least 11.2, and the Oracle Database client software version must be at least 11.1. The COMPATIBLE.ADVM attribute determines whether the Oracle ASM Dynamic Volume Manager feature can create an volume in a disk group.

Question #83 *Topic 1*

To enable the Database Smart Flash Cache, you configure the following parameters:
DB_FLASH_CACHE_FILE = /dev/flash_device_1 , /dev/flash_device_2

DB_FLASH_CACHE_SIZE=64G -
What is the result when you start up the database instance?

A. It results in an error because these parameter settings are invalid.

B. One 64G flash cache file will be used.

C. Two 64G flash cache files will be used.

D. Two 32G flash cache files will be used.

Correct Answer: *A*

Oracle2020 3 years ago
what is wrong?
upvoted 1 times

AGoodStuff 2 years, 11 months ago
please refer to Question #67,
DB_FLASH_CACHE_SIZE should set like this: DB_FLASH_CACHE_SIZE=64G,64G
upvoted 2 times

Question #84 *Topic 1*

You executed this command to create a password file:
$ orapwd file = orapworcl entries = 10 ignorecase = N
Which two statements are true about the password file? (Choose two.)

A. It will permit the use of uppercase passwords for database users who have been granted the SYSOPER role.

B. It contains username and passwords of database users who are members of the OSOPER operating system group.

C. It contains usernames and passwords of database users who are members of the OSDBA operating system group.

D. It will permit the use of lowercase passwords for database users who have granted the SYSDBA role.

E. It will not permit the use of mixed case passwords for the database users who have been granted the SYSDBA role.

Correct Answer: *AD*
* You can create a password file using the password file creation utility, ORAPWD.
* Adding Users to a Password File
When you grant SYSDBA or SYSOPER privileges to a user, that user's name and privilege information are added to the password file. If the server does not have an EXCLUSIVE password file (that is, if the initialization parameter REMOTE_LOGIN_PASSWORDFILE is NONE or SHARED, or the password file is missing),
Oracle Database issues an error if you attempt to grant these privileges.
A user's name remains in the password file only as long as that user has at least one of these two privileges. If you revoke both of these privileges, Oracle
Database removes the user from the password file.
* The syntax of the ORAPWD command is as follows:
ORAPWD FILE=filename [ENTRIES=numusers]
[FORCE={Y|N}] [IGNORECASE={Y|N}] [NOSYSDBA={Y|N}]
* IGNORECASE
If this argument is set to y, passwords are case-insensitive. That is, case is ignored when comparing the password that the user supplies during login with the password in the password file.

Question #85 *Topic 1*

Identify three valid methods of opening, pluggable databases (PDBs).

A. ALTER PLUGGABLE DATABASE OPEN ALL ISSUED from the root

B. ALTER PLUGGABLE DATABASE OPEN ALL ISSUED from a PDB

C. ALTER PLUGGABLE DATABASE PDB OPEN issued from the seed

D. ALTER DATABASE PDB OPEN issued from the root

E. ALTER DATABASE OPEN issued from that PDB

F. ALTER PLUGGABLE DATABASE PDB OPEN issued from another PDB

G. ALTER PLUGGABLE DATABASE OPEN issued from that PDB

Correct Answer: *AEG*

E: You can perform all ALTER PLUGGABLE DATABASE tasks by connecting to a PDB and running the corresponding ALTER DATABASE statement. This functionality is provided to maintain backward compatibility for applications that have been migrated to a CDB environment.

AG: When you issue an ALTER PLUGGABLE DATABASE OPEN statement, READ WRITE is the default unless a PDB being opened belongs to a CDB that is used as a physical standby database, in which case READ ONLY is the default.

You can specify which PDBs to modify in the following ways:

List one or more PDBs.

Specify ALL to modify all of the PDBs.

Specify ALL EXCEPT to modify all of the PDBs, except for the PDBs listed.

Question #86 *Topic 1*

You administer an online transaction processing (OLTP) system whose database is stored in Automatic Storage Management (ASM) and whose disk group use normal redundancy.
One of the ASM disks goes offline, and is then dropped because it was not brought online before DISK_REPAIR_TIME elapsed.
When the disk is replaced and added back to the disk group, the ensuing rebalance operation is too slow.
Which two recommendations should you make to speed up the rebalance operation if this type of failure happens again? (Choose two.)

A. Increase the value of the ASM_POWER_LIMIT parameter.

B. Set the DISK_REPAIR_TIME disk attribute to a lower value.

C. Specify the statement that adds the disk back to the disk group.

D. Increase the number of ASMB processes.

E. Increase the number of DBWR_IO_SLAVES in the ASM instance.

Correct Answer: *AD*
A: ASM_POWER_LIMIT specifies the maximum power on an Automatic Storage Management instance for disk rebalancing. The higher the limit, the faster rebalancing will complete. Lower values will take longer, but consume fewer processing and I/O resources.
D:
* Normally a separate process is fired up to do that rebalance. This will take a certain amount of time. If you want it to happen faster, fire up more processes. You tell ASM it can add more processes by increasing the rebalance power.
* ASMB

ASM Background Process -
Communicates with the ASM instance, managing storage and providing statistics
Incorrect:
Not B: A higher, not a lower, value of DISK_REPAIR_TIME would be helpful here.
Not E: If you implement database writer I/O slaves by setting the DBWR_IO_SLAVES parameter, you configure a single (master) DBWR process that has slave processes that are subservient to it. In addition, I/O slaves can be used to "simulate" asynchronous I/O on platforms that do not support asynchronous I/O or implement it inefficiently. Database I/O slaves provide non-blocking, asynchronous requests to simulate asynchronous I/O.

TuxBingo 4 years ago
I agree with the response, the trick is tha the disk was replaced and the rebalance process is to slow -- this is the key.
upvoted 2 times

Question #87 *Topic 1*

You are administering a database and you receive a requirement to apply the following restrictions:
1. A connection must be terminated after four unsuccessful login attempts by user.
2. A user should not be able to create more than four simultaneous sessions.
3. User session must be terminated after 15 minutes of inactivity.
4. Users must be prompted to change their passwords every 15 days.
How would you accomplish these requirements?

A. by granting a secure application role to the users

B. by creating and assigning a profile to the users and setting the REMOTE_OS_AUTHENT parameter to FALSE

C. By creating and assigning a profile to the users and setting the SEC_MAX_FAILED_LOGIN_ATTEMPTS parameter to 4

D. By Implementing Fine-Grained Auditing (FGA) and setting the REMOTE_LOGIN_PASSWORD_FILE parameter to NONE.

E. By implementing the database resource Manager plan and setting the SEC_MAX_FAILED_LOGIN_ATTEMPTS parameters to 4.

Correct Answer: *A*
You can design your applications to automatically grant a role to the user who is trying to log in, provided the user meets criteria that you specify. To do so, you create a secure application role, which is a role that is associated with a PL/SQL procedure (or PL/SQL package that contains multiple procedures). The procedure validates the user: if the user fails the validation, then the user cannot log in. If the user passes the validation, then the procedure grants the user a role so that he or she can use the application. The user has this role only as long as he or she is logged in to the application. When the user logs out, the role is revoked.
Incorrect:
Not B: REMOTE_OS_AUTHENT specifies whether remote clients will be authenticated with the value of the OS_AUTHENT_PREFIX parameter.
Not C, not E: SEC_MAX_FAILED_LOGIN_ATTEMPTS specifies the number of authentication attempts that can be made by a client on a connection to the server process. After the specified number of failure attempts, the connection will be automatically dropped by the server process.
Not D: REMOTE_LOGIN_PASSWORDFILE specifies whether Oracle checks for a password file.
Values:
shared
One or more databases can use the password file. The password file can contain SYS as well as non-SYS users. exclusive
The password file can be used by only one database. The password file can contain SYS as well as non-SYS users. none
Oracle ignores any password file. Therefore, privileged users must be authenticated by the operating system.
Note:
The REMOTE_OS_AUTHENT parameter is deprecated. It is retained for backward compatibility only.

Community vote distribution

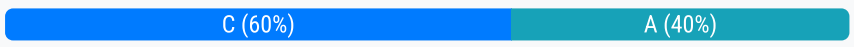

rcgenilo27 Highly Voted 4 years, 11 months ago
but C only satisfies 1 requirement
upvoted 8 times

vineetarneja 4 years, 3 months ago
Hi
And how A is satisfying all the requirements. The requirements mentioned in the question can only be meet through a profile which can be either assigned directly to the user or through a role but there is no mention like that in option A. so C seems to be the correct answer.
Regards
Vineet Arneja
vineetarneja1989@gmail.com
upvoted 3 times

khan Highly Voted 5 years, 3 months ago
Answer: C
upvoted 8 times

8111 Most Recent 10 months, 2 weeks ago
Selected Answer: A
ALTER ROLE role_name PROFILE 4_requirement;
upvoted 1 times

luizyto 1 year, 6 months ago

Selected Answer: A

Not C: SEC_MAX_FAILED_LOGIN_ATTEMPTS specifies the number of authentication attempts that can be made by a client on a connection to the server process. These login attempts can be for multiple user accounts in the same connection.
we need to apply on an user attemp not a client attemp.

upvoted 1 times

Ray520 2 years ago

A seems onnly viable option. The SEC_MAX_FAILED_LOGIN_ATTEMPTS is not a profile parameter option.

upvoted 1 times

JMAN1 2 years, 6 months ago

Selected Answer: C

Please refer to Question 2 in https://www.examtopics.com/exams/oracle/1z0-062/view/.
It is similar question and answer was assigning profile. So, I will go for C.

upvoted 3 times

jokerking 2 years, 9 months ago

I think C

upvoted 2 times

Modasser 3 years, 6 months ago

Answer should be A
Because: You can design your applications to automatically grant a role to the user who is trying to log in, provided the user meets criteria that you specify. To do so, you create a secure application role, which is a role that is associated with a PL/SQL procedure (or PL/SQL package that contains multiple procedures). The procedure validates the user: if the user fails the validation, then the user cannot log in. If the user passes the validation, then the procedure grants the user a role so that he orshe can use the application. The user has this role only as long as he or she is logged in to the application. When the user logs out, the role is revoked.
Incorrect:
Not B: REMOTE_OS_AUTHENT specifies whether remote clients will be authenticated with the value of the OS_AUTHENT_PREFIX parameter.
Not C, not E: SEC_MAX_FAILED_LOGIN_ATTEMPTS specifies the number of authentication attempts that can be made by a client on a connection to the server process. After the specified number of failure attempts, the connection will be automatically dropped by the server process.
Not D: REMOTE_LOGIN_PASSWORDFILE specifies whether Oracle checks for a password file.

upvoted 3 times

euler 4 years, 11 months ago

Correct answer is C, not A

upvoted 3 times

Question #88 *Topic 1*

A senior DBA asked you to execute the following command to improve performance:
SQL> ALTER TABLE subscribe log STORAGE (BUFFER_POOL recycle);
You checked the data in the SUBSCRIBE_LOG table and found that it is a large table containing one million rows.
What could be a reason for this recommendation?

A. The keep pool is not configured.

B. Automatic Workarea Management is not configured.

C. Automatic Shared Memory Management is not enabled.

D. The data blocks in the SUBSCRIBE_LOG table are rarely accessed.

E. All the queries on the SUBSCRIBE_LOG table are rewritten to a materialized view.

Correct Answer: *D*
The most of the rows in SUBSCRIBE_LOG table are accessed once a week.

vineetarneja Highly Voted 4 years, 3 months ago
Look at the explanation in the answer.. :D where is it written in the question that the table is access once per week. but he who provided the explanation knows everything. Looks God to me.. :D
upvoted 6 times

Yema Most Recent 3 years, 4 months ago
https://blog.toadworld.com/using_the_keep_and_recycle_pool

Any segments whose blocks tend to be accessed with less frequency should be assigned to the recycle pool so that it does not flush the other segments, either in the default cache or the keep pool. Segments that contain frequently accessed blocks should be assigned to the keep buffer pool so that the blocks of those segments will not be inadvertently removed, thus impacting performance.
upvoted 1 times

Question #89 *Topic 1*

Which three tasks can be automatically performed by the Automatic Data Optimization feature of Information lifecycle Management (ILM)?
(Choose three.)

A. Tracking the most recent read time for a table segment in a user tablespace

B. Tracking the most recent write time for a table segment in a user tablespace

C. Tracking insert time by row for table rows

D. Tracking the most recent write time for a table block

E. Tracking the most recent read time for a table segment in the SYSAUX tablespace

F. Tracking the most recent write time for a table segment in the SYSAUX tablespace

Correct Answer: *ABD*
Incorrect:
Not E, Not F When Heat Map is enabled, all accesses are tracked by the in-memory activity tracking module. Objects in the SYSTEM and SYSAUX tablespaces are not tracked.
* To implement your ILM strategy, you can use Heat Map in Oracle Database to track data access and modification.
Heat Map provides data access tracking at the segment-level and data modification tracking at the segment and row level.
* To implement your ILM strategy, you can use Heat Map in Oracle Database to track data access and modification. You can also use Automatic Data
Optimization (ADO) to automate the compression and movement of data between different tiers of storage within the database.
References:

Question #90 *Topic 1*

Which two partitioned table maintenance operations support asynchronous Global Index Maintenance in Oracle database 12c? (Choose two.)

A. ALTER TABLE SPLIT PARTITION

B. ALTER TABLE MERGE PARTITION

C. ALTER TABLE TRUNCATE PARTITION

D. ALTER TABLE ADD PARTITION

E. ALTER TABLE DROP PARTITION

F. ALTER TABLE MOVE PARTITION

Correct Answer: *CE*

Asynchronous Global Index Maintenance for DROP and TRUNCATE PARTITION

This feature enables global index maintenance to be delayed and decoupled from a DROP and TRUNCATE partition without making a global index unusable.

Enhancements include faster DROP and TRUNCATE partition operations and the ability to delay index maintenance to off-peak time.

References:

Question #91 *Topic 1*

You configure your database Instance to support shared server connections.
Which two memory areas that are part of PGA are stored in SGA instead, for shared server connection? (Choose two.)

A. User session data

B. Stack space

C. Private SQL area

D. Location of the runtime area for DML and DDL Statements

E. Location of a part of the runtime area for SELECT statements

Correct Answer: *AC*
A: PGA itself is subdivided. The UGA (User Global Area) contains session state information, including stuff like package-level variables, cursor state, etc. Note that, with shared server, the UGA is in the SGA. It has to be, because shared server means that the session state needs to be accessible to all server processes, as any one of them could be assigned a particular session. However, with dedicated server (which likely what you're using), the UGA is allocated in the PGA.
C: The Location of a private SQL area depends on the type of connection established for a session. If a session is connected through a dedicated server, private
SQL areas are located in the server process PGA. However, if a session is connected through a shared server, part of the private SQL area is kept in the SGA.
Note:
* System global area (SGA)
The SGA is a group of shared memory structures, known as SGA components, that contain data and control information for one Oracle Database instance. The
SGA is shared by all server and background processes. Examples of data stored in the SGA include cached data blocks and shared SQL areas.
* Program global area (PGA)
A PGA is a memory region that contains data and control information for a server process. It is nonshared memory created by Oracle Database when a server process is started. Access to the PGA is exclusive to the server process. There is one PGA for each server process. Background processes also allocate their own PGAs. The total memory used by all individual PGAs is known as the total instance PGA memory, and the collection of individual PGAs is referred to as the total instance PGA, or just instance PGA. You use database initialization parameters to set the size of the instance PGA, not individual PGAs.
References:

Question #92 *Topic 1*

Which two statements are true about Oracle Managed Files (OMF)? (Choose two.)

A. OMF cannot be used in a database that already has data files created with user-specified directions.

B. The file system directions that are specified by OMF parameters are created automatically.

C. OMF can be used with ASM disk groups, as well as with raw devices, for better file management.

D. OMF automatically creates unique file names for table spaces and control files.

E. OMF may affect the location of the redo log files and archived log files.

Correct Answer: *DE*
D: The database internally uses standard file system interfaces to create and delete files as needed for the following database structures:

Tablespaces -

Redo log files -

Control files -

Archived logs -

Block change tracking files -

Flashback logs -

RMAN backups -
Note:
* Using Oracle-managed files simplifies the administration of an Oracle Database. Oracle-managed files eliminate the need for you, the DBA, to directly manage the operating system files that make up an Oracle Database. With Oracle-managed files, you specify file system directories in which the database automatically creates, names, and manages files at the database object level. For example, you need only specify that you want to create a tablespace; you do not need to specify the name and path of the tablespace's datafile with the DATAFILE clause.
http://www.oracle-base.com/articles/9i/oracle-managed-files.php http://docs.oracle.com/cd/B10500_01/server.920/a96521/omf.htm
References:

Question #93 *Topic 1*

Which four actions are possible during an Online Data file Move operation? (Choose four.)

A. Creating and dropping tables in the data file being moved

B. Performing file shrink of the data file being moved

C. Querying tables in the data file being moved

D. Performing Block Media Recovery for a data block in the data file being moved

E. Flashing back the database

F. Executing DML statements on objects stored in the data file being moved

Correct Answer: *ACEF*
- You can now move On line Datafile without hove to stop Monoged Recovery and manually copy and rename Files. This can even be used to move Datafiles from or to ASM.
- New in Oracle Database 12c: FROM METAUNK. Physical Standby Database is in Active Data Guard Mode (opened READ ONLY and Managed Recovery is running):
It is now possible to online move a Datafile while Managed Recovery is running, ie. the Physical Standby Database is in Active Data Guard Mode. You con use this Command to move the Datafile
- A flashback operation does not relocate a moved data file to its previous location. If you move a data file online from one location to another and later flash back the database to a point in time before the move, then the Data file remains in the new location, but the contents of the Data file ore changed to the contents at the time specified in the flashback. Oracle0 Database Administrator's Guide 12c Release 1 (12.1)

Community vote distribution

ACDF (100%)

khan Highly Voted 5 years, 3 months ago

Answer : ACDF

upvoted 8 times

JMAN1 Most Recent 2 years, 6 months ago

Selected Answer: ACDF

Answer : ACDF

upvoted 1 times

charan94 3 years, 4 months ago

yes, Answer is ACDF . How can we flashback the database when the table is in online movement?

upvoted 1 times

TuxBingo 4 years ago

Right Answer is: ACDF
For who may think E is correct reference to Oracle Database 12c: New Features for Administrators Chapter 8 – Page# 50

An Online Move data file operation is not compatible when:
• The data file is an OFFLINE data file
• A concurrent flashback database operation is executing
• A media recovery is completing
• A file shrink operation or tablespace offline/drop operation involving the same file is performing

But it is compatible with:
• Block media recovery
• ALTER TABLESPACE READ ONLY or READ WRITE operations
• Data file extension operation
• Tablespace/database online backup mode involving the same file

upvoted 3 times

euler 4 years, 11 months ago

Answer : ACDF

upvoted 2 times

Question #94 *Topic 1*

Your multitenant container database (CDB) contains a pluggable database, HR_PDB. The default permanent tablespace in HR_PDB is USERDATA.
The container database (CDB) is open and you connect RMAN.
You want to issue the following RMAN command:
RMAN > BACKUP TABLESPACE hr_pdb:userdata;
Which task should you perform before issuing the command?

A. Place the root container in ARHCHIVELOG mode.

B. Take the user data tablespace offline.

C. Place the root container in the nomount stage.

D. Ensure that HR_PDB is open.

Correct Answer: *A*

Question #95 *Topic 1*

Identify three scenarios in which you would recommend the use of SQL Performance Analyzer to analyze impact on the performance of SQL statements.

A. Change in the Oracle Database version

B. Change in your network infrastructure

C. Change in the hardware configuration of the database server

D. Migration of database storage from non-ASM to ASM storage

E. Database and operating system upgrade

Correct Answer: *ACE*
Oracle 11g/12c makes further use of SQL tuning sets with the SQL Performance Analyzer, which compares the performance of the statements in a tuning set before and after a database change. The database change can be as major or minor as you like, such as:
* (E) Database, operating system, or hardware upgrades.
* (A,C) Database, operating system, or hardware configuration changes.
* Database initialization parameter changes.
* Schema changes, such as adding indexes or materialized views.
* Refreshing optimizer statistics.
* Creating or changing SQL profiles.

Question #96 *Topic 1*

Which two statements are true about the RMAN validate database command? (Choose two.)

A. It checks the database for intrablock corruptions.

B. It can detect corrupt pfiles.

C. It can detect corrupt spfiles.

D. It checks the database for interblock corruptions.

E. It can detect corrupt block change tracking files.

Correct Answer: *AC*
Block corruptions can be divided Into Interblock corruption and intrablock corruption. In intrablock corruption. th corruption occurs within the block itself and can be either physical or logical corruption. In interblock corruption, the corruption occurs between blocks and can only be logical corruption.
(key word) * The VALIDATE command checks for intrablock corruptions only. Only DBVERIFY and the ANALYZE statement detect Interblock corruption.
VALIDATE Command Output > List of Control File and SPFILE.
File TYPE > SPFILE or Control File.
Status > OK if no corruption, or FAILED If block corruption is found.
Blocks Failing The number of blocks that fail the corruption check. These blocks are newly corrupt.
Blocks Examined Total number of blocks in the file.
Oracle' Database Backup and Recovery User's Guide
12c Release 1 (12.1) - 16 Validating Database Files and Backups

Question #97 *Topic 1*

You install a non-RAC Oracle Database. During Installation, the Oracle Universal Installer (OUI) prompts you to enter the path of the Inventory directory and also to specify an operating system group name.
Which statement is true?

A. The ORACLE_BASE base parameter is not set.

B. The installation is being performed by the root user.

C. The operating system group that is specified should have the root user as its member.

D. The operating system group that is specified must have permission to write to the inventory directory.

Correct Answer: *D*
Note:

Providing a UNIX Group Name -
If you are installing a product on a UNIX system, the Installer will also prompt you to provide the name of the group which should own the base directory.
You must choose a UNIX group name which will have permissions to update, install, and deinstall Oracle software. Members of this group must have write permissions to the base directory chosen.
Only users who belong to this group are able to install or deinstall software on this machine.

Question #98 *Topic 1*

You are required to migrate your 11.2.0.3 database as a pluggable database (PDB) to a multitenant container database (CDB).
The following are the possible steps to accomplish this task:
1. Place all the user-defined tablespace in read-only mode on the source database.
2. Upgrade the source database to a 12c version.
3. Create a new PDB in the target container database.
4. Perform a full transportable export on the source database with the VERSION parameter set to 12 using the expdp utility.
5. Copy the associated data files and export the dump file to the desired location in the target database.
6. Invoke the Data Pump import utility on the new PDB database as a user with the DATAPUMP_IMP_FULL_DATABASE role and specify the full transportable import options.
7. Synchronize the PDB on the target container database by using the DBMS_PDS.SYNC_ODB function.
Identify the correct order of the required steps.

A. 2, 1, 3, 4, 5, 6

B. 1, 3, 4, 5, 6, 7

C. 1, 4, 3, 5, 6, 7

D. 2, 1, 3, 4, 5, 6, 7

E. 1, 5, 6, 4, 3, 2

Correct Answer: *C*
1. Set user tablespaces in the source database to READ ONLY.
2. From the Oracle Database 11g Release 2 {11.2.0.3) environment, export the metadata and any data residing in administrative tablespaces from the source database using the FULL=Y and TRANSPORTABLE=ALWAYS parameters. Note that the VER$ION=12 parameter is required only when exporting from an
Oracle Database llg Release 2 database:
3. Copy the tablespace data files from the source system to the destination system. Note that the log file from the export operation will list the data files required to be moved.
4. Create a COB on the destination system, including a PDB into which you will import the source database.
5. In the Oracle Database 12c environment, connect to the pre-created PDB and import the dump file. The act of importing the dump file will plug the tablespace data files into the destination PDB
Oracle White Paper - Upgrading to Oracle Database 12c -August 2013

TuxBingo 4 years ago
The correct aswer is C, because they area asking the correct order -- This is the trick
upvoted 3 times

fsamkoh 4 years, 2 months ago
B and C are true but C is the correct order
upvoted 1 times

NayruLove 4 years, 10 months ago
How can it be? If C is true than B must be true, too.
upvoted 1 times

5kong 4 years, 5 months ago
is B the correct order,too?
upvoted 1 times

Question #99 *Topic 1*

In your multitenant container database (CDB) with two pluggable database (PDBs). You want to create a new PDB by using SQL Developer.
Which statement is true?

A. The CDB must be open.

B. The CDB must be in the mount stage.

C. The CDB must be in the nomount stage.

D. Alt existing PDBs must be closed.

Correct Answer: *A*
* Creating a PDB
Rather than constructing the data dictionary tables that define an empty PDB from scratch, and then populating its Obj$ and Dependency$ tables, the empty PDB is created when the CDB is created. (Here, we use empty to mean containing no customer-created artifacts.) It is referred to as the seed PDB and has the name
PDB$Seed. Every CDB non-negotiably contains a seed PDB; it is non-negotiably always open in read-only mode. This has no conceptual significance; rather, it is just an optimization device. The create PDB operation is implemented as a special case of the clone PDB operation.
The size of the seed PDB is only about 1 gigabyte and it takes only a few seconds on a typical machine to copy it.

Community vote distribution
A (100%)

jackymak 2 years, 5 months ago

Selected Answer: A
Tested that the PDB cannot be created in mount mode.
upvoted 1 times

Question #100 *Topic 1*

Which two statements are true about the Oracle Direct Network File system (DNFS)? (Choose two.)

A. It utilizes the OS file system cache.

B. A traditional NFS mount is not required when using Direct NFS.

C. Oracle Disk Manager can manage NFS on its own, without using the operating kernel NFS driver.

D. Direct NFS is available only in UNIX platforms.

E. Direct NFS can load-balance I/O traffic across multiple network adapters.

Correct Answer: *CE*

E: Performance is improved by load balancing across multiple network interfaces (if available).

Note:

* To enable Direct NFS Client, you must replace the standard Oracle Disk Manager (ODM) library with one that supports Direct NFS Client.

Incorrect:

Not A: Direct NFS Client is capable of performing concurrent direct I/O, which bypasses any operating system level caches and eliminates any operating system write-ordering locks

Not B:

* To use Direct NFS Client, the NFS file systems must first be mounted and available over regular NFS mounts.
* Oracle Direct NFS (dNFS) is an optimized NFS (Network File System) client that provides faster and more scalable access to NFS storage located on NAS storage devices (accessible over TCP/IP).

Not D: Direct NFS is provided as part of the database kernel, and is thus available on all supported database platforms - even those that don't support NFS natively, like Windows.

Note:

* Oracle Direct NFS (dNFS) is an optimized NFS (Network File System) client that provides faster and more scalable access to NFS storage located on NAS storage devices (accessible over TCP/IP). Direct NFS is built directly into the database kernel - just like ASM which is mainly used when using DAS or SAN storage.
* Oracle Direct NFS (dNFS) is an internal I/O layer that provides faster access to large NFS files than traditional NFS clients.

Question #101 *Topic 1*

Examine the parameters for your database instance:

```
NAME                                 TYPE        VALUE
------------------------------------ ----------- ---------
optimizer_adaptive_reporting_only    boolean     FALSE
optimizer_capture_sql_plan_baselines boolean     FALSE
optimizer_dynamic_sampling           integer     2
optimizer_features_enable            string      12.1.0.1
```

Which three statements are true about the process of automatic optimization by using cardinality feedback? (Choose three.)

A. The optimizer automatically changes a plan during subsequent execution of a SQL statement if there is a huge difference in optimizer estimates and execution statistics.

B. The optimizer can re optimize a query only once using cardinality feedback.

C. The optimizer enables monitoring for cardinality feedback after the first execution of a query.

D. The optimizer does not monitor cardinality feedback if dynamic sampling and multicolumn statistics are enabled.

E. After the optimizer identifies a query as a re-optimization candidate, statistics collected by the collectors are submitted to the optimizer.

Correct Answer: *ACD*

C: During the first execution of a SQL statement, an execution plan is generated as usual.

D: if multi-column statistics are not present for the relevant combination of columns, the optimizer can fall back on cardinality feedback.

(not B)* Cardinality feedback. This feature, enabled by default in 11.2, is intended to improve plans for repeated executions.

optimizer_dynamic_sampling optimizer_features_enable

* dynamic sampling or multi-column statistics allow the optimizer to more accurately estimate selectivity of conjunctive predicates.

Note:

* OPTIMIZER_DYNAMIC_SAMPLING controls the level of dynamic sampling performed by the optimizer.

Range of values. 0 to 10 -

* Cardinality feedback was introduced in Oracle Database 11gR2. The purpose of this feature is to automatically improve plans for queries that are executed repeatedly, for which the optimizer does not estimate cardinalities in the plan properly. The optimizer may misestimate cardinalities for a variety of reasons, such as missing or inaccurate statistics, or complex predicates. Whatever the reason for the misestimate, cardinality feedback may be able to help.

Community vote distribution

ACE (67%) ACD (33%)

luizyto 1 year, 6 months ago

Selected Answer: ACD

In some cases, there are other techniques available to improve estimation; for instance, dynamic sampling or multi-column statistics allow the optimizer to more accurately estimate selectivity of conjunctive predicates. In cases where these techniques apply, statistics feedback is not enabled.

upvoted 1 times

luizyto 1 year, 6 months ago

Selected Answer: ACE

In some cases, there are other techniques available to improve estimation; for instance, dynamic sampling or multi-column statistics allow the optimizer to more accurately estimate selectivity of conjunctive predicates. In cases where these techniques apply, statistics feedback is not enabled.

upvoted 1 times

Ray520 2 years, 3 months ago

D is correct: cardinality feedback is called Statistics feedback in 12c. "In some cases, there are other techniques available to improve estimation; fo instance, dynamic sampling or multi-column statistics allow the optimizer to more accurately estimate selectivity of conjunctive predicates. In cas where these techniques apply, statistics feedback is not enabled. However, if multi-column statistics are not present for the relevant combination columns, the optimizer can fall back on statistics feedback."
Reference: https://blogs.oracle.com/optimizer/post/statistics-feedback-formerly-cardinality-feedback

upvoted 3 times

JMAN1 2 years, 6 months ago

Selected Answer: ACE

So given answer D is wrong. E is correct.
A C E.
Any comment please.

upvoted 1 times

JMAN1 2 years, 6 months ago

I Think A C E.
CF(cardinality feedback) is helping Dynamic sampling. In order to disable CF. You need to set _optimizer_use_feedback" to false at a session or statement level.
So given answer D is wrong. E is correct.

upvoted 1 times

Question #102 *Topic 1*

Which three statements are true when the listener handles connection requests to an Oracle 12c database instance with multithreaded architecture enabled In
UNIX? (Choose three.)

A. Thread creation must be routed through a dispatcher process

B. The local listener may spawn a now process and have that new process create a thread

C. Each Oracle process runs an SCMN thread.

D. Each multithreaded Oracle process has an SCMN thread.

E. The local listener may pass the request to an existing process which in turn will create a thread.

Correct Answer: *ADE*

Ray520 2 years, 3 months ago

Reference:
https://docs.oracle.com/cd/B10500_01/server.920/a96544/multithr.htm

In the architecture for multithreaded agents, each incoming connection request is processed by means of the three different kinds of threads:

A single monitor thread (SCMN thread) (option D)

The monitor thread is responsible for the following:

-Maintaining communication with the listener
-Monitoring the load on the process
-Starting and stopping threads when required
-Several dispatcher threads

The dispatcher threads are responsible for the following:

-Handling communication with the Oracle server (option A)
-Passing task requests on to the task threads

Several task threads

-The task threads handle requests from the Oracle processes.

When a SCMN thread registers the dispatcher threads it has created with all the listeners that are handling connections to this agent.
While the dispatcher for this SID is running, the listener does not start a new process when it gets an incoming connection. Instead, the listener hands over the connection to this same dispatcher. (option E)

upvoted 1 times

Question #103 *Topic 1*

Which three operations can be performed as multipartition operations in Oracle? (Choose three.)

A. Merge partitions of a list partitioned table

B. Drop partitions of a list partitioned table

C. Coalesce partitions of a hash-partitioned global index.

D. Move partitions of a range-partitioned table

E. Rename partitions of a range partitioned table

F. Merge partitions of a reference partitioned index

Correct Answer: *ABF*

Multipartition maintenance enables adding, dropping, truncate, merge, split operations on multiple partitions.

A: Merge Multiple Partitions:

The new "ALTER TABLE … MERGE PARTITIONS " help merge multiple partitions or subpartitions with a single statement. When merging multiple partitions, local and global index operations and semantics for inheritance of unspecified physical attributes are the same for merging two partitions.

B: Drop Multiple Partitions:

The new "ALTER TABLE … DROP PARTITIONS " help drop multiple partitions or subpartitions with a single statement.

Example:

view plaincopy to clipboardprint?

SQL> ALTER TABLE Tab_tst1 DROP PARTITIONS

Tab_tst1_PART5, Tab_tst1_PART6, Tab_tst1_PART7;

Table altered -

SQL>

Restrictions :

- You cant drop all partitions of the table.
- If the table has a single partition, you will get the error: ORA-14083: cannot drop the only partition of a partitioned.

Question #104 *Topic 1*

You are connected using SQL* Plus to a multitenant container database (CDB) with SYSDBA privileges and execute the following sequence statements:

```
SQL> CREATE PLUGGABLE DATABASE NEW_PDB ADMIN USER PDB_ADMIN IDENTIFIED BY SECRET;
Pluggable database created.

SQL> ALTER PLUGGABLE DATABASE NEW_PDB OPEN;
Pluggable database altered.
SQL> ALTER SESSION SET CONTAINER = NEW_PDB;
Session altered.
SQL> GRANT CONNECT TO PDB_ADMIN;
Grant succeeded.
SQL  CONNECT PDB_ADMIN/SECRET@LOCALHOST/NEW_PDB
Connected.
SQL> SELECT * FROM SESSION_PRIVS;

PRIVILEGE
----------------------------------------
CREATE SESSION
SET CONTAINER

SQL> ALTER SESSION SET CONTAINER = PDB$SEED;
```

What is the result of the last SET CONTAINER statement and why is it so?

A. It succeeds because the PDB_ADMIN user has the required privileges.

B. It fails because common users are unable to use the SET CONTAINER statement.

C. It fails because local users are unable to use the SET CONTAINER statement.

D. If fails because the SET CONTAINER statement cannot be used with PDB$SEED as the target pluggable database (PDB).

Correct Answer: *C*

Community vote distribution

C (100%)

suppa123 Highly Voted 4 years ago

Answer is C:

SQL> create pluggable database cdb1pdb3 admin user cdb1pdb3_admin identified by password file_name_convert= ('/oradata/CDB1/pdbseed','/oradata/CDB1/CDB1PDB3') ;

Pluggable database created.

SQL> alter pluggable database cdb1pdb3 open;

Pluggable database altered.

SQL> alter session set container=cdb1pdb3
2 /

Session altered.

SQL> grant connect to cdb1pdb3_admin;

Grant succeeded.

SQL> connect cdb1pdb3_admin/password@nonrac12c-oel6u5/cdb1pdb3
Connected.
SQL> select * from session_privs;

PRIVILEGE
———————————————-
SET CONTAINER
CREATE PLUGGABLE DATABASE
CREATE SESSION

SQL> alter session set container=PDB$SEED;
ERROR:
ORA-01031: insufficient privileges

SQL>

upvoted 5 times

jackymak Most Recent 2 years, 5 months ago

Selected Answer: C

Local user shouldn't able to access another PDBs or CDB

upvoted 2 times

NorthKorean 4 years, 4 months ago

I think the answer is D PDB$SEED is not alterable

upvoted 2 times

joe_doe 4 years, 2 months ago

But first,let me take .....

My apologize,but first,it'll check your authority to sure weather you have set container privileges

upvoted 1 times

joe_doe 4 years, 2 months ago

whether,it's whether

upvoted 1 times

Question #105 *Topic 1*

Examine the details of the Top 5 Timed Events in the following Automatic Workloads Repository (AWR) report:

Top 5 Timed Foreground Events

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
DB CPU		67		98.21	
db file sequential read	8,371	0	0	0.52	User I/O
latch row cache objects	16	0	8	0.19	Concurrency
latch shared pool	956	0	0	0.15	Concurrency
log file sync	25	0	2	0.06	Commit

What are three possible causes for the latch-related wait events?

A. The size of the shared pool is too small.

B. Cursors are not being shared.

C. A large number COMMITS are being performed.

D. There are frequent logons and logoffs.

E. The buffers are being read into the buffer cache, but some other session is changing the buffers.

Correct Answer: *ABD*

Community vote distribution

AE (100%)

Ray520 2 years, 3 months ago

reference: https://docs.oracle.com/database/121/TGDBA/pfgrf_instance_tune.htm#TGDBA94516

Possible causes of shared pool latch events:
Lack of statement reuse
Statements not using bind variables
Insufficient size of application cursor cache
Cursors closed explicitly after each execution
Frequent logins and logoffs
Underlying object structure being modified (for example truncate)
Shared pool too small

upvoted 1 times

JMAN1 2 years, 6 months ago

Selected Answer: AE

Table 10-3 Latch Wait Events
Possible Causes
• Lack of statement reuse
• Statements not using bind variables
• Insufficient size of application cursor cache
• Cursors closed explicitly after each execution
• (D) Frequent logins and logoffs
• (E) Underlying object structure being modified
• (A) Shared pool too small

So A, D, E are correct
B is possibly wrong

upvoted 3 times

Question #106 *Topic 1*

You enabled an audit policy by issuing the following statements:
SQL> AUDIT POLICY ORA_DATABASE_PARAMETER BY SCOTT;
SQL> AUDIT POLICY ORA_DATABASE_PARAMETER BY SYS, SYSTEM;
For which database users and for which executions is the audit policy now active? Select two.

A. SYS, SYSTEM

B. SCOTT

C. Only for successful executions

D. Only for failed executions

E. Both successful and failed executions

Correct Answer: *AE*
* The ORA_DATABASE_PARAMETER policy audits commonly used Oracle Database parameter settings. By default, this policy is not enabled.

DANG257315 Highly Voted 3 years, 7 months ago
The correct answer are A and E.
b is not correct bacause by default the SCOTT user is disable
upvoted 7 times

Ray_gk 2 years, 2 months ago
I agree with you. I would argue that E is correct based on an excerpt from the documentation for 12.1.

"Specify WHENEVER SUCCESSFUL to audit only SQL statements and operations that succeed.

Specify WHENEVER NOT SUCCESSFUL to audit only SQL statements and operations that fail or result in errors.

If you omit this clause, then Oracle Database performs the audit regardless of success or failure." --> E

I agree that user scott is disabled in 12.1 DB upon creating database as well so I would rule out B. The most appropriate answers seem to be A and E.
upvoted 1 times

yosiw96816 Highly Voted 3 years, 8 months ago
Agree AB
https://docs.oracle.com/database/121/SQLRF/statements_4008.htm#SQLRF56110
If multiple AUDIT ... BY ... statements are specified for the same unified audit policy, then the policy is enabled for the union of the users specified each statement.
upvoted 5 times

NorwayOracle 2 years, 4 months ago
Agree.
upvoted 1 times

dancymonkey Most Recent 3 years, 11 months ago
AB are correct
Both commands enable predefined policy
Note: You need to open a new session to see that 3 users get audit by this policy.

AUDIT POLICY ORA_DATABASE_PARAMETER BY C##AX;
AUDIT POLICY ORA_DATABASE_PARAMETER BY SYS, SYSTEM;
select * from audit_unified_enabled_policies;
select * from AUDIT_UNIFIED_POLICIES where policy_name = 'ORA_DATABASE_PARAMETER';

ALTER SYSTEM set DEFERRED_SEGMENT_CREATION = TRUE;
exec dbms_audit_mgmt.flush_unified_audit_trail;
select dbusername,sql_text, unified_audit_policies, event_timestamp, to_char(sysdate, 'HH MI SS AM') from unified_audit_trail order by event_timestamp desc
upvoted 2 times

Question #107 *Topic 1*

A redaction policy was added to the SAL column of the SCOTT.EMP table:

```
BEGIN
  DBMS_REDACT.ADD_POLICY(
    OBJECT_SCHEMA => 'SCOTT',
    OBJECT_NAME   => 'EMP',
    POLICY_NAME   => 'SCOTT_EMP',
    COLUMN_NAME   => 'SAL',
    EXPRESSION    => 'SYS_CONTEXT("SYS_SESSION_ROLES", "MGR") = "FALSE"');
END;
/
```

All users have their default set of system privileges.

For which three situations will data not be redacted? (Choose three.)

A. SYS sessions, regardless of the roles that are set in the session

B. SYSTEM sessions, regardless of the roles that are set in the session

C. SCOTT sessions, only if the MGR role is set in the session

D. SCOTT sessions, only if the MGR role is granted to SCOTT

E. SCOTT sessions, because he is the owner of the table

F. SYSTEM session, only if the MGR role is set in the session

Correct Answer: *ABD*

Ray520 2 years, 3 months ago

SYS and SYSTEM users can always bypass any existing Oracle Data Redaction policies, and will always be able to view data from tables (or views) that have Data Redaction policies defined on them.

A role cannot be granted to session unless the user that created the session has that role assigned to it. So, option D is correct.

reference: https://docs.oracle.com/database/121/ASOAG/security-considerations-for-using-oracle-data-redaction.htm#ASOAG10536

upvoted 1 times

lollo1234 3 years, 2 months ago

I think ABC:

The following example queries the SESSION_ROLES data dictionary view to show that RESOURCE is the only role currently enabled for the session
It then uses the SYS_CONTEXT function to show that the RESOURCE role is currently enabled for the session and the DBA role is not.
Ref: https://docs.oracle.com/database/121/SQLRF/functions199.htm#SQLRF06117

upvoted 2 times

jackymak 2 years, 3 months ago

SYS_SESSION_ROLES - Indicates whether a specified role is currently "enabled" for the session. This namespace is available starting with Oracle Database 11g Release 2 (11.2.0.4).

upvoted 1 times

melat 4 years ago

ABD:
https://docs.oracle.com/database/121/ASOAG/security-considerations-for-using-oracle-data-redaction.htm#ASOAG10536

upvoted 4 times

jackymak 2 years, 5 months ago

If so, I will prefer ABC. For C, the session is set to MGR that mean the user already granted the MGR role. But D, the user has only granted the role, which can owned and not using.

upvoted 1 times

Question #108 *Topic 1*

What is the result of executing a TRUNCATE TABLE command on a table that has Flashback Archiving enabled?

A. It fails with the ORA-665610 Invalid DDL statement on history-tracked message

B. The rows in the table are truncated without being archived.

C. The rows in the table are archived, and then truncated.

D. The rows in both the table and the archive are truncated.

Correct Answer: *C*

PaoloBL Highly Voted 2 years, 11 months ago

C is correct on 12.1

upvoted 5 times

JMAN1 2 years, 6 months ago

You are lying..

upvoted 1 times

Oracle2020 Highly Voted 2 years, 11 months ago

Answer: B

Explanation: You cannot roll back a TRUNCATE TABLE statement, nor can you use a FLASHBACK TABLE statement to retrieve the contents of a tab that has been truncated.

upvoted 5 times

Ray520 Most Recent 2 years, 3 months ago

The rows in the table are archived, and then truncated when FLASHBACK ARCHIVE is enabled. So if we want to get databack after a truncate, we cannot flashback the table though as TUNCATE is a DDL, but we can select the data and insert into a different table.:

insert into table1
select * from table1 as of timestamp to_timestamp(<your_time>);

upvoted 2 times

wangle0088 3 years, 1 month ago

answer is B

upvoted 4 times

Question #109 *Topic 1*

Which three activities are supported by the Data Recovery Advisor? (Choose three.)

A. Advising on block checksum failures

B. Advising on inaccessible control files

C. Advising on inaccessible block change tracking files

D. Advising on empty password files

E. Advising on invalid block header field values

Correct Answer: *ABE*
* Data Recovery Advisor can diagnose failures such as the following:
/ (B) Components such as datafiles and control files that are not accessible because they do not exist, do not have the correct access permissions, have been taken offline, and so on
/ (A, E) Physical corruptions such as block checksum failures and invalid block header field values
/ Inconsistencies such as a datafile that is older than other database files
/ I/O failures such as hardware errors, operating system driver failures, and exceeding operating system resource limits (for example, the number of open files)
* The Data Recovery Advisor automatically diagnoses corruption or loss of persistent data on disk, determines the appropriate repair options, and executes repairs at the user's request. This reduces the complexity of recovery process, thereby reducing the Mean Time To Recover (MTTR).

Question #110 *Topic 1*

You create a table with the PERIOD FOR clause to enable the use of the Temporal Validity feature of Oracle Database 12c.
Examine the table definition:

```
create table employees
(empno number, salary number,
deptid number, name varchar2(100),
period for employee_time);
```

Which three statements are true concerning the use of the Valid Time Temporal feature for the EMPLOYEES table? (Choose three.)

A. The valid time columns employee_time_start and employee_time_end are automatically created.

B. The same statement may filter on both transaction time and valid temporal time by using the AS OF TIMESTAMP and PERIOD FOR clauses.

C. The valid time columns are not populated by the Oracle Server automatically.

D. The valid time columns are visible by default when the table is described.

E. Setting the session valid time using DBMS_FLASHBACK_ARCHIVE.ENABLE_AT_VALID_TIME sets the visibility for data manipulation language (DML), data

Correct Answer: *ABC*

Community vote distribution

ABE (100%)

DANG257315 Highly Voted 3 years, 7 months ago

Answer: A,B, E
A: To implement Temporal Validity(TV), 12c offers the option to have two date columns in that table which is having TV enabled using the new clause Period For in the Create Table for the newly created tables or in the Alter Table for the existing ones. The columns that are used can be defined while creating the table itself and will be used in the Period For clause or you can skip having them in the table's definition in the case of which, the Period For clause would be creating them internally.
E: ENABLE_AT_VALID_TIME Procedure
This procedure enables session level valid time flashback.

upvoted 7 times

luizyto Most Recent 1 year, 9 months ago

Selected Answer: ABE

This procedure enables session level valid time flashback.

upvoted 1 times

jackymak 2 years, 5 months ago

Selected Answer: ABE

ABE they are right.

upvoted 2 times

JMAN1 2 years, 6 months ago

Official dump book says answer is ABE.

upvoted 1 times

Aleksandar 2 years, 10 months ago

The answer is A,B,E

upvoted 3 times

Taria24 3 years, 3 months ago

ans should be abc

upvoted 2 times

Question #111 *Topic 1*

Which three statements are true regarding the use of the Database Migration Assistant for Unicode (DMU)? (Choose three.)

A. A DBA can check specific tables with the DMU

B. The database to be migrated must be opened read-only.

C. The release of the database to be converted can be any release since 9.2.0.8.

D. The DMU can report columns that are too long in the converted characterset.

E. The DMU can report columns that are not represented in the converted characterset.

Correct Answer: *ADE*

A: In certain situations, you may want to exclude selected columns or tables from scanning or conversion steps of the migration process.

D: Exceed column limit -
The cell data will not fit into a column after conversion.

E: Need conversion -
The cell data needs to be converted, because its binary representation in the target character set is different than the representation in the current character set, but neither length limit issues nor invalid representation issues have been found.

* Oracle Database Migration Assistant for Unicode (DMU) is a unique next-generation migration tool providing an end-to-end solution for migrating your databases from legacy encodings to Unicode.

Incorrect:

Not C: The release of Oracle Database must be 10.2.0.4, 10.2.0.5, 11.1.0.7, 11.2.0.1, or later.

Question #112 *Topic 1*

Oracle Grid Infrastructure for a stand-alone server is installed on your production host before installing the Oracle Database server. The database and listener are configured by using Oracle Restart.
Examine the following command and its output:
$ crsctl config has
CRS-4622: Oracle High Availability Services auto start is enabled.
What does this imply?

A. When you start an instance on a high with SQL *Plus dependent listeners and ASM disk groups are automatically started.

B. When a database instance is started by using the SRVCTL utility and listener startup fails, the instance is still started.

C. When a database is created by using SQL* Plus, it is automatically added to the Oracle Restart configuration.

D. When you create a database service by modifying the SERVICE_NAMES initialization parameter, it is automatically added to the Oracle Restart configuration.

Correct Answer: *B*
About Startup Dependencies -
Oracle Restart ensures that Oracle components are started in the proper order, in accordance with component dependencies. For example, if database files are stored in Oracle ASM disk groups, then before starting the database instance, Oracle Restart ensures that the Oracle ASM instance is started and the required disk groups are mounted. Likewise, if a component must be shut down, Oracle Restart ensures that dependent components are cleanly shut down first.
Oracle Restart also manages the weak dependency between database instances and the Oracle Net listener (the listener): When a database instance is started,
Oracle Restart attempts to start the listener. If the listener startup fails, then the database is still started. If the listener later fails, Oracle Restart does not shut down and restart any database instances.
http://docs.oracle.com/cd/E16655_01/server.121/e17636/restart.htm#ADMIN12710

Ray520 2 years, 3 months ago
Reference: https://docs.oracle.com/cd/E18283_01/server.112/e17120/restart001.htm

Oracle Restart also manages the weak dependency between database instances and the Oracle Net listener (the listener): When a database instance is started, Oracle Restart attempts to start the listener. If the listener startup fails, then the database is still started. If the listener later fails Oracle Restart does not shut down and restart any database instances.
upvoted 1 times

Question #113 *Topic 1*

Your multitenant container database (CDB) contains some pluggable databases (PDBs), you execute the following command in the root container:

```
SQL> CREATE USER c##a_admin
     IDENTIFIED BY password
     DEFAULT TABLESPACE data_ts
     QUOTA 100M ON test_ts
     QUOTA 500K ON data_ts
     TEMPORARY TABLESPACE temp_ts
     PROFILE hr_profile;
```

Which two statements are true? (Choose two.)

A. Schema objects owned by the C# # A_ADMIN common user can be shared across all PDBs.

B. The C # # A_ADMIN user will be able to use the TEMP_TS temporary tablespace only in root.

C. The command will, create a common user whose description is contained in the root and each PDB.

D. The schema for the common user C # # A_ADMIN can be different in each container.

E. The command will create a user in the root container only because the container clause is not used.

Correct Answer: *CD*

bigswole 2 years, 9 months ago

https://docs.oracle.com/database/121/ADMQS/GUID-DA54EBE5-43EF-4B09-B8CC-FAABA335FBB8.htm

upvoted 2 times

jackymak 2 years, 4 months ago

For C,
A common user is a database user that has the same identity in the root and in every existing and future pluggable database (PDB). Every common user can connect to and perform operations within the root, and within any PDB in which it has privileges.
---
https://docs.oracle.com/database/121/ADMQS/GUID-DA54EBE5-43EF-4B09-B8CC-FAABA335FBB8.htm

upvoted 1 times

Question #114 *Topic 1*

You performed an incremental level 0 backup of a database:
RMAN > BACKUP INCREMENTAL LEVEL 0 DATABASE;
To enable block change tracking after the incremental level 0 backup, you issued this command:
SQL > ALTER DATABASE ENABLE BLOCK CHANGE TRACKING USING FILE
/mydir/rman_change_track.f;
To perform an incremental level 1 cumulative backup, you issued this command:
RMAN> BACKUP INCREMENTAL LEVEL 1 CUMULATIVE DATABASE;
Which three statements are true? (Choose three.)

A. Backup change tracking will sometimes reduce I/O performed during cumulative incremental backups.

B. The change tracking file must always be backed up when you perform a full database backup.

C. Block change tracking will always reduce I/O performed during cumulative incremental backups.

D. More than one database block may be read by an incremental backup for a change made to a single block.

E. The incremental level 1 backup that immediately follows the enabling of block change tracking will not read the change tracking file to discover changed

Correct Answer: *ADE*

Ray520 2 years, 3 months ago

reference:
https://docs.oracle.com/cd/B19306_01/backup.102/b14192/bkup004.htm#:~:text=A%20differential%20backup%2C%20which%20backs,incremen
%20backup%20at%20level%200

"Each data block in a datafile contains a system change number (SCN), which is the SCN at which the most recent change was made to the block. During an incremental backup, RMAN reads the SCN of each data block in the input file and compares it to the checkpoint SCN of the parent incremental backup. If the SCN in the input data block is greater than or equal to the checkpoint SCN of the parent, then RMAN copies the block.

"Block changing tracking improves the performance of incremental backups by recording changed blocks in the block change tracking file. Durin an incremental backup, instead of scanning all data blocks to identify which blocks have changed, RMAN uses this file to identify the changed blocks that need to be backed up."

upvoted 1 times

Question #115 *Topic 1*

You find this query being used in your Oracle 12c database:

```
select employee_id, first_name, salary
from hr.employees
order by employee_id
fetch first 20 percent rows only;
```

Which method a used by the optimizer to limit the rows being returned?

A. A filter is added to the table query dynamically using ROWNUM to limit the rows to 20 percent of the total rows

B. All the rows are returned to the client or middle tier but only the first 20 percent are returned to the screen or the application.

C. A view is created during execution and a filter on the view limits the rows to 20 percent of the total rows.

D. A TOP-N query is created to limit the rows to 20 percent of the total rows

Correct Answer: *C*

TuxBingo 4 years ago

I agree with the answer C:

SQL> select * from test fetch first 20 rows only;
Execution Plan
———————————————————————————————————-
| Id | Operation | Name | Rows | Bytes | Cost (%CPU)| Time |
———————————————————————————————————-
0	SELECT STATEMENT		20	880	3 (0)	00:00:01
* 1	VIEW		20	880	3 (0)	00:00:01
* 2	WINDOW NOSORT STOPKEY		1	31	3 (0)	00:00:01
3	TABLE ACCESS FULL	TEST	1	31	3 (0)	00:00:01

upvoted 4 times

TuxBingo 4 years ago

As you see in the explain plan, the view is incorporated.

upvoted 3 times

Question #116 *Topic 1*

Which three resources might be prioritized between competing pluggable databases when creating a multitenant container database plan (CDB plan) using Oracle
Database Resource Manager? (Choose three.)

A. Maximum Undo per consumer group

B. Maximum Idle time

C. Parallel server limit

D. CPU

E. Exadata I/O

F. Local file system I/O

Correct Answer: *CDE*

Community vote distribution

CDE (100%)

JMAN1 2 years, 6 months ago

Selected Answer: CDE

http://docs.oracle.com/database/121/ADMIN/dbrm.htm#ADMIN11852
27.3.1 CPU
27.3.2 EXADATA I/O
27.3.3 Parallel Execution Servers

upvoted 2 times

Question #117 *Topic 1*

You created an encrypted tablespace:

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

You then closed the encryption wallet because you were advised that this is secure.
Later in the day, you attempt to create the EMPLOYEES table in the SECURESPACE tablespace with the SALT option on the EMPLOYEE column.
Which is true about the result?

A. It creates the table successfully but does not encrypt any inserted data in the EMPNAME column because the wallet must be opened to encrypt columns with SALT.

B. It generates an error when creating the table because the wallet is closed.

C. It creates the table successfully, and encrypts any inserted data in the EMPNAME column because the wallet needs to be open only for tablespace creation.

D. It generates error when creating the table, because the salt option cannot be used with encrypted tablespaces.

Correct Answer: *B*

Ray520 2 years, 3 months ago

https://docs.oracle.com/cd/E11882_01/network.112/e40393/asotrans.htm#ASOAG10143

Closing the wallet disables all encryption and decryption operations. Any attempt to encrypt/decrypt data or access encrypted data results in the following error:

ORA-28365: wallet is not open

Note: By default TDE adds salt to cleartext before encrypting it. Adding salt means it adds some random strings to the data before it's encrypted, strengthen security.

upvoted 1 times

Question #118 *Topic 1*

On your Oracle Database, you issue the following commands to create indexes:
SQL > CREATE INDEX oe.ord_customer_ix1 ON oe.orders (customer_id, sales_rep_id) INVISIBLE;
SQL> CREATE BITMAP INDEX oe.ord_customer_ix2 ON oe.orders (customer_id, sales_rep_id);
Which two statements are true? (Choose two.)

A. Only the ORD_CUSTOMER_IX1 index created.

B. Both the indexes are updated when a row is inserted, updated, or deleted in the ORDERS table.

C. Both the indexes are created: however, only ORD_CUSTOMERS_IX1 is used by the optimizer for queries on the ORDERS table.

D. The ORD_CUSTOMER_IX1 index is not used by the optimizer even when the OPTIMIZER_USE_INVISIBLE_INDEXES parameters is set to true.

E. Both the indexes are created and used by the optimizer for queries on the ORDERS table.

F. Both the indexes are created: however, only ORD_CUSTOMERS_IX2 is used by the optimizer for queries on the ORDERS table.

Correct Answer: *BF*
Not A: Both indexes are created fine.
B: The invisible index ORD_CUSTOMERS_IX1 and the bitmap index are both updated by DML operations on the Orders table.
F: Since ORD_CUSTOMERS_IX1 is invisible only ORD_CUSTOMERS_IX2 is used by the query optimizer.
Not C,Not D,Not E:
* ord_customer_ix1 is an invisible index and is therefore not used by the optimizer.
* VISIBLE | INVISIBLE Use this clause to specify whether the index is visible or invisible to the optimizer. An invisible index is maintained by DML operations, but it is not be used by the optimizer during queries unless you explicitly set the parameter OPTIMIZER_USE_INVISIBLE_INDEXES to TRUE at the session or system level.
Note: Specify BITMAP to indicate that index is to be created with a bitmap for each distinct key, rather than indexing each row separately. Bitmap indexes store the rowids associated with a key value as a bitmap. Each bit in the bitmap corresponds to a possible rowid. If the bit is set, then it means that the row with the corresponding rowid contains the key value. The internal representation of bitmaps is best suited for applications with low levels of concurrent transactions, such as data warehousing.

Ray520 2 years, 3 months ago
https://oracle-base.com/articles/11g/invisible-indexes-11gr1
upvoted 1 times

Question #119 *Topic 1*

Which two statements are true when row archival management is enabled? (Choose two.)

A. The ORA_ARCHIVE_STATE column visibility is controlled by the ROW ARCHIVAL VISIBILITY session parameter.

B. The ORA_ARCHIVE_STATE column is updated manually or by a program that could reference activity tracking columns, to indicate that a row is no longer considered active.

C. The ROW ARCHIVAL VISIBILITY session parameter defaults to active rows only.

D. The ORA_ARCHIVE_STATE column is visible if referenced in the select list of a query.

E. The ORA_ARCHIVE_STATE column is updated automatically by the Oracle Server based on activity tracking columns, to Indicate that a row is no longer

Correct Answer: *CD*

mynameisladyviolette 1 year, 10 months ago

Option D:
To manage in-database archiving for a table, you must enable ROW ARCHIVAL for the table either CREATE TABLE/ALTER TABLE, for example:
CREATE TABLE residents (
id NUMBER,
street VARCHAR2(50),
CONSTRAINT tbb1_pk PRIMARY KEY (id))
ROW ARCHIVAL;

Enabling in-database archiving causes the addition of a system generated hidden
column called ORA_ARCHIVE_STATE
SELECT column_id, column_name, data_type, data_length, hidden_column FROM user_tab_cols
WHERE table_name = 'RESIDENTS'
ORDER BY column_id;

COLUMN_ID COLUMN_NAME DATA_TYPE DATA_LENGTH HID
---------- -------------------- -------------------- ---------- ---
1 ID NUMBER 22 NO
2 STREET VARCHAR2 50 NO
ORA_ARCHIVE_STATE VARCHAR2 4000 YES

Option C: By default, ORA_ARCHIVE_STATE column is populated with the value '0' for each row
When the value for ORA_ARCHIVE_STATE is set to '0' then that row is visible to applications, meaning that the row is active and visible via a standard query

upvoted 1 times

Ray520 2 years, 3 months ago

alter session set row archival visibility = all;
select employee_id, first_name, ora_archive_state from hr.employees;
ORA-00904: "ORA_ARCHIVE_STATE": invalid identifier

alter table hr.employees row archival;
select employee_id, first_name, ora_archive_state from hr.employees;---now output returned. So, ROW ARCHIVAL VISILITY does not control ora_archive_state visibility.

https://www.oracle.com/webfolder/technetwork/tutorials/obe/db/12c/r1/ilm/row_archival/row_archival.html

upvoted 1 times

JMAN1 2 years, 6 months ago

I think
https://oracle-base.com/articles/12c/in-database-archiving-12cr1
Answer is B C D or.... CD? (if hggz is right. but I do not know why B is wrong.)

upvoted 1 times

euler 4 years, 11 months ago

answer should be BCD, not CD.

upvoted 2 times

hggz 4 years, 4 months ago

B is wrong. The “OR by a program that could reference
activity tracking columns” is false

upvoted 4 times

Question #120 *Topic 1*

A warehouse fact table in your Oracle 12c Database is range-partitioned by month and accessed frequently with queries that span multiple partitions
The table has a local prefixed, range partitioned index.
Some of these queries access very few rows in some partitions and all the rows in other partitions, but these queries still perform a full scan for all accessed partitions.
This commonly occurs when the range of dates begins at the end of a month or ends close to the start of a month.
You want an execution plan to be generated that uses indexed access when only a few rows are accessed from a segment, while still allowing full scans for segments where many rows are returned.
Which three methods could transparently help to achieve this result? (Choose three.)

A. Using a partial local Index on the warehouse fact table month column with indexing disabled to the table partitions that return most of their rows to the queries.

B. Using a partial local Index on the warehouse fact table month column with indexing disabled for the table partitions that return a few rows to the queries.

C. Using a partitioned view that does a UNION ALL query on the partitions of the warehouse fact table, which retains the existing local partitioned column.

D. Converting the partitioned table to a partitioned view that does a UNION ALL query on the monthly tables, which retains the existing local partitioned column.

E. Using a partial global index on the warehouse fact table month column with indexing disabling for the table partitions that return most of their rows to the queries.

F. Using a partial global index on the warehouse fact table month column with indexing disabled for the table partitions that return a few rows to the queries.

Correct Answer: *ACE*
Note:
* Oracle 12c now provides the ability to index a subset of partitions and to exclude the others.
Local and global indexes can now be created on a subset of the partitions of a table. Partial Global indexes provide more flexibility in index creation for partitioned tables. For example, index segments can be omitted for the most recent partitions to ensure maximum data ingest rates without impacting the overall data model and access for the partitioned object.
Partial Global Indexes save space and improve performance during loads and queries. This feature supports global indexes that include or index a certain subset of table partitions or subpartitions, and exclude the others. This operation is supported using a default table indexing property. When a table is created or altered, a default indexing property can be specified for the table or its partitions.

Question #121 *Topic 1*

You use the segment advisor to help determine objects for which space may be reclaimed.
Which three statements are true about the advisor given by the segment advisor? (Choose three.)

A. It may advise the use of online table redefinition for tables in dictionary managed tablespace.

B. It may advise the use of segment shrink for tables in dictionary managed tablespaces it the no chained rows.

C. It may advise the use of online table redefinition for tables in locally managed tablespaces

D. It will detect and advise about chained rows.

E. It may advise the use of segment shrink for free list managed tables.

Correct Answer: *ACD*

Ray520 2 years, 3 months ago

On a system with sufficient resources for parallel execution, and in the case where the interim table is not partitioned, redefinition of a LONG column to a LOB column can be executed in parallel, provided that:

The segment used to store the LOB column in the interim table belongs to a locally managed tablespace with Automatic Segment Space Management (ASSM) enabled.

There is a simple mapping from one LONG column to one LOB column, and the interim table has only one LOB column.

upvoted 1 times

Ray520 2 years, 3 months ago

https://docs.oracle.com/cd/E18283_01/server.112/e17120/schema003.htm#:~:text=The%20Segment%20Advisor%20identifies%20segments,the%0data%20in%20the%20segment.

Only Local Tablespaces can have ASSM, and only tables with ASSM can have shrink operation be performed on them.
Dictionary Tablespaces have dictionary managed extent management, and only online table redefinition can be done on tables belonging to it fo segment space reclaim task.

upvoted 1 times

Question #122 *Topic 1*

You have altered a non-unique index to be invisible to determine if queries execute within an acceptable response time without using this index.
Which two are possible if table updates are performed which affect the invisible index columns? (Choose two.)

A. The index remains invisible.

B. The index is not updated by the DML statements on the indexed table.

C. The index automatically becomes visible in order to have it updated by DML on the table.

D. The index becomes unusable but the table is updated by the DML.

E. The index is updated by the DML on the table.

Correct Answer: *AE*
Unlike unusable indexes, an invisible index is maintained during DML statements.
Note:
* Oracle 11g allows indexes to be marked as invisible. Invisible indexes are maintained like any other index, but they are ignored by the optimizer unless the
OPTIMIZER_USE_INVISIBLE_INDEXES parameter is set to TRUE at the instance or session level. Indexes can be created as invisible by using the INVISIBLE keyword, and their visibility can be toggled using the ALTER INDEX command.

Question #123 *Topic 1*

In your multitenant container database (CDB) containing same pluggable databases (PDBs), you execute the following commands in the root container:

```
SQL> CREATE ROLE c##role1;
SQL> GRANT create view, create procedure to c##role1;
SQL> GRANT c##role1 to c##a_admin;
```

Which two statements are true? (Choose two.)

A. The C # # ROLE1 role is created in the root database and all the PDBs.

B. The C # # ROLE1 role is created only in the root database because the container clause is not used.

C. Privileges are granted to the C##A_ADMIN user only in the root database.

D. Privileges are granted to the C##A_ADMIN user in the root database and all PDBs.

E. The statement for granting a role to a user fails because the CONTAINER clause is not used.

Correct Answer: *AC*

* You can include the CONTAINER clause in several SQL statements, such as the CREATE USER, ALTER USER, CREATE ROLE, GRANT, REVOKE, and
ALTER SYSTEM statements.
* * CREATE ROLE with CONTAINER (optional) clause
/ CONTAINER = ALL
Creates a common role.
/ CONTAINER = CURRENT
Creates a local role in the current PDB.

Community vote distribution

AD (100%)

Ray_gk 2 years, 3 months ago

Both A and C are correct. Verified C in Oracle DB. C is correct as granting the privileges to a common role, while omitting the container clause, wil grant the privileges using default container=current clause.

Therefore, the grant create view, create procedure to c##role1; command will only grant these privileges to the common role for the current container, which is give as the root container in the question.

If the command including container=all, then D would be correct.

upvoted 1 times

Ray_gk 2 years, 3 months ago

Basically, the grant command in a multitenant DB uses the container=current clause by default

upvoted 1 times

Ray520 2 years, 3 months ago

https://docs.oracle.com/database/121/SQLRF/statements_6014.htm#SQLRF01311

The CONTAINER clause applies when you are connected to a CDB. However, it is not necessary to specify the CONTAINER clause because its defa values are the only allowed values.

-To create a common role, you must be connected to the root. You can optionally specify CONTAINER = ALL, which is the default when you are connected to the root.
-To create a local role, you must be connected to a PDB. You can optionally specify CONTAINER = CURRENT, which is the default when you are connected to a PDB.

Role created with C## prefix or whichever string set as COMMON_USER_PREFIX, when created in root, gets created in all PDBs.
Any User garnted access to the common role in CDB root, gets only assigned in CDB root.
So,
create role c##admin_tsrl;
grant create table,create session,create view to c##admin_tsrl;
grant c##admin_tsrl to C##A_ADMIN;---- only c##a_admin user (also a common user here, hence present in all PDBs as well), in the root gets assigned this role. The c##a_admin user in PDB doesn't get assigned this role.

upvoted 2 times

NorwayOracle 2 years, 4 months ago

Selected Answer: AD

A common role is a database role that exists in the root and in every existing and future pluggable database (PDB). Common roles are useful for cross-container operations, ensuring that a common user has a role in every container.

upvoted 1 times

jackymak 2 years, 4 months ago

Selected Answer: AD

Common user can assess to PDB, so I think common user is also owned the privileges in PDB

upvoted 1 times

lollo1234 3 years, 2 months ago

Confirm A, default Value for CONTAINER is ALL when connected to the root:

To create a common role, you must be connected to the root. You can optionally specify CONTAINER = ALL, which is the default when you are connected to the root.

Ref.: https://docs.oracle.com/database/121/SQLRF/statements_6014.htm#SQLRF01311 in Section "CONTAINER clause"

upvoted 1 times

hainx17 3 years ago

what about C?

upvoted 1 times

jackymak 2 years, 5 months ago

But the common user can execute the right in PDB, right?

upvoted 1 times

Question #124 *Topic 1*

The persistent configuration settings for RMAN have default for all parameters.
Identify four RMAN commands that produce a multi-section backup.

A. BACKUP TABLESPACE SYSTEM SECTION SIZE 100M;

B. BACKUP AS COPY TABLESPACE SYSTEM SECTION SIZE 100M;

C. BACKUP ARCHIVELOG ALL SECTION SIZE 25M;

D. BACKUP TABLESPACE "TEMP" SECTION SIZE 10M;

E. BACKUP TABLESPACE "UNDO" INCLUDE CURRENT CONTROLFILE SECTION SIZE 100M;

F. BACKUP SPFILE SECTION SIZE 1M;

G. BACKUP INCREMENTAL LEVEL 0 TABLESPACE SYSAUX SECTION SIZE 100M;

Correct Answer: *ABEG*

khan Highly Voted 5 years, 3 months ago

Answer: ABEG
More info on RMAN Multisection Backups

For backing up very large data files, RMAN provides multisection backups as a way to parallelize the backup operation within the file itself, such that sections of a file are backed up in parallel, rather than backing up on a per-file basis.

For example, a one TB data file can be sectioned into ten 100 GB backup pieces, with each section backed up in parallel, rather than the entire on TB file backed up as one file. The overall backup time for large data files can be dramatically reduced.

https://docs.oracle.com/database/121/VLDBG/GUID-FA70B6EF-0C8B-435B-98C9-CDEBCDCED145.htm#VLDBG1576

upvoted 5 times

jackymak 2 years, 5 months ago

Why are they multisection backup?
Because they are separated as 100M per section?
Or the tablespace are much larger than archive log, sp file and temp table?

upvoted 1 times

Ray520 Most Recent 2 years, 3 months ago

BACKUP ARCHIVELOG ALL SECTION SIZE 25M; This command does not fail, but only creates archivelog backups, not multi section backups.
BACKUP SPFILE SECTION SIZE 1M; This command also does not fail, but only creates normal backup, not multi section backups.

upvoted 2 times

Question #125 *Topic 1*

Flashback is enabled for your multitenant container database (CDB), which contains two pluggable database (PDBs). A local user was accidently dropped from one of the PDBs.
You want to flash back the PDB to the time before the local user was dropped. You connect to the CDB and execute the following commands:

SQL > SHUTDOWN IMMEDIATE -

SQL > STARTUP MOUNT -
SQL > FLASHBACK DATABASE to TIME "TO_DATE ('08/20/12' , 'MM/DD/YY')";
Examine following commands:
1. ALTER PLUGGABLE DATABASE ALL OPEN;
2. ALTER DATABASE OPEN;
3. ALTER DATABASE OPEN RESETLOGS;
Which command or commands should you execute next to allow updates to the flashback back schema?

A. Only 1

B. Only 2

C. Only 3

D. 3 and 1

E. 1 and 2

Correct Answer: *D*

Question #126 *Topic 1*

Examine the commands executed to monitor database operations:
$> conn sys oracle/oracle@prod as sysdba

SQL > VAR eid NUMBER -
SQL > EXEC: eid := DBMS_SQL_MONITOR.BEGIN_OPERATION (batch_job , FORCED_TRACKING => Y);
Which two statements are true? (Choose two.)

A. Database operations will be monitored only when they consume a significant amount of resource.

B. Database operations for all sessions will be monitored.

C. Database operations will be monitored only if the STATISTICS_LEVEL parameter is set to TYPICAL and CONTROL_MANAGEMENT_PACK_ACCESS is set DIAGNISTIC + TUNING.

D. Only DML and DDL statements will be monitored for the session.

E. All subsequent statements in the session will be treated as one database operation and will be monitored.

Correct Answer: *CE*
C: Setting the CONTROL_MANAGEMENT_PACK_ACCESS initialization parameter to DIAGNOSTIC+TUNING (default) enables monitoring of database operations. Real-Time SQL Monitoring is a feature of the Oracle Database Tuning Pack.
Note:
* The DBMS_SQL_MONITOR package provides information about Real-time SQL Monitoring and Real-time Database Operation Monitoring.
*(not B) BEGIN_OPERATION Function
starts a composite database operation in the current session.
/ (E) FORCE_TRACKING - forces the composite database operation to be tracked when the operation starts. You can also use the string variable 'Y'.
/ (not A) NO_FORCE_TRACKING - the operation will be tracked only when it has consumed at least 5 seconds of CPU or I/O time. You can also use the string variable 'N'.

Question #127 *Topic 1*

Which three statements are true about the working of system privileges in a multitenant control database (CDB) that has pluggable databases (PDBs)? (Choose three.)

A. System privileges apply only to the PDB in which they are used.

B. Local users cannot use local system privileges on the schema of a common user.

C. The granter of system privileges must possess the set container privilege.

D. Common users connected to a PDB can exercise privileges across other PDBs.

E. System privileges with the with grant option container all clause must be granted to a common user before the common user can grant privileges to other users.

Correct Answer: *ACE*

A, Not D: In a CDB, PUBLIC is a common role. In a PDB, privileges granted locally to PUBLIC enable all local and common users to exercise these privileges in this PDB only.

C: A user can only perform common operations on a common role, for example, granting privileges commonly to the role, when the following criteria are met:

The user is a common user whose current container is root.

The user has the SET CONTAINER privilege granted commonly, which means that the privilege applies in all containers.

The user has privilege controlling the ability to perform the specified operation, and this privilege has been granted commonly

Incorrect:

Note:

* Every privilege and role granted to Oracle-supplied users and roles is granted commonly except for system privileges granted to PUBLIC, which are granted locally.

Question #128 *Topic 1*

You are about to plug a multi-terabyte non-CDB into an existing multitenant container database (CDB) as a pluggable database (PDB).
The characteristics of the non-CDB are as follows:
Version: Oracle Database 12c Releases 1 64-bit
Character set: WE8ISO8859P15
National character set: AL16UTF16
O/S: Oracle Linux6 64-bit
The characteristics of the CDB are as follows:
Version: Oracle Database 12c Release 1 64-bit
Character set: AL32UTF8
O/S: Oracle Linux 6 64-bit
Which technique should you use to minimize down time while plugging this non-CDB into the CDB?

A. Transportable database

B. Transportable tablespace

C. Data Pump full export / import

D. The DBMS_PDB package

E. RMAN

Correct Answer: *C*

Ray_gk 2 years, 2 months ago
I agree with C. This exam is based on 12cR1. In 12cR1, Oracle did not yet support differing charactersets in the same container database. Howeve future releases do support this so I believe in the case of future releases D would be correct.
upvoted 1 times

Ray520 2 years, 3 months ago
Reference: https://support.oracle.com/epmos/faces/DocumentDisplay?_afrLoop=237882283796393&parent=EXTERNAL_SEARCH&sourceId=PROBLEM&id=1968706.1&_afrWindowMode=0&_adf.ctrl-state=9tgv5ju3b

In Oracle Database 12c, all pluggable databases (PDBs) in a container database (CDB) must have
* the same Database character set (NLS_CHARACTERSET) or the NLS_CHARACTERSET need to be a (Plug-in compatible) binary subset of the CDB NLS_CHARACTERSET
* the same National character set (NLS_NCHAR_CHARACTERSET) as the CDB's root container
in order to be able to plug in.
upvoted 2 times

Ray520 2 years, 3 months ago
We can use DMU to convert the non-CDB characterset first and then mograte. Export/import does not have character set dependency. Answe is Export/import
upvoted 2 times

katetel 2 years, 11 months ago
finally, the answer is d
upvoted 3 times

PaoloBL 2 years, 11 months ago
D would be faster but 12.1 does not support pdb with character set different from cdb
upvoted 1 times

charan94 3 years, 4 months ago
D is correct answer. Since, most of the time the the non cdb will be online i.e pre tasks and will be offline only while plugging in to as pdb.
upvoted 3 times

dam101 3 years, 4 months ago
Answer: C - ref https://asktom.oracle.com/pls/apex/f?p=100:11:0::::P11_QUESTION_ID:9537016000346410905
upvoted 3 times

Modasser 3 years, 6 months ago
Answer should be D
coz:
– If the non-CDB is version pre 12c or 12c you can still consider using DBMS_PDB package.
– If the non-CDB is version 11.2.0.3 or later you can still consider using Transport Database.
– If the non-CDB is version pre-11.2.0.3 you can still consider using transportable tablespaces.

upvoted 4 times

TuxBingo 4 years ago

The correct answer is D, check this link: http://dbaworkshop.blogspot.com/2013/08/How-to-create-a-PDB-from-a-Non-CDB-using-DBMS-PDB-package.html

upvoted 2 times

vineetarneja 4 years, 1 month ago

yes. the correct answer is D only. the size and characterset of the non-cdb doesn't matter while plugging it into a cdb.

upvoted 2 times

5kong 4 years, 4 months ago

Question #54

upvoted 2 times

NorthKorean 4 years, 4 months ago

Correct Answer: B
* Overview, example:
- Log into ncdb12c as sys
- Get the database in a consistent state by shutting it down cleanly.
- Open the database in read only mode
- Run DBMS_PDB.DESCRIBE to create an XML file describing the database.
- Shut down ncdb12c
- Connect to target CDB (CDB2)
- Check whether non-cdb (NCDB12c) can be plugged into CDB(CDB2)
- Plug-in Non-CDB (NCDB12c) as PDB(NCDB12c) into target CDB(CDB2).
- Access the PDB and run the noncdb_to_pdb.sql script.
- Open the new PDB in read/write mode.
* You can easily plug an Oracle Database 12c non-CDB into a CDB. Just create a PDB manifest file for the non-CDB, and then use the manifest file to create a cloned PDB in the CDB.
* Note that to plug in a non-CDB database into a CDB, the non-CDB database needs to be of version 12c as well. So existing 11g databases wi need to be upgraded to 12c before they can be part of a 12c CDB.

upvoted 2 times

hggz 4 years, 4 months ago

From your answer D is correct: use DBMS_PDB package

upvoted 4 times

Patrick9230 4 years, 4 months ago

B is error . Transportable tablespace have two condition.
Source and Target DB' Character Set,National Character Set must be same

upvoted 3 times

NorthKorean 4 years, 4 months ago

I just copied from #54

upvoted 1 times

Question #129 *Topic 1*

Your database has the SRV1 service configured for an application that runs on middle-tier application server. The application has multiple modules. You enable tracing at the service level by executing the following command:
SQL > exec DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE (SRV1);
The possible outcome and actions to aggregate the trace files are as follows:
1. The command fails because a module name is not specified.
2. A trace file is created for each session that is running the SRV1 service.
3. An aggregated trace file is created for all the sessions that are running the SRV1 service.
4. The trace files may be aggregated by using the trcess utility.
5. The trace files be aggregated by using the tkprof utility.
Identify the correct outcome and the step to aggregate by using tkprof utility?

A. 1

B. 2 and 4

C. 2 and 5

D. 3 and 4

E. 3 and 5

Correct Answer: *B*
Tracing information is present in multiple trace files and you must use the trcsess tool to collect it into a single file.
Incorrect:

Not 1: Parameter service_name -
Name of the service for which tracing is enabled.
module_name
Name of the MODULE. An optional additional qualifier for the service.
Note:
* The procedure enables a trace for a given combination of Service, MODULE and ACTION name. The specification is strictly hierarchical: Service Name or
Service Name/MODULE, or Service Name, MODULE, and ACTION name must be specified. Omitting a qualifier behaves like a wild-card, so that not specifying an ACTION means all ACTIONs. Using the ALL_ACTIONS constant achieves the same purpose.
* SERV_MOD_ACT_TRACE_ENABLE Procedure
This procedure will enable SQL tracing for a given combination of Service Name, MODULE and ACTION globally unless an instance_name is specified.
* DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE(
service_name IN VARCHAR2,
module_name IN VARCHAR2 DEFAULT ANY_MODULE,
action_name IN VARCHAR2 DEFAULT ANY_ACTION,
waits IN BOOLEAN DEFAULT TRUE,
binds IN BOOLEAN DEFAULT FALSE,
instance_name IN VARCHAR2 DEFAULT NULL);

Ray520 2 years, 3 months ago
https://docs.oracle.com/database/121/ARPLS/d_monitor.htm#ARPLS67174

This procedure will enable SQL tracing for a given combination of Service Name, MODULE and ACTION globally unless an instance_name is specified.
service_name: Name of the service for which SQL trace is enabled (mandatory field).
DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE(
service_name IN VARCHAR2,
module_name IN VARCHAR2 DEFAULT ANY_MODULE,
action_name IN VARCHAR2 DEFAULT ANY_ACTION,
waits IN BOOLEAN DEFAULT TRUE,
binds IN BOOLEAN DEFAULT FALSE,
instance_name IN VARCHAR2 DEFAULT NULL,
plan_stat IN VARCHAR2 DEFAULT NULL);

tracing information is present in multiple trace files and you must use the trcsess tool to collect it into a single file.

upvoted 1 times

Question #130 *Topic 1*

Your multitenant container database (CDB) contains pluggable databases (PDBs), you are connected to the HR_PDB. You execute the following command:
SQL > CREATE UNDO TABLESPACE undotb01
DATAFILE u01/oracle/rddb1/undotbs01.dbf SIZE 60M AUTOEXTEND ON;
What is the result?

A. It executes successfully and creates an UNDO tablespace in HR_PDB.

B. It falls and reports an error because there can be only one undo tablespace in a CDB.

C. It fails and reports an error because the CONTAINER=ALL clause is not specified in the command.

D. It fails and reports an error because the CONTAINER=CURRENT clause is not specified in the command.

E. It executes successfully but neither tablespace nor the data file is created.

Correct Answer: *E*
Interesting behavior in 12.1.0.1 DB of creating an undo tablespace in a PDB. With the new Multitenant architecture the undo tablespace resides at the CDB level and PDBs all share the same UNDO tablespace.
When the current container is a PDB, an attempt to create an undo tablespace fails without returning an error.

TuxBingo 4 years ago
The correct answer is E (12.1), remember between 12.1 and 12.2, most of the important features are in 12.2
upvoted 1 times

rdnkrkmz 5 years, 6 months ago
For 12.2 release,
A is correct answer. Because local_undo_enabled parameter allow to create local UNDO in PDB (tested 12.2)
upvoted 1 times

hggz 4 years, 4 months ago
Yes, A for 12.2, but E for 12.1
upvoted 2 times

Question #131 *Topic 1*

Which three statements are true about SQL plan directives? (Choose three.)

A. They are tied to a specific statement or SQL ID.

B. They instruct the maintenance job to collect missing statistics or perform dynamic sampling to generate a more optimal plan.

C. They are used to gather only missing statistics.

D. They are created for a query expression where statistics are missing or the cardinality estimates by the optimizer are incorrect.

E. They instruct the optimizer to create only column group statistics.

F. Improve plan accuracy by persisting both compilation and execution statistics in the SYSAUX tablespace.

Correct Answer: *BDF*

Ray520 2 years, 3 months ago

https://oracle-base.com/articles/12c/sql-plan-directives-12cr1

SQL Plan Directives are one of the functionalities that make up Adaptive Query Optimization in Oracle Database 12c:
-SQL plan directives are like "extra notes" for the optimizer, to remind it that it previously selected a suboptimal plan, typically because of incorre cardinality estimates
-Unlike SQL profiles, which are statement specific, SQL plan directives are linked to query expressions, so they can be used by several statements containing matching query expressions.
-Situations like missing histograms or missing extended statistics may result in SQL plan directives being generated.
-Situations like automatic reoptimization may result in SQL plan directives being written to the SGA and later persisted to the SYSAUX tablespace at which point they can be displayed using the DBA_SQL_PLAN_DIRECTIVES and DBA_SQL_PLAN_DIR_OBJECTS views.

upvoted 1 times

Question #132 *Topic 1*

You want to flash back a test database by five hours.

You issue this command:

SQL > FLASHBACK DATABASE TO TIMESTAMP (SYSDATE - 5/24);

Which two statements are true about this flashback scenario? (Choose two.)

A. The database must have multiplexed redo logs for the flashback to succeed.

B. The database must be MOUNTED for the flashback to succeed.

C. The database must use block change tracking for the flashback to succeed.

D. The database must be opened in restricted mode for the flashback to succeed.

E. The database must be opened with the RESETLOGS option after the flashback is complete.

F. The database must be opened in read-only mode to check if the database has been flashed back to the correct SCN.

Correct Answer: *BE*

Currently there are no comments in this discussion, be the first to comment!

Question #133 *Topic 1*

Examine these two statements:

```
SQL> CREATE BIGFILE TABLESPACE MRKT
  2  DATAFILE '/u01/app/oracle/oradata/orcl/mrkt.dbf' size 10M LOGGING
  3  EXTENT MANAGEMENT LOCAL SEGMENT SPACE MANAGEMENT AUTO;

Tablespace created.

SQL> ALTER DATABASE DEFAULT TABLESPACE MRKT;

Database altered.
```

Which three are true about the MRKT tablespace? (Choose three.)

A. The MRKT tablespace is created as a small file tablespace, because the file size is less than the minimum required for big file files.

B. The MRKT tablespace may be dropped if it has no contents.

C. Users who were using the old default tablespace will have their default tablespaces changed to the MRKT tablespace.

D. No more data files can be added to the tablespace.

E. The relative file number of the tablespace is not stored in rowids for the table rows that are stored in the MRKT tablespace.

Correct Answer: *CDE*

Ray520 2 years, 3 months ago

https://docs.oracle.com/cd/E18283_01/server.112/e17120/tspaces002.htm#:~:text=Creating%20a%20Bigfile%20Tablespace,with%20automatic%2egment%20space%20management.

bigfile tablespace is a tablespace with a single, but very large (up to 4G blocks) datafile.
A bigfile tablespace consists of a single data or temporary file which can be up to 128 TB.

upvoted 1 times

Question #134 *Topic 1*

In your database, you want to ensure that idle sessions that are blocking active are automatically terminated after a specified period of time.
How would you accomplish this?

A. Setting a metric threshold

B. Implementing Database Resource Manager

C. Enabling resumable timeout for user sessions

D. Decreasing the value of the IDLE_TIME resource limit in the default profile

Correct Answer: *B*

khan Highly Voted 5 years, 3 months ago
B is the answer.

https://docs.oracle.com/database/121/ADMIN/dbrm.htm#ADMIN027

Limit the amount of time that a session can be idle. This can be further defined to mean only sessions that are blocking other sessions.

Idle Time Limit
You can specify an amount of time that a session can be idle, after which it is terminated. You can also specify a more stringent idle time limit that applies to sessions that are idle and blocking other sessions
upvoted 14 times

SADIO Highly Voted 2 years, 9 months ago
It cannot be D because not all the users have the default profil as their profil.
A user can have an other profile as his default profil. So it won't have effect on that user.
So i think the answer is B
upvoted 5 times

Phoenix22 Most Recent 4 years, 4 months ago
And why is not D?
upvoted 1 times

soniat 3 years, 5 months ago
they key word is session blocking others. D will terminate idle session rather it is blocking or not
upvoted 2 times

Mamlouk_Med 5 years, 5 months ago
I think D too
upvoted 1 times

rdnkrkmz 5 years, 6 months ago
I think D
upvoted 1 times

Question #135 *Topic 1*

You Execute the Following command to create a password file in the database server:
$ orapwd file = +DATA/PROD/orapwprod entries = 5 ignorecase = N format = 12
Which two statements are true about the password file? (Choose two.)

A. It records the usernames and passwords of users when granted the DBA role.

B. It contains the usernames and passwords of users for whom auditing is enabled.

C. Is used by Oracle to authenticate users for remote database administration.

D. It records the usernames and passwords of all users when they are added to the OSDBA or OSOPER operating system groups.

E. It supports the SYSBACKUP, SYSDG, and SYSKM system privileges.

Correct Answer: *CE*

Question #136 *Topic 1*

Identify two situations in which the alert log file is updated.

A. Running a query on a table returns ORA-600: Internal Error.

B. Inserting a value into a table returns ORA-01722: invalid number.

C. Creating a table returns ORA-00955: name us already in used by an existing objects.

D. Inserting a value into a table returns ORA-00001: unique constraint (SYS.OK_TECHP) violated.

E. Rebuilding an index using ALTER INDEX . . . REBUILD fails with an ORA-01578: ORACLE data block corrupted (file # 14, block # 50) error.

Correct Answer: *AE*
The alert log is a chronological log of messages and errors, and includes the following items:
*All internal errors (ORA-600), block corruption errors (ORA-1578), and deadlock errors (ORA-60) that occur
* Administrative operations, such as CREATE, ALTER, and DROP statements and STARTUP, SHUTDOWN, and ARCHIVELOG statements
* Messages and errors relating to the functions of shared server and dispatcher processes
* Errors occurring during the automatic refresh of a materialized view
* The values of all initialization parameters that had nondefault values at the time the database and instance start
Note:
* The alert log file (also referred to as the ALERT.LOG) is a chronological log of messages and errors written out by an Oracle Database. Typical messages found in this file is: database startup, shutdown, log switches, space errors, etc. This file should constantly be monitored to detect unexpected messages and corruptions.

Question #137 *Topic 1*

Which two statements are true about Oracle Data Pump export and import operations? (Choose two.)

A. You can detach from a data pump export job and reattach later.

B. Data pump uses parallel execution server processes to implement parallel import.

C. Data pump import requires the import file to be in a directory owned by the oracle owner.

D. The master table is the last object to be exported by the data pump.

E. You can detach from a data pump import job and reattach later.

Correct Answer: *AB*

B: Data Pump can employ multiple worker processes, running in parallel, to increase job performance.

D: For export jobs, the master table records the location of database objects within a dump file set. / Export builds and maintains the master table for the duration of the job. At the end of an export job, the content of the master table is written to a file in the dump file set.

/ For import jobs, the master table is loaded from the dump file set and is used to control the sequence of operations for locating objects that need to be imported into the target database.

Community vote distribution

AE (100%)

hggz Highly Voted 4 years, 4 months ago

A,E is correct:
https://docs.oracle.com/database/121/SUTIL/GUID-EBC54840-22B7-486C-B1FD-CAFFB5FC2344.htm#SUTIL871
https://docs.oracle.com/database/121/SUTIL/GUID-0FD688BF-A25E-40D3-B835-E2CBF069F26E.htm#SUTIL945

D is not correct, there is no master table export, only master table content written to file.

What with B?

upvoted 7 times

soniat 3 years, 4 months ago

ABE is the correct answer

upvoted 3 times

luizyto Most Recent 1 year, 8 months ago

Selected Answer: AE

A and E

upvoted 1 times

Ray520 2 years ago

Selected Answer: AE

Both expdp and impdp jobs are detachable and attachable.
Regarding option B, it's not parallel execution server processes, to be exact "The value you specify for integer specifies the maximum number of processes of active execution operating on behalf of the import job. This execution set consists of a combination of worker processes and paralle I/O server processes."

upvoted 1 times

dfabac 2 years, 1 month ago

Selected Answer: AE

Answer is AE.

B is incorrect because parallelism on import process is not related to parallel execution server processes. Task is performed by database writer processes, ie:

Worker 1 Status:
Process Name: DW00
State: EXECUTING
Object Schema: XTP_AC
Object Name: SYS_C0063284986
Object Type: SCHEMA_EXPORT/TABLE/CONSTRAINT/CONSTRAINT
Completed Objects: 1,120
Worker Parallelism: 1

Worker 2 Status:
Process Name: DW01
State: WORK WAITING

Worker 3 Status:
Process Name: DW02
State: WORK WAITING
upvoted 1 times

Ray_gk 2 years, 3 months ago

A & E are correct.

You can both attach and reattach for export and import jobs.

D is incorrect because at the end of an export job, the content of the master table is written to a file.

I believe B is incorrect as datapump does not use parallel execution server process for parallel imports by default. We have to specifically provide the parallel paramter during the operation. However, I am not too sure on this rationale.

I believe we are all in agreement that C is incorrect for obvious reasons.
upvoted 1 times

PerkDizzzle 4 years, 10 months ago

B & D are correct

Ref: https://docs.oracle.com/database/121/SUTIL/GUID-6BDC1CC8-8596-402D-B016-602985B97AB6.htm#SUTIL808
upvoted 2 times

rdnkrkmz 5 years, 6 months ago

import and export job are both attachable. So A and D is correct (tested 11.2.0.4). But may B be correct, too.
Confused. If I see this question in exam, I mark AD.
upvoted 2 times

rcgenilo27 4 years, 10 months ago

You mean E?
upvoted 3 times

Question #138 *Topic 1*

Examine the query and its output executed In an RDBMS Instance:

```
SQL> SELECT * FROM v$pwfile_users;

USERNAME                       SYSDB SYSOP SYSAS SYSBA SYSDG SYSKM      CON_ID
------------------------------ ----- ----- ----- ----- ----- ----- ----------
SYS                            TRUE  TRUE  FALSE FALSE FALSE FALSE          0
C##B_ADMIN                     FALSE FALSE FALSE TRUE  FALSE FALSE          0
C##C_ADMIN                     FALSE FALSE FALSE FALSE TRUE  FALSE          0
C##A_ADMIN                     FALSE FALSE FALSE FALSE FALSE TRUE           0
C##D_ADMIN                     FALSE FALSE TRUE  FALSE FALSE FALSE          0
```

Which three statements are true about the users (other than sys) in the output? (Choose three.)

A. The C # # B_ADMIN user can perform all backup and recovery operations using RMAN only.

B. The C # # C_ADMIN user can perform the data guard operation with Data Guard Broker.

C. The C # # A_ADMIN user can perform wallet operations.

D. The C # # D_ADMIN user can perform backup and recovery operations for Automatic Storage Management (ASM).

E. The C # # B_ADMIN user can perform all backup and recovery operations using RMAN or SQL* Plus.

Correct Answer: *BDE*

B: SYSDG administrative privilege has ability to perform Data Guard operations (including startup and shutdown) using Data Guard Broker or dgmgrl.

D: SYSASM -
The new (introduced in 11g) SYSASM role to manage the ASM instance, variable extent sizes to reduce shared pool usage, and the ability of an instance to read from a specific disk of a diskgroup
E (Not A): SYSDBA is like a role in the sense that it is granted, but SYSDBA is a special built-in privilege to allow the DBA full control over the database
Incorrect:
Not C: SYSKM. SYSKM administrative privilege has ability to perform transparent data encryption wallet operations.
Note:
Use the V$PWFILE_USERS view to see the users who have been granted administrative privileges.

Community vote distribution

BCE (75%) BCD (25%)

Ray520 2 years ago

Selected Answer: BCE

SYSKM: TDE wallet operations
SYSBACKUP: All backup/recovery operations

upvoted 3 times

dfabac 2 years, 1 month ago

Selected Answer: BCD

Anwser is BCD
SYSBACKUP - This privilege allows a user to perform backup and recovery operations either from Oracle Recovery Manager (RMAN) or SQL*Plus.

upvoted 1 times

Ray520 2 years, 3 months ago

https://docs.oracle.com/database/121/ADMQS/GUID-90668C5A-3D99-45E0-BBC0-D2AA98D3F59D.htm

SYSBACKUP privilege encompasses all the privileges required to back up and recover the database. Those privileges are a subset of the privileges included in the SYSDBA administrative privilege.

upvoted 2 times

Question #139 *Topic 1*

In your Database, the TBS PERCENT USED parameter is set to 60 and the TBS PERCENT FREE parameter is set to 20.
Which two storage-tiering actions might be automated when using information Lifecycle Management (ILM) to automate data movement?
(Choose two.)

A. The movement of all segments to a target tablespace with a higher degree of compression, on a different storage tier, when the source tablespace exceeds TBS PERCENT USED

B. Setting the target tablespace to read-only

C. The movement of some segments to a target tablespace with a higher degree of compression, on a different storage tier, when the source tablespace exceeds TBS PERCENT USED

D. Setting the target tablespace offline

E. The movement of some blocks to a target tablespace with a lower degree of compression, on a different storage tier, when the source tablespace exceeds

Correct Answer: *BC*
The value for TBS_PERCENT_USED specifies the percentage of the tablespace quota when a tablespace is considered full. The value for TBS_PERCENT_FREE specifies the targeted free percentage for the tablespace. When the percentage of the tablespace quota reaches the value of
TBS_PERCENT_USED, ADO begins to move data so that percent free of the tablespace quota approaches the value of TBS_PERCENT_FREE. This action by
ADO is a best effort and not a guarantee.

Question #140 *Topic 1*

Which three statements are true about Flashback Database? (Choose three.)

A. Flashback logs are written sequentially, and are archived.

B. Flashback Database uses a restored control file to recover a database.

C. The Oracle database automatically creates, deletes, and resides flashback logs in the Fast Recovery Area.

D. Flashback Database can recover a database to the state that it was in before a reset logs operation.

E. Flashback Database can recover a data file that was dropped during the span of time of the flashback.

F. Flashback logs are used to restore to the blocks' before images, and then the redo data may be used to roll forward to the desired flashback time.

Correct Answer: *CDF*

TuxBingo 4 years ago

Agree with the answers C, D, and F.
– To return the database to the point in time immediately before the most recent OPEN RESETLOGS operation, you use FLASHBACK DATABASE TO BEFORE RESETLOGS.

A is incorrect as Flashback logs are written sequentially, but are not archived.

B is incorrect.
– If the database control file is restored from backup or re-created, all accumulated flashback log information is discarded. You cannot use FLASHBACK DATABASE to return to a point in time before the restore or re-creation of a control file.

upvoted 4 times

Question #141 *Topic 1*

Which statement is true about Enterprise Manager (EM) express in Oracle Database 12c?

A. By default, EM express is available for a database after database creation.

B. You can use EM express to manage multiple databases running on the same server.

C. You can perform basic administrative tasks for pluggable databases by using the EM express interface.

D. You cannot start up or shut down a database Instance by using EM express.

E. You can create and configure pluggable databases by using EM express.

Correct Answer: *D*
References:
http://www.oracle.com/technetwork/database/manageability/emx-intro-1965965.html

NorwayOracle 2 years, 9 months ago
Why not c?
upvoted 2 times

bigswole 2 years, 10 months ago
https://www.oracle.com/technetwork/database/manageability/emx-intro-1965965.html
upvoted 1 times

charan94 3 years, 3 months ago
A is not correct. Since EM is not available by default after db creation.
We need to enable through dbca .
D is correct.
upvoted 3 times

Modasser 3 years, 6 months ago
Answer: D
A incorrect coz, Enterprise Manager Database Express is available only when the database is open
upvoted 1 times

soniat 3 years, 5 months ago
A and D are correct the database is open after creation
upvoted 1 times

Isiee 3 years, 6 months ago
A is correct
upvoted 1 times

NorthKorean 4 years, 4 months ago
I saw a same question and the answer was different. I guess it was (a)
What is the correct answer for this?
upvoted 1 times

fsamkoh 4 years, 2 months ago
D is correct
upvoted 1 times

Question #142 *Topic 1*

Examine the following command;
ALTER SYSTEM SET enable_ddl_logging = TRUE;
Which statement is true?

A. Only the data definition language (DDL) commands that resulted in errors are logged in the alert log file.

B. All DDL commands are logged in the alert log file.

C. All DDL commands are logged in a different log file that contains DDL statements and their execution dates.

D. Only DDL commands that resulted in the creation of new segments are logged.

E. All DDL commands are logged in XML format in the alert directory under the Automatic Diagnostic Repository (ADR) home.

Correct Answer: *E*

Community vote distribution

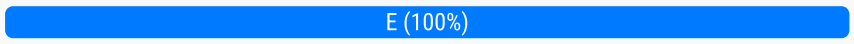

luizyto 1 year, 6 months ago

Selected Answer: E

.There are two DDL logs that contain the same information. One is an XML file, and the other is a text file. The DDL log is stored in the log/ddl subdirectory of the ADR home.

upvoted 1 times

Ray520 2 years, 3 months ago

https://docs.oracle.com/database/121/REFRN/GUID-6FBA6147-D545-4E7D-94F0-A97EC1C721AE.htm#REFRN10302

There are two DDL logs that contain the same information. One is an XML file, and the other is a text file. The DDL log is stored in the log/ddl subdirectory of the ADR home.

upvoted 1 times

bigswole 2 years, 9 months ago

it would appear that none are correct when oracle docs reads a subset of DDL is logged not "ALL"

"ENABLE_DDL_LOGGING enables or disables the writing of a subset of data definition language (DDL) statements to a DDL log."

-https://docs.oracle.com/database/121/REFRN/GUID-6FBA6147-D545-4E7D-94F0-A97EC1C721AE.htm#REFRN10302

C would be the closet to being correct.

upvoted 2 times

katetel 2 years, 11 months ago

E is correct, it's different. #297 is another question

upvoted 1 times

Modasser 3 years, 2 months ago

C is correct.
DDL command are loged in separate log files
There are two DDL logs that contain the same information. One is an XML file, and the other is a text file. The DDL log is stored in the log/ddl subdirectory of the ADR home.
https://docs.oracle.com/database/121/REFRN/GUID-6FBA6147-D545-4E7D-94F0-A97EC1C721AE.htm#REFRN10302

upvoted 2 times

NorthKorean 4 years, 3 months ago

#297 see

upvoted 1 times

NorthKorean 4 years, 3 months ago

There are two DDL logs that contain the same information. One is an XML file, and the other is a text file. The DDL log is stored in the log/ddl subdirectory of the ADR home. So E is incorrect

upvoted 2 times

yosiw96816 3 years, 6 months ago

E is correct.
#297 - D is "All DDL commands are logged in to a trace file in ADR home." which is incorrect (trace file).

upvoted 3 times

euler 4 years, 11 months ago

I think C is the right answer.

upvoted 4 times

rcgenilo27 4 years, 10 months ago

E is correct
https://docs.oracle.com/database/121/REFRN/GUID-6FBA6147-D545-4E7D-94F0-A97EC1C721AE.htm#REFRN10302

upvoted 6 times

Question #143 *Topic 1*

In which two scenarios do you use SQL* Loader to load data? (Choose two.)

A. Transform the data while it is being loaded into the database.

B. Use transparent parallel processing without having to split the external data first.

C. Load data into multiple tables during the same load statement.

D. Generate unique sequential key values in specified columns.

Correct Answer: *CD*

Community vote distribution

CD (67%) AB (33%)

luizyto 1 year, 8 months ago

Selected Answer: CD

However, in the following situations, use SQL*Loader for the best load performance:
- You want to load data remotely.
- Transformations are not required on the data, and the data does not need to be loaded in parallel.

upvoted 1 times

dfabac 2 years, 1 month ago

Selected Answer: CD

(https://docs.oracle.com/cd/B19306_01/server.102/b14215/ldr_concepts.htm#sthref519)

Choosing External Tables Versus SQL*Loader
The record parsing of external tables and SQL*Loader is very similar, so normally there is not a major performance difference for the same record format. However, due to the different architecture of external tables and SQL*Loader, there are situations in which one method is more appropria than the other.

In the following situations, use external tables for the best load performance:
- You want to transform the data as it is being loaded into the database.
- You want to use transparent parallel processing without having to split the external data first.

However, in the following situations, use SQL*Loader for the best load performance:
- You want to load data remotely.
- Transformations are not required on the data, and the data does not need to be loaded in parallel.

upvoted 1 times

Ray520 2 years, 3 months ago

C D is correct:
https://docs.oracle.com/cd/B19306_01/server.102/b14215/ldr_concepts.htm

upvoted 1 times

jackymak 2 years, 5 months ago

Selected Answer: AB

A and B
https://docs.oracle.com/cd/B19306_01/server.102/b14215/ldr_concepts.htm#sthref519

upvoted 1 times

jackymak 2 years, 4 months ago

Sorry, It is C & D.

upvoted 1 times

Ilse 2 years, 7 months ago

Answers: A & D

upvoted 1 times

Question #144 *Topic 1*

You are connected to a pluggable database (PDB) as a common user with DBA privileges.
The STATISTICS_LEVEL parameter is PDB_MODIFIABLE. You execute the following:
SQL > ALTER SYSTEM SET STATISTICS_LEVEL = ALL SID = * SCOPE = SPFILE;
Which is true about the result of this command?

A. The STATISTICS_LEVEL parameter is set to all whenever this PDB is re-opened.

B. The STATISTICS_LEVEL parameter is set to ALL whenever any PDB is reopened.

C. The STATISTICS_LEVEL parameter is set to all whenever the multitenant container database (CDB) is restarted.

D. Nothing happens; because there is no SPFILE for each PDB, the statement is ignored.

Correct Answer: *A*

Question #145 *Topic 1*

Which two are prerequisites for performing a flashback transaction? (Choose two.)

A. Flashback Database must be enabled.

B. Undo retention guarantee for the database must be configured.

C. EXECUTE privilege on the DBMS_FLASHBACK package must be granted to the user flashing back transaction.

D. Supplemental logging must be enabled.

E. Recycle bin must be enabled for the database.

F. Block change tracking must be enabled tor the database.

Correct Answer: *BD*

References:

http://searchoracle.techtarget.com/tip/How-to-perform-Oracle-Flashback-Transaction-Queries

https://docs.oracle.com/cd/E11882_01/appdev.112/e41502/adfns_flashback.htm#ADFNS610

Community vote distribution

CD (100%)

TuxBingo Highly Voted 4 years ago

C & D.

Oracle 12C Study Guide – Biju Thomas

Using Flashback Transaction
You can use the Flashback Transaction technology to undo a transaction and its dependent transactions. The DBMS_FLASHBACK.TRANSACTION_BACKOUT procedure is used to accomplish this task.
You must meet the following prerequisites to perform a Flashback Transaction on an Oracle Database 12c database:
■■ The database must be in ARCHIVELOG mode.
■■ Supplemental logging must be enabled in the database using ALTER DATABASE ADD SUPPLEMENTAL LOG DATA.
■■ A supplemental log data primary key should be created using the statement ALTER DATABASE ADD SUPPLEMENTAL LOG DATA (PRIMARY KEY) COLUMNS.
■■ The user performing the Flashback Transaction must have the SELECT ANY TRANSACTION privilege.
■■ The user should have the EXECUTE privilege on DBMS_FLASHBACK.
■■ The user should also have appropriate DML privileges on the tables (such as INSERT/UPDATE/DELETE).
Using EM Cloud Control, you can perform the Flashback Transaction. Choose Transactions from the Backup and Recovery menu under Availability.

upvoted 13 times

khan Highly Voted 5 years, 3 months ago

correct answers are CD.
https://docs.oracle.com/database/121/ADFNS/adfns_flashback.htm#ADFNS611

upvoted 9 times

luizyto Most Recent 1 year, 6 months ago

Selected Answer: CD

You can use the Flashback Transaction technology to undo a transaction and its dependent transactions. The DBMS_FLASHBACK.TRANSACTION_BACKOUT procedure is used to accomplish this task.

upvoted 1 times

masamma91 2 years, 4 months ago

C y D:
To allow access to the features in the DBMS_FLASHBACK package, grant the EXECUTE privilege on DBMS_FLASHBACK.
https://docs.oracle.com/cd/E11882_01/appdev.112/e41502/adfns_flashback.htm#ADFNS613

upvoted 1 times

Phoenix22 4 years, 4 months ago

I think that A have to be a MUST.
AC in my opinion.

upvoted 2 times

fsamkoh 4 years, 2 months ago

Should be A and C
upvoted 2 times

charan94 3 years, 3 months ago
A is not correct,
This question is related to flashbacking a transaction..
upvoted 2 times

Question #146 *Topic 1*

A database is stored in an Automatic Storage Management (ASM) disk group, disk group, DGROUP1 with SQL:

```
SQL> CREATE DISKGROUP dgroup1 NORMAL REDUNDANCY
   FAILGROUP controller1 DISK '/devices/diska1', '/devices/diska2'
   FAILGROUP controller2 DISK '/devices/diskb1', '/devices/diskb2';
```

There is enough free space in the disk group for mirroring to be done.
What happens if the CONTROLLER1 failure group becomes unavailable due to error of for maintenance?

A. Transactions and queries accessing database objects contained in any tablespace stored in DGROUP1 will fall.

B. Mirroring of allocation units will be done to ASM disks in the CONTROLLER2 failure group until the CONTROLLER1 for failure group is brought back online.

C. The data in the CONTROLLER1 failure group is copied to the controller2 failure group and rebalancing is initiated.

D. ASM does not mirror any data until the controller failure group is brought back online, and newly allocated primary allocation units (AU) are stored in the controller2 failure group, without mirroring.

E. Transactions accessing database objects contained in any tablespace stored in DGROUP1 will fail but queries will succeed.

Correct Answer: *D*

Ray520 2 years, 3 months ago
A. Transactions and queries accessing database objects contained in any tablespace stored in DGROUP1 will fall. ---Not true as NORMAL REDUNDANCY allows 2 way mirroring and diskgroup stays available even if one diskgroup fails.
C. The data in the CONTROLLER1 failure group is copied to the controller2 failure group and rebalancing is initiated.---Data mirrored is already available, rebalancing only evenly distributes file extents and space usage across all disks in the diskgroup.
E. Transactions accessing database objects contained in any tablespace stored in DGROUP1 will fail but queries will succeed. ----Not true (refer A)
upvoted 1 times

Ray520 2 years, 3 months ago
Disk loss in ASM can result from a number of reasons. To allow for sufficient time to recover from disk failures that do not involve the actual failu of a disk, ASM provides the ASM fast disk resync feature. By default, when a disk in an ASM disk group fails the disk will be taken offline automatically. The disk will be dropped some 3.6 hours later. As a result, you have only 3.6 hours by default to respond to a disk outage. If you correct the problem and the physical disk media is not corrupted, then ASM fast disk resync will quickly re-synchronize the disk when it comes ba online, correcting the problem very quickly. We an change the amount of time that Oracle will wait to automatically drop the disk by setting the disk_repair_time attribute.
upvoted 1 times

Question #147 *Topic 1*

On your Oracle 12c database, you Issue the following commands to create indexes
SQL > CREATE INDEX oe.ord_customer_ix1 ON oe.orders (customers_id, sales_rep_id) INVISIBLE;
SQL> CREATE BITMAP INDEX oe.ord_customer_ix2 ON oe.orders (customers_id, sales_rep_id);
Which two statements are correct? (Choose two.)

A. Both the indexes are created; however, only the ORD_COSTOMER index is visible.

B. The optimizer evaluates index access from both the Indexes before deciding on which index to use for query execution plan.

C. Only the ORD_CUSTOMER_IX1 index is created.

D. Only the ORD_CUSTOMER_IX2 index is created.

E. Both the indexes are updated when a new row is inserted, updated, or deleted In the orders table.

Correct Answer: *AE*
11G has a new feature called Invisible Indexes. An invisible index is invisible to the optimizer as default. Using this feature we can test a new index without effecting the execution plans of the existing sql statements or we can test the effect of dropping an index without dropping it.

khan Highly Voted 5 years, 3 months ago
The answer is right: A E.

However, the wording of option A should be:
"Both the indexes are created; however, only the ORD_COSTOMER_IX2 index is visible."
upvoted 9 times

jackymak 2 years, 5 months ago
Invisible index will be considered for optimizer
upvoted 1 times

jackymak 2 years, 4 months ago
I mean will NOT be
upvoted 1 times

5kong Most Recent 4 years, 4 months ago
Question #118
upvoted 2 times

Question #148 *Topic 1*

Your multitenant container database has three pluggable databases (PDBs): PDB1, PDB2, and PDB3.
Which two RMAN commands may be; used to back up only the PDB1 pluggable database? (Choose two.)

A. BACKUP PLUGGABLE DATABASE PDB1 while connected to the root container

B. BACKUP PLUGGABLE DATABASE PDB1 while connected to the PDB1 container

C. BACKUP DATABASE while connected to the PDB1 container

D. BACKUP DATABASE while connected to the boot container

E. BACKUP PLUGGABLE database PDB1 while connected to PDB2

Correct Answer: *AC*
To perform operations on a single PDB, you can connect as target either to the root or directly to the PDB.
* (A) If you connect to the root, you must use the PLUGGABLE DATABASE syntax in your RMAN commands. For example, to back up a PDB, you use the
BACKUP PLUGGABLE DATABASE command.
* (C)If instead you connect directly to a PDB, you can use the same commands that you would use when connecting to a non-CDB. For example, to back up a
PDB, you would use the BACKUP DATABASE command.
References:

Question #149 *Topic 1*

Identify three benefits of Unified Auditing.

A. Decreased use of storage to store audit trail rows in the database.

B. It improves overall auditing performance.

C. It guarantees zero-loss auditing.

D. The audit trail cannot be easily modified because it is read-only.

E. It automatically audits Recovery Manager (RMAN) events.

Correct Answer: *ABE*

A: Starting with 12c, Oracle has unified all of the auditing types into one single unit called Unified auditing. You dont have to turn on or off all of the different auidting types individually and as a matter of fact auditing is enabled by default right out of the box. The AUD$ and FGA$ tables have been replaced with one single audit trail table. All of the audit data is now stored in Secure Files table thus improving the overall management aspects of audit data itself.

B: Further the audit data can also be buffered solving most of the common performance related problems seen on busy environments.

E: Unified Auditing is able to collect audit data for Fine Grained Audit, RMAN, Data Pump, Label Security, Database Vault and Real Application Security operations.

Note:

* Benefits of the Unified Audit Trail

The benefits of a unified audit trail are many:

/ (B) Overall auditing performance is greatly improved. The default mode that unified audit works is Queued Write mode. In this mode, the audit records are batched in SGA queue and is persisted in a periodic way. Because the audit records are written to SGA queue, there is a significant performance improvement.

/ The unified auditing functionality is always enabled and does not depend on the initialization parameters that were used in previous releases

/ (A) The audit records, including records from the SYS audit trail, for all the audited components of your Oracle Database installation are placed in one location and in one format, rather than your having to look in different places to find audit trails in varying formats. This consolidated view enables auditors to co-relate audit information from different components. For example, if an error occurred during an INSERT statement, standard auditing can indicate the error number and the SQL that was executed. Oracle Database Vault-specific information can indicate whether this error happened because of a command rule violation or realm violation. Note that there will be two audit records with a distinct AUDIT_TYPE. With this unification in place, SYS audit records appear with AUDIT_TYPE set to
Standard Audit.

/ The management and security of the audit trail is also improved by having it in single audit trail.

/ You can create named audit policies that enable you to audit the supported components listed at the beginning of this section, as well as SYS administrative users. Furthermore, you can build conditions and exclusions into your policies.

* Oracle Database 12c Unified Auditing enables selective and effective auditing inside the Oracle database using policies and conditions. The new policy based syntax simplifies management of auditing within the database and provides the ability to accelerate auditing based on conditions.

* The new architecture unifies the existing audit trails into a single audit trail, enabling simplified management and increasing the security of audit data generated by the database.

Community vote distribution

BDE (100%)

Ray520 2 years ago

Selected Answer: BDE

Unified auditing:
-Improved performance: On Oracle Database 12c, with Unified Auditing and Conditional Auditing, you get the ability to configure precise, contex dependant logging which should reduce the performance overhead associated with database auditing
-Audit operations of all RDBMS and other components like RMAN, Oracle Data Pump using the new 12c Unified Auditing feature, consolidating a audit trails into a single unified audit trail table
-The unified audit trail resides in a read-only table in the AUDSYS schema in the SYSAUX tablespace

upvoted 1 times

Ray520 2 years, 3 months ago

Unified auditing:
-Improved performance: On Oracle Database 12c, with Unified Auditing and Conditional Auditing, you get the ability to configure precise, contex dependant logging which should reduce the performance overhead associated with database auditing and enable more effective analysis of audi logs.

-As part of the unified audit trail enhancement, a new schema, AUDSYS, will be used solely for storage of the unified audit trail data table.The existing audit data in the AUD$ and FGA_LOG$ system tables, audit metadata, and audit PL/SQL packages, will continue to reside in the SYS schema.
-Audit operations of all RDBMS and other components like RMAN, Oracle Data Pump using the new 12c Unified Auditing feature, consolidating a audit trails into a single unified audit trail table
-The unified audit trail resides in a read-only table in the AUDSYS schema in the SYSAUX tablespace

upvoted 1 times

duMaurier 3 years, 4 months ago

BDE confirmed,
https://docs.oracle.com/database/121/DBSEG/auditing.htm#DBSEG630
check section "What Is Unified Auditing" and "Benefits of the Unified Audit Trail"

upvoted 2 times

TuxBingo 4 years ago

I agree, the correct answer is:
B:https://www.vitalsofttech.com/overview-of-unified-auditing-in-oracle-database-12c/
D: https://blogs.oracle.com/imc/entry/oracle_database_12c_new_unified
E: http://www.oracle.com/webfolder/technetwork/tutorials/obe/db/12c/r1/security/sec_uni_audit/sec_uni_audit.html

upvoted 1 times

Mohit201920 4 years, 4 months ago

BDE CORRECT ANSWER

upvoted 4 times

hggz 4 years, 4 months ago

I agree

upvoted 2 times

joe_doe 4 years, 1 month ago

I stand by that

upvoted 1 times

Question #150 *Topic 1*

You upgraded from a previous Oracle database version to Oracle Database version to Oracle Database 12c. Your database supports a mixed workload. During the day, lots of insert, update, and delete operations are performed. At night, Extract, Transform, Load (ETL) and batch reporting jobs are run. The ETL jobs perform certain database operations using two or more concurrent sessions.
After the upgrade, you notice that the performance of ETL jobs has degraded. To ascertain the cause of performance degradation, you want to collect basic statistics such as the level of parallelism, total database time, and the number of I/O requests for the ETL jobs.
How do you accomplish this?

A. Examine the Active Session History (ASH) reports for the time period of the ETL or batch reporting runs.

B. Enable SQL tracing for the queries in the ETL and batch reporting queries and gather diagnostic data from the trace file.

C. Enable real-time SQL monitoring for ETL jobs and gather diagnostic data from the V$SQL_MONITOR view.

D. Enable real-time database operation monitoring using the DBMS_SQL_MONITOR.BEGIN_OPERATION function, and then use the

Correct Answer: *D*
* Monitoring database operations
Real-Time Database Operations Monitoring enables you to monitor long running database tasks such as batch jobs, scheduler jobs, and Extraction,
Transformation, and Loading (ETL) jobs as a composite business operation. This feature tracks the progress of SQL and PL/SQL queries associated with the business operation being monitored. As a DBA or developer, you can define business operations for monitoring by explicitly specifying the start and end of the operation or implicitly with tags that identify the operation.

khan Highly Voted 5 years, 3 months ago
Answer D
More Info..
Enable real-time database operation monitoring using the DBMS_SQL_MONITOR.BEGIN_OPERATION function, and then use the DBMS_SQL_MONITOR.REPORT_SQL_MONITOR function to view the required information.
upvoted 5 times

Question #151 *Topic 1*

Your multitenant container (CDB) contains two pluggable databases (PDB), HR_PDB and ACCOUNTS_PDB, both of which use the CDB tablespace. The temp file is called temp01.tmp.
A user issues a query on a table on one of the PDBs and receives the following error:
ERROR at line 1:
ORA-01565: error in identifying file /u01/app/oracle/oradata/CDB1/temp01.tmp
ORA-27037: unable to obtain file status
Identify two ways to rectify the error.

A. Add a new temp file to the temporary tablespace and drop the temp file that that produced the error.

B. Shut down the database instance, restore the temp01.tmp file from the backup, and then restart the database.

C. Take the temporary tablespace offline, recover the missing temp file by applying redo logs, and then bring the temporary tablespace online.

D. Shutdown the database instance, restore and recover the temp file from the backup, and then open the database with RESETLOGS.

E. Shut down the database instance and then restart the CDB and PDBs.

Correct Answer: *AE*
* Because temp files cannot be backed up and because no redo is ever generated for them, RMAN never restores or recovers temp files. RMAN does track the names of temp files, but only so that it can automatically re-create them when needed.
* If you use RMAN in a Data Guard environment, then RMAN transparently converts primary control files to standby control files and vice versa. RMAN automatically updates file names for data files, online redo logs, standby redo logs, and temp files when you issue RESTORE and RECOVER.

Question #152 *Topic 1*

Examine the following commands for redefining a table with Virtual Private Database (VPD) policies:

```
BEGIN
  DBMS_RLS.ADD_POLICY (
    object_schema     => 'hr',
    object_name       => 'employees',
    policy_name       => 'employees_policy',
    function_schema   => 'hr',
    policy_function   => 'auth_emp_dep_100',
    statement_types   => 'select, insert, update, delete'
  );
 END;

BEGIN
  DBMS_REDEFINITION.START_REDEF_TABLE (
    uname             => 'hr',
    orig_table        => 'employees',
    int_table         => 'int_employees',
    col_mapping       => NULL,
    options_flag      => DBMS_REDEFINITION.CONS_USE_PK,
    orderby_cols      => NULL,
    part_name         => NULL,
    copy_vpd_opt      => DBMS_REDEFINITION.CONS_VPD_AUTO);
END;
```

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

A. All the triggers for the table are disabled without changing any of the column names or column types in the table.

B. The primary key constraint on the EMPLOYEES table is disabled during redefinition.

C. VPD policies are copied from the original table to the new table during online redefinition.

D. You must copy the VPD policies manually from the original table to the new table during online redefinition.

Correct Answer: *BC*

C (not D): CONS_VPD_AUTO -

Used to indicate to copy VPD policies automatically

* DBMS_RLS.ADD_POLICY

/ The DBMS_RLS package contains the fine-grained access control administrative interface, which is used to implement Virtual Private Database

(VPD).DBMS_RLS is available with the Enterprise Edition only.

Note:

* CONS_USE_PK and CONS_USE_ROWID are constants used as input to the "options_flag" parameter in both the START_REDEF_TABLE Procedure and

CAN_REDEF_TABLE Procedure. CONS_USE_ROWID is used to indicate that the redefinition should be done using rowids while CONS_USE_PK implies that the redefinition should be done using primary keys or pseudo-primary keys (which are unique keys with all component columns having NOT NULL constraints).

* DBMS_REDEFINITION.START_REDEF_TABLE

To achieve online redefinition, incrementally maintainable local materialized views are used. These logs keep track of the changes to the master tables and are used by the materialized views during refresh synchronization.

* START_REDEF_TABLE Procedure

Prior to calling this procedure, you must manually create an empty interim table (in the same schema as the table to be redefined) with the desired attributes of the post-redefinition table, and then call this procedure to initiate the redefinition.

sherifelia Highly Voted 4 years, 8 months ago

Answer: AC

upvoted 7 times

khan Highly Voted 5 years, 3 months ago

Answer: AC

https://docs.oracle.com/database/121/ARPLS/d_redefi.htm#ARPLS67515

A. All the triggers of this tables will be disables
B. Primary key of this tables will not be disabled.

upvoted 6 times

Ray520 2 years, 3 months ago

The triggers of interim table stays disabled while redifinition in progress, not triggers of table being redefined. Nothing is said in document about table triggers being disabled. :

All cloned referential constraints involving the interim tables will be created disabled (they will be automatically enabled after the redefinition) and all triggers on interim tables will not fire till the redefinition is completed. After the redefinition is complete, the cloned objects will be renamed to the corresponding pre-redefinition names of the objects (from which they were cloned from).

upvoted 1 times

Question #153 *Topic 1*

Which two statements are true about the use of the procedures listed in the v$sysaux_occupants.move_procedure column? (Choose two.)

A. The procedure may be used for some components to relocate component data to the SYSAUX tablespace from its current tablespace.

B. The procedure may be used for some components to relocate component data from the SYSAUX tablespace to another tablespace.

C. All the components may be moved into SYSAUX tablespace.

D. All the components may be moved from the SYSAUX tablespace.

Correct Answer: *AB*

TuxBingo 4 years ago

Answer AB is correct: https://docs.oracle.com/database/121/ARPLS/d_redefi.htm#ARPLS67515

upvoted 2 times

Question #154 *Topic 1*

Which statement is true about Oracle Net Listener?

A. It acts as the listening endpoint for the Oracle database instance for all local and non-local user connections.

B. A single listener can service only one database instance and multiple remote client connections.

C. Service registration with the listener is performed by the process monitor (PMON) process of each database instance.

D. The listener.ora configuration file must be configured with one or more listening protocol addresses to allow remote users to connect to a database instance.

E. The listener.ora configuration file must be located in the ORACLE_HOME/network/admin directly.

Correct Answer: *C*

https://docs.oracle.com/database/121/CNCPT/process.htm

Community vote distribution

C (100%)

Havok Highly Voted 4 years, 6 months ago

C would not be correct for 12c, since PMON handled this process in earlier releases, but for 12c it is the LREG that handles service registration

upvoted 8 times

sherifelia Highly Voted 4 years, 8 months ago

D is correct
http://www.dba86.com/docs/oracle/11.2/network.112/e41945/listenercfg.htm#NETAG292

upvoted 5 times

luizyto Most Recent 1 year, 6 months ago

Selected Answer: C

I thing that this could be and old question and in exam 12c we will se the LREG service instead.

upvoted 1 times

Ray520 2 years, 3 months ago

Since the version of database is not mentioned, we can only consider option C as it's correct for versions before 12c. Rest all options are incorrect
"In releases before Oracle Database 12c, PMON performed the listener registration."

upvoted 1 times

Ray520 2 years, 3 months ago

When confused which option is correct, better to find out the incorrect options. Agree with @hggz

upvoted 1 times

fridaytar 3 years, 5 months ago

The listener forwards client requests to supported services. These services are dynamically registered with the listener. This dynamic registration feature is called service registration. The registration is performed by the Listener Registration (LREG) process. Dynamic service registration does not require any manual configuration in the listener.ora file.

upvoted 1 times

TuxBingo 4 years ago

Agree with Havok:
Listener Registration Process (LREG)

The listener registration process (LREG) registers information about the database instance and dispatcher processes with the Oracle Net Listener.

When an instance starts, LREG polls the listener to determine whether it is running. If the listener is running, then LREG passes it relevant parameters. If it is not running, then LREG periodically attempts to contact it.

Note:

In releases before Oracle Database 12c, PMON performed the listener registration.

upvoted 1 times

joe_doe 4 years, 1 month ago

Answer should be D， LREG do the register work instead of PMON in 12c

upvoted 3 times

Winterschill 4 years, 2 months ago

The Answer is 100% D guys...
Havok is correct on why C is WRONG.
Reference below will confirm LREG handles that proces.

It also confirms why D is correct.
https://docs.oracle.com/database/121/NETAG/listenercfg.htm#NETAG292

upvoted 5 times

hggz 4 years, 4 months ago

Answer is C.
But option C is a little bit different:
C. Service registration with the listener is performed by the listener registration process monitor (LREG) process of each database instance.

Not D.

upvoted 1 times

khan 5 years, 3 months ago

C. Service registration with the listener is performed by the listener registration process monitor (LREG) process of each database instance.

Reference:
http://docs.oracle.com/database/121/NETAG/listenercfg.htm

upvoted 1 times

hggz 4 years, 4 months ago

A is wrong because you can make local connections without going through a listener.
B is wrong because many databases can register with a listener.
D is wrong because you don't have to configure a listening protocol address(es) in istener.ora. If not specified, it will simply default to TCP por 1521.
E is wrong because Oracle will also look in $TNS_ADMIN directory or /etc.

upvoted 3 times

Question #155 *Topic 1*

You are administering a database stored in Automatic Storage Management (ASM). You use RMAN to back up the database and the MD_BACKUP command to back up the ASM metadata regularly. You lost an ASM disk group DG1 due to hardware failure.
In which three ways can you re-create the lost disk group and restore the data? (Choose three.)

A. Use the MD_RESTORE command to restore metadata for an existing disk group by passing the existing disk group name as an input parameter and use RMAN to restore the data.

B. Use the MKDG command to restore the disk group with the same configuration as the backed-up disk group and data on the disk group.

C. Use the MD_RESTORE command to restore the disk group with the changed disk group specification, failure group specification, name, and other attributes and use RMAN to restore the data.

D. Use the MKDG command to restore the disk group with the same configuration as the backed-up disk group name and same set of disks and failure group configuration, and use RMAN to restore the data.

E. Use the MD_RESTORE command to restore both the metadata and data for the failed disk group.

F. Use the MKDG command to add a new disk group DG1 with the same or different specifications for failure group and other attributes and use RMAN to restore

Correct Answer: *ACF*
AC (not E):
The md_restore command allows you to restore a disk group from the metadata created by the md_backup command. md_restore cant restore data, only metadata.

Ray520 2 years, 3 months ago
MKDG command Creates a disk group based on an XML configuration file hat specifies the name of the disk group, redundancy, attributes, and paths of the disks that form the disk group.

none of the given MGDG options are correct.
upvoted 1 times

Ray520 2 years, 3 months ago
https://docs.oracle.com/en/database/oracle/oracle-database/12.2/ostmg/asmcmd-diskgroup-commands.html#GUID-241C6676-FA7C-4505-ABA
BD45657CF4AA

md_restore restores Oracle ACFS metadata information.
To run the metadata restore for Oracle ACFS file systems, a script is created. This script must be run as root to format file systems.
Metadata for the contained Oracle ADVM volumes, such as stripe size and redundancy, is restored. Oracle ACFS file system configuration metada
such as compression and resize attributes, block size, and created snapshots, is also restored.
Oracle ACFS file System data and ACFS CRS Resource information is not restored.
When restoring Oracle ACFS encryption or security, file system extended attributes (xattrs) must be restored with the user data.
If using the --silent option while specifying a single disk group, and there is a dependency due to Oracle ADVM metadata or Oracle ACFS accelerators on another disk group, that dependency is dropped, and you must manually recreate that dependency later.
upvoted 1 times

Question #156 *Topic 1*

Your multitenant container database, CDB1, is running in ARCHIVELOG mode and has two pluggable databases, HR_PDB and ACCOUNTS_PDB. An RMAN backup exists for the database.
You issue the command to open ACCOUNTS_PDB and find that the USERDATA.DBF data file for the default permanent tablespace USERDATA belonging to
ACCOUNTS_PDB is corrupted.
What should you do before executing the commands to restore and recover the data file in ACCOUNTS_PDB?

A. Place CDB1 in the mount stage and then the USERDATA tablespace offline in ACCOUNTS_PDB.

B. Place CDB1 in the mount stage and issue the ALTER PLUGGABLE DATABASE accounts_pdb CLOSE IMMEDIATE command.

C. Issue the ALTER PLUGGABLE DATABASE accounts_pdb RESTRICTED command.

D. Take the USERDATA tablespace offline in ACCOUNTS_PDB.

Correct Answer: *D*
* You can take an online tablespace offline so that it is temporarily unavailable for general use. The rest of the database remains open and available for users to access data. Conversely, you can bring an offline tablespace online to make the schema objects within the tablespace available to database users. The database must be open to alter the availability of a tablespace.

Community vote distribution
B (100%)

jackymak 2 years, 5 months ago

Selected Answer: B
https://www.oracle.com/webfolder/technetwork/tutorials/obe/db/12c/r1/pdb/pdb_pitr/pdb_pitr.html#t4
upvoted 1 times

jackymak 2 years, 5 months ago

sorry, It should be D for restore table space.
https://oracle-base.com/articles/12c/multitenant-rman-backup-recovery-cdb-and-pdb-12cr1#tablespace-and-datafile-recovery
upvoted 1 times

khan 5 years, 3 months ago

Thought of sharing. In the Exam for this question answer option was different.
To select one of this.
ALTER TABLESPACE users OFFLINE IMMEDIATE;
ALTER TABLESPACE users OFFLINE TEMPORARY;
ALTER TABLESPACE users OFFLINE NORMAL;

https://docs.oracle.com/cd/B28359_01/server.111/b28310/tspaces005.htm#ADMIN11375
upvoted 3 times

Havok 4 years, 6 months ago

Would this tablespace be taken OFFLINE IMMEDIATE? since it is in ARCHIVELOG mode, and it was offline due to error.
upvoted 2 times

Khel 2 years, 10 months ago

@Khan, I believe the exam answer was B. Alter TABLESPACE users OFFLINE TEMPORARY
upvoted 2 times

Question #157 *Topic 1*

Which Oracle Database component is audited by default if the unified Auditing option is enabled?

A. Oracle Data Pump

B. Oracle Recovery Manager (RMAN)

C. Oracle Label Security

D. Oracle Database Vault

E. Oracle Real Application Security

Correct Answer: *B*

Question #158 *Topic 1*

Your multitenant container (CDB) containing three pluggable databases (PDBs) is running in ARCHIVELOG mode. You find that the SYSAUX tablespace is corrupted in the root container.
The steps to recover the tablespace are as follows:
1. Mount the CDB.
2. Close all the PDBs.
3. Open the database.
4. Apply the archive redo logs.
5. Restore the data file.
6. Take the SYSAUX tablespace offline.
7. Place the SYSAUX tablespace online.
8. Open all the PDBs with RESETLOGS.
9. Open the database with RESETLOGS.
10. Execute the command SHUTDOWN ABORT.
Which option identifies the correct sequence to recover the SYSAUX tablespace?

A. 6, 5, 4, 7

B. 10, 1, 2, 5, 8

C. 10, 1, 2, 5, 4, 9, 8

D. 10, 1, 5, 8, 10

Correct Answer: *A*
RMAN> ALTER TABLESPACE sysaux OFFLINE IMMEDIATE;
RMAN> RESTORE TABLESPACE sysaux;
RMAN> RECOVER TABLESPACE sysaux;
RMAN> ALTER TABLESPACE sysaux ONLINE;
* Example:
While evaluating the 12c beta3 I was not able to do the recover while testing "all pdb files lost".
Cannot close the pdb as the system datafile was missing
So only option to recover was:
Shutdown cdb (10)
startup mount; (1)
restore pluggable database
recover pluggable database
alter database open;
alter pluggable database name open;
Oracle support says: You should be able to close the pdb and restore/recover the system tablespace of PDB.
* Inconsistent backups are usually created by taking online database backups. You can also make an inconsistent backup by backing up data files while a database is closed, either:
/ Immediately after the crash of an Oracle instance (or, in an Oracle RAC configuration, all instances)
/ After shutting down the database using SHUTDOWN ABORT
Inconsistent backups are only useful if the database is in ARCHIVELOG mode and all archived redo logs created since the backup are available.
* Open the database with the RESETLOGS option after finishing recovery:
SQL> ALTER DATABASE OPEN RESETLOGS;

Currently there are no comments in this discussion, be the first to comment!

Question #159 *Topic 1*

Which three are direct benefits of the multiprocess, multithreaded architecture of Oracle Database 12c when it is enabled? (Choose three.)

A. Reduced logical I/O

B. Reduced virtual memory utilization

C. Improved parallel Execution performance

D. Improved Serial Execution performance

E. Reduced physical I/O

F. Reduced CPU utilization

Correct Answer: *BCF*
* Multiprocess and Multithreaded Oracle Database Systems
Multiprocess Oracle Database (also called multiuser Oracle Database) uses several processes to run different parts of the Oracle Database code and additional
Oracle processes for the userseither one process for each connected user or one or more processes shared by multiple users. Most databases are multiuser because a primary advantage of a database is managing data needed by multiple users simultaneously.
Each process in a database instance performs a specific job. By dividing the work of the database and applications into several processes, multiple users and applications can connect to an instance simultaneously while the system gives good performance.
* In previous releases, Oracle processes did not run as threads on UNIX and Linux systems. Starting in Oracle Database 12c, the multithreaded Oracle Database model enables Oracle processes to execute as operating system threads in separate address spaces.

Question #160 *Topic 1*

In order to exploit some new storage tiers that have been provisioned by a storage administrator, the partitions of a large heap table must be moved to other tablespaces in your Oracle 12c database?
Both local and global partitioned B-tree Indexes are defined on the table.
A high volume of transactions access the table during the day and a medium volume of transactions access it at night and during weekends.
Minimal disrupt ion to availability is required.
Which three statements are true about this requirement? (Choose three.)

A. The partitions can be moved online to new tablespaces.

B. Global indexes must be rebuilt manually after moving the partitions.

C. The partitions can be compressed in the same tablespaces.

D. The partitions can be compressed in the new tablespaces.

E. Local indexes must be rebuilt manually after moving the partitions.

Correct Answer: *ACD*
A: You can create and rebuild indexes online. Therefore, you can update base tables at the same time you are building or rebuilding indexes on that table. You can perform DML operations while the index build is taking place, but DDL operations are not allowed. Parallel execution is not supported when creating or rebuilding an index online.
D: Moving (Rebuilding) Index-Organized Tables
Because index-organized tables are primarily stored in a B-tree index, you can encounter fragmentation as a consequence of incremental updates. However, you can use the ALTER TABLE...MOVE statement to rebuild the index and reduce this fragmentation.
C: If a table can be compressed in the new tablespace, also it can be compressed in the same tablespace.
Incorrect:
Not B, not E: Local and Global indexes can be automatically rebuild with UPDATE INDEXES when you move the table.

Question #161 *Topic 1*

Which two statements are true about the Automatic Database Diagnostic Monitor (ADDM)? (Choose two.)

A. The ADDM requires at least four AWR snapshots for analysis

B. The ADDM runs after each AWR snapshot is collected automatically by MMON

C. The results of the ADDM analysis are stored in the Automatic Workload Repository (AWR)

D. The ADDM analysis provides only diagnostics information but does not provide recommendations

E. The ADDM calls other advisors if required, but does not provide recommendations about the advisors

Correct Answer: *BC*

charan94 3 years, 3 months ago

B,C is correct
Since the snapshot means the report pulled between two snap ids

upvoted 1 times

NorthKorean 4 years, 4 months ago

Isn't B wrong ? ADDM works once for two snapshots

upvoted 1 times

joe_doe 4 years, 1 month ago

it's wrong for 1st snapshot,ADDM will compare 1st & 2nd\ 2nd & 3rd .....n\n+1,so it's for every snapshots,but not for 1st snapshot only

upvoted 2 times

Question #162 *Topic 1*

In your production database, data manipulation language (DML) operations are executed on the SALES table.
You have noticed some dubious values in the SALES table during the last few days. You are able to track users, actions taken, and the time of the action for this particular period but the changes in data are not tracked. You decide to keep track of both the old data and new data in the table long with the user information.
What action would you take to achieve this task?

A. Apply fine-grained auditing.

B. Implement value-based auditing.

C. Impose standard database auditing to audit object privileges.

D. Impose standard database auditing to audit SQL statements.

Correct Answer: *B*

TuxBingo Highly Voted 4 years ago

One tip for this question for understand better is: They need track the data changes with old and new values, and we can do only with "value-based-auditing" check this link.
http://www.ocptechnology.com/value-based-auditing/

Correct answer is B

upvoted 7 times

JMAN1 2 years, 6 months ago

Hi. TuxBingo Thank you for helping me understanding questions's answer.

upvoted 1 times

jackymak 2 years, 5 months ago

http://oraclebykkbakshi.blogspot.com/2017/08/value-based-auditing-in-oracle-11g.html
It like a trigger to record all the change in another table.

upvoted 1 times

Question #163 *Topic 1*

The user SCOTT owns the CUST table that is placed in the SALES tablespace. The user SCOTT opens a session and executes commands as follows:
SQL> INSERT INTO cust VALUES(101, 'JACK');
1 row created.
SQL> INSERT INTO cust VALUES(102, 'SMITH');
1 row created.
As a DBA, you execute the following command from another session:
ALTER TABLESPACE sales READ ONLY;
Which statement is true regarding the effect of this command on the transaction in Scott's session?

A. The command fails as a transaction is still pending.

B. The transaction in Scott's session is rolled back and the tablespace becomes readonly.

C. The command waits and the user SCOTT can execute data manipulation language (DML) statements only as part of the current transaction.

D. The command hangs until all transactions on the objects in the tablespace commit or rollback, and then the tablespace is placed in readonly mode.

Correct Answer: *B*

Community vote distribution

D (100%)

NorthKorean Highly Voted 4 years, 4 months ago

D is ,surely

upvoted 9 times

khan Highly Voted 5 years, 3 months ago

Answer D.
The command hangs until all transactions on the objects in the tablespace commit or rollback, and then the tablespace is placed in readonly moc

upvoted 7 times

luizyto Most Recent 1 year, 8 months ago

Selected Answer: D

specify READ ONLY to place the tablespace in transition read-only mode. In this state, existing transactions can complete (commit or roll back), bu no further DML operations are allowed to the tablespace except for rollback of existing transactions that previously modified blocks in the tablespace.

upvoted 1 times

jackymak 2 years, 5 months ago

Selected Answer: D

it should have behavior like shutdown normail

upvoted 1 times

duMaurier 2 years, 7 months ago

Selected Answer: D

Tested the scenario it should be D.

upvoted 2 times

Clozoerie 2 years, 12 months ago

Answer is D

upvoted 2 times

TuxBingo 4 years ago

It is D !!!
Test it yourself, you will get a ORA-00372 for every further DML that modifys cust table and the transaction is not terminated nor rolled back automatically.

upvoted 3 times

sherifelia 4 years, 8 months ago

Answer is D

upvoted 6 times

PerkDizzzle 4 years, 10 months ago

"Which statement is true regarding the effect of this command on the transaction in Scott's session?"

Answer is correctly B, it's asking specifically about Scott's session.
upvoted 3 times

loveroba 3 years, 8 months ago
right,checked
upvoted 2 times

charan94 3 years, 3 months ago
Even it is scott's session, The answer is D only..
upvoted 2 times

Question #164 *Topic 1*

You plan to implement the distributed database system in your company. You invoke Database Configuration Assistant (DBCA) to create a database on the server. During the installation, DBCA prompts you to specify the Global Database Name.
What must this name be made up of?

A. It must be made up of a database name and a domain name.

B. It must be made up of the value in ORACLE_SID and HOSTNAME.

C. It must be made up of the value that you plan to assign for INSTANCE_NAME and HOSTNAME.

D. It must be made up of the value that you plan to assign for ORACLE_SID and SERVICE_NAMES.

Correct Answer: *A*
Using the DBCA to Create a Database (continued)
3. Database Identification: Enter the Global Database Name in The form database_name.domain_name, and the system identifier (SID). The SID defaults lo the database name and uniquely identifies the instance associated with the database.
4. Management Options: Use this page to set up your database so that it can be managed with Oracle Enterprise Manager. Select the default: "Configure the
Database with Enterprise Manager." Optionally, this page allows you to configure alert notifications and daily disk backup area settings.
Note: Yon must configure the listener before you can configure Enterprise Manager (as shown earlier).

Question #165 *Topic 1*

Which two statements are true about standard database auditing? (Choose two.)

A. DDL statements can be audited.

B. Statements that refer to standalone procedure can be audited.

C. Operations by the users logged on as SYSDBA cannot be audited.

D. Only one audit record is ever created for a session per audited statement even though it is executed more than once.

Correct Answer: *AB*

TuxBingo 4 years ago

Here the explanation of why is A and B:

ydisconzi
November 5, 2019 at 3:38 pm

A: true

The SQL statements that you can audit are in the following categories:
DDL statements. For example, enabling the auditing of tables (AUDIT TABLE) audits all CREATE and DROP TABLE statements

https://docs.oracle.com/cd/B28359_01/server.111/b28337/tdpsg_auditing.htm#TDPSG50523

B: true
SQL statements inside PL/SQL program units are individually audited, as necessary, when the program unit is executed.
https://docs.oracle.com/cd/B19306_01/network.102/b14266/auditing.htm#i1008362

upvoted 2 times

Question #166 *Topic 1*

You executed the following command to create a password file in the database server:
$ orapwd file = orapworcl entries = 5 ignorecase=N
Which statement describes the purpose of the above password file?

A. It records usernames and passwords of users when granted the DBA role

B. It contains usernames and passwords of users for whom auditing is enabled

C. It is used by Oracle to authenticate users for remote database administrator

D. It records usernames and passwords of all users when they are added to OSDBA or OSOPER operating groups

Correct Answer: *A*

Community vote distribution

C (100%)

rdnkrkmz Highly Voted 5 years, 6 months ago

Answer should be C or
In A option if the word DBA -> SYSDBA changed, then A can be alse true

upvoted 14 times

sherifelia Highly Voted 4 years, 8 months ago

Answer is C

upvoted 9 times

JMAN1 Most Recent 2 years, 6 months ago

Selected Answer: C

Answer C. same question .

upvoted 1 times

charan94 3 years, 3 months ago

Answer is C
Since the password file usage comes in to picture only when the remote users with sysdba..etc priviliges try to connect to database...

upvoted 1 times

charan94 3 years, 3 months ago

Answer is C
Since the password file usage comes in to picture only when the remote users with sysdba..etc priviliges try to connect to database...

upvoted 1 times

Question #167 *Topic 1*

Which three statements are true about space usage alerts? (Choose three.)

A. Alerts are issued only when the critical threshold for space available in a tablespace is breached.

B. The sum of active extents and allocated user quotas is considered to compute space usage for an undo tablespace.

C. Database alerts can provide warnings about low space availability at both tablespace and segment levels.

D. Alerts are not issued for locally managed tablespaces that are offline or in read-only mode.

E. A newly created locally managed tablespace is automatically assigned the default threshold values defined for a database.

Correct Answer: *ADE*
References:
https://docs.oracle.com/cd/B28359_01/server.111/b28310/schema001.htm#ADMIN10120

khan Highly Voted 5 years, 3 months ago
Correct Answer: CDE
https://docs.oracle.com/database/121/ADMIN/schema.htm#ADMIN014
upvoted 7 times

PerkDizzzle Highly Voted 4 years, 10 months ago
CDE correct answers
upvoted 7 times

Ray520 Most Recent 2 years, 3 months ago
Alerts can be 2 types: By space used (%) or By free space (MB). So option A is incorrect.
You can reclaim unused space. Segment Advisor, is an Oracle Database component that identifies segments that have space available for reclamation.
upvoted 1 times

Question #168 *Topic 1*

Which three statements are true about Oracle Data Pump? (Choose three.)

A. IMPDP can be used to change target data file names, schemas, and tablespaces during import.

B. The DBMS_DATAPUMP PL/SQL package can be used independently of Data Pump clients to perform export and import operations.

C. EXPDP and IMPDP are the client components of Oracle Data Pump.

D. Oracle Data Pump export and import operations can be performed only by users with the SYSDBA privilege.

E. IMPDP always use the conventional path insert method to import data.

Correct Answer: *ABC*
References:
https://docs.oracle.com/cd/E11882_01/server.112/e22490/dp_overview.htm#SUTIL2880

Question #169 *Topic 1*

Your database instance has started using an SPFILE.

Examine the RMAN configuration settings:

```
CONFIGURE RETENTION POLICY TO REDUNDANCY 1; # default
CONFIGURE BACKUP OPTIMIZATION OFF; # default
CONFIGURE DEFAULT DEVICE TYPE TO DISK; # default
CONFIGURE CONTROLFILE AUTOBACKUP ON;
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '%F'; default
```

You execute the command:

RMAN> BACKUP AS COPY TABLESPACE TEST;

Which three types of files are backed up by using this command? (Choose three.)

A. online redo log files

B. control file

C. SPFILE

D. archived redo log files

E. data file(s)

F. PFILE

Correct Answer: *BCE*

References: http://www.juliandyke.com/Research/RMAN/BackupCommand.php

TuxBingo Highly Voted 4 years ago

Correct answer is BCE: if you see the previous configuration of RMAN, the controlfile autobackup is ON so:
If the CONFIGURE CONTROLFILE AUTOBACKUP command is set to ON (by default it is OFF), then RMAN automatically backs up the control file ar
server parameter file after every backup and after database structural changes.

upvoted 8 times

NorthKorean Highly Voted 4 years, 4 months ago

This BACKUP AS COPY command allows you to copy a database, tablespaces, datafiles, archived redo logs as well as control files. ANSWER IS BDE

upvoted 7 times

jackymak Most Recent 2 years, 5 months ago

B,D,E
backup as copy ... will get the database image. It should mean all of data in data file. There is no need to backup redo log for recovery.

upvoted 1 times

JMAN1 2 years, 6 months ago

This BACKUP AS COPY command allows you to copy a database, tablespaces, datafiles, archived redo logs as well as control files.

RMAN> backup as copy current controlfile format 'c:\control.ctl';
RMAN> backup controlfilecopy 'c:\control.ctl';

RMAN> backup as copy tablespace users format 'c:\users01.dbf';
RMAN> backup as copy datafile 4 format 'c:\datafile4.dbf';
RMAN> backup as copy spfile format 'c:\image_copy_spfile.ora';
RMAN> backup as copy archivelog all format 'c:\arch_%U.arc';

spfile, pfile are datafiles.

BDE.

upvoted 1 times

JMAN1 2 years, 6 months ago

Sorry... BCDE looks like are all answer... It is so confusing... I want to quit studying :(((..

upvoted 2 times

khalilshahin01 3 years, 11 months ago

BCE using controlfile autobackup will include spfile
https://docs.oracle.com/cd/E11882_01/backup.112/e10642/rcmbckba.htm#BRADV89517

upvoted 4 times

Question #170 *Topic 1*

You have the following entry in the tnsnames.ors of your hq.us.example.com host machine:

```
ORCL =
  (DESCRIPTION =
  (ADDRESS_LIST =
  (ADDRESS = (PROTOCOL = TCP) (HOST = hq.us.example.com) (PORT = 1521))
)
CONNECTED_DATA =
  (SERVICE_NAME = ORCL.us.example.com)
)
)
```

You issue the following command at the command prompt:

Sqlplus HR/HR@ORCL -

Which statement is true about the connection to the ORCL database instance?

A. The connection succeeds, provided the NAMES.DEFAULT_DOMAIN parameter is set to us.example.com in the sqlnet.ora file on the client side.

B. The connection fails because the net service name does not have the suffix us.example.com.

C. The connection succeeds, provided the SERVICE_NAMES initialization parameter is set to ORCL.

D. The connection succeeds, provided the ORCL.us.example.com database service is registered with a listener, the listener is up, and the database is open.

Correct Answer: *D*

Question #171 *Topic 1*

Which three statements are true about Oracle Restart? (Choose three.)

A. It can be configured to automatically attempt to restart various components after a hardware or software failure.

B. While starting any components, it automatically attempts to start all dependencies first and in proper order.

C. It can be configured to automatically restart a database in case of normal shutdown of the database instance.

D. It can be used to only start Oracle components.

E. It runs periodic check operations to monitor the health of Oracle components.

Correct Answer: *BDE*

Community vote distribution

ABE (100%)

khan Highly Voted 5 years, 3 months ago

Correct Answer: ABE
https://docs.oracle.com/database/121/ADMIN/restart.htm#ADMIN12708

upvoted 16 times

sherifelia Highly Voted 4 years, 8 months ago

Correct Answer: ABE

upvoted 10 times

luizyto Most Recent 1 year, 6 months ago

Selected Answer: ABE

E: Oracle Restart runs periodic check operations to monitor the health of these components. If a check operation fails for a component, the component is shut down and restarted.

upvoted 1 times

yarsalan 2 years, 2 months ago

Correct Answer BDE

upvoted 1 times

Ray520 2 years, 3 months ago

Not D: Individual components can be restarted manually using SRVCTL. Oracle restart is used for starting/stopping components in an orderly fashion. So, randomly stopping/starting any individual component has to be done manually (advised to use SRVCTL for that).
A is definitely true: "When you install Oracle Restart, various Oracle components can be automatically restarted after a hardware or software failu
or whenever your database host computer restarts"
https://docs.oracle.com/cd/E18283_01/server.112/e17120/restart001.htm#CIHGAEHA

upvoted 1 times

jackymak 2 years, 5 months ago

Selected Answer: ABE

For A:
https://docs.oracle.com/database/121/ADMIN/restart.htm#ADMIN-GUID-AEB55E5C-48C7-4542-9384-492204F501AE

upvoted 1 times

Oracle2020 2 years, 11 months ago

I agree ABE: https://docs.oracle.com/cd/E18283_01/server.112/e17120/restart001.htm

upvoted 3 times

joe_doe 3 years, 11 months ago

confirm： ABE

upvoted 5 times

Question #172 *Topic 1*

You want to schedule a job to rebuild a tables indexes after a bulk insert, which must be scheduled as soon as a file containing data arrives on the system.
What would you do to accomplish this?

A. Create a file watcher and an event-based job for bulk insert and then create another job to rebuild indexes on the table.

B. Create a file watcher for the bulk inserts and then create a job to rebuild indexes.

C. Create a job array and add a job for bulk insert and a job to rebuild indexes to the job array.

D. Create an event-based job for the file arrival event, then create a job for bulk insert, and then create a job to rebuild indexes.

Correct Answer: *A*

Question #173 *Topic 1*

You execute this command:

```
SQL> CREATE TABLESPACE lmtbsb DATAFILE '/u02/oracle/data/lmtbsb01.dbf' SIZE 50M
     EXTENT MANAGEMENT LOCAL;
```

Which two statements are true about segment space management for segments in this tablespace? (Choose two.)

A. Space utilization inside segments is mapped by bitmaps.

B. Segments are automatically shrunk and compressed when rows are deleted.

C. The PCTFREE storage parameter has no effect on segments created in this tablespace.

D. The PCTUSED storage parameter has no effect on segments created in this tablespace.

Correct Answer: *AD*

TuxBingo 4 years ago

Read the full article: https://oracle-base.com/articles/9i/automatic-segment-free-space-management and you will understand why the asnwer is A and D

upvoted 3 times

JMAN1 2 years, 6 months ago

TuxBingo. You are such a good study roadmap to me.

upvoted 2 times

Question #174 *Topic 1*

Which two statements are true about the (PMON) background process in Oracle Database 12c? (Choose two.)

A. It records checkpoint information in the control file.

B. It frees unused temporary segments.

C. It kills sessions that exceed idle time.

D. It registers database services with all local and remote listeners known to the database instance.

E. It frees resources held by abnormally terminated processes.

Correct Answer: *DE*

References:

https://docs.oracle.com/database/122/CNCPT/process-architecture.htm#CNCPT9840

Community vote distribution

CE (100%)

sherifelia Highly Voted 4 years, 8 months ago

Answer CE
https://docs.oracle.com/database/122/CNCPT/process-architecture.htm#CNCPT9840

upvoted 11 times

LazyLH 4 years, 2 months ago

Yep, CE
PMON periodically performs cleanup of all the following:
Processes that died abnormally
Sessions that were killed
Detached transactions that have exceeded their idle timeout
Detached network connections which have exceeded their idle timeout
In addition, PMON monitors, spawns, and stops the following as needed:
Dispatcher and shared server processes
Job queue processes
Pooled server processes for database resident connection pooling
Restartable background processes

https://docs.oracle.com/database/121/REFRN/GUID-86184690-5531-405F-AA05-BB935F57B76D.htm#REFRN104

upvoted 3 times

khan Highly Voted 5 years, 3 months ago

Which two statements about the Process Monitor (PMON) process are true?

A. PMON performs process recovery when a client process fails.
B. During instance startup, PMON takes care of instance recovery.
C. PMON performs listener registration.
D. PMON restarts background and dispatcher processes when they fail.
E. PMON resolves failures of distributed transactions.

Answer: AD

upvoted 5 times

Havok 4 years, 6 months ago

In 12c, the LREG performs the Listener Registration... PMON was for previous releases.

upvoted 4 times

luizyto Most Recent 1 year, 6 months ago

Selected Answer: CE

PMON is responsible for cleaning up the database buffer cache and freeing resources that the client process was using. For example, PMON rese the status of the active transaction table, releases locks that are no longer required, and removes the process ID from the list of active processes.

upvoted 1 times

Ray520 2 years, 3 months ago

D and E are correct. Since database version is not mentione din question, the option D is valid. Prior to 12c, PMON registers database to Listener services.
PMON does not kill process, but cleans up resources after a process is killed. "The PMON group oversees cleanup of the buffer cache and the release of resources used by a client process."
So C is not correct.
https://docs.oracle.com/en/database/oracle/oracle-database/12.2/cncpt/process-architecture.html#GUID-B5CA9579-53DB-442C-A85F-F21FD334833A

upvoted 1 times

Ray520 2 years, 3 months ago

It seems I am wrong. Confused between the workd cleanup and kill. Document https://docs.oracle.com/database/121/REFRN/GUID-86184690-5531-405F-AA05-BB935F57B76D.htm#REFRN104
states:
PMON periodically performs cleanup of all the following:
Detached transactions that have exceeded their idle timeout
Detached network connections which have exceeded their idle timeout

So, I'm not sure if that means kills a session or not.
Also, since question states version as 12c (I missed that :/), so PMON does not register database to listener, LREG does. So, D is incorrect.

upvoted 1 times

Oracle2020 2 years, 11 months ago

I agree C,E

upvoted 2 times

Mohit201920 4 years, 4 months ago

CE IS CORRECT

upvoted 3 times

Sugar 4 years, 8 months ago

@khan where did you get these items you listed?

upvoted 1 times

Question #175 *Topic 1*

You have successfully taken a database backup by using the command:
RMAN> BACKUP AS BACKUPSET DATABASE;
Now you execute this command:
RMAN> BACKUP INCREMENTAL LEVEL 1 DATABASE;
What is the outcome?

A. It fails because an incremental level 1 backup always searches for an image copy as level 0 backup.

B. It fails because an incremental level 0 backup does not exist.

C. It takes a backup of blocks that have been formatted since the last full database backup.

D. It takes an incremental level 0 backup of the database.

E. It first takes an incremental level 0 backup and then an incremental level 1 backup.

Correct Answer: *E*
References:
https://docs.oracle.com/cd/B19306_01/backup.102/b14192/bkup004.htm
(4.4.1.2)

Community vote distribution
E (100%)

SKY_ming Highly Voted 4 years, 1 month ago
D IS RIGHT
upvoted 9 times

dad Highly Voted 3 years, 9 months ago
correct answer is D :If there is no level 0 backup when you attempt a level 1 backup, RMAN will realize this and rather than returning an error will fact perform a level 0 backup. (i just copy it from oracle doc)
upvoted 5 times

jackymak 2 years, 5 months ago
https://docs.oracle.com/cd/B19306_01/backup.102/b14192/bkup004.htm#sthref391
If compatibility is >=10.0.0, RMAN copies all blocks changed since the file was created, and stores the results as a level 1 backup.
It sound like E.
upvoted 1 times

Ray520 Most Recent 2 years ago
Selected Answer: E
Takes level 0 then level 1 backup
upvoted 1 times

Ray520 2 years, 3 months ago
D is correct. Sample spool for given commands (ran BACKUP AS BACKUPSET DATABASE first):
BACKUP INCREMENTAL LEVEL 1 DATABASE;
Starting backup at 18-APR-22
using channel ORA_DISK_1
no parent backup or copy of datafile 10 found
no parent backup or copy of datafile 8 found
no parent backup or copy of datafile 7 found
no parent backup or copy of datafile 9 found
channel ORA_DISK_1: starting incremental level 0 datafile backup set -----Level 0 backup being taken
....
Finished backup at 18-APR-22
upvoted 2 times

Ray520 2 years ago
Correction: E is correct as verified, takes level 0 first then level 1 backup
upvoted 1 times

Ray_gk 2 years, 3 months ago
E is actually correct. I verified using 12.1.0.2. My RMAN output is too large to post as a comment but I will explain my output.

I took a backupset using the first command in the questions, which was done successfully.

Upon completion of the previous command, I took an incremental level 1 backup. My initial output started as "no parent backup or copy of dataf

x found" meaning that a previous level 0 backup of the specified datafiles were not detected.

RMAN immediately began initiating a level 0 backup "channel ORA_DISK_1: starting compressed incremental level 0 datafile backup set" for each datafile without a level 0 backup taken.

A level 1 backup began immediately after the completion of the level 0 backups, which I observed in the output "channel ORA_DISK_1: starting compressed incremental level 1 datafile backup set"

In this case, E is correct

upvoted 2 times

Oracle2020 2 years, 11 months ago

I think E:
The only difference between a level 0 incremental backup and a full backup is that a full backup is never included in an incremental strategy.

A level 1 incremental backup can be either of the following types:

A differential backup, which backs up all blocks changed after the most recent incremental backup at level 1 or 0

A cumulative backup, which backs up all blocks changed after the most recent incremental backup at level 0

upvoted 1 times

dancymonkey 3 years, 11 months ago

E is right. Did the test on my own HW

upvoted 3 times

charan94 3 years, 3 months ago

But taking level 0 backup and level 1 backup back to back has no use...

upvoted 1 times

TuxBingo 4 years ago

If no level 0 backup is available, then the behavior depends upon the compatibility mode setting. If compatibility is >=10.0.0, RMAN copies all blocks changed since the file was created, and stores the results as a level 1 backup. In other words, the SCN at the time the incremental backup i taken is the file creation SCN. If compatibility <10.0.0, RMAN generates a level 0 backup of the file contents at the time of the backup, to be consistent with the behavior in previous releases.

upvoted 2 times

Mohit201920 4 years, 4 months ago

SHOULD BE C

upvoted 3 times

Question #176

Topic 1

Which two actions does an incremental checkpoint perform? (Choose two.)

A. It signals CKPT to write the checkpoint position to the data file headers.

B. It writes the checkpoint position to the data file headers.

C. It advances the checkpoint position in the checkpoint queue.

D. It writes the checkpoint position to the control file.

Correct Answer: *CD*

References:

http://www.dba-oracle.com/t_incremental_checkpoint.htm

Community vote distribution

CD (100%)

luizyto 1 year, 6 months ago

Selected Answer: CD

An incremental checkpoint is a type of thread checkpoint partly intended to avoid writing large numbers of blocks at online redo log switches. DBW checks at least every three seconds to determine whether it has work to do. When DBW writes dirty buffers, it advances the checkpoint position, causing CKPT to write the checkpoint position to the control file, but not to the data file headers.

upvoted 1 times

Ray520 2 years, 3 months ago

Both A and B are incorrect. Answer is CD:

Incremental checkpoints

An incremental checkpoint is a type of thread checkpoint partly intended to avoid writing large numbers of blocks at online redo log switches. DBW checks at least every three seconds to determine whether it has work to do. When DBW writes dirty buffers, it advances the checkpoint position, causing CKPT to write the checkpoint position to the control file, but not to the data file headers.

https://docs.oracle.com/database/121/CNCPT/startup.htm#CNCPT1301

upvoted 1 times

brg850 2 years, 8 months ago

https://docs.oracle.com/database/121/CNCPT/startup.htm#CNCPT89045
An incremental checkpoint is a type of thread checkpoint partly intended to avoid writing large numbers of blocks at online redo log switches. DBW checks at least every three seconds to determine whether it has work to do.
When DBW writes dirty buffers, it advances the checkpoint position, causing CKPT to write the checkpoint position to the control file, but not to t data file headers.

upvoted 2 times

gio 2 years, 9 months ago

BD
CKPT writes SCN in control file
DBWn writes SCN in data files header

upvoted 2 times

jackymak 2 years, 5 months ago

If so, it should not be B

upvoted 1 times

Clozoerie 3 years ago

answer should be AD

upvoted 1 times

cweather328 4 years, 5 months ago

https://saruamit4.wordpress.com/2014/11/01/checkpoint-scn/

upvoted 1 times

Question #177 *Topic 1*

You want to prevent a group of users in your database from performing long-running transactions that consume huge amounts of space in the undo tablespace. If the quota for these users is exceeded during execution of a data manipulation language (DML) statement, the operation should abort and return an error.
However, queries should still be allowed, even if users have exceeded the undo space limitation.
How would you achieve this?

A. Specify the maximum amount of quota a user can be allocated in the undo tablespace.

B. Decrease the number of Interested Transaction List (ITL) slots for the segments on which these users perform transactions.

C. Implement a profile for these users.

D. Implement a Database Resource Manager plan.

Correct Answer: *D*

Community vote distribution

D (100%)

NorwayOracle 2 years, 4 months ago

Selected Answer: D

From the documentation (https://docs.oracle.com/database/121/ADMIN/dbrm.htm#ADMIN11853):

27.3.6 Undo Pool
You can specify an undo pool for each consumer group. An undo pool controls the total amount of undo for uncommitted transactions that can b generated by a consumer group.

When the total undo generated by a consumer group exceeds its undo limit, the current DML statement generating the undo is terminated. No other members of the consumer group can perform further data manipulation until undo space is freed from the pool.

upvoted 2 times

Grace1_2239478234 2 years, 4 months ago

https://docs.oracle.com/cd/E11882_01/server.112/e25494/undo.htm#ADMIN013

upvoted 1 times

jackymak 2 years, 5 months ago

The question said "a group of user".
Implement on resource will affect all user, so it should implement on profile level ISN'T IT?
I think C is more suitable

upvoted 1 times

jackymak 2 years, 4 months ago

It's D

upvoted 1 times

Question #178 *Topic 1*

A database instance is started by using an SPFILE. The database is configured in ARCHIVELOG mode and the control file autobackup is configured. Daily full database backups are performed by using RMAN.
You lost all control files due to media failure.
Given the steps to recover from the error in random order:
1. Shut down the instance, if it is not already down.
2. Restore the control file from autobackup to a new location.
3. Start the database instance to NOMOUNT state.
4. Recover the database to the point of failure of the control file.
5. Open the database with the RESETLOGS option.
6. Mount the database.
7. Update the SPFILE with the new location of the control file by using the ALTER SYSTEM command.
Identify the correct sequence of the required steps.

A. 1, 3, 2, 6, 7, 4, 5

B. 1, 3, 7, 2, 6, 4, 5

C. 1, 3, 2, 4, 5

D. 1, 2, 6, 4, 5

E. 1, 6, 2, 4, 5

Correct Answer: *A*

dancymonkey Highly Voted 3 years, 11 months ago
Sorry, I did the research

https://docs.oracle.com/database/121/RCMRF/rcmsynta2008.htm#RCMRF149
Example 3-25 Restoring a Control File Autobackup to a Nondefault Location

```
CONNECT TARGET /
STARTUP FORCE NOMOUNT
SET DBID 36508508; # required when restoring control file in NOCATALOG mode
RUN
{
SET CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK
TO '/disk1/prod_cf_auto_%F';
RESTORE CONTROLFILE TO '/tmp/cf_auto.dbf' FROM AUTOBACKUP
MAXSEQ 20 MAXDAYS 150;
ALTER DATABASE MOUNT;
RESTORE DATABASE;
RECOVER DATABASE;
}
ALTER DATABASE OPEN RESETLOGS;
```

The correct answer must be B
upvoted 9 times

loveroba 3 years, 7 months ago
all controlfiles are lost,so you must restore new controlfile and update the spfile before the next startup mount ,checked
upvoted 2 times

jackymak 2 years, 5 months ago
the spfile is only for startup the instance (ie to nomount mode).
The db is already in nomount mode then it is no reason the modify the spfile?
upvoted 1 times

Ray520 Most Recent 2 years ago
Steps as verified:
Stop database
startup nomout.
restore controlfile from autobackup.
startup mount.