



Oracle

1Z0-067

Upgrade Oracle9i/10g/11g OCA to Oracle Database 12c OCP

Questions & Answers PDF

For More Information:

<https://www.certswarrior.com/>

Features:

- 90 Days Free Updates
- 30 Days Money Back Guarantee
- Instant Download Once Purchased
- 24/7 Online Chat Support
- Its Latest Version

Question: 1

Which two statements are true about scheduling operations in a pluggable database (PDB)?

- A. Scheduler jobs for a PDB can be defined only at the container database (CDB) level.
- B. A job defined in a PDB runs only if that PDB is open.
- C. Scheduler attribute setting is performed only at the CDB level.
- D. Scheduler objects created by users can be exported or imported using Data Pump.
- E. Scheduler jobs for a PDB can be created only by common users.

Answer: B,D

Explanation:

In general, all scheduler objects created by the user can be exported/imported into the PDB using data pump. Predefined scheduler objects will not get exported and that means that any changes made to these objects by the user will have to be made once again after the database has been imported into the pluggable database. However, this is how import/export works currently. A job defined in a PDB will run only if a PDB is open.

Question: 2

A complete database backup to media is taken for your database every day. Which three actions would you take to improve backup performance?

- A. Set the backup_tape_io_slaves parameter to true.
- B. Set the dbwr_io_slaves parameter to a nonzero value if synchronous I/O is in use.
- C. Configure large pool if not already done.
- D. Remove the rate parameter, if specified, in the allocate channel command.
- E. Always use RMAN compression for tape backups rather than the compression provided by media manager.
- F. Always use synchronous I/O for the database.

Answer: B,C,D

Explanation:

Tuning RMAN Backup Performance: Procedure Many factors can affect backup performance. Often, finding the solution to a slow backup is a process of trial and error. To get the best performance for a backup, follow the suggested steps in this section: Step 1: Remove RATE Parameters from Configured and Allocated Channels Step 2: If You Use Synchronous Disk I/O, Set DBWR_IO_SLAVES Step 3: If You Fail to Allocate Shared Memory, Set LARGE_POOL_SIZE Step 4: Tune RMAN Tape Streaming Performance Bottlenecks Step 5: Query V\$ Views to Identify Bottlenecks

References:

<https://docs.oracle.com/database/121/BRADV/rcmtunin.htm#BRADV172>

Question: 3

For which three pieces of information can you use the RMAN list command?

- A. stored scripts in the recovery catalog
- B. available archived redo log files
- C. backup sets and image copies that are obsolete
- D. backups of tablespaces
- E. backups that are marked obsolete according to the current retention policy

Answer: A,B,D

Explanation:

About the LIST Command: The primary purpose of the LIST command is to list backup and copies. For example, you can list: -Backups and proxy copies of a database, tablespace, datafile, archived redo log, or control file -Backups that have expired -Backups restricted by time, path name, device type, tag, or recoverability -Archived redo log files and disk copies.

References:

http://docs.oracle.com/cd/B28359_01/backup.111/b28270/rcmreprt.htm#BRADV89585

Question: 4

You notice performance degradation in your production Oracle 12c database. You want to know what caused this performance difference.

Which method or feature should you use?

- A. Database Replay
- B. Automatic Database Diagnostic Monitor (ADDM) Compare Period report
- C. Active Session History (ASH) report
- D. SQL Performance Analyzer

Answer: B

Explanation:

References: http://docs.oracle.com/cd/E24628_01/server.121/e17635/tdpnt_degrade.htm

Question: 5

Which three statements are true about a job chain?

- 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.

Answer: A,B,E

Explanation:

Chains are the means by which you can implement dependency based scheduling, in which jobs are started depending on the outcomes of one or more previous jobs.

DBMS_SCHEDULER.DEFINE_CHAIN_STEP DBMS_SCHEDULER.DEFINE_CHAIN_EVENT_STEP

References:

http://docs.oracle.com/cd/B28359_01/server.111/b28310/scheduse009.htm#ADMIN12

Question: 6

Because of logical corruption of data in a table, you want to recover the table from an RMAN backup to a specified point in time.

Examine the steps to recover this table from an RMAN backup:

1. Determine which backup contains the table that needs to be recovered.
 2. Issue the recover table RMAN command with an auxiliary destination defined and the point in time specified.
 3. Import the Data Pump export dump file into the auxiliary instance.
 4. Create a Data Pump export dump file that contains the recovered table on a target database.
- Identify the required steps in the correct order.

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

Answer: D

Explanation:

Because according to oracle PDFs if you run restore table ... auxiliary a impede and rename can be included. So there is no reason to make the import manually if it can be already included in step 2.

References:

<https://docs.oracle.com/database/121/BRADV/rcmresind.htm#BRADV689>

Question: 7

Examine the command:

SQL> RECOVER DATABASE USING BACKUP CONTROLFILE UNTIL CANCEL;

In which two scenarios is this command required?

- A. The current online redo log file is missing.
- B. A data file belonging to a noncritical tablespace is missing.
- C. All the control files are missing.
- D. The database backup is older than the control file backup.

E. All the data files are missing.

Answer: A,C

Explanation:

References:

<http://searchoracle.techtarget.com/answer/Recover-database-using-backup-controlfile-until-cancel>

Question: 8

Which two are prerequisites for setting up Flashback Data Archive?

- A. Fast Recovery Area should be defined.
- B. Undo retention guarantee should be enabled.
- C. Supplemental logging should be enabled.
- D. Automatic Undo Management should be enabled.
- E. All users using Flashback Data Archive should have unlimited quota on the Flashback Data Archive tablespace.
- F. The tablespace in which the Flashback Data Archive is created should have Automatic Segment Space Management (ASSM) enabled.

Answer: D,F

Explanation:

There are a number of restrictions for flashback archives: The tablespaces used for a flashback archive must use local extent management and automatic segment space management. The database must use automatic undo management.

References:

http://www.dba-oracle.com/t_11g_new_enabling_fdba.htm

Question: 9

The environmental variable oracle_Base is set to /u01/app/oracle and oracle_home is set to /u01/app/oracle/product/12.1.0/db 1.

You want to check the diagnostic files created as part of the Automatic Diagnostic Repository (ADR).

Examine the initialization parameters set in your database.

NAME TYPE VALUE

audit_file_deststring/u01/app/oracle/admin/eml2rep/adump
background_dump_deststring
core_dump_deststring
db_create_file_deststring
db_recovery_file_deststring/u01/app/oracle/fast_recovery_area
diagnostic_deststring

What is the location of the ADR base?

- A. It is set to /u01/app/oracle/product:/12.1.0/db_1/log.
- B. It is set to /u01/app/oracle/admin/en12.1.0/adump.
- C. It is set to /u01/app/oracle.
- D. It is set to /u01/app/oracle/flash_recovery_area.

Answer: C

Explanation:

The Automatic Diagnostic Repository (ADR) is a directory structure that is stored outside of the database. It is therefore available for problem diagnosis when the database is down.

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

If environment variable ORACLE_BASE is set, DIAGNOSTIC_DEST is set to the directory designated by ORACLE_BASE.

If environment variable ORACLE_BASE is not set, DIAGNOSTIC_DEST is set to ORACLE_HOME/log.

References:

http://docs.oracle.com/cd/B28359_01/server.111/b28310/diag001.htm#ADMIN11008

Question: 10

You want to export the pluggable database (PDB) hr_pdb1 from the multitenant container database (CDB)CDB1 and import it into the cdb2 CDB as the emp_pdb1 PDB.

Examine the list of possible steps required to perform the task:

- 1.Create a PDB named emp_pdb1.
 - 2.Export the hr_pdb1 PDB by using the full clause.
 - 3.Open the emp_pdb1 PDB.
 - 4.Mount the emp_pdb1 PDB.
 - 5.Synchronize the emp_pdb1 PDB in restricted mode.
 - 6.Copy the dump file to the Data Pump directory.
 - 7.Create a Data Pump directory in the emp_pdb1 PDB.
 - 8.Import data into emp_pdb1 with the full and remap clauses.
 - 9.Create the same tablespaces in emp_pdb1 as in hr_pdb1 for new local user objects.
- Identify the required steps in the correct order.

- A. 2, 1, 3, 7, 6, and 8
- B. 2, 1, 4, 5, 3, 7, 6, 9, and 8
- C. 2, 1, 3, 7, 6, 9, and 8
- D. 2, 1, 3, 5, 7, 6, and 8

Answer: C

Explanation:

Because the step 2 says that you perform an expdp with the full clause and you don't need to create the tablespaces when you perform the impdp. FULL=yes will export tablespace definitions. So no need step 9.

References:

https://docs.oracle.com/cd/B10501_01/server.920/a96652/ch01.htm



CERTSWARRIOR

FULL PRODUCT INCLUDES:

Money Back Guarantee



Instant Download after Purchase



90 Days Free Updates



PDF Format Digital Download



24/7 Live Chat Support



Latest Syllabus Updates



For More Information – Visit link below:

<http://www.certswarrior.com>

Discount Coupon Code:

CERTSWARRIOR10

We Accept

PayPal