

1Z0-053 Dumps

Oracle Database 11g: Administration II

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



NEW QUESTION 1

- (Topic 1)

Which two are the uses of the ASM metadata backup and restore (AMBR) feature? (Choose two.)

- A. It can be used to back up all data on ASM disks.
- B. It can be used to recover the damaged ASM disk group along with the data.
- C. It can be used to gather information about a pre-existing ASM disk group with disk paths, disk name, failure groups, attributes, templates, and alias directory structure.
- D. It can be used to re-create the ASM disk group with its attributes.

Answer: CD

NEW QUESTION 2

- (Topic 1)

The following command is executed to shut down an Automatic Storage Management (ASM) instance:

```
SQL>SHUTDOWN ABORT;
```

Which two statements describe the consequences of the above command? (Choose two.)

- A. The disk groups are orderly dismounted.
- B. The ASM instance requires recovery when it is started.
- C. The database instance that are currently clients of the ASM instance are aborted.
- D. The CSS daemon stops and has to be restarted before the ASM instance is restarted.

Answer: BC

NEW QUESTION 3

- (Topic 1)

The ORACLE_SID environment variable is set to +ASM. ASMLIB is not used in the configuration. You executed the following command to startup the Automatic Storage Management (ASM) instance.

```
SQL> STARTUP;
```

Which two activities are performed during a successful start up operation? (Choose two.)

- A. The databases configured to use the ASM instance are mounted
- B. The disk groups are mounted as per the ASM_DISKGROUPS initialization parameter
- C. ASM starts the Oracle Cluster Synchronization Services (CSS) daemon if it is not started
- D. ASM discovers and examines the contents of all files that are in the paths specified in the ASM_DISKGROUPS initialization parameters

Answer: BC

Explanation:

Refer to Starting Up an ASM Instance. To start up an ASM instance, you must:

1. To connect to the ASM instance with SQL*Plus, you must set the ORACLE_SID environment variable to the ASM SID.
2. The initialization parameter file, which can be a server parameter file, must contain: INSTANCE_TYPE = ASM
3. The STARTUP command, tries to mount the disk groups specified by the initialization parameter ASM_DISKGROUPS.

If ASM_DISKGROUPS is blank, the ASM instance starts and warns that no disk groups were mounted. You can then mount disk groups with the ALTER DISKGROUP...MOUNT command.

The Cluster Synchronization Services (CSS) daemon is required to enable synchronization between ASM and its client database instances. The CSS daemon is normally started (and configured to start upon reboot) when you use Database Configuration Assistant (DBCA) to create your database. If you did not use DBCA to create the database, you must ensure that the CSS daemon is running before you start the ASM instance.

NEW QUESTION 4

- (Topic 1)

You are managing an Oracle Database 11g database with the ASM storage. The database is having big file tablespaces. You want files to open faster and less memory to be used in the shared pool to manage the extent maps.

What configuration would you effect to achieve your objective? (Choose all that apply.)

- A. Set the ASM compatibility attribute for the ASM disk group to 11.1.0.
- B. Set the RDBMS compatibility attribute for the ASM disk group to 11.1.0.
- C. Set the COMPATIBLE initialization parameter for the ASM instance to 11.1.0.
- D. Set the COMPATIBLE initialization parameter for the database instance to 11.1.0.

Answer: AD

NEW QUESTION 5

- (Topic 1)

In your database, the LDAP_DIRECTORY_SYSAUTH initialization parameter has been set to YES and the users who need to access the database as DBAs have been granted SYSDBA enterprise role in Oracle Internet Directory (OID). SSL and the password file have been configured. A user SCOTT with the SYSDBA privilege tries to connect to the

database instance from a remote machine using the command:

```
$ SQLPLUS scott/tiger@DB01 AS SYSDBA
```

Which DB01 is the net service name.

Which authentication method would be used first?

- A. authentication by password file
- B. authentication by using certificates over SSL
- C. authentication by using the Oracle Internet Directory
- D. authentication by using the local OS of the database server

Answer: A

NEW QUESTION 6

- (Topic 1)

What are the recommendations for Oracle Database 11g installation to make it Optimal Flexible Architecture (OFA)-compliant? (Choose all that apply.)

- A. ORACLE_BASE should be set explicitly.
- B. An Oracle base should have only one Oracle home created in it.
- C. Flash recovery area and data file location should be on separate disks.
- D. Flash recovery area and data file location should be created under Oracle base in a non-Automatic Storage Management (ASM) setup.

Answer: ACD

NEW QUESTION 7

- (Topic 1)

You want to perform the following operations for the DATA ASM disk group:

- ? Verify the consistency of the disk.
- ? Cross-check all the file extent maps and allocation tables for consistency.
- ? Check whether the alias metadata directory and file directory are linked correctly.
- ? Check that ASM metadata directories do not have unreachable allocated blocks.

Which command accomplishes these tasks?

- A. ALTER DISKGROUP data CHECK;
- B. ALTER DISKGROUP data CHECK DISK;
- C. ALTER DISKGROUP data CHECK FILE;
- D. ALTER DISKGROUP data CHECK DISK IN FAILURE GROUP 1;

Answer: A

Explanation:

Syntax: ALTER DISKGROUP <disk_group_id> CHECK [REPAIR | NOREPAIR];

The check_diskgroup_clause lets you verify the internal consistency of Oracle ASM disk group metadata. The disk group must be mounted. Oracle ASM displays summary errors and writes the details of the detected errors in the alert log.

The CHECK keyword performs the following operations:

- ? Checks the consistency of the disk.
- ? Cross checks all the file extent maps and allocation tables for consistency.
- ? Checks that the alias metadata directory and file directory are linked correctly.
- ? Checks that the alias directory tree is linked correctly.
- ? Checks that Oracle ASM metadata directories do not have unreachable allocated blocks.

Refer to here

NEW QUESTION 8

- (Topic 1)

What is the result of increasing the value of the parameter ASM_POWER_LIMIT during a rebalance operation?

- A. The ASM rebalance operation will likely consume fewer resources and complete in a shorter amount of time.
- B. The ASM rebalance operation will consume fewer resources and complete in a longer amount of time.
- C. The ASM rebalance operation will be parallelized and should complete in a shorter amount of time.
- D. There is no ASM_POWER_LIMIT setting used in ASM.
- E. None of the above

Answer: C

NEW QUESTION 9

- (Topic 1)

ASM supports all but which of the following file types? (Choose all that apply.)

- A. Database files
- B. SPFILES
- C. Redo-log files
- D. Archived log files
- E. RMAN backup sets
- F. Password files
- G. init.ora files

Answer: FG

Explanation:

What Types of Files Does Oracle ASM Support?

Table 7-1 File Types Supported by Automatic Storage Management

File Type	Default Templates
Control files	CONTROLFILE
Data files	DATAFILE
Redo log files	ONLINELOG
Archive log files	ARCHIVELOG
Temporary files	TEMPFILE
Data file backup pieces	BACKUPSET
Data file incremental backup pieces	BACKUPSET
Archive log backup piece	BACKUPSET
Data file copy	DATAFILE
Persistent initialization parameter file (SPFILE)	PARAMETERFILE
Flashback logs	FLASHBACK
Change tracking file	CHANGETRACKING
Data Pump dumpset	DUMPSET
Automatically generated control file backup	AUTOBACKUP
Cross-platform transportable data files	XTRANSPORT
Flash file	FLASHFILE
Oracle ASM Persistent initialization parameter file (SPFILE)	ASMPARAMETERFILE
Oracle ASM Persistent initialization parameter file (SPFILE) backup	ASMPARAMETERFILEBACKUP
Oracle Cluster Registry file	OCRFILE
Oracle ASM Dynamic Volume Manager volumes	n/a

NEW QUESTION 10

- (Topic 1)

Which background process of a database instance, using Automatic Storage Management (ASM), connects as a foreground process into the ASM instance?

- A. ASMB
- B. PMON
- C. RBAL
- D. SMON

Answer: A

Explanation:

ASMB (ASM Background Process): Communicates with the ASM instance, managing storage and providing statistics, runs in ASM instances when the ASMCMD cp command runs or when the database instance first starts if the server parameter file is stored in ASM. ASMB also runs with Oracle Cluster Registry on ASM.
RBAL (ASM Rebalance Master Process): In an ASM instance, it coordinates rebalance activity for disk groups. In a database instances, it manages ASM disk groups.

PMON (Process Monitor): Monitors the other background processes and performs process recovery when a server or dispatcher process terminates abnormally.

SMON (System Monitor Process): Performs critical tasks such as instance recovery and dead transaction recovery, and maintenance tasks such as temporary space reclamation, data dictionary cleanup, and undo tablespace management

NEW QUESTION 10

- (Topic 1)

When an ASM instance receives a SHUTDOWN NORMAL command, what command does it pass on to all database instances that rely on the ASM instances disk groups?

- A. TRANSACTIONAL
- B. IMMEDIATE
- C. ABORT
- D. NORMAL

Answer: A

NEW QUESTION 13

- (Topic 1)

Which type of database file is spread across all disks in a disk group?

- A. All types of files are spread across all disks in the disk group.
- B. Datafiles
- C. Redo log files
- D. Archived redo log files
- E. Control files

Answer: A

NEW QUESTION 18

- (Topic 1)

You are managing an Oracle Database 11g instance with ASM storage. The ASM instance is down. To know the details of the disks in the DATA disk group, you issued the following ASMCMD command:

```
ASMCMD> lsdisk -l -d DATA
```

Which statement is true regarding the outcome of this command?

- A. The command succeeds but it retrieves only the disk names.
- B. The command produces an error because the ASM instance is down.
- C. The command succeeds but it shows only the status of the ASM instance.
- D. The command succeeds and retrieves information by scanning the disk headers based on anASM_DISKSTRING value.

Answer: D

Explanation:

See details at Options for the Lsdisk command.

-l Scans disk headers for information rather than extracting the information from an Oracle ASM instance. This option forces non-connected mode.

NEW QUESTION 20

- (Topic 1)

What are three benefits of using ASM? (Choose three.)

- A. Ease of disk administration and maintenance
- B. Load balancing across physical disks
- C. Software RAID-1 data redundancy with double or triple mirrors
- D. Automatic recovery of failed disks

Answer: ABC

NEW QUESTION 21

- (Topic 1)

You have three production databases, HRDB, FINDB, and ORGDB, that use the same ASM instance. At the end of the day, while all three production database instances are running, you execute the following command on the ASM instance:

```
SQL> shutdown immediate;
```

What is the result of executing this command?

- A. The ASM instance is shut down, but the other instances are still running.
- B. It results in an error because other database instances are connected to it.
- C. All the instances, including the ASM instance, are shut down in the IMMEDIATE mode.
- D. HRDB, FINDB, and ORGDB instances are shut down in the ABORT mode and the ASM instance is shut down in the IMMEDIATE mode.

Answer: B

NEW QUESTION 23

- (Topic 2)

What does it mean if a backup is expired?

- A. The backup set has exceeded the retention criteria set in RMAN and is eligible for removal.
- B. The backup set has one or more invalid blocks in it and is not usable for recovery.
- C. The backup set contains one or more tablespaces no longer in the database.
- D. The backup set contains one or more missing backup set pieces.
- E. The backup set is from a previous version of RMAN and was not upgraded.

Answer: D

NEW QUESTION 27

- (Topic 2)

What command is used to reset a database to a previous incarnation?

- A. reset incarnation
- B. incarnation reset
- C. reset database to incarnation
- D. reset database incarnation
- E. reset databse incarnation number

Answer: C

NEW QUESTION 30

- (Topic 2)

If a log file becomes corrupted, it may cause the database to stale. How would you correct such a situation?

- A. Recover the online redo log from backup.
- B. Delete and re-create the log file.
- C. Use the ALTER DATABASE CLEAR LOGFILE command to clear the log file.
- D. Shut down the database and restart it.
- E. Shut down the database and then mount it.
- F. Clear the log file with the ALTER DATABASE CLEAR LOGFILE command and then restart the database with ALTER DATABASE OPEN RESETLOGS.

Answer: C

Explanation:

If you don't use the ALTER DATABASE CLEAR LOGFILE command to clear log online, after the database restarted, the log file will be INVALID. So the simple way is to clear the log file directly.

NEW QUESTION 32

- (Topic 2)

Archived redo logs can be copied to more than one destination by Oracle.

- A. True
- B. False

Answer: A

NEW QUESTION 36

- (Topic 2)

Your database has a backup that was taken yesterday (Tuesday) between 13:00 and 15:00 hours. This is the only backup you have. You have lost all the archived redo logs generated since the previous Monday, but you have archived redo logs available from the previous Sunday and earlier. You now need to restore your backup due to database loss.

To which point can you restore your database?

- A. 13:00 on Tuesday.
- B. 15:00 on Tuesday.
- C. Up until the last available archived redo log on Sunday.
- D. To any point; all the redo should still be available in the online redo logs.
- E. The database is not recoverable.

Answer: E

NEW QUESTION 38

- (Topic 2)

What does the SCN represent?

- A. The system change number, which is a point in time relative to transactions within a given database.
- B. A number that represents time
- C. Thus, at 1300 hours, the SCN is the same on all databases.
- D. The security change number, which represents the security code that is needed to access any database structure.
- E. A conversion factor that converts internal database time to external clock time.
- F. UTC time in the database, providing a standardized way of tracking time in Oracle.

Answer: A

Explanation:

SCN

System Change Number. A database ordering primitive.

The value of an SCN is the logical point in time at which changes are made to a database. Refer to v11.2 document

NEW QUESTION 41

- (Topic 2)

You want to use the automatic management of backup and recovery operations features for your database.

Which configuration must you set?

- A. Enable the flash recovery area and specify it as the archived redo log destination.
- B. Disable the flash recovery area and start the database instance in ARCHIVELOG mode.
- C. Enable the flash recovery area but do not specify it as the archived redo log destination.
- D. Disable the flash recovery area and start the database instance in NOARCHIVELOG mode.

Answer: A

NEW QUESTION 45

- (Topic 2)

What are the different logging modes available in Oracle Database 11g? (Choose two.)

- A. NOLOG mode
- B. NOARCHIVELOG mode
- C. LOGGING mode
- D. HOTDATABASE mode
- E. ARCHIVELOG mode

Answer: BE

NEW QUESTION 50

- (Topic 2)

Every Sunday the Unix system administrator has a job that executes a full backup of the entire Unix system your database is on. Is this backup usable for backup and recovery of your database?

- A. Yes, if the database is in ARCHIVELOG mode.
- B. Yes, if the database is in NOARCHIVELOG mode.
- C. No, the backup is not usable in any way.
- D. Only if the ENABLE_ONLINE_BACKUP parameter is set to TRUE.

Answer: C

NEW QUESTION 51

- (Topic 2)

Your developers have asked you to restore the development database, which is in NOARCHIVELOG mode, back to last Tuesday the 20th. Your last backup is from Monday the 19th.

What do you do?

- A. Restore the 19th's backup, restore all archived redo logs, recover the database to the 20th, and open the database.
- B. Tell them that their request cannot be met with the current backup strategy.
- C. Restore the 19th's backup, apply the online redo logs, and open the database.
- D. Switch the database into ARCHIVELOG mode, restore the 19th's backup, restore all archived redo logs, and recover the database to the 20th.
- E. Use the recover database command to roll back the database from today to the 19th of the month.

Answer: B

NEW QUESTION 54

- (Topic 2)

You have the following requirements in relation to the detection of block corruption for your database instance:

- ? Check for logical self-consistency of data blocks when modified in memory.
- ? Checksums are calculated before and after the block change.
- ? Checks are performed for the lost writes to the physical standby database.

Which method would help you perform the above checks automatically?

- A. Set the DB_SECUREFILE parameter to PERMITTED.
- B. Set the DB_ULTRA_SAFE parameter to DATA_ONLY.
- C. Set the DB_LOCK_CHECKSUM parameter to TYPICAL.
- D. Set the DB_LOST_WRITE_PROTECT parameter to TYPICAL.

Answer: B

Explanation:

Parameter type	String
Syntax	DB_ULTRA_SAFE = { OFF DATA_ONLY DATA_AND_INDEX }
Default value	OFF
Modifiable	No
Basic	No

DB_ULTRA_SAFE sets the default values for other parameters that control protection levels.

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

OFF

When any of DB_BLOCK_CHECKING, DB_BLOCK_CHECKSUM, or DB_LOST_WRITE_PROTECT are explicitly set, no changes are made.

DATA_ONLY

? DB_BLOCK_CHECKING will be set to MEDIUM.

? DB_LOST_WRITE_PROTECT will be set to TYPICAL.

? DB_BLOCK_CHECKSUM will be set to FULL.

DATA_AND_INDEX

DB_BLOCK_CHECKING will be set to FULL. DB_LOST_WRITE_PROTECT will be set to TYPICAL. DB_BLOCK_CHECKSUM will be set to FULL.

Parameter type	String
Syntax	DB_BLOCK_CHECKING = { FALSE OFF LOW MEDIUM TRUE FULL }
Default value	FALSE
Modifiable	ALTER SYSTEM
Basic	No

DB_BLOCK_CHECKING specifies whether or not Oracle performs block checking for database blocks.

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

OFF or FALSE

No block checking is performed for blocks in user tablespaces. However, semantic block checking for SYSTEM tablespace blocks is always turned on.

LOW

Basic block header checks are performed after block contents change in memory (for example, after UPDATE or INSERT statements, on-disk reads, or inter-instance block transfers in Oracle RAC).

MEDIUM

All LOW checks and full semantic checks are performed for all objects except indexes (whose contents can be reconstructed by a drop+rebuild on encountering a corruption).

FULL or TRUE

All LOW and MEDIUM checks and full semantic checks are performed for all objects.

NEW QUESTION 58

- (Topic 2)

Which is the correct command to put the database in ARCHIVELOG mode?

- A. alter database archivelog
- B. alter system enable archivelog mode
- C. alter database enable archive
- D. alter database archivelog enable

E. None of the above

Answer: A

NEW QUESTION 61

- (Topic 2)

If you issue the command shutdown abort prior to trying to put the database in ARCHIVELOG mode, what will be the result when you issue the command alter database archivelog?

- A. The alter database archivelog command will fail.
- B. The alter database archivelog inconsistent command must be used to put the database in ARCHIVELOG mode.
- C. The alter database archivelog command will succeed.
- D. The alter database archivelog command will ask if you want to make the database consistent first.
- E. There is no alter database archivelog command.
- F. The correct command is alter database alterlogging.

Answer: A

Explanation:

Before you change database to archivelog mode, you need to have a clean database shutdown.

NEW QUESTION 63

- (Topic 2)

You have not configured Oracle Managed Files (OMF) in your database. You do not want to scan the entire datafile every time an incremental backup is performed. You decide to enable the block change tracking feature. Which statement should you use to enable the block change tracking feature?

- A. ALTER DATABASE ENABLE BLOCK CHANGE TRACKING;
- B. ALTER SYSTEM ENABLE BLOCK CHANGE TRACKING USING FILE <path>;
- C. ALTER DATABASE ENABLE BLOCK CHANGE TRACKING USING FILE <path>;
- D. ALTER SYSTEM ENABLE BLOCK CHANGE TRACKING;

Answer: C

NEW QUESTION 64

- (Topic 2)

How many individual archive-log destination directories are supported by Oracle Database 11g?

- A. 7
- B. 1
- C. 10
- D. 11
- E. 21

Answer: C

NEW QUESTION 67

- (Topic 2)

Which of the following parameters defines the location where Oracle should create archived redo logs?

- A. LOG_ARCHIVE_1
- B. LOG_DESTINATION_1
- C. LOG_ARCHIVED_DESTINATION_1
- D. LOG_ARCHIVE_DEST_1
- E. LOG_ARCHIVE_SOURCE_1

Answer: D

NEW QUESTION 68

- (Topic 2)

Which command would you use to determine what database backups are currently available for restore?

- A. list database backup;
- B. report database backup;
- C. list backup of database;
- D. list summary backup;
- E. report backup of database;

Answer: C

NEW QUESTION 71

- (Topic 3)

What privileges must be granted to allow an account to create the recovery catalog? (Choose all that apply.)

- A. RECOVERY_CATALOG_OWNER
- B. DBA
- C. RESOURCE
- D. SELECT ANY DICTIONARY

E. CONNECT

Answer: AC

NEW QUESTION 76

- (Topic 3)

Identify two advantages of using a recovery catalog in-load of the control File of the target database Recovery Manager (RMAN). (Choose two.)

- A. You can use RMAN stored scripts.
- B. Recovery is faster if data is stored in catalog in addition to the control file.
- C. You can store backup Information of all registered databases in one place.
- D. Database backups are automatically deleted when they are older than the specified time period.

Answer: AC

NEW QUESTION 78

- (Topic 3)

While performing a regular check on your recovery catalog you realized that the catalog database is running out of space and you do not have options to increase the space. However, you have another database where more space is available and you want to move your existing recovery catalog to this database.

The options that can be considered while moving the recovery catalog are as follows:

1. Using one of the Oracle expdp utilities to export the catalog data
2. Creating a recovery catalog user and granting the necessary privileges in the other database
3. Creating the recovery catalog using the CREATE CATALOG command
4. Using the corresponding impdp utility to import the catalog data into the other database
5. Registering the target database in the new catalog database using the REGISTER DATABASE command.

Identify the option with the correct sequence for moving the recovery catalog.

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

Answer: B

Explanation:

The exp/imp tools can export and import the complete data structure and data extents to the destination database, so that you don't need to do create catalog and register database.

NEW QUESTION 79

- (Topic 4)

The Oracle Database 11g database is running in the ARCHIVELOG mode. The archived redo log files are stored on three locations. The Flash Recovery Area is one of the locations. The details are given below:

```
LOG_ARCHIVE_DEST_1 = 'LOCATION = /disk1/archive' LOG_ARCHIVE_DEST_2 = 'SERVICE = stdb1'
```

```
DB_RECOVERY_FILE_DEST = '/u01/oradata'
```

Examine the following RMAN command issued to set the deletion policy for archived log files:

```
RMAN> CONFIGURE ARCHIVELOG DELETION POLICY TO BACKED UP 2 TIMES TO
```

```
sbt;
```

Which statement is true regarding what this command accomplishes?

- A. All archived redo log files backed up at least twice to tape are deleted.
- B. All archived redo log files backed up at least once to tape are eligible for deletion.
- C. All archived redo log files backed up at least twice to tape are deleted from the flash recovery area.
- D. All archived redo log files in local archiving destinations and the flash recovery area backed up at least twice to tape are eligible for deletion.

Answer: D

NEW QUESTION 83

- (Topic 4)

Which type of backup backs up only data blocks modified since the most recent backup at the same level or lower?

- A. Differential incremental backup
- B. Different incremental backup
- C. Cumulative backup
- D. Cumulative incremental backup

Answer: A

NEW QUESTION 87

- (Topic 4)

You want the ability to recovery any time within the last seven days and therefore you configured the recovery window retention policy using the command:

```
RMAN> CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 7 DAYS;
```

After configuring the recovery window, you performed the database backup as follows:

- A. Backup RB1 at log sequence number 12871 on 5th Jan
- B. Backup RB2 at log sequence number 15622 on 12th Jan
- C. Backup RB3 at log sequence 16721 on 15th Jan
- D. On 20th Jan when the log sequence number was 18112 you realize that there is a need to a point in time at the beginning of the recovery window
- E. You have all archived redo log files to date.

Answer: D

NEW QUESTION 88

- (Topic 4)

You configured the default backup device type as disk for RMAN backups. In your database, because of business requirements, you have to take a simultaneous duplicate backup of the data files when the RMAN BACKUP command is used.

What must you set using the RMAN CONFIGURE command to achieve this?

- A. MAXSETSIZE TO 2;
- B. DEVICE TYPE DISK PARALLELISM 2;
- C. RETENTION POLICY TO REDUNDANCY 2;
- D. DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 2;

Answer: D

Explanation:

Duplexing Backup Sets with CONFIGURE BACKUP COPIES ([Link](#))

NEW QUESTION 89

- (Topic 4)

Which type of backup contains only the blocks that have changed since the last level 0 incremental backup?

- A. a cumulative level 1 backup
- B. a differential level 1 backup
- C. a full backup
- D. a whole backup

Answer: A

NEW QUESTION 90

- (Topic 4)

What command would you issue to enable automated backups of control files?

- A. alter database controlfile autobackup on
- B. alter system controlfile autobackup on
- C. configure controlfile autobackup on
- D. enable controlfile autobackup

Answer: C

NEW QUESTION 95

- (Topic 4)

Which of the following most closely represents an image copy?

- A. Unix cp command of a file
- B. Bit-by-bit copy of a file
- C. Windows COPY command of a file
- D. All of the above

Answer: D

NEW QUESTION 98

- (Topic 5)

Compressed backups work with which of the following commands?

- A. copy as backup
- B. backup as copy
- C. backup
- D. copy

Answer: C

Explanation:

The backup compression only works with backup set, NONE of image copies can work with compression.

NEW QUESTION 100

- (Topic 5)

You execute the following RMAN command to perform the backup operation:

```
RMAN> RUN
```

```
{  
ALLOCATE CHANNEL c1 DEVICE TYPE disk MAXOPENFILES 8; BACKUP DATABASE FILESPERSET 4;  
}
```

What is the multiplexing level in the preceding backup process?

- A. 4
- B. 8

C. 7

Answer: A

NEW QUESTION 103

- (Topic 5)

You want to perform an RMAN backup of database as a copy. Which two factors will you consider while performing the backup operation? (Choose two).

- A. The backup as copy can only be taken to disk
- B. The backup as copy can only be taken to tape
- C. Backup can be performed only when the instance is shutdown
- D. Backup will constitute all used and unused blocks in the database

Answer: AD

NEW QUESTION 107

- (Topic 5)

Which of the following commands will fail?

- A. report schema;
- B. report need backup;
- C. report need backup days 3;
- D. report user;
- E. report obsolete;

Answer: D

NEW QUESTION 112

- (Topic 5)

What is the impact of the results of the output of the following command?

```
RMAN> report unrecoverable database;
```

Report of files that need backup due to unrecoverable operations File Type of Backup Required Name

```
-----  
4 full or incremental C:\ORACLE\ORADATA\ORCL\USERS01.DBF
```

- A. There are no backup sets with any backups of the users01.dbf datafile.
- B. The users01.dbf datafile has had unrecoverable operations occur in i
- C. It will need to be backed up or some data loss is possible during a recovery.
- D. The users01.dbf datafile is corrupted.
- E. The users01.dbf datafile backup exceeds the retention criteria.
- F. The last backup of the users01.dbf datafile failed and must be rerun.

Answer: D

NEW QUESTION 117

- (Topic 5)

Which backup option defines a user-defined name for a backup?

- A. FORMAT
- B. NAME
- C. TAG
- D. FORMAT U%

Answer: C

NEW QUESTION 122

- (Topic 5)

Why would you execute the report obsolete command?

- A. To list all backups that were no longer available for restore operations
- B. To list all backups that had aged beyond the RMAN retention criteria
- C. To list all backup set pieces listed in control-file or recovery-catalog metadata that are not on the backup media
- D. To list all datafiles that are no longer part of the database and thus do not need to be backed up
- E. To list all archived redo logs that are no longer needed for any database recovery

Answer: B

Explanation:

You can report backup sets, backup pieces, and data file copies that are obsolete that is, not needed to meet a specified retention policy by specifying the OBSOLETE keyword.

NEW QUESTION 127

- (Topic 5)

Examine the following RMAN script:

```
RMAN> run { debug on;  
allocate channel c1 type disk; backup datafile 5;  
}
```

Which statement describes the purpose of the script?

- A. The data file is checked for physical corruption and backed up if found clean.
- B. The backup of data file 5 is performed and the interactive messages during the backup are suppressed.
- C. The existing backup for the data file is checked and the backup is performed if there are changes in the data file after the last backup.
- D. The backup of data file 5 is performed and all SQL statements that are executed during RMAN compilation and their results are displayed

Answer: D

NEW QUESTION 128

- (Topic 5)

Which is the correct command to back up the database, back up the archived redo logs, and then remove the backed-up archived redo logs?

- A. backup database
- B. backup database and archivelogs
- C. backup database plus archivelogs
- D. backup database plus archivelog delete input
- E. backup database and archivelog delete input

Answer: D

NEW QUESTION 131

- (Topic 5)

Why would you run the delete obsolete command? (Choose all that apply.)

- A. To remove missing backup set pieces physically from disk
- B. To remove metadata related to backup set pieces in the control file and the recovery catalog
- C. To mark as deleted records in the control file and the recovery catalog associated with obsolete backup sets
- D. To delete backup set pieces associated with backups that are no longer needed due to retention criteria
- E. To remove old versions of RMAN backups

Answer: CD

Explanation:

Deleting Expired RMAN Backups and Copies

If you run CROSSCHECK, and if RMAN cannot locate the files, then it updates their records in the RMAN repository to EXPIRED status. You can then use the DELETE EXPIRED command to remove records of expired backups and copies from the RMAN repository.

The DELETE EXPIRED command issues warnings if any files marked as EXPIRED actually exist. In rare cases, the repository can mark a file as EXPIRED even though it exists. For example, a directory containing a file is corrupted at the time of the crosscheck, but is later repaired, or the media manager was not configured properly and reported some backups as not existing when they really existed.

To delete expired repository records:

If you have not performed a crosscheck recently, then issue a CROSSCHECK command. For example, issue:

```
CROSSCHECK BACKUP;
```

Delete the expired backups. For example, issue: DELETE EXPIRED BACKUP;

Deleting Obsolete RMAN Backups Based on Retention Policies

The RMAN DELETE command supports an OBSOLETE option, which deletes backups that are no longer needed to satisfy specified recoverability requirements. You can delete files that are obsolete according to the configured default retention policy, or another retention policy that you specify as an option to the DELETE OBSOLETE command. As with other forms of the DELETE command, the files deleted are removed from backup media, deleted from the recovery catalog, and marked as DELETED in the control file.

If you specify the DELETE OBSOLETE command with no arguments, then RMAN deletes all obsolete backups defined by the configured retention policy. For example:

```
DELETE OBSOLETE;
```

NEW QUESTION 134

- (Topic 6)

You are working on a 24x7 database. You want to design a backup strategy for your database that uses user managed backups. You want to be able to perform all backups while the database remains online.

Which statement about performing user-managed backups in a 24x7 environment is true?

- A. You must have change tracking enabled in your database
- B. Your database must be running in NOARCHIVELOG mode
- C. To back up a tablespace, it must be in backup mode
- D. To back up a tablespace, it must first be taken offline

Answer: C

NEW QUESTION 139

- (Topic 6)

While working on a data problem, Curt, Bill, Ben, Mike, and Matt introduced a vast amount of corrupted data into the database. Pablo has discovered this problem and he needs you to recover the database to the point in time prior to the introduction of the corruption. The logical corruption was introduced at 6:30 p.m. on September 6, 2008.

Which of the following would be the correct commands to use to restore the database to a point in time before the corruption?

- A. restore database until time '06-SEP-2008 06:30:00'); recover database until time '06-SEP-2008 06:30:00'); alter database open;
- B. restore database until time '06-SEP-2008 06:30:00'); recover database until time '06-SEP-2008 06:30:00'); alter database open resetlogs;
- C. restore database until time '06-SEP-2008 18:29:55'); recover database until time '06-SEP-2008 18:29:55'); alter database open resetlogs;
- D. restore database until time '06-SEP-2008 18:29:55'); alter database open resetlogs;
- E. restore database until time '06-SEP-2008 18:29:55'); recover database; alter database open resetlogs;

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E

Answer: C

NEW QUESTION 144

- (Topic 6)

You realize that the control file is damaged in your production database. After restoring the control file from autobackup, what is the next step that you must do to proceed with the database recovery?

- A. Mount the database
- B. Open the database in NORMAL mode
- C. Open the database in RESTRICTED mode
- D. Open the database with the RESETLOGS option

Answer: A

Explanation:

After restoring the control file, it is time to load the restored control file onto the database instance, to issue STARTUP MOUNT command to load it.

NEW QUESTION 149

- (Topic 6)

Which is NOT a valid way of backing up a control file?

- A. Backing up the control file to trace
- B. Copying the existing control file of the database to the backup location during a hot backup
- C. Copying the existing control file of the database to the backup location during a cold backup
- D. Creating a backup control file
- E. Using the create controlfile command

Answer: B

NEW QUESTION 150

- (Topic 6)

You have backed up your database using image copies. You have lost the SYSTEM tablespace and need to restart your database as quickly as possible. What is the correct solution?

- A. Restore the SYSTEM tablespace from the last backup set and then recover the database.
- B. Restore the SYSTEM tablespace image copy using the restore command and then restore the database.
- C. Use the switch datafile command to instantly switch to the datafile copy, recover the tablespace, and open the database.
- D. The database is not recoverable in this situation with image copies.
- E. Manually copy the datafile image copy to the correct location and then manually restore the database from SQL*Plus.

Answer: C

NEW QUESTION 153

- (Topic 6)

You need to restore your database back to 9/30/2008 at 18:00. In what order would you run the following commands to complete this task?

- A. restore controlfile until time '09/30/2008:18:00:00';
- B. restore database until time '09/30/2008:18:00:00';
- C. restore spfile until time '09/30/2008:18:00:00';
- D. recover database until time '09/30/2008:18:00:00';
- E. alter database open resetlogs;
- F. alter database open;
- G. b, d, e
- H. b, d, f
- I. c, a, b, d, e
- J. c, a, b, d, f
- K. a, b, d, e

Answer: A

NEW QUESTION 155

- (Topic 6)

You have lost all your SYSTEM tablespace datafiles (system_01.dbf and system_02.dbf) and the database has crashed. What would be the appropriate order of operations to correct the situation?

- A. Mount the database with the startup mount command.
- B. Take the SYSTEM data file offline with the alter database command.
- C. Restore the SYSTEM_01.dbf data file from backup media with the required archived redo logs.
- D. Restore all SYSTEM tablespace-related datafiles from backup media.
- E. Issue the recover tablespace SYSTEM command.
- F. Issue the recover data file SYSTEM_01.dbf command.
- G. Open the database with the alter database open command.
- H. Open the database with the alter database open RESETLOGS command.
- I. a, c, f, g
- J. b, d, e, h
- K. a, b, c, f, g
- L. d, a, e, g
- M. b, c, f, e, g

Answer: D

Explanation:

Because there is NO controlfile damaged, and there is NO PITR recover, you don't need to use RESETLOGS option.

NEW QUESTION 158

- (Topic 6)

The database is running in the ARCHIVELOG mode. It has three redo log groups with one member each. One of the redo log groups has become corrupted. You have issued the following command during the recovery of a damaged redo log file:

```
ALTER DATABASE CLEAR UNARCHIVED LOGFILE GROUP 3;
```

Which action should you perform immediately after using this command?

- A. You should perform a log switch
- B. You should make a backup of the database
- C. You should switch the database to the NOARCHIVELOG mode
- D. You should shut down the database instance and perform a complete database recovery

Answer: B

NEW QUESTION 163

- (Topic 6)

Your database is up and running and one of your three control files is accidentally erased. You start RMAN and run the following command:

```
RESTORE CONTROLFILE FROM AUTOBACKUP;
```

Which of the following statements is true? (Choose all that apply.)

- A. The command restores only the missing control file.
- B. The command restores all the control files.
- C. The command fails because the database is running.
- D. This is the correct way to address this problem.
- E. This is not the correct way to address this problem.

Answer: CE

Explanation:

During the database running, the control files are locked by the database instance, you must shutdown the database and startup at NOMOUNT status to restore a missing control file.

And you have to open database with RESETLOGS option, due to control file change.

NEW QUESTION 164

- (Topic 6)

What is the correct order of steps to perform an online database backup?

- a: alter database begin backup; b: alter database end backup;
- c: Back up the database datafiles. d: Back up the archive log files.
- e: alter system switch logfile;

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

Answer: E

Explanation:

(link)

To back up online read/write tablespaces in an open database:

1. Use the DBA_DATA_FILES data dictionary view to identify all of the data files in the tablespace.
2. ALTER TABLESPACE users BEGIN BACKUP;
3. Back up the online data files of the online tablespace with operating system commands.
4. ALTER TABLESPACE users END BACKUP;
5. ALTER SYSTEM ARCHIVE LOG CURRENT; Archive the unarchived redo logs so that the redo required to recover the tablespace backup is archived.

NEW QUESTION 167

- (Topic 6)

Your database is running in ARCHIVELOG mode. One of the data files, USERDATA01.dbf, in the USERS tablespace is damaged and you need to recover the file until the point of failure. The backup for the datafile is available.

Which three files would be used in the user-managed recovery process performed by the database administrator (DBA)? (Choose Three)

- A. Redo logs
- B. Control file
- C. The latest backup of only the damaged data file
- D. The latest backup of all the data file in the USERS tablespace
- E. Temporary files of temporary tablespace
- F. Archive Logs since the latest backup to point of failure

Answer: ACF

NEW QUESTION 169

- (Topic 6)

A database is running In ARCHIVELOG mode. It has two online redo log groups and each group has one member.

A LGWR Input/output (I/O) falls due to permanent media failure that has resulted In the loss of redo log file and the LWGR terminates causing the instance to crash. The steps to recover from the loss of a current redo log group member in the random order are as follow.

- 1) Restore the corrupted redo log group.
- 2) Restore from a whole database backup.
- 3) Perform incomplete recovery.
- 4) Relocate by renaming the member of the damaged online redo log group to a new location.
- 5) Open the database with the RESETLOGS option.
- 6) Restart the database instance.
- 7) Issue a checkpoint and clear the log.

Identify the option with the correct sequential steps to accomplish the task efficiently.

- A. 1, 3, 4, and 5
- B. 7, 3, 4. and 5
- C. 2, 3, 4, and 5
- D. 7, 4, 3. and 5
- E. Only 6 is required

Answer: C

Explanation:

Recovering After Losing All Members of an Online Redo Log Group

If a media failure damages all members of an online redo log group, then different scenarios can occur depending on the type of online redo log group affected by the failure

and the archiving mode of the database.

If the damaged online redo log group is current and active, then it is needed for crash recovery; otherwise, it is not. Table 30-4 outlines the various recovery scenarios.

If the Group Is...	Then...	And You Should...
Inactive	It is not needed for crash recovery	Clear the archived or unarchived group.
Active	It is needed for crash recovery	Attempt to issue a checkpoint and clear the log; if impossible, then you must either use Flashback Database or restore a backup and perform incomplete recovery up to the most recent available redo log.
Current	It is the redo log that the database is currently writing to	Attempt to clear the log; if impossible, then you must either use Flashback Database or restore a backup and perform incomplete recovery up to the most recent available redo log.

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

NEW QUESTION 172

- (Topic 6)

You have discovered that one of three control files has been lost. What steps would you follow to recover that control file?

- A. Shut down the database.
- B. Restore a control-file copy from backup media.
- C. Use the create control file command to create a new control file.
- D. Copy the backup control file into place.
- E. Create a new copy of the control file from one of the surviving control files.
- F. Recover the database using the recover database using BACKUP CONTROLFILE command.
- G. Start up the database.
- H. a, b, f, g
- I. c, f, g
- J. a, d, f, g
- K. a, f, g
- L. a, e, g

Answer: E

NEW QUESTION 176

- (Topic 6)

Which of the following statements is true when the database is in ARCHIVELOG mode and tablespaces are in hot backup mode?

- A. Archive log generation is suspended until the tablespaces are taken out of hot backup mode.
- B. Datafiles are not written to during hot backups.
- C. Changes to the database are cached during the backup and not written to the datafiles to ensure that the datafiles are consistent when recovered.
- D. The datafile headers are not updated during the backup.
- E. The way data is written to the online redo logs is unchanged during the backup.

Answer: D

NEW QUESTION 181

- (Topic 7)

You are using recovery Manager (RMAN) with a recovery catalog to backup up your production database. The backups and the archived redo log files are copied to a tape drive on a daily basis. The database was open and transactions were recorded in the redo logs. Because of fire in the building you lost your servers having the production database and the recovery catalog database. The archive log files generated after the last backup are intact on one of the remote locations. While performing a disaster recovery of the production database what is the next step that you must perform after restoring the data files and applying archived redo logs?

- A. Open the database in NORMAL mode
- B. Open the database in read-only mode
- C. Open the database in RESTRICTED mode
- D. Open the database with the RESETLOGS option

Answer: D

Explanation:

Recovering the Database After a Disaster

The procedure for disaster recovery is similar to the procedure for recovering the database with a backup control file in NOCATALOG mode. If you are restoring the database to a new host, then you should also review the considerations described in "Restoring a Database on a New Host".

This scenario assumes that the Linux server on which your database was running has been damaged beyond repair. Fortunately, you backed up the database to Oracle Secure Backup and have the tapes available. The scenario assumes the following:

- ? Oracle Database is already installed on the new host.
- ? You are restoring the database to a new Linux host with the same directory structure as the old host.
- ? You have one tape drive containing backups of all the data files and archived redo logs through log 1124, and autobackups of the control file and server parameter file.
- ? You do not use a recovery catalog with the database.

To recover the database on the new host:

1. If possible, restore or re-create all relevant network files such as tnsnames.ora and listener.ora and a password file.
2. Start RMAN and connect to the target database instance.

At this stage, no initialization parameter file exists. If you have set ORACLE_SID and ORACLE_HOME, then you can use operating system authentication to connect as SYSDBA. For example, start RMAN as follows:

```
% rman
RMAN> CONNECT TARGET
/
```

3. Specify the DBID for the target database with the SET DBID command, as described in "Restoring the Server Parameter File".

For example, enter the following command: SET DBID 676549873;

4. Run the STARTUP NOMOUNT command.

When the server parameter file is not available, RMAN attempts to start the instance with a dummy server parameter file.

5. Allocate a channel to the media manager and then restore the server parameter file from autobackup. For example, enter the following command to restore the server parameter file from Oracle Secure Backup:

```
RUN
{
ALLOCATE CHANNEL c1 DEVICE TYPE sbt; RESTORE SPFILE FROM AUTOBACKUP;
}
```

6. Restart the instance with the restored server parameter file. STARTUP FORCE NOMOUNT;
7. Write a command file to perform the restore and recovery operation, and then execute the command file.

The command file should do the following:

- a. Allocate a channel to the media manager.
- b. Restore a control file autobackup (see "Performing Recovery with a Backup Control File and No Recovery Catalog").
- c. Mount the restored control file.
- d. Catalog any backups not recorded in the repository with the CATALOG command.
- e. Restore the data files to their original locations. If volume names have changed, then run SET NEWNAME commands before the restore operation and perform a switch after the restore operation to update the control file with the new locations for the data files, as shown in the following example.
- f. Recover the data files. RMAN stops recovery when it reaches the log sequence number specified.

```
RMAN> RUN
{
# Manually allocate a channel to the media manager ALLOCATE CHANNEL t1 DEVICE TYPE sbt;
# Restore autobackup of the control file. This example assumes that you have
# accepted the default format for the autobackup name. RESTORE CONTROLFILE FROM AUTOBACKUP;
# The set until command is used in case the database
# structure has changed in the most recent backups, and you want to
# recover to that point in time. In this way RMAN restores the database
# to the same structure that the database had at the specified time. ALTER DATABASE MOUNT;
SET UNTIL SEQUENCE 1124 THREAD 1; RESTORE DATABASE;
RECOVER DATABASE;
}
```

The following example of the RUN command shows the same scenario except with new file names for the restored data files:

```

RMAN> RUN
{
# If you must restore the files to new locations,
# use SET NEWNAME commands:
SET NEWNAME FOR DATAFILE 1 TO '/dev/vgd_1_0/rlvt5_500M_1'; SET NEWNAME FOR DATAFILE 2 TO '/dev/vgd_1_0/rlvt5_500M_2'; SET NEWNAME FOR
DATAFILE 3 TO '/dev/vgd_1_0/rlvt5_500M_3'; ALLOCATE CHANNEL t1 DEVICE TYPE sbt;
RESTORE CONTROLFILE FROM AUTOBACKUP; ALTER DATABASE MOUNT;
SET UNTIL SEQUENCE 124 THREAD 1; RESTORE DATABASE;
SWITCH DATAFILE ALL; # Update control file with new location of data files. RECOVER DATABASE;
}
8. If recovery was successful, then open the database and reset the online logs: ALTER DATABASE OPEN RESETLOGS;

```

NEW QUESTION 182

- (Topic 7)

Identify two situations in which you can use Data Recovery Advisor for recovery. (Choose two.)

- A. The user has dropped an important table that needs to be recovered.
- B. The database files are corrupted when the database is open.
- C. You are not able to start up the database instance because the required database files are missing.
- D. The archived log files are missing for which backup is not available.

Answer: BC

NEW QUESTION 184

- (Topic 7)

David managed to accidentally delete the datafiles for database called DSL. He called Heber and Heber tried to help but he managed to delete the control files of the database. Heber called Bill and Bill saved the day.

They are using a recovery catalog for this database.

What steps did Bill perform to recover the database and in what order?

- A. Restored the control file with the RMAN restore controlfile command.
- B. Mounted the DSL instance with the alter database mount command.
- C. Restored the datafiles for the DSL database with the RMAN restore command.
- D. Opened the DSL database with the alter database open resetlogs command.
- E. Recovered the datafiles for the DSL database with the RMAN recover command.
- F. Started the DSL instance.
- G. Connected to the recovery catalog with RMAN.
- H. a, b, c, d, e, f, g
- I. b, c, d, g, f, e, a
- J. g, f, a, b, c, e, d
- K. c, a, d, b, f, e, g
- L. g, f, a, b, e, c, d

Answer: C

Explanation:

About Recovery with a Backup Control File

If all copies of the current control file are lost or damaged, then you must restore and mount a backup control file. You must then run the RECOVER command, even if no data files have been restored, and open the database with the RESETLOGS option. If some copies of the current control file are usable, however, then you can follow the procedure in "Responding to the Loss of a Subset of the Current Control Files" and avoid the recovery and RESETLOGS operation.

When RMAN is connected to a recovery catalog, the recovery procedure with a backup control file is identical to recovery with a current control file. The RMAN metadata missing from the backup control file is available from the recovery catalog. The only exception is if the database name is not unique in the catalog, in which case you must use SET DBID command before restoring the control file.

1. Start RMAN and connect to a target database.
2. Start the target instance without mounting the database. RMAN>STARTUP NOMOUNT;
3. Restore the control file
RMAN> SET DBID 320066378; # (Optional) If the database name is not unique, you need to specify the DBID
RMAN> RUN
{
SET CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO
'autobackup_format';
RESTORE CONTROLFILE FROM AUTOBACKUP;
}
4. Start the target instance with mounting the database. RMAN>STARTUP MOUNT;
5. Restore the data files; RMAN>RESTORE DATABASE;
6. Recover the database; RMAN>RECOVER DATABASE;
7. Open the database with RESETLOGS option; RMAN> ALTER DATABASE OPEN RESETLOGS;

NEW QUESTION 187

- (Topic 7)

You are using Recovery Manager (RMAN) for backup and recovery operations with a recovery catalog. You have been taken database backups every evening. On November 15, 2007, at 11:30 AM, you were informed that the USER_DATA tablespace was accidentally dropped. On investigation, you found that the tablespace existed until 11:00 AM, and important transactions were done after that.

So you decided to perform incomplete recovery until 11:00 AM. All the archive logs needed to perform recovery are intact. In NOMOUNT state you restored the control file that has information about the USER_DATA tablespace from the latest backup. Then you mounted the database. Identify the next set of commands that are required to accomplish the task?

```
A. RMAN> run
{
  SET UNTIL TIME 'Nov 15 2007 11:00:00';
  RESTORE DATABASE;
  RECOVER DATABASE;
}

B. RMAN> run
{
  SET UNTIL TIME 'Nov 15 2007 11:00:00';
  RESTORE DATABASE;
  RECOVER DATABASE USING BACKUP CONTROLFILE;
}

C. RMAN> run
{
  RESTORE DATABASE;
  RECOVER DATABASE UNTIL TIME 'Nov 15 2007 11:00:00';
}

D. RMAN> run
{
  RESTORE TABLESPACE user_data;
  RECOVER TABLESPACE user_data UNTIL TIME 'Nov 15 2007 11:00:00';
}
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: A

NEW QUESTION 191

- (Topic 7)

While performing the backup of the flash recovery area, you notice that one of the archived redo log files is missing. You have multiple destinations for archived redo log files. What implications does it have on the backup of the flash recovery area?

- A. The backup fails because one of the archived redo log files is missing.
- B. The backup succeeds but it would be without the missing archived log.
- C. During backup, you are prompted for the alternative destination for the missing archived redo log file.
- D. The backup succeeds because it fails over to one of the alternative archived redo log destinations.

Answer: D

NEW QUESTION 195

- (Topic 8)

Examine the following command that is used to duplicate a database on the same host:

```
RMAN> RUN
{
  ALLOCATE AUXILIARY CHANNEL aux 1 DEVICE TYPE DISK; DUPLICATE TARGET DATABASE TO auxdb SKIP READONLY;
}
```

Which two statements describe the effect after the database is duplicated successfully? (Choose two)

- A. The data files of the read-only tablespaces in the target database are not duplicated
- B. The read-only tablespaces in the target database are still defined in new the database
- C. The read-only tablespaces in the target database are changed to online after duplication
- D. The data files of the read-only tablespaces in the target database get duplicated
- E. The read-only tablespaces in the target database are not defined in the new database

Answer: AB

NEW QUESTION 196

- (Topic 8)

Which of the following are prerequisite steps to transport a database? (Choose all that apply.)

- A. Query the V\$TRANSPORTABLE_PLATFORMS view in the source database to determine if the intended destination is listed.
- B. Verify that there are no restrictions or limitations that the source or destination database may encounter.
- C. Verify that the source and destination have the same Oracle version, critical updates, patch-set version, and patch- set exceptions.
- D. Determine if you will perform the conversion on the source or destination platform.
- E. None of the above.

Answer: ABD

NEW QUESTION 198

- (Topic 8)

Which two statements are correct about database transportation? (Choose two.)

- A. The source and target platforms must be the same
- B. Redo logs, control files and temp files are also transported
- C. The transported database must have the same database identifier (DBID) as the source database and cannot be changed
- D. The COMPATIBLE parameter must be set to 10.0.0.0 or higher and the database must be opened in readonly mode before being transported
- E. Recovery Manager (RMAN) is used to convert the necessary data files of the database if the target platform is different and has different endian format

Answer: DE

NEW QUESTION 200

- (Topic 8)

Which two statements are true about a duplicate database that is created by using the DUPLICATE command in RMAN? (Choose two.)

- A. It is a copy or a subset of the target database.
- B. It is opened in RESTRICT mode after a duplicating operation.
- C. It is created by using backups and archived redo log files from the target database.
- D. It is created with the same database identifier (DBID) as that of target database.

Answer: AC

Explanation:

Duplicating a Database

Overview of RMAN Database Duplication

Database duplication is the use of the DUPLICATE command to copy all or a subset of the data in a source database. The duplicate database (the copied database) functions entirely independently from the source database (the database being copied).

Purpose of Database Duplication

If you copy a database with operating system utilities instead of the DUPLICATE command, then the DBID of the copied database remains the same as the original database. To register the copy database in the same recovery catalog with the original, you must change the DBID with the DBNEWID utility (see Oracle Database Utilities). In contrast, the DUPLICATE command automatically assigns the duplicate database a different DBID so that it can be registered in the same recovery catalog as the source database.

Backup-Based Duplication

In backup-based duplication, RMAN creates the duplicate database by using pre-existing RMAN backups and copies. This technique of duplication uses one of the following mutually exclusive subtechniques: Duplication without a target database connection, RMAN obtains metadata about backups from a recovery catalog.

Duplication without a target database connection and without a recovery catalog. RMAN obtains metadata about where backups and copies reside from BACKUP LOCATION. Duplication with a target database connection. RMAN obtains metadata about backups

from the target database control file or from the recovery catalog. How RMAN Duplicates a Database

For backup-based duplication, the principal work of the duplication is performed by the auxiliary channels. These channels correspond to a server session on the auxiliary instance on the destination host. For active database duplication the primary work is performed by target channels. RMAN must perform database point-in-time recovery, even when no explicit point in time is provided for duplication. Point-in-time recovery is required because the online redo log files in the source database are not backed up and cannot be applied to the duplicate database. The farthest point of recovery of the duplicate database is the most recent redo log file archived by the source database.

As part of the duplicating operation, RMAN automates the following steps:

1. Creates a default server parameter file for the auxiliary instance if the following conditions are true:
 - ? Duplication does not involve a standby database.
 - ? Server parameter files are not being duplicated.
 - ? The auxiliary instance was not started with a server parameter file.
2. Restores from backup or copies from active database the latest control file that satisfies the UNTIL clause requirements.
3. Mounts the restored or copied backup control file from the active database.
4. Uses the RMAN repository to select the backups for restoring the data files to the auxiliary instance. This step applies to backup-based duplication.
5. Restores and copies the duplicate data files and recovers them with incremental backups and archived redo log files to a noncurrent point in time.
6. Shuts down and restarts the database instance in NOMOUNT mode.
7. Creates a new control file, which then creates and stores the new DBID in the data files.
8. Opens the duplicate database with the RESETLOGS option and creates the online redo log for the new database.

NEW QUESTION 204

- (Topic 8)

The following query will provide what information about transportable tablespaces for the current database? (Choose all that apply.)

```
SELECT d.platform_name "Source", t.platform_name "Compatible Targets",
       endian_format
FROM v$transportable_platform t, v$database d
WHERE t.endian_format =
       (select endian_format
        from v$transportable_platform t, v$database d
        where d.platform_name = platform_name);
```

- A. The list of target platforms having the same endian format as the source database
- B. The list of target platforms requiring endian conversion
- C. The list of target platforms that will not require endian conversion
- D. The list of all target platforms that can receive transportable tablespaces from the source database
- E. None of the above

Answer: AC

NEW QUESTION 207

- (Topic 8)

You are managing the APPPROD database as a DBA. You plan to duplicate this database in the same system with the name DUPDB. You issued the following RMAN commands to create a duplicate database:

```
RMAN> CONNECT target sys/sys@APPPROD
RMAN> DUPLICATE TARGET DATABASE TO dupdb FROM ACTIVE DATABASE
      DB_FILE_NAME_CONVERT '/oracle/oradata/prod/', '/scratch/oracle/oradata/
dupdb/'
      SPILE
      PARAMETER_VALUE_CONVERT '/oracle/oradata/prod/', '/scratch/oracle/oradata/
dupdb/'
      SET SGA_MAX_SIZE = '300M'
      SET SGA_TARGET = '250M'
      SET LOG_FILE_NAME_CONVERT '/oracle/oradata/prod/redo/', '/scratch/oracle/
oradata/dupdb/redo/';
```

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

- A. The source database should be open.
- B. The target database should be in ARCHIVELOG mode if it is open.
- C. RMAN should be connected to both the instances as SYSDBA.
- D. The target database backups should be copied to the source database backup directories.
- E. The password file must exist for the source database and have the same SYS user password as the target.

Answer: BCE

NEW QUESTION 212

- (Topic 9)

Which options must you configure while performing an automated Tablespace Point-in- Time Recovery (TSPITR) by using Recovery Manager (RMAN)?

- A. New channels for restore and recovery tasks
- B. New name for the data files of the tablespace
- C. Auxiliary name for the data files of the tablespace
- D. Auxiliary destinations for an auxiliary set of data files

Answer: D

NEW QUESTION 216

- (Topic 9)

Because of a logical corruption in your production database, you wanted to perform Tablespace Point in Time Recovery (TSPITR). But before you start the recovery, you queried the TS_PITR_OBJECTS_TO_BE_DROPPED view and realized that there are a large number of objects that would be dropped when you start the recovery by using this method. You want to preserve these objects. Which option must you use to perform TSPITR and preserve the object?

- A. Perform Export before TSPITR and Import after TSPITR
- B. Move objects to another schema that has the same tablespace assigned
- C. Perform Incomplete Recovery before TSPITR with the Log Sequence Number (LSN)
- D. Perform Incomplete Recovery before TSPITR with the System Change Number (SCN)

Answer: A

NEW QUESTION 217

- (Topic 9)

What RMAN command is used to execute a tablespace point-in-time recovery?

- A. recover
- B. duplicate
- C. restore
- D. copy
- E. None of the above

Answer: A

NEW QUESTION 222

- (Topic 9)

Why should you back up a duplicated tablespace after a TSPITR is complete?

- A. The tablespace cannot be duplicated or restored to any point in time after the duplication.
- B. The tablespace cannot be duplicated or restored to the point in time before the duplication.
- C. The entire database cannot be restored after a TSPITR, so a backup is required.
- D. You cannot bring the tablespace online until its been backed up.
- E. There is no requirement to do so, as RMAN will back up the tablespace after the TSPITR.

Answer: B

NEW QUESTION 227

- (Topic 10)

You perform a backup using the following BACKUP command: RMAN> BACKUP AS COMPRESSED BACKUPSET DATABASE;
Which statement is true of this command?

- A. A different procedure is required to restore a database from compressed backups
- B. The AS COMPRESSED clause of the BACKUP command provided by RMAN is used to create compressed backup sets and image copies.
- C. Using this command to create backups minimizes the bandwidth consumed
- D. Using this command to create backups improves the performance of the backup process

Answer: C

NEW QUESTION 231

- (Topic 10)

You want to take the backup of the USERS tablespace. It has a single data file of 900 MB. You have tape drives of 300 MB each. The SBT channel is configured for the RMAN. To accomplish the backup, you issued the following RMAN command:

RMAN> BACKUP SECTION SIZE 300M TABLESPACE users;

Which two statements are true regarding the execution of the above command? (Choose two.)

- A. The RMAN parallelizes the backup although the parallelism is not set for a channel.
- B. The backup piece size will be limited to 300 MB.
- C. The operation is accomplished using the default channel available.
- D. Three channels for the tape drive must be configured by setting the parallelism to three.

Answer: BC

Explanation:

SECTION SIZE sizeSpec Specifies the size of each backup section produced during a data file backup.

By setting this parameter, RMAN can create a multisection backup. In a multisection backup, RMAN creates a backup piece that contains one file section, which is a contiguous range of blocks in a file. All sections of a multisection backup are the same size. You can create a multisection backup for a data file, but not a data file copy.

File sections enable RMAN to create multiple steps for the backup of a single large data file. RMAN channels can process each step independently and in parallel, with each channel producing one section of a multisection backup set.

If you specify a section size that is larger than the size of the file, then RMAN does not use multisection backup for the file. If you specify a small section size that would produce more than 256 sections, then RMAN increases the section size to a value that results in exactly 256 sections.

Depending on where you specify this parameter in the RMAN syntax, you can specify different section sizes for different files in the same backup job.

Note: You cannot use SECTION SIZE with MAXPIECESIZE or with INCREMENTAL LEVEL 1.

NEW QUESTION 233

- (Topic 11)

You discover that your Recycle Bin contains two tables with the same name, MY_TABLE.

You also have a table named MY_TABLE in your schema. You execute the following statement:

FLASHBACK TABLE my_table TO BEFORE DROP RENAME TO my_table2; What will be the result of executing this statement?

- A. One of the tables is recovered from the Recycle Bin using a First In First Out (FIFO) approach.
- B. One of the tables is recovered from the Recycle Bin using a Last In First Out (LIFO) approach.
- C. Both the tables are recovered from the Recycle Bin with one table renamed to MY_TABLE2 and the other to a system-generated name.
- D. None of the tables are recovered from the Recycle Bin, and the statement returns an error.

Answer: B

NEW QUESTION 237

- (Topic 11)

View the Exhibit and examine the data manipulation language (DML) operations that you performed on the NEWEMP table. Note that the first two updated are not listed by the Flashback Versions Query.

What could be the reason? Exhibit:

```
SQL> UPDATE newemp SET sal=sal+100 WHERE ename='FORD';
1 row updated.
SQL> UPDATE newemp SET sal=sal+100 WHERE ename='FORD';
1 row updated.
SQL> ALTER TABLE newemp DROP COLUMN comm;
Table altered.
SQL> COMMIT;
Commit complete.
SQL> UPDATE newemp SET sal=sal+100 WHERE ename='FORD';
1 row updated.
SQL> COMMIT;
Commit complete.
SQL> SELECT versions_xid AS XID,
       2 versions_startscn AS START_SCN,
       3 versions_endscn AS END_SCN,
       4 versions_operation AS OPERATION, sal
       5 FROM newemp VERSIONS BETWEEN SCN MINVALUE AND MAXVALUE
       6 WHERE ename='FORD';
```

XID	START_SCN	END_SCN	OPERATION	SAL
07002E00B1030000	1705446		U	3300
		1705446		3200

- A. The first two updated were not explicitly committed.
- B. ALTER TABLE caused the recycle bin to release the space.
- C. The data definition language (DDL) operation caused a log switch.
- D. Flashback Versions Query stops producing versions of rows that existed before a change in the table structure.

Answer: C

NEW QUESTION 238

- (Topic 11)

Which of the following can be used in conjunction with a Flashback Versions Query to filter the results? (Choose all that apply.)

- A. A range of SCN values
- B. A list of SCN values
- C. A starting and ending timestamp
- D. Minimum and maximum sequence values
- E. A list of sequence values

Answer: AC

NEW QUESTION 241

- (Topic 11)

A user named Arren is executing this query:

```
select table_name, operation, undo_sql from  
flashback_transaction_query t, (select versions_xid as xid  
from employees versions between scn minvalue and maxvalue  
where employee_id = 123) e where t.xid = e.xid;
```

When the query runs, he receives an ORA-01031: insufficient privileges error. Since the user owns the employees table, you know that it is not the problem. Which of the following SQL statements will correct this problem?

- A. GRANT SELECT ANY TRANSACTION TO ARREN;
- B. GRANT SELECT ON FLASHBACK_TRANSACTION_QUERY TO ARREN;
- C. GRANT SELECT ANY TRANSACTION TO ARREN;
- D. GRANT FLASHBACK TO ARREN;
- E. GRANT SELECT ANY VIEW TO ARREN;

Answer: A

NEW QUESTION 243

- (Topic 11)

The RECYCLEBIN parameter is set to ON for your database. You drop a table, PRODUCTS, from the SCOTT schema. Which two statements are true regarding the outcome of this action? (Choose two)

- A. All the related indexes and views are automatically dropped
- B. The flashback drop feature can recover only the table structure
- C. Only the related indexes are dropped whereas views are invalidated
- D. The flashback drop feature can recover both the table structure and its data

Answer: CD

NEW QUESTION 246

- (Topic 11)

Which pseudocolumn could you use to identify a unique row in a Flashback Versions Query?

- A. XID
- B. VERSIONS_PK
- C. VERSIONS_XID
- D. VERSIONS_UNIQUE

Answer: C

NEW QUESTION 251

- (Topic 11)

On Friday at 11:30 am you decided to flash back the database because of a user error that occurred at 8:30 am. Which option must you use to check whether a flashback operation can recover the database to the specified time?

- A. Check the alert log file
- B. Query the V\$FLASHBACK_DATABASE_LOG view
- C. Query the V\$RECOVERY_FILE_DEST_SIZE view
- D. Query the V\$FLASHBACK_DATABASE_STAT view
- E. Check the value assigned for the UNDO_RETENTION parameter

Answer: B

Explanation:

To query the V\$FLASHBACK_DATABASE_LOG to get the lowest SCN or the nearest TIMESTAMP to decide the recovery possibility.

NEW QUESTION 252

- (Topic 11)

Which method would you use to undo the changes made by a particular transaction without affecting the changes made by other transactions?

- A. Point-in-time recovery
- B. Execute the ROLLBACK command with transaction number
- C. Flashback the database to before the transaction was committed
- D. Determine all the necessary undo SQL statements from FLASHBACK_TRANSACTION_QUERY and use them for recovery

Answer: D

NEW QUESTION 255

- (Topic 11)

You executed the following query:

```
SELECT oldest_flashback_scn, oldest_flashback_time FROM V$FLASHBACK_DATABASE_LOG;
```

Considering that all the redo logs are available, what information can you derive from the output of the preceding query?

- A. The time when the last flashback operation in your database was performed
- B. The time when the first flashback operation in your database was performed
- C. The approximate time and the lowest system change number (SCN) to which you can flash back your database
- D. The system change number (SCN) and the time when the Flashback Database was enabled in the database instance

Answer: C

Explanation:

V\$FLASHBACK_DATABASE_LOG displays information about the flashback data. Use this view to help estimate the amount of flashback space required for the current workload.

Column	Datatype	Description
OLDEST_FLASHBACK_SCN	NUMBER	Lowest system change number (SCN) in the flashback data, for any incarnation
OLDEST_FLASHBACK_TIME	DATE	Time of the lowest SCN in the flashback data, for any incarnation
RETENTION_TARGET	NUMBER	Target retention time (in minutes)
FLASHBACK_SIZE	NUMBER	Current size (in bytes) of the flashback data
ESTIMATED_FLASHBACK_SIZE	NUMBER	Estimated size of flashback data needed for the current target retention

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

NEW QUESTION 258

- (Topic 12)

Which three actions are required to configure the Flashback Database? (Choose three.)

- A. set Flash Recovery Area
- B. enable Flashback logging
- C. create FLASHBACK tablespace
- D. start the database in the ARCHIVELOG mode
- E. start the database in the NOARCHIVELOG mode

Answer: ABD

NEW QUESTION 260

- (Topic 12)

Which are the two prerequisites before setting up Flashback Data Archive? (Choose two.)

- A. Flash recovery area must be defined
- B. Undo retention guarantee must be enabled.
- C. Database must be running in archivelog mode.
- D. Automatic undo management must be enabled.
- E. The tablespace in which the Flashback Data Archive is created must have automatic segment space Management (ASSM).

Answer: DE

NEW QUESTION 262

- (Topic 12)

When setting up the Flashback Data Archive, which of these key parameters are required? (Choose all that apply.)

- A. Tablespace name
- B. Storage quota
- C. Retention
- D. Table name
- E. Create a default archive

Answer: AC

NEW QUESTION 263

- (Topic 12)

A user performs an update on a table. Shortly after committing the transaction, they realize that they had an error in their WHERE clause causing the wrong rows to be updated. Which Flashback option would allow you to undo this transaction and restore the table to its previous state?

- A. Flashback Drop
- B. Flashback Query
- C. Flashback Versions Query
- D. Flashback Transaction Query
- E. Flashback Table

Answer: E

Explanation:

FLASHBACK TABLE

Purpose

Use the FLASHBACK TABLE statement to restore an earlier state of a table in the event of human or application error. The time in the past to which the table can be flashed back is dependent on the amount of undo data in the system. Also, Oracle Database cannot restore a table to an earlier state across any DDL operations that change the structure of the table.

NEW QUESTION 268

- (Topic 12)

You are working as a DBA on the decision support system.

There is a business requirement to track and store all transactions for at least three years for a few tables in the database.

Automatic undo management is enabled in the database. Which configuration should you use to accomplish this task?

- A. Enable Flashback Data Archive for the tables.
- B. Specify undo retention guarantee for the database.
- C. Enable supplemental logging for the database.
- D. Query V\$UNDOSTAT to determine the amount of undo that will be generated and create an undo tablespace for that size.
- E. Create Flashback Data Archive on the tablespace on which the tables are stored.

Answer: A**NEW QUESTION 269**

- (Topic 12)

You plan to execute the following command to perform a Flashback Database operation in your database:

```
SQL> FLASHBACK DATABASE TO TIMESTAMP (SYSDATE - 5/24);
```

Which two statements are true about this? (Choose two)

- A. The database must have multiplexed redo log files
- B. The database must be in the MOUNT state to execute the command
- C. The database must be in the NOMOUNT state to execute the command
- D. The database must be opened in RESTRICTED mode before this operation
- E. The database must be opened with the RESETLOGS option after the flashback operation

Answer: BE**NEW QUESTION 270**

- (Topic 12)

To clean up old records that are in a Flashback Data Archive and are past the retention period, what must the DBA do?

- A. TRUNCATE the archive table.
- B. DROP the Flashback Data Archive.
- C. Nothing; expired rows are automatically removed.
- D. Nothing; expired rows are moved to an archive table.
- E. Delete entries from the archive where the metadata date retained is greater than the retention period.

Answer: C**NEW QUESTION 271**

- (Topic 13)

Which of the following information will be gathered by the SQL Test Case Builder for the problems pertaining to SQL-related problems? (Choose all that apply.)

- A. ADR diagnostic files
- B. all the optimizer statistics
- C. initialization parameter settings
- D. PL/SQL functions, procedures, and packages
- E. the table and index definitions and actual data

Answer: BCD**Explanation:**

The information gathered by SQL Test Case Builder includes:

- ? the query being executed,
- ? table and index definitions (but not the actual data),
- ? PL/SQL functions, procedures, and packages,
- ? optimizer statistics,
- ? and initialization parameter settings.

NEW QUESTION 276

- (Topic 13)

Examine the following values of the initialization parameters in the database having the SID ORCL:

```
BACKGROUND_DUMP_DEST=/u01/app/oracle/product/11.1.0/db_1/bdump USER_DUMP_DEST=/u01/app/oracle/product/11.1.0/db_1/udump
```

```
CORE_DUMP_DEST=/u01/app/oracle/product/11.1.0/db_1/cdump DIAGNOSTIC_DEST=
```

The environment variables have the following value: ORACLE_BASE=/u01/app/oracle

```
ORACLE_HOME=/u01/app/oracle/product/11.1.0/db_1
```

What is the location of the Automatic Diagnostic Repository (ADR) home?

- A. /u01/app/oracle/product/11.1.0/db_1
- B. /u01/app/oracle

- C. \$ORACLE_HOME/bdump
- D. \$ORACLE_HOME/log

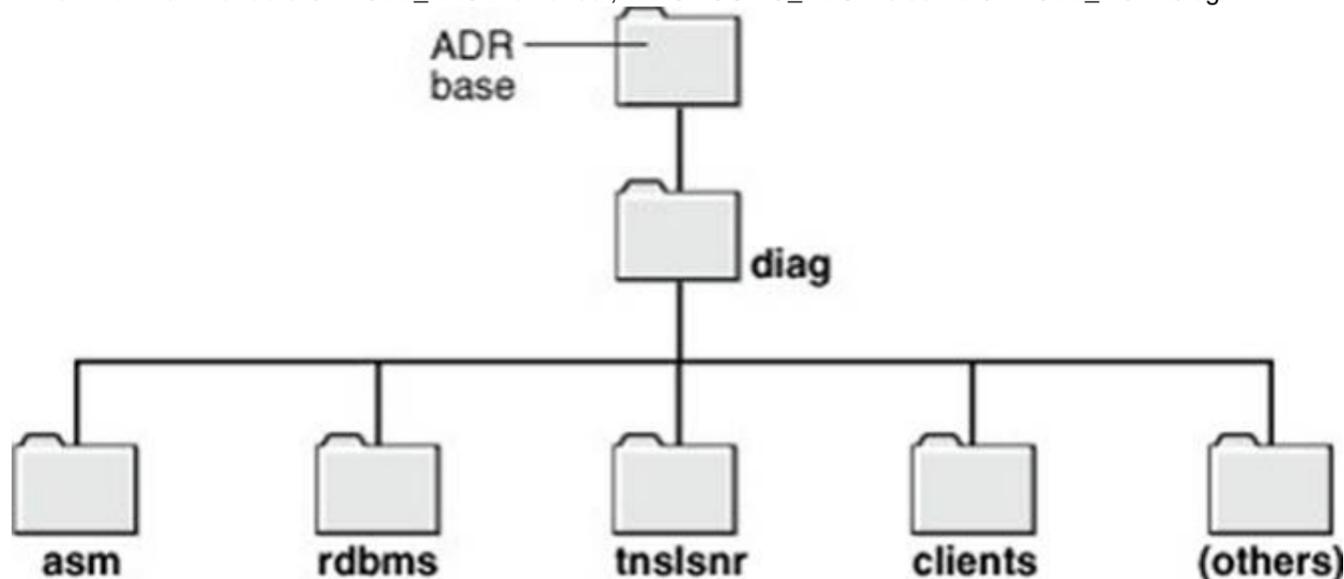
Answer: B

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.



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

NEW QUESTION 280

- (Topic 13)

You need to perform a block media recovery on the tools01.dbf data file in the SALES database by using Recovery Manager (RMAN).

Which two are the prerequisites to perform this operation? (Choose two)

- A. You must configure block change tracking file
- B. You must have first level 1 backups for RMAN to restore blocks
- C. You must ensure that the SALES database is mounted or open
- D. You must have full or level 0 backups for RMAN to restore blocks
- E. You must take the tools01.dbf data file offline before you start a block media recovery

Answer: CD

Explanation:

Prerequisites for Block Media Recovery (link)

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

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

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

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

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

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

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

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

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

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

NEW QUESTION 284

- (Topic 13)

During the installation of Oracle Database 11g, you do not set ORACLE_BASE explicitly. You selected the option to create a database as part of the installation. How would this environment variable setting affect the installation?

- A. The installation terminates with an error.
- B. The installation proceeds with the default value without warnings and errors.
- C. The installation proceeds with the default value but it would not be an OFA-compliant database.
- D. The installation proceeds with the default value but a message would be generated in the alert log file.

Answer: D

NEW QUESTION 288

- (Topic 13)

Which of the following is a prerequisite for running DBMS_TDB.CHECK_DB to a successful completion?

- A. The database must be in read-write mode.
- B. The database must have no external files.
- C. The database must open in read-only mode.
- D. The database must be mounted but not opened.

Answer: C

NEW QUESTION 291

- (Topic 13)

Which of the following is NOT an advantage of block media recovery (BMR)?

- A. Reduced MTTR.
- B. Datafiles remain offline while corrupt blocks are repaired.
- C. Datafiles remain online while corrupt blocks are repaired.
- D. A and C

Answer: B

Explanation:

Overview of Block Media Recovery (link) Purpose of Block Media Recovery

You can use block media recovery to recover one or more corrupt data blocks within a data file. Block media recovery provides the following advantages over data file media recovery:

? Lowers the mean time to recover (MTTR) because only blocks needing recovery are restored and recovered

? Enables affected data files to remain online during recovery

Without block media recovery, if even a single block is corrupt, then you must take the data file offline and restore a backup of the data file. You must apply all redo generated for the data file after the backup was created. The entire file is unavailable until media recovery completes. With block media recovery, only the blocks actually being recovered are unavailable during the recovery.

Block media recovery is most useful for physical corruption problems that involve a small, well-known number of blocks. Block-level data loss usually results from intermittent, random I/O errors that do not cause widespread data loss, and memory corruptions that are written to disk. Block media recovery is not intended for cases where the extent of data loss or corruption is unknown and the entire data file requires recovery. In such cases, data file media recovery is the best solution.

NEW QUESTION 292

- (Topic 13)

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

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

Answer: AC

Explanation:

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

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

? If environment variable ORACLE_BASE is not set, DIAGNOSTIC_DEST is set to

ORACLE_HOME/log

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

NEW QUESTION 293

- (Topic 14)

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

MEMORY_MAX_TARGET=0 MEMORY_TARGET=500M PGA_AGGREGATE_TARGET=90M

SGA_TARGET=270M

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

- A. The MEMORY_MAX_TARGET parameter is automatically set to 500 MB.
- B. The value of the MEMORY_MAX_TARGET parameter remains zero till it is changed manually.
- C. The PGA_AGGREGATE_TARGET and SGA_TARGET parameters are automatically set to zero.
- D. The lower bounds of PGA_AGGREGATE_TARGET and SGA_TARGET parameters are set to 90 MB and 270 MB, respectively.

Answer: AD

NEW QUESTION 294

- (Topic 14)

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

- A. Setting a nonzero value for SGA_TARGET
- B. Configuring Automatic Shared Memory Management
- C. Configuring Automatic Memory Management
- D. Setting a nonzero value for SGA_MAX_SIZE and PGA_AGGREGATE_TARGET
- E. None of the above

Answer: C

NEW QUESTION 297

- (Topic 14)

Observe the structure of the table employees: The table contains 8475 records.

DESC employees

Name	Null?	Type
emp_no	NOT NULL	VARCHAR(5) PRIMARY KEY
emp_fname	NOT NULL	VARCHAR(15)
emp_lname	NOT NULL	VARCHAR(15)
emp_dob	NOT NULL	DATE
emp_dtojoin	NOT NULL	DATE
emp_salary	NOT NULL	NUMBER(5,2)

One of the employees wants to know the names of all employees of the company. For this, he fires the following query:

```
SELECT * FROM EMPLOYEES ORDER BY emp_fname;
```

Since the operation performed on executing the query cannot fit into memory, it requires disk space to complete the operation.

Which of the following types of segments will Oracle allocate to complete the operation and to provide the required result?

- A. Rollback segment
- B. Temporary segment
- C. Data segment
- D. Index segment

Answer: B

NEW QUESTION 301

- (Topic 14)

Examine the parameter setting in your database:

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

```
SQL> SHOW PARAMETER SGA_MAX_SIZE
```

NAME	TYPE	VALUE
sga_max_size	big integer	808M

Which statement is correct about the database?

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

Answer: D

NEW QUESTION 302

- (Topic 14)

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

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

Answer: ABC

NEW QUESTION 307

- (Topic 14)

Note the following parameters settings in your database:

```
SGA_MAX_SIZE      = 1024M
SGA_TARGET        = 700M
DB_8K_CACHE_SIZE = 124M
LOG_BUFFER        = 200M
```

You issued the following command to increase the value of DB_8K_CACHE_SIZE: SQL> ALTER SYSTEM SET DB_8K_CACHE_SIZE=140M; What would happen?

- A. It will fail because DB_8K_CACHE_SIZE parameter cannot be changed dynamically
- B. It will be successful only if the memory is available from the auto tuned components
- C. It will fail because an increase in DB_8K_CACHE_SIZE cannot be accommodated within SGA_TARGET
- D. It will fail because an increase in the DB_8K_CACHE_SIZE cannot be accommodated within SGA_MAX_SIZE

Answer: D

Explanation:

This is an actually mathematic question.

NEW QUESTION 310

- (Topic 15)

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

Only profiles that significantly improve SQL performance were implemented.

Select	SQL Text	Parsing Schema	SQL ID	Statistics	SQL Profile	Index	Restructure SQL	Miscellaneous	Error Date
<input type="radio"/>	SELECT NULL AS table_cat, Lower...	SYSMAN	351qjn3w9uflh	✓	(99.9%) ✓				7/12/07
<input checked="" type="radio"/>	SELECT EXECUTION_ID, STATUS, STATUS_DET...	SYSMAN	lyk8b9986atk7		(60%) ✓	(97.9%) ✓			7/12/07
<input type="radio"/>	SELECT /*+ INDEX(sqlobj\$ (signature cate...	SYS	8b75qwpan202v					✓	7/12/07
<input type="radio"/>	select OBJOID, CLSOID, RUNTIME, PRI, JO...	SYS	8vf1dhwgk1xy5					✓	7/12/07
<input type="radio"/>	select smontabv.cnt, smontab.time_mp, ...	SYS	4q8mr2bvy6qr					✓	7/12/07
<input type="radio"/>	select l.ts#,l.file#,l.block#,nv (l.boj...	SYS	1qu8t96d2bdmu					✓	7/12/07
<input type="radio"/>	select obj#, dataobj#, part#, hbounden...	SYS	130dvvr5s8bqn					✓	7/12/07
<input type="radio"/>	select privilege#,level from sysauth\$ co...	SYS	0b6b2pawb74n					✓	7/12/07
<input type="radio"/>	select value(p\$) from "XDB"."XDB\$RESOURC...	SYS	23y48d28wkg2r					✓	7/12/07
<input type="radio"/>	SELECT obj_type, plan_id, name, flags, L...	SYS	0n1napsmccz0c					✓	7/12/07

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

Answer: C

NEW QUESTION 313

- (Topic 15)

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

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

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

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

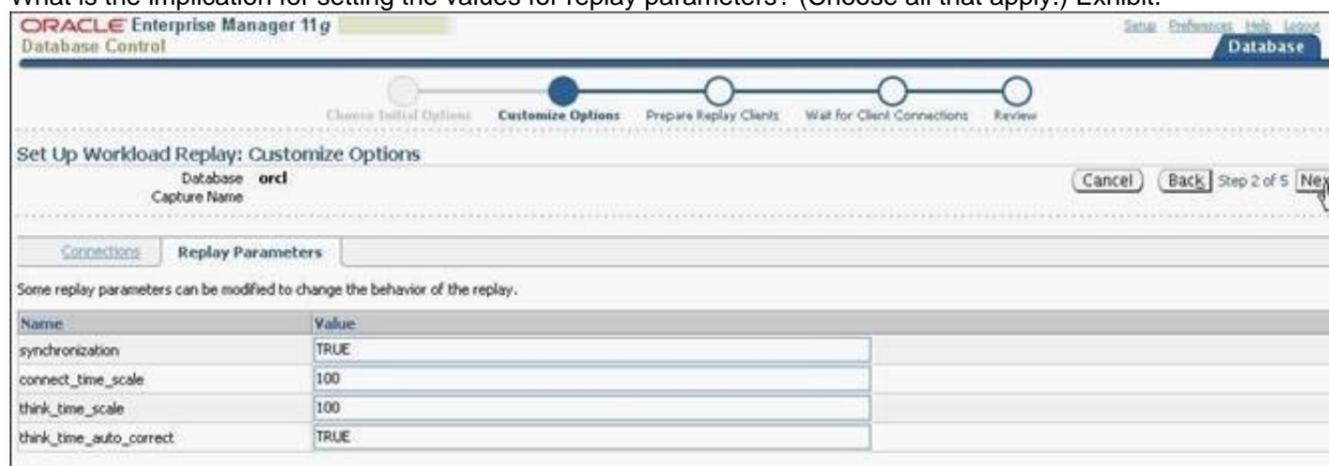
Answer: B

NEW QUESTION 315

- (Topic 15)

View the Exhibit to examine the replay settings for replay parameters.

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



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

Answer: ACD

NEW QUESTION 319

- (Topic 15)

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

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

Answer: CE

Explanation:

9.3 Workload Capture Restrictions

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

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

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

NEW QUESTION 322

- (Topic 15)

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

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

Answer: BC

Explanation:

Online Patching

Regular patches typically contain .o (object) files and/or .a (archive) libraries, and therefore require a relink of the RDBMS binary. Online patches, however, contain .so files, which are dynamic/shared libraries, and do not require a relink of the RDBMS binary. Consequently, since a relink is not needed, you can apply or roll

back online patches while the RDBMS instance is running. This simplifies administration, because no downtime is needed, and also results in a much quicker turnaround time for installing or de-installing Online Patches.

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

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

- ? Linux x86
- ? Linux x86_64
- ? HP-UX Itanium
- ? Solaris SPARC 64-bit
- ? Solaris AMD 64-bit
- ? AIX (AIX 6.1 and later)

About Patch Conflicts:

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

NEW QUESTION 324

- (Topic 15)

View the Exhibit for some of the parameter settings. You start a session and issue the following command:

SQL>CREATE INDEX emp_ename ON emp(ename) TABLESPACE users INVISIBLE;

What is the outcome of the above command? Exhibit:

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

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

Answer: A

NEW QUESTION 326

- (Topic 15)

You plan to have a larger moving window size for the default system-defined moving window baseline because you want to use the adaptive threshold.

Which statement factors in this consideration while increasing the size of the moving window?

- A. The collection level for the AWR should be set to BASIC.
- B. The moving window size must be less than Undo Retention.
- C. The moving window size should be greater than the Automatic Workload Repository (AWR) retention period.
- D. The moving window size should be equal to or less than the Automatic Workload Repository (AWR) retention period.

Answer: D

Explanation:

Moving Window Baseline ([link](#))

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

NEW QUESTION 330

- (Topic 15)

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

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

Answer: B

NEW QUESTION 333

- (Topic 16)

Which steps are mandatory to enable Direct NFS?

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

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

Answer: B

NEW QUESTION 337

- (Topic 16)

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

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

Answer: A

NEW QUESTION 341

- (Topic 16)

You need to perform an online table redefinition of an existing SALES table to partition it into two tablespaces TBS1 and TBS2. The SALES table has a materialized view, materialized log, indexes, referential integrity constraint, and triggers with the PRECEDES clause existing on it. What action is required for dependent objects when you perform online table redefinition?

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

Answer: A

Explanation:

When performing the online table redefinition, you will:

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

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

The following are the end results of the redefinition process:

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

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

Note:

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

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

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

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

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

NEW QUESTION 342

- (Topic 16)

Observe the following PL/SQL block: BEGIN

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

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

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

Answer: D

NEW QUESTION 343

- (Topic 16)

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

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

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

- A. 1, 2, 3

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

Answer: C

NEW QUESTION 345

- (Topic 16)

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

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

Answer: DE

NEW QUESTION 350

- (Topic 16)

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

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

What would be the outcome of this command?

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

Answer: C

NEW QUESTION 354

- (Topic 16)

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

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

Answer: ACD

NEW QUESTION 359

- (Topic 16)

Evaluate the following statements:

```
CREATE TABLE purchase_orders ( po_id NUMBER(4),  
po_date TIMESTAMP, supplier_id NUMBER(6), po_total NUMBER(8,2),  
CONSTRAINT order_pk PRIMARY KEY(po_id)) PARTITION BY RANGE(po_date) (  
PARTITION Q1 VALUES LESS THAN (TO_DATE('1-apr-2007','d-mon-yyyy'), PARTITION Q2 VALUES LESS THAN (TO_DATE('1-jul-2007','d-mon-yyyy'),  
PARTITION Q3 VALUES LESS THAN (TO_DATE('1-oct-2007','d-mon-yyyy'), PARTITION Q4 VALUES LESS THAN (TO_DATE('1-jan-2008','d-mon-yyyy')));  
CREATE TABLE purchase_order_items ( po_id NUMBER(4) NOT NULL,  
product_id NUMBER(6) NOT NULL, unit_price NUMBER(8,2),  
quantity NUMBER(8),  
CONSTRAINT po_items_fk FOREIGN KEY (po_id) REFERENCES purchase_orders(po_id)) PARTITION BY REFERENCE(po_items_fk);
```

What are the two consequences of the above statements? (Choose two.)

- A. Partitions of PURCHASE_ORDER_ITEMS have system-generated names.
- B. Both PURCHASE_ORDERS and PURCHASE_ORDER_ITEMS tables are created with four partitions each.
- C. Partitions of the PURCHASE_ORDER_ITEMS table exist in the same tablespaces as the partitions of the PURCHASE_ORDERS table.
- D. The PURCHASE_ORDER_ITEMS table inherits the partitioning key from the parent table by automatically duplicating the key columns.
- E. Partition maintenance operations performed on the PURCHASE_ORDER_ITEMS table are automatically reflected in the PURCHASE_ORDERS table.

Answer: BC

NEW QUESTION 360

- (Topic 16)

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

Which of these actions is appropriate?

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

Answer: B

NEW QUESTION 361

- (Topic 17)

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

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

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

Which statement describes the consequence? Exhibit:

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

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

Answer: A

Explanation:

Requisition of Calibrate I/O (link)

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

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

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

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

NEW QUESTION 364

- (Topic 17)

Your database initialization parameter file has the following entry: SEC_MAX_FAILED_LOGIN_ATTEMPTS=3

Which statement is true regarding this setting?

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

Answer: A

NEW QUESTION 369

- (Topic 17)

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

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

Answer: ADE

NEW QUESTION 372

- (Topic 17)

Evaluate the following block of code:

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

What is the outcome of the above code?

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

Answer: C

NEW QUESTION 377

- (Topic 17)

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

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

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

Answer: B

Explanation:

9.1.1.2 FILESYSTEMIO_OPTIONS Initialization Parameter

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

FILESYSTEMIO_OPTIONS can be set to one of the following values:

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

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

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

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

disk_asynch_io = true filesystemio_options = asynch

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

NEW QUESTION 380

- (Topic 17)

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

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

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

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

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

Answer: BC

Explanation:

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

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

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

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

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

NEW QUESTION 384

- (Topic 17)

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

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

Answer: ABD

NEW QUESTION 385

- (Topic 17)

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

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

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

PL/SQL procedure successfully completed.

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

PL/SQL procedure successfully completed.

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

PL/SQL procedure successfully completed.

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

PL/SQL procedure successfully completed.
```

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

Answer: D

NEW QUESTION 390

- (Topic 17)

To control the execution of a server process when it is receiving bad packets from a potentially malicious client, you set the SEC_PROTOCOL_ERROR_FURTHER_ACTION initialization parameter as follows:

SQL> ALTER SYSTEM SET SEC_PROTOCOL_ERROR_FURTHER_ACTION = Drop,10;

What is the significance of this setting?

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

Answer: B

NEW QUESTION 394

- (Topic 17)

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

Request	WebCache		Web Server		Application Server		Database Service		Server Time Details
	Count	HR	Count	Average Time (ms)	Count	Average Time (ms)	Count	Average Time (ms)	
...	1	n/a	2	108,874.17	2	327.05	1	1,144.50	
...	11	n/a	10	2,548.05	1	1,878.05	1	1,002.00	
...	31	n/a	38	4,428.17	31	818.71	2	2,212.38	
...	1	n/a	3	232.85	6	58.05	2	2,651.51	
...	21	n/a	24	4,554.20	24	2,048.53	1	1,045.00	
...	51	n/a	54	2,056.62	32	446.05	1	1,768.00	
...	1	n/a	1	1,936.63	1	1,711.05	1	1,757.00	
...	64	n/a	64	2,301.12	40	438.45	1	1,711.05	
...	4	n/a	5	7,276.21	2	1,329.50	1	1,602.00	
...	1	n/a	1	2,090.00	1	1,251.05	1	1,452.00	

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

Answer: A

NEW QUESTION 395

- (Topic 18)

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

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

Answer: D

Explanation:

The bulk load operation can be done by setting up a file watcher, then raise the event to inform the scheduler, it is kind of application based event. After the bulk load job, you can use Scheduler raised event, e.g. JOB COMPLETE, to fire the index rebuild job. There are two kinds of events consumed by the Scheduler:
 ? Events raised by your application
 An application can raise an event to be consumed by the Scheduler. The Scheduler reacts to the event by starting a job. For example, when an inventory tracking system notices that the inventory has gone below a certain threshold, it can raise an event that starts an inventory replenishment job. See "Starting Jobs with Events Raised by Your Application".
 ? File arrival events raised by a file watcher
 You can create a file watcher—a Scheduler object introduced in Oracle Database 11g Release 2—to watch for the arrival of a file on a system. You can then configure a job to start when the file watcher detects the presence of the file. For example, a data warehouse for a chain of stores loads data from end-of-day revenue reports uploaded from the point-of-sale systems in the stores. The data warehouse load job starts each time a new end-of-day report arrives.

NEW QUESTION 396

- (Topic 19)

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

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

Answer: ACE

NEW QUESTION 397

- (Topic 19)

How many different calendars does Oracle 11g support?

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

Answer: B**NEW QUESTION 399**

- (Topic 19)

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

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

Answer: BCE**NEW QUESTION 403**

- (Topic 19)

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

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

Answer: B**NEW QUESTION 406**

- (Topic 20)

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

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

Answer: E**NEW QUESTION 410**

- (Topic 20)

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

- A. All objects in the database
- B. Tables owned by SYS and SYSTEM
- C. Data dictionary tables
- D. NCHAR columns
- E. CHAR columns

Answer: E**NEW QUESTION 412**

- (Topic 20)

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

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

Answer: F**Explanation:**

Parameter type String

Syntax NLS_LANGUAGE = language

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

Modifiable ALTER SESSION

Range of values Any valid language name Basic Yes

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

NEW QUESTION 416

- (Topic 20)

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

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

Answer: BC

Explanation:

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

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

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

BINARY

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

LINGUISTIC

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

ANSI

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

NEW QUESTION 419

- (Topic 20)

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

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

Answer: B

NEW QUESTION 422

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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