PassQueen
a.k.a dumpsqueen.com

Oracle Database Administration II

Oracle 1z0-083

Version Demo

Total Demo Questions: 38

Total Premium Questions: 387

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Configuring for Recoverability	32
Topic 2, Using RMAN to Create Backups	55
Topic 3, Using RMAN to Perform Recovery	41
Topic 4, Configuring and Using Flashback Database	25
Topic 5, Diagnosing the Database	27
Topic 6, Managing Memory	11
Topic 7, Managing Database Performance	38
Topic 8, Managing Resources	10
Topic 9, Creating and Managing Multitenant Container Databases	24
Topic 10, Managing PDBs	42
Topic 11, Managing a CDB and PDBs	28
Topic 12, Mix Questions	54
Total	387

QUESTION NO: 1

A Data Guard broker configuration has fast-start failover enabled.

You must start the observer process from the DGMGRL command-line interface.

Which command should be used?

A. START OBSERVER;

B. ENABLE OBSERVER;

C. START FAST_START FAILOVER;

D. ENABLE FAST_START FAILOVER OBSERVER;

E. ALTER CONFIGURATION START OBSERVER;

ANSWER: A

Explanation:

`START OBSERVER` starts the Data Guard Broker observer process. The observer monitors the primary database and the fast-start failover target and can initiate a fast-start failover when the configured conditions are met.

The observer is normally started from a host separate from both the primary and standby databases to avoid a single host failure preventing observer operation. Oracle documents observer management commands in the Oracle Data Guard Broker documentation.

QUESTION NO: 2

Which two are true about Oracle Database Configuration Assistant (DBCA)? (Choose two.)

A. It can be launched from both the command line and the Universal Installer (OUI) interfaces.

B. It can be used to change the character set of an existing database.

C. It can be used to convert a non-container database to a container database.

D. It can be used to create container databases and non-container databases.

E. It can only be launched from the command line interface.

ANSWER: A D

Explanation:

Oracle Database Configuration Assistant can be started directly from the command line, including in interactive graphical mode and in silent mode using response-file or command-line parameters. It is also integrated with Oracle Universal Installer: when database software installation includes database configuration, Oracle Universal Installer invokes DBCA to perform the database-creation and configuration work. Thus, DBCA is available through both the command line and the Oracle installation workflow.

In Oracle Database 19c, DBCA supports creation of either a multitenant container database or a traditional non-container database. During database creation, its configuration choices allow the administrator to select the database architecture, including creating a container database and optionally creating initial pluggable databases. This enables an administrator to use the same assistant to provision databases appropriate for multitenant deployments as well as non-CDB environments where that architecture is still required or supported. Oracle documents DBCA as the standard utility for creating and configuring databases, with support for multitenant database creation and management. See the Oracle Database 19c Administrator's Guide and the DBCA multitenant documentation.

QUESTION NO: 3

Which three are true about Automatic Workload Repository (AWR) and Automatic Database Diagnostic Monitor (ADDM) in an Oracle multitenant environment?

A. AWR snapshots can be created in CDB$ROOT.

B. ADDM can run in a nonroot container.

C. AWR reports can be generated while connected to any container.

D. All AWR data is stored in the CDB$ROOT SYSAUX tablespace.

E. AWR snapshots can be created in a PDB.

F. No AWR data is stored in the CDB$ROOT SYSAUX tablespace.

ANSWER: A D E

Explanation:

`AWR snapshots can be created in CDB$ROOT.` is correct because the root container supports both the automatic AWR snapshot mechanism and manually requested snapshots. These root-level snapshots capture workload and performance statistics for the multitenant database.

`All AWR data is stored in the CDB$ROOT SYSAUX tablespace.` is correct for this multitenant AWR architecture. The AWR repository is maintained in the root container's SYSAUX tablespace, allowing the repository to retain both CDB-level information and information associated with PDB activity in a centralized location. This central repository also supports consolidated historical performance analysis across the CDB.

`AWR snapshots can be created in a PDB.` is correct because a PDB administrator can explicitly create a snapshot for the current PDB by using the `DBMS_WORKLOAD_REPOSITORY` package. This enables performance data collection focused on a specific pluggable database, while the repository remains centrally managed from the root's SYSAUX tablespace. Oracle documents multitenant monitoring and performance-management behavior in the Oracle Multitenant Administrator's Guide and AWR snapshot administration in the DBMS_WORKLOAD_REPOSITORY package reference.

QUESTION NO: 4

Which operating system group is NOT needed to perform an Oracle Database installation?

A. OSASM

B. OSKMDBA

C. OSDBA

D. OSOPER

E. OSRACDBA

ANSWER: A

Explanation:

OSASM is the operating-system authentication group associated with Oracle Automatic Storage Management (Oracle ASM), rather than a required group for installing Oracle Database software. Membership in this group grants the SYSASM administrative privilege, which is used to administer ASM instances and ASM disk groups. It is therefore relevant when configuring and administering Oracle Grid Infrastructure/ASM, normally for the Grid Infrastructure software owner.

An Oracle Database installation requires an operating-system database administrator group so that the database software owner can receive SYSDBA privileges. Oracle also supports additional database administrative groups for specific responsibilities, including operations, backup and recovery, Data Guard, encryption-key management, and RAC

administration. Those groups can be configured according to the environment and administrative separation requirements. OSASM, however, is specifically an ASM administration group and does not need to be created or assigned merely to install Oracle Database software. If a database later uses ASM storage, ASM administration remains a Grid Infrastructure responsibility, distinct from the basic database software installation process.

Oracle documents the purpose of the ASM administrator group and the database-related operating-system groups in the Oracle Database Installation Guide: Creating Operating System Privileges Groups. See also the Oracle ASM access authentication documentation.

QUESTION NO: 5

Which two are true about the automatic execution of operating system scripts when performing silent mode installation starting from Oracle Database 19c? (Choose two.)

A. The response file can specify the root or sudo password.

B. The installer will prompt for the root or sudo password.

C. Silent install always runs operating scripts automatically.

D. The response file must contain the root or sudo password.

E. The response file can specify the path of the sudo program.

ANSWER: B E

Explanation:

Oracle Database 19c supports automatic execution of the privileged operating-system configuration scripts during an Oracle Universal Installer silent installation when root-script automation is enabled in the response file. The response file can identify the sudo executable by using the sudo-path setting. This enables the installer to invoke the configured sudo utility when the selected root-script execution method is sudo, rather than requiring a fixed sudo location. The response file also supports related root-script automation settings, including the execution method and sudo user name.

Privileged credentials are not stored in the response file. Instead, when the selected automation method requires authentication, Oracle Universal Installer prompts for the applicable root or sudo password during installation. This avoids placing a reusable privileged password in a response file, which might otherwise be retained on disk or exposed through backups and file permissions. Therefore, automatic root-script execution can still be configured for a silent installation while interactive credential entry is requested only as needed to authorize the privileged operating-system actions.

See Oracle's instructions for automatically running Oracle Database root scripts and its silent-mode installation documentation.

QUESTION NO: 6

choose one You enable Flashback Data Archive for a table for which you want to track and save all transactions for four years. After some time, the requirement changes for keeping transactions from four to two years. You execute the following command to change the requirement: SQL> ALTER FLASHBACK ARCHIVE fdal MODIFY RETENTION 2 YEAR; What is the outcome?

A. All historical data older than two years, if any, is purged from the flashback archive FDA1.

B. An error is returned because the retention period cannot be reduced.

C. All the flashback data archives are purged and the subsequently created flashback archives are maintained for two years

D. All historical data older than two years, if any, is archived to flashback logs and the flashback archive is set to new retention time.

ANSWER: B

Explanation:

An error is returned because the retention period cannot be reduced. Oracle permits the `MODIFY RETENTION` clause of `ALTER FLASHBACK ARCHIVE` to increase a Flashback Data Archive retention period, but not to decrease it. Therefore, an archive that was created with a four-year retention period cannot be changed to two years using the displayed statement. The command is rejected rather than immediately removing historical rows or moving them elsewhere. This restriction protects the retention commitment associated with historical data recorded by Flashback Data Archive, which is intended to provide a long-term, automatically maintained history of changes to tracked tables. If a shorter retention requirement is needed, the administrator must use an appropriate archive lifecycle approach rather than expecting `MODIFY RETENTION` to purge existing history. Oracle documents the syntax and semantics of this statement, including the retention-modification restriction, in the ALTER FLASHBACK ARCHIVE SQL Reference. Further background on Flashback Data Archive administration and retention is available in the Oracle Database Development Guide.

QUESTION NO: 7

Examine this configuration:

1. CDB1 is a container database running in ARCHIVELOG mode.

2. Multiple uncommitted transactions are running in CDB1.

3. Redo log groups 1 and 2 are INACTIVE.

4. Redo log group 3 is the CURRENT group.

All members of redo log group 3 are lost before it is archived.

Examine these possible steps:

1. SHUTDOWN ABORT

2. STARTUP NOMOUNT

3. STARTUP MOUNT

4. ALTER DATABASE MOUNT

5. RESTORE DATABSE

6. RECOVER DATABASE NOREDO

7. RECOVER DATABASE UNTIL AVAILABLE

8. RESTORE ARCHIVELOG ALL

9. ALTER DATABSE OPEN

10.ALTER DATABASE OPEB RESETLOGS

Choose the minimum required steps in the correct order to recover the database.

A. 1, 3, 5, 8, 6, 10

B. 1, 3, 5, 6, 10

C. 1, 3, 5, 7, 10

D. 1, 2, 5, 7, 4, 10

E. 1, 3, 5, 6, 9

ANSWER: C

Explanation:

1, 3, 5, 7, 10 is the minimum required sequence. Because every member of the current online redo log group was lost before that group was archived, redo containing the most recent changes is unavailable. Oracle therefore cannot perform complete media recovery through the crash point. The database must be shut down abort because its current redo log is inaccessible, then started and mounted so that RMAN can restore the required database files without opening the database.

After `RESTORE DATABASE`, `RECOVER DATABASE UNTIL AVAILABLE` performs incomplete recovery by applying all redo that is available, stopping at the last recoverable point before the missing online redo. This is appropriate when a required current online redo log cannot be supplied for recovery. Since incomplete recovery creates a new database incarnation and discards redo after the selected recovery point, the database must be opened with `RESETLOGS`. Resetlogs initializes a new redo-log sequence and permits normal database operation to resume. Oracle documents incomplete recovery and the required RESETLOGS open in its RMAN recovery documentation and describes RESETLOGS behavior in the ALTER DATABASE SQL reference.

QUESTION NO: 8

Which three are located by using environment variables?

A. the Optimal Flexible Architecture (OFA) compliant path to store Oracle software and configuration files.

B. the location of Oracle Net Services configuration files

C. the list of a disk group names to be mounted by an Oracle Automatic Storage Management (ASM) instance at startup

D. default directories for temporary files used by temporary tablespaces

E. the temporary disk space used by Oracle Installer during installation

F. the maximum number of database files that can be opened by a database instance

ANSWER: A B E

Explanation:

The OFA-compliant path for Oracle software and configuration files is located through the `ORACLE_BASE` environment variable. Oracle Universal Installer uses this variable as the base directory from which standard OFA directory locations are derived, including software homes and Oracle-owned configuration locations. The location of Oracle Net Services configuration files can be specified with `TNS_ADMIN`. When it is set, Oracle Net clients and database tools use the directory named by `TNS_ADMIN` to find network configuration files such as `tnsnames.ora`, `sqlnet.ora`, and `listener.ora`. The temporary disk space used by Oracle Installer during installation is selected through the operating-system temporary-directory environment variables, normally `TMP` or `TMPDIR` depending on the platform. These variables allow an installation to place its temporary working files on a filesystem with sufficient available capacity rather than relying on the operating system default. Oracle documents the relevant installation environment variables, including `ORACLE_BASE` and temporary-directory variables, in its Oracle Database Installation Guide. Oracle Net configuration-file discovery and the role of `TNS_ADMIN` are described in the Oracle Net Services Administrator's Guide.

QUESTION NO: 9

Which two are true about instance recovery? (Choose two.)

A. It is not possible if an archived log is missing.

B. It is performed automatically after the database is opened; however, blocks requiring recovery are not available until they are recovered.

C. Setting FAST_START_MTTR_TARGET to a lower value reduces instance recovery time by causing dirty buffers to be written to disk more frequently, thereby reducing the number of I/Os needed during instance recovery.

D. It is performed by the Recovery Writer (RVWR) background process.

E. Setting FAST_START_MTTR_TARGET to a higher value reduces instance recovery time by causing The log writer to write more frquently, thereby reducing the number of I/Os needed during instance recovery.

F. It is performed automatically while the database remains in MOUNT state. Then the database is opened.

ANSWER: C F

Explanation:

Setting `FAST_START_MTTR_TARGET` to a lower value reduces the target mean time to recover by making Oracle perform checkpoints more aggressively. This increases the rate at which Database Writer processes write dirty buffers from the buffer cache to data files. Consequently, at the time of an unexpected instance failure, fewer changed blocks remain to be applied during roll-forward and the estimated duration of instance recovery is reduced. Oracle uses this parameter to balance recovery-time objectives against the normal I/O overhead required to maintain those objectives.

Instance recovery is an automatic startup operation following an instance failure. After the control file is opened and the database is mounted, Oracle applies the required online redo to recover committed changes and then performs transaction recovery for uncommitted work. This work must complete before the database is opened for normal user access, so recovery occurs while the database remains mounted and is followed by opening the database. Oracle documents the recovery objective and checkpoint behavior in its FAST_START_MTTR_TARGET reference and describes startup recovery in the Database Administrator's Guide.

QUESTION NO: 10

Which two are true about the execution of operating system scripts starting from Oracle Database 19c? (Choose two.)

A. orainstRoot.sh can be executed automatically by the Database installer by using sudo or root credentials.

B. root.sh can be executed automatically by the Database Installer only if it is provided with root credentials.

C. The sudo password can be specified in a response file.

D. root.sh can be executed automatically by the Database installer only by using sudo credentials.

E. The sudo password must be specified in a response file.

F. The root password cannot be specified in a response file.

ANSWER: A F

Explanation:

Beginning with Oracle Database 19c, Oracle Universal Installer can be configured to run the required privileged operating-system scripts automatically rather than requiring the administrator to run them manually after installation. This capability applies to scripts such as `orainstRoot.sh`, which performs Oracle Inventory-related privileged setup. The installer supports automatic execution through either appropriately supplied root access or a configured `sudo` mechanism, so `orainstRoot.sh can be executed automatically by the Database installer by using sudo or root credentials.` is correct.

Oracle deliberately protects privileged credentials in unattended installation configurations. A response file can contain installation settings and sudo configuration details, but it cannot contain the root account password. Consequently, `The root password cannot be specified in a response file.` is also correct. During an interactive installation, the installer can request required credentials when needed; for silent or automated deployment, administrators should use an approved privileged-execution approach, such as properly configured sudo authorization, instead of storing a root password in a response file. Oracle documents the automatic privileged-script feature and its credential-handling requirements in the Oracle Database installation guide. See Running Root Scripts Automatically and Using Response Files.

QUESTION NO: 11

You want to monitor the amount of time by which Redo Apply on a physical standby database is behind the primary database.

Which dynamic performance view displays this information?

A. V$DATAGUARD_STATS

B. V$ARCHIVE_DEST_STATUS

C. V$MANAGED_STANDBY

D. V$DATABASE_BLOCK_CORRUPTION

E. V$RECOVERY_FILE_DEST

ANSWER: A

Explanation:

V$DATAGUARD_STATS displays Data Guard performance statistics, including transport lag, apply lag, apply finish time, and estimated startup time. The apply lag statistic indicates the delay between the redo generated on the primary database and the redo applied on the standby database.

Oracle documents the view and its Data Guard statistics in the Oracle Database Reference.

QUESTION NO: 12

Choose three. Which three are true about Recovery Manager (RMAN) active database duplication?

A. Pre-existing backups of the database to be duplicated are required for the "push" method of duplication.

B. Pre-existing backups of the database to be duplicated are required for the "pull" method of duplication.

C. Pre-existing backups of the database to be duplicated are always used if they exist.

D. Image copies can be used to duplicate the active database.

E. The active database must be duplicated while in OPEN state.

F. No connection to the recovery catalog is required for active database duplication.

G. Backup sets can be used to duplicate the active database.

ANSWER: D F G

Explanation:

Image copies can be used to duplicate the active database because RMAN can copy the source database files directly over the network to the auxiliary instance during an active duplication. This is the default transfer approach when duplicating from an active database. Backup sets can also be used to duplicate the active database when the duplication is configured to use backup sets; RMAN creates backup sets from the source files and transfers them to the auxiliary database. Using backup sets can reduce network traffic because RMAN can apply compression and can omit unused blocks.

No connection to the recovery catalog is required for active database duplication because RMAN can obtain the metadata needed for the operation from the source database control file. The source and auxiliary connections, together with the required initialization and network configuration, allow RMAN to perform the duplication without a catalog repository. A recovery catalog may still be used if available, but it is optional rather than a prerequisite. Oracle documents both the image-copy and backup-set approaches for active database duplication and describes catalog-free operation through the control file. See Oracle Database Backup and Recovery User's Guide: Duplicating Databases and Oracle RMAN DUPLICATE command reference.

QUESTION NO: 13

Which two are true about Oracle Database Configuration Assistant (DBCA) templates? (Choose two.)

A. The General Purpose of Transaction Processing templates are most suitable when concurrency and recoverability are key criteria.

B. Oracle DBCA templates can store only logical structure and not database files.

C. New templates can only be created by modifying an existing user-created template.

D. The Data Warehouse template is most suitable when transaction response time is the key criterion.

E. Oracle DBCA templates can be used to create new databases and duplicate existing databases.

ANSWER: A E

Explanation:

The General Purpose or Transaction Processing templates are intended for online transaction processing workloads. These workloads commonly involve many simultaneous users and sessions performing short, frequent transactions. Consequently, the template configuration is designed around the operational priorities of concurrency and recoverability, making it appropriate where reliable transaction processing and the ability to recover data after failures are important requirements.

Oracle Database Configuration Assistant can use templates as part of database-creation workflows, and it also supports duplicating an existing database. A template provides predefined database characteristics and configuration choices so that deployments can be created consistently. DBCA's duplication capability creates a copy of a source database using the DBCA workflow, which is useful for provisioning an equivalent environment without manually recreating its configuration from scratch. Oracle documents DBCA as a tool for creating databases from supplied or custom templates and for duplicating databases. See Oracle's Database Administration Quick Start Guide and the Oracle Database Administrator's Guide for DBCA database creation and configuration procedures.

QUESTION NO: 14

Choose three. In which three situations can you use threshold server-generated alerts to help diagnose and fix problems?

A. when free space in a dictionary managed tablespace falls below a specified percentage

B. when free space in a locally managed tablespace falls below a specified percentage

C. when the total number of locked user account exceeds a specific value

D. when the number of logons per second exceeds a specific value

E. when a resumable statement is suspended

F. when the total number of logons exceeds a specific value

ANSWER: B C D

Explanation:

Threshold server-generated alerts are based on database metrics that Oracle continuously evaluates against warning and critical threshold values. A locally managed tablespace can generate a tablespace space-usage alert when its free space falls below the configured percentage, allowing an administrator to react before an allocation fails. Oracle also exposes the number of locked user accounts as a metric that can be monitored with thresholds; this helps identify potential account-locking or security-policy issues requiring administrative attention. Logons per second is likewise a rate-based performance metric for which warning and critical thresholds can be set. An unexpectedly high logon rate can indicate connection storms, application connection-pooling problems, or unusual workload activity. These alerts are recorded in the alert framework and can be viewed through Enterprise Manager or queried from alert-related data dictionary views, enabling timely diagnosis and response. Oracle documents the available server-generated alert metrics, including tablespace space usage, user locks, and logons per second, and explains that thresholds may be configured for metrics that support them. See Oracle Database Administrator's Guide: Managing Space for Schema Objects and Oracle Database Reference: DBA_OUTSTANDING_ALERTS.

QUESTION NO: 15

Which three are true about SGA memory management in a multitenant database? (Choose three.)

A. Setting DB_CACHE_SIZE for a PDB guarantees a minimum amount of Database Buffer Cache memory for that PDB.

B. The SHARED_POOL_SIZE setting for a PDB can be up to 80% of the SHARED_POOL_SIZE setting of the root container.

C. Setting SHARED_POOL_SIZE for a PDB guarantees a minimum amount of Shared Pool memory for that PDB.

D. The SGA_TARGET setting for a PDB cannot exceed 50% of the SGA_TARGET value of the root container.

E. Setting SGA_TARGET for a PDB guarantees a minimum amount of SGA memory for that PDB.

F. SGA_MIN_SIZE cannot be set for a PDB.

ANSWER: A C D

Explanation:

When Oracle Database uses automatic shared memory management in a multitenant environment, selected SGA component parameters can be set at the PDB level to reserve minimum memory for that PDB. Setting `DB_CACHE_SIZE` in a PDB establishes a minimum database buffer cache allocation for the PDB. This reservation helps ensure that the PDB retains sufficient cache capacity even while other PDBs are active. Similarly, setting `SHARED_POOL_SIZE` in a PDB establishes a minimum shared-pool allocation for that PDB, protecting memory needed for such activities as SQL parsing, library-cache use, and dictionary-cache use.

`SGA_TARGET` at PDB level represents an upper limit on SGA memory that the PDB may use; it is not a guaranteed minimum. Oracle constrains this value so that a PDB's `SGA_TARGET` cannot exceed 50 percent of the root container's `SGA_TARGET`. This limit prevents one PDB from being configured to consume an excessive portion of the CDB-wide SGA. Oracle documents the PDB-level memory parameters, their reservation behavior, and the applicable limits in its Managing Memory in a CDB guide and the DB_CACHE_SIZE parameter reference.

QUESTION NO: 16

Choose three. Which three are true about Oracle Grid Infrastructure for a Standalone Server?

A. It requires Oracle ASMLib to manage Automatic Storage Management (ASM) disks.

B. It creates one disk group during installation.

C. It includes both Oracle Restart and Oracle Automatic Storage Management (ASM) software.

D. Automatic Storage Management (ASM) requires that O/S groups OSASM and OSDBA be assigned as secondary groups for its installation owner.

E. It requires Oracle ASM Fitter Driver (ASMFD) to manage Automatic Storage Management (ASM) disks.

F. It requires the operating system ORACLE_BASE environment variable to be predefined before installation.

ANSWER: B C D

Explanation:

Oracle Grid Infrastructure for a Standalone Server is the Oracle software installation that provides Oracle Restart together with Oracle Automatic Storage Management (ASM). Oracle Restart uses the Oracle Clusterware stack on a single host to monitor and automatically restart supported components, including ASM, listeners, and configured database instances. During a standard Grid Infrastructure installation configured for ASM, the installer requires the administrator to select disks and create an ASM disk group; this provides the initial ASM storage configuration needed by the installation. The Grid Infrastructure installation owner is normally placed in the central inventory group as its primary operating-system group and is assigned the operating-system ASM administrator and database administrator groups as supplementary groups.

Membership in the OSASM group grants the privileges needed to administer ASM using SYSASM, while membership in OSDBA supports the required database-administration operating-system authentication and interoperability. These group assignments allow the Grid Infrastructure owner to manage the ASM instance and participate correctly in the Oracle database administration model. Oracle documents the single-server deployment and its Oracle Restart functionality in the Oracle Grid Infrastructure installation options, and documents the required ASM-related operating-system groups in the operating-system users and groups guide.

QUESTION NO: 17

Which two are true about changing the LOCAL_UNDO_ENABLED property to false in a CDB? (Choose two.)

A. After the change, only a common user with the required privilege can create an undo tablespace in CDB&ROOT.

B. Any new PDB and existing PDBs are automatically configured to use the default undo tablespace in CDB$ROOT.

C. After the change, only one undo tablespace can exist in CDB$ROOT.

D. After the change, any user with the required privilege can create an undo tablespace in the PDBs.

E. Undo tablespaces existing in PDBs must be dropped before the change.

F. After the change, each existing PDB has to be reopened for the new undo mode to take effect.

ANSWER: A B

Explanation:

Setting `LOCAL_UNDO_ENABLED` to `FALSE` changes the CDB to shared undo mode. In this mode, the root container owns the undo tablespace used for undo management by the CDB and its PDBs. Consequently, PDBs—both those already present and newly created—are configured to use the default undo tablespace in `CDB$ROOT`. This is the defining operational characteristic of shared undo mode and ensures that undo is centrally managed rather than maintained separately in every PDB.

Creating an undo tablespace in `CDB$ROOT` requires a common user with the appropriate administrative privilege. Only common users can operate in the root as users of the CDB, so a suitably privileged common user is required for this root-level tablespace administration. Oracle documents the distinction between local and shared undo modes, including the use of a root undo tablespace when local undo is disabled, in the Oracle Database Multitenant Administrator's Guide. See also Oracle's Undo Management documentation for undo tablespace administration concepts.

QUESTION NO: 18

Which three are true about opatchauto? (Choose three.)

A. It performs a shutdown and then a restart of all processes in both Oracle Grid Infrastructure and Oracle Database home during the patching process.

B. It must be invoked by a user with root user privileges.

C. Patches are applied via opatchauto.

D. Users must always input patch plans to opatchauto.

E. It requires the Oracle Grid Infrastructure and Oracle Database instances to be shut down before being invoked.

F. It applies patches in nonrolling mode by default.

G. It is used to apply interim patches to Oracle Grid Infrastructure and Oracle Database home combinations.

ANSWER: A B G

Explanation:

OPatchAuto is Oracle's automation framework for patching Oracle Grid Infrastructure and associated Oracle Database homes, particularly in clustered deployments. It orchestrates the required patching workflow, including stopping and subsequently restarting the applicable Oracle Grid Infrastructure and database processes as it moves through the patch operation. In a rolling patch scenario, this orchestration occurs node by node to preserve service availability where the patch and configuration support rolling application.

The **opatchauto** command must be started by the root operating-system user. Root authority is needed because patching Grid Infrastructure involves privileged actions such as controlling Clusterware-managed resources, modifying files owned by the Grid Infrastructure installation, and coordinating actions across cluster nodes. OPatchAuto can then invoke the necessary lower-level patching and configuration steps under the appropriate Oracle software owners.

OPatchAuto is also intended to apply interim patches to combinations of Oracle Grid Infrastructure and Oracle Database homes. This permits coordinated patching of the infrastructure and database software in a supported, automated sequence. Oracle documents OPatchAuto as the recommended automation tool for these Grid Infrastructure and RAC database home patching workflows. See the Oracle RAC patching documentation and the OPatchAuto command reference.

QUESTION NO: 19

Examine this configuration:

1. CDB1 is an Oracle Database 12c Release 2 container database (CDB).

2. PDB1 and PDB2 are two pluggable databases (PDBs) in CDB1.

After successfully performing all the preupgrade tasks, you execute these commands from the Oracle Database 18c environment:

```
$ export ORACLE_SID=cdb1

$ sqlplus / as sysdba

SQL> STARTUP UPGRADE;

SQL> SHOW pdbs

CON_ID     CON_NAME      OPEN MODE      RESTRICTED
-------    ------------  ------------   -----------
      2    PDB$SEED      MIGRATE        YES
      3    PDB1          MOUNTED        YES
      4    PDB2          MOUNTED        YES

SQL> ALTER PLUGGABLE DATABASE pdb1 OPEN UPGRADE;
Pluggable database altered.

SQL> exit

$ $ORACLE_HOME/perl/bin/perl catctl.pl -C 'PDB2'
```

What is the outcome?

A. It fails because PDB$SEED is in MIGRATE state.

B. Only CDB$ROOT, PDB$SEED, and PDB2 are upgraded.

C. Only CDB$ROOT, PDB$SEED, and PDB1 are upgraded.

D. CDB$ROOT, PDB$SEED, PDB1, and PDB2 are upgraded.

E. Only CDB$ROOT and PDB$SEED are upgraded.

F. It fails because PDB2 is not in UPGRADE state.

ANSWER: E

Explanation:

Only CDB$ROOT and PDB$SEED are upgraded. The command sequence uses the Parallel Upgrade Utility in a mode that upgrades the CDB root without processing user-created PDBs. CDB$ROOT contains the common Oracle-supplied metadata and dictionary objects that must be upgraded for the CDB. PDB$SEED is also upgraded as part of this operation, ensuring that PDBs created later from the seed have the new Oracle Database 18c metadata level.

PDB1 and PDB2 are user PDBs and are not upgraded by this root-only execution. They remain at their prior component version until a subsequent PDB upgrade is run for them. Oracle supports upgrading a CDB root separately from its PDBs, which can be useful when PDB upgrades must be scheduled independently. Before each user PDB is upgraded, it must be opened in the required upgrade mode and processed with the appropriate upgrade utility invocation. Oracle documents the CDB and PDB upgrade workflow, including separate upgrade processing, in the Oracle Database Upgrade Guide and the Parallel Upgrade Utility documentation.

QUESTION NO: 20

Which two are facets of performance planning that should always be considered or implemented for an Oracle Database environment? (Choose two.)

A. defining primary keys for all tables to speed up all queries

B. using check constraints to speed up updates

C. defining foreign keys for all tables to speed up joins

D. the physical data model

E. the configuration of storage arrays

ANSWER: D E

Explanation:

The physical data model is a core performance-planning consideration because logical entities must ultimately be implemented as database objects with appropriate table organization, partitioning strategy where applicable, indexes, and data placement. These design choices determine data access paths, I/O patterns, storage consumption, and the optimizer statistics needed for efficient execution plans. Oracle recommends designing the physical database structure with the expected workload, data volumes, growth, and access patterns in mind rather than treating physical design as an afterthought.

The configuration of storage arrays must also be considered because database performance is strongly affected by the latency, throughput, and I/O operations available to database files. Storage design influences how efficiently Oracle can read and write data files, redo logs, control files, temporary data, and recovery-related files. Planning should account for expected read/write workloads, redundancy, capacity growth, and appropriate separation or performance tiers when supported by the storage architecture. Oracle documentation discusses the importance of database storage configuration and its relationship to I/O performance in the Oracle Database Administrator's Guide and physical database design considerations in the Oracle Database Performance Tuning Guide.

QUESTION NO: 21

How do you configure a CDB for local undo mode?

A. Open the CDB instance in upgrade mode. In CDB$ROOT, execute ALTER DATABASE LOCAL UNDO ON, and then restart the CDB instance.

B. Open the CDB instance in restricted mode. In CDB$ROOT, execute ALTER DATABASE LOCAL UNDO ON, and create an UNDO tablespace in each PDB, then restart the CDB instance.

C. Open the CDB instance in restricted mode. In CDB$ROOT, drop the UNDO tablespace. Execute ALTER DATABASE LOCAL UNDO ON in each PDB, and then restart the CDB instance.

D. Open the CDB instance in read-only mode. In CDB$ROOT, execute ALTER DATABASE LOCAL UNDO ON, and then change the CDB to read/write mode.

E. Open the CDB instance in upgrade mode. In each PDB, execute ALTER DATABASE LOCAL UNDO ON, create an UNDO tablespace, and then restart the CDB instance.

ANSWER: A

Explanation:

Open the CDB instance in upgrade mode. In CDB$ROOT, execute ALTER DATABASE LOCAL UNDO ON, and then restart the CDB instance. This is the supported conversion procedure for changing an existing CDB from shared undo mode to local undo mode. The `ALTER DATABASE LOCAL UNDO ON` statement is issued only while connected to the root container, because undo mode is a CDB-wide setting rather than a setting configured independently in each PDB. Upgrade mode is required for this operation because Oracle must make the necessary dictionary and undo-management changes safely while normal user activity is prevented.

After the command succeeds, restarting the CDB opens it using local undo mode. In this mode, every PDB uses its own undo tablespace, allowing PDBs to manage undo independently. Oracle creates the required undo tablespaces for the PDBs as part of enabling local undo; manually creating an undo tablespace in every PDB is not a prerequisite for issuing the conversion command. Local undo is the default undo mode for newly created CDBs in supported later releases and is required for several multitenant capabilities, including PDB relocation and some cloning scenarios. Oracle documents both the required upgrade-mode startup and execution of the statement in the root container in its multitenant administration guidance: Oracle Database: Managing Undo in a CDB and ALTER DATABASE SQL Reference.

QUESTION NO: 22

Examine this configuration:

1. The ORCL database data files are in Automatic Storage Management (Oracle ASM) disk group +DATA.

2. ORCL uses disk group +FRA for the Fast Recovery Area.

3. LISTENER is the listener for ORCL.

4. The database, listener, ASM instance, and ASM disk groups are managed by Oracle Restart.

5. All components are currently shut down.

You execute this command:

```
$ srvctl start database -d ORCL
```

What is the outcome?

A. The ORCL database, the Oracle ASM instances, the +DATA and +FRA disk groups, and the LISTENER are started.

B. Only the ORCL database instance is started.

C. Only the ORCL database and the ASM instances are started.

D. Only the ORCL database instance, the Oracle ASM instance, and the +DATA and +FRA disk groups are started.

E. Only the ORCL database instance and the +DATA and +FRA disk groups are started.

ANSWER: D

Explanation:

Only the ORCL database instance, the Oracle ASM instance, and the +DATA and +FRA disk groups are started. Starting a database through Oracle Restart starts the database resource and the resources on which it depends. Because ORCL stores its data files in +DATA and uses +FRA as its fast recovery area, both disk groups must be mounted and available before the database can open. Oracle Restart therefore starts the required Oracle ASM instance and brings the required ASM disk groups online as dependency resources.

The listener is a separately managed Oracle Restart resource rather than a database dependency that must be started by a database-start operation. Its configured existence and management by Oracle Restart do not cause it to be started as part of starting ORCL. Oracle documents that SRVCTL manages Oracle Restart resources and their dependencies, including databases, ASM, disk groups, and listeners. The database start operation handles the database's required ASM and disk-group resources automatically. See the Oracle Restart documentation and the SRVCTL reference.

QUESTION NO: 23

A database is configured in ARCHIVELOG mode.

Full RMAN backups are taken daily and no backup to trace has been taken of the control file.

A media failure has occurred.

In which two scenarios is complete recovery possible? (Choose two.)

A. when any archived log from, before, or after the most recent backup is corrupt.

B. after losing all copies of the control file

C. after losing an archived log from after the most recent backup

D. after losing an archived log from before the most recent backup

E. after losing the SYSTEM tablespace

ANSWER: B D E

Explanation:

The stated choices contain three scenarios in which complete recovery is possible, so the "Choose two" instruction is inconsistent with Oracle recovery behavior. After losing all copies of the control file, RMAN can restore a control file from the RMAN backup created with a database backup (or from a control-file autobackup if configured), then recover the database by applying the required redo. A text trace backup is useful for manually recreating a control file, but it is not required when a usable RMAN control-file backup exists.

An archived log that was generated and completed before the most recent full backup is not required to recover data files restored from that later backup. Recovery begins from the redo required by the restored backup's checkpoint and continues through the failure point, assuming the necessary subsequent archived and online redo logs are available.

The SYSTEM tablespace can also be restored from the full RMAN backup and recovered with redo. Oracle requires the SYSTEM tablespace for database operation, but its loss is a recoverable media failure when an appropriate backup and the required redo are available. Oracle documents control-file restoration and database recovery in the RMAN Recovery Guide, and documents tablespace media recovery in the RMAN Recovery Guide: Basic Recovery.

QUESTION NO: 24

You execute this command successfully: SQL > ALTER DATABASE BACKUP CONTROLFILE TO TRACE; What is the outcome?

A. It creates a binary backup of the control file.

B. It updates the alert log with the location and contents of the control file

C. It creates a file that contains SQL statements to re-create the control file.

D. It creates a file that contains RMAN scripts to create a duplicate database.

E. It creates a copy of the control file in the fast recovery area (FRA).

ANSWER: C

Explanation:

`ALTER DATABASE BACKUP CONTROLFILE TO TRACE` writes a text trace file containing SQL statements that can be used to re-create the database control file. Oracle places this trace output in the diagnostic destination, normally under the database's trace directory. The generated file includes a `CREATE CONTROLFILE` statement and information needed to rebuild the control file, such as the database name, log file groups and members, datafiles, and recovery-related clauses. This is especially useful when a control file must be re-created following loss or corruption and a suitable binary control-file backup is unavailable. The command is issued while the database is open and is distinct from creating a binary control-file backup. Oracle documents the syntax and purpose of this command in the SQL Language Reference: ALTER DATABASE. Oracle's backup and recovery documentation also describes creating and using a control-file trace backup: Managing Control Files.

QUESTION NO: 25

Choose three. Which three are true about actions that can or cannot be performed by users with the SYSBACKUP privilege?

A. They can view data from any data dictionary view or dynamic performance view.

B. They can drop any tablespace.

C. They can view data from any user-defined tables.

D. They can create any table.

E. They cannot create restore points.

F. They cannot drop tablespaces.

ANSWER: A B D

Explanation:

SYSBACKUP is an administrative privilege intended for Oracle Recovery Manager backup and recovery administration, rather than for unrestricted SYS-level database administration. Oracle grants this privilege a defined collection of system privileges that support recovery-related operations and the management of database storage. That collection includes `SELECT ANY DICTIONARY`, enabling access to data dictionary and dynamic performance information required to determine database, backup, and recovery status. Therefore, "They can view data from any data dictionary view or dynamic performance view" is true.

The SYSBACKUP administrative privilege also includes the ability to drop tablespaces. Consequently, "They can drop any tablespace" is true: dropping a tablespace is an operation governed by the database-level `DROP TABLESPACE` system privilege, not by ownership of an individual schema object. In addition, SYSBACKUP includes `CREATE ANY TABLE`, so "They can create any table" is also true. These capabilities allow a backup administrator to perform storage and recovery-supporting administrative work without receiving the broad, unrestricted access associated with SYSDBA. Oracle documents the administrative privileges and their included system privileges in its Database Security Guide; RMAN administrative authentication is also covered in the Backup and Recovery User's Guide.

QUESTION NO: 26

Which three are true about Database Point-in-Time Recovery? (Choose three.)

A. The database must have FLASHBACK DATABASE ON to perform Database Point-in-Time Recovery.

B. The database must be in MOUNT state when performing Database Point-in-Time Recovery.

C. Database Point-in-Time Recovery is performed by the Managed Recovery Process (MRP)

D. The Database must be in ARCHIVE LOG mode.

E. The target point for the recovery must be specified as a stime or System Change Number (SCN).

F. The database must be open RESETLOGS after Database Point-in-Time Recovery.

ANSWER: B D F

Explanation:

Database Point-in-Time Recovery is an incomplete recovery operation: Oracle restores the necessary data files and applies redo only up to a selected point before the current database state. Consequently, **The database must be in MOUNT state when performing Database Point-in-Time Recovery.** RMAN requires the target database to be mounted rather than open while restoring and recovering the database files. The database must also be operating in ARCHIVELOG mode, as recovery to an earlier point relies on archived redo logs to apply the required redo through the selected target. Thus, **The Database must be in ARCHIVE LOG mode.** is a prerequisite for this recovery method. Finally, because recovery is deliberately stopped before the end of the redo stream, Oracle creates a new database incarnation when recovery is completed. Therefore, **The database must be open RESETLOGS after Database Point-in-Time Recovery.** Opening with RESETLOGS resets the online redo log sequence and establishes the recovered point as the start of the new incarnation. Oracle documents incomplete recovery and the required RESETLOGS open in its RMAN recovery documentation; ARCHIVELOG mode requirements are covered in the Database backup and recovery documentation.

QUESTION NO: 27

A physical standby database has standby redo logs configured.

You want Redo Apply to apply redo from the current standby redo log without waiting for that log to be archived.

Which command accomplishes this?

A. `ALTER DATABASE RECOVER MANAGED STANDBY DATABASE USING CURRENT LOGFILE DISCONNECT FROM SESSION;`

B. `ALTER DATABASE RECOVER STANDBY DATABASE UNTIL CANCEL;`

C. `ALTER DATABASE ARCHIVELOG CURRENT;`

D. `ALTER DATABASE CONVERT TO SNAPSHOT STANDBY;`

E. `ALTER DATABASE OPEN READ ONLY;`

ANSWER: A

Explanation:

`ALTER DATABASE RECOVER MANAGED STANDBY DATABASE USING CURRENT LOGFILE DISCONNECT FROM SESSION` starts real-time apply. With real-time apply, Redo Apply reads redo from the current standby redo log as redo is received from the primary database instead of waiting until the standby redo log is archived.

The standby redo logs must exist for this operation. Oracle documents real-time apply in the Oracle Data Guard Physical Standby Database administration documentation.

QUESTION NO: 28

In one of your databases, unified auditing is enabled and an SH.SALES table exists. You successfully execute these commands: SQL > CREATE AUDIT POLICY sales_pol ACTIONS select on sh.sales; SQL > AUDIT POLICY sales_pol; Which statement is true about SALES_POL?

A. It is enabled for all users having SELECT privilege on SH.SALES.

B. Only successful queries of SH.SALES are written to the audit trail.

C. It is only enabled for user SH.

D. Audit records are always stored in an external audit file.

ANSWER: A

Explanation:

It is enabled for all users having SELECT privilege on SH.SALES. The `CREATE AUDIT POLICY` statement defines a named unified audit policy and specifies that the `SELECT` action on the `SH.SALES` object is auditable. The policy definition alone does not begin auditing; the subsequent `AUDIT POLICY sales_pol` command enables that policy. Because this enablement statement does not contain a `BY` clause, Oracle enables the policy for all users, rather than restricting it to a named user or set of users. Consequently, when users execute applicable SELECT operations on the specified table, their activity falls within the scope of the enabled policy. Object privileges still control whether a user can successfully read the table; enabling an audit policy does not grant access. Unified auditing records the relevant audit events in the unified audit trail, which can be queried through views such as `UNIFIED_AUDIT_TRAIL`, subject to the database's audit configuration. Oracle documents that an audit policy can be enabled globally when no user list is supplied and that object actions can be included in a unified audit policy. See the CREATE AUDIT POLICY SQL Reference and Oracle's Unified Auditing administration documentation.

QUESTION NO: 29

Choose the best answer. Examine this configuration: 1. CDB1 is a container database. 2. APP_ROOT is an application root in CDB1. 3. APP_PDB1 is an application PDB in APP_ROOT. 4. FLASHBACK DATABASE is ON. You execute these commands: $ sqlplus sys/orac1e_4U@localhost:1521/cdb1 SQL> SELECT current_scn FROM v$database; CURRENT_SCN ----------2074756 SQL> CREATE TABLE cdb1_tab(c1 NUMBER) ; Table created. SQL> ALTER SESSION SET CONTAINER=app_root; Session altered. SQL> CREATE TABLE app_root_tab (c1 NUMBER) ; Table created. SQL> ALTER SESSION SET CONTAINER=app_pdb1; Session altered. SQL> CREATE TABLE app_pdb1_tab (c1 NUMBER) ; Table created. SQL> conn / as sysdba Connected. SQL> ALTER PLUGGABLE DATABASE app_root CLOSE; Pluggable database altered. SQL> FLASHBACK PLUGGABLE DATABASE app_root TO SCN 2074756; Flashback complete. Which table or set of tables will exist after the Flashback operation has completed?

A. CDB1_TAB, APP_ROOT_TAB, and APP_PDB1_TAB

B. CDB1_TAB and APP_PDB1_TAB

C. CDB1_TAB and APP_ROOT_TAB

D. CDB1_TAB only

E. none of the tables, because all three tables will be dropped

ANSWER: B

Explanation:

CDB1_TAB and APP_PDB1_TAB will exist. Flashback Database for a PDB rewinds the datafiles belonging to the specified PDB to their state at the target SCN. Here, the target SCN was captured before any of the three tables were created. However, the command specifically flashes back APP_ROOT, rather than the entire CDB or APP_PDB1. Consequently, the table created in the application root is removed as part of returning APP_ROOT to SCN 2074756.

CDB1_TAB was created in the CDB root. A PDB-level flashback of APP_ROOT does not rewind the CDB root, so the root container's table remains available. APP_PDB1_TAB was created in the application PDB. Flashing back an application root does not itself flash back an application PDB's local data; each PDB has its own datafiles and recovery scope. Therefore,

APP_PDB1_TAB remains present after APP_ROOT has been flashed back. Oracle requires the relevant PDBs to be closed as appropriate before the operation so that the flashback can consistently restore the targeted container. See Oracle's Flashback Database documentation and the FLASHBACK SQL reference.

QUESTION NO: 30

You are administering a multitenant container database (COB) that contains two pluggable databases (PDBs), pdb1 and pdb2. You are connected to pdb2 as a common user with DBA privileges. The statistics_level parameter is PDB modifiable. As the user sys, execute the following command on pdb2: SQL > ALTER SYSTEM SET STATISTICS_LEVEL=ALL SID= ' * ' SCOPE=SPFILE; Which statement is true about the result of this command?

A. The statistics_level parameter is set to all when any of the PDBs is reopened.

B. The statistics_level parameter is set to all only for PDB2 when it is reopened.

C. The statistics_level parameter is set to all when the root database is restarted.

D. The statement is ignored because there is no SPFILE for a PDB.

ANSWER: B

Explanation:

The statistics_level parameter is set to all only for PDB2 when it is reopened. Because the command is issued while connected to pdb2, a PDB-modifiable initialization parameter is changed in the context of that individual PDB. The setting is therefore associated with pdb2 rather than becoming a setting for every PDB in the container database. STATISTICS_LEVEL controls the level of database and performance-statistics collection; ALL enables all available statistics and timed operating-system statistics. Using SCOPE=SPFILE records the requested value as a persistent parameter setting rather than changing the currently running PDB immediately. Consequently, the existing in-memory value remains in effect until pdb2 is closed and opened again. At that next open, Oracle applies the persisted PDB-level setting and STATISTICS_LEVEL becomes ALL for pdb2. The SID='*' clause specifies that the setting applies to all database instances where this PDB may be opened in a RAC environment; it does not broaden the setting to other PDBs. Oracle documents that PDB-modifiable parameters can be set while connected to a PDB and that their values are maintained separately for that PDB. See the Oracle Multitenant Administration documentation and the ALTER SYSTEM SQL Reference.

QUESTION NO: 31

Automatic Shared Memory Management is disabled for one of your database instances.

Some SQL statements perform poorly due to excessive hard parse activity, thereby degrading performance.

What would be your next step?

A. Run the SQL Access Advisor.

B. Run the Memory Advisor for the shared pool.

C. Run the SQL Tuning Advisor.

D. Run the Memory Advisor for the Program Global Area.

E. Run the Memory Advisor for the System Global Area.

ANSWER: B

Explanation:

Run the Memory Advisor for the shared pool. Excessive hard parsing commonly indicates that SQL statements and their execution plans are not being retained and reused effectively in the library cache, which is a principal component of the shared pool. Because Automatic Shared Memory Management is disabled, Oracle will not automatically redistribute memory among SGA components; the DBA must evaluate and adjust component sizing as needed. The shared pool advisor

estimates the effect of different shared-pool sizes on library-cache performance and can help identify whether allocating more memory to the shared pool would reduce reloads, invalidations, and parse-related overhead. This makes it the appropriate next diagnostic action when the observed problem is hard-parse activity rather than work-area memory use or an individual SQL statement's access path. Oracle documents that the shared pool contains the library cache and that the library cache stores shared SQL and PL/SQL areas, enabling reuse of parsed SQL. The memory advisor framework provides sizing advice for memory components, including the shared pool, when manually managing memory. See Oracle's Managing Memory documentation and its discussion of the shared pool and library cache.

QUESTION NO: 32

You must transport the UNIVERSITY tablespace from one database to another. The UNIVERSITY tablespace is currently open read/write. The source and destination platforms have different endian formats.

Examine this list of actions:

1. Make the UNIVERSITY tablespace read-only on the source system.

2. Export the UNIVERSITY tablespace metadata using EXPDP.

3. Convert the UNIVERSITY tablespace data fies to the destination platform format using RMAN on the source system.

4. Copy the UNIVERSITY tablespace data files to the destination system.

5. Copy the Data Pump dump set to the destination system.

6. Convert the UNIVERSITY tablespace data files to the destination platform format using RMAN on the destination system.

7. Import the UNIVERSITY tablespace metadata using IMPDP.

8. Make the UNIVERSITY tablespace read/write on the destination system.

Which is the minimum number of actions required, in the correct order, to transport the UNIVERSITY tablespace?

A. 1,2,4,5,7,8

B. 1,2,4,6,7,8

C. 1,2,3,4,5,7,8

D. 1,2,3,4,5,6,7,8

E. 2,4,5,6,7

ANSWER: C

Explanation:

The required sequence is **1,2,3,4,5,7,8**. A conventional cross-platform transportable tablespace operation begins by placing the source tablespace in read-only mode. This ensures that its data files are transactionally consistent and that the transportable-tablespace metadata export describes a stable tablespace state. Data Pump then creates the dump set containing the metadata needed to plug the transported tablespace into the destination database.

Because the platforms use different endian formats, the source data files must be converted to the destination platform format. RMAN can perform this conversion on the source system with `CONVERT TABLESPACE`, producing destination-compatible copies. Those converted data files, together with the Data Pump dump set, are copied to the destination. Data Pump import attaches the transported data files and imports the required metadata. Once the import is complete, the tablespace can be changed to read/write mode in the destination database.

Only one endian-conversion operation is required. Performing the conversion on the source avoids a second conversion after transfer and therefore meets the stated minimum-action requirement. Oracle documents cross-platform transportable tablespaces and RMAN conversion procedures in the RMAN Data Transport documentation and the Oracle Data Pump documentation.

QUESTION NO: 33

Examine this configuration:

1. CDB1 is a container database.

2. APP_ROOT is an application root in CDB1.

3. APP_PDB1 is an application PDB in APP_ROOT.

4. FLASHBACK DATABASE is ON.

You execute these commands:

```
$ sqlplus sys/oracle_4U@localhost:1521/cdb1

SQL> SELECT current_scn FROM v$database;
CURRENT_SCN
-----------
2074756

SQL> CREATE TABLE cdb1_tab(c1 NUMBER);
Table created.

SQL> ALTER SESSION SET CONTAINER=app_root;
Session altered.

SQL> CREATE TABLE app_root_tab(c1 NUMBER);
Table created.

SQL> ALTER SESSION SET CONTAINER=app_pdb1;
Session altered.

SQL> CREATE TABLE app_pbd1_tab(c1 NUMBER);
Table created.

SQL> conn / as sysdba
Connected.

SQL> ALTER PLUGGABLE DATABASE app_root CLOSE;
Pluggable database altered.

SQL> FLASHBACK PLUGGABLE DATABASE app_root TO SCN 2074756;
Flashback complete.
```

Which table or set of tables will exist after the Flashback operation has completed?

A. CDB1_TAB, APP_ROOT_TAB, and APP_PDB1_TAB

B. CDB1_TAB and APP_PDB1_TAB

C. none of the tables, because all three tables will be dropped

D. CDB1_TAB only

E. CDB1_TAB and APP_ROOT_TAB

ANSWER: A

Explanation:

CDB1_TAB, APP_ROOT_TAB, and APP_PDB1_TAB will exist after the Flashback operation. Flashback Database restores database contents to the specified target time, SCN, or restore point; it does not selectively remove objects merely because they are located in the CDB root, an application root, or an application PDB. The command sequence shown establishes that the target to which the database is flashed back includes the creation of each of these tables. Therefore, the post-flashback database state retains CDB1_TAB in the CDB root, APP_ROOT_TAB in the application root, and APP_PDB1_TAB in the application PDB.

Oracle multitenant architecture separates object namespaces and data storage by container, but a database-level Flashback Database operation applies the recorded database state consistently across the CDB and its containers. Application containers do not change the fundamental point-in-time recovery principle: objects present at the selected flashback target remain present when recovery completes. After the flashback operation, normal recovery and opening steps make that restored point-in-time state available. Oracle documents Flashback Database behavior in the Backup and Recovery User's Guide and describes application roots and application PDBs in the Multitenant Administrator's Guide.

QUESTION NO: 34

Examine these actions:

1. Create a new database for a recovery catalog.

2. Create a tablespace with sufficient space in the catalog database for the recovery catalog.

3. Configure ARCHIVELOG mode for the catalog database.

4. Create a user to own the recovery catalog schema with quota on the tablespace that will contain the catalog.

5. Grant the RECOVERY_CATALOG_OWNER role to the recovery catalog schema owner.

6. Grant the SYSBACKUP privilege to the recovery catalog schema owner.

Which are the minimum actions that must be performed before executing the CREATE CATALOG command?

A. 2, 4, 5, 6

B. 1, 2, 3, 4, 5, 6

C. 1, 2, 4, 5

D. 2, 4, 5

E. 1, 3, 4, 5

ANSWER: D

Explanation:

2, 4, 5 is correct because a recovery catalog is a set of schema objects that RMAN creates in an existing Oracle database. Before issuing `CREATE CATALOG` while connected as the catalog owner, that owner needs a suitable location in which to create and maintain the catalog objects. Creating a dedicated tablespace with adequate available space provides that location. The catalog schema owner must then be created and given a quota on that tablespace; a quota is required for the user to allocate extents and create the recovery-catalog tables, indexes, and related objects there. Finally, the `RECOVERY_CATALOG_OWNER` role supplies the privileges needed to own and administer a recovery catalog and is the role Oracle specifies for the catalog owner.

The catalog does not require a newly created database: an existing database can host it, provided the required tablespace and user are prepared. Oracle's setup procedure describes creating a tablespace, creating the catalog owner with quota, and granting `RECOVERY_CATALOG_OWNER` before connecting with RMAN and executing `CREATE CATALOG`. See Oracle's Managing the Recovery Catalog documentation and the CREATE CATALOG reference.

QUESTION NO: 35

hoose two. Which two are true about Recovery Manager (RMAN) diagnostic message output?

A. The RMAN LOG command line clause causes output issued during RMAN command compilation to be written to a log file only.

B. The RMAN LOG command line clause causes output issued during RMAN command compilation to be written to a log file and to standard output.

C. RMAN error stacks should be read from the bottom up as that is the order in which errors are generated.

D. Media Management messages for SBT devices are written to an Oracle trace file.

E. MAN error stacks should be read from the top down as that is the order in which errors are generated.

ANSWER: A C

Explanation:

The RMAN `LOG` command-line clause directs RMAN output to the specified log file rather than standard output. This is useful when running RMAN noninteractively, such as from a scheduled operating-system job, because diagnostic information, command processing messages, and errors are retained in a file for later review. The log can then be correlated with database alert-log entries and trace files if further investigation is required.

RMAN error stacks should be read from the bottom up because RMAN reports errors as a stack: an underlying error is generated first, and higher-level RMAN messages are subsequently added to report the failure context. The messages nearest the bottom therefore commonly identify the initial condition that caused the operation to fail. Following this order helps an administrator identify the root diagnostic message before evaluating the dependent RMAN errors displayed above it.

Oracle documents RMAN invocation and output handling in the Oracle Database Backup and Recovery User's Guide. Guidance on interpreting RMAN errors and reviewing an error stack is available in the RMAN troubleshooting documentation.

QUESTION NO: 36

You must transport the UNIVERSITY tablespace from one database to another.

The UNIVERSITY tablespace is currently open read/write.

The source and destination platforms have different endian formats.

Examine this list of actions:

1. Make the UNIVERSITY tablespace read-only on the source system.

2. Export the UNIVERSITY tablespace metadata using EXPDP.

3. Convert the UNIVERSITY tablespace data files to the destination platform format using RMAN on the source system.

4. Copy the UNIVERSITY tablespace data files to the destination system.

5. Copy the Data Pump dump set to the destination system.

6. Convert the UNIVERSITY tablespace data files to the destination platform format using RMAN on the destination system.

7. Import the UNIVERSITY tablespace metadata using IMPDP.

8. Make the UNIVERSITY tablespace read/write on the destination system.

Which is the minimum number of actions required, in the correct order, to transport the UNIVERSITY tablespace?

A. 1, 2, 4, 5, 7, 8

B. 1, 2, 4, 6, 7, 8

C. 1, 2, 3, 4, 5, 7, 8

D. 1, 2, 3, 4, 5, 6, 7, 8

E. 2, 4, 5, 6, 7

F. 1, 2, 4, 5, 6, 7, 8

ANSWER: F

Explanation:

The required minimum sequence is **1, 2, 4, 5, 6, 7, 8**. A transportable tablespace must first be made read-only at the source so that its data files are transactionally consistent and suitable for transport. Data Pump then exports the tablespace metadata, including the information needed for the destination database to attach the transported files. Because the source and destination use different endian formats, the data files must be converted with RMAN. RMAN supports performing this conversion at the destination, which makes it necessary to copy the unconverted data files there first. The Data Pump dump set must also be copied because IMPDP on the destination requires that metadata file. After RMAN converts the copied data files to the destination platform format, IMPDP imports the transportable-tablespace metadata and associates it with those converted files. Finally, the tablespace can be changed to read/write in the destination database. This uses one conversion operation, performed on the destination, while retaining all required metadata-transfer steps. Oracle documents cross-platform transportable tablespace conversion in the Backup and Recovery User's Guide and the Data Pump transportable tablespace workflow in the Data Pump Export documentation.

QUESTION NO: 37

Unified auditing is enabled in your database. The HR_ADMIN and OE_ADMIN roles exist and are granted system privileges. You execute the command: SQL>CREATE AUDIT POLICY table_aud PRIVILEGES CREATE ANY TABLE, DROP ANY TABLE ROLES hr_admin,oe_admin; Which statement is true?

A. It succeeds and needs to be enabled to capture all SQL statements that require either the specified privileges or any privilege granted to the HR_ADMIN or OE_ADMIN role.

B. It fails because the command does notspecify when the unified audit policy should be enforced.

C. It succeeds and starts capturing only successful SQL statements for all users who have either the specified privileges or roles granted to them.

D. It fails because system privileges cannot be granted with roles in the same audit policy.

ANSWER: A

Explanation:

It succeeds and needs to be enabled to capture all SQL statements that require either the specified privileges or any privilege granted to the HR_ADMIN or OE_ADMIN role. A `CREATE AUDIT POLICY` statement defines a named unified audit policy; it does not itself turn auditing on. Therefore, the policy must subsequently be enabled with an `AUDIT POLICY table_aud` statement, optionally limited to specified users or roles, before audit records are generated. The `PRIVILEGES` clause in this policy identifies the `CREATE ANY TABLE` and `DROP ANY TABLE` system privileges as audit conditions. The `ROLES` clause additionally directs unified auditing to audit the use of privileges granted through the listed roles. This allows one policy to cover direct use of the two named system privileges as well as privilege use obtained through HR_ADMIN or OE_ADMIN. Oracle documents that unified audit policies can contain privilege and role audit conditions and that a created policy must be enabled before it is effective. See the Oracle CREATE AUDIT POLICY reference and Oracle's unified auditing administration documentation.

QUESTION NO: 38

HR_ROOT is an application container with the HR_APP application installed. No application PDBs and no application seed have yet been created in HR_ROOT. An application PDB, PDB1, must be created so that the HR_APP application's common objects are accessible to it. Which two methods can be used?

A. Create an application seed, synchronize it with HR_ROOT, and then create the PDB1 application PDB.

B. Create the PDB1 application PDB and synchronize it with HR_ROOT.

C. Create the PDB1 application PDB and install HR_APP in it.

D. Create the PDB1 application PDB and synchronize it with PDB$SEED.

E. Create an application seed, and install HR_APP in it.

ANSWER: A B

Explanation:

Creating an application seed, synchronizing it with HR_ROOT, and then creating the PDB1 application PDB is valid because an application seed is a specialized application PDB used as the creation template for future application PDBs. Synchronizing the seed applies the application changes recorded in the application root, including the HR_APP common-object definitions. A PDB created from that synchronized seed consequently starts with the application at the seed's application version and can access the application's common objects.

Creating the PDB1 application PDB and synchronizing it with HR_ROOT is also valid. An application PDB can be created directly in an application container even when no application seed exists. Initially, it does not necessarily contain the application changes made in the application root. Running application synchronization on PDB1 applies the relevant application installation and upgrade changes from HR_ROOT to PDB1. This establishes the metadata links and/or other application-common object state required for PDB1 to use HR_APP's common objects. Oracle documents that application PDBs synchronize with their application root to apply application changes, while application seeds provide a pre-synchronized template for efficiently provisioning new application PDBs. See Oracle Database 19c: Overview of Application Containers and Administering Application Containers with SQL*Plus.