

## Lead2Pass: Latest Free Oracle 1Z0-060 Dumps (131-140) Download!

### QUESTION 131

In which two scenarios do you use SQL\*Loader to load data?

- A. Transform the data while it is being loaded into the database.
- B. Use transparent parallel processing without having to split the external data first.
- C. Load data into multiple tables during the same load statement.
- D. Generate unique sequential key values in specified columns.

Answer: AD

Explanation:

You can use SQL\*Loader to do the following:

/ (A) Manipulate the data before loading it, using SQL functions. / (D) Generate unique sequential key values in specified columns.  
etc:

/ Load data into multiple tables during the same load session.

/ Load data across a network. This means that you can run the SQL\*Loader client on a different system from the one that is running the SQL\*Loader server.

/ Load data from multiple datafiles during the same load session.

/ Specify the character set of the data.

/ Selectively load data (you can load records based on the records' values).

/ Use the operating system's file system to access the datafiles.

/ Load data from disk, tape, or named pipe.

/ Generate sophisticated error reports, which greatly aid troubleshooting.

/ Load arbitrarily complex object-relational data.

/ Use secondary datafiles for loading LOBs and collections.

/ Use either conventional or direct path loading. While conventional path loading is very flexible, direct path loading provides superior loading performance.

Note:

\*SQL\*Loader loads data from external files into tables of an Oracle database. It has a powerful data parsing engine that puts little limitation on the format of the data in the datafile.

### QUESTION 132

You are connected to a pluggable database (PDB) as a common user with DBA privileges.

The STATISTICS\_LEVEL parameter is PDB\_MODIFIABLE. You execute the following:

```
SQL > ALTER SYSTEM SET STATISTICS_LEVEL = ALL SID = '*' SCOPE = SPFILE;
```

Which is true about the result of this command?

- A. The STATISTICS\_LEVEL parameter is set to all whenever this PDB is re-opened.
- B. The STATISTICS\_LEVEL parameter is set to ALL whenever any PDB is reopened.
- C. The STATISTICS\_LEVEL parameter is set to all whenever the multitenant container database (CDB) is restarted.
- D. Nothing happens; because there is no SPFILE for each PDB, the statement is ignored.

Answer: C

Explanation:

Note:

\*In a container architecture, the parameters for PDB will inherit from the root database. That means if statistics\_level=all in the root that will cascade to the PDB databases.

You can over ride this by using Alter system set, if that parameter is pdb modifiable, there is a new column in v\$system\_parameter for the same.

### QUESTION 133

Which two are prerequisites for performing a flashback transaction?

- A. Flashback Database must be enabled.
- B. Undo retention guarantee for the database must be configured.
- C. EXECUTE privilege on the DBMS\_FLASHBACK package must be granted to the user flashing back transaction.
- D. Supplemental logging must be enabled.
- E. Recycle bin must be enabled for the database.
- F. Block change tracking must be enabled for the database.

Answer: BC

Explanation:

B: Specify the RETENTION GUARANTEE clause for the undo tablespace to ensure that unexpired undo data is not discarded.

C: You must have the EXECUTE privilege on the DBMS\_FLASHBACK package.

Note:

\*Use Flashback Transaction to roll back a transaction and its dependent transactions while the database remains online. This recovery operation uses undo data to create and run the corresponding compensating transactions that return the affected data to its original state. (Flashback Transaction is part of DBMS\_FLASHBACK package.)

### QUESTION 134

A database is stored in an Automatic Storage Management (ASM) disk group, disk group, DGROUP1 with SQL:

```
SQL> CREATE DISKGROUP dgroup1 NORMAL REDUNDANCY
      FAILGROUP controller1 DISK '/devices/disk1', '/devices/disk2'
      FAILGROUP controller2 DISK '/devices/disk1', 'Lead2pass.com'
```

There is enough free space in the disk group for mirroring to be done.

What happens if the CONTROLLER1 failure group becomes unavailable due to error or for maintenance?

- A. Transactions and queries accessing database objects contained in any tablespace stored in DGROUP1 will fail.
- B. Mirroring of allocation units will be done to ASM disks in the CONTROLLER2 failure group until the CONTROLLER1 failure group is brought back online.
- C. The data in the CONTROLLER1 failure group is copied to the controller2 failure group and rebalancing is initiated.
- D. ASM does not mirror any data until the controller failure group is brought back online, and newly allocated primary allocation units (AU) are stored in the controller2 failure group, without mirroring.
- E. Transactions accessing database objects contained in any tablespace stored in DGROUP1 will fail but queries will succeed.

Answer: B

Explanation:

```
CREATE DISKGROUP NORMAL REDUNDANCY
```

\*For Oracle ASM to mirror files, specify the redundancy level as NORMAL REDUNDANCY (2-way mirroring by default for most file types) or HIGH REDUNDANCY (3-way mirroring for all files).

### QUESTION 135

On your Oracle 12c database, you Issue the following commands to create indexes

```
SQL > CREATE INDEX oe.ord_customer_ix1 ON oe.orders (customers_id, sales_rep_id) INVISIBLE;
```

```
SQL> CREATE BITMAP INDEX oe.ord_customer_ix2 ON oe.orders (customers_id, sales_rep_id);
```

Which two statements are correct?

- A. Both the indexes are created; however, only the ORD\_COSTOMER index is visible.
- B. The optimizer evaluates index access from both the Indexes before deciding on which index to use for query execution plan.
- C. Only the ORD\_CUSTOMER\_IX1 index is created.
- D. Only the ORD\_CUSTOMER\_IX2 index is created.
- E. Both the indexes are updated when a new row is inserted, updated, or deleted In the orders table.

Answer: AE

Explanation:

11G has a new feature called Invisible Indexes. An invisible index is invisible to the optimizer as default. Using this feature we can test a new index without effecting the execution plans of the existing sql statements or we can test the effect of dropping an index without dropping it.

#### QUESTION 136

Your multitenant container database has three pluggable databases (PDBs): PDB1, PDB2, and PDB3. Which two RMAN commands may be; used to back up only the PDB1 pluggable database?

- A. BACKUP PLUGGABLE DATABASE PDB1 while connected to the root container
- B. BACKUP PLUGGABLE DATABASE PDB1 while connected to the PDB1 container
- C. BACKUP DATABASE while connected to the PDB1 container
- D. BACKUP DATABASE while connected to the boot container
- E. BACKUP PLUGGABLE database PDB1 while connected to PDB2

Answer: AC

Explanation:

To perform operations on a single PDB, you can connect as target either to the root or directly to the PDB.

\* (A)If you connect to the root, you must use the PLUGGABLE DATABASE syntax in your RMAN commands. For example, to back up a PDB, you use the BACKUP PLUGGABLE DATABASE command.

\* (C)If instead you connect directly to a PDB, you can use the same commands that you would use when connecting to a non-CDB. For example, to back up a PDB, you would use the BACKUP DATABASE command.

#### QUESTION 137

Identify three benefits of Unified Auditing.

- A. Decreased use of storage to store audit trail rows in the database.
- B. It improves overall auditing performance.
- C. It guarantees zero-loss auditing.
- D. The audit trail cannot be easily modified because it is read-only.
- E. It automatically audits Recovery Manager (RMAN) events.

Answer: ABE

Explanation:

A:Starting with 12c, Oracle has unified all of the auditing types into one single unit called Unified auditing. You don't have to turn on or off all of the different auditing types individually and as a matter of fact auditing is enabled by default right out of the box. The AUD\$ and FGA\$ tables have been replaced with one single audit trail table. All of the audit data is now stored in Secure Files table thus improving the overall management aspects of audit data itself.

B:Further the audit data can also be buffered solving most of the common performance related problems seen on busy environments.

E:Unified Auditing is able to collect audit data for Fine Grained Audit, RMAN, Data Pump, Label Security, Database Vault and Real Application Security operations.

Note:

\*Benefits of the Unified Audit Trail

The benefits of a unified audit trail are many:

/(B)Overall auditing performance is greatly improved. The default mode that unified audit works is Queued Write mode. In this mode, the audit records are batched in SGA queue and is persisted in a periodic way. Because the audit records are written to SGA queue, there is a significant performance improvement.

/The unified auditing functionality is always enabled and does not depend on the initialization parameters that were used in previous releases

/(A)The audit records, including records from the SYS audit trail, for all the audited components of your Oracle Database installation are placed in one location and in one format, rather than your having to look in different places to find audit trails in varying formats. This consolidated view enables auditors to co-relate audit information from different components. For example, if an error occurred during an INSERT statement, standard auditing can indicate the error number and the SQL that was executed. Oracle Database Vault-specific information can indicate whether this error happened because of a command rule violation or realm violation. Note that there will be two audit records with a distinct AUDIT\_TYPE. With this unification in place, SYS audit records appear with AUDIT\_TYPE set to Standard Audit.

/The management and security of the audit trail is also improved by having it in single audit trail.

/You can create named audit policies that enable you to audit the supported components listed at the beginning of this section, as well as SYS administrative users. Furthermore, you can build conditions and exclusions into your policies.

\*Oracle Database 12c Unified Auditing enables selective and effective auditing inside the Oracle database using policies and conditions. The new policy based syntax simplifies management of auditing within the database and provides the ability to accelerate auditing based on conditions.

\*The new architecture unifies the existing audit trails into a single audit trail, enabling simplified management and increasing the security of audit data generated by the database.

#### QUESTION 138

You upgraded from a previous Oracle database version to Oracle Database version to Oracle Database 12c. Your database supports a mixed workload. During the day, lots of insert, update, and delete operations are performed. At night, Extract, Transform, Load (ETL) and batch reporting jobs are run. The ETL jobs perform certain database operations using two or more concurrent sessions. After the upgrade, you notice that the performance of ETL jobs has degraded. To ascertain the cause of performance degradation, you want to collect basic statistics such as the level of parallelism, total database time, and the number of I/O requests for the ETL jobs.

How do you accomplish this?

- A. Examine the Active Session History (ASH) reports for the time period of the ETL or batch reporting runs.
- B. Enable SQL tracing for the queries in the ETL and batch reporting queries and gather diagnostic data from the trace file.
- C. Enable real-time SQL monitoring for ETL jobs and gather diagnostic data from the V\$SQL\_MONITOR view.
- D. Enable real-time database operation monitoring using the DBMS\_SQL\_MONITOR.BEGIN\_OPERATION function, and then use the DBMS\_SQL\_MONITOR.REPORT\_SQL\_MONITOR function to view the required information.

Answer: D

Explanation:

\*Monitoring database operations

Real-Time Database Operations Monitoring enables you to monitor long running database tasks such as batch jobs, scheduler jobs, and Extraction, Transformation, and Loading (ETL) jobs as a composite business operation. This feature tracks the progress of SQL and PL/SQL queries associated with the business operation being monitored. As a DBA or developer, you can define business operations for monitoring by explicitly specifying the start and end of the operation or implicitly with tags that identify the operation.

### QUESTION 139

Your multitenant container (CDB) contains two pluggable databases (PDB), HR\_PDB and ACCOUNTS\_PDB, both of which use the CDB tablespace. The temp file is called temp01.tmp.

A user issues a query on a table on one of the PDBs and receives the following error:

ERROR at line 1:

ORA-01565: error in identifying file '\u01/app/oracle/oradata/CDB1/temp01.tmp'

ORA-27037: unable to obtain file status

Identify two ways to rectify the error.

- A. Add a new temp file to the temporary tablespace and drop the temp file that that produced the error.
- B. Shut down the database instance, restore the temp01.tmp file from the backup, and then restart the database.
- C. Take the temporary tablespace offline, recover the missing temp file by applying redo logs, and then bring the temporary tablespace online.
- D. Shutdown the database instance, restore and recover the temp file from the backup, and then open the database with RESETLOGS.
- E. Shut down the database instance and then restart the CDB and PDBs.

Answer: AE

Explanation:

\* Because temp files cannot be backed up and because no redo is ever generated for them, RMAN never restores or recovers temp files. RMAN does track the names of temp files, but only so that it can automatically re-create them when needed.

\*If you use RMAN in a Data Guard environment, then RMAN transparently converts primary control files to standby control files and vice versa. RMAN automatically updates file names for data files, online redo logs, standby redo logs, and temp files when you issue RESTORE and RECOVER.

### QUESTION 140

Examine the following commands for redefining a table with Virtual Private Database (VPD) policies:

```
BEGIN
  DBMS_RLS.ADD_POLICY (
    object_schema => 'hr',
    object_name   => 'employees',
    policy_name   => 'employees_policy',
    function_schema => 'hr',
    policy_function => 'auth_emp_dep_100',
    statement_types => 'select, insert, update, delete'
  );
END;

BEGIN
  DBMS_REDEFINITION.START_REDEF_TABLE (
    uname          => 'hr',
    orig_table     => 'employees',
    int_table      => 'int_employees',
    col_mapping    => NULL,
    options_flag   => DBMS_REDEFINITION.CONS_USE_PK,
    orderby_cols  => NULL,
    part_name      => NULL,
    copy_vpd_opt  => DBMS_REDEFINITION.CONS_VPD_AUTO
  );
END;
```

Which two statements are true about redefining the table?

- A. All the triggers for the table are disabled without changing any of the column names or column types in the table.
- B. The primary key constraint on the EMPLOYEES table is disabled during redefinition.
- C. VPD policies are copied from the original table to the new table during online redefinition.
- D. You must copy the VPD policies manually from the original table to the new table during online redefinition.

Answer: BC

Explanation:

#### C (not D):CONS\_VPD\_AUTO

Used to indicate to copy VPD policies automatically

##### \* DBMS\_RLS.ADD\_POLICY

/The DBMS\_RLS package contains the fine-grained access control administrative interface, which is used to implement Virtual Private Database (VPD).DBMS\_RLS is available with the Enterprise Edition only.

Note:

\*CONS\_USE\_PK and CONS\_USE\_ROWID are constants used as input to the "options\_flag" parameter in both the START\_REDEF\_TABLE Procedure and CAN\_REDEF\_TABLE Procedure. CONS\_USE\_ROWID is used to indicate that the redefinition should be done using rowids while CONS\_USE\_PK implies that the redefinition should be done using primary keys or pseudo- primary keys (which are unique keys with all component columns having NOT NULL constraints).

##### \* DBMS\_REDEFINITION.START\_REDEF\_TABLE

To achieve online redefinition, incrementally maintainable local materialized views are used. These logs keep track of the changes to the master tables and are used by the materialized views during refresh synchronization.

##### \*START\_REDEF\_TABLE Procedure

Prior to calling this procedure, you must manually create an empty interim table (in the same schema as the table to be redefined) with the desired attributes of the post-redefinition table, and then call this procedure to initiate the redefinition.

If you want to pass the Oracle 12c 1Z0-060 exam successfully, recommend to read latest Oracle 12c 1Z0-060 Dumps full version.

<http://www.lead2pass.com/1z0-060.html>