

2014 Latest Citrix 1Z0-053 Exam Demo Free Download!

QUESTION 1 As part of archiving the historical data, you want to transfer data from one database to another database, which is on another server. All tablespaces in the source database are read/write and online. The source and target databases use the same compatibility level and character sets. View the Exhibit and examine the features in the source and target database.

Source:	
SQL> SELECT tp.e	
2 FROM v\$tran	
3 v\$databas	
4 WHERE tp.pl	
ENDIAN_FORMAT	P
Little	M
Target:	
SQL> SELECT tp.e	
2 FROM v\$tran	
3 v\$databas	
4 WHERE tp.pl	
ENDIAN_FORMAT	P
Little	I

Which of the following steps are required to transport a tablespace from the database to the target database: 1.Make the tablespace read-only at the source database. 2.Export metadata from the source database. 3.Convert data filed by using Recovery Manager (RMAN). 4.Transfer the dump file and data filed to the target machine. 5.Import metadata at the target database. 6.Make the tablespace read/write at the target database. A. 2, 4, and 5 B. All the steps C. 2, 3, 4 and 5 D. 1, 2, 4, 5 and 6 Answer: D

Explanation: http://docs.oracle.com/cd/E11882_01/server.112/e25494/tspaces.htm#ADMIN11400

read_only--expdp--4-import-read_wrtie If both platforms have the same endianness, no conversion is necessary. Otherwise you must do a conversion of the tablespace set either at the source or destination database.

a. If no endianness conversion of the tablespace set is needed, or if you already converted the tablespace set:

Run the following query to determine the location of DATA_PUMP_DIR:

```
SELECT * FROM DBA_DIRECTORIES WHERE DIRECTORY_NAME = 'DATA_PUMP_DIR';
```

OWNER	DIRECTORY_NAME	DIRECTORY_PATH
SYS	DATA_PUMP_DIR	/u01/app/oracle/oradata/SID/dump/

b. Transport the data files to the location of the existing data files of the destination database.

On the UNIX and Linux platforms, this location is typically /u01/app/oracle/oradata/*SID* or +DISKGROUP/*SID*/datafile/.

Note:

If you converted the data files, obtain the new names and locations of the data files from the CONVERT TABLESPACE command output, as shown in Step 3 of [Task 3: Generate a Transportable Tablespace Set](#).

QUESTION 2 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 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.

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

- 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.
- Restores from backup or copies from active database the latest control file that satisfies the UNTIL clause requirements.
- Mounts the restored or copied backup or copies from the active database.
- Uses the RMAN repository to select the backups for restoring the data files to the auxiliary instance. This step applies to backup-based duplication.
- Restores and copies the duplicate data files and recovers them with incremental backups and archived redo log files to a noncurrent point in time.
- Shuts down and restarts the database instance in NOMOUNT mode.
- Creates a new control file, which then creates and stores the new DBID in the data files.
- Opens the duplicate database with the RESETLOGS option and creates the online redo log for the new database.

QUESTION 3 A user receives the following error while performing a large volume of inserts into a table:

ERROR at line 1:
ORA-01536: space quota exceeded

The issue is resolved by increasing the space quota on the USERS tablespace for the user. But the user may perform such transaction in the future. You want to ensure that the command waits rather than produce an error when such an event occurs the next time.

What can you do to achieve this before running the command in the future? A. Set RESUMABLE_TIMEOUT for the instance. B.

Set the RESOURCE_LIMIT Parameter to TRUE. C. Enable the database instance to use asynchronous commit. D. Set the LOG_CHECKPOINT_TIMEOUT parameter to a nonzero value for the database instance. Answer: A Explanation:

http://docs.oracle.com/cd/E11882_01/server.112/e25494/schema.htm#ADMIN11584 QUESTION 4 A database has three online redo log groups with one member each. A redo log member with the status ACTICE is damages while the database is running. What is the first step you should take to solve this problem? A. Attempt to Issue a checkpoint. B. Restart the database using the RESETLOGS option. C. Drop the redo log number and create it in a different location. D. Perform and incomplete recovery up to the most recent available redo log. Answer: A

http://docs.oracle.com/cd/E11882_01/backup.112/e10642/osadvscs.htm#BRADV90049

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 whether the group is 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 needed for recovery scenarios.

Table 30-4 Recovering After the Loss of an Online Redo Log Group

If the Group Is...	Then...	And You Should...
Inactive	It is not needed for crash recovery.	Clear the checkpoint of the inactive group.
Active	It is needed for crash recovery.	Attempt to issue a checkpoint and clear the checkpoint. If the database must either use Flashback Database or 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, the Database or restore a backup and perform incomplete recovery up to the most recent available redo log.

QUESTION 5 How is block-change tracking enabled? A. With alter database enable block change tracking B. With alter system enable block change tracking C. With an init.ora parameter change D. With an spfile parameter change Answer: A Explanation:

You can enable block change tracking when the database is either open or mounted. To enable block change tracking: Start SQL*Plus and connect to a target database with administrator privileges. Ensure that the DB_CREATE_FILE_DEST initialization parameter is set. SHOW PARAMETER DB_CREATE_FILE_DEST If the parameter is not set, and if the database is open, then you can set the parameter with the following form of the ALTER SYSTEM statement: ALTER SYSTEM SET

DB_CREATE_FILE_DEST = '/disk1/bct' SCOPE=BOTH SID='*'; Enable block change tracking. Execute the following ALTER DATABASE statement: ALTER DATABASE ENABLE BLOCK CHANGE TRACKING; You can also create the change tracking file in a location that you choose yourself by using the following form of SQL statement: ALTER DATABASE ENABLE BLOCK CHANGE TRACKING USING FILE '/mydir/rman_change_track.f' REUSE; The REUSE option tells Oracle Database to overwrite any existing block change tracking file with the specified name.

http://docs.oracle.com/cd/E16655_01/backup.121/e17630/rcmbckba.htm#BRADV8125 QUESTION 6 What type of backup is stored in a proprietary RMAN format? A. Backup set B. Image copy C. Backup section D. Backup group Answer: A

Explanation: http://docs.oracle.com/cd/E11882_01/backup.112/e10642/glossary.htm#BRADV90108 Backup set A backup of one or more data files, control files, server parameter files, and archived redo log files. Each backup set consists of one or more binary files. Each binary file is called a backup piece. Backup pieces are written in a proprietary format that can only be created or restored by RMAN. Backup sets are produced by the RMAN BACKUP command. A backup set usually consists of only one backup piece. RMAN divides the contents of a backup set among multiple backup pieces only if you limit the backup piece size using the MAXPIECESIZE option of the ALLOCATE CHANNEL or CONFIGURE CHANNEL command. QUESTION 7 Which dynamic view displays the status of block-change tracking? A. V\$BLOCK_CHANGE B. V\$BLOCK_CHANGE_TRACKING C. V\$BLOCKCHANGE D. V\$BLOCK_TRACKING Answer: B Explanation:

V\$BLOCK_CHANGE_TRACKING displays the status of block change tracking for the database.
Status of block change tracking in the database:

- DISABLED - Block change tracking is disabled
- TRANSITION - Block change tracking is in the process of transitioning between the enabled and disabled states. The TRANSITION state is usually never observed, because it only exists while enabling or disabling block change tracking. This state might be observed if the database crashed while enabling or disabling block change tracking, in which case it will be cleaned up automatically the next time that the database is opened.
- ENABLED - Block change tracking is enabled

http://docs.oracle.com/cd/E16655_01/server.121/e17615/refrn30030.htm#REFRN30030

Checking Whether Change Tracking Is Enabled

You can query the `V$BLOCK_CHANGE_TRACKING` view to determine whether change tracking is enabled, and if it is, the file name of the block change tracking file.

To determine whether change tracking is enabled:

Enter the following query in SQL*Plus (sample output included):

```
COL STATUS FORMAT A8
```

```
COL FILENAME FORMAT A60
```

```
SELECT STATUS, FILENAME
```

```
FROM V$BLOCK_CHANGE_TRACKING;
```

```
STATUS FILENAME
```

```
ENABLED /disk1/bct/RDBMS/change_tracking/o1_mf_2f71np5j_chg
```

Braindump2go.com

http://docs.oracle.com/cd/E16655_01/backup.121/e17630/rcmbckba.htm#BRADV8125 QUESTION 8 In a database with the database character set of US7ASCII and a national character set of UTF-8, which datatypes would be capable of storing Unicode data by default? A. VARCHAR2 B. CHAR C. NVARCHAR2 D. CLOB E. LONG Answer: C Explanation: NVARCHAR2 Data Type The NVARCHAR2 data type is a Unicode-only data type. When you create a table with an NVARCHAR2 column, you supply the maximum number of characters it can hold. Oracle subsequently stores each value in the column exactly as you specify it, provided the value does not exceed the maximum length of the column. The maximum length of the column is determined by the national character set definition. Width specifications of character data type NVARCHAR2 refer to the number of characters. The maximum column size allowed is 4000 bytes QUESTION 9 The database is configured in ARCHIVELOG mode. The database needs to be up 24 X 7. You want to perform user managed backup for the data files of the HR_DATA tablespace. To accomplish the task, you issued the following command: SQL> ALTER TABLESPACE hr_data BEGIN BACKUP; Which two statements are true in this scenario? (Choose two.) A. No transaction on the tablespace is allowed but you can perform queries. B. The tablespace will automatically come out of backup mode when the file copy is complete. C. The checkpoint change number is frozen in headers of the data files until the file is removed from backup mode. D. The database writes the before image of an entire block to the redo stream before modifying block for the first time. Answer: CD Explanation:

http://docs.oracle.com/cd/E11882_01/backup.112/e10642/osbackup.htm#BRADV90002 Making User-Managed Backups of Online Read/Write Tablespaces You must put a read/write tablespace in backup mode to make user-managed data file backups when the tablespace is online and the database is open. The ALTER TABLESPACE ... BEGIN BACKUP statement places a tablespace in backup mode. In backup mode, the database copies whole changed data blocks into the redo stream. After you take the tablespace out of backup mode with the ALTER TABLESPACE ... END BACKUP or ALTER DATABASE END BACKUP statement, the database advances the data file checkpoint SCN to the current database checkpoint SCN. When restoring a data file backed up in this way, the database asks for the appropriate set of redo log files to apply if recovery is needed. The redo logs contain all changes required to recover the data files and make them consistent. QUESTION 10 You executed the following commands: SQL> ALTER SESSION SET OPTIMIZER_USE_PENDING_STATISTICS = false; SQL> EXECUTE DBMS_STATS.SET_TABLE_PREFS('SH', 'CUSTOMERS', 'PUBLISH', 'false'); SQL> EXECUTE DBMS_STATS.GATHER_TABLE_STATS('SH', 'CUSTOMERS'); Which statement is correct regarding the above statistics collection on the SH.CUSTOMERS table in the above session? A. The statistics are stored in the pending statistics table in the data dictionary. B. The statistics are treated as the current statistics by the optimizer for all sessions. C. The statistics are treated as the current statistics by the optimizer for the current sessions only. D. The statistics are temporary and used by the optimizer for all sessions until this session terminates. Answer: A Explanation:

http://docs.oracle.com/cd/E11882_01/server.112/e41573/stats.htm#PFGRF94725 By default, the optimizer uses the published statistics stored in the data dictionary views. If you want the optimizer to use the newly collected pending statistics, then set the initialization parameter OPTIMIZER_USE_PENDING_STATISTICS to TRUE (the default value is FALSE), and run a workload against the table or schema: ALTER SESSION SET OPTIMIZER_USE_PENDING_STATISTICS = TRUE; QUESTION 11 In Oracle 11g, which recommendations does the SQL Access Advisor generate? (Choose all that apply.) A. partitioning recommendations B. statistics collection recommendations C. index creation recommendations D. materialized view recommendations E. materialized view log recommendations Answer: ACDE Explanation:

http://docs.oracle.com/cd/E11882_01/server.112/e41573/advisor.htm#PFGRF008 8.1 Overview of SQL Access Advisor Materialized views, partitions, and indexes are essential when tuning a database to achieve optimum performance for complex, data-intensive queries. SQL Access Advisor helps you achieve your performance goals by recommending the proper set of materialized views, materialized view logs, partitions, and indexes for a given workload. Understanding and using these structures is essential when optimizing SQL as they can result in significant performance improvements in data retrieval. The advantages, however, do not come without a cost. Creation and maintenance of these objects can be time consuming, and space requirements can

be significant. In particular, partitioning of an unpartitioned base table is a complex operation that must be planned carefully.

QUESTION 12 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 http://docs.oracle.com/cd/E11882_01/server.112/e41573/sql_overview.htm#PFGRF16001

Building SQL Test Cases

For many SQL-related problems, obtaining a reproducible test case makes it easier to resolve the problem. The Oracle Database contains the SQL Test Case Builder, which automates the somewhat difficult and time-consuming task of reproducing as much information as possible about a problem and the environment in which it occurred. The SQL Test Case Builder captures information regarding a SQL problem, stores it in a file, and then generates a test case so that you can reproduce and test the problem on a separate database. After the test case is reproduced, you can enable support personnel to reproduce and troubleshoot the problem.

The information gathered by SQL Test Case Builder includes the query being executed, table and index definitions, PL/SQL functions, procedures, and packages, optimizer statistics, and initialization parameters.

QUESTION 13 You enabled Flashback Data Archive on the INVENTORY table. Which DDL operation is supported on the table after enabling Flashback Data Archive? A. Drop the table. B. Partition the table C. Truncate the table. D. Add a column to the table. E. Rename a column in the table. Answer: D Explanation:

http://docs.oracle.com/cd/E11882_01/appdev.112/e41502/adfns_flashback.htm#ADFNS640

DDL Statements on Tables Enabled for Flashback Data Archive

Flashback Data Archive supports only these DDL statements:

- ALTER TABLE statement that does any of the following:
 - Adds, drops, renames, or modifies a column
 - Adds, drops, or renames a constraint
 - Drops or truncates a partition or subpartition operation
- TRUNCATE TABLE statement
- RENAME statement that renames a table

Flashback Data Archive does not support DDL statements that involve split, merge, or coalesce partitions of subpartitions to LOB columns.

For example, the following DDL statements cause error ORA-55610 when used on a table enabled for Flashback Data Archive:

- ALTER TABLE statement that includes an UPGRADE TABLE clause, with or without an INCLUDING DATA clause
- ALTER TABLE statement that moves or exchanges a partition or subpartition operation
- DROP TABLE statement

If you must use unsupported DDL statements on a table enabled for Flashback Data Archive, use the DBMS_FLASHBACK package to disassociate the base table from its Flashback Data Archive. To reassociate the Flashback Data Archive with the base table, use the DBMS_FLASHBACK_ARCHIVE.REASSOCIATE_FBA procedure.

QUESTION 14 You have a range-partitioned table in your database. Each partition in the table contains the sales data for a quarter. The partition related to the current quarter is modified frequently and other partitions undergo fewer data manipulations. The preferences for the table are set to their default values. You collect statistics for the table using the following command in regular intervals: SQL> EXECUTE DBMS_STATS.GATHER_TABLE_STATS('SH','SALES',GRANULARITY=>'GLOBAL'); You need to collect statistics more quickly. What can you do to achieve this? A. Set DYNAMIC_SAMPLING to level 4. B. Set the STATISTICS_LEVEL parameter to BASIC. C. Set the INCREMENTAL value to TRUE for the partition table. D. Increase the value of STALE_PERCENT for the partition table. Answer: C Explanation:

http://docs.oracle.com/cd/E11882_01/appdev.112/e40758/d_stats.htm#ARPLS68595 Passing your Oracle 1Z0-053 Exam by using the latest Oracle 1Z0-053 Exam Demo Full Version: <http://www.braindump2go.com/1z0-053.html>