



IBM

A2090-544

Accessment: DB2 9.7 Advanced DBA for LUW - Assessment

QUESTION: 95

A database named DB1 has a database partition group named MULTIPG that has been defined across partitions 3, 4 and 5. The DBA issues the following series of commands/SQL statements: CONNECT TO db1;
SET CURRENTDBPARTITIONNUM 4;
ALTER DATABASE PARTITION GROUP multigp ADD DBPARTITIONNUM (6); Which partition will DB2 use as a model for table spaces on partition 6?

- A. catalog partition
- B. partition 3
- C. partition 4
- D. partition 5

Answer: B

QUESTION: 96

How can you change an existing non-partitioned index for a partitioned table to a partitioned index?

- A. Use the ALTER INDEX command.
- B. Use the DROP INDEX and CREATE INDEX commands.
- C. Use the REORG INDEX command.
- D. Use the ALTER INDEX and REORG INDEX commands.

Answer: B

QUESTION: 97

A DBA would like to examine repartitioning options for a partitioned database named PRODDb. Which tool can be used to provide recommendations on repartitioning?

- A. Configuration Advisor
- B. Partition Assistant
- C. DesignAdvisor
- D. Workload Manager

Answer: C

QUESTION: 98

Which option of the REORG TABLE statement will release blocks that contain no

rows from a multi-dimensional clustering (MDC) table?

- A. RECLAIM EXTENTS ONLY
- B. RELEASE FREE EXTENTS ONLY
- C. RECLAIM FREE BLOCKS
- D. RELEASE FREE BLOCKS ONLY

Answer: A

QUESTION: 99

The DBA needs to create a table with key columns YEARDAY, YEAR, and DAY. This table needs to be partitioned by column YEARDAY with three months per data partition. Additionally, data needs to be organized so that all rows within any three month date range are clustered together based on 12 months of data. Which CREATE TABLE statement will accomplish this objective?

- A. CREATE TABLE tab1 (yearday INT, year CHAR(2), day INT)
PARTITION BY RANGE (year) (STARTING 201001 ENDING 201012 EVERY 3)
ORGANIZE BY DIMENSIONS (year, day))
- B. CREATE TABLE tab1 (yearday INT, year CHAR(2), day INT)
PARTITION BY RANGE (yearday) (STARTING 201001 ENDING 201012 EVERY 3)
ORGANIZE BY DIMENSIONS (year, day))
- C. CREATE TABLE tab1 (yearday INT, year CHAR(2), day INT)
PARTITION BY RANGE (yearday) (STARTING 201001 ENDING 201012 EVERY 3)
ORGANIZE BY DIMENSIONS (year))
- D. CREATE TABLE tab1 (yearday INT, year CHAR(2), day INT)
PARTITION BY RANGE (yearday) (STARTING 201001 ENDING 201012 EVERY 3)
ORGANIZE BY DIMENSIONS (day))

Answer: C

QUESTION: 100

Given the following statement:

```
CREATE TABLE t1 (id INT, order_date DATE, name VARCHAR (40))  
IN tblsp1, tblsp2, tblsp3, tblsp4  
PARTITION BY RANGE (order_date)  
(STARTING FROM ('2010-01-01') ENDING ('2010-03-31') EVERY 1 MONTH)
```

Which statement will successfully add a new partition to table T1 for April 2010?

- A. Export data from the main partitioned table; create a table with the same definition as the partitioned table; attach to the main partitioned table; import data into the main partitioned table; run the SET INTEGRITY statement.
- B. Create a table with the same definition as the partitioned table; attach to the main partitioned table; data will be automatically redistributed across all table partitions.
- C. Create a table with the same definition as the partitioned table; load the new table with data; attach the new table to the main partitioned table; run the SET INTEGRITY statement.
- D. Export data from the main partitioned table; drop the main partitioned table; create the main partitioned table with new partitioning range; import data into the main partitioned table.

Answer: C

QUESTION: 101

Click the Exhibit button.

```

put on assessment exam ?use previous question for
exam)

CATALOG TCP/IP NODE node001 REMOTE
hostnm01.ca.xyzcorp.com SERVER 446;
CATALOG DATABASE dsample AT NODE node001
AUTHENTICATION DCS;
CATALOG DCS DATABASE dsample AS bigdb001;
TERMINATE;
CONNECT TO dsample USER tsouser USING tsopwd;

```

In a DB2 Connect scenario between a Windows client and a z/OS server, the commands in the exhibit produced an SQL30061N error while trying to connect to the BIGDB001 database in the DB2T subsystem on the z/OS server: SQL30061N The database alias or database name was not found at the remote node. Explanation: The database name is not an existing database at the remote database node. The statement cannot be processed. The z/OS DBA has confirmed that the LOCATION NAME is DB2T. Which change would enable the command in the exhibit to execute correctly?

- A. CATALOG DATABASE bigdb001 AS dsample AT NODE db2t AUTHENTICATION DCS
- B. CATALOG DCS DATABASE dsample AS db2t
- C. CATALOG DCS DATABASE db2t AS dsample
- D. CATALOG DATABASE db2t AS dsample AT NODE node001 AUTHENTICATION DCS

Answer: B

QUESTION: 102

Which series of items are required to establish a TCP/IP connection to a DB2 for z/OS database from a DB2 for Linux, UNIX, or Windows server?

- A. hostname, logical unit (LU) name, target databasename, local database name, node name
- B. hostname, logical unit (LU) name, subsystem ID, local database name, node name
- C. hostname, port number/protocol, subsystem ID, local database name, node name
- D. hostname, port number/protocol, target database name, local database name, node name

Answer: D

QUESTION: 103

A Windows user with a local database is not able to connect to a zOS database. The System Database Directory and the Node Directory have the correct entries. Other users can connect to the database from their workstations without problems. What should be done resolve the problem?

- A. Uncatalog and re-catalog the directory entries for the zOS database and try again.
- B. Catalog the zOS database in the Local Database Directory.
- C. Catalog the zOS database in the Database Connection Services Directory.
- D. Catalog the zOS database in the Distributed Data Facility Directory.

Answer: C

QUESTION: 104

Which statement is true about clustering indexes in a partitioned database?

- A. The partitioning key column(s) should be added to the clustering index with the INCLUDE option.
- B. The clustering index column(s) should be suffixed by the partitioning key column(s).
- C. The clustering index column(s) should be prefixed by the partitioning key column(s).
- D. The partitioning key column(s) should be added immediately after the column in the clustering index with the highest cardinality.

Answer: C

QUESTION: 105

If LOGSECOND=-1 then which of the following statements is true?

- A. You still use LOGPRIMARY and LOGFILSIZ configuration parameters to specify how many log files the database manager should keep in the active log path.
- B. You must configure the OVERFLOWLOGPATH configuration parameter if your active log path is not on a raw device.
- C. The database will reserve active log space for transactions that may need to roll back and write log records.
- D. An infinite number of overflow log files will be used with archival logging.

Answer: A

Download Full Version From <https://www.certkillers.net>



DON'T KNOW
OR NO PREFERENCE

Pass your exam at First Attempt....Guaranteed!