

**100%** Money Back  
**Guarantee**

**Vendor:**IBM

**Exam Code:**C2090-614

**Exam Name:**DB2 10.1 Advanced DBA for Linux UNIX  
and Windows

**Version:**Demo

### QUESTION 1

A database named TEST was created and two buffer pools named BP\_GENERAL and BP\_AUDITING were created as follows:

```
CREATE BUFFERPOOL bp_general SIZE 100000 PAGESIZE 8K;  
CREATE BUFFERPOOL bp_auditing SIZE 100000 PAGESIZE 8K;
```

Then, two table spaces were created as follows:

```
CREATE TABLESPACE ts1 BUFFERPOOL bp_general;  
CREATE TABLESPACE ts2 BUFFERPOOL bp_auditing;
```

The majority of the database tables were created in table space TS1, while one large table was created in table space TS2. A set of triggers that would populate the table in table space TS2 each time changes were made to the database were then defined and the database was put into production. At first, ad-hoc queries performing random I/O against the table in table space TS2 ran fine, but performance has gotten worse as the table has gotten larger. What should be done to improve query performance?

- A. Reduce the size of the BP\_AUDITING buffer pool.
- B. Increase the size of the BP\_AUDITING buffer pool.
- C. Change the page size of the BP\_AUDITING buffer pool to 16 or 32 K.
- D. Replace the BP\_AUDITING buffer pool with a block-based buffer pool.

Correct Answer: A

---

### QUESTION 2

In which situation would you NOT use a redirected restore?

- A. To redefine the paths of a defined storage group.
- B. To move table space containers onto different physical devices.
- C. To restore a backup image from one operating system to another.
- D. To restore a backup image to a target machine that is different from the source machine.

Correct Answer: C

---

### QUESTION 3

A storage group named HOT\_SG was created as follows:

```
CREATE STOGROUP hot_sg
ON '/mnt/ssd_path1', '/mnt/ssd_path2'
OVERHEAD 0.75
DEVICE READ RATE 500
DATA TAG 1
SET AS DEFAULT
```

What does the DATA TAG 1 clause do?

- A. It indicates that the multiple storage paths referenced in the HOT\_SG storage group definition exist in one physical location.
- B. It tells the DB2 Optimizer that database objects stored in the HOT\_SG storage group are to be given preference over other objects when data access plans are generated.
- C. It indicates that the value 1 is to be assigned to data stored in the HOT\_SG storage group, which Workload Manager (WLM) can then use to determine the processing priority to assign to database activities that interact with this data.
- D. It assigns a unique identifier to the instance of the HOT\_SG storage group being defined; this identifier can then be used when moving table spaces to and from the storage group, when dropping the storage group, or when commenting on the storage group.

Correct Answer: C

---

#### QUESTION 4

What are two supported methods for providing a low latency, high speed interconnect between DB2 members and the cluster caching facility (CF) in a DB2 pureScale environment? (Choose two.)

- A. InfiniBand with RDMA (IB)
- B. SCSI RDMA Protocol (SRP)
- C. iSCSI Extensions for RDMA (iSER)
- D. RDMA over Converged Ethernet (RoCE)
- E. Internet Wide Area RDMA Protocol (iWARP)

Correct Answer: AD

---

#### QUESTION 5

Which procedure is used to Explain an SQL statement using information obtained from the package cache or the package cache event monitor?

- A. EXPLAIN\_FROM\_DATA
- B. EXPLAIN\_FROM\_ACTIVITY

- C. EXPLAIN\_FROM\_SECTION
- D. EXPLAIN\_FROM\_CATALOG

Correct Answer: C

---

#### **QUESTION 6**

Which statement about Range Clustered Tables (RCT) is correct?

- A. Storage space is not pre-allocated for an RCT.
- B. The ALTER TABLE statement can be used to change the physical structure of an RCT.
- C. When defining a range of key values for an RCT, a minimum value is optional.
- D. The minimum value of a key column for an RCT cannot be a negative value.

Correct Answer: C

---

#### **QUESTION 7**

Which statement about DB2 Merge Backup is FALSE?

- A. It can be used to merge offline backups to build full offline table space backups.
- B. It can be used to merge SNAPSHOT and TSM tape backup images to create a full database backup image.
- C. It can be used merge multiple offline database backups to build a cumulative single offline database backup.
- D. It can be used to merge backups on a machine that is different from the machine where the backup has been created in stand-alone mode.

Correct Answer: B

---

#### **QUESTION 8**

What must be done to convert a populated base table to a multidimensional clustering (MDC) table?

- A. Export the data, create an MDC table, and then import or load the data into the MDC table.
- B. Execute the ALTER TABLE ... ORGANIZE BY DIMENSIONS statement to convert the table.
- C. Create an MDC table and then move the data in the base table to the MDC table with the db2move command.
- D. Export the data, execute the ALTER TABLE ... ORGANIZE BY DIMENSIONS statement to convert the table, and then import or load the data into the converted table.

Correct Answer: A

---

**QUESTION 9**

Which statement about partitioned indexes is FALSE?

- A. Both system- and user-generated indexes can be partitioned.
- B. Each index partition contains references only to data in its corresponding data partition.
- C. A partitioned index is made up of a set of index partitions, each of which contains the index entries for a single data partition
- D. When attaching new partitions, the SET INTEGRITY statement must be run before data in the newly attached partition will be reflected in the corresponding partition index.

Correct Answer: D

---

**QUESTION 10**

What contributes to the efficient use of block-based buffer pools?

- A. The application does not use sequential prefetching.
- B. The application performs a lot of random I/O operations.
- C. The table space extent size is equal to the block size specified for the buffer pool.
- D. The table space prefetch size is equal to the block size specified for the buffer pool.

Correct Answer: C

---

**QUESTION 11**

What must be done to convert an existing table into an Insert-Time-Clustered (ITC) table?

- A. Call the ADMIN\_MOVE\_TABLE procedure.
- B. Call the ADMIN\_REVALIDATE\_DB\_OBJECTS procedure.
- C. Issue a REORG TABLE statement with the ORGANIZE BY INSERT TIME option.
- D. Issue an ALTER TABLE statement with the ORGANIZE BY INSERT TIME option.

Correct Answer: A

---

**QUESTION 12**

Which statement about value compression is correct?

- A. To use value compression, you must have a license for the DB2 Storage Optimization Feature.
- B. Value compression uses a page-level dictionary-based compression algorithm to compress data based on repetition of values found in a single column.
- C. Value compression involves removing duplicate entries for a value, only storing one copy, and keeping track of the location of any references to the stored copy.
- D. When value compression is used, the byte count of compressed NULLs, zeros, and blanks might be larger than that of the uncompressed version of the same values.

Correct Answer: C