



# C2090-614<sup>Q&As</sup>

DB2 10.1 Advanced DBA for Linux UNIX and Windows

## Pass IBM C2090-614 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/c2090-614.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by IBM Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



**QUESTION 1**

On which two operating systems can DB2 pureScale run? (Choose two.)

- A. AIX
- B. Linux
- C. HP-UX
- D. Solaris
- E. Windows

Correct Answer: AB

---

**QUESTION 2**

A DBA wants to create a table named SALES that has the following characteristics:

Data for the first half of the year will reside in table space TBSP1; the index for this data will reside in table space TBSP21.

Data for the second half of the year will reside in table space TBSP2; the index for this data will reside in table space TBSP22. Which clause must be used with the CREATE TABLE statement to define this table?

- A. DISTRIBUTE BY HASH
- B. PARTITION BY RANGE
- C. ORGANIZE BY DIMENSIONS
- D. ORGANIZE BY KEY SEQUENCE

Correct Answer: B

---

**QUESTION 3**

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 4**

Which statement about compression dictionaries is FALSE?

- A. Page-level dictionaries are stored in hidden rows within the table.
- B. Table-level dictionaries are stored in hidden rows within the table.
- C. Page-level dictionaries are automatically created or recreated by the database manager as pages are filled.
- D. Table-level dictionaries are automatically created or recreated as by the database manager as pages are filled.

Correct Answer: D

---

**QUESTION 5**

What statement regarding Asynchronous index cleanup (AIC) is correct?

- A. AIC removes all invalid index entries that refer to the attached partitions.
- B. AIC is initiated if a range-partitioned table contains one or more nonpartitioned indexes.
- C. AIC activates a newly attached partition after the index entries have been synchronized.
- D. AIC will not be initiated if a range-partitioned table has dependent materialized query tables (MQTs).

Correct Answer: B

---

**QUESTION 6**

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

---

**QUESTION 7**

Given the following output:

APPL_ID	UOW_ID	USER_CPU_TIME
*N2.CB2INST1.0B5A12222841	1	92782334234
*N2.CB2INST1.0B5A12725841	2	326

Which Explain procedure can be used to find the most expensive SQL statement in terms of CPU cost over a period of time?

- A. EXPLAIN\_FROM\_DATA
- B. EXPLAIN\_FROM\_ACTIVITY
- C. EXPLAIN\_FROM\_SECTION
- D. EXPLAIN\_FROM\_CATALOG

Correct Answer: B

**QUESTION 8**

Given the following CREATE TABLE statement:

```
CREATE TABLE students
(schoo_id INT NOT NULL,
program_id INT NOT NULL,
student_num INT NOT NULL,
student_id INT NOT NULL,
first_name CHAR(30),
last_name CHAR(30),
gpa FLOAT)
ORGANIZE BY KEY SEQUENCE
(student_id STARTING FROM 1 ENDING AT 1000000)
ALLOW OVERFLOW
```

Which statement is correct?

- A. A range-clustered table is created and only values from 1 to 1000000 can be stored in the STUDENT\_ID column.
- B. A range-partitioned table is created and only values from 1 to 1000000 can be stored in the STUDENT\_ID column.
- C. A range-clustered table is created and values outside the range of 1 to 1000000 can be stored in the STUDENT\_ID column.
- D. A range-partitioned table is created and values outside the range of 1 to 1000000 can be stored in the STUDENT\_ID



column.

Correct Answer: C

---

#### QUESTION 9

Which is NOT a valid method for prefetching data?

- A. Smart data prefetching
- B. Read-ahead prefetching
- C. Consecutive prefetching
- D. Sequential detection prefetching

Correct Answer: C

---

#### QUESTION 10

Which statement about Insert-Time-Clustered (ITC) is correct?

- A. Views cannot be created on ITC tables.
- B. ITC tables do not support row compression.
- C. Block indexes cannot be created for ITC tables.
- D. ITC tables can release their empty pages to the operating system.

Correct Answer: C

---

#### QUESTION 11

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 12

In which scenario should a DBA consider creating a multidimensional clustering (MDC) table instead of a regular base table?

- A. To improve access performance for tables in an OLTP database.
- B. To improve index performance in a data warehouse environment.
- C. To improve query performance in a data warehouse environment.
- D. To improve insert/delete performance for tables in an OLTP database.

Correct Answer: C

---

### QUESTION 13

You recently added more storage to a table space and a rebalance operation was started automatically. How can you pause the rebalance operation?

- A. Set the dynamic instance registry variable UTIL\_IMPACT\_LIM to 0.
- B. Set the dynamic instance registry variable UTIL\_IMPACT\_LIM to 100.
- C. Issue a LIST UTILITIES statement with the `\\SUSPEND \\` option specified.
- D. Issue an ALTER TABLESPACE statement with the `\\REBALANCE SUSPEND\\` option specified.

Correct Answer: D

---

### QUESTION 14

If the following commands complete successfully for a single partition database: `CREATE DATABASE mydb;`



CREATE BUFFERPOOL mybp; Which statement is true?

- A. Self tuning is enabled for the MYBP buffer pool.
- B. The initial size for the MYBP buffer pool is zero pages.
- C. The database has to be restarted before the MYBP buffer pool will become active.
- D. The page size for the MYBP buffer pool can be changed using the ALTER BUFFERPOOL statement.

Correct Answer: A

---

### QUESTION 15

Prior to migrating to DB2 10.1, a DBA executed the following command on a daily basis:

```
db2_all "
```