



# C2090-101<sup>Q&As</sup>

IBM Big Data Engineer





**Pass IBM C2090-101 Exam with 100% Guarantee**

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

<https://www.passapply.com/c2090-101.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

The number of partitions created by DynamicPartitions in Hive can be controlled by which of the following?

- A. hive.exec.max.dynamic.partitions.pernode
- B. hive.exec.max.dynamic.partitions
- C. hive.exec.max.created.files
- D. All of the above

Correct Answer: A

Reference: <https://resources.zaloni.com/blog/partitioning-in-hive>

---

### QUESTION 2

A Resilient Distributed Dataset supports which of the following?

- A. Creating a new dataset from an old one
- B. Returning a computed value to the driver program
- C. Both "Creating a new dataset from an old one" and "Returning a computed value to the driver program"
- D. Neither "Creating a new dataset from an old one" nor "Returning a computed value to the driver program"

Correct Answer: C

Reference: <https://spark.apache.org/docs/latest/rdd-programming-guide.html> (RDD operations)

---

### QUESTION 3

For which level of Apache Spark data locality is the data closest to the code that processes it?

- A. ANY
- B. NO\_PREF
- C. BACK\_LOCAL
- D. NODE\_LOCAL
- E. PROCESS\_LOCAL

Correct Answer: E

Reference: [https://www.ibm.com/support/knowledgecenter/SSZU2E\\_2.4.1/performance\\_tuning/application\\_spark\\_parameters.html](https://www.ibm.com/support/knowledgecenter/SSZU2E_2.4.1/performance_tuning/application_spark_parameters.html)

---



#### QUESTION 4

Which of the following techniques is NOT employed by Big SQL to improve performance?

- A. Query Optimization
- B. Predicate Push down
- C. Compression efficiency
- D. Load data into DB2 and return the data

Correct Answer: A

Reference: [https://www.ibm.com/support/knowledgecenter/en/SSZLC2\\_7.0.0/com.ibm.commerce.developer.soa.doc/refs/rsdperformancespaces.htm](https://www.ibm.com/support/knowledgecenter/en/SSZLC2_7.0.0/com.ibm.commerce.developer.soa.doc/refs/rsdperformancespaces.htm)

---

#### QUESTION 5

Which of the following statements is TRUE regarding the use of Data Click to load data into BigInsights?

- A. Data Click uses Map Reduce to create a Hive table
- B. Only relational databases are supported as sources for Data Click
- C. Only whole tables can be moved from a relational database source, not individual columns within the table
- D. If the metadata that you specify in InfoSphere Data Click activities changes, you need to reimport the metadata by using InfoSphere Metadata Asset Manager

Correct Answer: D

Reference: [https://www.ibm.com/support/knowledgecenter/SSZJPZ\\_11.3.0/com.ibm.swg.im.iis.dataclick.doc/topics/createblueprint.html](https://www.ibm.com/support/knowledgecenter/SSZJPZ_11.3.0/com.ibm.swg.im.iis.dataclick.doc/topics/createblueprint.html)

[Latest C2090-101 Dumps](#)

[C2090-101 Exam Questions](#)

[C2090-101 Braindumps](#)