



C2090-424^{Q&As}

InfoSphere DataStage v11.3

Pass IBM C2090-424 Exam with 100% Guarantee

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

<https://www.passapply.com/c2090-424.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

Which of the following statements about checkpoint/restart is TRUE?

- A. The parallel framework does not support checkpoints inside jobs
- B. You must set the environment variable \$APT_DISABLE_COMBINATION to TRUE
- C. To perform checkpoint inside a parallel job, select the checkpoint option in each stage
- D. If the checkpoint/restart is turned on for a sequence job, all components of this sequence job must perform checkpoint recording

Correct Answer: A

QUESTION 2

A job is designed with 3 Join stages in sequence.

Join_1 keys are "Account", "Customer" and "PO_Date"

Join_2 keys are "Account", "Customer" and "Txn"

Join_3 keys are "Account", "Product_CD" and "Txn"

Which of the following will lead to the best run time performance?

- A. Left links are Hash partitioned on "Account" for all joins Right links are Hash partitioned on "Account" for all joins
- B. Both left and right links are Hash partitioned as follows: Join_1 by "Account", "Customer" and "PO_Date" Join_2 by "Account", "Customer" and "Txn" Join_3 by "Account", "Product_CD" and "Txn"
- C. Both left and right links are Hash partitioned as follows: Join_1 by "Account" and "Customer" Join_2 by "Account" and "Customer" Join_3 by "Account"
- D. Left links are Hash partitioned on "Account" for the first two joins and on "Product_CD" for the third. Right links are Hash partitioned on "Account" for all joins

Correct Answer: A

QUESTION 3

Your job aggregates data before writing the results out to a DB2 table. You have since learned that this job will be required to aggregate very large amounts of data. To ensure high performance you create a new, non-default configuration file for the job that defines a node pool named "highperfnodes" that includes only nodes with very high processing power.

What additional design changes do you need to make to the job to ensure that it aggregates the data using the "highperfnodes" pool? (Choose two.)

- A. Select the "highperfnodes" pool in the Execution tab of the Job Properties window



- B. Add the \$APT_CONFIG_FILE environment variable as an additional job parameter
- C. Add a Sort stage before the Aggregator stage to sort the data before the data is aggregated
- D. Change properties in the Aggregator stage to force it to use the nodes in the appropriate pool
- E. Change properties in the Aggregator stage to force it to use the new non-default configuration file

Correct Answer: AD

QUESTION 4

Which of the following will determine the running row number within a group?

- A. Insert an Aggregator stage, hashing on the key column, prior to the Transformer stage
- B. Use a Surrogate Key Generator stage to generate the running row number prior to the Transformer stage
- C. Use a Remove Duplicates stage prior to the Transformer stage based on the key column, with the Duplicate set to retain = last option
- D. Use a Sort stage prior to the Transformer stage based on the key break column. Then, use the LastRowinGroup() function to determine the last row in a group

Correct Answer: D

QUESTION 5

Which is NOT a valid native operator for OSH?

- A. Sort
- B. Copy
- C. Import
- D. Generator

Correct Answer: A

[Latest C2090-424 Dumps](#)

[C2090-424 Exam Questions](#)

[C2090-424 Brindumps](#)