



Vendor: IBM


Exam Code: C2090-421

Exam Name: InfoSphere DataStage v8.5

Version: DEMO

QUESTION 1

A customer is complaining of poor performance from the job shown in the exhibit, which aggregates and joins on ZIP code. Currently all partitioning is set to "Auto" and sort/partition insertion is allowed.



What would give the best performance improvement for the job?

- A. Hash partition and sort on ZIP code on the input links to the Aggregator and Join stages.
- B. Hash partition and sort on ZIP code prior to the Copy stage, and use same partitioning on the Aggregator and Join stages.
- C. Use same partitioning on all input links in the job; create the input data set by hash-partitioning its input on ZIP code.
- D. Hash partition and sort on ZIP code prior to the Copy stage and use entire partitioning on the output of the Aggregator stage.

Answer: B

QUESTION 2

Which two tasks can the Slowly Changing Dimensions (SCD) stage perform? (Choose two.)

- A. Look up whether a record with a matching business key value exists in a dimension table. If it does, add new values for selected fields to values lists for those fields.
- B. Look up whether a record with a matching business key value exists in a fact table. If it does not, retrieve a new surrogate key value and insert a new row into the fact table.
- C. Look up whether a record with a matching business key value exists in a dimension table. If it does not, retrieve a new surrogate key value and insert a new row into the dimension table.
- D. Look up whether a record with a matching business key value exists in a dimension table. If it does, mark the record as not-current, and generate a new record with new values for selected fields.

Answer: CD

QUESTION 3

In which two situations is it appropriate to use a Sparse Lookup? (Choose two.)

- A. When accessing DB2 data using the DB2 API stage.
- B. When the output of the Lookup stage needs to be hashed partitioned.
- C. When reference data is significantly larger than the streaming data (100:1).
- D. When invoking a stored procedure within a database per row in the streaming link.

Answer: CD

QUESTION 4

You set environment variable \$APT_ORACLE_LOAD_OPTIONS=PTIONS(DIRECT=TRUE, PARALLELTRUE)?for loading index organized tables. You set environment variable \$APT_ORACLE_LOAD_OPTIONS=?PTIONS(DIRECT=TRUE, PARALLELTRUE)?for loading index organized tables. Which statement is accurate regarding the resulting effect of this environment variable setting?

- A. Oracle load will fail when executed.
- B. Oracle load will run in parallel and bypass DataStage Director.
- C. The Oracle load will run in parallel using simple direct path mode.
- D. The Oracle database stage will run in parallel using conventional path mode.

Answer: A

QUESTION 5

The derivation for a stage variable is: Upcase(input_column1):": Upcase(input_column2). Suppose that input_column1 contains a NULL value. Which behavior is expected?

- A. The job aborts.
- B. NULL is written to the target stage variable.
- C. The input row is either dropped or rejected depending on whether the Transformer has a reject link.
- D. The target stage variable is populated with spaces or zeros depending on the stage variable data type.

Answer: B

QUESTION 6

You have a job that reads in Sequential File followed by a Transformer stage. When you run this job, which partitioning method will be used by default?

- A. Hash
- B. Same
- C. Random
- D. Round Robin

Answer: D

QUESTION 7

Which job design technique can be used to give unique names to sequential output files that are used in multi-instance jobs?

- A. Use parameters to identify file names.
- B. Generate unique file names by using a macro.
- C. Use DSJobInvocationID to generate a unique filename.
- D. Use a Transformer stage variable to generate the name.

Answer: A

QUESTION 8

Which two statements are correct when referring to an Aggregator Stage? (Choose two.)

- A. Use Sort method for a limited number of distinct key values.
- B. Use Hash method for a limited number of distinct key values.
- C. Use Sort method with a large number of distinct key-column values.
- D. Use Hash method with a large number of distinct key-column values.

Answer: BC

QUESTION 9

Which two statements are true about the use of named node pools? (Choose two.)

- A. Named node pools can allow separation of buffering from sorting disks.
- B. Clustered environments must have named node pools for data processing.
- C. Using appropriately named node pools forces DataStage to use named pipes between stages.
- D. Named node pools constraints will limit stages to be executed only on the nodes defined in the node pools.

Answer: AD

QUESTION 10

You are asked by management to document all jobs written to make future maintenance easier. Which two statements are true about annotations? (Choose two.)

- A. The full job description can be identified within the Description Annotation stage.
- B. The background for the Description Annotation stage can be changed for each unique stage.
- C. The Description Annotation stage contains both the short and full descriptions for the job.
- D. The Description Annotation stage can be added several times at different locations to identify business logic.

Answer: AC

QUESTION 11

You have run ten instances of the same job the previous evening. You want to examine the job logs for all instances but can only find five of them. How can you avoid this in the future for this job?

- A. Change the Auto-purge settings in Administrator.
- B. Change the Auto-purge settings for the job in Director.
- C. Set the \$APT_AUTOPURGE_LOG environment variable to False.
- D. Set the \$APT_AUTOLOG_PURGE environment variable to False.

Answer: B

QUESTION 12

In which two situations would you use the Web Services Client stage? (Choose two.)

- A. You need to create a WSDL.
- B. You want to deploy a service.
- C. You do not need both input and output links in a single web service operation.
- D. You need the Web service to act as either a data source or a data target during an operation.

Answer: CD

QUESTION 13

You are describing to your customer how to work with packages. There are two steps that must be carried out in order to move the package to a target project. Which two steps must be carried out to move the DataStage objects to the Production system? (Choose two.)

- A. Build
- B. Export
- C. Deploy
- D. Package
- E. Compile

Answer: AC

Thank You for Trying Our Product

PassLeader Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.passleader.com/all-products.html>



Microsoft



ORACLE



CITRIX



JUNIPER
NETWORKS



EMC²
where information lives®

10% Discount Coupon Code: STNAR2014