

Exam 000-418 study material

Made available by Examsexpert.com



Free 000-418 Exam Preparation Questions

Exam 000-418: InfoSphere DataStage v8.0

Question: 1

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

Answer: C

Question: 2

Which two statements are true regarding access to a MQ queue? (Choose two.)

- A. MQ stage can be used to read and write messages from a remote MQ message queue.
- B. MQ connector stage is capable of connecting to Queue manager in both Server and Client mode.
- C. MQ stage connects to a queue manager in Client mode only.
- D. MQ stage and MQ connector stages support destructive reads.

Answer: B, D

Question: 3

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

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

Answer: A, C

Question: 4

Which three methods can be used to import metadata from a Web Services Description Language (WSDL) document? (Choose three.)

- A. Web Service Function Definitions
- B. Orchestrate Schema Definitions
- C. XML Table Definitions
- D. Job Stage Column tab properties entered using "Load" feature
- E. Web Services WSDL Definitions

Answer: A, C, E

Question: 5

In a job using Sequential File stage to write a sequential file, the ContactName column is nullable, but the "Null field value" property is left unspecified. The ContactName column has a Char(50) data type. What will happen if NULL is written to the ContactName column?

- A. The row containing the NULL will be rejected.
- B. The job will abort.
- C. The empty string ("") will be written to the ContactName column.
- D. The column will be padded with the character specified by the \$APT_STRING_PADCHAR character.

Answer: A

Question: 6

A scenario requires selecting only the most recent transactions for each of 2 million unique customers from a 20 million row DB2 source table containing order history. Which parallel job design would satisfy this functional requirement?

- A. Using the DB2 API stage, select all rows. Use a parallel Sort stage with customer number and order date sort keys in ascending order, then Remove Duplicates with Last Duplicate to retain.
- B. Using the Dynamic Relational stage, use custom SQL to select all DISTINCT customer numbers from the order history table.
- C. Select all rows using the ODBC Connector stage, use a Sort Aggregator on customer number key to select the maximum order date.
- D. Using the DB2 Enterprise stage, select all rows. Perform a unique Sort using customer number and order date sort keys in ascending order.

Answer: A

Question: 7

Your job is using the DB2 Enterprise stage as a target. In this job you are inserting data into a single existing table on a remote database. You are using Run-Time Column Propagation (RCP) on all of the links. Which two job design configuration actions would you consider to effectively use RCP? (Choose two.)

- A. In the DB2 Enterprise stage use Connection properties to manage column mappings.
- B. Use a Modify stage prior to the DB2 Enterprise stage in order to drop the columns not found in the table.
- C. Use DB2 partitioning properties on the input link to the DB2 Enterprise stage to control column layout in partitions.
- D. In the DB2 Enterprise stage set the option named "Silently Drop Columns Not In Table" to "True".

Answer: B, D

Question: 8

Your job design calls for using a target ODBC connector stage. The target database is found on a remote server. The target table you are writing into contains a single column primary key. What is the one required property for the Connection properties?

- A. Username
- B. Data source
- C. Password
- D. Remote server

Answer: B

Question: 9

You have a job named MyJob in project ProdProject and directory /xmlReportDir. Which two return an XML report containing complete performance data for MyJob? (Choose two.)

- A. dsjob -report ProdProject MyJob DETAIL > /xmlReportDir/MyJob.xml
- B. \$APT_PERFORMANCE_DATA=/xmlReportDir
- C. \$APT_DUMP_SCORE=true
- D. dsjob -report ProdProject MyJob XML > /xmlReportDir/MyJob.xml

Answer: B, D

Question: 10

What are two tasks that can create DataStage projects? (Choose two.)

- A. Install the DataStage engine.
- B. Create a project in the IBM Information Server Console.
- C. Add new projects from DataStage Administrator.
- D. Export and import a DataStage project from DataStage Director.

Answer: A, C

Question: 11

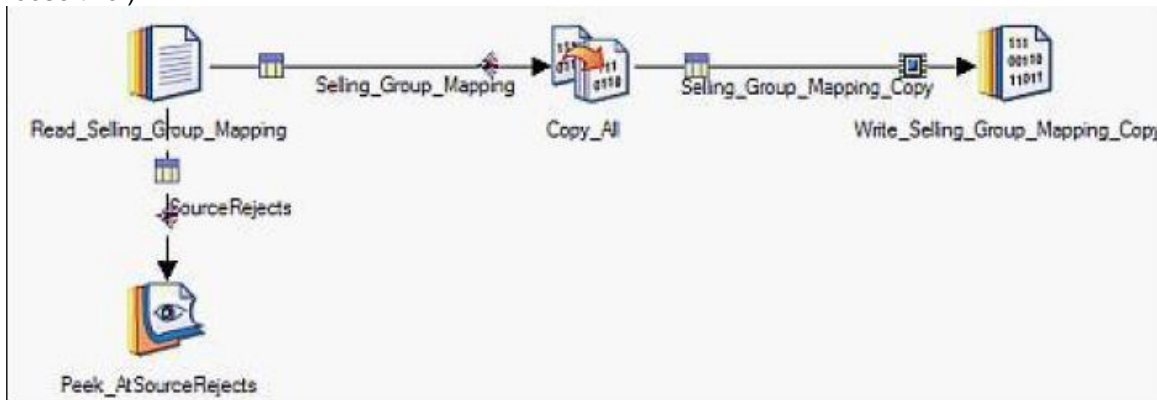
You are assigned to write a job that splits a single Data Set into three Data Sets based on character values on the input record. Normally you would use a transformer link constraint to achieve this but you discover that the condition for each output link is supplied at run-time. Which stage would allow you to parameterize the entire condition for splitting the input data set?

- A. Modify
- B. Filter
- C. Match Frequency
- D. Split Vector

Answer: B

Question: 12

Click the Exhibit button. In the exhibit, data is being read from a Sequential File stage (Read_Selling_Group_Mapping) and written to a target Data Set stage (Write_Selling_Group_Mapping_Copy). All columns are being mapped source to target in the Copy stage. After the job is executed using a 2-node configuration file, you notice that the Data Set has only one partition. What are two actions that can be done to this job to ensure Data Set has two partitions? (Choose two.)



- A. Disable Runtime Column Propagation (RCP) and set "Force" option in the Copy stage to "TRUE".
- B. Enable Runtime Column Propagation (RCP) and set "Force" option in the Copy stage to "FALSE".
- C. Enable Runtime Column Propagation (RCP) and set "Force" option in the Copy stage to "TRUE".
- D. Disable Runtime Column Propagation (RCP) and set "Force" option in the Data Set stage to "TRUE".

Answer: A, C

Question: 13

A 100 MB input data set has even distribution across 400 unique key values. When run with a 4-node configuration file, which two changes could improve sort performance for this scenario? (Choose two.)

- A. Set \$APT_TSORT_STRESS_BLOCKSIZE to 50 MB.
- B. Specify Round Robin partitioning on the input link to the Sort stage.
- C. Specify "Restrict Memory Usage" to 60 MB on the Sort stage properties.
- D. Alter the Sort stage properties to perform a Stable sort.

Answer: A, C

Question: 14

You are assigned to write a job that reads in a Sequential File, applies some business logic, and writes the results to a database. However, the layout of the input file will change from run to run. All of the columns will remain the same, but the order of the columns will change. Which feature of the Sequential File stage would allow you to accomplish this task?

- A. Filter
- B. First Line Is Column Names
- C. Reject Mode
- D. Schema File

Answer: D

Question: 15

Which requirement must be met to read from a database in parallel using the ODBC Enterprise stage?

- A. Specify the partitioncol property.
- B. ODBC Enterprise stage always reads in parallel.
- C. Configure environment variable \$APT_PARTITION_COUNT.
- D. Configure environment variable \$APT_MAX_TRANSPORT_BLOCK_SIZE.

Answer: A

Question: 16

Which two partitioning methods are keyless? (Choose two.)

- A. Round Robin
- B. Entire
- C. Hash
- D. Modulus

Answer: A, B

Question: 17

Which two statements are true for named node pools? (Choose two.)

- A. Using appropriately named node pools can allow separation of buffering from sorting disks.
- B. Using appropriately named node pools forces DataStage to use named pipes between stages.
- C. Using appropriately named node pools constraints will limit stages to be executed only on the nodes defined in the node pools.
- D. Using appropriately named node pools is only available in clustered environments.

Answer: A, C

Question: 18

Your customer, understanding the value of end-to-end parallelism, requires all jobs to have parallel input and output. You are also required to use an input/output format that is archivable, and compatible between product versions. What achieves these goals without ignoring the parallelism of the configuration file?

- A. Schema File
- B. Data Sets
- C. Multiple File options
- D. File Sets

Answer: D

For complete [Exam 000-418 Training kits and Self-Paced Study Material](http://www.Examsexpert.com/000-418.html)

Visit:

<http://www.Examsexpert.com/000-418.html>



For Latest 000-418 Exam Questions and study guides- visit- <http://www.Examsexpert.com/000-418.html>