

VCE4Plus



Everything you need to prepare, learn & pass your certification exam easily.

Pass Your Next Certification Exam Fast!

365 days free updates. First attempt guaranteed success.

Choose the version that fits your needs	PDF Version	Desktop Test Engine	Online Test Engine
Latest and Up-to-Date exam dumps with real exam questions answers.	✓	✓	✓
Get 12-Months free updates without any extra charges.	✓	✓	✓
Experience same exam environment before appearing in the certification exam.	✗	✓	✓
100% exam passing guarantee in the first attempt.	✓	✓	✓
20% discount on more than one license and 30% discount on 5+ license purchases.	✗	✓	✓
100% secure purchase on SSL.	✓	✓	✓
Completely private purchase without sharing your personal info with anyone.	✓	✓	✓

<http://www.vce4plus.com>

Accurate exam material ensure you pass for sure by your first attempt - VCE4Plus

Exam : **000-303**

Title : IBM Certified Solution
Developer - IBM InfoSphere
DataStage

Vendor : IBM

Version : DEMO

NO.1 You are using the Complex Flat File stage as a source in your job. What are two types of data specifically supported by the Complex Flat File stage for your job? (Choose two.)

- A. XML files.
- B. Mainframe data sets with VSAM files.
- C. Mainframe data files with DB2 tables.
- D. Data from files that contain Hadoop record types.
- E. Data from flat files that contain multiple record types.

Answer: B,E

NO.2 Identify two different types of custom stages you can create to extend the Parallel job syntax? (Choose two.)

- A. Input stage
- B. Basic stage
- C. Group stage
- D. Custom stage
- E. Wrapped stage

Answer: D,E

NO.3 When using Runtime Column Propagation, which two stages require a schema file? (Choose two.)

- A. Peek stage
- B. Pivot stage
- C. Column Import stage
- D. DB2 Connector stage
- E. Sequential File stage

Answer: C,E

NO.4 What is the result of running the following command: dsjob -report DSProject ProcData

- A. Generates a report about the ProcData job, including information about its stages and links.
- B. Returns a report of the last run of the ProcData job in a DataStage project named DSProject.
- C. Runs the DataStage job named ProcData and returns performance information, including the number of rows processed.
- D. Runs the DataStage job named ProcData and returns job status information, including whether the job aborted or ran without warnings.

Answer: B

NO.5 Your job is using the Sequential File stage as a source. Identify two stage property configurations that will allow you to accomplish parallel reads? (Choose two.)

- A. Read Method property set to 'File Pattern' and the environment variable \$APT_IMPORT_PATTERN_USES_FILESET is set to 'FALSE'.
- B. Read Method property set to 'File Pattern' and the 'Readers per Node' property value set to greater than 1.
- C. Read Method property is set to 'Parallel' with the Key property set to a Input column that has the file names.

- D. Read Method property set to 'Specific File(s)' with only one file specified and the 'Readers Per Node' property value set to greater than 1.
- E. Read Method property set to 'Specific File(s)' with more than one file specified and each file specified has the same format/column properties.

Answer: D,E

NO.6 Which statement is true when using the SaveInputRecord() function in a Transformer stage.

- A. You can only use the SaveInputRecord() function in Loop variable derivations.
- B. You can access the saved queue records using Vector referencing in Stage variable derivations.
- C. You must retrieve all saved queue records using the GetSavedInputRecord() function within Loop variable derivations.
- D. You must retrieve all saved queue records using the GetSavedInputRecord() function within Stage variable derivations.

Answer: C

NO.7 What two project environment variables can be considered in your parallel jobs to support your optimization strategy of partitioning and sorting? (Choose two.)

- A. \$APT_NO_PART_INSERTION
- B. \$APT_OPT_SORT_INSERTION
- C. \$APT_RESTRICT_SORT_USAGE
- D. \$APT_PARTITION_FLUSH_COUNT
- E. \$APT_TSORT_STRESS_BLOCKSIZE

Answer: A,E

NO.8 The effective use of naming conventions means that objects need to be spaced appropriately on the DataStage Designer canvas. For stages with multiple links, expanding the icon border can significantly improve readability. This approach takes extra effort at first, so a pattern of work needs to be identified and adopted to help development. Which feature of Designer can improve development speed?

- A. Palette Layout
- B. Repository Layout
- C. Snap to Grid Feature
- D. Show Performance Statistics Feature

NO.9 A DataStage job uses an Inner Join to combine data from two source parallel datasets that were written to disk in sort order based on the join key columns. Which two methods could be used to dramatically improve performance of this job? (Choose two.)

- A. Disable job monitoring.
- B. Set the environment variable \$APT_SORT_INSERTION_CHECK_ONLY.
- C. Unset the Preserve Partitioning flag on the output of each parallel dataset.
- D. Explicitly specify hash partitioning and sorting on each input to the Join stage.
- E. Add a parallel sort stage before each Join input, specifying the "Don't Sort, Previously Grouped" sort key mode for each key.

Answer: B,E

NO.10 When using a Sequential File stage as a source what are the two reject mode property options? (Choose two.)

- A. Set
- B. Fail
- C. Save
- D. Convert
- E. Continue

Answer: B,E