



IBM 000-415

Exam Name: IBM WebSphere IIS DataStage Enterprise Edition v7.5

Q & A : 141 Q&As

Pdf Demo

Quality and Value for the 000-415 Exam

[Just4Exams Practice Exams](#) for IBM IBM certifications I 000-415 are written to the highest standards of technical accuracy, using only certified subject matter experts and published authors for development.

100% Guarantee to Pass Your 000-415 Exam

If you do not pass the IBM certifications I 000-415 exam on your first attempt using our Just4Exams **000-415 testing engine and pdf study guide**, we will give you a FULL REFUND of your purchasing fee.

Downloadable, Interactive 000-415 Testing engines and PDF Version

Our Exam Preparation Material provides you everything you will need to take a [IBM certifications I certification](#) examination. Details are researched and produced by [IBM Certification](#) Experts who are constantly using industry experience to produce precise, and logical.

Free 000-415 Exams:

This is demo only, this pdf do not include the questions and answers picture

Exam : IBM 000-415

Title : IBM WebSphere IIS DataStage Enterprise Edition v7.5

1. Which task is performed by the DataStage JobMon daemon?

- A. graceful shutdown of DataStage engine
- B. writes the job's OSH script to the job log
- C. automatically sets all environment variables
- D. provides a snapshot of a job's performance

Answer: D

2. You have a parallel job that does not scale beyond two nodes. After investigation you find that the data has been partitioned on gender and you have an Aggregator stage that is accumulating totals based on gender using a sort method. Which technique should you use to

allow the job to scale?

- A. Add an additional column for partitioning to result in additional data partitions.
- B. Change the aggregation method to hash to eliminate the blocking sort operation.
- C. Change the preserve partitioning option on the stage ahead of the aggregator to clear partitioning.
- D. Change the partitioning key to one that results in more partitions, add a second Aggregator stage that re-aggregates based on gender.

Answer: D

3. What is the purpose of the uv command in a UNIX DataStage server?

- A. Start and stop the DataStage engine.
- B. Report DataStage client connections.
- C. Cleanup resources from a failed DataStage job.
- D. Provide read access to a DataStage EE configuration file.

Answer: A

4. Which two stages allow field names to be specified using job parameters? (Choose two.)

- A. Filter stage
- B. Modify stage
- C. Funnel stage
- D. Transformer stage

Answer: AB

5. Click the Exhibit button. In the exhibit, which item describes how you would make a stage run only one way parallel?

- A. constrain to node pool app1
- B. constrain to node pool app2
- C. constrain to node pool /temp
- D. constrain to node pool /orch/n4/d1

Answer: B

6. Which three UNIX kernel parameters have minimum requirements for DataStage installations? (Choose three.)

- A. NOPROC no process limit
- B. NOFILES number of open files
- C. MAXPERM disk cache threshold
- D. SHMMAX maximum shared memory segment size
- E. MAXUPROC maximum number of processes per user

Answer: BDE

7. Which two describe a DataStage EE installation in a clustered environment? (Choose two.)

- A. The C++ compiler must be installed on all cluster nodes.
- B. Transform operators must be copied to all nodes of the cluster.
- C. A remote shell must be configured to support communication between the conductor and section leader nodes.
- D. The DataStage parallel engine must be installed or accessible in the same directory on all machines in the cluster.

Answer: CD

8. Click the Exhibit button. In the exhibit, what is "bigdata"?

- A. a node
- B. a resource
- C. a node pool
- D. a resource pool

Answer: D

9. Which three defaults are set in DataStage Administrator? (Choose three.)

- A. default SMTP mail server name
- B. project level defaults for environment variables
- C. project level default for Auto-purge of job log entries
- D. project level default for Runtime Column Propagation
- E. default prompting options, such as Autosave job before compile

Answer: BCD

10. Your input rows contain customer data from a variety of locations. You want to select just those rows from a specified location based on a parameter value. You are trying to decide whether to use a Transformer or a Filter stage to accomplish this. Which statement is true?

- A. You cannot use a Filter stage because you cannot use parameters in a Filter stage Where clause.

- B. The Filter stage will yield better performance because it has less overhead than a Transformer stage.
- C. The Transformer stage will yield better performance because the Filter stage Where clause is interpreted at runtime.
- D. You cannot use the Transformer stage because you cannot use parameters in a Transformer stage constraint.

Answer: C

11. How are transaction commit operations handled by the DB2/UDB Enterprise stage?

- A. Commit operations can be defined globally by setting APT_TRANSACTION_ROWS variable.
- B. Commit operations can only be defined by the number of rows since the start of a transaction.
- C. Commit operations can only be defined by the number of seconds since the start of a transaction.
- D. Transaction commits can be controlled by defining the number of rows per transaction or by a specific time period defined by the number of seconds elapsed between commits.

Answer: D

12. Which feature in DataStage will allow you determine all the jobs that are using a Shared Container?

- A. Usage Analysis
- B. Impact Analysis
- C. Extended Job view
- D. Reporting Assistant

Answer: A

13. Which three features of datasets make them suitable for job restart points? (Choose three.)

- A. They are persistent.
- B. They are partitioned.
- C. They are indexed for fast data access.
- D. They are compressed to minimize storage space.
- E. They use datatypes that are in the parallel engine internal format.

Answer: ABE

14. Which three accurately describe the differences between a DataStage server root installation and a non-root installation? (Choose three.)

- A. A root installation enables auto-start on reboot.
- B. A non-root installation enables auto-start on reboot.
- C. A root installation will start DataStage services in impersonation mode.
- D. A root installation must specify the user "dsadm" as the DataStage administrative user.
- E. A non-root installation inherits the permissions of the user who starts the DataStage services.

Answer: ACE

15. Click the Exhibit button. In the exhibit, how many parallel partitions will each DataStage EE operator run by default?

- A. two
- B. three
- C. four
- D. five

Answer: B

16. Which two statements are true of the Merge stage? (Choose two.)

- A. The Merge stage supports only two input links.
- B. The Merge stage supports inner, left outer, and right outer joins.
- C. All the inputs to the Merge stage must be sorted by the key column(s).
- D. The columns that are used to merge the incoming data must be identically named.

Answer: CD

17. Click the Exhibit button. Which two hardware configurations are shown in the exhibit? (Choose two.)

- A. grid
- B. SMP
- C. cluster
- D. single processor

Answer: AC

18. Which three privileges must the user possess when running a parallel job? (Choose three.)

- A. read access to APT_ORCHHOME
- B. read/write permissions to APT_ORCHHOME

- C. read/write permissions to the UNIX/etc directory
- D. read/write access to disk and scratch disk resources
- E. execute permissions on local copies of programs and scripts

Answer: ADE

19. Which two tasks will create DataStage projects? (Choose two.)

- A. Install the DataStage engine.
- B. Copy a project in DataStage Administrator.
- C. Add new projects from DataStage Administrator.
- D. Export and import a DataStage project from DataStage Manager.

Answer: AC

20. Which "Reject Mode" option in the Sequential File stage will write records to a reject link?

- A. Fail
- B. Drop
- C. Output
- D. Continue

Answer: C

More [000-415 Braindumps](#) Information

Related 000-415 Exams

000-974	000-223	000-377	000-015	000-990
000-417	000-669	000-316	000-667	000-913
000-936	000-431	000-084	000-994	000-111
000-908	000-017	000-113	000-207	000-011

Other IBM Exams

000-385	000-371	000-890	000-060	000-021
000-378	000-779	000-746	000-431	000-740
000-285	000-543	000-002	000-288	000-100
000-349	000-026	000-273	000-718	000-209