

Exam : IBM 000-415

**Title : IBM WebSphere IIS
DataStage Enterprise Edition
v7.5**

Version : Demo

Important Note, Please Read Carefully

Other examsoon products

[All examsoon IT Exam Products](#)

Our products of Offline Testing Engine

Use the offline Testing engine product to practice the questions in an exam environment.

Build a foundation of knowledge which will be useful also after passing the exam.

[examsoon Testing Engine](#)

Guarantee

Examsoon provides the most competitive quality of all exams for the customers, we guarantee your success at the first attempt with only our Certification Question&Answers, if you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!.

Features

1. Comprehensive questions with complete details
2. Instant Downloadable in PDF form.
3. Verified Answers Researched by Industry Experts
4. Questions accompanied by exhibits
5. Drag and Drop questions as experienced in the Actual Exams.
6. These questions and answers are backed by our GUARANTEE.
7. Questions updated on regular basis.
8. Like actual certification exams our product is in multiple-choice questions (MCQs).

1. 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

2. 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

3. 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

4. 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

5. 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

6. 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

7. 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

8. 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

9. 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

10. 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

11. 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

12. 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

13. 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

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

```
{
  node 'n1' {
    fastrname "s1"
    pool "" "n1" "s1" "app2" "sort"
    resource disk "/orch/n1/d1" {}
    resource disk "/orch/n1/d2" {"bigdata"}
    resource scratchdisk "/temp" {"sort"}
  }
  node 'n2' {
    fastrname "s2"
    pool "" "n2" "s2" "app1"
    resource disk "/orch/n2/d1" {}
    resource disk "/orch/n2/d2" {"bigdata"}
    resource scratchdisk "/temp" {}
  }
  node 'n3' {
    fastrname "s3"
    pool "" "n3" "s3" "app1"
    resource disk "/orch/n3/d1" {}
    resource scratchdisk "/temp" {}
  }
  node 'n4' {
    fastrname "s4"
    pool "n4" "s4" "app1"
    resource disk "/orch/n4/d1" {}
    resource scratchdisk "/temp" {} }
}
```

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

Answer: B

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

```
{
  node 'n1' {
    fastrname "s1"
    pool "" "n1" "s1" "app2" "sort"
    resource disk "/orch/n1/d1" {}
    resource disk "/orch/n1/d2" {"bigdata"}
    resource scratchdisk "/temp" {"sort"}
  }
  node 'n2' {
    fastrname "s2"
    pool "" "n2" "s2" "app1"
    resource disk "/orch/n2/d1" {}
    resource disk "/orch/n2/d2" {"bigdata"}
    resource scratchdisk "/temp" {}
  }
  node 'n3' {
    fastrname "s3"
    pool "" "n3" "s3" "app1"
    resource disk "/orch/n3/d1" {}
    resource scratchdisk "/temp" {}
  }
  node 'n4' {
    fastrname "s4"
    pool "n4" "s4" "app1"
    resource disk "/orch/n4/d1" {}
    resource scratchdisk "/temp" {} }
}
```

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

Answer: D

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

```
{
  node 'n1' {
    fastrname "s1"
    pool "" "n1" "s1" "app2" "sort"
    resource disk "/orch/n1/d1" {}
    resource disk "/orch/n1/d2" {"bigdata"}
    resource scratchdisk "/temp" {"sort"}
  }
  node 'n2' {
    fastrname "s2"
    pool "" "n2" "s2" "app1"
    resource disk "/orch/n2/d1" {}
    resource disk "/orch/n2/d2" {"bigdata"}
    resource scratchdisk "/temp" {}
  }
  node 'n3' {
    fastrname "s3"
    pool "" "n3" "s3" "app1"
    resource disk "/orch/n3/d1" {}
    resource scratchdisk "/temp" {}
  }
  node 'n4' {
    fastrname "s4"
    pool "n4" "s4" "app1"
    resource disk "/orch/n4/d1" {}
    resource scratchdisk "/temp" {} }
}
```

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

Answer: AC

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

```
{
  node 'n1' {
    fastrname "s1"
    pool "" "n1" "s1" "app2" "sort"
    resource disk "/orch/n1/d1" {}
    resource disk "/orch/n1/d2" {"bigdata"}
    resource scratchdisk "/temp" {"sort"}
  }
  node 'n2' {
    fastrname "s2"
    pool "" "n2" "s2" "app1"
    resource disk "/orch/n2/d1" {}
    resource disk "/orch/n2/d2" {"bigdata"}
    resource scratchdisk "/temp" {}
  }
  node 'n3' {
    fastrname "s3"
    pool "" "n3" "s3" "app1"
    resource disk "/orch/n3/d1" {}
    resource scratchdisk "/temp" {}
  }
  node 'n4' {
    fastrname "s4"
    pool "n4" "s4" "app1"
    resource disk "/orch/n4/d1" {}
    resource scratchdisk "/temp" {} }
}
```

- 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

18. 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

19. 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

20. 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

Examsoon.com was founded in 2006. The safer,easier way to help you pass any IT Certification exams . We provide high quality IT Certification exams practice questions and answers(Q&A). Especially [Adobe](#), [Apple](#), [Citrix](#), [Comptia](#), [EMC](#), [HP](#), [Juniper](#), [LPI](#), [Nortel](#), [Oracle](#), [SUN](#), [Vmware](#) and so on. And help you pass any IT Certification exams at the first try.

we guarantee your success at the first attempt with only our Certification Question&Answers, if you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

You can reach us at any of the email addresses listed below.

E-MAIL: [examsoon\(at\)hotmail.Com](mailto:examsoon(at)hotmail.Com)