

## **000-418 IBM certifications III IBM IBM Websphere Datastage V.8.0**


**Practice Exam:** 000-418 Exams

**Exam Number/Code:** 000-418

**Exam Name:** IBM Websphere Datastage V.8.0

**Questions and Answers:** 140 Q&As

( [IBM certifications III](#) )



Exam : [000-418](#)

"IBM Websphere Datastage V.8.0", also known as 000-418 exam, is a IBM certification. With the complete collection of questions and answers, TestInside has assembled to take you through 140 Q&As to your 000-418 Exam preparation. In the 000-418 exam resources, you will cover every field and category in IBM Certification helping to ready you for your successful IBM Certification.

Quality and Value for the 000-418 Exam TestInside Practice Exams for IBM **IBM certifications III** Certification 000-418 are written to the highest standards of technical accuracy, using only certified subject matter experts and published authors for development.

### ***TestInside provide the professional Q&A.***

1. We offer free update service for three month.

After you purchase our product, we will offer free update in time for three month.

2. High quality and Value for the 000-418 Exam.

000-418 simulation test questions, including the examination question and the answer, complete by our senior IT lecturers and the IBM certifications III product experts, included the current newest 000-418 examination questions.

3. 100% Guarantee to Pass Your IBM certifications III exam and get your IBM certifications III Certification.

If you do not pass the IBM Certification 000-418 exam (IBM Websphere Datastage V.8.0) on your first attempt using our TestInside testing engine and pdf file, we will give you a FULL REFUND of your purchasing fee.

### ***use TestInside 000-418 Q&A ensure you pass the exam at your first try.***

TestInside professional provide IBM certifications III 000-418 the newest Q&A, completely covers 000-418 test original topic. With our complete IBM certifications III resources, you will minimize your IBM certifications III cost and be ready to pass your 000-418 tests on Your First Try, 100% Money Back Guarantee included!

[IBM 000-418](#) Test belongs to one of the IBM certifications III certified test, if needs to obtain the IBM certifications III certificate, you also need to participate in other related test, the details you may visit the [IBM certifications III](#) certified topic, in there, you will see all related IBM certifications III certified subject of examination.

### ***TestInside Testing Engine Features***

Comprehensive questions and answers about 000-418 exam

000-418 exam questions accompanied by exhibits

Verified Answers Researched by Industry Experts and almost 100% correct

000-418 exam questions updated on regular basis

Same type as the certification exams, 000-418 exam preparation is in multiple-choice questions (MCQs).

Tested by multiple times before publishing

Try free 000-418 exam demo before you decide to buy it in Test-Inside.com.

Note: This pdf demo do not include the question's picture.

Exam : IBM 000-418

Title : IBM Websphere Datastage V.8.0

1. You are working on a job in which a sequential file cannot be read in parallel. In an attempt to improve job performance, you first define a single large string column for the non-parallel sequential file read. Which stage may be used to parse the large string in parallel?

- A. the Column Import stage
- B. the Column Export stage
- C. the Make Vector stage
- D. the Split Vector stage

Answer: A

2. A parallel job combines rows from a source DB2 table with historical information maintained in two separate Oracle tables. Only rows in the DB2 source whose key values match either Oracle table are output to a target Teradata table. Both Oracle tables have identical column definitions and are stored in the same Oracle instance. Which two design techniques would satisfy this requirement? (Choose two.)

- A. Using a master DB2 Enterprise stage, merge by dropping unmatched masters against a single Oracle Enterprise stage with custom SQL with UNION ALL.
- B. Combine the inputs from the DB2 Enterprise stage and two Oracle Enterprise stages using the Sort options of the Funnel stage defined on the key columns.
- C. Use a separate Oracle Enterprise stage for each source table to a Funnel stage and then perform an inner join with rows from a DB2 Enterprise stage.
- D. Use a Lookup stage to combine the DB2 Enterprise input with each Oracle Enterprise reference link using range partitioning to limit each reference by historical data values.

Answer: AC

3. Which two methods can be used for adding messages to a message handler? (Choose two.)

- A. Import message handler from existing message handler dsx.
- B. Drag and drop a message from the job log onto the message handler.
- C. Type in the message rule by hand.
- D. Use the add rule to message handler interface.

Answer: CD

4. Which three lookup types may be performed in the Lookup stage? (Choose three.)

- A. Equality match
- B. Negative match
- C. Range on stream link
- D. Range on the reject link
- E. Range on the reference link

Answer: ACE

5. You have created a parallel job in which there are several stages that you want to be able to re-use in other jobs. You decided to create a parallel shared container from these stages. Identify two things that are true about this shared container. (Choose two.)

- A. It can be used in sequencer jobs.
- B. It can take advantage of Run Time Column Propagation (RCP).

- C. It can be used in Transformer stage derivations.
- D. It can have job parameters to resolve stage property values.

Answer: BD

6. What describes the column information specified in Orchestrate schemas? (Choose two.)

- A. C++ data types, such as string[max 145]
- B. column properties, such as nullability
- C. SQL data types, such as Char(20)
- D. record format information, such as record delimiter

Answer: AB

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

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

Answer: BC

8. Click the Exhibit button.

You are asked to convert a relational source, as shown in the exhibit, into three separate XML files.

How would you accomplish this task?

- A. Set "Output mode" setting "Aggregate all rows" under "Transformation Settings" of the Output link of an XML Output stage.
- B. Set "Output mode" setting "Use Trigger" on column "Customer ID" under "Transformation Settings" of the Output link of an XML Output stage.
- C. Set "Output mode" setting "Separate rows" on column "Customer ID" under "Transformation Settings" of the Output link of an XML Output stage.
- D. Set "Output mode" setting "Single row" under "Transformation Settings" of the Output link of an XML Output stage.

Answer: D

9. A DataStage job is sourcing a flat file which contains a VARCHAR field. This field needs to be mapped to a target field which is a date. Which will accomplish this?

- A. Use a Column Exporter to perform the type conversion.
- B. DataStage handles the type conversion automatically.
- C. Use the TimestampFromDate function in a Transformer.
- D. Use the Modify stage to perform the type conversion.

Answer: D

10. Which three statements are true about File Sets? (Choose three.)

- A. File sets are partitioned.
- B. File sets are unpartitioned.
- C. File sets are stored as a single file.
- D. File sets are readable by external applications.
- E. File sets are stored as header file and data files.

Answer: ADE

11. Click the Exhibit button.

In the exhibit, a Funnel stage has two Input Links. Input 1(Seq\_File) comes from a sequential file with set "Readers per Node" set to "2" , Input 2(Dataset) comes from a Data Set created in parallel (three-way parallel). Within the Funnel stage, the funnel type is set to "Sequence". The parallel configuration file contains three nodes.

How many instances of the Funnel stage run in parallel?

- A. 1
- B. 2

C. 3

D. 6

Answer: C

12. Click the Exhibit button.

In the exhibit, how many parallel partitions will each DataStage operator run without specifying a "node pool and resource constraint" or actual "node constraint"?

A. one

B. two

C. three

D. four

Answer: C

13. When tuning a parallel process, it is necessary to measure the amount of system resources that are used by each instance of a stage. Which two methods enable the collection of CPU time used by each instance of a stage? (Choose two.)

A. Set the environment variable \$APT\_PM\_PLAYER\_TIMING=true.

B. Invoke vmstat before the job run and after the job completes.

C. Select the Record job performance data check box from Job Properties.

D. Set the environment variable \$DS\_MAKE\_JOB\_REPORT=2.

Answer: AC

14. Which three actions can improve sort performance in a DataStage job? (Choose three.)

A. Specify only the key columns which are necessary.

B. Use the stable-sort option to avoid the random ordering of non-key data.

C. Minimize the number of sorts used within a job flow.

D. Adjusting the "Restrict Memory Usage" option in the Sort stage.

E. Run the job sequentially so that only one sort process is invoked.

Answer: ACD

15. Your job uses the MQ connector stage to read messages from an MQ queue. The job should retrieve the message ID into the MessageID field and parse the payload into two fields: Name is to get the first ten characters, Description is to get the remaining characters.

What will accomplish this?

A. First column is MessageID as Binary 24; second column is Name as Binary 10; select WSMG.MSPAYLOAD data element; third column is Description as VarBinary 200; select WSMG.MSPAYLOAD data element.

B. First column is MessageID; select the WSMQ.MSGID data element for the Message ID field; second column is Description as VarBinary 200; third column is Name as Binary 10.

C. First column is MessageID; select the WSMQ.MSGID data element for the Message ID field; second column is Name; select WSMG.MSPAYLOAD data element; third column is Description; select WSMG.MSPAYLOAD data element.

D. First column is MessageID; select the WSMQ.MSGID data element for the Message ID field; second column is Name as Binary 10; third column is Description as VarBinary 200.

Answer: D

16. What is the lowest CPU cost partitioning method for parallel stage to parallel stage?

A. Range

B. Modulus

C. Entire

D. Same

Answer: D

17. A job design reads from a complex flat file, performs some transformations on the data, and outputs the results to

a WISD output stage. What are two ways that parameter values can be passed to this job at run-time? (Choose two.)

- A. Pass the parameter values at the time of the service request.
- B. Change the properties of the information provider and redeploy.
- C. Include the parameter values in the data.
- D. Execute a DSSetParam with the values at job execution time.

Answer: AB

18. Your job is to setup credential mappings for DataStage developers within DataStage. Which two statements are true? (Choose two.)

- A. You must be a Information Server Suite administrator to complete this task.
- B. You can create Information Server Suite users and groups in the Web console.
- C. You can create new Information Server Suite users by using the DataStage Administrator.
- D. You can create new users in the operating system level and map these credentials within DataStage Administrator.

Answer: AB

19. Which two steps are required to change from a normal lookup to a sparse lookup in an ODBC Enterprise stage? (Choose two.)

- A. Change the lookup option in the stage properties to "Sparse".
- B. Replace columns at the beginning of a SELECT statement with a wildcard asterisk (\*).
- C. Establish a relationship between the key field column in the source stage with the database table field.
- D. Sort the data on the reference link.

Answer: AC

20. When invoking a job from a third-party scheduler, it is often desirable to invoke a job and wait for its completion in order to return the job's completion status. Which three commands would invoke a job named "BuildWarehouse" in project DevProject and wait for the job's completion? (Choose three.)

- A. dsjob -run -log DevProject BuildWarehouse
- B. dsjob -run -jobstatus DevProject BuildWarehouse
- C. dsjob -run -userstatus DevProject BuildWarehouse
- D. dsjob -run DevProject BuildWarehouse
- E. dsjob -run -wait DevProject BuildWarehouse

Answer: BCE

### [More 000-418 Information](#)

#### **Related 000-418 Exams**

[000-101](#) *Virtualization Technical Support for AIX and Linux*

[000-100](#) *ENTERPRISE TECHNICAL SUPPORT FOR AIX AND LINUX*

[000-330](#) *Power Systems Technical Support for AIX and Linux*

[000-418](#) *IBM Websphere Datastage V.8.0*

[000-962](#) *Storage Sales for HighEnd Tape Version 1*

[COG-480](#) *IBM Cognos 8 Planning Professional*

[000-136](#) *Rational Application Developer v7*

[000-960](#) *IBM Storage Sales Exam*

[000-210](#) *IBM Storage Networking Solutions, Version 4*

[000-433](#) *IBM Tivoli Storage Manager V5.5 Implementation*

[000-432](#) *IBM Tivoli Storage Manager V5.5 Operator*

[COG-300](#) *IBM Cognos TM1 Analyst*

[000-314](#) *IBM Web Sphere Message Broker V6.1. System Administration*

000-019 *IBM Tivoli Network Manager IP Edition V3.8 Implementation*

000-R18 *IBM SurePOS ACE V5 Technical Professional*

000-233 *aix installation,backup and system rec0ver*

000-071 *ib m eserver x series technical principle v5*

000-209 *SVC/Virtullization V.1*

000-M22 *IBM Rational Application Developer Tech Sales Mastery V1*

000-M15 *IBM Rational Quality Mgt Technical Sales Mastery Test v1*

### **Other IBM Exams**

<u>000-375</u>	<u>000-025</u>	<u>000-P03</u>	<u>000-061</u>	<u>000-M07</u>	<u>000-842</u>	<u>006-002</u>	<u>000-931</u>
<u>000-002</u>	<u>000-443</u>	<u>LOT-957</u>	<u>000-348</u>	<u>000-389</u>	<u>COG-145</u>	<u>LOT-710</u>	<u>000-P01</u>
<u>000-204</u>	<u>000-664</u>	<u>000-883</u>	<u>COG-300</u>				