



**Vendor:** IBM

**Exam Code:** 000-421

**Exam Name:** InfoSphere DataStage v8.5

**Version:** DEMO

1. You are responsible for deploying objects into your customers production environment. To ensure the stability of the production system the customer does not permit compilers on production machines. They have also protected the project and only development machines have the required compiler. What two options will allow jobs with a parallel transformer to execute in the customers production machines? (Choose two.)

- A. Add \$APT\_COMPILE\_OPT=-portable
- B. Set \$APT\_COPY\_TRANSFORM\_OPERATOR
- C. Export the jobs with Information Server Manager with the executables.
- D. Create a package with Information Server Manager and select the option to include executables.

**Answer:** C,D

2 You have been asked to delete a shared container from the project by your customer. Before you do this you want to make sure it will not impact other objects in the project. How will you ensure that deleting the shared container will not cause a failure when jobs are recompiled?

- A. Select the shared container, Right-click on the Where used command.
- B. Select the shared container, Right-click on the Where used (deep) command.
- C. Select the shared container, Right-click on the Dependencies (deep) command.
- D. Advanced find, set Dependencies Of field equal to the container name, and Check the Option to Include nested results for Dependency searches.

**Answer:** B

3 You are working on a project that contains a large number of jobs contained in many folders. You would like to review the jobs created by the former developer of the project. How can you find these jobs?

- A. Filter jobs in Director Client's Repository window.
- B. Sort the jobs by date in the Repository window.
- C. Use the Advanced Find feature contained in the Designer interface.
- D. While selecting the top folder in the project, choose Find Dependencies.

**Answer:** C

4 When you configure a domain for source code integration a source control workspace is created. What are two reasons for the source control workspace? (Choose two.)

- A. default directory for tool exports.
- B. directory for deployment package files
- C. local transfer area for assets being submitted to the source control system
- D. provides a place to store the archive created by IS source code control integration

**Answer:** C,D

5 You are about to begin major changes to jobs in a project. You want to conveniently identify job changes on an ad hoc basis. What two tasks will allow you to identify changes to your jobs? (Choose two.)

- A. Import the original job from a .dsx export.

- B. Select the job, then right click Compare within.
- C. Select the job, then right click Cross Project Compare.
- D. Before making a change to a job make a copy of the job in a different category folder.

**Answer:** B,D

6. You are responsible for the project's Source Code Repository. When a developer notifies you that changes are ready to deploy you must first check them in. You will use the Information Server Source Code integration features to check in the changed assets. How will you identify the DataStage assets the developer modified in order to send them to the workspace?

- A. From the Information Server Manager Select the project then right-click on Synchronize with source code control system.
- B. Use Designer client Advanced Find, specify a modification date and the developer's username in the Modified by fields.
- C. From Information Server Manager use Search with advanced options, specify a date modified and the developer's user name in the Modified by fields.
- D. From the workspace perspective select the project Right-click then Refresh from repository, specify a modification date and the developer's username in the Modified by fields.

**Answer:** C

7. Where are project level message handlers defined?

- A. DSENV
- B. Director Client
- C. Designer Client
- D. Administrator Client

**Answer:** D

8. Which three of the following options does the dsjob command have? (Choose three.)

- A. Stopping a job
- B. Setting an alias for a job
- C. Specifying an appropriate log file
- D. Listing projects, jobs, stages, links, and parameters


**Answer:** A,B,D

9. You would like to pass values into parameters that will be used in a variety of downstream activity stages within a job sequence. What are three valid ways to do this? (Choose three.)

- A. Use local parameters.
- B. Use environment variables.
- C. Place a parameter set stage on the job sequence.
- D. Check the "Propagate Parameters" checkbox in the Sequence Job properties.
- E. Use the UserVariablesActivity Stage to populate the local parameters from an outside source such as a file.

**Answer:** A,B,E

10. Click on the Exhibit button.



Provided you have enough system resources, what is the maximum number of jobs that could be running concurrently in this image?

- A. 2
- B. 3
- C. 4
- D. 5

**Answer: B**

11. You are experiencing performance issues for a given job. You are assigned the task of understanding what is happening at run time for that job. What are the first two steps you should take to understand the job performance issues? (Choose two.)

- A. Review the objectives of the job.
- B. Run job with \$APT\_TRACE\_RUN set to true.
- C. Run job with \$APT\_DUMP\_SCORE set to true.
- D. Replace Transformer stages with custom operators.

**Answer: A,C**

12. Click on the Exhibit button.

```

##I IFSC 004000 14:51:50(000) <main_program>
This step has 1 data set:
ds0: {op0[1p] (sequential generator)
      eOther(APT_HashPartitioner { key={ value=a }
    })->eCollectAny
      op1[2p] (parallel APT_CombinedOperatorController:tsort)}
It has 2 operators:
op0[1p] {(sequential generator)
  on nodes {
    lemond.torrent.com[op0,p0]
  }}
op1[2p] {(parallel APT_CombinedOperatorController:
  (tsort)
  (peek)
  ) on nodes {
    lemond.torrent.com[op1,p0]
    lemond.torrent.com[op1,p1]
  }}

```

- A. 3 processes on 2 nodes
- B. 3 processes on 3 nodes
- C. 2 processes on 2 nodes
- D. 2 processes on 3 nodes

**Answer: A**

13. Click the exhibit button.

```

dsadm      7176  7175  5 13:12 ?      00:00:00 /opt/IBM/InformationServer/Serve
r/PXEngine/bin/osh -f RT_SC9/OshScript.osh -monitorport 13400 -pf RT_SC9/jpfile
-impexp charset UTF-8 -string charset UTF-8 -input charset UTF-8 -output charset
dsadm      7177  7134  3 13:12 ?      00:00:00 phantom DSD.OshMonitor RowGen 71
dsadm      7215      1 1 13:12 ?      00:00:00 /opt/IBM/InformationServer/Serve
r/PXEngine/bin/osh -APT_PMsectionLeaderFlag EDServer 10000 0 30 node1 EDServer 1
dsadm      7216      1 0 13:12 ?      00:00:00 /opt/IBM/InformationServer/Serve
r/PXEngine/bin/osh -APT_PMsectionLeaderFlag EDServer 10000 1 30 node2 EDServer 1
dsadm      7217  7215  1 13:12 ?      00:00:00 /opt/IBM/InformationServer/Serve
r/PXEngine/bin/osh -APT_PMsectionLeaderFlag EDServer 10000 0 30 node1 EDServer 1
dsadm      7218  7215  5 13:12 ?      00:00:00 /opt/IBM/InformationServer/Serve
r/PXEngine/bin/osh -APT_PMsectionLeaderFlag EDServer 10000 0 30 node1 EDServer 1
dsadm      7219  7216  2 13:12 ?      00:00:00 /opt/IBM/InformationServer/Serve
r/PXEngine/bin/osh -APT_PMsectionLeaderFlag EDServer 10000 1 30 node2 EDServer 1
dsadm      7220  7216  5 13:12 ?      00:00:00 /opt/IBM/InformationServer/Serve
r/PXEngine/bin/osh -APT_PMsectionLeaderFlag EDServer 10000 1 30 node2 EDServer 1
dsadm      7221  7216  9 13:12 ?      00:00:00 /opt/IBM/InformationServer/Serve
r/PXEngine/bin/osh -APT_PMsectionLeaderFlag EDServer 10000 1 30 node2 EDServer 1

```

You submit a job from DataStage Director and then log onto your DataStage Linux server to issue the command "ps -ef | grep ds" and receive the following screen: Which process is a player?

- A. 7117
- B. 7215
- C. 7216
- D. 7217

**Answer: D**

14. What would help debug the run-time environment?

- A. \$APT\_VERBOSE
- B. \$OSH\_PRINT\_SCHEMAS
- C. Check the message filters for the job.
- D. View detailed runtime stats through Job monitor in Director Client.

**Answer: B**

15.To better understand the data flowing through a job, it is desirable to have structural information about the datasets captured in the job log. Which environment variable can provide this information?

- A. \$OSH\_STDOUT\_MSG
- B. \$OSH\_PRINT\_SCHEMAS
- C. \$APT\_PM\_PLAYER\_MEMORY
- D. \$APT\_NO\_PART\_INSERTION

**Answer: B**

16.DataStageoffers database connectivity through connectors, native parallel and plug-in stage types.

Which two statements are correct? (Choose two.)

- A. ODBC API is a plug-in stage.
- B. Next to the connector stage it is best to use the native parallel database stages.
- C. The connector stage offers better functionality and performance and is the best to use.
- D. For maximum parallel performance, scalability, and features it is best to use the native parallel database stages.

**Answer: B,C**

17.Which two property areas must be configured when using the ODBC connector stage as a target in your job design? (Choose two.)

- A. Define columns for the output link.
- B. Specify the remote server property.
- C. Specify properties for the input link.
- D. Define the connection properties to an ODBC data source.

**Answer: C,D**

18.You set environment variable

\$APT\_ORACLE\_LOAD\_OPTIONS=PTIONS(DIRECT=TRUE, PARALLEL=TRUE)?for loading index organized tables.You set environment variable \$APT\_ORACLE\_LOAD\_OPTIONS=?PTIONS(DIRECT=TRUE, PARALLEL=TRUE)?for loading index organized tables. Which statement is accurate regarding the resulting effect of this environment variable setting?

- A. Oracle load will fail when executed.
- B. Oracle load will run in parallel and bypassDataStage Director.
- C. The Oracle load will run in parallel using simple direct path mode.
- D. The Oracle database stage will run in parallel using conventional path mode.

**Answer: A**

19. Which two statements are accurate regarding usage of database stages? (Choose two.)

- A. Plug-in stages match columns by name.
- B. Native database stages match columns by name.
- C. DataStage provides native Oracle database stages.
- D. Database stages cannot create tables and must be done externally.

**Answer:** B,C

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

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

**Answer:** C,D