

Alibaba

ACP-BigData1

ACP Big Data Professional

Questions And Answers PDF Format:

For More Information – Visit link below:
<https://www.certsgrade.com/>

Version = Product



Latest Version: 6.0

Question: 1

In DataWorks, a task is instantiated before a scheduled task is run, that is, an instance is generated and executed to run the scheduled task. Instances displayed in the Task Operation & Management(O&M) view of the O&M Center are automatically scheduled by tasks.

We can perform the _____ operations on these instances.

(Number of correct answers: 2)

Response:

- A. Test
- B. Data population
- C. Viewing node run logs
- D. Re-running and resuming scheduling

Answer: C D

Question: 2

When we use the MaxCompute tunnel command to upload the log.txt file to the t_log table, the t_log is a partition table and the partitioning column is (p1 string, p2 string).

Which of the following commands is correct?

Response:

- A. tunnel upload log.txt t_log/(p1="b1",p2="b2")
- B. tunnel upload log.txt t_log/p1="b1",p2="b2"
- C. tunnel upload log.txt t_log/p1="b1"/p2="b2"
- D. tunnel upload log.txt t_log(p1="b1",p2="b2")

Answer: B

Question: 3

When MaxCompute saves data to the system, the _____ file format is used to reduce invalid disk read operations.

Response:

- A. Column-store
- B. Row-store
- C. Key-value pair storage

D. Document storage

Answer: A

Question: 4

When odpscmd is used to connect to a project in MaxCompute, the command _____ can be executed to view the size of the space occupied by table table_a.

Response:

- A. select size from table_a;
- B. size table_a;
- C. desc table_a;
- D. show table table_a;

Answer: C

Question: 5

Authorization in MaxCompute refers to assigning some permissions of some objects to some specified users. Which three factors of the following must be included?

(Number of correct answers: 3)

Response:

- A. Subject, which may be a user or role
- B. Object, which may be a table or resource
- C. Action, such as read, write, etc.
- D. Effect, such as accept, reject, etc.

Answer: A B C

Question: 6

The distributed system consists of hundreds or even thousands of storage machines built from inexpensive commodity parts and is accessed by a comparable number of client machines.

The quantity and quality of the components virtually guarantee that some are not functional at any given time and some will not recover from their current failures.

Based on the information above, when you try to design a distributed system, the first thing to consider would be _____.

Response:

- A. component failures

- B. files are huge by traditional standards
- C. data locality
- D. metadata management

Answer: A

Question: 7

About the job and execution plan, in E-MapReduce, which of the following descriptions are correct?

(Number of correct answers: 3)

Response:

- A. In E-MapReduce, to create a job is to create a configuration about how to run the job. A job cannot be run directly. The configuration of a job must contain the jar package to be run for the job, the input and output addresses of data, and some running parameters.
- B. You cannot view job logs on the worker nodes in E-MapReduce.
- C. The execution plan is a bond that associates the job and the cluster
- D. Through the execution plan, multiple jobs can be combined into a job sequence and prepare a running cluster for the job (or automatically create a temporary cluster or associate an existing cluster).

Answer: A C D

Question: 8

Jack is analyzing natural languages by using MaxCompute. He has collected 1000 news articles from portal websites with each piece of news saved as a record.

Now, he wants to extract terms from these records (i.e. split contents of each record into multiple independent terms and save them in a separate spreadsheet) with each term saved as an individual record. Jack is familiar with user-defined functions (UDF), and he wants to compile a UDF to extract terms for him.

From the perspective of functionality, please help him decide which type of UDF is more suitable for this scenario?

Response:

- A. User Defined Scalar Function
- B. User Defined Table Valued Function
- C. User Defined Aggregation Function
- D. User Defined Group Function

Answer: B

Question: 9

In DataWorks, SQL Task1 is a daily scheduled periodic task. Task1 will use the data of a partition in Table2 (the Table2 partition field name is ds, in the format of yyyyymmdd) every time Task1 runs.

The Table2 partition value is the first day of the month of the current business date (for example, the run time is May 10, 2018, then the business date will be 20180509, and the Table2 partition value will be 20180501).

To use the time parameters provided by the scheduling system to define the Table2 partitions so that the time value is automatically replaced at every scheduled run of Task1, how should we configure Task1?

Response:

- A. In Task1 code, Table 2 partition ds='\${var}01' and the parameter configuration: var=\${yyyyymm-1}.
- B. In Task1 code, Table 2 partition ds='\${var}' and the parameter configuration: var=\${yyyyymmdd}.
- C. In Task1 code, Table 2 partition ds='\${var}' and the parameter configuration: var=\${yyyyymm}01.
- D. In Task1 code, Table 2 partition ds='\${var}' and the parameter configuration: var=\${yyyyymm01}.

Answer: A

For More Information – **Visit link below:**
<https://www.certsgrade.com/>

PRODUCT FEATURES

-  **100% Money Back Guarantee**
-  **90 Days Free updates**
-  **Special Discounts on Bulk Orders**
-  **Guaranteed Success**
-  **50,000 Satisfied Customers**
-  **100% Secure Shopping**
-  **Privacy Policy**
-  **Refund Policy**

16 USD Discount Coupon Code: NB4XKTMZ



Visit us at: <https://www.certsgrade.com/pdf/acp-bigdata1>