



Alibaba ACP-BigData1

Alibaba ACP Big Data Certification Questions & Answers

Get Instant Access to Vital
Exam Acing Materials |
Study Guide | Sample
Questions | Practice Test

ACP-BIGDATA1
[Alibaba Big Data \(ACP\)](#)

60 Questions Exam – 65/100 Cut Score – Duration of 120 minutes




Table of Contents:

Discover More about the ACP-BigData1 Certification ...	2
Alibaba ACP-BigData1 Certification Details:	2
ACP-BigData1 Syllabus:.....	2
Broaden Your Knowledge with Alibaba ACP-BigData1 Sample Questions:	5
Avail the Study Guide to Pass Alibaba ACP-BigData1 Exam:	8
Career Benefits:	9

Discover More about the ACP-BigData1 Certification

Are you interested in passing the Alibaba ACP-BigData1 exam? First discover, who benefits from the ACP-BigData1 certification. The ACP-BigData1 is suitable for a candidate if he wants to learn about Big data. Passing the ACP-BigData1 exam earns you the Alibaba Big Data (ACP) title.

While preparing for the ACP-BigData1 exam, many candidates struggle to get the necessary materials. But do not worry; your struggling days are over. The ACP-BigData1 PDF contains some of the most valuable preparation tips and the details and instant access to useful ACP-BigData1 study materials [just at one click](#).

Alibaba ACP-BigData1 Certification Details:

Exam Name	Alibaba ACP Big Data (ACP-BigData1)
Exam Code	ACP-BigData1
Exam Price	\$120 USD
Duration	120 minutes
Number of Questions	60
Passing Score	65/100
Recommended Training / Books	ACT82002 Alibaba Cloud MaxCompute & DataWorks Deep Dive
Schedule Exam	PEARSON VUE
Sample Questions	Alibaba ACP Big Data Sample Questions
Recommended Practice	Alibaba Big Data (ACP) Practice Test

ACP-BigData1 Syllabus:

Section	Objectives
MaxCompute	- Familiar with big data computing services basic concepts, including project, table, partition, resources, task, etc.

Section	Objectives
	<ul style="list-style-type: none"> - Understand big data computing services including the composition of the structure and function of each component. - Master the characteristics, advantages and application scenarios of Alibaba Cloud big data computing services. - Know how to connect and use the computing services, including the use of client odpscmd, management console, Java SDK, etc. - Know how to do the big data computing service data upload and download, can use tunnel command line tools, understand the Tunnel SDK - Know how to use SQL commands for large data computing, including DDL, DML, and common built-in functions? - Familiar with user-defined functions, including UDF, UDAF, and UDTF, able to write simple custom functions. - Familiar with the MapReduce programming framework, can setup IntelliJ IDEA integrated development environment, and write a simple MapReduce program. - Understand the Graph programming framework, including basic concepts, processing procedures, can write a simple Graph program. - Familiar with the concept and practical operation of the security and permission management of MaxCompute, including users, roles, authorization (ACL & Policy), project space protection, external and security level, etc.
DataWorks	<ul style="list-style-type: none"> - Familiar with the basic functions of DataWorks, including data Integration, data development, data management, operation & maintenance center, organization management and project management. - Understand the basic features of DataWorks, including role isolation, environment isolation, etc. - Has knowledge about how to leverage project management and organizational management modules to build data analysis environment. - Proficient in the design and development of data development

Section	Objectives
	<p>module of DataWorks, including construction table, task development, resource upload, data upload, new functions, etc.</p> <ul style="list-style-type: none"> - Able to use DataWorks' data development module for workflow task and node task development and design, can configure appropriate dependencies and periodic scheduling. - Able to use the data management module for data management, including lineage analysis, application and authorization of use of table, etc. - Able to fix the basic problems by identifying and locating the problems in the process.
E-MapReduce	<ul style="list-style-type: none"> - Know about the basic distributed system theory, like the concept of distributed file system and distributed computing framework. - Know how the common components in Hadoop ecosystem work, e.g. distributed file system (HDFS), computation framework (MapReduce), resource management component (YARN) and resource coordination component (Zookeeper). - Familiar with the basic concepts of each component of EMapReduce, including YARN, Spark, Zookeeper, Kafka, etc. - Familiar with Auto Scaling features, product advantages and common application scenarios.
Alibaba Big Data Ecosystem tools	<ul style="list-style-type: none"> - Has knowledge about the Alibaba Cloud Machine Learning Platform for AI (PAI). - Know about the Alibaba Cloud Steaming data process solution and products (Realtime Compute, Apache Flink). - Familiar with the basic concepts of Quick BI and has knowledge about workflow of how to use Quick BI get better insight of data. - Understand the features and application scenarios of other related products, including Alibaba Cloud RDS, distributed relational database DRDS, Table Store, Analytic database, Data Transmission Service (DTS), Realtime Compute and DataV, etc. - Understand how DataWorks data integration synchronizes data with other related products?

Broaden Your Knowledge with Alibaba ACP-BigData1 Sample Questions:

Question: 1

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

(Number of correct answers: 3)

- 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: 2

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?

- a) In Task1 code, Table 2 partition ds='\${var}01' and the parameter configuration: var=\${yyyymm-1}.
- b) In Task1 code, Table 2 partition ds='\${var}' and the parameter configuration: var=\${yyyymmdd}.
- c) In Task1 code, Table 2 partition ds='\${var}' and the parameter configuration: var=\${yyyymm}01.
- d) In Task1 code, Table 2 partition ds='\${var}' and the parameter configuration: var=\${yyyymm01}.

Answer: a

Question: 3

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

- a) Column-store
- b) Row-store
- c) Key-value pair storage
- d) Document storage

Answer: a

Question: 4

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)

- a) Test
- b) Data population
- c) Viewing node run logs
- d) Re-running and resuming scheduling

Answer: c, d

Question: 5

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 _____.

- a) component failures
- b) files are huge by traditional standards
- c) data locality
- d) metadata management

Answer: a

Question: 6

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)

- 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: 7

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?

- a) User Defined Scalar Function
- b) User Defined Table Valued Function
- c) User Defined Aggregation Function
- d) User Defined Group Function

Answer: b

Question: 8

MaxCompute supports label-based security (Label Security), which is a mandatory access control (MAC) strategy of the project level. Its introduction is to make the project administrator control user access to sensitive data more flexibly.

Which granularity of sensitive data can be controlled by Label Security?

- a) Table;
- b) Partition;
- c) Row;
- d) Column;

Answer: d

Question: 9

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.

- a) select size from table_a;
- b) size table_a;
- c) desc table_a;
- d) show table table_a;

Answer: c

Question: 10

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?

- 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

Avail the Study Guide to Pass Alibaba ACP-BigData1 Exam:

- Find out about the ACP-BigData1 syllabus topics. Visiting the official site offers an idea about the exam structure and other important study resources. Going through the syllabus topics help to plan the exam in an organized manner.
- Once you are done exploring the [ACP-BigData1 syllabus](#), it is time to plan for studying and covering the syllabus topics from the core. Chalk out the best plan for yourself to cover each part of the syllabus in a hassle-free manner.
- A study schedule helps you to stay calm throughout your exam preparation. It should contain your materials and thoughts like study hours, number of topics for daily studying mentioned on it. The best bet to clear the exam is to follow your schedule rigorously.
- The candidate should not miss out on the scope to learn from the ACP-BigData1 training. Joining the Alibaba provided training for ACP-

BigData1 exam helps a candidate to strengthen his practical knowledge base from the certification.

- Learning about the probable questions and gaining knowledge regarding the exam structure helps a lot. Go through the [ACP-BigData1 sample questions](#) and boost your knowledge
- Make yourself a pro through online practicing the syllabus topics. ACP-BigData1 practice tests would guide you on your strengths and weaknesses regarding the syllabus topics. Through rigorous practicing, you can improve the weaker sections too. Learn well about time management during exam and become confident gradually with practice tests.

Career Benefits:

- Passing the ACP-BigData1 exam, helps a candidate to prosper highly in his career. Having the certification on the resume adds to the candidate's benefit and helps to get the best opportunities.

Here Is the Trusted Practice Test for the ACP-BigData1 Certification

VMExam.Com is here with all the necessary details regarding the ACP-BigData1 exam. We provide authentic practice tests for the ACP-BigData1 exam. What do you gain from these practice tests? You get to experience the real exam-like questions made by industry experts and get a scope to improve your performance in the actual exam. Rely on VMExam.Com for rigorous, unlimited two-month attempts on the [ACP-BigData1 practice tests](#), and gradually build your confidence. Rigorous practice made many aspirants successful and made their journey easy towards grabbing the Alibaba Big Data (ACP).

Start Online practice of ACP-BigData1 Exam by visiting URL

<https://www.vmexam.com/alibaba/acp-bigdata1-alibaba-acp-big-data>