

AWS BDS-C00 Study Guide PDF

GRAB THE AWS BIG DATA SPECIALTY CERTIFICATION PDF QUESTIONS & ANSWERS

Exam Summary – Syllabus – Questions

BDS-C00

AWS Certified Big Data - Specialty

65 Questions Exam – 750 / 1000 Cut Score – Duration of 170 minutes

www.VMExam.com

Table of Contents

Get an Overview of the BDS-C00 Certification:.....	3
Why Should You Earn the AWS BDS-C00 Certification?	3
What Is the AWS BDS-C00 Big Data Specialty Certification Exam Structure?	3
Enhance Knowledge with BDS-C00 Sample Questions:	4
What Study Guide Works Best in Acing the AWS BDS- C00 Big Data Specialty Certification?	8
Explore the Syllabus Topics and Learn from the Core:	8
Make Your Schedule:	8
Get Expert Advice from the Training:	8
Get Access to the PDF Sample Questions:	8
Avoid Dumps and Utilize the AWS BDS-C00 Practice Test:.....	8

Get an Overview of the BDS-C00 Certification:

Who should take the [BDS-C00 exam](#)? This is the first question that comes to a candidate's mind when preparing for the Big Data Specialty certification. The BDS-C00 certification is suitable for candidates who are keen to earn knowledge on the Specialty and grab their AWS Certified Big Data - Specialty. When it is about starting the preparation, most candidates get confused regarding the study materials and study approach. But BDS-C00 study guide PDF is here to solve the problem. BDS-C00 PDF combines some effective sample questions and offers valuable tips to pass the exam with ease.

Why Should You Earn the AWS BDS-C00 Certification?

There are several reasons why one should grab the BDS-C00 certification.

- The Big Data Specialty certification proves to be one of the most recognized certifications.
- The certification badge proves the knowledge of the candidate regarding subject matters and makes his resume presentable to potential candidates.
- Thus earning the [AWS Certified Big Data - Specialty](#) is a powerful qualification for a prosperous career.

What Is the AWS BDS-C00 Big Data Specialty Certification Exam Structure?

Exam Name	AWS Certified Big Data - Specialty (Big Data Specialty)
Exam Code	BDS-C00
Exam Price	\$300 USD
Duration	170 minutes
Number of Questions	65
Passing Score	750 / 1000
Schedule Exam	PEARSON VUE
Sample Questions	AWS BDS-C00 Sample Questions
Recommended Practice	AWS Certified Big Data - Specialty Practice Test

Enhance Knowledge with BDS-C00 Sample Questions:

Question: 1

A data engineer needs to architect a data warehouse for an online retail company to store historic purchases. The data engineer needs to use Amazon Redshift.

To comply with PCI:DSS and meet corporate data protection standards, the data engineer must ensure that data is encrypted at rest and that the keys are managed by a corporate on-premises HSM.

Which approach meets these requirements in the most cost-effective manner?

- a) Create a VPC, and then establish a VPN connection between the VPC and the on-premises network. Launch the Amazon Redshift cluster in the VPC, and configure it to use your corporate HSM.
- b) Use the AWS CloudHSM service to establish a trust relationship between the CloudHSM and the corporate HSM over a Direct Connect connection. Configure Amazon Redshift to use the CloudHSM device.
- c) Configure the AWS Key Management Service to point to the corporate HSM device, and then launch the Amazon Redshift cluster with the KMS managing the encryption keys.
- d) Use AWS Import/Export to import the corporate HSM device into the AWS Region where the Amazon Redshift cluster will launch, and configure Redshift to use the imported HSM.

Answer: a

Question: 2

A data engineer needs to collect data from multiple Amazon Redshift clusters within a business and consolidate the data into a single central data warehouse. Data must be encrypted at all times while at rest or in flight.

What is the most scalable way to build this data collection process?

- a) Run an ETL process that connects to the source clusters using SSL to issue a SELECT query for new data, and then write to the target data warehouse using an INSERT command over another SSL secured connection.
- b) Use AWS KMS data key to run an UNLOAD ENCRYPTED command that stores the data in an unencrypted S3 bucket; run a COPY command to move the data into the target cluster.
- c) Run an UNLOAD command that stores the data in an S3 bucket encrypted with an AWS KMS data key; run a COPY command to move the data into the target cluster.
- d) Connect to the source cluster over an SSL client connection, and write data records to Amazon Kinesis Firehose to load into your target data warehouse.

Answer: b

Question: 3

A mobile application collects data that must be stored in multiple Availability Zones within five minutes of being captured in the app.

What architecture securely meets these requirements?

- a) The mobile app should write to an S3 bucket that allows anonymous PutObject calls.
- b) The mobile app should authenticate with an Amazon Cognito identity that is authorized to write to an Amazon Kinesis Firehose with an Amazon S3 destination.
- c) The mobile app should authenticate with an embedded IAM access key that is authorized to write to an Amazon Kinesis Firehose with an Amazon S3 destination.
- d) The mobile app should call a REST-based service that stores data on Amazon EBS. Deploy the service on multiple EC2 instances across two Availability Zones.

Answer: b

Question: 4

A company needs to deploy a data lake solution for their data scientists in which all company data is accessible and stored in a central S3 bucket.

The company segregates the data by business unit, using specific prefixes. Scientists can only access the data from their own business unit.

The company needs a single sign-on identity and management solution based on Microsoft Active Directory (AD) to manage access to the data in Amazon S3.

Which method meets these requirements?

- a) Use AWS IAM Federation functions and specify the associated role based on the users' groups in AD.
- b) Create bucket policies that only allow access to the authorized prefixes based on the users' group name in Active Directory.
- c) Deploy the AD Synchronization service to create AWS IAM users and groups based on AD information.
- d) Use Amazon S3 API integration with AD to impersonate the users on access in a transparent manner.

Answer: a

Question: 5

An administrator has a 500-GB file in Amazon S3. The administrator runs a nightly COPY command into a 10-node Amazon Redshift cluster. The administrator wants to prepare the data to optimize performance of the COPY command.

How should the administrator prepare the data?

- a) Compress the file using gz compression.
- b) Split the file into 500 smaller files.
- c) Convert the file format to AVRO.
- d) Split the file into 10 files of equal size.

Answer: b

Question: 6

A customer needs to load a 550-GB data file into an Amazon Redshift cluster from Amazon S3, using the COPY command.

The input file has both known and unknown issues that will probably cause the load process to fail. The customer needs the most efficient way to detect load errors without performing any cleanup if the load process fails.

Which technique should the customer use?

- a) Compress the input file before running COPY.
- b) Write a script to delete the data from the tables in case of errors.
- c) Split the input file into 50-GB blocks and load them separately.
- d) Use COPY with NOLOAD parameter.

Answer: d

Question: 7

A company logs data from its application in large files and runs regular analytics of these logs to support internal reporting for three months after the logs are generated.

After three months, the logs are infrequently accessed for up to a year. The company also has a regulatory control requirement to store application logs for seven years.

Which course of action should the company take to achieve these requirements in the most cost-efficient way?

- a) Store the files in S3 Glacier with a Deny Delete vault lock policy for archives less than seven years old and a vault access policy that restricts read access to the analytics IAM group and write access to the log writer service role.
- b) Store the files in S3 Standard with a lifecycle policy to transition the storage class to Standard - IA after three months. After a year, transition the files to Glacier and add a Deny Delete vault lock policy for archives less than seven years old.
- c) Store the files in S3 Standard with lifecycle policies to transition the storage class to Standard - IA after three months and delete them after a year. Simultaneously store the files in Amazon Glacier with a Deny Delete vault lock policy for archives less than seven years old.
- d) Store the files in S3 Standard with a lifecycle policy to remove them after a year. Simultaneously store the files in Amazon S3 Glacier with a Deny Delete vault lock policy for archives less than seven years old.

Answer: c

Question: 8

An organization needs a data store to handle the following data types and access patterns:

- Key-value access pattern
- Complex SQL queries and transactions
- Consistent reads
- Fixed schema

Which data store should the organization choose?

- Amazon S3
- Amazon Kinesis
- Amazon DynamoDB
- Amazon RDS

Answer: d

Question: 9

An administrator decides to use the Amazon Machine Learning service to classify social media posts that mention your company into two categories:

posts that require a response and posts that do not. The training dataset of 10,000 posts contains the details of each post, including the timestamp, author, and full text of the post. You are missing the target labels that are required for training.

Which two options will create valid target label data?

- Ask the social media handling team to review each post and provide the label.
- Use the sentiment analysis NLP library to determine whether a post requires a response.
- Use the Amazon Mechanical Turk web service to publish Human Intelligence Tasks that ask Turk workers to label the posts.
- Using the a priori probability distribution of the two classes, use Monte-Carlo simulation to generate the labels.

Answer: a, c

Question: 10

A web application emits multiple types of events to Amazon Kinesis Streams for operational reporting. Critical events must be captured immediately before processing can continue, but informational events do not need to delay processing.

What is the most appropriate solution to record these different types of events?

- Log all events using the Kinesis Producer Library.
- Log critical events using the Kinesis Producer Library, and log informational events using the PutRecords API method.
- Log critical events using the PutRecords API method, and log informational events using the Kinesis Producer Library.
- Log all events using the PutRecords API method.

Answer: c

What Study Guide Works Best in Acing the AWS BDS-C00 Big Data Specialty Certification?

The BDS-C00 study guide is a combination of some proven study tips and the combination of all valuable study materials like sample questions, syllabus and practice tests in one place.

Explore the Syllabus Topics and Learn from the Core:

If you are determined to earn success in the Big Data Specialty exam, getting in full touch of the [syllabus](#) is mandatory. During preparation, you might not like all syllabus sections or topics, but try to get at least the fundamental knowledge from the sections you don't like. The more you possess knowledge on all syllabus sections, the more is the chance to attempt maximum number of questions during the actual exam.

Make Your Schedule:

Studying and completing the syllabus becomes easier, if you work on the syllabus topics after making a schedule. Your syllabus must mention what areas you want to cover and within what time. Once you make a schedule and follow it regularly, syllabus completion becomes easier and preparation becomes smoother.

Get Expert Advice from the Training:

Do not forget to join the AWS BDS-C00 training if it is providing any. Training enhances the practical knowledge of a candidate, which helps them to work well in the practical field during projects.

Get Access to the PDF Sample Questions:

If your study material is in a [PDF format](#) or the materials are mobile-friendly, what could be better than that? Get access to the free sample questions and keep enhancing your knowledge beyond the syllabus.

Avoid Dumps and Utilize the AWS BDS-C00 Practice Test:

Why should you rely on practice tests? The reason is simple: you must get familiar with the exam pattern before reaching the exam hall. An aspirant aware of the exam structure and time management during the exam preparation can perform well in the actual exam and attempt the maximum number of questions during the exam.

Many aspirants prefer to read from dumps, but they miss out on the self assessment method. Therefore, BDS-C00 practice tests always stand out to be the better choice than dumps PDF.

Avail the Proven BDS-C00 Practice Test for Success!!!

Do you want to pass the BDS-C00 exam on your first attempt? Stop worrying; we, VMExam.com are here to provide you the best experience during your AWS Certified Big Data - Specialty preparation. Try out our free mock tests to get a glimpse of our quality study materials, and build your confidence with the premium [BDS-C00 practice tests](#). Our expert-designed questions help you to improve performance and pass the exam on your first attempt.