



Free Amazon BDS-C00 Practice Questions

Shared by Brewer on 12-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

A real-time bidding company is rebuilding their monolithic application and is focusing on serving real-time data. A large number of reads and writes are generated from thousands of concurrent users who follow items and bid on the company's sale offers.

The company is experiencing high latency during special event spikes, with millions of concurrent users.

The company needs to analyze and aggregate a part of the data in near real time to feed an internal dashboard.

What is the BEST approach for serving and analyzing data, considering the constraint of the row latency on the highly demanded data?

Options:

- A- Use Amazon Aurora with Multi Availability Zone and read replicas. Use Amazon ElastiCache in front of the read replicas to serve read-only content quickly. Use the same database as datasource for the dashboard.
- B- Use Amazon DynamoDB to store real-time data with Amazon DynamoDB. Accelerator to serve content quickly. use Amazon DynamoDB Streams to replay all changes to the table, process and stream to Amazon Elasti search Service with AWS Lambda.
- C- Use Amazon RDS with Multi Availability Zone. Provisioned IOPS EBS volume for storage. Enable up to five read replicas to serve read-only content quickly. Use Amazon EMR with Sqoop to import Amazon RDS data into HDFS for analysis.
- D- Use Amazon Redshift with a DC2 node type and a multi-mode cluster. Create an Amazon EC2 instance with pgpool1 installed. Create an Amazon ElastiCache cluster and route read requests through pgpool1, and use Amazon Redshift for analysis.

Answer:

B

Explanation:

1. <https://rockset.com/blog/live-dashboards-dynamodb-streams-lambda-elasticache/>
2. <https://aws.amazon.com/blogs/compute/indexing-amazon-dynamodb-content-with-amazon-elasticsearch-service-using-aws-lambda/> 3. <https://aws.amazon.com/blogs/startups/combining-dynamodb-and-amazon-elasticsearch-with-lambda/>
4. https://d1.awsstatic.com/whitepapers/Big_Data_Analytics_Options_on_AWS.pdf?did=wp_card&trk=wp_card

Question 2

Question Type: MultipleChoice

You are deploying an application to collect votes for a very popular television show. Millions of users will submit votes using mobile devices. The votes must be collected into a durable, scalable, and highly available data store for real-time public tabulation. Which service should you use?

Options:

- A- Amazon DynamoDB
- B- Amazon Redshift
- C- Amazon Kinesis
- D- Amazon Simple Queue Service

Answer:

C

Question 3

Question Type: MultipleChoice

The project you are working on currently uses a single AWS CloudFormation template to deploy its AWS infrastructure, which supports a multi-tier web application. You have been tasked with organizing the AWS CloudFormation resources so that they can be maintained in the future, and so that different departments such as Networking and Security can review the architecture before it goes to Production.

How should you do this in a way that accommodates each department, using their existing workflows?

Options:

- A- Organize the AWS CloudFormation template so that related resources are next to each other in the template, such as VPC subnets and routing rules for Networking and Security groups and IAM information for Security
- B- Separate the AWS CloudFormation template into a nested structure that has individual templates for the resources that are to be governed by different departments, and use the outputs from the networking and security stacks for the application template that you control

C- Organize the AWS CloudFormation template so that related resources are next to each other in the template for each department's use, leverage your existing continuous integration tool to constantly deploy changes from all parties to the Production environment, and then run tests for validation

D- Use a custom application and the AWS SDK to replicate the resources defined in the current AWS CloudFormation template, and use the existing code review system to allow other departments to approve changes before altering the application for future deployments

Answer:

B

Question 4

Question Type: MultipleChoice

Fill in the blanks: A ____ is a storage device that moves data in sequences of bytes or bits (blocks). Hint: These devices support random access and generally use buffered I/O.

Options:

A- block map

B- storage block

C- mapping device

D- block device

Answer:

D

Question 5

Question Type: MultipleChoice

You have been tasked with implementing an automated data backup solution for your application servers that run on Amazon EC2 with Amazon EBS volumes. You want to use a distributed data store for your backups to avoid single points of failure and to increase the durability of the data. Daily backups should be retained for 30 days so that you can restore data within an hour.

How can you implement this through a script that a scheduling daemon runs daily on the application servers?

Options:

- A- Write the script to call the ec2-create-volume API, tag the Amazon EBS volume with the current data time group, and copy backup data to a second Amazon EBS volume. Use the ec2-describe-volumes API to enumerate existing backup volumes. Call the ec2-delete-volume API to prune backup volumes that are tagged with a date-time group older than 30 days
- B- Write the script to call the Amazon Glacier upload archive API, and tag the backup archive with the current date-time group. Use the list vaults API to enumerate existing backup archives. Call the delete vault API to prune backup archives that are tagged with a date-time group older than 30 days
- C- Write the script to call the ec2-create-snapshot API, and tag the Amazon EBS snapshot with the current date-time group. Use the ec2-describe-snapshot API to enumerate existing Amazon EBS snapshots. Call the ec2-delete-snapshot API to prune Amazon EBS snapshots that are tagged with a date-time group older than 30 days
- D- Write the script to call the ec2-create-volume API, tag the Amazon EBS volume with the current date-time group, and use the ec2-copy-snapshot API to backup data to the new Amazon EBS volume. Use the ec2-describe-snapshot API to enumerate existing backup volumes. Call the ec2-delete-snapshot API to prune backup Amazon EBS volumes that are tagged with a date-time group older than 30 days

Answer:

C

Question 6

Question Type: MultipleChoice

Can I detach the primary (eth0) network interface when the instance is running or stopped?

Options:

- A- Yes, You can.
- B- No. You cannot
- C- Depends on the state of the interface at the time

Answer:

B

Question 7

Question Type: MultipleChoice

You have a video Trans coding application running on Amazon EC2. Each instance pools a queue to find out which video should be Trans coded, and then runs a Trans coding process.

If this process is interrupted, the video will be Trans coded by another instance based on the queuing system. You have a large backlog of videos which need to be Trans coded and would like to reduce this backlog by adding more instances. You will need these instances only until the backlog is reduced. Which type of Amazon EC2 instance should you use to reduce the backlog in the most cost-effective way?

Options:

- A- Dedicated instances
- B- Spot instances
- C- On-demand instances
- D- Reserved instances

Answer:

B

Question 8

Question Type: MultipleChoice

A company is running a batch analysis every hour on their main transactional DB running on an RDS MySQL instance to populate their central Data Warehouse running on Redshift. During the execution of the batch their transactional applications are very slow. When the batch completes they need to update the top management dashboard with the new data. The dashboard is produced by another system running on-premises that is currently started when a manually-sent email notifies that an update is required. The on-premises system cannot be modified because it is managed by another team. How would you optimize this scenario to solve performance issues and automate the process as much as possible?

Options:

- A- Replace RDS with Redshift for the batch analysis and SNS to notify the on-premises system to update the dashboard
- B- Replace RDS with Redshift for the batch analysis and SQS to send a message to the on-premises system to update the dashboard
- C- Create an RDS Read Replica for the batch analysis and SNS to notify the on-premises system to update the dashboard
- D- Create an RDS Read Replica for the batch analysis and SQS to send a message to the on-premises system to update the dashboard.

Answer:

C

Question 9

Question Type: MultipleChoice

A web-hosting company is building a web analytics tools to capture clickstream data from all of the websites hosted within its platform and to provide near-real-time business intelligence. This entire system is built on AWS services. The web-hosting company is interested in using Amazon Kinesis to collect this data and perform sliding window analytics. What is the most reliable and fault-tolerant technique to get each website to send data to Amazon Kinesis with every click?

Options:

- A- After receiving a request each web server sends it to Amazon Kinesis using the Amazon Kinesis PutRecord API. Use the SessionID as a partition key and set up a loop to retry until a success response is received
- B- After receiving a request each web server sends it to Amazon Kinesis using the Amazon Kinesis Producer Library addRecord method
- C- Each web server buffers the request until the count reaches 500 and sends them to Amazon Kinesis using the Amazon Kinesis PutRecord API call
- D- After receiving a request each web server sends it to Amazon Kinesis using the Amazon Kinesis PutRecord API. Use the exponential back off algorithm for retries until a successful response is received

Answer:

A

Explanation:

<https://docs.aws.amazon.com/streams/latest/dev/kinesis-producer-adv-retries-rate-limiting.html>

Question 10

Question Type: MultipleChoice

An administrator tries to use the Amazon Machine Learning service to classify social media posts that mention the administrator's company into posts that requires 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. The administrator is missing the target labels that are required for training.

Which Amazon Machine Learning model is the most appropriate for the task?

Options:

- A- Unary classification model, where the target class is the require-response post
- B- Binary classification model, where the two classes are require-response and does-not-require-response
- C- Multi-class prediction model, with two classes require-response and does-not-require response
- D- Regression model where the predicted value is the probability that the post requires a response

Answer:

D

To Get Premium Files for BDS-C00 Visit

<https://www.p2pexams.com/products/bds-c00>

For More Free Questions Visit

<https://www.p2pexams.com/amazon/pdf/bds-c00>

20%
DISCOUNT

P2P
exams