



# BDS-C00<sup>Q&As</sup>

AWS Certified Big Data - Speciality (BDS-C00)

**Pass Amazon BDS-C00 Exam with 100% Guarantee**

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/aws-certified-big-data-specialty.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon  
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





### QUESTION 1

By default what are ENIs that are automatically created and attached to instances using the EC2 console set to do when the attached instance terminates?

- A. Remain as is
- B. Terminate
- C. Hibernate
- D. Pause

Correct Answer: B

---

### QUESTION 2

An advertising organization uses an application to process a stream of events that are received from clients in multiple unstructured formats.

The application does the following:

Transforms the events into a single structured format and streams them to Amazon Kinesis for real-time analysis.

Stores the unstructured raw events from the log files on local hard drivers that are rotated and uploaded to Amazon S3.

The organization wants to extract campaign performance reporting using an existing Amazon redshift cluster.

Which solution will provide the performance data with the LEAST number of operations?

- A. Install the Amazon Kinesis Data Firehose agent on the application servers and use it to stream the log files directly to Amazon Redshift.
- B. Create an external table in Amazon Redshift and point it to the S3 bucket where the unstructured raw events are stored.
- C. Write an AWS Lambda function that triggers every hour to load the new log files already in S3 to Amazon redshift.
- D. Connect Amazon Kinesis Data Firehose to the existing Amazon Kinesis stream and use it to stream the event directly to Amazon Redshift.

Correct Answer: B

---



### QUESTION 3

A company needs to deploy services to an AWS region which they not previously used. The company currently has an AWS identity and Access Management (IAM) role for their Amazon EC2 instances, which permits the instance to have access to Amazon DynamoDB. The company wants their EC2 instances in the new region to have the same privileges. How should the company achieve this?

- A. Create a new IAM role and associated policies within the new region
- B. Assign the existing IAM role to the Amazon EC2 instances in the new region
- C. Copy the IAM role and associated policies to the new region and attach it to the instances
- D. Create the Amazon Machine Image of the instance and copy it to the desired region using the AMI Copy feature

Correct Answer: B

---

### QUESTION 4

When you put objects in Amazon S3, what is the indication that an object was successfully stored?

- A. A HTTP 200 result code and MD5 checksum, taken together, indicate that the operation was successful
- B. A success code is inserted into the S3 object metadata
- C. Amazon S3 is engineered for 99.999999999% durability. Therefore there is no need to confirm that data was inserted.
- D. Each S3 account has a special bucket named `_s3_logs`. Success codes are written to this bucket with a timestamp and checksum

Correct Answer: A

---

### QUESTION 5

You are building a mobile app for consumers to post cat pictures online. You will be storing the images in AWS S3. You want to run the system very cheaply and simply. Which one of these options allows you to build a photo sharing application without needing to worry about scaling expensive uploads processes, authentication/authorization and so forth?

- A. Build the application out using AWS Cognito and web identity federation to allow users to log in using Facebook or Google Accounts. Once they are logged in, the secret token passed to that user is used to directly access resources on AWS, like AWS S3. (Amazon Cognito is a superset of the functionality provided by web identity federation. Referlink)
- B. Use JWT or SAML compliant systems to build authorization policies. Users log in with a username and password, and are given a token they can use indefinitely to make calls against the photo infrastructure.
- C. Use AWS API Gateway with a constantly rotating API Key to allow access from the client-side. Construct a custom build of the SDK and include S3 access in it.
- D. Create an AWS OAuth Service Domain and grant public signup and access to the domain. During setup, add at least



one major social media site as a trusted Identity Provider for users

Correct Answer: A

---

#### QUESTION 6

A social media customer has data from different data sources including RDS running MySQL, Redshift, and Hive on EMR. To support better analysis, the customer needs to be able to analyze data from different data sources and to combine the results.

What is the most cost-effective solution to meet these requirements?

- A. Load all data from a different database/warehouse to S3. Use Redshift COPY command to copy data to Redshift for analysis.
- B. Install Presto on the EMR cluster where Hive sits. Configure MySQL and PostgreSQL connector to select from different data sources in a single query.
- C. Spin up an Elasticsearch cluster. Load data from all three data sources and use Kibana to analyze.
- D. Write a program running on a separate EC2 instance to run queries to three different systems. Aggregate the results after getting the responses from all three systems.

Correct Answer: B

---

#### QUESTION 7

You have started a new job and are reviewing your company's infrastructure on AWS. You notice one web application where they have an Elastic Load Balancer (ELB) in front of web instances in an Auto Scaling Group. When you check the metrics for the ELB in CloudWatch, you see four healthy instances in Availability Zone (AZ) A and zero in AZ B. There are zero unhealthy instances. What do you need to fix to balance the instances across AZs?

- A. Set the ELB to only be attached to another AZ.
- B. Make sure Auto Scaling is configured to launch in both AZs.
- C. Make sure your AMI is available in both AZs.
- D. Make sure the maximum size of the Auto Scaling Group is greater than 4.

Correct Answer: B

---

#### QUESTION 8

A company is building a new application in AWS. The architect needs to design a system to collect application log events. The design should be a repeatable pattern that minimizes data loss if an application instance fails, and keeps a durable copy of a log data for at least 30 days.

What is the simplest architecture that will allow the architect to analyze the logs?

- A. Write them directly to a Kinesis Firehose. Configure Kinesis Firehose to load the events into an Amazon Redshift



cluster for analysis.

B. Write them to a file on Amazon Simple Storage Service (S3). Write an AWS Lambda function that runs in response to the S3 event to load the events into Amazon Elasticsearch Service for analysis.

C. Write them to the local disk and configure the Amazon CloudWatch Logs agent to load the data into CloudWatch Logs and subsequently into Amazon Elasticsearch Service.

D. Write them to CloudWatch Logs and use an AWS Lambda function to load them into HDFS on an Amazon Elastic MapReduce (EMR) cluster for analysis.

Correct Answer: B

---

### QUESTION 9

A new algorithm has been written in Python to identify SPAM e-mails. The algorithm analyzes the free text contained within a sample set of 1 million e-mails stored on Amazon S3. The algorithm must be scaled across a production dataset of 5 PB, which also resides in Amazon S3 storage.

Which AWS service strategy is best for this use case?

A. Copy the data into Amazon ElastiCache to perform text analysis on the in-memory data and export the results of the model into Amazon Machine Learning.

B. Use Amazon EMR to parallelize the text analysis tasks across the cluster using a streaming program step.

C. Use Amazon Elasticsearch Service to store the text and then use the Python Elasticsearch Client to run analysis against the text index.

D. Initiate a Python job from AWS Data Pipeline to run directly against the Amazon S3 text files.

Correct Answer: C

Reference: <https://aws.amazon.com/blogs/database/indexing-metadata-in-amazon-elasticsearch-service-using-aws-lambda-and-python/>

---

### QUESTION 10

A user wants to make so that whenever the CPU utilization of the AWS EC2 instance is above 90%, the redlight of his bedroom turns on. Which of the below mentioned AWS services is helpful for this purpose?

A. AWS CloudWatch + AWS SES

B. AWS CloudWatch + AWS SNS

C. It is not possible to configure the light with the AWS infrastructure services

D. AWS CloudWatch and a dedicated software turning on the light

Correct Answer: B

---



### QUESTION 11

You run a web application with the following components Elastic Load Balancer (ELB), 3 Web/Application servers, 1 MySQL RDS database with read replicas, and Amazon Simple Storage Service (Amazon S3) for static content. Average response time for users is increasing slowly.

What three CloudWatch RDS metrics will allow you to identify if the database is the bottleneck? Choose 3 answers

- A. The number of outstanding IOs waiting to access the disk
- B. The amount of write latency
- C. The amount of disk space occupied by binary logs on the master.
- D. The amount of time a Read Replica DB Instance lags behind the source DB Instance
- E. The average number of disk I/O operations per second.

Correct Answer: ABD

---

### QUESTION 12

Using only AWS services. You intend to automatically scale a fleet of stateless of stateless web servers based on CPU and network utilization metrics. Which of the following services are needed? Choose 2 answers

- A. Auto Scaling
- B. Amazon Simple Notification Service
- C. AWS Cloud Formation
- D. CloudWatch
- E. Amazon Simple Workflow Service

Correct Answer: AD

---

### QUESTION 13

True or False: When you add a rule to a DB security group, you do not need to specify port number or protocol.

- A. Depends on the RDMS used
- B. TRUE
- C. FALSE

Correct Answer: B

---

**QUESTION 14**

A gas company needs to monitor gas pressure in their pipelines. Pressure data is streamed from sensors placed throughout the pipelines to monitor the data in real time. When an anomaly is detected, the system must send a notification to open valve. An Amazon Kinesis stream collects the data from the sensors and an anomaly Kinesis stream triggers an AWS Lambda function to open the appropriate valve.

Which solution is the MOST cost-effective for responding to anomalies in real time?

- A. Attach a Kinesis Firehose to the stream and persist the sensor data in an Amazon S3 bucket. Schedule an AWS Lambda function to run a query in Amazon Athena against the data in Amazon S3 to identify anomalies. When a change is detected, the Lambda function sends a message to the anomaly stream to open the valve.
- B. Launch an Amazon EMR cluster that uses Spark Streaming to connect to the Kinesis stream and Spark machine learning to detect anomalies. When a change is detected, the Spark application sends a message to the anomaly stream to open the valve.
- C. Launch a fleet of Amazon EC2 instances with a Kinesis Client Library application that consumes the stream and aggregates sensor data over time to identify anomalies. When an anomaly is detected, the application sends a message to the anomaly stream to open the valve.
- D. Create a Kinesis Analytics application by using the RANDOM\_CUT\_FOREST function to detect an anomaly. When the anomaly score that is returned from the function is outside of an acceptable range, a message is sent to the anomaly stream to open the valve.

Correct Answer: A

**QUESTION 15**

Can the string value of `\\Key\\` be prefixed with laws?

- A. No
- B. Only for EC2 not S3
- C. Yes
- D. Only for S3 not EC

Correct Answer: A

[BDS-C00 PDF Dumps](#)

[BDS-C00 Exam Questions](#)

[BDS-C00 Braindumps](#)