BDS-C00 Q&As

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.pass4itsure.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

A company uses Amazon Redshift for its enterprise data warehouse. A new on-premises PostgreSQL OLTP DB must be integrated into the data warehouse. Each table in the PostgreSQL DB has an indexed last_modified timestamp column. The data warehouse has a staging layer to load source data into the data warehouse environment for further processing.

The data lag between the source PostgreSQL DB and the Amazon Redshift staging layer should NOT exceed four hours.

What is the most efficient technique to meet these requirements?

A. Create a DBLINK on the source DB to connect to Amazon Redshift. Use a PostgreSQL trigger on the source table to capture the new insert/update/delete event and execute the event on the Amazon Redshift staging table.

B. Use a PostgreSQL trigger on the source table to capture the new insert/update/delete event and write it to Amazon Kinesis Streams. Use a KCL application to execute the event on the Amazon Redshift staging table.

C. Extract the incremental changes periodically using a SQL query. Upload the changes to multiple Amazon Simple Storage Service (S3) objects, and run the COPY command to load to the Amazon Redshift staging layer.

D. Extract the incremental changes periodically using a SQL query. Upload the changes to a single Amazon Simple Storage Service (S3) object, and run the COPY command to load to the Amazon Redshift staging layer.

Correct Answer: C

QUESTION 2

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

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

A. Binary classification model, where the target class is the require-response post

B. Binary classification model, where the two classes are the require-response post and does-not-requireresponse

C. Multi-class prediction model, with two classes: require-response post and does-not-require-response

D. Regression model where the predicted value is the probability that the post requires a response

Correct Answer: A

QUESTION 3

What does Amazon ELB stand for?

A. Elastic Linux Box.

B. Encrypted Linux Box.

C. Encrypted Load Balancing.

D. Elastic Load Balancing.

Correct Answer: D

QUESTION 4

What is one key difference between an Amazon EBS-backed and an instance-store backed instance?

A. Amazon EBS-backed instances can be stopped and restarted

B. Instance-store backed instances can be stopped and restarted

C. Auto scaling requires using Amazon EBS-backed instances

D. Virtual Private Cloud requires EBS backed instances

Correct Answer: A

QUESTION 5

A system engineer for a company proposes digitalization and backup of large archives for customers. The systems engineer needs to provide users with a secure storage that makes sure that data will never be tampered with once it has been uploaded.

How should this be accomplished?

A. Create an Amazon Glacier Vault. Specify a "Deny" Vault Lock policy on this Vault to block "glacier:DeleteArchive".

B. Create an Amazon S3 bucket. Specify a "Deny" bucket policy on this bucket to block "s3:DeleteObject".

C. Create an Amazon Glacier Vault. Specify a "Deny" vault access policy on this Vault to block "glacier:DeleteArchive".

D. Create secondary AWS Account containing an Amazon S3 bucket. Grant "s3:PutObject" to the primary account.

Correct Answer: C

Reference: https://docs.aws.amazon.com/amazonglacier/latest/dev/vault-lock-policy.html