

Latest Version: 48.0

Question: 1

A company stores sales transaction data in Amazon DynamoDB tables. To detect anomalous behaviors and respond quickly, all changes to the items stored in the DynamoDB tables must be logged within 30 minutes.

Which solution meets the requirements?

- A. Copy the DynamoDB tables into Apache Hive tables on Amazon EMR every hour and analyze them (or anomalous behaviors). Send Amazon SNS notifications when anomalous behaviors are detected.
- B. Use AWS CloudTrail to capture all the APIs that change the DynamoDB tables. Send SNS notifications when anomalous behaviors are detected using CloudTrail event filtering.
- C. Use Amazon DynamoDB Streams to capture and send updates to AWS Lambda. Create a Lambda function to output records to Amazon Kinesis Data Streams. Analyze any anomalies with Amazon Kinesis Data Analytics. Send SNS notifications when anomalous behaviors are detected.
- D. Use event patterns in Amazon CloudWatch Events to capture DynamoDB API call events with an AWS Lambda function as a target to analyze behavior. Send SNS notifications when anomalous behaviors are detected.

Answer: C

Question: 2

A company has a complex web application that leverages Amazon CloudFront for global scalability and performance. Over time, users report that the web application is slowing down.

The company's operations team reports that the CloudFront cache hit ratio has been dropping steadily. The cache metrics report indicates that query strings on some URLs are inconsistently ordered and are specified sometimes in mixed-case letters and sometimes in lowercase letters.

Which set of actions should the solutions architect take to increase the cache hit ratio as quickly as possible?

- A. Deploy a Lambda@Edge function to sort parameters by name and force them to be lowercase. Select the CloudFront viewer request trigger to invoke the function.
- B. Update the CloudFront distribution to disable caching based on query string parameters.
- C. Deploy a reverse proxy after the load balancer to post-process the emitted URLs in the application to force the URL strings to be lowercase.
- D. Update the CloudFront distribution to specify casing-insensitive query string processing.

Answer: C

Question: 3

A company is running an Apache Hadoop cluster on Amazon EC2 instances. The Hadoop cluster stores approximately 100 TB of data for weekly operational reports and allows occasional access for data scientists to retrieve data. The company needs to reduce the cost and operational complexity for storing and serving this data.

Which solution meets these requirements in the MOST cost-effective manner?

- A. Move the Hadoop cluster from EC2 instances to Amazon EMR. Allow data access patterns to remain the same.
- B. Write a script that resizes the EC2 instances to a smaller instance type during downtime and resizes the instances to a larger instance type before the reports are created.
- C. Move the data to Amazon S3 and use Amazon Athena to query the data for reports. Allow the data scientists to access the data directly in Amazon S3.
- D. Migrate the data to Amazon DynamoDB and modify the reports to fetch data from DynamoDB. Allow the data scientists to access the data directly in DynamoDB.

Answer: A

Question: 4

A company has an application that sells tickets online and experiences bursts of demand every 7 days. The application has a stateless presentation layer running on Amazon EC2, an Oracle database to store unstructured data catalog information, and a backend API layer. The front-end layer uses an Elastic Load Balancer to distribute the load across nine On-Demand Instances over three Availability Zones (AZs). The Oracle database is running on a single EC2 instance. The company is experiencing performance issues when running more than two concurrent campaigns. A solutions architect must design a solution that meets the following requirements:

- Address scalability issues.
- Increase the level of concurrency.
- Eliminate licensing costs.
- Improve reliability.

Which set of steps should the solutions architect take?

- A. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce costs. Convert the Oracle database into a single Amazon RDS reserved DB instance.
- B. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce costs. Create two additional copies of the database instance, then distribute the databases in separate AZs.
- C. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce costs. Convert the tables in the Oracle database into Amazon DynamoDB tables.
- D. Convert the On-Demand Instances into Spot Instances to reduce costs for the front end. Convert the tables in the Oracle database into Amazon DynamoDB tables.

Answer: C

Question: 5

A company wants to retire its Oracle Solaris NFS storage arrays. The company requires rapid data migration over its internet network connection to a combination of destinations for Amazon S3, Amazon Elastic File System (Amazon EFS), and Amazon FSx for Windows File Server. The company also requires a full initial copy, as well as incremental transfers of changes until the retirement of the storage arrays. All data must be encrypted and checked for integrity.

What should a solutions architect recommend to meet these requirements?

- A. Configure CloudEndure. Create a project and deploy the CloudEndure agent and token to the storage array. Run the migration plan to start the transfer.
- B. Configure AWS DataSync. Configure the DataSync agent and deploy it to the local network. Create a transfer task and start the transfer.
- C. Configure the aws S3 sync command. Configure the AWS client on the client side with credentials. Run the sync command to start the transfer.
- D. Configure AWS Transfer (or FTP). Configure the FTP client with credentials. Script the client to connect and sync to start the transfer.

Answer: B