SAP-C01 Q&As

AWS Certified Solutions Architect - Professional (SAP-C01)

Pass Amazon SAP-C01 Exam with 100% Guarantee

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

https://www.passapply.com/aws-solution-architect-professional.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 runs a highly available data collection application on Amazon EC2 in the eu- north-1 Region. The application collects data from end-user devices and writes records to an Amazon Kinesis data stream and a set of AWS Lambda functions that process the records The company persists the output of the record processing to an Amazon S3 bucket in eu-north-1. The company uses the data in the S3 bucket as a data source for Amazon Athena

A. In each of the Iwo new Regions set up the Lambda functions to run in a VPC Set up an S3 gateway endpoint in that VPC

B. Turn on S3 Transfer Acceleration on the S3 bucket in eu-north-1 Change the application to use the new S3 accelerated endpoint when the application uploads data to the S3 bucket

C. Create an S3 bucket in each of the two new Regions Set the application in each new Region to upload to its respective S3 bucket Set up S3 Cross-Region Replication to replicate data to the S3 bucket in eunorth-1

D. Increase the memory requirements of the Lambda functions to ensure that they have multiple cores available Use the multipart upload feature when the application uploads data to Amazon S3 Lambda

Correct Answer: A

QUESTION 2

A company is hosting a three-tier web application in an on-premises environment. Due to a recent surge in traffic that resulted in downtime and a significant financial impact, company management has ordered that the application be moved to AWS. The application is written in .NET and has a dependency on a MySQL database. A solutions architect must design a scalable and highly available solution to meet the demand of 200,000 daily users.

Which steps should the solutions architect take to design an appropriate solution?

A. Use AWS Elastic Beanstalk to create a new application with a web server environment and an Amazon RDS MySQL Multi-AZ DB instance. The environment should launch a Network Load Balancer (NLB) in front of an Amazon EC2 Auto Scaling group in multiple Availability Zones. Use an Amazon Route 53 alias record to route traffic from the company\\'s domain to the NLB.

B. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones. The stack should launch a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy. Use an Amazon Route 53 alias record to route traffic from the company\\'s domain to the ALB.

C. Use AWS Elastic Beanstalk to create an automatically scaling web server environment that spans two separate Regions with an Application Load Balancer (ALB) in each Region. Create a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a cross-Region read replica. Use Amazon Route 53 with a geoproximity routing policy to route traffic between the two Regions.

D. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon ECS cluster of Spot instances spanning three Availability Zones. The stack should launch an Amazon RDS MySQL DB instance with a Snapshot deletion policy. Use an Amazon Route 53 alias record to route traffic from the company\\'s domain to the ALB.

Correct Answer: B

QUESTION 3

A retail company runs a business-critical web service on an Amazon Elastic Container Service (Amazon ECS) cluster that runs on Amazon EC2 instances. The web service receives POST requests from end users and writes data to a MySQL database that runs on a separate EC2 instance. The company needs to ensure that data loss does not occur.

The current code deployment process includes manual updates of the ECS service. During a recent deployment, end users encountered intermittent 502 Bad Gateway errors in response to valid web requests.

The company wants to implement a reliable solution to prevent this issue from recurring. The company also wants to automate code deployments. The solution must be highly available and must optimize cost-effectiveness.

Which combination of steps will meet these requirements? (Choose three.)

A. Run the web service on an ECS cluster that has a Fargate launch type. Use AWS CodePipeline and AWS CodeDeploy to perform a blue/green deployment with validation testing to update the ECS service.

B. Migrate the MySQL database to run on an Amazon RDS for MySQL Multi-AZ DB instance that uses Provisioned IOPS SSD (io2) storage.

C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an event source to receive the POST requests from the web service. Configure an AWS Lambda function to poll the queue. Write the

data to the database.

D. Run the web service on an ECS cluster that has a Fargate launch type. Use AWS CodePipeline and AWS CodeDeploy to perform a canary deployment to update the ECS service.

E. Configure an Amazon Simple Queue Service (Amazon SQS) queue. Install the SQS agent on the containers that run in the ECS cluster to poll the queue. Write the data to the database.

F. Migrate the MySQL database to run on an Amazon RDS for MySQL Multi-AZ DB instance that uses General Purpose SSD (gp3) storage.

Correct Answer: BCD

QUESTION 4

A company is running an application in the AWS Cloud. The company has several third- party services that integrate with the application through a RESTful API. The API is a serverless implementation with an Amazon API Gateway regional API endpoint that integrates with several different AWS Lambda functions.

The application\\'s data is nonrelational and is stored in an Amazon DynamoDB table. The application and the API are running in the eu-west-1 Region. The company needs the API to also be available in the useast-1 Region. All data must be available in both Regions. A solutions architect already has deployed all the Lambda functions in us-east-1

Which additional steps should the solutions architect take to meet these requirements? (Select TWO.)

A. Deploy a second API Gateway regional API endpoint in us-east-1. Create Lambda integration with the functions in us-east-1.

B. Enable DynamoDB Streams on the table in eu-west-1. Replicate all changes to a DynamoDB table in us-east-1

C. Modify the DynamoDB table to be a global table in eu-west-1 and in us-east-1.

D. Change the API Gateway API endpoint in eu-west-1 to an edge-optimized endpoint.Create Lambda integration with the functions in both Regions.

E. Create a DynamoDB read replica in us-east-1.

Correct Answer: AC

QUESTION 5

A company has a Microsoft SQL Server database in its data center and plans to migrate data to Amazon Aurora MySQL. The company has already used the AWS Schema Conversion Tool to migrate triggers, stored procedures and other schema objects to Aurora MySQL. The database contains 1 TB of data and grows less than 1 MB per day. The company\\'s data center is connected to AWS through a dedicated 1Gbps AWS Direct Connect connection.

The company would like to migrate data to Aurora MySQL and perform reconfigurations with minimal downtime to the applications.

Which solution meets the company\\'s requirements?

A. Shut down applications over the weekend. Create an AWS DMS replication instance and task to migrate existing data from SQL Server to Aurora MySQL. Perform application testing and migrate the data to the new database endpoint.

B. Create an AWS DMS replication instance and task to migrate existing data and ongoing replication from SQL Server to Aurora MySQL. Perform application testing and migrate the data to the new database endpoint.

C. Create a database snapshot of SQL Server on Amazon S3. Restore the database snapshot from Amazon S3 to Aurora MySQL. Create an AWS DMS replication instance and task for ongoing replication from SQL Server to Aurora MySQL. Perform application testing and migrate the data to the new database endpoint.

D. Create a SQL Server native backup file on Amazon S3. Create an AWS DMS replication instance and task to restore the SQL Server backup file to Aurora MySQL. Create another AWS DMS task for ongoing replication from SQL Server to Aurora MySQL. Perform application testing and migrate the data to the new database endpoint.

Correct Answer: B