Free Questions for SAP-C02

Shared by Boyle on 04-10-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

An events company runs a ticketing platform on AWS. The company's customers configure and schedule their events on the platform The events result in large increases of traffic to the platform The company knows the date and time of each customer's events

The company runs the platform on an Amazon Elastic Container Service (Amazon ECS) cluster The ECS cluster consists of Amazon EC2 On-Demand Instances that are in an Auto Scaling group. The Auto Scaling group uses a predictive scaling policy

The ECS cluster makes frequent requests to an Amazon S3 bucket to download ticket assets The ECS cluster and the S3 bucket are in the same AWS Region and the same AWS account Traffic between the ECS cluster and the S3 bucket flows across a NAT gateway

The company needs to optimize the cost of the platform without decreasing the platform's availability

Which combination of steps will meet these requirements? (Select TWO)

Options:

A- Create a gateway VPC endpoint for the S3 bucket

B- Add another ECS capacity provider that uses an Auto Scaling group of Spot Instances Configure the new capacity provider strategy to have the same weight as the existing capacity provider strategy

C- Create On-Demand Capacity Reservations for the applicable instance type for the time period of the scheduled scaling policies

D- Enable S3 Transfer Acceleration on the S3 bucket

E- Replace the predictive scaling policy with scheduled scaling policies for the scheduled events

Answer:

A, B

Explanation:

Gateway VPC Endpoint for S3:

Create a gateway VPC endpoint for Amazon S3 in your VPC. This allows instances in your VPC to communicate with Amazon S3 without going through the internet, reducing data transfer costs and improving security.

Add Spot Instances to ECS Cluster:

Add another ECS capacity provider that uses an Auto Scaling group of Spot Instances. Configure this new capacity provider to share the load with the existing On-Demand Instances by setting an appropriate weight in the capacity provider strategy. Spot Instances offer significant cost savings compared to On-Demand Instances.

Configure Capacity Provider Strategy:

Adjust the ECS service's capacity provider strategy to utilize both On-Demand and Spot Instances effectively. This ensures a balanced distribution of tasks across both instance types, optimizing cost while maintaining availability.

By implementing a gateway VPC endpoint for S3 and incorporating Spot Instances into the ECS cluster, the company can significantly reduce operational costs without compromising on the availability or performance of the platform.

Reference

AWS Cost Optimization Blog on VPC Endpoints

AWS ECS Documentation on Capacity Providers

Question 2

Question Type: MultipleChoice

A company that develops consumer electronics with offices in Europe and Asia has 60 TB of software images stored on premises in Europe The company wants to transfer the images to an Amazon S3 bucket in the ap-northeast-1 Region New software images are created daily and must be encrypted in transit The company needs a solution that does not require custom development to automatically transfer all existing and new software images to Amazon S3

What is the next step in the transfer process?

Options:

A- Deploy an AWS DataSync agent and configure a task to transfer the images to the S3 bucket

B- Configure Amazon Kinesis Data Firehose to transfer the images using S3 Transfer Acceleration

C- Use an AWS Snowball device to transfer the images with the S3 bucket as the target

D- Transfer the images over a Site-to-Site VPN connection using the S3 API with multipart upload

Answer:

A

Explanation:

Deploy AWS DataSync Agent:

Install the DataSync agent on your on-premises environment. This can be done by downloading the agent as a virtual appliance and deploying it on VMware ESXi, Hyper-V, or KVM hypervisors.

Configure Source and Destination Locations:

Set up the source location pointing to your on-premises storage where the software images are currently stored.

Configure the destination location to point to your Amazon S3 bucket in the ap-northeast-1 Region.

Create and Schedule DataSync Tasks:

Create a DataSync task to automate the transfer process. This task will specify the source and destination locations and set options for how the data should be transferred.

Schedule the task to run at intervals that suit your data transfer requirements, ensuring new images are transferred as they are created.

Encryption in Transit:

AWS DataSync automatically encrypts data in transit using TLS, ensuring that your data is secure during the transfer process.

Monitoring and Management:

Use the DataSync console or the AWS CLI to monitor the progress of your data transfers and manage the tasks.

AWS DataSync is an efficient solution that automates and accelerates the process of transferring large amounts of data to AWS, handling encryption, data integrity checks, and optimizing network usage without requiring custom development.

Reference

AWS Storage Blog on DataSync40.

AWS DataSync Documentation41.

Question 3

Question Type: **MultipleChoice**

A software as a service (SaaS) company has developed a multi-tenant environment. The company uses Amazon DynamoDB tables that the tenants share tor the storage layer. The company uses AWS Lambda functions for the application services.

The company wants to offer a tiered subscription model that is based on resource consumption by each tenant Each tenant is identified by a unique tenant ID that is sent as part of each request to the Lambda functions The company has created an AWS Cost and Usage Report (AWS CUR) in an AWS account The company wants to allocate the DynamoDB costs to each tenant to match that tenant"s resource consumption

Which solution will provide a granular view of the DynamoDB cost for each tenant with the LEAST operational effort?

Options:

A- Associate a new lag that is named tenant ID with each table in DynamoDB Activate the tag as a cost allocation tag m the AWS Billing and Cost Management console Deploy new Lambda function code to log the tenant ID in Amazon CloudWatch Logs Use the AWS CUR to separate DynamoDB consumption cost for each tenant ID

B- Configure the Lambda functions to log the tenant ID and the number of RCUs and WCUs consumed from DynamoDB for each transaction to Amazon CloudWatch Logs Deploy another Lambda function to calculate the tenant costs by using the logged capacity units and the overall DynamoDB cost from the AWS Cost Explorer API Create an Amazon EventBridge rule to invoke the calculation Lambda function on a schedule.

C- Create a new partition key that associates DynamoDB items with individual tenants Deploy a Lambda function to populate the new column as part of each transaction Deploy another Lambda function to calculate the tenant costs by using Amazon Athena to calculate the number of tenant items from DynamoDB and the overall DynamoDB cost from the AWS CUR Create an Amazon EventBridge rule to invoke the calculation Lambda function on a schedule

D- Deploy a Lambda function to log the tenant ID the size of each response, and the duration of the transaction call as custom metrics to Amazon CloudWatch Logs Use CloudWatch Logs Insights to query the custom metrics for each tenant. Use AWS Pricing Calculator to obtain the overall DynamoDB costs and to calculate the tenant costs

Answer:

B

Explanation:

Log Tenant ID and RCUs/WCUs:

Update the AWS Lambda functions to log the tenant ID and the number of Read Capacity Units (RCUs) and Write Capacity Units (WCUs) consumed from DynamoDB for each transaction. This data will be logged to Amazon CloudWatch Logs.

Calculate Tenant Costs:

Deploy an additional Lambda function that reads the logs from CloudWatch Logs, calculates the RCUs and WCUs used by each tenant, and then uses the AWS Cost Explorer API to retrieve the overall cost of DynamoDB usage. This function will then allocate the costs to each tenant based on their usage.

Scheduled Cost Calculation:

Create an Amazon EventBridge rule to trigger the cost calculation Lambda function at regular intervals (e.g., daily or hourly). This ensures that cost allocation is continuously updated and tenants are billed accurately based on their consumption.

This solution minimizes operational effort by automating the cost allocation process and ensuring that the company can accurately bill tenants based on their resource consumption.

Reference

AWS Cost Explorer Documentation

Amazon CloudWatch Logs Documentation

AWS Lambda Documentation

Question 4

Question Type: **MultipleChoice**

A startup company recently migrated a large ecommerce website to AWS The website has experienced a 70% increase in sates Software engineers are using a private GitHub repository to manage code The DevOps team is using Jenkins for builds and unit testing The engineers need to receive notifications for bad builds and zero downtime during deployments The engineers also need to ensure any changes to production are seamless for users and can be rolled back in the event of a major issue

The software engineers have decided to use AWS CodePipeline to manage their build and deployment process

Which solution will meet these requirements'?

Options:

A- Use GitHub websockets to trigger the CodePipeline pipeline Use the Jenkins plugin for AWS CodeBuild to conduct unit testing Send alerts to an Amazon SNS topic for any bad builds Deploy in an in-place all-at-once deployment configuration using AWS CodeDeploy

B- Use GitHub webhooks to trigger the CodePipelme pipeline Use the Jenkins plugin for AWS CodeBuild to conduct unit testing Send alerts to an Amazon SNS topic for any bad builds Deploy in a blue'green deployment using AWS CodeDeploy

C- Use GitHub websockets to trigger the CodePipelme pipeline. Use AWS X-Ray for unit testing and static code analysis Send alerts to an Amazon SNS topic for any bad builds Deploy in a blue/green deployment using AWS CodeDeploy.

D- Use GitHub webhooks to trigger the CodePipeline pipeline Use AWS X-Ray for unit testing and static code analysis Send alerts to an Amazon SNS topic for any bad builds Deploy in an m-place. all-at-once deployment configuration using AWS CodeDeploy

Answer:

B

Explanation:

GitHub Webhooks to Trigger CodePipeline:

Configure GitHub webhooks to trigger the AWS CodePipeline pipeline. This ensures that every code push to the repository automatically triggers the pipeline, initiating the build and deployment process.

Unit Testing with Jenkins and AWS CodeBuild:

Use Jenkins integrated with the AWS CodeBuild plugin to perform unit testing. Jenkins will manage the build process, and the results will be handled by CodeBuild.

Notifications for Bad Builds:

Configure Amazon SNS (Simple Notification Service) to send alerts for any failed builds. This keeps the engineering team informed of build issues immediately, allowing for quick resolutions.

Blue/Green Deployment with AWS CodeDeploy:

Utilize AWS CodeDeploy with a blue/green deployment strategy. This method reduces downtime and risk by running two identical production environments (blue and green). CodeDeploy shifts traffic between these environments, allowing you to test in the new environment (green) while the old environment (blue) remains live. If issues arise, you can quickly roll back to the previous environment.

This solution provides seamless, zero-downtime deployments, and the ability to quickly roll back changes if necessary, fulfilling the requirements of the startup company.

Reference

AWS DevOps Blog on Integrating Jenkins with AWS CodeBuild and CodeDeploy32.

Plain English Guide to AWS CodePipeline with GitHub33.

Jenkins Plugin for AWS CodePipeline34.

Question 5

Question Type: MultipleChoice

A company wants to migrate virtual Microsoft workloads from an on-premises data center to AWS The company has successfully tested a few sample workloads on AWS. The company also has created an AWS Site-to-Site VPN connection to a VPC A solutions architect needs to generate a total cost of ownership (TCO) report for the migration of all the workloads from the data center

Simple Network Management Protocol (SNMP) has been enabled on each VM in the data center The company cannot add more VMs m the data center and cannot install additional software on the VMs The discovery data must be automatically imported into AWS Migration Hub

Which solution will meet these requirements?

Options:

A- Use the AWS Application Migration Service agentless service and the AWS Migration Hub Strategy Recommendations to generate the TCO report

B- Launch a Windows Amazon EC2 instance Install the Migration Evaluator agentless collector on the EC2 instance Configure Migration Evaluator to generate the TCO report

C- Launch a Windows Amazon EC2 instance. Install the Migration Evaluator agentless collector on the EC2 instance. Configure Migration Hub to generate the TCO report

D- Use the AWS Migration Readiness Assessment tool inside the VPC Configure Migration Evaluator to generate the TCO report

Answer:

A

Explanation:

AWS Application Migration Service:

AWS Application Migration Service (MGN) facilitates the migration of virtual machines (VMs) to AWS without installing additional software on the VMs. This agentless service helps in seamlessly migrating workloads to AWS.

AWS Migration Hub Strategy Recommendations:

AWS Migration Hub Strategy Recommendations offer insights and guidance for planning and implementing migration strategies. It helps in generating a Total Cost of Ownership (TCO) report by automatically importing discovery data from the VMs.

Generating the TCO Report:

The combined use of AWS Application Migration Service and Migration Hub Strategy Recommendations enables the automatic import of discovery data and the generation of an accurate TCO report, ensuring a smooth and cost-effective migration process.

Reference

AWS Migration Hub Strategy Recommendations (AWS Documentation).

Question 6

Question Type: MultipleChoice

A company runs an application in (he cloud that consists of a database and a website Users can post data to the website, have the data processed, and have the data sent back to them in an email Data is stored in a MySQL database running on an Amazon EC2 instance The database is running in a VPC with two private subnets The website is running on Apache Tomcat in a single EC2 instance in a different VPC with one public subnet There is a single VPC peering connection between the database and website VPC.

The website has suffered several outages during the last month due to high traffic

Which actions should a solutions architect take to increase the reliability of the application? (Select THREE.)

Options:

A- Place the Tomcat server in an Auto Scaling group with multiple EC2 instances behind an Application Load Balancer

B- Provision an additional VPC peering connection

C- Migrate the MySQL database to Amazon Aurora with one Aurora Replica

D- Provision two NAT gateways in the database VPC.

E- Move the Tomcat server to the database VPC

F- Create an additional public subnet in a different Availability Zone in the website VPC

Answer:

A, C, F

Explanation:

Auto Scaling Group with Application Load Balancer:

Moving the Tomcat server to an Auto Scaling group ensures that the number of instances adjusts dynamically based on the traffic load. An Application Load Balancer (ALB) distributes incoming traffic across multiple instances, improving the application's reliability and availability.

Migrate to Amazon Aurora with Replica:

Migrating the MySQL database to Amazon Aurora and adding an Aurora Replica enhances the database's scalability and availability. Aurora is optimized for performance, and replicas help distribute read traffic, reducing the load on the primary instance.

Additional Public Subnet:

Creating an additional public subnet in a different Availability Zone enhances fault tolerance. This ensures that the website remains accessible even if one Availability Zone experiences issues.

Reference

AWS Well-Architected Framework

Amazon Aurora Documentation

Question 7

Question Type: MultipleChoice

A company runs a software-as-a-service
Which solution meets these requirements'?

Options:

A- Create an Amazon CloudWatch alarm action that triggers a Lambda function to add an Amazon RDS for MySQL read replica when resource utilization hits a threshold

B- Migrate the database to Amazon Aurora, and add a read replica Add a database connection pool outside of the Lambda handler function

C- Migrate the database to Amazon Aurora and add a read replica Use Amazon Route 53 weighted records

D- Migrate the database to Amazon Aurora and add an Aurora Replica Configure Amazon RDS Proxy to manage database connection pools

Answer:

D

Explanation:

Migrate to Amazon Aurora:

Amazon Aurora is a MySQL-compatible, high-performance database designed to provide higher throughput than standard MySQL. Migrating the database to Aurora will enhance the performance and scalability of the database, especially under heavy workloads.

Add Aurora Replica:

Aurora Replicas provide read scalability and improve availability. Adding an Aurora Replica allows read operations to be distributed, thereby reducing the load on the primary instance and improving response times during peak periods.

Configure Amazon RDS Proxy:

Amazon RDS Proxy acts as an intermediary between the application and the Aurora database, managing connection pools efficiently. RDS Proxy reduces the overhead of opening and closing database connections, thus maintaining fewer active connections to the database and handling surges in database connections from the Lambda functions more effectively.

This configuration reduces the database's resource usage and improves its ability to handle high volumes of concurrent connections.

Reference

AWS Database Blog on RDS Proxy (Amazon Web Services, Inc.).

AWS Compute Blog on Using RDS Proxy with Lambda (Amazon Web Services, Inc.).