



Free Questions for SAA-C03 by dumpshq

Shared by Lawson on 23-09-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company is running several business applications in three separate VPCs within the us-east-1 Region. The applications must be able to communicate between VPCs. The applications also must be able to consistently send hundreds to gigabytes of data each day to a latency-sensitive application that runs in a single on-premises data center.

A solutions architect needs to design a network connectivity solution that maximizes cost-effectiveness

Which solution meets those requirements?

Options:

- A-** Configure three AWS Site-to-Site VPN connections from the data center to AWS. Establish connectivity by configuring one VPN connection for each VPC.
- B-** Launch a third-party virtual network appliance in each VPC. Establish an IPsec VPN tunnel between the Data center and each virtual appliance.
- C-** Set up three AWS Direct Connect connections from the data center to a Direct Connect gateway in us-east-1. Establish connectivity by configuring each VPC to use one of the Direct Connect connections.
- D-** Set up one AWS Direct Connect connection from the data center to AWS. Create a transit gateway, and attach each VPC to the transit gateway. Establish connectivity between the Direct Connect connection and the transit gateway.

Answer:

D

Explanation:

<https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-aws-transit-gateway.html>

Question 2

Question Type: MultipleChoice

A company wants to measure the effectiveness of its recent marketing campaigns. The company performs batch processing on csv files of sales data and stores the results i an Amazon S3 bucket once every hour. The S3 bi petabytes of objects. The company runs one-time queries in Amazon Athena to determine which products are most popular on a particular date for a particular region Queries sometimes fail or take longer than expected to finish.

Which actions should a solutions architect take to improve the query performance and reliability? (Select TWO.)

Options:

- A- Reduce the S3 object sizes to less than 126 MB
- B- Partition the data by date and region n Amazon S3
- C- Store the files as large, single objects in Amazon S3.
- D- Use Amazon Kinesis Data Analytics to run the Queries as pan of the batch processing operation
- E- Use an AWS duo extract, transform, and load (ETL) process to convert the csv files into Apache Parquet format.

Answer:

C, E

Question 3

Question Type: MultipleChoice

A company runs a global web application on Amazon EC2 instances behind an Application Load Balancer The application stores data in Amazon Auror

a. The company needs to create a disaster recovery solution and can tolerate up to 30 minutes of downtime and potential data loss. The solution does not need to handle the load when the primary infrastructure is healthy

What should a solutions architect do to meet these requirements?

Options:

- A-** Deploy the application with the required infrastructure elements in place Use Amazon Route 53 to configure active-passive failover
Create an Aurora Replica in a second AWS Region
- B-** Host a scaled-down deployment of the application in a second AWS Region Use Amazon Route 53 to configure active-active failover
Create an Aurora Replica in the second Region
- C-** Replicate the primary infrastructure in a second AWS Region Use Amazon Route 53 to configure active-active failover Create an Aurora database that is restored from the latest snapshot
- D-** Back up data with AWS Backup Use the backup to create the required infrastructure in a second AWS Region Use Amazon Route 53 to configure active-passive failover Create an Aurora second primary instance in the second Region

Answer:

C

Question 4

Question Type: MultipleChoice

An ecommerce company has an order-processing application that uses Amazon API Gateway and an AWS Lambda function. The application stores data in an Amazon Aurora PostgreSQL database. During a recent sales event, a sudden surge in customer orders occurred. Some customers experienced timeouts and the application did not process the orders of those customers A solutions architect

determined that the CPU utilization and memory utilization were high on the database because of a large number of open connections. The solutions architect needs to prevent the timeout errors while making the least possible changes to the application.

Which solution will meet these requirements?

Options:

- A-** Configure provisioned concurrency for the Lambda function. Modify the database to be a global database in multiple AWS Regions.
- B-** Use Amazon RDS Proxy to create a proxy for the database. Modify the Lambda function to use the RDS Proxy endpoint instead of the database endpoint.
- C-** Create a read replica for the database in a different AWS Region. Use query string parameters in API Gateway to route traffic to the read replica.
- D-** Migrate the data from Aurora PostgreSQL to Amazon DynamoDB by using AWS Database Migration Service (AWS DMS). Modify the Lambda function to use the DynamoDB table.

Answer:

D

Question 5

Question Type: MultipleChoice

A company's application is having performance issues. The application is stateful and needs to complete memory-intensive tasks on Amazon EC2 instances. The company used AWS CloudFormation to deploy infrastructure and used the M5 EC2 Instance family. As traffic increased, the application performance degraded. Users are reporting delays when they attempt to access the application.

Which solution will resolve these issues in the MOST operationally efficient way?

Options:

- A-** Replace the EC2 instances with T3 EC2 instances that run in an Auto Scaling group. Make the changes by using the AWS Management Console.
- B-** Modify the CloudFormation templates to run the EC2 instances in an Auto Scaling group. Increase the desired capacity and the maximum capacity of the Auto Scaling group manually when an increase is necessary.
- C-** Modify the CloudFormation templates. Replace the EC2 instances with R5 EC2 instances. Use Amazon CloudWatch built-in EC2 memory metrics to track the application performance for future capacity planning.
- D-** Modify the CloudFormation templates. Replace the EC2 instances with R5 EC2 instances. Deploy the Amazon CloudWatch agent on the EC2 instances to generate custom application latency metrics for future capacity planning.

Answer:

D

Question 6

Question Type: MultipleChoice

A company is building a containerized application on premises and decides to move the application to AWS. The application will have thousands of users soon after it is deployed. The company is unsure how to manage the deployment of containers at scale. The company needs to deploy the containerized application in a highly available architecture that minimizes operational overhead.

Which solution will meet these requirements?

Options:

- A-** Store container images in an Amazon Elastic Container Registry (Amazon ECR) repository. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the AWS Fargate launch type to run the containers. Use target tracking to scale automatically based on demand.
- B-** Store container images in an Amazon Elastic Container Registry (Amazon ECR) repository. Use an Amazon Elastic Container Service (Amazon ECS) cluster with the Amazon EC2 launch type to run the containers. Use target tracking to scale automatically based on demand.
- C-** Store container images in a repository that runs on an Amazon EC2 instance. Run the containers on EC2 instances that are spread across multiple Availability Zones. Monitor the average CPU utilization in Amazon CloudWatch. Launch new EC2 instances as needed.
- D-** Create an Amazon EC2 Amazon Machine Image (AMI) that contains the container image. Launch EC2 instances in an Auto Scaling group across multiple Availability Zones. Use an Amazon CloudWatch alarm to scale out EC2 instances when the average CPU utilization threshold is breached.

Answer:

A

Question 7

Question Type: MultipleChoice

A company has an event-driven application that invokes AWS Lambda functions up to 800 times each minute with varying runtimes. The Lambda functions access data that is stored in an Amazon Aurora MySQL OB cluster. The company is noticing connection timeouts as user activity increases. The database shows no signs of being overloaded. CPU, memory, and disk access metrics are all low.

Which solution will resolve this issue with the LEAST operational overhead?

Options:

- A-** Adjust the size of the Aurora MySQL nodes to handle more connections. Configure retry logic in the Lambda functions for attempts to connect to the database.
- B-** Set up Amazon ElastiCache for Redis to cache commonly read items from the database. Configure the Lambda functions to connect to ElastiCache for reads.
- C-** Add an Aurora Replica as a reader node. Configure the Lambda functions to connect to the reader endpoint of the OB cluster rather than to the writer endpoint.

D- Use Amazon ROS Proxy to create a proxy. Set the DB cluster as the target database. Configure the Lambda functions to connect to the proxy rather than to the DB cluster.

Answer:

D

Question 8

Question Type: MultipleChoice

A company is planning to build a high performance computing (HPC) workload as a service solution that is hosted on AWS. A group of 16 Amazon EC2 Linux Instances requires the lowest possible latency for node-to-node communication. The instances also need a shared block device volume for high-performing storage.

Which solution will meet these requirements?

Options:

A- Use a cluster placement group. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon EBS) volume to all the instances by using Amazon EBS Multi-Attach.

- B-** Use a cluster placement group. Create shared file systems across the instances by using Amazon Elastic File System (Amazon EFS)
- C-** Use a partition placement group. Create shared tile systems across the instances by using Amazon Elastic File System (Amazon EFS).
- D-** Use a spread placement group. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon EBS) volume to all the instances by using Amazon EBS Multi-Attach

Answer:

A

Question 9

Question Type: MultipleChoice

A company has an ecommerce checkout workflow that writes an order to a database and calls a service to process the payment. Users are experiencing timeouts during the checkout process. When users resubmit the checkout form, multiple unique orders are created for the same desired transaction.

How should a solutions architect refactor this workflow to prevent the creation of multiple orders?

Options:

- A-** Configure the web application to send an order message to Amazon Kinesis Data Firehose. Set the payment service to retrieve the message from Kinesis Data Firehose and process the order.
- B-** Create a rule in AWS CloudTrail to invoke an AWS Lambda function based on the logged application path request Use Lambda to query the database, call the payment service, and pass in the order information.
- C-** Store the order in the database. Send a message that includes the order number to Amazon Simple Notification Service (Amazon SNS). Set the payment service to poll Amazon SNS. retrieve the message, and process the order.
- D-** Store the order in the database. Send a message that includes the order number to an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Set the payment service to retrieve the message and process the order. Delete the message from the queue.

Answer:

D

Question 10

Question Type: MultipleChoice

A company hosts a two-tier application on Amazon EC2 instances and Amazon RDS. The application's demand varies based on the time of day. The load is minimal after work hours and on weekends. The EC2 instances run in an EC2 Auto Scaling group that is configured with a minimum of two instances and a maximum of five instances. The application must be available at all times, but the company is concerned about overall cost.

Which solution meets the availability requirement MOST cost-effectively?

Options:

- A-** Use all EC2 Spot Instances. Stop the RDS database when it is not in use.
- B-** Purchase EC2 Instance Savings Plans to cover five EC2 instances. Purchase an RDS Reserved DB Instance
- C-** Purchase two EC2 Reserved Instances Use up to three additional EC2 Spot Instances as needed. Stop the RDS database when it is not in use.
- D-** Purchase EC2 Instance Savings Plans to cover two EC2 instances. Use up to three additional EC2 On-Demand Instances as needed. Purchase an RDS Reserved DB Instance.

Answer:

D

To Get Premium Files for SAA-C03 Visit

<https://www.p2pexams.com/products/saa-c03>

For More Free Questions Visit

<https://www.p2pexams.com/amazon/pdf/saa-c03>

