

Amazon

SOA-C02 Exam

AWS Certified SysOps Administrator - Associate

**Questions & Answers
Demo**

Version: 18.0

Topic 1, Mix Questions

Question: 1

A Sysops administrator creates an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that uses AWS Fargate. The cluster is deployed successfully. The Sysops administrator needs to manage the cluster by using the kubectl command line tool.

Which of the following must be configured on the Sysops administrator's machine so that kubectl can communicate with the cluster API server?

- A. The kubeconfig file
- B. The kube-proxy Amazon EKS add-on
- C. The Fargate profile
- D. The eks-connector.yaml file

Answer: A

Explanation:

The kubeconfig file is a configuration file used to store cluster authentication information, which is required to make requests to the Amazon EKS cluster API server. The kubeconfig file will need to be configured on the SysOps administrator's machine in order for kubectl to be able to communicate with the cluster API server.

<https://aws.amazon.com/blogs/developer/running-a-kubernetes-job-in-amazon-eks-on-aws-fargate-using-aws-stepfunctions/>

Question: 2

A Sysops administrator needs to configure automatic rotation for Amazon RDS database credentials. The credentials must rotate every 30 days. The solution must integrate with Amazon RDS.

Which solution will meet these requirements with the LEAST operational overhead?

A. Store the credentials in AWS Systems Manager Parameter Store as a secure string. Configure automatic rotation with a rotation interval of 30 days.

B. Store the credentials in AWS Secrets Manager. Configure automatic rotation with a rotation interval of 30 days.

C. Store the credentials in a file in an Amazon S3 bucket. Deploy an AWS Lambda function to automatically rotate the credentials every 30 days.

D. Store the credentials in AWS Secrets Manager. Deploy an AWS Lambda function to automatically

rotate the credentials every 30 days.

Answer: B

Explanation:

Storing the credentials in AWS Secrets Manager and configuring automatic rotation with a rotation interval of 30 days is the most efficient way to meet the requirements with the least operational overhead. AWS Secrets Manager automatically rotates the credentials at the specified interval, so there is no need for an additional AWS Lambda function or manual rotation. Additionally, Secrets Manager is integrated with Amazon RDS, so the credentials can be easily used with the RDS database.

Question: 3

A company has an application that runs only on Amazon EC2 Spot Instances. The instances run in an Amazon EC2 Auto Scaling group with scheduled scaling actions.

However, the capacity does not always increase at the scheduled times, and instances terminate many times a day. A Sysops administrator must ensure that the instances launch on time and have fewer interruptions.

Which action will meet these requirements?

A. Specify the capacity-optimized allocation strategy for Spot Instances. Add more instance types to the Auto Scaling group.

B. Specify the capacity-optimized allocation strategy for Spot Instances. Increase the size of the instances in the Auto Scaling group.

C. Specify the lowest-price allocation strategy for Spot Instances. Add more instance types to the Auto Scaling group.

D. Specify the lowest-price allocation strategy for Spot Instances. Increase the size of the instances in the Auto Scaling group.

Answer: A

Explanation:

Specifying the capacity-optimized allocation strategy for Spot Instances and adding more instance types to the Auto Scaling group is the best action to meet the requirements. Increasing the size of the instances in the Auto Scaling group will not necessarily help with the launch time or reduce interruptions, as the Spot Instances could still be interrupted even with larger instance sizes.

Question: 4

A company stores its data in an Amazon S3 bucket. The company is required to classify the data and find any sensitive personal information in its S3 files.

Which solution will meet these requirements?

A. Create an AWS Config rule to discover sensitive personal information in the S3 files and mark them as noncompliant.

B. Create an S3 event-driven artificial intelligence/machine learning (AI/ML) pipeline to classify sensitive personal information by using Amazon Recognition.

C. Enable Amazon GuardDuty. Configure S3 protection to monitor all data inside Amazon S3.

D. Enable Amazon Macie. Create a discovery job that uses the managed data identifier.

Answer: D

Explanation:

Amazon Macie is a security service designed to help organizations find, classify, and protect sensitive data stored in Amazon S3. Amazon Macie uses machine learning to automatically discover, classify, and protect sensitive data in Amazon S3. Creating a discovery job with the managed data identifier will allow Macie to identify sensitive personal information in the S3 files and classify it accordingly. Enabling AWS Config and Amazon GuardDuty will not help with this requirement as they are not designed to automatically classify and protect data.

Question: 5

A company has an application that customers use to search for records on a website. The application's data is stored in an Amazon Aurora DB cluster. The application's usage varies by season and by day of the week.

The website's popularity is increasing, and the website is experiencing slower performance because of increased load on the DB cluster during periods of peak activity. The application logs show that the performance issues occur when users are searching for information. The same search is rarely performed multiple times.

A SysOps administrator must improve the performance of the platform by using a solution that maximizes resource efficiency.

Which solution will meet these requirements?

A. Deploy an Amazon ElastiCache for Redis cluster in front of the DB cluster. Modify the application to check the cache before the application issues new queries to the database. Add the results of any queries to the cache.

B. Deploy an Aurora Replica for the DB cluster. Modify the application to use the reader endpoint for search operations. Use Aurora Auto Scaling to scale the number of replicas based on load. Most Voted

C. Use Provisioned IOPS on the storage volumes that support the DB cluster to improve performance sufficiently to support the peak load on the application.

D. Increase the instance size in the DB cluster to a size that is sufficient to support the peak load on the application. Use Aurora Auto Scaling to scale the instance size based on load.

Answer: A

Explanation:

Step-by-Step

Understand the Problem:

The application experiences slower performance during peak activity due to increased load on the Amazon Aurora DB cluster.

Performance issues occur primarily during search operations.

The goal is to improve performance and maximize resource efficiency.

Analyze the Requirements:

The solution should improve the performance of the platform.

It should maximize resource efficiency, which implies cost-effective and scalable options.

Evaluate the Options:

Option A: Deploy an Amazon ElastiCache for Redis cluster.

ElastiCache for Redis is a managed in-memory caching service that can significantly reduce the load on the database by caching frequently accessed data.

By modifying the application to check the cache before querying the database, repeated searches for the same information will be served from the cache, reducing the number of database reads.

This is efficient and cost-effective as it reduces database load and improves response times.

Option B: Deploy an Aurora Replica and use Auto Scaling.

Adding Aurora Replicas can help distribute read traffic and improve performance.

Aurora Auto Scaling can adjust the number of replicas based on the load.

However, this option may not be as efficient in terms of resource usage compared to caching because it still involves querying the database.

Option C: Use Provisioned IOPS.

Provisioned IOPS can improve performance by providing fast and consistent I/O.

This option focuses on improving the underlying storage performance but doesn't address the inefficiency of handling repeated searches directly.

Option D: Increase the instance size and use Auto Scaling.

Increasing the instance size can provide more resources to handle peak loads.

Aurora Auto Scaling can adjust instance sizes based on the load.

This option can be costly and may not be as efficient as caching in handling repeated searches.

Select the Best Solution:

Option A is the best solution because it leverages caching to reduce the load on the database, which directly addresses the issue of repeated searches causing performance problems. Caching is generally more resource-efficient and cost-effective compared to scaling database instances or storage.

Reference:

[Amazon ElastiCache for Redis Documentation](#)

[Amazon Aurora Documentation](#)

[AWS Auto Scaling](#)

Using ElastiCache for Redis aligns with best practices for improving application performance by offloading repetitive read queries from the database, leading to faster response times and more efficient resource usage.