DUMPSARENA

AWS Certified SysOps Administrator - Associate (SOA-C02)

Amazon AWS SOA-C02

Version Demo

Total Demo Questions: 15

Total Premium Questions: 305

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Mix Questions	301
Topic 2, Simulation	4
Total	305

QUESTION NO: 1

A company has an Auto Scaling group of Amazon EC2 instances that scale based on average CPU utilization. The Auto Scaling group events log indicates an InsufficientInstanceCapacity error.

Which actions should a SysOps administrator take to remediate this issue? (Select TWO.

A. Change the instance type that the company is using.

B. Configure the Auto Scaling group in different Availability Zones.

C. Configure the Auto Scaling group to use different Amazon Elastic Block Store (Amazon EBS) volume sizes.

D. Increase the maximum size of the Auto Scaling group.

E. Request an increase in the instance service quota.

ANSWER: A B

QUESTION NO: 2

A SysOps administrator has an AWS CloudFormation template of the company's existing infrastructure in us-west-2. The administrator attempts to use the template to launch a new stack in eu-west-1, but the stack only partially deploys, receives an error message, and then rolls back.

Why would this template fail to deploy? (Select TWO.)

A. The template referenced an IAM user that is not available in eu-west-1.

B. The template referenced an Amazon Machine Image (AMI) that is not available in eu-west-1.

C. The template did not have the proper level of permissions to deploy the resources.

D. The template requested services that do not exist in eu-west-1.

E. CloudFormation templates can be used only to update existing services.

ANSWER: B D

QUESTION NO: 3

A company runs a stateless application that is hosted on an Amazon EC2 instance. Users are reporting performance issues. A SysOps administrator reviews the Amazon CloudWatch metrics for the application and notices that the instance's CPU utilization frequently reaches 90% during business hours.

What is the MOST operationally efficient solution that will improve the application's responsiveness?

A. Configure CloudWatch logging on the EC2 instance. Configure a CloudWatch alarm for CPU utilization to alert the SysOps administrator when CPU utilization goes above 90%.

B. Configure an AWS Client VPN connection to allow the application users to connect directly to the EC2 instance private IP address to reduce latency.

C. Create an Auto Scaling group, and assign it to an Application Load Balancer. Configure a target tracking scaling policy that is based on the average CPU utilization of the Auto Scaling group.

D. Create a CloudWatch alarm that activates when the EC2 instance's CPU utilization goes above 80%. Configure the alarm to invoke an AWS Lambda function that vertically scales the instance.

ANSWER: C

QUESTION NO: 4

A company's VPC has connectivity to an on-premises data center through an AWS Site-to-Site VPN. The company needs Amazon EC2 instances in the VPC to send DNS queries for example com to the DNS servers in the data center.

Which solution will meet these requirements?

A. Create an Amazon Route 53 Resolver inbound endpoint Create a conditional forwarding rule on the on-primes DNS servers to forward DNS requests for example.com to the inbound endpoints.

B. Create an Amazon Route 53 Resolver inbound endpoint Create a forwarding rule on the resolver that sends all queries for example.com to the on-premises DNS servers. Associate this rule with the VPC.

C. Create an Amazon Route 53 Resolver outbound endpoint Create a conditional forwarding rule on the on-premises DNS servers to forward DNS requests for example.com to the outbound endpoints

D. Create an Amazon Route 53 Resolver outbound endpoint. Create a forwarding rule on the resolver that sends all queries for exarrc4e.com to the on-premises DNS servers Associate this rule with the VPC.

ANSWER: C

QUESTION NO: 5

A company is running a serverless application on AWS Lambda The application stores data in an Amazon RDS for MySQL DB instance Usage has steadily increased and recently there have been numerous "too many connections" errors when the Lambda function attempts to connect to the database The company already has configured the database to use the maximum max_connections value that is possible

What should a SysOps administrator do to resolve these errors'?

A. Create a read replica of the database Use Amazon Route 53 to create a weighted DNS record that contains both databases

B. Use Amazon RDS Proxy to create a proxy Update the connection string in the Lambda function

C. Increase the value in the max_connect_errors parameter in the parameter group that the database uses

D. Update the Lambda function's reserved concurrency to a higher value

ANSWER: B

Explanation:

https://aws.amazon.com/blogs/compute/using-amazon-rds-proxy-with-aws-lambda/

RDS Proxy acts as an intermediary between your application and an RDS database. RDS Proxy establishes and manages the necessary connection pools to your database so that your application creates fewer database connections. Your Lambda functions interact with RDS Proxy instead of your database instance. It handles the connection pooling necessary for scaling many simultaneous connections created by concurrent Lambda functions. This allows your Lambda applications to reuse existing connections, rather than creating new connections for every function invocation.

Check "Database proxy for Amazon RDS" section in the link to see how RDS proxy help Lambda handle huge connections to RDS MySQL https://aws.amazon.com/blogs/compute/using-amazon-rds-proxy-with-aws-lambda/

QUESTION NO: 6 - (SIMULATION)

If your AWS Management Console browser does not show that you are logged in to an AWS account, close the browser and relaunch the

console by using the AWS Management Console shortcut from the VM desktop.

If the copy-paste functionality is not working in your environment, refer to the instructions file on the VM desktop and use Ctrl+C, Ctrl+V or Command-C , Command-V.

Configure Amazon EventBridge to meet the following requirements.

1. use the us-east-2 Region for all resources,

2. Unless specified below, use the default configuration settings.

3. Use your own resource naming unless a resource

name is specified below.

4. Ensure all Amazon EC2 events in the default event

bus are replayable for the past 90 days.

5. Create a rule named RunFunction to send the exact message every 1 5 minutes to an existing AWS Lambda function named LogEventFunction.

6. Create a rule named SpotWarning to send a notification to a new standard Amazon SNS topic named TopicEvents whenever an Amazon EC2 Spot Instance is interrupted. Do NOT create any topic subscriptions. The notification must match the following structure:

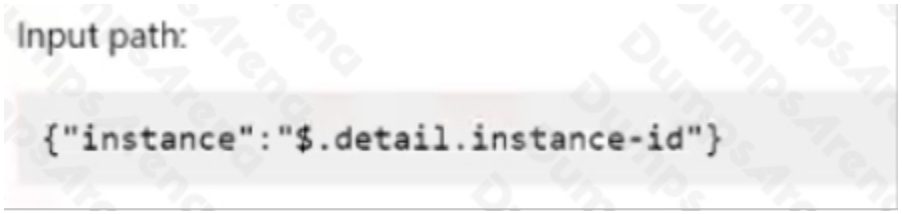

Input Path:

{“instance” : “$.detail.instance-id”}

Input template:

“ The EC2 Spot Instance has been on account.

ANSWER: Seetheforsolution.

Explanation:

Here are the steps to configure Amazon EventBridge to meet the above requirements:

Note:

QUESTION NO: 7

A large company is using AWS Organizations to manage hundreds of AWS accounts across multiple AWS Regions. The company has turned on AWS Config throughout the organization.

The company requires all Amazon S3 buckets to block public read access. A SysOps administrator must generate a monthly report that shows all the S3 buckets and whether they comply with this requirement.

Which combination of steps should the SysOps administrator take to collect this data? {Select TWO).

A. Create an AWS Config aggregator in an aggregator account. Use the organization as the source. Retrieve the compliance data from the aggregator.

B. Create an AWS Config aggregator in each account. Use an S3 bucket in an aggregator account as the destination. Retrieve the compliance data from the S3 bucket

C. Edit the AWS Config policy in AWS Organizations. Use the organization's management account to turn on the s3-bucket-public-read-prohibited rule for the entire organization.

D. Use the AWS Config compliance report from the organization's management account. Filter the results by resource, and select Amazon S3.

E. Use the AWS Config API to apply the s3-bucket-public-read-prohibited rule in all accounts for all available Regions.

ANSWER: C D

QUESTION NO: 8

A company uses an AWS CloudFormation template to provision an Amazon EC2 instance and an Amazon RDS DB instance A SysOps administrator must update the template to ensure that the DB instance is created before the EC2 instance is launched

What should the SysOps administrator do to meet this requirement?

A. Add a wait condition to the template Update the EC2 instance user data script to send a signal after the EC2 instance is started

B. Add the DependsOn attribute to the EC2 instance resource, and provide the logical name of the RDS resource

C. Change the order of the resources in the template so that the RDS resource is listed before the EC2 instance resource

D. Create multiple templates Use AWS CloudFormation StackSets to wait for one stack to complete before the second stack is created

ANSWER: B

Explanation:

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-dependson.html

Syntax The DependsOn attribute can take a single string or list of strings. "DependsOn" : [String, ...] Example The following template contains an AWS::EC2::Instance resource with a DependsOn attribute that specifies myDB, an AWS::RDS::DBInstance. When CloudFormation creates this stack, it first creates myDB, then creates Ec2Instance.

QUESTION NO: 9

A company has an existing web application that runs on two Amazon EC2 instances behind an Application Load Balancer (ALB) across two Availability Zones The application uses an Amazon RDS Multi-AZ DB Instance Amazon Route 53 record sets route requests tor dynamic content to the load balancer and requests for static content to an Amazon S3 bucket Site visitors are reporting extremely long loading times.

Which actions should be taken to improve the performance of the website? (Select TWO)

A. Add Amazon CloudFront caching for static content

B. Change the load balancer listener from HTTPS to TCP

C. Enable Amazon Route 53 latency-based routing

D. Implement Amazon EC2 Auto Scaling for the web servers

E. Move the static content from Amazon S3 to the web servers

ANSWER: A D

QUESTION NO: 10

A company is using Amazon Elastic Container Sen/ice (Amazon ECS) to run a containerized application on Amazon EC2 instances. A SysOps administrator needs to monitor only traffic flows between the ECS tasks.

Which combination of steps should the SysOps administrator take to meet this requirement? (Select TWO.)

A. Configure Amazon CloudWatch Logs on the elastic network interface of each task.

B. Configure VPC Flow Logs on the elastic network interface of each task.

C. Specify the awsvpc network mode in the task definition.

D. Specify the bridge network mode in the task definition.

E. Specify the host network mode in the task definition.

ANSWER: A E

QUESTION NO: 11

A company hosts several write-intensive applications. These applications use a MySQL database that runs on a single Amazon EC2 instance. The company asks a SysOps administrator to implement a highly available database solution that is ideal for multi-tenant workloads.

Which solution should the SysOps administrator implement to meet these requirements?

A. Create a second EC2 instance for MySQL. Configure the second instance to be a read replica.

B. Migrate the database to an Amazon Aurora DB cluster. Add an Aurora Replica.

C. Migrate the database to an Amazon Aurora multi-master DB cluster.

D. Migrate the database to an Amazon RDS for MySQL DB instance.

ANSWER: C

QUESTION NO: 12

A company has multiple Amazon EC2 instances that run a resource-intensive application in a development environment. A SysOps administrator is implementing a solution to stop these EC2 instances when they are not in use.

Which solution will meet this requirement?

A. Assess AWS CloudTrail logs to verify that there is no EC2 API activity. Invoke an AWS Lambda function to stop the EC2 instances.

B. Create an Amazon CloudWatch alarm to stop the EC2 instances when the average CPU utilization is lower than 5% for a 30-minute period.

C. Create an Amazon CloudWatch metric to stop the EC2 instances when the VolumeReadBytes metric is lower than 500 for a 30-minute period.

D. Use AWS Config to invoke an AWS Lambda function to stop the EC2 instances based on resource configuration changes.

ANSWER: B

Explanation:

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/UsingAlarmActions.html#AddingStopActions

QUESTION NO: 13

A SysOps administrator is deploying a test site running on Amazon EC2 instances. The application requires both incoming and outgoing connectivity to the internet.

Which combination of steps are required to provide internet connectivity to the EC2 instances? (Choose two.)

A. Add a NAT gateway to a public subnet.

B. Attach a private address to the elastic network interface on the EC2 instance.

C. Attach an Elastic IP address to the internet gateway.

D. Add an entry to the route table for the subnet that points to an internet gateway.

E. Create an internet gateway and attach it to a VPC.

ANSWER: D E

Explanation:

https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Internet_Gateway.html

QUESTION NO: 14

A SysOps administrator is testing an application mat is hosted on five Amazon EC2 instances The instances run in an Auto Scaling group behind an Application Load Balancer (ALB) High CPU utilization during load testing is causing the Auto Scaling group to scale out. The SysOps administrator must troubleshoot to find the root cause of the high CPU utilization before the Auto Scaling group scales out.

Which action should the SysOps administrator take to meet these requirements?

A. Enable instance scale-in protection.

B. Place the instance into the Standby stale.

C. Remove the listener from the ALB

D. Suspend the Launch and Terminate process types.

ANSWER: A

QUESTION NO: 15

An ecommerce company uses an Amazon ElastiCache for Memcached cluster for in-memory caching of popular product queries on the shopping site. When viewing recent Amazon CloudWatch metrics data for the ElastiCache cluster, the SysOps administrator notices a large number of evictions.

Which of the following actions will reduce these evictions? (Choose two.)

A. Add an additional node to the ElastiCache cluster.

B. Increase the ElastiCache time to live (TTL).

C. Increase the individual node size inside the ElastiCache cluster.

D. Put an Elastic Load Balancer in front of the ElastiCache cluster.

E. Use Amazon Simple Queue Service (Amazon SQS) to decouple the ElastiCache cluster.

ANSWER: A C

Explanation:

https://d1.awsstatic.com/training-and-certification/docs-sysops-associate/AWS-Certified-SysOps-Administrator-Associate_Sample-Questions_C02.pdf