Exam SOA-C02 Question | SOA-C02 Best Vce

What's more, part of that PracticeVCE SOA-C02 dumps now are free: https://drive.google.com/open?id=1B-O3mA3Frznqk6LqQeGHgx0wqU3Aut92

It is a prevailing belief for many people that practice separated from theories are blindfold. Our SOA-C02 learning quiz is a salutary guidance helping you achieve success. The numerous feedbacks from our clients praised and tested our strength on this career, thus our SOA-C02 practice materials get the epithet of high quality and accuracy. We are considered the best ally to our customers who want to pass their SOA-C02 exam by their first attempt and achieve the certification successfully!

Amazon SOA-C02 Certification is recognized globally and is highly valued by employers, making it an excellent investment for IT professionals. Obtaining this certification can help professionals differentiate themselves from their peers, increase their earning potential, and open up new career opportunities. Additionally, the certification is valid for three years, after which professionals can renew it by passing a recertification exam or completing a qualifying continuing education program.

>> Exam SOA-C02 Question <<

Get Special 30% EXTRA Discount on SOA-C02 Dumps By PracticeVCE

This SOA-C02 exam material contains all kinds of actual Amazon SOA-C02 exam questions and practice tests to help you to ace your exam on the first attempt. A steadily rising competition has been noted in the tech field. Countless candidates around the globe aspire to be Amazon SOA-C02 individuals in this field.

Amazon SOA-C02 (AWS Certified SysOps Administrator - Associate) certification exam is designed for IT professionals who are responsible for deploying, managing, and operating systems on the AWS platform. SOA-C02 exam tests the candidates' understanding of various AWS services and their ability to manage and monitor them efficiently. AWS Certified SysOps Administrator - Associate (SOA-C02) certification helps professionals validate their skills and knowledge in the field of AWS, enabling them to advance their careers.

Amazon AWS Certified SysOps Administrator - Associate (SOA-C02) Sample Questions (Q274-Q279):

NEW QUESTION # 274
A company is trying to connect two applications. One application runs in an on-premises data center that has a hostname of hostl .onprem.private. The other application runs on an Amazon EC2 instance that has a hostname of hostl.awscloud.private. An AWS Site-to-Site VPN connection is in place between the on-premises network and AWS.
The application that runs in the data center tries to connect to the application that runs on the EC2 instance, but DNS resolution fails.
A SysOps administrator must implement DNS resolution between on-premises and AWS resources.
Which solution allows the on-premises application to resolve the EC2 instance hostname?

- A. Set up an Amazon Route 53 outbound resolver endpoint. Associate the resolver with the AWS Region of the EC2 instance. Configure the on-premises DNS resolver to forward awscloud.private DNS queries to the outbound resolver endpoint.
- B. Set up an Amazon Route 53 outbound resolver endpoint with a forwarding rule for the onprem.private hosted zone. Associate the resolver with the AWS Region of the EC2 instance. Configure the on-premises DNS resolver to forward onprem.private DNS queries to the outbound resolver endpoint.
- C. Set up an Amazon Route 53 inbound resolver endpoint. Associate the resolver with the VPC of the EC2 instance. Configure the on-premises DNS resolver to forward awscloud.private DNS queries to the inbound resolver endpoint.
- D. Set up an Amazon Route 53 inbound resolver endpoint with a forwarding rule for the onprem.private hosted zone. Associate the resolver with the VPC of the EC2 instance. Configure the on-premises DNS resolver to forward onprem.private DNS queries to the inbound resolver endpoint.

Answer: D

Explanation:
Step-by-Step
Understand the Problem:
There are two applications, one in an on-premises data center and the other on an Amazon EC2 instance.
DNS resolution fails when the on-premises application tries to connect to the EC2 instance.
The goal is to implement DNS resolution between on-premises and AWS resources.
Analyze the Requirements:
Need to resolve the hostname of the EC2 instance from the on-premises network.
Utilize the existing AWS Site-to-Site VPN connection for DNS queries.
Evaluate the Options:
Option A: Set up an Amazon Route 53 inbound resolver endpoint with a forwarding rule for the onprem.private hosted zone.
This allows DNS queries from on-premises to be forwarded to Route 53 for resolution.
The resolver endpoint is associated with the VPC, enabling resolution of AWS resources.
Option B: Set up an Amazon Route 53 inbound resolver endpoint without specifying the forwarding rule.
This option does not address the specific need to resolve onprem.private DNS queries.
Option C: Set up an Amazon Route 53 outbound resolver endpoint.
Outbound resolver endpoints are used for forwarding DNS queries from AWS to on-premises, not vice versa.
Option D: Set up an Amazon Route 53 outbound resolver endpoint without specifying the forwarding rule.
Similar to Option C, this does not meet the requirement of resolving on-premises queries in AWS.
Select the Best Solution:
Option A: Setting up an inbound resolver endpoint with a forwarding rule for onprem.private and associating it with the VPC ensures that DNS queries from on-premises can resolve AWS resources effectively.
Reference:
Amazon Route 53 Resolver
Integrating AWS and On-Premises Networks with Route 53
Using an Amazon Route 53 inbound resolver endpoint with a forwarding rule ensures that on-premises applications can resolve EC2 instance hostnames effectively.

NEW QUESTION # 275

A SysOps administrator is configuring an Auto Scaling group of Amazon EC2 instances for an application. The average CPU utilization of the instances in the Auto Scaling group must remain at approximately 40% when the load on the application changes. Which solution will meet this requirement in the MOST operationally efficient manner?

- A. Configure a step scaling policy. Create an Amazon CloudWatch alarm that enters ALARM state when CPU utilization is greater than 40%. Associate the alarm with the scaling policy.
- B. Create a scheduled scaling action. Configure the action to run at times when the application typically experiences an increase in traffic.
- C. Configure a target tracking scaling policy. Specify a target value of 40 for average CPU utilization.
- D. Configure a simple scaling policy. Create an Amazon CloudWatch alarm that enters ALARM state when CPU utilization is greater than 40%. Associate the alarm with the scaling policy.

Answer: C

Explanation:
Target tracking automatically adjusts the Auto Scaling group size to hold the specified metric (average CPU utilization) close to the target (40%), without requiring you to define individual alarms or scaling steps. This "set-and-forget" approach is the most operationally efficient way to maintain the desired CPU utilization as load fluctuates.

NEW QUESTION # 276

A SysOps administrator Is troubleshooting an AWS Cloud Formation template whereby multiple Amazon EC2 instances are being created The template is working In us-east-1. but it is failing In us-west-2 with the error code:
How should the administrator ensure that the AWS Cloud Formation template is working in every region?

- A. Copy the source region's Amazon Machine Image (AMI) to the destination region and assign it the same ID.
- B. Modify the AWS CloudFormation template by including the AMI IDs in the "Mappings" section. Refer to the proper mapping within the template for the proper AMI ID.
- C. Edit the AWS CloudFormatton template to offer a drop-down list of all AMIs to the user by using the aws :: EC2:: ami :: imageiD control.
- D. Edit the AWS CloudFormatton template to specify the region code as part of the fully qualified AMI ID.

Answer: B

Explanation:
To ensure that the AWS CloudFormation template works in every region, you should use the Mappings section to specify region-specific AMI IDs. This allows the template to dynamically reference the correct AMI ID based on the region where the stack is being deployed.
Steps:
* Add Mappings to the Template:
* Define the AMI IDs for each region in the Mappings section of the CloudFormation template.
json
Copy code
{
"Mappings": {
"RegionMap": {
"us-east-1": { "AMI": "ami-0123456789abcdef0" },
"us-west-2": { "AMI": "ami-abcdef0123456789" }
}
}
}
* Reference the Mapping in the Template:
* Use the Fn::FindInMap function to reference the correct AMI ID based on the region.
json
Copy code
{
"Resources": {
"MyEC2Instance": {
"Type": "AWS::EC2::Instance",
"Properties": {

```
"ImageId": { "Fn::FindInMap": [ "RegionMap", { "Ref": "AWS::Region" }, "AMI" ] },
}
}
}
}
```
* Deploy the Template:
* Deploy the CloudFormation stack in any region, and it will use the correct AMI ID.
Reference: AWS CloudFormation Mappings

NEW QUESTION # 277

A global company handles a large amount of personally identifiable information (Pll) through an internal web portal. The company's application runs in a corporate data center that is connected to AWS through an AWS Direct Connect connection. The application stores the Pll in Amazon S3. According to a compliance requirement, traffic from the web portal to Amazon S3 must not travel across the internet.

What should a SysOps administrator do to meet the compliance requirement?

- A. Modify the application to use the S3 path-style endpoint.
- B. Configure AWS Network Firewall to redirect traffic to the internal S3 address.
- C. Set up a range of VPC network ACLs to redirect traffic to the Internal S3 address.
- D. Provision an interface VPC endpoint for Amazon S3. Modify the application to use the interface endpoint.

Answer: D

Explanation:
Using the interface endpoint, applications in your on-premises data center can easily query S3 buckets over AWS Direct Connect or Site-to-Site VPN. https://aws.amazon.com/blogs/architecture/choosing-your-vpc-endpoint-strategy-for-amazon-s3/

NEW QUESTION # 278

A company has an infernal web application that runs on Amazon EC2 instances behind an Application Load Balancer.
The instances run in an Amazon EC2 Auto Scaling group in a single Availability Zone.
A SysOps administrator must make the application highly available.
Which action should the SysOps administrator take to meet this requirement?

- A. Update the Auto Scaling group to launch new instances in a second Availability Zone in the same AWS Region.
- B. Increase the maximum number of instances in the Auto Scaling group to meet the capacity that is required at peak usage.
- C. Update the Au o Scaling group to launch new instances in an Availability Zone in a second AWS Region.
- D. Increase the minimum number of i stances in the Auto Scaling group to meet the capacity that is required at peak usage.

Answer: A

Explanation:
An Auto Scaling group can contain EC2 instances in one or more Availability Zones within the same Region. However, Auto Scaling groups cannot span multiple Regions.
https://docs.aws.amazon.com/autoscaling/ec2/userguide/auto-scaling-benefits.html

NEW QUESTION # 279

......