Pdf Amazon SOA-C02 Version | Regualer SOA-C02 Update

P.S. Free & New SOA-C02 dumps are available on Google Drive shared by ITPassLeader: https://drive.google.com/open?id=1jWyEHiER0xckI31WZYrq_8oZXeWre-vD

We really take the requirements of our worthy customers into account. Perhaps you know nothing about our SOA-C02 study guide. Our free demos of our SOA-C02 learning questions will help you know our study materials comprehensively. As we have three different kinds of the SOA-C02 Practice Braindumps, accordingly we have three kinds of the free demos as well. They are a small part of the questions and answers of the SOA-C02 learning quiz.

Amazon SOA-C02 (AWS Certified SysOps Administrator - Associate) exam is designed for IT professionals who are interested in validating their technical expertise in deploying, managing, and operating scalable, highly available, and fault-tolerant systems on the Amazon Web Services (AWS) platform. AWS Certified SysOps Administrator - Associate (SOA-C02) certification exam is intended for individuals who have a solid understanding of AWS services and best practices, as well as experience working with AWS services in production environments.

>> Pdf Amazon SOA-C02 Version <<

Pdf SOA-C02 Version - Provide Valid Material to pass AWS Certified SysOps Administrator - Associate (SOA-C02)

The SOA-C02 Exam software's user-friendly interface is made to uproot potential problems. Once you will try the demo of SOA-C02 exam questions, you will be well- acquainted with the software and its related features. Also SOA-C02 exam comes with various self-assessment features like timed exam, randomization questions, and multiple questions types, test history and score etc. Which means it enables you to customize the question type and you may practice random questions in order to enhance your skills and expertise. You may keep attempting the same questions many a time also.

AWS-SysOps Exam Syllabus Topics:

Section	Objectives	Weight

Networking and Content Delivery	- Implement networking features and connectivity • Configure a VPC (for example, subnets, route tables, network ACLs, security groups, NAT gateway, internet gateway) • Configure private connectivity (for example, Systems Manager Session Manager, VPC endpoints, VPC peering, VPN) • Configure AWS network protection services (for example, AWS WAF, AWS Shield) - Configure domains, DNS services, and content delivery • Configure Route 53 hosted zones and records • Implement Route 53 routing policies (for example, geolocation, geoproximity) • Configure DNS (for example, Route 53 Resolver) • Configure Amazon CloudFront and S3 origin access identity (OAI) • Configure S3 static website hosting - Troubleshoot network connectivity issues • Interpret VPC configurations (for example, subnets, route tables, network ACLs, security groups) • Collect and interpret logs (for example, VPC Flow Logs, Elastic Load Balancer access logs, AWS WAF web ACL logs, CloudFront logs) • Identify and remediate CloudFront caching issues • Troubleshoot hybrid and private connectivity issues	18%
Cost and Performance Optimization	- Implement cost optimization strategies • Implement cost allocation tags • Identify and remediate underutilized or unused resources by using AWS services and tools (for example, Trusted Advisor, AWS Compute Optimizer, Cost Explorer) • Configure AWS Budgets and billing alarms • Assess resource usage patterns to qualify workloads for EC2 Spot Instances • Identify opportunities to use managed services (for example, Amazon RDS, AWS Fargate, EFS) - Implement performance optimization strategies • Recommend compute resources based on performance metrics • Monitor Amazon EBS metrics and modify configuration to increase performance efficiency • Implement S3 performance features (for example, S3 Transfer Acceleration, multipart uploads) • Monitor RDS metrics and modify the configuration to increase performance efficiency (for example, Performance Insights, RDS Proxy) • Enable enhanced EC2 capabilities (for example, enhanced network adapter, instance store, placement groups)	12%

Reliability and Business Continuity	- Implement scalability and elasticity • Create and maintain AWS Auto Scaling plans • Implement caching • Implement Amazon RDS replicas and Amazon Aurora Replicas • Implement loosely coupled architectures • Differentiate between horizontal scaling and vertical scaling - Implement high availability and resilient environments • Configure Elastic Load Balancer and Amazon Route 53 health checks • Differentiate between the use of a single Availability Zone and Multi-AZ deployments (for example, Amazon EC2 Auto Scaling groups, Elastic Load Balancing, Amazon FSx, Amazon RDS) • Implement fault-tolerant workloads (for example, Amazon Elastic File System [Amazon EFS], Elastic IP addresses) • Implement Route 53 routing policies (for example, failover, weighted, latency based) - Implement backup and restore strategies • Automate snapshots and backups based on use cases (for example, RDS snapshots, AWS Backup, RTO and RPO, Amazon Data Lifecycle Manager, retention policy) • Restore databases (for example, point-in-time restore, promote read replica) • Implement versioning and lifecycle rules • Configure Amazon S3 Cross-Region Replication • Execute disaster recovery procedures	16%
Deployment, Provisioning, and Automation	- Provision and maintain cloud resources • Create and manage AMIs (for example, EC2 Image Builder) • Create, manage, and troubleshoot AWS CloudFormation • Provision resources across multiple AWS Regions and accounts (for example, AWS Resource Access Manager, CloudFormation StackSets, IAM cross-account roles) • Select deployment scenarios and services (for example, blue/green, rolling, canary) • Identify and remediate deployment issues (for example, service quotas, subnet sizing, CloudFormation and AWS OpsWorks errors, permissions) - Automate manual or repeatable processes • Use AWS services (for example, OpsWorks, Systems Manager, CloudFormation) to automate deployment processes • Implement automated patch management • Schedule automated tasks by using AWS services (for example, EventBridge, AWS Config)	18%

Amazon AWS Certified SysOps Administrator - Associate (SOA-C02) Sample Questions (Q536-Q541):

NEW QUESTION # 536
A company has a hybrid environment. The company has set up an AWS Direct Connect connection between the company's on-premises data center and a workload that runs in a VPC.
The company uses Amazon Route 53 for DNS on AWS. The company uses a private hosted zone to manage DNS names for a set of services that are hosted on AWS.
The company wants the on-premises servers to use Route 53 for DNS resolution of the private hosted zone.
Which solution will meet these requirements?

- A. Create a Route 53 outbound endpoint. Ensure that security groups and routing allow the traffic from the VPC. Configure the DNS server on the on-premises network to conditionally forward DNS queries for the private hosted zone's domain name to the IP addresses of the outbound endpoint.
- B. Edit the private hosted zone in Route 53 with a PTR record that references the on-premises DNS servers. Configure the DNS server on the on-premises network to conditionally forward DNS queries for the private hosted zone's domain name to

the base of the VPC CIDR IPv4 network range, plus two.

- C. Edit the private hosted zone in Route 53 with a TXT record that references the on-premises DNS servers. Configure the DNS server on the on-premises network to conditionally forward DNS queries for the private hosted zone's domain name to the base of the VPC CIDR IPv4 network range, plus two.
- D. Create a Route 53 inbound endpoint. Ensure that security groups and routing allow the traffic from the on-premises data center. Configure the DNS server on the on-premises network to conditionally forward DNS queries for the private hosted zone's domain name to the IP addresses of the inbound endpoint.

Answer: D

NEW QUESTION # 537
A company uses Amazon Elasticsearch Service (Amazon ES) to analyze sales and customer usage data.
Members of the company's geographically dispersed sales team are traveling. They need to log in to Kibana by using their existing corporate credentials that are stored in Active Directory. The company has deployed Active Directory Federation Services (AD FS) to enable authentication to cloud services.
Which solution will meet these requirements?

- A. Enable Active Directory user authentication in Kibana. Create an IP-based custom domain access policy in Amazon ES that includes the Active Directory server's IP address.
- B. Establish a trust relationship with Kibana on the Active Directory server. Enable Active Directory user authentication in Kibana. Add the Active Directory server's IP address to Kibana.
- C. Deploy an Amazon Cognito user pool. Configure Active Directory as an external identity provider for the user pool. Enable Amazon Cognito authentication for Kibana on Amazon ES.
- D. Configure Active Directory as an authentication provider in Amazon ES. Add the Active Directory server's domain name to Amazon ES. Configure Kibana to use Amazon ES authentication.

Answer: C

Explanation:
https://aws.amazon.com/blogs/security/how-to-enable-secure-access-to-kibana-using-aws-single-sign-on/
https://docs.aws.amazon.com/elasticsearch-service/latest/developerguide/es-cognito-auth.html

NEW QUESTION # 538
A company hosts a web application on an Amazon EC2 instance. The web server logs are published to Amazon CloudWatch Logs. The log events have the same structure and include the HTTP response codes that are associated with the user requests. The company needs to monitor the number of times that the web server returns an HTTP 404 response.
What is the MOST operationally efficient solution that meets these requirements?

- A. Create a script that runs a CloudWatch Logs Insights query that counts the number of 404 codes in the log events during the past hour.
- B. Create an AWS Lambda function that runs a CloudWatch Logs Insights query that counts the number of 404 codes in the log events during the past hour.
- C. Create a CloudWatch Logs metric filter that counts the number of times that the web server returns an HTTP 404 response.
- D. Create a CloudWatch Logs subscription filter that counts the number of times that the web server returns an HTTP 404 response.

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/MonitoringLogData.html

NEW QUESTION # 539
A company analyzes sales data for its customers. Customers upload files to one of the company's Amazon S3 buckets, and a message is posted to an Amazon Simple Queue Service (Amazon SQS) queue that contains the object Amazon Resource Name (ARN). An application that runs on an Amazon EC2 instance polls the queue and processes the messages. The processing time depends on the size of the file.

Customers are reporting delays in the processing of their files. A SysOps administrator decides to configure Amazon EC2 Auto Scaling as the first step. The SysOps administrator creates an Amazon Machine Image (AMI) that is based on the existing EC2 instance. The SysOps administrator also creates a launch template that references the AMI.
How should the SysOps administrator configure the Auto Scaling policy to improve the response time?

- A. Create an Auto Scaling policy based on the ApproximateNumberOfMessagesDelayed metric to scale the number of instances based on the number of messages in the queue that have been delayed.
- B. Create a custom metric based on the ASGAverageCPUUtilization metric and the GroupPendingInstances metric from the Auto Scaling group.
  Modify the application to calculate the metric and post the metric to Amazon CloudWatch once each minute.
  Create an Auto Scaling policy based on this metric to scale the number of instances.
- C. Add several different instance sizes in the launch template.
  Create an Auto Scaling policy based on the ApproximateNumberOfMessagesVisible metric to select the size of the instance based on the number of messages in the queue.
- D. Create a custom metric based on the ApproximateNumberOfMessagesVisible metric and the number of instances in the InService state in the Auto Scaling group.
  Modify the application to calculate the metric and post the metric to Amazon CloudWatch once each minute.Create an Auto Scaling policy based on this metric to scale the number of instances.

Answer: D

Explanation:
When there are delays in processing files due to a high volume of messages in the queue, adding more instances using Auto Scaling can help to reduce the processing time. The ApproximateNumberOfMessagesVisible metric is a good indicator of the workload on the EC2 instances. By creating an Auto Scaling policy based on this metric, the number of instances can be scaled up or down depending on the number of messages in the queue.
https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-target-tracking-metric- math.html#metric-math-sqs-queue-backlog

NEW QUESTION # 540
A company has users that deploy Amazon EC2 instances that have more disk performance capacity than is required. A SysOps administrator needs to review all Amazon Elastic Block Store (Amazon EBS) volumes that are associated with the instances and create cost optimization recommendations based on IOPS and throughput.
What should the SysOps administrator do to meet these requirements in the MOST operationally efficient way?

- A. Install the fio tool onto the EC2 instances and create a cfg file to approximate the required workloads.Use the benchmark results to gauge whether the provisioned EBS volumes are of the most appropriate type.
- B. Opt in to AWS Compute Optimizer. Allow sufficient time for metrics to be gathered. Review the Compute Optimizer findings for EBS volumes.
- C. Use the monitoring graphs in the EC2 console to view metrics for EBS volumes Review the consumed space against the provisioned space on each volume. Identify any volumes that have low utilization.
- D. Stop the EC2 instances from the EC2 console. Change the EC2 instance type to Amazon EBS- optimized. Start the EC2 instances

Answer: B

Explanation:
* AWS Compute Optimizer Overview:
* AWS Compute Optimizer analyzes the configuration and utilization of AWS resources, including EBS volumes, to provide cost-optimization recommendations.
* Steps to Use AWS Compute Optimizer for EBS Volumes:
* Enable Compute Optimizer:
* Open the Compute Optimizer Console.
* Enable the service for your account.
* Allow Metrics Collection:
* Allow sufficient time (up to 12 hours) for Compute Optimizer to gather metrics on your EBS volumes.
* Review Recommendations:
* Go to the Compute Optimizer dashboard.
* Navigate to the EBS volume recommendations.
* Review the findings for underutilized or overprovisioned volumes.

* Why Other Options Are Incorrect:
* A: Manually reviewing EC2 monitoring graphs is less efficient and prone to errors compared to Compute Optimizer.
* B: Changing instance types to EBS-optimized without assessing performance is unnecessary and unrelated to cost optimization.
* D: Installing the fio tool and benchmarking is a time-intensive, manual process that does not align with operational efficiency.
References:
* AWS Compute Optimizer Documentation
* EBS Volume Optimization

NEW QUESTION # 541
.....