Valid Braindumps SOA-C03 Pdf - Valid SOA-C03 Vce

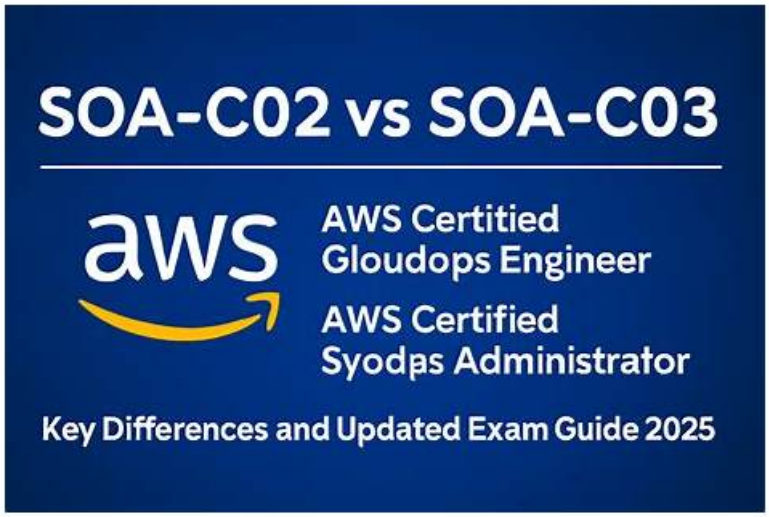

P.S. Free 2026 Amazon SOA-C03 dumps are available on Google Drive shared by Pass4Leader: https://drive.google.com/open?id=1GTK0Dmp1b8LvXZjpcxTgwtvfCmsxycvX

The price for SOA-C03 exam dumps are reasonable, and no matter you are an employee or a student, you can afford it. In addition, you can try free demo before buying, so that you can have a deeper understanding for SOA-C03 exam dumps. In order to build up your confidence for SOA-C03 Exam Materials, we are pass guarantee and money back guarantee. If you fail to pass the exam, we will give you full refund. You can enjoy the right of free update for 365 days, the update version will be sent you automatically.

Amazon SOA-C03 Exam Syllabus Topics:

Topic	Details
Topic 1	• Reliability and Business Continuity: This section measures the skills of System Administrators and focuses on maintaining scalability, elasticity, and fault tolerance. It includes configuring load balancing, auto scaling, Multi-AZ deployments, implementing backup and restore strategies with AWS Backup and versioning, and ensuring disaster recovery to meet RTO and RPO goals.
Topic 2	• Networking and Content Delivery: This section measures skills of Cloud Network Engineers and focuses on VPC configuration, subnets, routing, network ACLs, and gateways. It includes optimizing network cost and performance, configuring DNS with Route 53, using CloudFront and Global Accelerator for content delivery, and troubleshooting network and hybrid connectivity using logs and monitoring tools.
Topic 3	• Security and Compliance: This section measures skills of Security Engineers and includes implementing IAM policies, roles, MFA, and access controls. It focuses on troubleshooting access issues, enforcing compliance, securing data at rest and in transit using AWS KMS and ACM, protecting secrets, and applying findings from Security Hub, GuardDuty, and Inspector.
Topic 4	• Deployment, Provisioning, and Automation: This section measures the skills of Cloud Engineers and covers provisioning and maintaining cloud resources using AWS CloudFormation, CDK, and third-party tools. It evaluates automation of deployments, remediation of resource issues, and managing infrastructure using Systems Manager and event-driven processes like Lambda or S3 notifications.
Topic 5	• Monitoring, Logging, Analysis, Remediation, and Performance Optimization: This section of the exam measures skills of CloudOps Engineers and covers implementing AWS monitoring tools such as CloudWatch, CloudTrail, and Prometheus. It evaluates configuring alarms, dashboards, and notifications, analyzing performance metrics, troubleshooting issues using EventBridge and Systems Manager, and applying strategies to optimize compute, storage, and database performance.

Quiz Trustable SOA-C03 - Valid Braindumps AWS Certified CloudOps Engineer - Associate Pdf

The best way for candidates to know our AWS Certified CloudOps Engineer - Associate SOA-C03 training dumps is downloading our free demo. We provide free PDF demo for each exam. This free demo is a small part of the official complete Amazon SOA-C03 training dumps. The free demo can show you the quality of our exam materials. You can download any time before purchasing.

Amazon AWS Certified CloudOps Engineer - Associate Sample Questions (Q140-Q145):

NEW QUESTION # 140
A company's Amazon EC2 instance with high CPU utilization is a t3.large instance running a test web app. The company determines the app would run better on a compute-optimized large instance.
What should the CloudOps engineer do?

- A. Stop the EC2 instance. Change the instance type to a compute optimized instance. Start the EC2 instance.
- B. Migrate the EC2 instance to a compute optimized instance by using AWS VM Import/Export.
- C. Enable hibernation on the EC2 instance. Change the instance type to a compute optimized instance. Disable hibernation on the EC2 instance.
- D. Change the instance type to a compute optimized instance while the EC2 instance is running.

Answer: A

Explanation:
As described in the AWS Cloud Operations and EC2 Management documentation, changing an instance type (e.g., from T3 to C5) requires that the instance be stopped first. Once stopped, the engineer can modify the instance type through the AWS Management Console, CLI, or API, then start the instance again to apply changes.
This process preserves the root volume, networking configuration, and data, making it an operationally safe and efficient way to upgrade to a different instance family.
Changing the instance type while running (Option D) is unsupported. VM Import/Export (Option A) is for external VM migration. Hibernation (Option B) does not apply to type changes.
Thus, Option C is correct - stopping the instance, changing its type, and restarting it meets AWS best practices.

NEW QUESTION # 141
A company hosts a critical legacy application on two Amazon EC2 instances that are in one Availability Zone. The instances run behind an Application Load Balancer (ALB). The company uses Amazon CloudWatch alarms to send Amazon Simple Notification Service (Amazon SNS) notifications when the ALB health checks detect an unhealthy instance. After a notification, the company's engineers manually restart the unhealthy instance. A CloudOps engineer must configure the application to be highly available and more resilient to failures. Which solution will meet these requirements?

- A. Create an Amazon Machine Image (AMI) from a healthy instance. Launch additional instances from the AMI in the same Availability Zone. Add the new instances to the ALB target group.
- B. Increase the size of each instance. Create an Amazon EventBridge rule. Configure the EventBridge rule to restart the instances if they enter a failed state.
- C. Create an Amazon Machine Image (AMI) from a healthy instance. Launch an additional instance from the AMI in the same Availability Zone. Add the new instance to the ALB target group. Create an AWS Lambda function that runs when an instance is unhealthy. Configure the Lambda function to stop and restart the unhealthy instance.
- D. Create an Amazon Machine Image (AMI) from a healthy instance. Create a launch template that uses the AMI. Create an Amazon EC2 Auto Scaling group that is deployed across multiple Availability Zones. Configure the Auto Scaling group to add instances to the ALB target group.

Answer: D

Explanation:
High availability requires removing single-AZ risk and eliminating manual recovery. The AWS Reliability best practices state to design for multi-AZ and automatic healing: Auto Scaling "helps maintain application availability and allows you to automatically add

or remove EC2 instances" (AWS Auto Scaling User Guide). The Reliability Pillar recommends to "distribute workloads across multiple Availability Zones" and to "automate recovery from failure" (AWS Well-Architected Framework - Reliability Pillar). Attaching the Auto Scaling group to an ALB target group enables health-based replacement: instances failing load balancer health checks are replaced and traffic is routed only to healthy targets. Using an AMI in a launch template ensures consistent, repeatable instance configuration (AWS EC2 Launch Templates). Options A and C keep all instances in a single Availability Zone and rely on manual or ad-hoc restarts, which do not meet high-availability or resiliency goals. Option B only scales vertically and adds a restart rule; it neither removes the single-AZ failure domain nor provides automated replacement. Therefore, creating a multi-AZ EC2 Auto Scaling group with a launch template and attaching it to the ALB target group (Option D) is the CloudOps-aligned solution for resilience and business continuity.
References (AWS CloudOps Documents / Study Guide):
* AWS Certified CloudOps Engineer - Associate (SOA-C03) Exam Guide: Domain 2 - Reliability and Business Continuity
* AWS Well-Architected Framework - Reliability Pillar
* Amazon EC2 Auto Scaling User Guide - Health checks and replacement
* Elastic Load Balancing User Guide - Target group health checks and ALB integration
* Amazon EC2 Launch Templates - Reproducible instance configuration

NEW QUESTION # 142
A company runs thousands of Amazon EC2 instances that are based on the Amazon Linux 2 Amazon Machine Image (AMI). A SysOps administrator must implement a solution to record commands and output from any user that needs an interactive session on one of the EC2 instances. The solution must log the data to a durable storage location. The solution also must provide automated notifications and alarms that are based on the log data.
Which solution will meet these requirements with the MOST operational efficiency?

- A. Configure command session logging on each EC2 instance. Configure the unified Amazon CloudWatch agent to send session logs to Amazon CloudWatch Logs. Set up query filters and alerts by using Amazon Athena.
- B. Require all users to use AWS Systems Manager Session Manager when they need command line access to an EC2 instance. Configure Session Manager to stream session logs to Amazon CloudWatch Logs.
  Set up a metric filter and a metric alarm for relevant security findings in CloudWatch Logs.
- C. Configure command session logging on each EC2 instance. Require all users to use AWS Systems Manager Run Command documents when they need command line access to an EC2 instance.
  Configure the unified Amazon CloudWatch agent to send session logs to Amazon CloudWatch Logs.Set up CloudWatch alarms that are based on Amazon Athena query results.
- D. Require all users to use a central bastion host when they need command line access to an EC2 instance.
  Configure the unified Amazon CloudWatch agent on the bastion host to send session logs to Amazon CloudWatch Logs. Set up a metric filter and a metric alarm for relevant security findings in CloudWatch Logs.

Answer: B

Explanation:
Comprehensive and Detailed Explanation From Exact Extract of AWS CloudOps Documents:
The most operationally efficient solution is C because AWS Systems Manager Session Manager is purpose- built for secure, auditable interactive access to EC2 instances at scale-without managing bastion hosts or distributing SSH keys. Session Manager can be configured to log session activity, including commands and output, to durable destinations such as Amazon CloudWatch Logs (and optionally Amazon S3). This directly satisfies the requirement to record interactive sessions and store logs durably.
For automated notifications and alarms, CloudWatch Logs supports metric filters that transform matching log patterns into CloudWatch metrics. Those metrics can then drive CloudWatch alarms and notifications (for example, via Amazon SNS). This is a standard CloudOps pattern: centralize logs, derive metrics from security-relevant patterns, and alert automatically.
Option A and D require installing and operating agents and building a more complex analytics path (Athena queries for alerting), which is less efficient and introduces more moving parts across thousands of instances.
Option B adds a bastion host dependency that becomes an operational burden (scaling, patching, hardening, HA) and a potential choke point. Session Manager reduces these burdens by using SSM Agent already installed, IAM-based access control, and centralized logging/monitoring integrations.
References:
AWS Systems Manager User Guide - Session Manager and session logging to CloudWatch Logs/S3 Amazon CloudWatch Logs User Guide - Metric filters and alarms from log patterns AWS SysOps Administrator Study Guide - Centralized logging, auditing, and operational monitoring

NEW QUESTION # 143

A company runs a high performance computing (HPC) data-processing application on Amazon EC2 instances in one Availability Zone within a development environment. The application uses a dataset that the company stores on an Amazon S3 general purpose bucket in the same AWS Region as the EC2 instances.
A SysOps administrator must improve the application's performance for retrieval of objects from Amazon S3.
Which solution will meet these requirements?

- A. Create a second general purpose S3 bucket in the same Region. Copy the objects from the original bucket to the new bucket. Use the S3 Express One Zone storage class to store the objects in the new bucket. Update the application to use an S3 Regional endpoint.
- B. Enable S3 Transfer Acceleration for the S3 bucket. Create an S3 access point for the bucket. Update the application to use the access point.
- C. Create an S3 Lifecycle configuration for the S3 bucket to move all objects to the S3 Express One Zone storage class. Update the application to use an S3 Regional endpoint.
- D. Create an S3 directory bucket in the same Availability Zone. Import objects from the original bucket to the new bucket. Use the S3 Express One Zone storage class to store the objects in the new bucket.Update the application to use an S3 Zonal endpoint.

Answer: D

Explanation:
Comprehensive and Detailed Explanation From Exact Extract of AWS CloudOps Documents:
The correct answer is D because Amazon S3 Express One Zone with directory buckets and zonal endpoints is specifically designed for single-Availability Zone, high-performance workloads such as HPC, machine learning, and analytics applications running on Amazon EC2. AWS CloudOps documentation states that S3 Express One Zone delivers single-digit millisecond latency and up to 10x higher request performance compared to general purpose S3 buckets when data is accessed from the same Availability Zone. An S3 directory bucket is required to use the S3 Express One Zone storage class. These buckets are explicitly associated with a single Availability Zone and use zonal endpoints, which eliminate cross-AZ network hops and significantly reduce latency. Importing the data from the existing general purpose bucket ensures compatibility while achieving maximum throughput and lowest latency. Option A is incorrect because S3 Transfer Acceleration is optimized for long-distance, internet-based transfers, not for in-Region HPC workloads. Option B is incorrect because lifecycle policies cannot move objects into S3 Express One Zone, and S3 Express One Zone does not use Regional endpoints. Option C is incorrect because general purpose buckets do not support zonal endpoints and therefore cannot achieve the same performance benefits.
AWS CloudOps performance optimization guidance clearly identifies S3 directory buckets with S3 Express One Zone and zonal endpoints as the optimal architecture for high-throughput, low-latency workloads in a single Availability Zone.
References:
Amazon S3 User Guide - S3 Express One Zone and Directory Buckets
AWS SysOps Administrator Study Guide - Storage Performance Optimization AWS Well-Architected Framework - Performance Efficiency Pillar

NEW QUESTION # 144
A company is running an application on premises and wants to use AWS for data backup. All of the data must be available locally. The backup application can write only to block-based storage that is compatible with the Portable Operating System Interface (POSIX).
Which backup solution will meet these requirements?

- A. Configure the backup software to use Amazon S3 Glacier Flexible Retrieval as the target for the data backups.
- B. Configure the backup software to use Amazon S3 as the target for the data backups.
- C. Use AWS Storage Gateway, and configure it to use gateway-stored volumes.
- D. Use AWS Storage Gateway, and configure it to use gateway-cached volumes.

Answer: C

Explanation:
The Storage Gateway service enables hybrid cloud backup by presenting local block storage that synchronizes with AWS cloud storage. For scenarios where all data must remain available locally while still backed up to AWS, the correct mode is gateway-stored volumes.
"Use stored volumes if you want to keep all your data locally while asynchronously backing up point-in-time snapshots to Amazon S3 for durable storage." These volumes expose an iSCSI interface compatible with POSIX file systems, allowing direct use by on-premises backup software.
Gateway-cached volumes (Option C) store primary data in AWS with limited local cache, violating the "all data must be available

locally" requirement. Options A and B are object-based storage solutions, not compatible with POSIX or block-based backup applications.
Therefore, Option D fully satisfies CloudOps reliability and continuity best practices by ensuring local availability, cloud durability, and POSIX compatibility for backups.

NEW QUESTION # 145
......

There is almost no innovative and exam-oriented format that can be compared with the precision and relevance of the actual AWS Certified CloudOps Engineer - Associate exam questions, you get with Pass4Leader brain dumps PDF. As per the format of the SOA-C03 Exam, our experts have consciously created a questions and answers pattern. It saves your time by providing you direct and precise information that will help you cover the syllabus contents within no time.

Valid SOA-C03 Vce: https://www.pass4leader.com/Amazon/SOA-C03-exam.html