[2025-November-NewBraindump2go SOA-C03 VCE Questions Free[Q1-Q40

2025/October Latest Braindump2go SOA-C03 Exam Dumps with PDF and VCE Free Updated Today! Following are some new Braindump2go SOA-C03 Real Exam Questions!QUESTION 1A medical research company uses an Amazon Bedrock powered AI assistant with agents and knowledge bases to provide physicians quick access to medical study protocols. The company needs to generate audit reports that contain user identities, usage data for Bedrock agents, access data for knowledge bases, and interaction parameters. Which solution will meet these requirements? A. Use AWS CloudTrail to log API events from generative AI workloads. Store the events in CloudTrail Lake, Use SQL-like queries to generate reports.B. Use Amazon CloudWatch to capture generative AI application logs. Stream the logs to Amazon OpenSearch Service. Use an OpenSearch dashboard visualization to generate reports.C. Use Amazon CloudWatch to log API events from generative AI workloads. Send the events to an Amazon S3 bucket. Use Amazon Athena queries to generate reports.D. Use AWS CloudTrail to capture generative AI application logs. Stream the logs to Amazon Managed Service for Apache Flink. Use SQL queries to generate reports. Answer: AExplanation: As per AWS Cloud Operations, Bedrock, and Governance documentation, AWS CloudTrail is the authoritative service for capturing API activity and audit trails across AWS accounts. For Amazon Bedrock, CloudTrail records all user-initiated API calls, including interactions with agents, knowledge bases, and generative AI model parameters. Using CloudTrail Lake, organizations can store, query, and analyze CloudTrail events directly without needing to export data. CloudTrail Lake supports SQL-like queries for generating audit and compliance reports, enabling the company to retrieve information such as user identity, API usage, timestamp, model or agent ID, and invocation parameters. In contrast, CloudWatch focuses on operational metrics and log streaming, not API-level identity data. OpenSearch or Flink would add unnecessary complexity and cost for this use case. Thus, the AWS-recommended CloudOps best practice is to leverage CloudTrail with CloudTrail Lake to maintain auditable, queryable API activity for Bedrock workloads, fulfilling governance and compliance requirements.QUESTION 2A company needs to enforce tagging requirements for Amazon DynamoDB tables in its AWS accounts. A CloudOps engineer must implement a solution to identify and remediate all DynamoDB tables that do not have the appropriate tags. Which solution will meet these requirements with the LEAST operational overhead? A. Create a custom AWS Lambda function to evaluate and remediate all DynamoDB tables. Create an Amazon EventBridge scheduled rule to invoke the Lambda function.B. Create a custom AWS Lambda function to evaluate and remediate all DynamoDB tables. Create an AWS Config custom rule to invoke the Lambda function.C. Use the required-tags AWS Config managed rule to evaluate all DynamoDB tables for the appropriate tags. Configure an automatic remediation action that uses an AWS Systems Manager Automation custom runbook.D. Create an Amazon EventBridge managed rule to evaluate all DynamoDB tables for the appropriate tags. Configure the EventBridge rule to run an AWS Systems Manager Automation custom runbook for remediation. Answer: CExplanation: According to the AWS Cloud Operations, Governance, and Compliance documentation, AWS Config provides managed rules that automatically evaluate resource configurations for compliance. The "required-tags" managed rule allows CloudOps teams to specify mandatory tags (e.g., Environment, Owner, CostCenter) and automatically detect non-compliant resources such as DynamoDB tables. Furthermore, AWS Config supports automatic remediation through AWS Systems Manager Automation runbooks, enabling correction actions (for example, adding missing tags) without manual intervention. This automation minimizes operational overhead and ensures continuous compliance across multiple accounts. Using a custom Lambda function (Options A or B) introduces unnecessary management complexity, while EventBridge rules alone (Option D) do not provide resource compliance tracking or historical visibility. Therefore, Option C provides the most efficient, fully managed, and compliant CloudOps solution.QUESTION 3A company's website runs on an Amazon EC2 Linux instance. The website needs to serve PDF files from an Amazon S3 bucket. All public access to the S3 bucket is blocked at the account level. The company needs to allow website users to download the PDF files. Which solution will meet these requirements with the LEAST administrative effort?A. Create an IAM role that has a policy that allows s3:list* and s3:get* permissions. Assign the role to the EC2 instance. Assign a company employee to download requested PDF files to the EC2 instance and deliver the files to website users. Create an AWS Lambda function to periodically delete local files.B. Create an Amazon CloudFront distribution that uses an origin access control (OAC) that points to the S3 bucket. Apply a bucket policy to the bucket to allow connections from the CloudFront distribution. Assign a company employee to provide a download URL that contains the distribution URL and the object path to users when users request PDF files.C. Change the S3 bucket permissions to allow public access on the source S3 bucket. Assign a company employee to provide a PDF file URL to users when users request the PDF files.D. Deploy an EC2 instance that has an IAM instance profile to a public subnet. Use a signed URL from the EC2 instance to provide temporary access to the S3 bucket for website users. Answer: BExplanation: Per the AWS Cloud Operations, Networking, and Security documentation, the best practice for serving private S3 content securely to end users is to use Amazon CloudFront with Origin Access Control (OAC).OAC

enables CloudFront to access S3 buckets privately, even when Block Public Access settings are enabled at the account level. This allows content to be delivered globally and securely without making the S3 bucket public. The bucket policy explicitly allows access only from the CloudFront distribution, ensuring that users can retrieve PDF files only via CloudFront URLs. This configuration offers: Automatic scalability through CloudFront caching, Improved security via private access control, Minimal administration effort with fully managed services. Other options require manual handling or make the bucket public, violating AWS security best practices. Therefore, Option B--using CloudFront with Origin Access Control and a restrictive bucket policy-- provides the most secure, efficient, and low-maintenance CloudOps solution.QUESTION 4A financial services company stores customer images in an Amazon S3 bucket in the us-east-1 Region. To comply with regulations, the company must ensure that all existing objects are replicated to an S3 bucket in a second AWS Region. If an object replication fails, the company must be able to retry replication for the object. What solution will meet these requirements? A. Configure Amazon S3 Cross-Region Replication (CRR). Use Amazon S3 live replication to replicate existing objects.B. Configure Amazon S3 Cross-Region Replication (CRR). Use S3 Batch Replication to replicate existing objects.C. Configure Amazon S3 Cross-Region Replication (CRR). Use S3 Replication Time Control (S3 RTC) to replicate existing objects.D. Use S3 Lifecycle rules to move objects to the destination bucket in a second Region. Answer: BExplanation: Per the AWS Cloud Operations and S3 Data Management documentation, Cross-Region Replication (CRR) automatically replicates new objects between S3 buckets across Regions. However, CRR alone does not retroactively replicate existing objects created before replication configuration. To include such objects, AWS introduced S3 Batch Replication. S3 Batch Replication scans the source bucket and replicates all existing objects that were not copied previously. Additionally, it can retry failed replication tasks automatically, ensuring regulatory compliance for complete dataset replication.S3 Replication Time Control (S3 RTC) guarantees predictable replication times for new objects only-- it does not cover previously stored data. S3 Lifecycle rules (Option D) move or transition objects between storage classes or buckets, but not in a replication context. Therefore, the correct solution is to use S3 Cross-Region Replication (CRR) combined with S3 Batch Replication to ensure all current and future data is synchronized across Regions with retry capability.QUESTION 5A CloudOps engineer has created a VPC that contains a public subnet and a private subnet. Amazon EC2 instances that were launched in the private subnet cannot access the internet. The default network ACL is active on all subnets in the VPC, and all security groups allow outbound traffic. Which solution will provide the EC2 instances in the private subnet with access to the internet?A. Create a NAT gateway in the public subnet. Create a route from the private subnet to the NAT gateway.B. Create a NAT gateway in the public subnet. Create a route from the public subnet to the NAT gateway.C. Create a NAT gateway in the private subnet. Create a route from the public subnet to the NAT gateway.D. Create a NAT gateway in the private subnet. Create a route from the private subnet to the NAT gateway. Answer: AExplanation: According to the AWS Cloud Operations and Networking documentation, instances in a private subnet do not have a direct route to the internet gateway and thus require a NAT gateway for outbound internet access. The correct configuration is to create a NAT gateway in the public subnet, associate an Elastic IP address, and then update the private subnet's route table to send all 0.0.0.0/0 traffic to the NAT gateway. This enables instances in the private subnet to initiate outbound connections while keeping inbound traffic blocked for security. Placing the NAT gateway inside the private subnet (Options C or D) prevents connectivity because it would not have a route to the internet gateway. Configuring routes from the public subnet to the NAT gateway (Option B) does not serve private subnet traffic. Hence, Option A follows AWS best practices for enabling secure, managed, outbound-only internet access from private resources.QUESTION 6A company's architecture team must receive immediate email notifications whenever new Amazon EC2 instances are launched in the company's main AWS production account. What should a CloudOps engineer do to meet this requirement?A. Create a user data script that sends an email message through a smart host connector. Include the architecture team's email address in the user data script as the recipient. Ensure that all new EC2 instances include the user data script as part of a standardized build process.B. Create an Amazon Simple Notification Service (Amazon SNS) topic and a subscription that uses the email protocol. Enter the architecture team's email address as the subscriber. Create an Amazon EventBridge rule that reacts when EC2 instances are launched. Specify the SNS topic as the rule's target.C. Create an Amazon Simple Queue Service (Amazon SQS) queue and a subscription that uses the email protocol. Enter the architecture team's email address as the subscriber. Create an Amazon EventBridge rule that reacts when EC2 instances are launched. Specify the SQS queue as the rule's target.D. Create an Amazon Simple Notification Service (Amazon SNS) topic. Configure AWS Systems Manager to publish EC2 events to the SNS topic. Create an AWS Lambda function to poll the SNS topic. Configure the Lambda function to send any messages to the architecture team's email address. Answer: BExplanation: As per the AWS Cloud Operations and Event Monitoring documentation, the most efficient method for event-driven notification is to use Amazon EventBridge to detect specific EC2 API events and trigger a Simple Notification Service (SNS) alert. EventBridge continuously monitors AWS service events, including RunInstances, which signals the creation of new EC2 instances. When such an event occurs, EventBridge sends it to an

SNS topic, which then immediately emails subscribed recipients -- in this case, the architecture team. This combination provides real-time, serverless notifications with minimal management. SQS (Option C) is designed for queue-based processing, not direct user alerts. User data scripts (Option A) and custom polling with Lambda (Option D) introduce unnecessary operational complexity and latency. Hence, Option B is the correct and AWS-recommended CloudOps design for immediate launch notifications.QUESTION 7A company runs an application on Amazon EC2 that connects to an Amazon Aurora PostgreSQL database. A developer accidentally drops a table from the database, causing application errors. Two hours later, a CloudOps engineer needs to recover the data and make the application functional again. Which solution will meet this requirement? A. Use the Aurora Backtrack feature to rewind the database to a specified time, 2 hours in the past.B. Perform a point-in-time recovery on the existing database to restore the database to a specified point in time, 2 hours in the past.C. Perform a point-in-time recovery and create a new database to restore the database to a specified point in time, 2 hours in the past. Reconfigure the application to use a new database endpoint.D. Create a new Aurora cluster. Choose the Restore data from S3 bucket option. Choose log files up to the failure time 2 hours in the past. Answer: CExplanation: In the AWS Cloud Operations and Aurora documentation, when data loss occurs due to human error such as dropped tables, Point-in-Time Recovery (PITR) is the recommended method for restoration. PITR creates a new Aurora cluster restored to a specific time before the failure. The restored cluster has a new endpoint that must be reconfigured in the application to resume normal operations. AWS does not support performing PITR directly on an existing production database because that would overwrite current data. Aurora Backtrack (Option A) applies only to Aurora MySQL, not PostgreSQL. Option B is incorrect because PITR cannot be executed in place. Option D refers to an import process from S3, which is unrelated to time-based recovery. Hence, Option C is correct and follows the AWS CloudOps standard recovery pattern for PostgreSQL workloads.QUESTION 8A company is using an Amazon Aurora MySQL DB cluster that has point-in-time recovery, backtracking, and automatic backup enabled. A CloudOps engineer needs to roll back the DB cluster to a specific recovery point within the previous 72 hours. Restores must be completed in the same production DB cluster. Which solution will meet these requirements?A. Create an Aurora Replica. Promote the replica to replace the primary DB instance.B. Create an AWS Lambda function to restore an automatic backup to the existing DB cluster.C. Use backtracking to rewind the existing DB cluster to the desired recovery point.D. Use point-in-time recovery to restore the existing DB cluster to the desired recovery point. Answer: C Explanation: As documented in AWS Cloud Operations and Database Recovery, Aurora Backtrack allows you to rewind the existing database cluster to a chosen point in time without creating a new cluster. This feature supports fine-grained rollback for accidental data changes, making it ideal for scenarios like table deletions or logical corruption. Backtracking maintains continuous transaction logs and permits rewinding within a configurable window (up to 72 hours). It does not require creating a new cluster or endpoint, and it preserves the same production environment, fulfilling the operational requirement for in-place recovery. In contrast, Point-in-Time Recovery (Option D) always creates a new cluster, while replica promotion (Option A) and Lambda restoration (Option B) are unrelated to immediate rollback operations. Therefore, Option C, using Aurora Backtrack, best meets the requirement for same-cluster restoration and minimal downtime.QUESTION 9An application runs on Amazon EC2 instances that are in an Auto Scaling group. A CloudOps engineer needs to implement a solution that provides a central storage location for errors that the application logs to disk. The solution must also provide an alert when the application logs an error. What should the CloudOps engineer do to meet these requirements?A. Deploy and configure the Amazon CloudWatch agent on the EC2 instances to log to a CloudWatch log group. Create a metric filter on the target CloudWatch log group. Create a CloudWatch alarm that publishes to an Amazon Simple Notification Service (Amazon SNS) topic that has an email subscription.B. Create a cron job on the EC2 instances to identify errors and push the errors to an Amazon CloudWatch metric filter. Configure the filter to publish to an Amazon Simple Notification Service (Amazon SNS) topic that has an SMS subscription.C. Deploy an AWS Lambda function that pushes the errors directly to Amazon CloudWatch Logs. Configure the Lambda function to run every time the log file is updated on disk.D. Create an Auto Scaling lifecycle hook that invokes an EC2-based script to identify errors. Configure the script to push the error messages to an Amazon CloudWatch log group when the EC2 instances scale in. Create a CloudWatch alarm that publishes to an Amazon Simple Notification Service (Amazon SNS) topic that has an email subscription when the number of error messages exceeds a threshold. Answer: AExplanation: The AWS Cloud Operations and Monitoring documentation specifies that the Amazon CloudWatch Agent is the recommended tool for collecting system and application logs from EC2 instances. The agent pushes these logs into a centralized CloudWatch Logs group, providing durable storage and real-time monitoring. Once the logs are centralized, a CloudWatch Metric Filter can be configured to search for specific error keywords (for example, "ERROR" or "FAILURE"). This filter transforms matching log entries into custom metrics. From there, a CloudWatch Alarm can monitor the metric threshold and publish notifications to an Amazon SNS topic, which can send email or SMS alerts to subscribed recipients. This combination provides a fully automated, managed, and serverless solution for log aggregation and error alerting. It eliminates the need for manual

cron jobs (Option B), custom scripts (Option D), or Lambda-based log streaming (Option C).QUESTION 10A company's security policy prohibits connecting to Amazon EC2 instances through SSH and RDP. Instead, staff must use AWS Systems Manager Session Manager. Users report they cannot connect to one Ubuntu instance, even though they can connect to others. What should a CloudOps engineer do to resolve this issue?A. Add an inbound rule for port 22 in the security group associated with the Ubuntu instance.B. Assign the AmazonSSMManagedInstanceCore managed policy to the EC2 instance profile for the Ubuntu instance.C. Configure the SSM Agent to log in with a user name of "ubuntu".D. Generate a new key pair, configure Session Manager to use this new key pair, and provide the private key to the users. Answer: BExplanation: According to AWS Cloud Operations and Systems Manager documentation, Session Manager requires that each managed instance be associated with an IAM instance profile that grants Systems Manager core permissions. The required permissions are provided by the AmazonSSMManagedInstanceCore AWS-managed policy. If this policy is missing or misconfigured, the Systems Manager Agent (SSM Agent) cannot communicate with the Systems Manager service, causing connection failures even if the agent is installed and running. This explains why other instances work--those instances likely have the correct IAM role attached. Enabling port 22 (Option A) violates the company's security policy, while configuring user names (Option C) and key pairs (Option D) are irrelevant because Session Manager operates over secure API channels, not SSH keys. Therefore, the correct resolution is to attach or update the instance profile with the AmazonSSMManagedInstanceCore policy, restoring Session Manager connectivity.QUESTION 11A company deploys an application on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB). The company wants to protect the application from SQL injection attacks. Which solution will meet this requirement? A. Deploy AWS Shield Advanced in front of the ALB. Enable SQL injection filtering.B. Deploy AWS Shield Standard in front of the ALB. Enable SQL injection filtering.C. Deploy a vulnerability scanner on each EC2 instance. Continuously scan the application code.D. Deploy AWS WAF in front of the ALB. Subscribe to an AWS Managed Rule for SQL injection filtering. Answer: DExplanation: The AWS Cloud Operations and Security documentation confirms that AWS WAF (Web Application Firewall) is designed to protect web applications from application-layer threats, including SQL injection, cross-site scripting (XSS), and other OWASP Top 10 vulnerabilities. When integrated with an Application Load Balancer, AWS WAF inspects incoming traffic using rule groups. The AWS Managed Rules for SQL Injection Protection provide preconfigured, continuously updated filters that detect and block malicious SQL patterns. AWS Shield (Standard or Advanced) defends against DDoS attacks, not application-layer SQL attacks, and vulnerability scanners (Option C) only detect, not prevent, exploitation. Thus, Option D provides the correct, managed, and automated protection aligned with AWS best practices.QUESTION 12An AWS Lambda function is intermittently failing several times a day. A CloudOps engineer must find out how often this error occurred in the last 7 days. Which action will meet this requirement in the MOST operationally efficient manner?A. Use Amazon Athena to query the Amazon CloudWatch logs that are associated with the Lambda function.B. Use Amazon Athena to query the AWS CloudTrail logs that are associated with the Lambda function.C. Use Amazon CloudWatch Logs Insights to query the associated Lambda function logs.D. Use Amazon OpenSearch Service to stream the Amazon CloudWatch logs for the Lambda function. Answer: CExplanation: The AWS Cloud Operations and Monitoring documentation states that Amazon CloudWatch Logs Insights provides a purpose-built query engine for analyzing and visualizing log data directly within CloudWatch. For Lambda, all invocation results (including errors) are automatically logged to CloudWatch Logs. By querying these logs with CloudWatch Logs Insights, the CloudOps engineer can efficiently count the number of "ERROR" or "Exception" occurrences over the past 7 days using simple SQL-like commands. This method is serverless, cost-efficient, and real-time. Athena (Options A and B) would require exporting data to Amazon S3, and OpenSearch (Option D) adds unnecessary operational complexity. Thus, Option C provides the most efficient and native AWS CloudOps approach for rapid Lambda error analysis.QUESTION 13A company is storing backups in an Amazon S3 bucket. These backups must not be deleted for at least 3 months after creation. What should the CloudOps engineer do? A. Configure an IAM policy that denies the s3:DeleteObject action for all users. Three months after an object is written, remove the policy.B. Enable S3 Object Lock on a new S3 bucket in compliance mode. Place all backups in the new S3 bucket with a retention period of 3 months.C. Enable S3 Versioning on the existing S3 bucket. Configure S3 Lifecycle rules to protect the backups.D. Enable S3 Object Lock on a new S3 bucket in governance mode. Place all backups in the new S3 bucket with a retention period of 3 months. Answer: B Explanation: Per the AWS Cloud Operations and Data Protection documentation, S3 Object Lock enforces write- once-read-many (WORM) protection on objects for a defined retention period. There are two modes: Compliance mode: Even the root user cannot delete or modify objects during the retention period. Governance mode: Privileged users with special permissions can override lock settings. For regulatory or audit requirements that prohibit deletion, Compliance mode is the correct choice. When configured with a 3-month retention period, all backup objects are protected from deletion until expiration, ensuring compliance with data retention mandates. Versioning (Option C) alone does not prevent deletion. IAM-based restrictions (Option A) lack time-based enforcement

and require manual intervention. Governance mode (Option D) is less strict and unsuitable for regulatory retention. Thus, Option B is the correct CloudOps solution for immutable S3 backups.QUESTION 14A 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. Migrate the EC2 instance to a compute optimized instance by using AWS VM Import/Export.B. Enable hibernation on the EC2 instance. Change the instance type to a compute optimized instance. Disable hibernation on the EC2 instance. C. Stop the EC2 instance. Change the instance type to a compute optimized instance. Start the EC2 instance.D. Change the instance type to a compute optimized instance while the EC2 instance is running. Answer: C 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.QUESTION 15A company's CloudOps engineer monitors multiple AWS accounts in an organization and checks each account's AWS Health Dashboard. After adding 10 new accounts, the engineer wants to consolidate health alerts from all accounts. Which solution meets this requirement with the least operational effort? A. Enable organizational view in AWS Health.B. Configure the Health Dashboard in each account to forward events to a central AWS CloudTrail log.C. Create an AWS Lambda function to query the AWS Health API and write all events to an Amazon DynamoDB table.D. Use the AWS Health API to write events to an Amazon DynamoDB table. Answer: AExplanation: The AWS Cloud Operations and Governance documentation defines that enabling Organizational View in AWS Health allows the management account in AWS Organizations to view and aggregate health events from all member accounts. This feature provides a single-pane-of-glass view of service health issues, account-specific events, and planned maintenance across the organization -- without requiring additional automation or data pipelines. Alternative options (B, C, and D) require custom integration and ongoing maintenance. CloudTrail does not natively forward AWS Health events, and custom Lambda or DynamoDB approaches increase complexity. Therefore, Option A -- enabling the Organizational View feature in AWS Health -- is the most operationally efficient and AWS-recommended solution. OUESTION 16A company that uses AWS Organizations recently implemented AWS Control Tower. The company now needs to centralize identity management. A CloudOps engineer must federate AWS IAM Identity Center with an external SAML 2.0 identity provider (IdP) to centrally manage access to all AWS accounts and cloud applications. Which prerequisites must the CloudOps engineer have so that the CloudOps engineer can connect to the external IdP? (Select TWO.)A. A copy of the IAM Identity Center SAML metadataB. The IdP metadata, including the public X.509 certificateC. The IP address of the IdPD. Root access to the management accountE. Administrative permissions to the member accounts of the organizationAnswer: AB Explanation: According to the AWS Cloud Operations and Identity Management documentation, when configuring federation between IAM Identity Center (formerly AWS SSO) and an external SAML 2.0 identity provider, two key prerequisites are required: The IAM Identity Center SAML metadata file -- This is uploaded to the external IdP to establish trust, define SAML endpoints, and enable identity federation. The IdP metadata (including the public X.509 certificate) -- This information is imported into IAM Identity Center to validate authentication assertions and encryption signatures. IAM Identity Center and the IdP exchange this metadata to mutually establish secure, bidirectional federation. Network-level details such as IP addresses (Option C) are unnecessary. Root access (Option D) or permissions to member accounts (Option E) are not required; only Control Tower or IAM administrative permissions in the management account are needed for setup. Thus, the correct answer is A and B -- the SAML metadata from both sides is required for federation. QUESTION 17A company hosts an FTP server on EC2 instances. AWS Security Hub sends findings to Amazon EventBridge when the FTP port becomes publicly exposed in attached security groups. A CloudOps engineer needs an automated, event-driven remediation solution to remove public access from security groups. Which solution will meet these requirements? A. Configure the existing EventBridge event to stop the EC2 instances that have the exposed port. B. Create a cron job for the FTP server to invoke an AWS Lambda function. Configure the Lambda function to modify the security group of the identified EC2 instances and to remove the instances that allow public access.C. Create a cron job for the FTP server that invokes an AWS Lambda function. Configure the Lambda function to modify the server to use SFTP instead of FTP.D. Configure the existing EventBridge event to invoke an AWS Lambda function. Configure the function to remove the security group rule that allows public access. Answer: DExplanation: Per the AWS Cloud Operations and Security Automation documentation, Security Hub integrates with Amazon EventBridge to publish findings in real time. These events can trigger automated responses using AWS Lambda functions or AWS Systems Manager Automation runbooks. In this scenario, the correct CloudOps approach is

to configure the existing EventBridge rule to invoke a Lambda function that inspects the event payload, identifies the affected security group, and removes the offending inbound rule (e.g., port 21 open to 0.0.0.0/0). This event-driven remediation provides continuous compliance and eliminates manual intervention. Cron jobs (Options B and C) contradict event-driven design and add operational overhead. Stopping instances (Option A) doesn't address the root cause -- the insecure security group. Thus, Option D aligns with AWS best practices for automated security remediation through EventBridge and Lambda.QUESTION 18A company has an application running on EC2 that stores data in an Amazon RDS for MySQL Single- AZ DB instance. The application requires both read and write operations, and the company needs failover capability with minimal downtime. Which solution will meet these requirements?A. Modify the DB instance to be a Multi-AZ DB instance deployment.B. Add a read replica in the same Availability Zone where the DB instance is deployed.C. Add the DB instance to an Auto Scaling group that has a minimum capacity of 2 and a desired capacity of 2.D. Use RDS Proxy to configure a proxy in front of the DB instance. Answer: A Explanation: According to the AWS Cloud Operations and Database Reliability documentation, Amazon RDS Multi- AZ deployments provide high availability and automatic failover by maintaining a synchronous standby replica in a different Availability Zone. In the event of instance failure, planned maintenance, or Availability Zone outage, Amazon RDS automatically promotes the standby to primary with minimal downtime (typically less than 60 seconds). The failover is transparent to applications because the DB endpoint remains the same. By contrast, read replicas (Option B) are asynchronous and do not provide automated failover. Auto Scaling (Option C) applies to EC2, not RDS. RDS Proxy (Option D) improves connection management but does not add redundancy. Thus, Option A -- converting the RDS instance into a Multi-AZ deployment -- delivers the required high availability and business continuity with minimal operational effort.QUESTION 19A company has an AWS CloudFormation template that includes an AWS::EC2::Instance resource and a custom resource (Lambda function). The Lambda function fails because it runs before the EC2 instance is launched. Which solution will resolve this issue? A. Add a DependsOn attribute to the custom resource. Specify the EC2 instance in the DependsOn attribute.B. Update the custom resource's service token to point to a valid Lambda function.C. Update the Lambda function to use the cfn-response module to send a response to the custom resource.D. Use the Fn::If intrinsic function to check for the EC2 instance before the custom resource runs. Answer: AExplanation: The AWS Cloud Operations and Infrastructure-as-Code documentation specifies that when using AWS CloudFormation, resources are created in parallel by default unless explicitly ordered using DependsOn.If a custom resource (Lambda) depends on another resource (like an EC2 instance) to exist before execution, a DependsOn attribute must be added to enforce creation order. This ensures the EC2 instance is launched and available before the custom resource executes its automation logic. Updating the service token (Option B) doesn't affect order of execution. The cfn-response module (Option C) handles callback communication but not sequencing. Fn::If (Option D) is for conditional creation, not dependency control. Therefore, Option A is correct -- adding a DependsOn attribute guarantees that CloudFormation provisions the EC2 instance before executing the Lambda custom resource.QUESTION 20A company uses an Amazon Simple Queue Service (Amazon SQS) queue and Amazon EC2 instances in an Auto Scaling group with target tracking for a web application. The company collects the ASGAverageNetworkIn metric but notices that instances do not scale fast enough during peak traffic. There are a large number of SQS messages accumulating in the queue. A CloudOps engineer must reduce the number of SQS messages during peak periods. Which solution will meet this requirement? A. Define and use a new custom Amazon CloudWatch metric based on the SQS ApproximateNumberOfMessagesDelayed metric in the target tracking policy.B. Define and use Amazon CloudWatch metric math to calculate the SQS queue backlog for each instance in the target tracking policy.C. Define and use step scaling by specifying a ChangeInCapacity value for the EC2 instances.D. Define and use simple scaling by specifying a ChangeInCapacity value for the EC2 instances. Answer: BExplanation: According to the AWS Cloud Operations and Auto Scaling documentation, scaling applications that consume Amazon SQS messages should be driven by queue backlog per instance, not by general system metrics such as network traffic or CPU. The correct approach is to calculate a custom metric using CloudWatch metric math that divides the SQS metric ApproximateNumberOfMessagesVisible by the number of active EC2 instances in the Auto Scaling group. This "backlog per instance" value represents the average number of messages waiting to be processed by each instance. Then, the CloudOps engineer can create a target tracking policy that automatically scales out or in based on maintaining a desired backlog threshold. This approach ensures dynamic, workload-driven scaling behavior that reacts in near real time to message volume. Step and simple scaling (Options C and D) require manual thresholds and do not automatically balance the load per instance. Thus, Option B--using CloudWatch metric math to define queue backlog per instance for target tracking--is the most effective and AWS-recommended CloudOps practice.OUESTION 21A CloudOps engineer has created an AWS Service Catalog portfolio and shared it with a second AWS account in the company, managed by a different CloudOps engineer. Which action can the CloudOps engineer in the second account perform?A. Add a product from the imported portfolio to a local portfolio. B. Add new products to the imported portfolio.C. Change the launch role for the products contained in the imported portfolio.D.

documentation, when a portfolio is shared across AWS accounts, the recipient account imports the shared portfolio. The recipient CloudOps engineer cannot modify the original products or their configurations but can: Add products from the imported portfolio into their local portfolios for deployment, Control end-user access in the recipient account, and Manage local constraints or permissions. However, the recipient cannot edit, delete, or reconfigure the shared products (Options B, C, and D). The source (owner) account retains full administrative control over products, launch roles, and lifecycle policies. This model aligns with AWS CloudOps principles of centralized governance with distributed self- service deployment across multiple accounts. Thus, Option A is correct--imported portfolios allow the recipient to add products to a local portfolio but not alter the shared configuration.QUESTION 22A company is migrating a legacy application to AWS. The application runs on EC2 instances across multiple Availability Zones behind an Application Load Balancer (ALB). The target group routing algorithm is set to weighted random, and the application requires session affinity (sticky sessions). After deployment, users report random application errors that were not present before migration, even though target health checks are passing. Which solution will meet this requirement? A. Set the routing algorithm of the target group to least outstanding requests.B. Turn on anomaly mitigation for the target group.C. Turn off the cross-zone load balancing attribute of the target group.D. Increase the deregistration delay attribute of the target group.Answer: AExplanation: According to the AWS Cloud Operations and Elastic Load Balancing documentation, Application Load Balancer (ALB) supports multiple routing algorithms to distribute requests among targets:Round robin (default)Least outstanding requests (LOR)Weighted randomWhen applications require session affinity, AWS recommends using "least outstanding requests" as the load balancing algorithm because it reduces latency, distributes load evenly, and ensures consistent target responsiveness during high traffic. Using weighted random routing with sticky sessions can cause sessions to be routed inconsistently if one target's capacity fluctuates, leading to session mismatches and application errors -- especially when user sessions rely on instance-specific state. Disabling cross-zone balancing (Option C) or adjusting deregistration delay (Option D) does not address routing inconsistency. Anomaly mitigation (Option B) protects against target performance degradation, not sticky-session misrouting. Therefore, the correct solution is Option A -- changing the target group's routing algorithm to least outstanding requests ensures smoother, predictable session handling and resolves random application errors.QUESTION 23A CloudOps engineer must manage the security of an AWS account. Recently, an IAM user's access key was mistakenly uploaded to a public code repository. The engineer must identify everything that was changed using this compromised key. How should the CloudOps engineer meet these requirements? A. Create an Amazon EventBridge rule to send all IAM events to an AWS Lambda function for analysis.B. Ouery Amazon EC2 logs by using Amazon CloudWatch Logs Insights for all events initiated with the compromised access key within the suspected timeframe. C. Search AWS CloudTrail event history for all events initiated with the compromised access key within the suspected timeframe. D. Search VPC Flow Logs for all events initiated with the compromised access key within the suspected timeframe. Answer: C Explanation: According to the AWS Cloud Operations and Security documentation, AWS CloudTrail is the authoritative service for recording API activity across all AWS services within an account. When an access key is compromised, CloudTrail logs all API requests made using that key, including details such as: The user identity (access key ID) that made the request, The service, operation, resource, and timestamp affected, and The source IP address and region of the request. By searching the CloudTrail event history for the specific access key ID, the CloudOps engineer can identify every action performed by that key during the suspected breach window.QUESTION 24A company runs custom statistical analysis software on a cluster of Amazon EC2 instances. The software is highly sensitive to network latency between nodes, although network throughput is not a limitation. Which solution will minimize network latency?A. Place all the EC2 instances into a cluster placement group.B. Configure and assign two Elastic IP addresses for each EC2 instance.C. Configure jumbo frames on all the EC2 instances in the cluster.D. Place all the EC2 instances into a spread placement group in the same AWS Region. Answer: AExplanation: The AWS Cloud Operations and Compute documentation explains that placement groups control how EC2 instances are physically arranged within AWS data centers to optimize network performance. Cluster placement groups place instances physically close together within a single Availability Zone, connected through high-bandwidth, low-latency networking (ideal for tightly coupled, HPC, or distributed workloads). Spread placement groups distribute instances across distinct racks or Availability Zones for fault tolerance, increasing latency. Partition placement groups separate instances into partitions for isolation, not latency reduction. Therefore, to minimize latency for workloads such as computational clusters, the CloudOps engineer should use a cluster placement group. This placement ensures single-digit microsecond latency and enhanced packet rate performance between instances. Elastic IPs (Option B) do not influence internal networking. Jumbo frames (Option C) can marginally improve throughput but do not reduce propagation latency. Spread placement (Option D) increases distance, worsening latency. Hence, Option A -- using a cluster placement group -- delivers the lowest possible network latency and is AWS's best-practice design for HPC-style clusters.QUESTION 25A company is using AWS CloudTrail and

Customize the products in the imported portfolio. Answer: AExplanation: Per the AWS Cloud Operations and Service Catalog

wants to ensure that SysOps administrators can easily verify that the log files have not been deleted or changed. Which action should a SysOps administrator take to meet this requirement?A. Grant administrators access to the AWS Key Management Service (AWS KMS) key used to encrypt the log files.B. Enable CloudTrail log file integrity validation when the trail is created or updated.C. Turn on Amazon S3 server access logging for the bucket storing the log files.D. Configure the S3 bucket to replicate the log files to another bucket. Answer: BExplanation: Cloud Trail can produce digest files and sign log files to detect tampering. Enabling log file integrity validation ensures that any modification or deletion of delivered log files can be detected by verification against the digests and signatures, providing verifiable evidence of tampering or non-tampering for each log file. This is the standard mechanism used to assure SysOps that logs have not been altered after delivery.QUESTION 26A company runs several production workloads on Amazon EC2 instances. A SysOps administrator discovered that a production EC2 instance failed a system health check. The SysOps administrator recovered the instance manually. The SysOps administrator wants to automate the recovery task of EC2 instances and receive notifications whenever a system health check fails. Detailed monitoring is activated for all of the company's production EC2 instances. Which of the following is the MOST operationally efficient solution that meets these requirements? A. For each production EC2 instance, create an Amazon CloudWatch alarm for Status Check Failed: System. Set the alarm action to recover the EC2 instance. Configure the alarm notification to be published to an Amazon Simple Notification Service (Amazon SNS) topic.B. On each production EC2 instance, create a script that monitors the system health by sending a heartbeat notification every minute to a central monitoring server. If an EC2 instance fails to send a heartbeat, run a script on the monitoring server to stop and start the EC2 instance and to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic.C. On each production EC2 instance, create a script that sends network pings to a highly available endpoint by way of a cron job. If the script detects a network response timeout, invoke a command to reboot the EC2 instance.D. On each production EC2 instance, configure an Amazon CloudWatch agent to collect and send logs to a log group in Amazon CloudWatch Logs. Create a CloudWatch alarm that is based on a metric filter that tracks errors. Configure the alarm to invoke an AWS Lambda function to reboot the EC2 instance and send a notification email. Answer: AExplanation: EC2 status checks run every minute and expose a System check metric that indicates hardware or system-level impairment. CloudWatch alarms can trigger on Status Check Failed: System and automatically recover the instance (or take a defined action) without manual intervention, satisfying both the automatic recovery and alerting requirements. This approach minimizes manual steps and scales across many instances, leveraging built-in AWS health signals and standard notification channels. It is also consistent with AWS guidance on using CloudWatch alarms to recover impaired instances and to notify via SNS.OUESTION 27A company has an application that uses Amazon ElastiCache for Memcached to cache query responses to improve latency. However, the application's users are reporting slow response times. A SysOps administrator notices that the Amazon CloudWatch metrics for Memcached evictions are high. Which actions should the SysOps administrator take to fix this issue? (Choose two.)A. Flush the contents of ElastiCache for Memcached.B. Increase the ConnectionOverhead parameter value.C. Increase the number of nodes in the cluster.D. Increase the size of the nodes in the cluster.E. Decrease the number of nodes in the cluster.Answer: CDExplanation: High eviction counts in Memcached occur when the total cache capacity is insufficient for the working set. Increasing the node size provides more memory per node, allowing more data to be cached before evictions occur. This directly reduces eviction pressure on a per-node basis. Adding more nodes increases overall cluster memory and parallel cache capacity, allowing the aggregate cache to hold more items concurrently and reducing the likelihood that frequently accessed data is evicted. This scalability approach is standard practice when facing memory-bound eviction issues in ElastiCache Memcached clusters.QUESTION 28A company's ecommerce application is running on Amazon EC2 instances that are behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. Customers report that the website is occasionally down. When the website is down, it returns an HTTP 500 (server error) status code to customer browsers. The Auto Scaling group's health check is configured for EC2 status checks, and the instances appear healthy. Which solution will resolve the problem?A. Replace the ALB with a Network Load Balancer.B. Add Elastic Load Balancing (ELB) health checks to the Auto Scaling group.C. Update the target group configuration on the ALB. Enable session affinity (sticky sessions).D. Install the Amazon CloudWatch agent on all instances. Configure the agent to reboot the instances. Answer: BExplanation: In this scenario, the EC2 instances pass their EC2 status checks, indicating that the operating system is responsive. However, the application hosted on the instance is failing intermittently, returning HTTP 500 errors. This demonstrates a discrepancy between the instance-level health and the application-level health. Auto Scaling groups should incorporate Elastic Load Balancing (ELB) health checks instead of relying solely on EC2 status checks. The ELB health check probes the application endpoint (for example, HTTP or HTTPS target group health checks), ensuring that the application itself is functioning correctly. When an instance fails an ELB health check, Amazon EC2 Auto Scaling will automatically mark the instance as unhealthy and replace it with a new one, ensuring continuous availability and performance optimization."Implement monitoring and health checks using ALB and EC2 Auto Scaling integration.

Application Load Balancer health checks allow Auto Scaling to terminate and replace instances that fail application-level health checks, ensuring consistent application performance.""When you enable the ELB health check type for your Auto Scaling group, Amazon EC2 Auto Scaling considers both EC2 status checks and Elastic Load Balancing health checks to determine instance health. If an instance fails the ELB health check, it is automatically replaced."Therefore, the correct answer is B, as it ensures proper application-level monitoring and remediation using ALB-integrated ELB health checks--a core CloudOps operational practice for proactive incident response and availability assurance. OUESTION 29A 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: DExplanation: 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.QUESTION 30An Amazon EC2 instance is running an application that uses Amazon Simple Queue Service (Amazon SQS) queues. A CloudOps engineer must ensure that the application can read, write, and delete messages from the SQS queues. Which solution will meet these requirements in the MOST secure manner? A. Create an IAM user with an IAM policy that allows the sqs:SendMessage permission, the sqs:ReceiveMessage permission, and the sqs:DeleteMessage permission to the appropriate queues. Embed the IAM user's credentials in the application's configuration.B. Create an IAM user with an IAM policy that allows the sqs:SendMessage permission, the sqs:ReceiveMessage permission, and the sqs:DeleteMessage permission to the appropriate queues. Export the IAM user's access key and secret access key as environment variables on the EC2 instance.C. Create and associate an IAM role that allows EC2 instances to call AWS services. Attach an IAM policy to the role that allows sqs:* permissions to the appropriate queues.D. Create and associate an IAM role that allows EC2 instances to call AWS services. Attach an IAM policy to the role that allows the sqs:SendMessage permission, the sqs:ReceiveMessage permission, and the sqs:DeleteMessage permission to the appropriate queues. Answer: DExplanation: The most secure pattern is to use an IAM role for Amazon EC2 with the minimum required permissions. AWS guidance states: "Use roles for applications that run on Amazon EC2 instances" and "grant least privilege by allowing only the actions required to perform a task." By attaching a role to the instance, short-lived credentials are automatically provided through the instance metadata service; this removes the need to create long-term access keys or embed secrets. Granting only sqs:SendMessage, sqs:ReceiveMessage, and sqs:DeleteMessage against the specific SQS queues enforces least privilege and aligns with CloudOps security controls. Options A and B rely on IAM user access keys, which contravene best practices for workloads on EC2 and increase credential- management risk. Option C uses a role but grants sqs:*, violating least-privilege principles. Therefore, Option D meets the security requirement with scoped, temporary credentials and precise permissions.QUESTION 31A company runs an application that logs user data to an Amazon CloudWatch Logs log group. The company discovers that personal information the application has logged is visible in plain text in the CloudWatch logs. The company needs a solution to redact personal information in the logs by default. Unredacted information must be available only to the company's security team. Which solution will meet these requirements?A. Create an Amazon S3 bucket.

Create an export task from appropriate log groups in CloudWatch. Export the logs to the S3 bucket. Configure an Amazon Macie scan to discover personal data in the S3 bucket. Invoke an AWS Lambda function to move identified personal data to a second S3 bucket. Update the S3 bucket policies to grant only the security team access to both buckets. B. Create a customer managed AWS KMS key. Configure the KMS key policy to allow only the security team to perform decrypt operations. Associate the KMS key with the application log group.C. Create an Amazon CloudWatch data protection policy for the application log group. Configure data identifiers for the types of personal information that the application logs. Ensure that the security team has permission to call the unmask API operation on the application log group.D. Create an OpenSearch domain. Create an AWS Glue workflow that runs a Detect PII transform job and streams the output to the OpenSearch domain. Configure the CloudWatch log group to stream the logs to AWS Glue, Modify the OpenSearch domain access policy to allow only the security team to access the domain. Answer: C Explanation: CloudWatch Logs data protection provides native redaction/masking of sensitive data at ingestion and query. AWS documentation states it can "detect and protect sensitive data in logs" using data identifiers, and that authorized users can "use the unmask action to view the original data." Creating a data protection policy on the log group masks PII by default for all viewers, satisfying the requirement to redact personal information. Granting only the security team permission to invoke the unmask API operation ensures that unredacted content is restricted. Option B (KMS) encrypts at rest but does not redact fields; encryption alone does not prevent plaintext visibility to authorized readers. Options A and D add complexity and latency, move data out of CloudWatch, and do not provide default inline redaction/unmask controls in CloudWatch itself. Therefore, the CloudOps-aligned, managed solution is to use CloudWatch Logs data protection with appropriate data identifiers and unmask permissions limited to the security team.QUESTION 32A multinational company uses an organization in AWS Organizations to manage over 200 member accounts across multiple AWS Regions. The company must ensure that all AWS resources meet specific security requirements. The company must not deploy any EC2 instances in the ap-southeast-2 Region. The company must completely block root user actions in all member accounts. The company must prevent any user from deleting AWS CloudTrail logs, including administrators. The company requires a centrally managed solution that the company can automatically apply to all existing and future accounts. Which solution will meet these requirements?A. Create AWS Config rules with remediation actions in each account to detect policy violations. Implement IAM permissions boundaries for the account root users.B. Enable AWS Security Hub across the organization. Create custom security standards to enforce the security requirements. Use AWS CloudFormation StackSets to deploy the standards to all the accounts in the organization. Set up Security Hub automated remediation actions.C. Use AWS Control Tower for account governance. Configure Region deny controls, Use Service Control Policies (SCPs) to restrict root user access.D. Configure AWS Firewall Manager with security policies to meet the security requirements. Use an AWS Config aggregator with organization-wide conformance packs to detect security policy violations. Answer: CExplanation: AWS CloudOps governance best practices emphasize centralized account management and preventive guardrails. AWS Control Tower integrates directly with AWS Organizations and provides "Region deny controls" and "Service Control Policies (SCPs)" that apply automatically to all existing and newly created member accounts. SCPs are organization-wide guardrails that define the maximum permissions for accounts. They can explicitly deny actions such as launching EC2 instances in a specific Region, or block root user access. To prevent CloudTrail log deletion, SCPs can also include denies on cloudtrail:DeleteTrail ands3:DeleteObject actions targeting the CloudTrail log S3 bucket. These SCPs ensure that no user, including administrators, can violate the compliance requirements."Use AWS Control Tower to establish a secure, compliant, multi-account environment with preventive guardrails through service control policies and detective controls through AWS Config."This approach meets all stated needs: centralized enforcement, automatic propagation to new accounts, region-based restrictions, and immutable audit logs. Options A, B, and D either detect violations reactively or lack complete enforcement and automation across future accounts.QUESTION 33A company uses AWS Organizations to create and manage many AWS accounts. The company wants to deploy new IAM roles in each account. Which action should the SysOps administrator take to deploy the new roles in each of the organization's accounts?A. Create a service control policy (SCP) in the organization to add the new IAM roles to each account.B. Deploy an AWS CloudFormation change set to the organization with a template to create the new IAM roles.C. Use AWS CloudFormation StackSets to deploy a template to each account to create the new IAM roles.D. Use AWS Config to create an organization rule to add the new IAM roles to each account. Answer: C Explanation:StackSets enables you to create, update, or delete CloudFormation stacks across multiple AWS accounts and regions from a single administrator account, which is ideal for applying identical IAM role configurations across all accounts in an organization. This approach minimizes manual effort, ensures consistency, and scales with the number of accounts. It is the recommended pattern for cross-account resource provisioning within AWS Organizations.QUESTION 34A company has a microservice that runs on a set of Amazon EC2 instances. The EC2 instances run behind an Application Load Balancer (ALB). A CloudOps engineer must use Amazon Route 53 to create a record that maps the ALB URL to example.com. Which type of record

will meet this requirement?A. An A recordB. An AAAA recordC. An alias recordD. A CNAME recordAnswer: C Explanation: An alias record is the recommended Route 53 record type to map domain names (e.g., example.com) to AWS-managed resources such as an Application Load Balancer. Alias records are extension types of A or AAAA records that support AWS resources directly, providing automatic DNS integration and no additional query costs." Use alias records to map your domain or subdomain to an AWS resource such as an Application Load Balancer, CloudFront distribution, or S3 website endpoint." A and AAAA records are used for static IP addresses, not load balancers. CNAME records cannot be used at the root domain (e.g., example.com). Thus, Option C is correct as it meets CloudOps networking best practices for scalable, managed DNS resolution to ALBs.QUESTION 35Application A runs on Amazon EC2 instances behind a Network Load Balancer (NLB). The EC2 instances are in an Auto Scaling group and are in the same subnet that is associated with the NLB. Other applications from an on-premises environment cannot communicate with Application A on port 8080. To troubleshoot the issue, a CloudOps engineer analyzes the flow logs. The flow logs include the following records: ACCEPT from 192.168.0.13:59003 172.31.16.139:8080REJECT from 172.31.16.139:8080 192.168.0.13:59003What is the reason for the rejected traffic?A. The security group of the EC2 instances has no Allow rule for the traffic from the NLB.B. The security group of the NLB has no Allow rule for the traffic from the on-premises environment.C. The ACL of the on-premises environment does not allow traffic to the AWS environment.D. The network ACL that is associated with the subnet does not allow outbound traffic for the ephemeral port range. Answer: DExplanation: VPC Flow Logs show the request arriving and being ACCEPTed on dstport 8080 and the corresponding response being REJECTed on the return path to the client's ephemeral port (59003). AWS networking guidance states that security groups are stateful (return traffic is automatically allowed) while network ACLs are stateless and require explicit inbound and outbound rules for both directions. CloudOps operational guidance for VPC networking further notes that when you allow an inbound request (for example, TCP 8080) through a subnet's network ACL, you must also allow the outbound ephemeral port range (typically 1024?5535) for the response traffic; otherwise, the return packets are dropped and appear as REJECT in flow logs. The observed pattern--request accepted to 8080, response rejected to 59003--matches a missing outbound ephemeral-range allow on the subnet's NACL. Therefore, the cause is the subnet NACL, not security groups or on-premises ACLs. The remediation is to add an outbound ALLOW rule on the NACL for the appropriate ephemeral TCP port range back to the on-premises CIDR (and the corresponding inbound rule if asymmetric).QUESTION 36A company runs a website on Amazon EC2 instances. Users can upload images to an Amazon S3 bucket and publish the images to the website. The company wants to deploy a serverless image-processing application that uses an AWS Lambda function to resize the uploaded images. The company's development team has created the Lambda function. A CloudOps engineer must implement a solution to invoke the Lambda function when users upload new images to the S3 bucket. Which solution will meet this requirement?A. Configure an Amazon Simple Notification Service (Amazon SNS) topic to invoke the Lambda function when a user uploads a new image to the S3 bucket.B. Configure an Amazon CloudWatch alarm to invoke the Lambda function when a user uploads a new image to the S3 bucket.C. Configure S3 Event Notifications to invoke the Lambda function when a user uploads a new image to the S3 bucket.D. Configure an Amazon Simple Queue Service (Amazon SQS) queue to invoke the Lambda function when a user uploads a new image to the S3 bucket. Answer: CExplanation: Use Amazon S3 Event Notifications with AWS Lambda to trigger image processing on object creation. S3 natively supports invoking Lambda for events such as s3:ObjectCreated:*, providing a serverless, low-latency pipeline without managing additional services. AWS operational guidance states that "Amazon S3 can directly invoke a Lambda function in response to object-created events," allowing you to pass event metadata (bucket/key) to the function for resizing and writing results back to S3. This approach minimizes operational overhead, scales automatically with upload volume, and integrates with standard retry semantics. SNS or SQS can be added for fan-out or buffering patterns, but they are not required when the requirement is simply "invoke the Lambda function on upload." CloudWatch alarms do not detect individual S3 object uploads and cannot directly satisfy per-object triggers. Therefore, configuring S3 Lambda event notifications meets the requirement most directly and aligns with CloudOps best practices for event-driven, serverless automation.QUESTION 37A company hosts a production MySQL database on an Amazon Aurora single-node DB cluster. The database is queried heavily for reporting purposes. The DB cluster is experiencing periods of performance degradation because of high CPU utilization and maximum connections errors. A CloudOps engineer needs to improve the stability of the database. Which solution will meet these requirements? A. Create an Aurora Replica node. Create an Auto Scaling policy to scale replicas based on CPU utilization. Ensure that all reporting requests use the read-only connection string.B. Create a second Aurora MySQL single-node DB cluster in a second Availability Zone. Ensure that all reporting requests use the connection string for this additional node.C. Create an AWS Lambda function that caches reporting requests. Ensure that all reporting requests call the Lambda function.D. Create a multi-node Amazon ElastiCache cluster. Ensure that all reporting requests use the ElastiCache cluster. Use the database if the data is not in the cache. Answer: AExplanation: Amazon Aurora supports up to 15 Aurora Replicas

that share the same storage volume and provide read scaling and improved availability. Official guidance states that replicas "offload read traffic from the writer" and that you should direct read-only workloads to the reader endpoint, reducing CPU pressure and connection counts on the primary. Aurora also supports Replica Auto Scaling through Application Auto Scaling policies using metrics such as CPU utilization or connections to add or remove replicas automatically. This design addresses both high CPU and maximum connections by moving reporting traffic to read replicas while keeping a single write primary for OLTP. Option B creates a separate cluster with independent storage, increasing operational overhead and data synchronization complexity. Options C and D introduce application-layer caching changes that may not guarantee data freshness or relieve the write node directly. Therefore, adding read replicas and routing reporting to the reader endpoint, with auto scaling based on load, is the least intrusive, CloudOps-aligned way to stabilize performance.QUESTION 38A CloudOps engineer configures an application to run on Amazon EC2 instances behind an Application Load Balancer (ALB) in a simple scaling Auto Scaling group with the default settings. The Auto Scaling group is configured to use the RequestCountPerTarget metric for scaling. The CloudOps engineer notices that the RequestCountPerTarget metric exceeded the specified limit twice in 180 seconds. How will the number of EC2 instances in this Auto Scaling group be affected in this scenario?A. The Auto Scaling group will launch an additional EC2 instance every time the RequestCountPerTarget metric exceeds the predefined limit.B. The Auto Scaling group will launch one EC2 instance and will wait for the default cooldown period before launching another instance.C. The Auto Scaling group will send an alert to the ALB to rebalance the traffic and not add new EC2 instances until the load is normalized.D. The Auto Scaling group will try to distribute the traffic among all EC2 instances before launching another instance. Answer: BExplanation: With simple scaling policies, an Auto Scaling group performs one scaling activity when the alarm condition is met, then observes a default cooldown period (300 seconds) before another scaling activity of the same type can begin. CloudOps guidance explains that cooldown prevents rapid successive scale-outs by allowing time for the newly launched instance(s) to register with the load balancer and impact the metric. Even if the alarm breaches multiple times during the cooldown window, the group waits until the cooldown completes before evaluating and acting again. In this case, although RequestCountPerTarget exceeded the threshold twice within 180 seconds, the group will launch a single instance and then wait for cooldown before any additional scale-out can occur. Options A, C, and D do not reflect the behavior of simple scaling with cooldowns; A describes step/target-tracking-like behavior, and C/D are not Auto Scaling mechanics.QUESTION 39A company uses Amazon ElastiCache (Redis OSS) to cache application data. A CloudOps engineer must implement a solution to increase the resilience of the cache. The solution also must minimize the recovery time objective (RTO). Which solution will meet these requirements? A. Replace ElastiCache (Redis OSS) with ElastiCache (Memcached). B. Create an Amazon EventBridge rule to initiate a backup every hour. Restore the backup when necessary.C. Create a read replica in a second Availability Zone. Enable Multi-AZ for the ElastiCache (Redis OSS) replication group.D. Enable automatic backups. Restore the backups when necessary. Answer: CExplanation: For high availability and fast failover, ElastiCache for Redis supports replication groups with Multi-AZ and automatic failover. CloudOps guidance states that a primary node can be paired with one or more replicas across multiple Availability Zones; if the primary fails, Redis automatically promotes a replica to primary in seconds, thereby minimizing RTO. This architecture maintains in-memory data continuity without waiting for backup restore operations. Backups (Options B and D) provide durability but require restore and re-warm procedures that increase RTO and may impact application latency. Switching engines (Option A) to Memcached does not provide Redis replication/failover semantics and would not inherently improve resilience for this use case. Therefore, creating a read replica in a different AZ and enabling Multi-AZ with automatic failover is the prescribed CloudOps pattern to increase resilience and achieve the lowest practical RTO for Redis caches.QUESTION 40An AWS CloudFormation template creates an Amazon RDS instance. This template is used to build up development environments as needed and then delete the stack when the environment is no longer required. The RDS-persisted data must be retained for further use, even after the CloudFormation stack is deleted. How can this be achieved in a reliable and efficient way?A. Write a script to continue backing up the RDS instance every five minutes.B. Create an AWS Lambda function to take a snapshot of the RDS instance, and manually invoke the function before deleting the stack.C. Use the Snapshot Deletion Policy in the CloudFormation template definition of the RDS instance.D. Create a new CloudFormation template to perform backups of the RDS instance, and run this template before deleting the stack. Answer: CExplanation: AWS CloudFormation supports the DeletionPolicy attribute to control what happens to a resource when a stack is deleted. For Amazon RDS DB instances, setting DeletionPolicy: Snapshot instructs CloudFormation to retain a final DB snapshot automatically at stack deletion. CloudOps best practice recommends using this native mechanism for data retention and auditability, avoiding manual scripts or out-of-band processes. Options A, B, and D introduce operational overhead and potential human error. With DeletionPolicy set to Snapshot, the environment can be repeatedly created and torn down while preserving data states for later restoration with minimal manual steps. This aligns with IaC principles--declarative, repeatable, and reliable--and supports efficient lifecycle management of ephemeral

development stacks. Resources From: 1.2025 Latest Braindump2go SOA-C03 Exam Dumps (PDF & VCE) Free Share: https://www.braindump2go.com/soa-c03.html2.2025 Latest Braindump2go SOA-C03 PDF and SOA-C03 VCE Dumps Free Share: https://drive.google.com/drive/folders/1TzvKjcQz5M-E9pY9BMRFzhgzv6IrI2zB?usp=sharing3.2025 Free Braindump2go SOA-C03 Exam Questions Download:

https://www.braindump2go.com/free-online-pdf/SOA-C03-PDF-Dumps(1-40).pdfFree Resources from Braindump2go,We Devoted to Helping You 100% Pass All Exams!