



AWS Certified CloudOps EngineerAssociate

Amazon Web Services SOA-C03

Version Demo

Total Demo Questions: 23

Total Premium Questions: 239

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

PASSQUEEN.COM

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Monitoring, Logging, and Troubleshooting	35
Topic 2, Reliability and Business Continuity	36
Topic 3, Deployment, Provisioning, and Automation	43
Topic 4, Security and Compliance	51
Topic 5, Networking and Content Delivery	48
Topic 6, Cost and Performance Optimization	25
Topic 7, Mix Questions	1
Total	239

QUESTION NO: 1

A company uses memory-optimized Amazon EC2 instances behind a Network Load Balancer (NLB) to run an application. The company launched the EC2 instances from an AWS-provided Red Hat Enterprise Linux (RHEL) AMI.

A CloudOps engineer must monitor RAM utilization in 5-minute intervals. The CloudOps engineer must ensure that the EC2 instances scale in and out appropriately based on incoming load.

Which solution will meet these requirements?

- A.** Configure detailed monitoring for the EC2 instances. Configure the Amazon CloudWatch agent on the EC2 instances. Create an EC2 Auto Scaling group and Auto Scaling policy that is based on the **mem_active** metric.
- B.** Configure detailed monitoring for the EC2 instances. Use the **mem_used_percent** metric that the detailed monitoring feature provides. Create an IAM role that allows the CloudWatch agent to upload data. Create an EC2 Auto Scaling group and Auto Scaling policy that is based on the **mem_used_percent** metric.
- C.** Configure basic monitoring for the EC2 instances. Configure the Amazon CloudWatch agent on the EC2 instances. Create an IAM role that allows the CloudWatch agent to upload data. Create an EC2 Auto Scaling group and Auto Scaling policy that is based on the **mem_used_percent** metric.
- D.** Configure basic monitoring for the EC2 instances. Use the standard **mem_used_percent** metric for monitoring. Create an EC2 Auto Scaling group and Auto Scaling policy that is based on the **mem_used_percent** metric.

ANSWER: C

Explanation:

Configure basic monitoring for the EC2 instances. Configure the Amazon CloudWatch agent on the EC2 instances. Create an IAM role that allows the CloudWatch agent to upload data. Create an EC2 Auto Scaling group and Auto Scaling policy that is based on the **mem_used_percent** metric. This solution meets the monitoring and scaling requirements because Amazon EC2 does not automatically publish guest operating system memory utilization as a native Amazon CloudWatch metric. The CloudWatch agent runs on each RHEL instance, collects memory data from the operating system, and publishes it to CloudWatch as a custom metric. The agent's memory measurements include **mem_used_percent**, which provides a percentage-based RAM utilization value suitable for an alarm and an Auto Scaling policy.

Basic EC2 monitoring publishes EC2 service metrics at five-minute intervals, matching the requested interval. The CloudWatch agent can likewise be configured with a 300-second collection interval for the memory metric. An instance profile role must grant permissions for the agent to publish CloudWatch metrics; AWS commonly provides the **CloudWatchAgentServerPolicy** managed policy for this purpose. The Auto Scaling group can then use the custom memory metric in a target tracking or step scaling policy, allowing capacity to change in response to sustained memory demand. See [CloudWatch agent metrics collected](#) and [Amazon EC2 Auto Scaling target tracking](#).

QUESTION NO: 2

A company has a non-production application that runs on an Amazon EC2 instance. The Amazon CloudWatch agent is installed on the EC2 instance. The application includes a process that randomly overuses temporary disk space and fills disks to 100% capacity.

A CloudOps engineer needs to automate a reboot of the EC2 instance after the disks reach 100% capacity.

Which solution will meet this requirement in the MOST operationally efficient way?

- A.** Create a CloudWatch alarm for the EC2 instance. Create an Amazon EventBridge event rule that reacts to the CloudWatch alarm and reboots the EC2 instance.
- B.** Create a CloudWatch alarm for the EC2 instance. Create an Amazon SES notification that reacts to the CloudWatch alarm and reboots the EC2 instance.
- C.** Create an AWS Lambda function to reboot the EC2 instance. Create a CloudWatch alarm that uses Amazon EventBridge to invoke the Lambda function.
- D.** Create an AWS Lambda function to reboot the EC2 instance. Use EC2 health checks to invoke the Lambda function.

ANSWER: A

Explanation:

Create a CloudWatch alarm for the EC2 instance. Create an Amazon EventBridge event rule that reacts to the CloudWatch alarm and reboots the EC2 instance. This is the most operationally efficient solution because the CloudWatch agent collects operating-system-level disk utilization and publishes it to CloudWatch as a metric. An alarm can evaluate that metric and enter the `ALARM` state when the disk reaches 100% utilization. CloudWatch alarm state changes are emitted as events, which Amazon EventBridge can match with a rule. The rule can then invoke the Amazon EC2 `RebootInstances` API operation for the affected instance through an AWS service target, without maintaining custom application code or infrastructure.

This event-driven design provides automated remediation immediately after the threshold breach, while retaining CloudWatch alarm history and EventBridge event visibility for operational auditing. The EventBridge execution role should be restricted to the specific instance and the `ec2:RebootInstances` action, following least-privilege practices. Because the workload is non-production and the requirement explicitly calls for a reboot after full disk utilization is detected, this direct alarm-to-event-to-EC2 remediation path is appropriate. See the [CloudWatch agent metrics documentation](#) and [Amazon EventBridge targets documentation](#).

QUESTION NO: 3

A company's developers manually install software modules on **Amazon EC2 instances** to deploy new versions of a service. A security audit finds that instances contain **inconsistent and unapproved modules**.

A CloudOps engineer must create a new **instance image** that contains **only approved software**.

Which solution will meet these requirements?

- A.** Use Amazon Detective to continuously find and uninstall unauthorized modules from the instances.
- B.** Use Amazon GuardDuty to create and deploy an Amazon Machine Image (AMI) that includes only the approved modules.
- C.** Use AWS Systems Manager Run Command to install the approved modules on all running instances during an in-place update.
- D.** Use EC2 Image Builder to create and test an Amazon Machine Image (AMI) that includes only the approved modules. Update the deployment workflow to use the new AMI.

ANSWER: D

Explanation:

Use EC2 Image Builder to create and test an Amazon Machine Image (AMI) that includes only the approved modules. Update the deployment workflow to use the new AMI. is correct because EC2 Image Builder is a managed AWS service specifically designed to automate the creation, customization, validation, and distribution of hardened machine images. The CloudOps engineer can define an image recipe containing the approved software components and configuration steps, then use an image pipeline to build a repeatable AMI. This establishes an approved, version-controlled baseline rather than relying on developers to make manual changes to running instances.

Image Builder also supports testing during the image build process. Test components can verify that required software is present, expected services and configurations are functional, and the resulting image meets the organization's standards before it is released. The pipeline can produce new AMI versions when its recipe, components, or source image changes. Updating the deployment workflow to launch instances from the validated AMI ensures that new service deployments begin from the consistent approved software baseline required by the audit. AWS documents Image Builder as a service for automating the build, test, and deployment of secure images: [EC2 Image Builder User Guide](#). For AMI lifecycle and usage details, see [Amazon Machine Images \(AMIs\)](#).

QUESTION NO: 4

A CloudOps engineer is maintaining a web application that uses an Amazon CloudFront web distribution, an Application Load Balancer (ALB), Amazon RDS, and Amazon EC2 in a VPC. All services have logging enabled. The CloudOps engineer needs to investigate HTTP Layer 7 status codes from the web application.

Which log sources contain the status codes? (Select TWO.)

- A. VPC Flow Logs
- B. AWS CloudTrail logs
- C. ALB access logs
- D. CloudFront access logs
- E. RDS logs

ANSWER: C D

Explanation:

ALB access logs and **CloudFront access logs** contain the HTTP status information needed to investigate Layer 7 behavior. Application Load Balancers operate at the application layer and their access-log entries include both the `elb_status_code`, which is the response status code sent by the load balancer to the client, and the `target_status_code`, which is the status code returned by the registered target when available. This enables investigation of client-facing outcomes and responses from the EC2-hosted application.

CloudFront standard access logs record each viewer request handled by the distribution. Their log fields include the HTTP status code returned to the viewer, allowing the engineer to analyze the client experience at CloudFront edge locations. Together, these sources provide visibility into responses delivered by the CDN and the load balancer/application tier, including status codes such as 200, 404, 403, and 5xx responses. AWS documents the ALB access-log status-code fields at [Application Load Balancer access logs](#) and the CloudFront status field at [CloudFront standard access logs](#).

QUESTION NO: 5

A company uses default settings to create an AWS Lambda function. The function needs to access an Amazon RDS database that is in a private subnet of a VPC

C.

The function has the correct IAM permissions to access the database. The private subnet has appropriate routing configurations and is accessible from within the VPC

C.

However, the Lambda function is unable to connect to the RDS instance.

What is the likely reason the Lambda function cannot connect to the RDS instance?

- A. The company did not set the RDS instance as the destination for the Lambda function in the function configuration.
- B. The Lambda function configuration did not deploy the function in the same VPC that contains the RDS instance.
- C. The VPC where the Lambda function is deployed is not peered with the VPC where the RDS instance is deployed.
- D. The security group for the Lambda function does not allow outbound access to the RDS instance.

ANSWER: B

Explanation:

The Lambda function configuration did not deploy the function in the same VPC that contains the RDS instance. Lambda functions created with the default network configuration are not attached to a customer VPC. Although Lambda can invoke AWS services through AWS-managed networking, it does not automatically have private network connectivity to resources

such as an RDS DB instance located only in a private subnet of a customer VPC. IAM permissions authorize AWS API actions, but they do not establish the TCP/IP network path required for a database connection.

To access the private RDS instance, the function must be configured for VPC access. This configuration specifies subnets and security groups for Lambda, and Lambda creates and manages Hyperplane elastic network interfaces to provide connectivity into that VPC. The selected subnets must have network reachability to the DB subnet, and the Lambda security group must permit the required database traffic; likewise, the RDS security group must allow inbound traffic from the Lambda security group on the database port. AWS documents this VPC attachment requirement for connecting Lambda to resources in a VPC, including private databases. See [Configuring a Lambda function to access resources in a VPC](#) and [Working with a DB instance in a VPC](#).

QUESTION NO: 6

A company runs a three-tier web application on AWS. The application includes web servers, application servers, and database servers. The application servers process requests from the web servers. The company wants to ensure high availability of the application. Therefore, the company needs to monitor the health of the application servers and route traffic only to healthy instances.

Which solution will meet these requirements?

- A. Create an Application Load Balancer (ALB) in front of the application servers with health checks for the application servers.
- B. Create an Amazon Route 53 health check for the application servers. Attach a Network Load Balancer (NLB) in front of the application servers.
- C. Create an AWS Lambda function that restarts an application server. Configure an Amazon CloudWatch alarm to monitor the health of the application servers. Run the function when an application is unhealthy.
- D. Create an Amazon CloudWatch metric to monitor the health of the application servers. Route traffic by using a Network Load Balancer (NLB).

ANSWER: A

Explanation:

Create an Application Load Balancer (ALB) in front of the application servers with health checks for the application servers. An ALB distributes incoming application-layer HTTP and HTTPS requests across registered targets, such as Amazon EC2 instances, based on listener and target-group rules. A target group includes health-check settings that cause the ALB to periodically send requests to each application server's health-check endpoint. When a target fails the configured health-check threshold, the load balancer marks it unhealthy and stops routing new requests to that target. Once the server again meets the healthy threshold, it can resume receiving traffic automatically. This directly delivers both required capabilities: continuous application-server health monitoring and traffic routing only to healthy servers. Deploying the application servers across multiple Availability Zones and registering them in the target group further supports a highly available architecture because the ALB can distribute traffic among healthy targets in those zones. Health-check paths, ports, intervals, timeouts, and success codes can be configured to reflect the application's actual readiness criteria. See the [Application Load Balancer target group health checks documentation](#) and the [Application Load Balancer guide](#).

QUESTION NO: 7

A company hosts a **static website** in **Amazon S3** behind an **Amazon CloudFront distribution**. When new versions are deployed, users sometimes do not see updated content immediately.

Which solution will meet this requirement?

- A. Configure the CloudFront distribution to add a custom Cache-Control header to requests for content from the S3 bucket.
- B. Modify the distribution settings to specify the protocol as HTTPS only.
- C. Attach the CachingOptimized managed cache policy to the distribution.
- D. Create a CloudFront invalidation.

ANSWER: D

Explanation:

Create a CloudFront invalidation. is the appropriate action because CloudFront stores copies of S3 origin objects in its edge caches. A deployment can successfully replace an object in the S3 bucket while CloudFront edge locations continue serving the previously cached object until its configured cache lifetime expires. An invalidation removes the specified existing objects from CloudFront edge caches. On the next viewer request after the invalidation has completed, CloudFront retrieves the current object version from the S3 origin and caches that version according to the distribution's cache behavior and origin response headers.

The invalidation path should include every changed object that needs an immediate refresh. For example, a deployment can invalidate `/index.html` and other changed asset paths, or use `/*` when the deployment requires all cached website objects to be refreshed. Invalidation can be created in the CloudFront console, through the AWS CLI or SDK, or integrated into the deployment pipeline so cache refresh occurs consistently after publishing the updated site. AWS documents invalidations as the mechanism for removing objects from CloudFront edge caches before their normal expiration time. See [Invalidating files in the CloudFront Developer Guide](#) and [Managing how long content stays in the cache](#).

QUESTION NO: 8

A company runs a database on an Amazon EC2 instance with an Amazon EBS General Purpose SSD (gp2) volume. During periods of sustained database activity, the application experiences increasing I/O latency. Amazon CloudWatch shows that the volume's BurstBalance metric decreases to 0% during these periods. The workload requires IOPS and throughput that exceed the baseline performance of the current volume.

Which combination of actions will resolve the performance issue? (Select TWO.)

- A. Modify the EBS volume from gp2 to gp3.
- B. Configure the gp3 volume with provisioned IOPS and throughput that meet the database workload requirements.
- C. Increase the EBS snapshot frequency so that the volume can replenish burst credits more quickly.
- D. Create an Amazon CloudWatch alarm that reboots the EC2 instance when BurstBalance reaches 0%.
- E. Change the EBS volume type to Magnetic.

ANSWER: A B

Explanation:

Modify the EBS volume from gp2 to gp3. Unlike gp2 volumes, gp3 volumes do not use burst credits to provide their baseline performance. This removes the dependency on BurstBalance that is causing the sustained-workload performance degradation.

Configure the gp3 volume with provisioned IOPS and throughput that meet the measured workload requirements. gp3 allows IOPS and throughput to be configured independently of volume size within supported limits. Increasing the EC2 instance size alone does not address an EBS volume that is constrained by depleted gp2 burst credits.

QUESTION NO: 9

A CloudOps engineer wants to provide access to AWS services by attaching an IAM policy to multiple IAM users. The CloudOps engineer also wants to be able to change the policy and create new versions.

Which combination of actions will meet these requirements? (Select TWO.)

- A. Add the users to an IAM service-linked role. Attach the policy to the role.
- B. Add the users to an IAM user group. Attach the policy to the group.
- C. Create an AWS managed policy.

D. Create a customer managed policy.

E. Create an inline policy.

ANSWER: B D

Explanation:

“Add the users to an IAM user group. Attach the policy to the group.” meets the requirement to grant a common set of permissions to multiple IAM users through one centrally managed assignment point. IAM groups are collections of users, and permissions granted through a policy attached to the group apply to each group member. This lets the CloudOps engineer manage membership separately from the shared permissions configuration; users added later receive the group’s permissions automatically.

“Create a customer managed policy.” meets the requirements for policy reuse, modification, and version management. Customer managed policies are standalone IAM resources that can be attached to multiple identities, including IAM groups. Unlike an inline policy, the same policy can be reused without duplicating its JSON document. IAM supports policy versioning for customer managed policies: changes can be saved as a new version, one version is designated as the default version, and IAM retains up to five versions at a time. Updating the default version changes the permissions effective for the group and therefore its members. This combination provides centralized, scalable access management while preserving controlled policy updates. See the [IAM managed policies documentation](#) and the [IAM policy versioning documentation](#).

QUESTION NO: 10

A company manages a set of AWS accounts by using AWS Organizations. The company’s security team wants to use a native AWS service to regularly scan all AWS accounts against the Center for Internet Security (CIS) AWS Foundations Benchmark.

What is the MOST operationally efficient way to meet these requirements?

A. Designate a central security account as the AWS Security Hub administrator account. Use scripts to invite and accept member accounts.

B. Run the CIS AWS Foundations Benchmark by using Amazon Inspector.

C. Designate a central security account as the Amazon GuardDuty administrator account and configure CIS scans.

D. Designate an AWS Security Hub administrator account, automatically enroll new organization accounts, and enable CIS AWS Foundations Benchmark.

ANSWER: D

Explanation:

Designating an AWS Security Hub administrator account, automatically enrolling new organization accounts, and enabling the CIS AWS Foundations Benchmark provides centralized, continuous standards-based security checks with minimal account-by-account administration. AWS Security Hub includes the CIS AWS Foundations Benchmark security standard, which evaluates AWS configurations against CIS recommendations and generates findings for controls that do not meet the benchmark requirements. The delegated Security Hub administrator can manage standards and findings across the organization from a central security account.

With AWS Organizations integration and automatic account enrollment enabled, Security Hub can automatically add eligible new organization accounts as they are created. This removes the operational overhead of manually inviting accounts, accepting invitations, and maintaining custom onboarding scripts. Security Hub then aggregates security findings centrally, giving the security team an organization-wide view of benchmark compliance and a consistent basis for remediation workflows. The CIS standard must be enabled in each applicable Region because Security Hub standards and findings are Regional. For implementation details, see the [AWS Security Hub CIS AWS Foundations Benchmark documentation](#) and [automatically enabling Security Hub for organization accounts](#).

QUESTION NO: 11

A company uses AWS Organizations to manage its AWS accounts. A CloudOps engineer must create a backup strategy for all Amazon EC2 instances across all the company's AWS accounts.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Deploy an AWS Lambda function to each account to run EC2 instance snapshots on a scheduled basis.
- B. Create an AWS CloudFormation stack set in the management account to add an AutoBackup=True tag to every EC2 instance.
- C. Use AWS Backup in the management account to deploy policies for all accounts and resources.
- D. Use a service control policy (SCP) to run EC2 instance snapshots on a scheduled basis in each account.

ANSWER: C

Explanation:

Use AWS Backup in the management account to deploy policies for all accounts and resources. AWS Backup integrates with AWS Organizations to provide centrally governed, organization-wide data protection. After enabling trusted access for AWS Backup, the organization's management account or a delegated administrator can create and attach backup policies to the organization root, organizational units, or individual member accounts. These policies define backup plans, including schedules, backup vaults, lifecycle settings, and resource selections based on tags. For EC2 workloads, AWS Backup creates and manages the underlying Amazon EBS snapshots required to protect supported instances and volumes.

This approach provides a consistent backup standard across all accounts while allowing inherited policy settings to be centrally maintained as accounts and organizational units evolve. It also supplies centralized reporting, monitoring, auditability, and compliance-oriented controls, avoiding per-account operational work. Resource assignments can be defined through tags, enabling newly tagged EC2 resources to be included automatically in the organization's backup strategy. AWS documents AWS Backup policies as the mechanism for centrally applying backup plans across an organization, and AWS Backup as a fully managed centralized service for automating data protection. See [AWS Backup policies in AWS Organizations](#) and [What is AWS Backup?](#).

QUESTION NO: 12

A company uses AWS CloudFormation to manage a stack of Amazon EC2 instances on AWS. A CloudOps engineer needs to keep the instances and all of the instances' data, even if someone deletes the stack.

Which solution will meet these requirements?

- A. Set the DeletionPolicy attribute to Snapshot for the EC2 instance resource in the CloudFormation template.
- B. Automate backups by using Amazon Data Lifecycle Manager (Amazon DLM).
- C. Create a backup plan in AWS Backup.
- D. Set the DeletionPolicy attribute to Retain for the EC2 instance resource in the CloudFormation template.

ANSWER: D

Explanation:

Set the DeletionPolicy attribute to Retain for the EC2 instance resource in the CloudFormation template. A CloudFormation `DeletionPolicy` controls what CloudFormation does with a resource when the stack is deleted or when a resource is removed from the template. With `Retain`, CloudFormation removes the EC2 instance from stack management but does not terminate the physical instance. Because the instance continues running or remains stopped rather than being terminated, its attached instance storage configuration and EBS-backed data remain available with the retained instance. This directly satisfies the requirement to preserve both the EC2 instances and their data even if the stack is deleted. The retained resources can subsequently be managed outside the deleted stack or imported into another CloudFormation stack if appropriate. AWS documents that the Retain policy preserves the resource or content when the stack is deleted, while CloudFormation proceeds with deletion of the remainder of the stack. See the [AWS CloudFormation DeletionPolicy reference](#) and the [AWS::EC2::Instance resource reference](#).

QUESTION NO: 13

A 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: C

Explanation:

Create a read replica in a second Availability Zone. Enable Multi-AZ for the ElastiCache (Redis OSS) replication group. This configuration provides high availability by maintaining a replica of the primary cache node in another Availability Zone and configuring the replication group for automatic failover. If the primary node, its Availability Zone, or the underlying infrastructure becomes unavailable, ElastiCache promotes an available read replica to become the new primary. The promotion is managed by the service, so the application can resume writes against the new primary without a manual restore workflow.

This directly minimizes RTO because recovery is based on an already-running, in-memory replica rather than creating a new cache cluster and restoring data from a snapshot. Using a second Availability Zone also protects the cache tier from an Availability Zone-scoped failure, improving resilience beyond a single-node or single-AZ deployment. For Redis OSS replication groups, Multi-AZ with automatic failover is the AWS high-availability mechanism intended to maintain service continuity when the primary node fails. AWS documents the behavior and requirements for this configuration in the [ElastiCache automatic failover documentation](#) and describes replication groups and their replicas in the [ElastiCache Redis OSS replication documentation](#).

QUESTION NO: 14

A company uses AWS Organizations to manage multiple AWS accounts. The company must evaluate whether Amazon EBS volumes are encrypted in all current and future member accounts. The central security account must have a consolidated view of compliance results.

Which combination of actions will meet these requirements? (Select TWO.)

- A. Create an AWS Config organization rule that uses the AWS managed `encrypted-volumes` rule.
- B. Create an AWS Config aggregator in the central security account and authorize aggregation of data from the organization.
- C. Create an Amazon CloudWatch dashboard in the central security account that displays the EBS volume count from each account.
- D. Create an AWS CloudTrail organization trail and search the trail for `CreateVolume` API calls.
- E. Create an IAM service-linked role in every member account that denies all Amazon EBS API operations.

ANSWER: A B

Explanation:

Create an AWS Config organization rule that uses the AWS managed `encrypted-volumes` rule. An organization rule can be deployed from the organization management account or a delegated administrator account to member accounts. The rule evaluates whether attached EBS volumes are encrypted and can be configured to deploy automatically to accounts that join the organization.

Create an AWS Config aggregator in the central security account and authorize aggregation of data from the organization. An aggregator collects AWS Config configuration and compliance data from multiple accounts and Regions, providing the required centralized view. An aggregator does not itself enable AWS Config recording or create compliance rules, so it must be used with the organization rule rather than instead of the rule.

QUESTION NO: 15

A 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: B

Explanation:

Assign the **AmazonSSMManagedInstanceCore** managed policy to the EC2 instance profile for the Ubuntu instance. For an Amazon EC2 instance to be available as an AWS Systems Manager managed node and accept Session Manager connections, the SSM Agent must be able to authenticate to and communicate with Systems Manager services. AWS recommends attaching the **AmazonSSMManagedInstanceCore** AWS managed policy to the IAM role associated with the instance profile. This policy supplies the core permissions that SSM Agent needs, including permissions for Systems Manager messaging channels used to establish interactive sessions.

Because Session Manager does not require an inbound SSH or RDP rule, a public IP address, or an SSH key pair, it supports the company's policy of preventing direct remote-access protocols. The fact that users can connect to other instances makes an instance-specific configuration issue likely. Ensuring that this Ubuntu instance has an attached instance profile whose IAM role includes **AmazonSSMManagedInstanceCore** restores the required Systems Manager authorization, provided the agent is installed and the instance has network access to Systems Manager endpoints. See the [AWS Systems Manager instance permissions documentation](#) and the [AWS Session Manager User Guide](#).

QUESTION NO: 16

A CloudOps engineer is investigating why an application on an Amazon EC2 instance cannot connect to a database server on TCP port 5432. The engineer needs to determine whether network traffic is being rejected and must be able to investigate the results by using a managed query service.

Which combination of steps will meet these requirements? (Select TWO.)

- A. Create VPC Flow Logs for the EC2 instance network interface. Publish the flow logs to an Amazon CloudWatch Logs log group.
- B. Use CloudWatch Logs Insights to query flow-log records for destination port 5432 and an action of `REJECT`.
- C. Create an AWS CloudTrail trail and query the trail for rejected TCP packets on port 5432.
- D. Enable Amazon RDS Performance Insights and query the database wait events for rejected network packets.
- E. Create an Amazon CloudWatch dashboard that displays the EC2 instance CPU utilization metric.

ANSWER: A B

Explanation:

Create VPC Flow Logs for the network interface of the EC2 instance and publish the flow logs to an Amazon CloudWatch Logs log group. VPC Flow Logs record IP traffic to and from supported network interfaces and include an action field that identifies whether traffic was accepted or rejected. Publishing to CloudWatch Logs provides centralized, near-real-time log storage for the investigation.

Use CloudWatch Logs Insights to query the flow-log records for traffic with destination port 5432 and an action of REJECT. The query can also filter on the EC2 instance private IP address, the database server address, and the relevant time range. VPC Flow Logs do not identify the specific security group or network ACL rule that caused a rejection, but they provide evidence that traffic was rejected and the source, destination, ports, and protocol involved.

QUESTION NO: 17

A SysOps administrator needs to implement a solution that protects credentials for an Amazon RDS for MySQL DB instance. The solution must rotate the credentials automatically one time every week.

Which combination of steps will meet these requirements? (Select TWO.)

- A. Configure an RDS proxy to store the credentials.
- B. Add the credentials to AWS Secrets Manager.
- C. Add the credentials to AWS Systems Manager Parameter Store.
- D. Create an AWS Lambda function to rotate the credentials.
- E. Create an AWS Systems Manager Automation runbook to rotate the credentials.

ANSWER: B D

Explanation:

"Add the credentials to AWS Secrets Manager." is correct because AWS Secrets Manager is designed to securely store, manage, and retrieve sensitive values such as Amazon RDS database usernames and passwords. A secret can be encrypted with AWS Key Management Service, accessed through tightly scoped IAM permissions, and used by applications without placing database credentials in application code or configuration files. Secrets Manager provides direct support for Amazon RDS credentials, including RDS for MySQL.

"Create an AWS Lambda function to rotate the credentials." is correct because Secrets Manager uses a Lambda rotation function to perform the rotation workflow. For an RDS for MySQL secret, the rotation function changes the database password and updates the stored secret value so that the secret remains synchronized with the DB instance. The administrator configures rotation on the Secrets Manager secret and specifies a schedule of every seven days, satisfying the weekly automatic rotation requirement. AWS provides rotation-function templates for supported RDS database engines, which can reduce the effort to implement the required workflow. The Lambda function runs as part of Secrets Manager rotation rather than requiring manual password-change operations. See the [AWS Secrets Manager rotation documentation](#) and the [Lambda rotation function documentation](#).

QUESTION NO: 18

A CloudOps engineer needs to control access to groups of Amazon EC2 instances using AWS Systems Manager Session Manager. Specific tags on the EC2 instances have already been added.

Which additional actions should the CloudOps engineer take to control access? (Select TWO.)

- A. Attach an IAM policy to the users or groups that require access to the EC2 instances.
- B. Attach an IAM role to control access to the EC2 instances.
- C. Create a placement group for the EC2 instances and add a specific tag.
- D. Create a service account and attach it to the EC2 instances that need to be controlled.
- E. Create an IAM policy that grants access to any EC2 instances with a tag specified in the Condition element.

ANSWER: A E

Explanation:

“Attach an IAM policy to the users or groups that require access to the EC2 instances.” is required because Session Manager authorizes the identity that requests a session through IAM. The policy must allow the relevant Systems Manager API actions, principally `ssm:StartSession`, so the authorized users or groups can initiate sessions without needing inbound SSH or RDP access. Assigning this permission to the appropriate human identities supports centralized, least-privilege administration.

“Create an IAM policy that grants access to any EC2 instances with a tag specified in the Condition element.” provides the tag-based scope required by the scenario. A Session Manager policy can use an EC2 instance ARN as the resource for `ssm:StartSession` and an `ssm:resourceTag/tag-key` condition. For example, a condition can require an `Environment` tag with a designated value. IAM evaluates the target managed instance’s tags when a user starts a session, allowing sessions only to the intended group of instances. This enables access changes to be managed by applying or changing tags rather than maintaining separate policies for every individual instance.

AWS documents tag-based Session Manager access controls and policy examples in the [Session Manager restricted-access examples](#) and documents the [use of resource tags to restrict Session Manager access](#).

QUESTION NO: 19

A CloudOps engineer is configuring an Amazon CloudFront distribution to use an SSL/TLS certificate. The CloudOps engineer must ensure automatic certificate renewal.

Which combination of steps will meet this requirement? (Select TWO.)

- A. Use a certificate issued by AWS Certificate Manager (ACM).
- B. Use a certificate issued by a third-party certificate authority (CA).
- C. Configure CloudFront to automatically renew the certificate when the certificate expires.
- D. Configure email validation for the certificate.
- E. Configure DNS validation for the certificate.

ANSWER: A E

Explanation:

Use a certificate issued by AWS Certificate Manager (ACM) and configure DNS validation for the certificate. ACM manages the lifecycle of Amazon-issued public certificates, including renewing eligible certificates automatically before they expire. For use with a CloudFront distribution, the ACM certificate must be requested or imported in the US East (N. Virginia) Region, although the essential renewal requirement is that it is an Amazon-issued ACM certificate rather than an imported certificate.

DNS validation supports unattended renewal because ACM can continue to verify domain control by querying the CNAME validation record in the domain’s DNS configuration. The validation CNAME record must remain in place and resolvable through the certificate lifecycle. ACM uses this persistent DNS proof to renew the certificate without requiring a person to respond to an approval message. After ACM renews an eligible certificate associated with CloudFront, CloudFront can continue using the renewed certificate without the engineer manually replacing the certificate on the distribution.

AWS documents the managed renewal behavior for ACM-issued certificates and the requirement to retain DNS validation records for renewal. See the [AWS Certificate Manager managed renewal documentation](#) and the [ACM DNS validation documentation](#).

QUESTION NO: 20

A company needs to monitor its website 's availability to end users. The company needs a solution to provide an Amazon Simple Notification Service (Amazon SNS) notification if the website 's uptime decreases to less than 99%. The monitoring must provide an accurate view of the user experience on the website.

Which solution will meet these requirements?

- A.** Create an Amazon CloudWatch alarm that is based on the website's logs that are published to a CloudWatch Logs log group. Configure the alarm to publish an SNS notification if the number of HTTP 4xx and 5xx errors exceeds a specified threshold.
- B.** Create an Amazon CloudWatch alarm that is based on the website 's published metrics in CloudWatch. Configure the alarm to publish an SNS notification based on anomaly detection.
- C.** Create an Amazon CloudWatch Synthetics heartbeat monitoring canary. Associate the canary with the website's URL. Create a CloudWatch alarm for the canary. Configure the alarm to publish an SNS notification if the value of the SuccessPercent metric is less than 99%.
- D.** Create an Amazon CloudWatch Synthetics broken link checker monitoring canary. Associate the canary with the website's URL. Create a CloudWatch alarm for the canary. Configure the alarm to publish an SNS notification if the value of the SuccessPercent metric is less than 99%.

ANSWER: C

Explanation:

Create an Amazon CloudWatch Synthetics heartbeat monitoring canary. Associate the canary with the website's URL. Create a CloudWatch alarm for the canary. Configure the alarm to publish an SNS notification if the value of the SuccessPercent metric is less than 99%.

A CloudWatch Synthetics heartbeat canary is designed to continuously verify endpoint availability by making scheduled HTTP requests to the website URL. Because it performs active, external checks rather than relying solely on application-generated data, it measures whether the site can actually be reached and used from the canary's execution environment. This directly supports the requirement to monitor availability from an end-user-experience perspective.

Each canary run reports CloudWatch metrics in the `CloudWatchSynthetics` namespace. The `SuccessPercent` metric expresses the percentage of successful canary executions during the selected alarm evaluation period. A CloudWatch alarm can evaluate that metric against a threshold of 99 percent and invoke an Amazon SNS topic when the alarm enters the `ALARM` state. This gives the operations team an automated notification when measured website uptime falls below the required service level. AWS documents heartbeat canaries as the canary type for monitoring an endpoint's availability and responsiveness, and documents the canary success metrics and alarm integration at [CloudWatch Synthetics canaries](#) and [Synthetics metrics](#).

QUESTION NO: 21

A company uses AWS Organizations to manage multiple AWS accounts. The company must record management API activity from all current and future member accounts in a central logging account. The company must also protect the delivered log files from deletion or modification for 7 years, including from administrators in member accounts.

Which combination of actions will meet these requirements? (Select TWO.)

- A.** Create an AWS CloudTrail organization trail that delivers management events to an Amazon S3 bucket in the central logging account.
- B.** Create an individual AWS CloudTrail trail in each member account and deliver the logs to an S3 bucket in that same account.
- C.** Enable Amazon S3 Object Lock in compliance mode on the central logging bucket. Configure a retention period of 7 years.
- D.** Configure an S3 Lifecycle rule to expire the log objects after 7 years.
- E.** Configure Amazon CloudWatch Logs retention for 7 years in each member account.

ANSWER: A C

Explanation:

Create an AWS CloudTrail organization trail from the organization management account or a delegated administrator account. An organization trail applies across the organization and can automatically include newly created member accounts. The trail can deliver management events to an Amazon S3 bucket in the central logging account.

Enable Amazon S3 Object Lock in compliance mode on the destination bucket and configure a 7-year retention period for the log objects. Compliance mode prevents protected object versions from being deleted or overwritten before the retention period expires, including by users with administrative permissions. A trail that is separately created in each account would create unnecessary operational overhead and would not automatically cover future accounts.

QUESTION NO: 22

A company deploys AWS infrastructure in a VPC that has an internet gateway. The VPC has public subnets and private subnets. An Amazon RDS for MySQL DB instance is deployed in a private subnet. An AWS Lambda function uses the same private subnet and connects to the DB instance to query data.

A developer modifies the Lambda function to require the function to publish messages to an Amazon Simple Queue Service (Amazon SQS) queue. After these changes, the Lambda function times out when it tries to publish messages to the SQS queue.

Which solutions will resolve this issue? (Select TWO.)

- A. Reconfigure the Lambda function so that the function is not connected to the VPC.
- B. Deploy an RDS proxy. Configure the Lambda function to connect to the DB instance through the proxy.
- C. Deploy a NAT gateway. Update the private subnet's route table to route all traffic to the NAT gateway.
- D. Create an interface VPC endpoint for Amazon SQS in the VPC.
- E. Create a gateway endpoint for Amazon SQS in the VPC.

ANSWER: C D

Explanation:

A Lambda function attached to private subnets uses the networking and route tables of those subnets for outbound connections. Although the function can reach the private RDS DB instance, it needs an additional network path to reach the regional Amazon SQS service endpoint. Deploying a NAT gateway and updating the private subnet route table to send internet-bound traffic through that NAT gateway provides this path. The NAT gateway must reside in a public subnet and have a route to the internet gateway; the private-subnet route should direct the default route through the NAT gateway. This permits the function to access the public SQS endpoint while retaining access to the private database.

Creating an interface VPC endpoint for Amazon SQS also resolves the issue by providing private connectivity from the VPC to SQS through AWS PrivateLink. With private DNS enabled, calls made to the normal regional SQS endpoint resolve to the endpoint's private network interfaces, so traffic stays within the AWS network and does not require internet or NAT access. The endpoint security group must allow HTTPS traffic from the Lambda function's security group or subnet address range. AWS documents Lambda VPC internet access at [Lambda VPC internet access](#) and SQS interface endpoint support at [Amazon SQS VPC endpoints](#).

QUESTION NO: 23

A company has a stateful web application that is hosted on Amazon EC2 instances in an Auto Scaling group. The instances run behind an Application Load Balancer (ALB) that has a single target group. The ALB is configured as the origin in an Amazon CloudFront distribution. Users are reporting random logouts from the web application.

Which combination of actions should a CloudOps engineer take to resolve this problem? (Select TWO.)

- A. Change to the least outstanding requests algorithm on the ALB target group.
- B. Configure cookie forwarding in the CloudFront distribution cache behavior.
- C. Configure header forwarding in the CloudFront distribution cache behavior.

D. Enable group-level stickiness on the ALB listener rule.

E. Enable sticky sessions on the ALB target group.

ANSWER: B E

Explanation:

Configuring cookie forwarding in the CloudFront distribution cache behavior ensures that the viewer's session-related cookies are included in requests that CloudFront sends to the Application Load Balancer origin. For a stateful application, these cookies can carry the application session identifier and the load balancer's affinity information. CloudFront must be configured to forward the cookies required by the application, either through cache behavior cookie settings or the appropriate cache policy and origin request policy configuration. This preserves the information needed to associate subsequent requests with an existing user session.

Enabling sticky sessions on the ALB target group then provides session affinity between a client and a specific registered EC2 target. ALB target-group stickiness can use load-balancer-generated cookies or application-based cookies to bind requests to the target handling that user's stateful session. With CloudFront forwarding the relevant cookies and the target group honoring stickiness, requests from a logged-in user continue to reach the same instance for the duration of the configured affinity period. This prevents a request from being routed to another Auto Scaling group instance that does not have the user's in-memory session state. See the AWS documentation for [CloudFront cookie handling](#) and [ALB sticky sessions](#).