Number: SOA-C02
Passing Score: 800
Time Limit: 120
File Version: 31.0

# Exam Code: SOA-C02

# Exam Name: AWS Certified SysOps Administrator - Associate

**Exam A**

**QUESTION 1**
SIMULATION
You need to update an existing AWS CloudFormation stack. If needed, a copy to the CloudFormation template is available in an Amazon SB bucket named cloudformation-bucket
1. Use the us-east-2 Region for all resources.
2. Unless specified below, use the default configuration settings.
3. update the Amazon EQ instance named Devinstance by making the following changes to the stack named 1700182:
a) Change the EC2 instance type to us-east-t2.nano.
b) Allow SSH to connect to the EC2 instance from the IP address range 192.168.100.0/30.
c) Replace the instance profile IAM role with IamRoleB.
4. Deploy the changes by updating the stack using the CFServiceR01e role.
5. Edit the stack options to prevent accidental deletion.
6. Using the output from the stack, enter the value of the ProdInstanceId in the text box below:

A. See the for solution.

**Correct Answer: A**
**Section:**
**Explanation:**
Here are the steps to update an existing AWS CloudFormation stack:
Log in to the AWS Management Console and navigate to the CloudFormation service in the us-east-2 Region.
Find the existing stack named 1700182 and click on it.
Click on the "Update" button.
Choose "Replace current template" and upload the updated CloudFormation template from the Amazon S3 bucket named "cloudformation-bucket"
In the "Parameter" section, update the EC2 instance type to us-east-t2.nano and add the IP addressrange 192.168.100.0/30 for SSH access.
Replace the instance profile IAM role with IamRoleB.
In the "Capabilities" section, check the checkbox for "IAM Resources"
Choose the role CFServiceR01e and click on "Update Stack"
Wait for the stack to be updated.
Once the update is complete, navigate to the stack and click on the "Stack options" button, and select "Prevent updates to prevent accidental deletion"
To get the value of the ProdInstanceId , navigate to the "Outputs" tab in the CloudFormation stack and find the key "ProdInstanceId". The value corresponding to it is the value that you need to enter in the text box below.
Note:
You can use AWS CloudFormation to update an existing stack.
You can use the AWS CloudFormation service role to deploy updates.
You can refer to the AWS CloudFormation documentation for more information on how to update and manage stacks: https://aws.amazon.com/cloudformation/

**QUESTION 2**
A company recently acquired another corporation and all of that corporation's AWS accounts. A financial analyst needs the cost data from these accounts. A SysOps administrator uses Cost Explorer to generate cost and usage reports. The SysOps administrator notices that "No Tagkey" represents 20% of the monthly cost.
What should the SysOps administrator do to tag the "No Tagkey" resources?

A. Add the accounts to AWS Organizations. Use a service control policy (SCP) to tag all the untagged resources.
B. Use an AWS Config rule to find the untagged resources. Set the remediation action to terminate the resources.

C. Use Cost Explorer to find and tag all the untagged resources.
D. Use Tag Editor to find and taq all the untaqqed resources.

**Correct Answer: D**
**Section:**
**Explanation:**
"You can add tags to resources when you create the resource. You can use the resource's service console or API to add, change, or remove those tags one resource at a time. To add tags to—or edit or delete tags of—multiple resources at once, use Tag Editor. With Tag Editor, you search for the resources that you want to tag, and then manage tags for the resources in your search results." https://docs.aws.amazon.com/ARG/latest/userguide/tag-editor.html

**QUESTION 3**
A SysOps administrator noticed that the cache hit ratio for an Amazon CloudFront distribution is less than 10%. Which collection of configuration changes will increase the cache hit ratio for the distribution? (Select TWO.)

A. Ensure that only required cookies, query strings, and headers are forwarded in the Cache Behavior Settings.
B. Change the Viewer Protocol Policy to use HTTPS only.
C. Configure the distribution to use presigned cookies and URLs to restrict access to the distribution.
D. Enable automatic compression of objects in the Cache Behavior Settings.
E. Increase the CloudFront time to live (TTL) settings in the Cache Behavior Settings.

**Correct Answer: A, E**
**Section:**
**Explanation:**
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/cache-hitratio.html#cache-hit-ratio-http-streaming

**QUESTION 4**
A SysOps administrator is setting up an automated process to recover an Amazon EC2 instance in the event of an underlying hardware failure. The recovered instance must have the same private IP address and the same Elastic IP address that the original instance had. The SysOps team must receive an email notification when the recovery process is initiated. Which solution will meet these requirements?

A. Create an Amazon CloudWatch alarm for the EC2 instance, and specify the StatusCheckFailed_Instance metric. Add an EC2 action to the alarm to recover the instance. Add an alarm notification to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the SysOps team email address to the SNS topic.
B. Create an Amazon CloudWatch alarm for the EC2 instance, and specify the StatusCheckFailed_System metric. Add an EC2 action to the alarm to recover the instance. Add an alarm notification to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the SysOps team email address to the SNS topic.
C. Create an Auto Scaling group across three different subnets in the same Availability Zone with a minimum, maximum, and desired size of 1. Configure the Auto Scaling group to use a launch template that specifies the private IP address and the Elastic IP address. Add an activity notification for the Auto Scaling group to send an email message to the SysOps team through Amazon Simple Email Service (Amazon SES).
D. Create an Auto Scaling group across three Availability Zones with a minimum, maximum, and desired size of 1. Configure the Auto Scaling group to use a launch template that specifies the private IP address and the Elastic IP address. Add an activity notification for the Auto Scaling group to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the SysOps team email address to the SNS topic.

**Correct Answer: A**
**Section:**
**Explanation:**
Reference: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-cloudwatch-createalarm.html

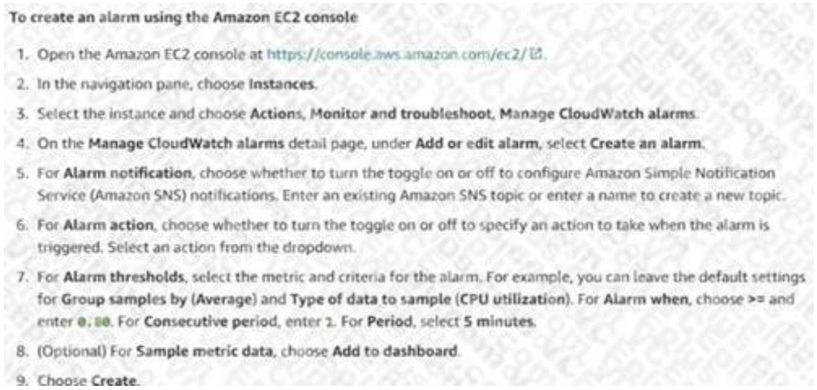

**To create an alarm using the Amazon EC2 console**

1. Open the Amazon EC2 console at https://console.aws.amazon.com/ec2/.
2. In the navigation pane, choose **Instances**.
3. Select the instance and choose **Actions**, **Monitor and troubleshoot**, **Manage CloudWatch alarms**.
4. On the **Manage CloudWatch alarms** detail page, under **Add or edit alarm**, select **Create an alarm**.
5. For **Alarm notification**, choose whether to turn the toggle on or off to configure Amazon Simple Notification Service (Amazon SNS) notifications. Enter an existing Amazon SNS topic or enter a name to create a new topic.
6. For **Alarm action**, choose whether to turn the toggle on or off to specify an action to take when the alarm is triggered. Select an action from the dropdown.
7. For **Alarm thresholds**, select the metric and criteria for the alarm. For example, you can leave the default settings for **Group samples by (Average)** and **Type of data to sample (CPU utilization)**. For **Alarm when**, choose **>=** and enter **0.80**. For **Consecutive period**, enter **1**. For **Period**, select **5 minutes**.
8. (Optional) For **Sample metric data**, choose **Add to dashboard**.
9. Choose **Create**.

**QUESTION 5**
A company has an Amazon Route 53 private hosted zone in its AWS account. The private hosted zone is connected to the company's on-premises data center by an AWS Direct Connect connection. Virtual machines (VMs) in the on- premises data center need to resolve DNS queries that exist in the private hosted zone.
What is the MOST operationally efficient solution that meets this requirement?

A. Create a Route 53 inbound resolver. Configure the on-premises VMs to use the inbound resolver.
B. Create a Route 53 outbound resolver. Configure the on-premises VMs to use the outbound resolver.
C. Configure the security group on the Route 53 private hosted zone by adding an inbound rule for the on-premises CIDR range.
D. Configure a Route 53 public hosted zone. Create an NS record for the private hosted zone. Query the public hosted zone from the on-premises VMs.

**Correct Answer: D**
**Section:**
**Explanation:**
Reference: https://aws.amazon.com/blogs/security/how-to-centralize-dns-management-in-a-multi-account-environment/

**QUESTION 6**
A development team recently deployed a new version of a web application to production. After the release, penetration testing revealed a cross-site scripting vulnerability that could expose user data. Which AWS service will mitigate this issue?

A. AWS Shield Standard
B. AWS WAF
C. Elastic Load Balancing
D. Amazon Cognito

**Correct Answer: B**
**Section:**
**Explanation:**
Reference: https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-xss-match.html

**QUESTION 7**
A SysOps administrator has enabled AWS CloudTrail in an AWS account. If CloudTrail is disabled, it must be re-enabled immediately. What should the SysOps administrator do to meet these requirements WITHOUT writing custom code?

A. Add the AWS account to AWS Organizations. Enable CloudTrail in the management account.
B. Create an AWS Config rule that is invoked when CloudTrail configuration changes. Apply the AWSConfigureCloudTrailLogging automatic remediation action.

C. Create an AWS Config rule that is invoked when CloudTrail configuration changes. Configure the rule to invoke an AWS Lambda function to enable CloudTrail.
D. Create an Amazon EventBridge (Amazon CloudWatch Event) hourly rule with a schedule pattern to run an AWS Systems Manager Automation document to enable CloudTrail.

**Correct Answer: B**
**Section:**

**QUESTION 8**
A Sysops administrator creates an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that uses AWS Fargate. The cluster is deployed successfully. The Sysops administrator needs to manage the cluster by using the kubect1 command line tool.
Which of the following must be configured on the Sysops administrator's machine so that kubect1can communicate with the cluster API server?

A. The kubeconfig file
B. The kube-proxy Amazon EKS add-on
C. The Fargate profile
D. The eks-connector.yaml file

**Correct Answer: A**
**Section:**
**Explanation:**
The kubeconfig file is a configuration file used to store cluster authentication information, which is required to make requests to the Amazon EKS cluster API server. The kubeconfig file will need to be configured on the SysOps administrator's machine in order for kubectl to be able to communicate with the cluster API server. https://aws.amazon.com/blogs/developer/running-a-kubernetes-job-in-amazon-eks-on-aws-fargateusing-aws-stepfunctions/

**QUESTION 9**
A Sysops administrator needs to configure automatic rotation for Amazon RDS database credentials.
The credentials must rotate every 30 days. The solution must integrate with Amazon RDS.
Which solution will meet these requirements with the LEAST operational overhead?

A. Store the credentials in AWS Systems Manager Parameter Store as a secure string. Configure automatic rotation with a rotation interval of 30 days.
B. Store the credentials in AWS Secrets Manager. Configure automatic rotation with a rotation interval of 30 days.
C. Store the credentials in a file in an Amazon S3 bucket. Deploy an AWS Lambda function to automatically rotate the credentials every 30 days.
D. Store the credentials in AWS Secrets Manager. Deploy an AWS Lambda function to automatically rotate the credentials every 30 days.

**Correct Answer: B**
**Section:**
**Explanation:**
Storing the credentials in AWS Secrets Manager and configuring automatic rotation with a rotation interval of 30 days is the most efficient way to meet the requirements with the least operational overhead. AWS Secrets Manager automatically rotates the credentials at the specified interval, so there is no need for an additional AWS Lambda function or manual rotation. Additionally, Secrets Manager is integrated with Amazon RDS, so the credentials can be easily used with the RDS database.

**QUESTION 10**
A company has an application that runs only on Amazon EC2 Spot Instances. The instances run in an Amazon EC2 Auto Scaling group with scheduled scaling actions. However, the capacity does not always increase at the scheduled times, and instances terminate many times a day. A Sysops administrator must ensure that the instances launch on time and have fewer interruptions. Which action will meet these requirements?

A. Specify the capacity-optimized allocation strategy for Spot Instances. Add more instance types to the Auto Scaling group.
B. Specify the capacity-optimized allocation strategy for Spot Instances. Increase the size of the instances in the Auto Scaling group.
C. Specify the lowest-price allocation strategy for Spot Instances. Add more instance types to the Auto Scaling group.
D. Specify the lowest-price allocation strategy for Spot Instances. Increase the size of the instances in the Auto Scaling group.

**Correct Answer: A**
**Section:**
**Explanation:**
Specifying the capacity-optimized allocation strategy for Spot Instances and adding more instance types to the Auto Scaling group is the best action to meet the requirements. Increasing the size of the instances in the Auto Scaling group will not necessarily help with the launch time or reduce interruptions, as the Spot Instances could still be interrupted even with larger instance sizes.

**QUESTION 11**
A company stores its data in an Amazon S3 bucket. The company is required to classify the data and find any sensitive personal information in its S3 files. Which solution will meet these requirements?

A. Create an AWS Config rule to discover sensitive personal information in the S3 files and mark them as noncompliant.
B. Create an S3 event-driven artificial intelligence/machine learning (AI/ML) pipeline to classify sensitive personal information by using Amazon Recognition.
C. Enable Amazon GuardDuty. Configure S3 protection to monitor all data inside Amazon S3.
D. Enable Amazon Macie. Create a discovery job that uses the managed data identifier.

**Correct Answer: D**
**Section:**
**Explanation:**
Amazon Macie is a security service designed to help organizations find, classify, and protect sensitive data stored in Amazon S3. Amazon Macie uses machine learning to automatically discover, classify, and protect sensitive data in Amazon S3. Creating a discovery job with the managed data identifier will allow Macie to identify sensitive personal information in the S3 files and classify it accordingly. Enabling AWS Config and Amazon GuardDuty will not help with this requirement as they are not designed to automatically classify and protect data.

**QUESTION 12**
A company wants to apply an existing Amazon Route 53 private hosted zone to a new VPC to allow for customized resource name resolution within the VPC. The Syspps administrator created the VPC and added the appropriate resource record sets to the private hosted zone.
Which step should the SysOps administrator take to complete the setup?

A. Associate the Route 53 private hosted zone with the VPC.
B. Create a rule in the default security group for the VPC that allows traffic to the Route 53 Resolver.
C. Ensure the VPC network ACLs allow traffic to the Route 53 Resolver.
D. Ensure there is a route to the Route 53 Resolver in each of the VPC route tables.

**Correct Answer: A**
**Section:**
**Explanation:**
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/hosted-zones-private.html
To apply an existing Amazon Route 53 private hosted zone to a new VPC, the appropriate step is to associate the private hosted zone with the new VPC. This allows the resources within the VPC to use the custom DNS settings defined in the private hosted zone. Option A is the correct step to ensure that DNS queries from the new VPC are resolved using the specified private hosted zone. Detailed steps for this process can be found in the AWS Route 53 documentation on associating hosted zones with VPCs Associating Hosted Zones with VPCs.

**QUESTION 13**
A company runs its web application on multiple Amazon EC2 instances that are part of an Auto Scaling group. The company wants the Auto Scaling group to scale out as soon as CPU utilization rises above 50% for the instances.
How should a SysOps administrator configure the Auto Scaling group to meet these requirements?

A. Configure the Auto Scaling group to scale based on events.
B. Configure the Auto Scaling group to scale based on a schedule.
C. Configure the Auto Scaling group to scale dynamically based on demand.
D. Configure the Auto Scaling group to use predictive scaling.

**Correct Answer: C**
**Section:**
**Explanation:**
To ensure that the Auto Scaling group scales out when CPU utilization rises above 50%, the administrator should configure the Auto Scaling group to dynamically scale based on demand. This is achieved by setting up a scaling policy that triggers based on specific CloudWatch alarms---like CPU utilization exceeding 50%. This dynamic scaling method directly responds to changes in workload, ensuring that resources are allocated efficiently and promptly as demand increases. Option C is the correct answer, aligning with best practices for managing EC2 Auto Scaling based on real-time metrics. Further guidance is available in AWS documentation on dynamic scaling Dynamic Scaling for EC2.

**QUESTION 14**
A SysOps administrator needs to design a disaster recovery (DR) plan for an application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The application uses an Amazon Aurora PostgreSQL database. The recovery time objective (RTO) and recovery point objective (RPO) are 15 minutes each.
Which combination of steps should the SysOps administrator take to meet these requirements MOST cost-effectively? (Select TWO.)

A. Configure Aurora backups to be exported to the DR Region.
B. Configure the Aurora cluster to replicate data to the DR Region by using the Aurora global database option.
C. Configure the DR Region with an ALB and an Auto Scaling group. Use the same configuration as in the primary Region.
D. Configure the DR Region with an ALB and an Auto Scaling group. Set the Auto Scaling group's minimum capacity, maximum capacity, and desired capacity to 1.
E. Manually launch a new ALB and a new Auto Scaling group by using AWS CloudFormation during a failover activity.

**Correct Answer: B, C**
**Section:**
**Explanation:**
For a disaster recovery (DR) plan with a 15-minute RTO and RPO, the most cost-effective steps include:
B: Configuring the Aurora cluster to replicate data to the DR region using the Aurora global database option. This allows continuous replication with typically low replication lag, meeting the 15-minute RPO requirement efficiently.
C: Pre-configuring the DR region with an ALB and an Auto Scaling group using the same configuration as the primary region. This ensures readiness and quick failover, aligning with the 15-minute RTO target.
These steps provide a robust disaster recovery setup that minimizes downtime and data loss while optimizing costs by using built-in AWS functionalities and avoiding over-provisioning. More information can be found in the AWS documentation on Aurora Global Databases Aurora Global Databases and disaster recovery planning AWS Disaster Recovery.

**QUESTION 15**
A company migrates a write-once, read-many (WORM) drive to an Amazon S3 bucket that has S3 Object Lock configured in governance mode. During the migration, the company copies unneeded data to the S3 bucket.
A SysOps administrator attempts to delete the unneeded data from the S3 bucket by using the AWS CLI. However, the SysOps administrator receives an error.
Which combination of steps should the SysOps administrator take to successfully delete the unneeded data? (Select TWO.)

A. Increase the Retain Until Date.
B. Assume a role that has the s3:BypassLegalRetention permission.
C. Assume a role that has the s3:BypassGovernanceRetention permission.
D. Include the x-amz-bypass-governance-retention:true header in the request when issuing the delete command.
E. Include the x-amz-bypass-legal-retention:true header in the request when issuing the delete command.

**Correct Answer: C, D**
**Section:**
**Explanation:**
When using Amazon S3 Object Lock configured in governance mode, deleting objects before their retention period ends requires specific permissions. To bypass these governance restrictions, the administrator must:
C: Assume a role that has the s3:BypassGovernanceRetention permission. This permission allows the role to override the governance mode restrictions.
D: Include the x-amz-bypass-governance-retention:true header in the delete request. This header is necessary to programmatically bypass the governance retention settings when making a delete request via the AWS CLI or SDK. These steps enable the deletion of objects under governance mode retention without waiting for the retention period to expire, addressing the need to remove unintended data uploads effectively. For further details, refer to the AWS documentation on S3 Object Lock Amazon S3 Object Lock.

**QUESTION 16**
A company is deploying an ecommerce application to an AWS Region that is located in France. The company wants users from only France to be able to access the first version of the application. The company plans to add more countries for the next version of the application. A SysOps administrator needs to configure the routing policy in Amazon Route 53.
Which solution will meet these requirements?

A. Use a geoproximity routing policy. Select France as the location in the record.
B. Use a geolocation routing policy. Select France as the location in the record.
C. Use an IP-based routing policy. Select all IP addresses that are allocated to France in the record.
D. Use a geoproximity routing policy. Select all IP addresses that are allocated to France in the record.

**Correct Answer: B**
**Section:**
**Explanation:**
To restrict access to an application based on geographic location (France in this case), the appropriate routing policy in Amazon Route 53 is geolocation routing. This policy allows you to specify traffic routing based on the geographic location of your users:
B: Use a geolocation routing policy. Select France as the location in the record. This ensures that only DNS queries originating from France are routed to the application, fulfilling the requirement to limit access to users within France initially. More information about setting up geolocation routing can be found in the AWS Route 53 documentation on geolocation routing Amazon Route 53 Geolocation Routing.

**QUESTION 17**
A company has an on-premises DNS solution and wants to resolve DNS records in an Amazon Route 53 private hosted zone for example.com. The company has set up an AWS Direct Connect connection for network connectivity between the on-premises network and the VPC. A SysOps administrator must ensure that an on-premises server can query records in the example.com domain.
What should the SysOps administrator do to meet these requirements?

A. Create a Route 53 Resolver inbound endpoint Attach a security group to the endpoint to allow inbound traffic on TCP/UDP port 53 from the on-premises DNS servers.
B. Create a Route 53 Resolver inbound endpoint. Attach a security group to the endpoint to allow outbound traffic on TCP/UDP port 53 to the on-premises DNS servers.
C. Create a Route 53 Resolver outbound endpoint. Attach a security group to the endpoint to allow inbound traffic on TCP/UDP port 53 from the on-premises DNS servers.
D. Create a Route 53 Resolver outbound endpoint. Attach a security group to the endpoint to allow outbound traffic on TCP/UDP port 53 to the on-premises DNS servers.

**Correct Answer: A**
**Section:**
**Explanation:**
To allow on-premises servers to resolve DNS records in an Amazon Route 53 private hosted zone via AWS Direct Connect, the following step should be taken:
A: Create a Route 53 Resolver inbound endpoint and attach a security group that allows inbound traffic on TCP/UDP port 53 from the on-premises DNS servers. This setup enables the on-premises DNS servers to forward DNS queries to AWS for the domains managed by Route 53. The inbound resolver endpoint acts as a bridge between the on-premises network and AWS for DNS resolution. Additional guidance on setting up Route 53 Resolver endpoints can be found in AWS documentation Route 53 Resolver.

**QUESTION 18**
A company has an AWS Lambda function in Account

A. The Lambda function needs to read the objects in an Amazon S3 bucket in Account B. A SysOps administrator must create corresponding 1AM roles in both accounts. Which solution will meet these requirements?
B. In Account A, create a Lambda execution role to assume the role in Account B. In Account B. create a role that the function can assume to gain access to the S3 bucket.
C. In Account A, create a Lambda execution role that provides access to the S3 bucket. In Account B. create a role that the function can assume.
D. In Account A. create a role that the function can assume. In Account B, create a Lambda execution role that provides access to the S3 bucket.
E. In Account A. create a role that the function can assume to gain access to the S3 bucket. In Account B. create a Lambda execution role to assume the role in Account A.

**Correct Answer: A**
**Section:**
**Explanation:**

For a Lambda function in Account A to access an S3 bucket in Account B, the correct IAM roles setup includes:
A: In Account A, create a Lambda execution role that has permissions to assume another role in Account B. In Account B, create a role with permissions to access the S3 bucket and trust the Lambda execution role from Account A to assume it. This configuration allows the Lambda function to assume the cross-account role and access the S3 bucket as needed, maintaining security and proper access control. AWS documentation provides more details on setting up cross-account roles for Lambda access AWS Lambda Permissions.

**QUESTION 19**
A company's SysOps administrator maintains a highly available environment. The environment includes Amazon EC2 instances and an Amazon RDS Multi-AZ database. The EC2 instances are in an Auto Scaling group behind an Application Load Balancer.
Recently, the company conducted a failover test. The SysOps administrator needs to decrease the failover time of the RDS database by at least 10%.
Which solution will meet this requirement?

A. Increase the RDS instance size.
B. Modify the RDS cluster to run in a single Availability Zone.
C. Create a read replica in another AWS Region. Promote the read replica in case of failure.
D. Create an RDS proxy. Point the application to the proxy endpoint.

**Correct Answer: A**
**Section:**
**Explanation:**
https://aws.amazon.com/blogs/database/improving-application-availability-with-amazon-rds-proxy/

**QUESTION 20**
A company has an application that collects notifications from thousands of alarm systems. The notifications include alarm notifications and information notifications. The information notifications include the system arming processes, disarming processes, and sensor status.
All notifications are kept as messages in an Amazon Simple Queue Service (Amazon SQS) queue. Amazon EC2 instances that are in an Auto Scaling group process the messages. A SysOps administrator needs to implement a solution that prioritizes alarm notifications over information notifications.
Which solution will meet these requirements?

A. Adjust the Auto Scaling group to scale faster when a high number of messages is in the queue.
B. Use the Amazon Simple Notification Service (Amazon SNS) fanout feature with Amazon SQS to send the notifications in parallel to all the EC2 instances.
C. Add an Amazon DynamoDB stream to accelerate the message processing.
D. Create a queue for alarm notifications and a queue for information notifications. Update the application to collect messages from the alarm notifications queue first.

**Correct Answer: D**
**Section:**
**Explanation:**
To prioritize alarm notifications over information notifications in an Amazon SQS environment, creating separate queues for each type of notification is the best approach. This allows the application processing the queues to prioritize fetching messages from the alarm notifications queue first. Option D directly addresses the need for prioritization without complicating the existing infrastructure. This setup is supported by best practices for using Amazon SQS, where using multiple queues to segregate message types is a recommended strategy Amazon SQS Best Practices.

**QUESTION 21**
A company uses AWS CloudFormation to manage a stack of Amazon EC2 instances on AWS. A SysOps administrator 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.

**Correct Answer: D**
**Section:**
**Explanation:**
To prevent the EC2 instances and their data from being deleted when a CloudFormation stack is deleted:
DeletionPolicy Attribute: In the CloudFormation template that defines the EC2 instances, set the DeletionPolicy attribute to Retain for each EC2 instance resource. This setting ensures that when the CloudFormation stack is deleted, the EC2 instances are not terminated.
Impact of the Retain Policy: With this policy, all data on the instance, such as data on its attached EBS volumes, remains intact even after the stack deletion. The resources will remain in your AWS account and will need to be managed or deleted manually thereafter.
This approach is directly supported by AWS CloudFormation and provides a simple and effective way to protect EC2 instances and their data during stack deletions.

**QUESTION 22**
A company has an application that runs behind an Application Load Balancer (ALB) in the us-west-2 Region. An Amazon Route 53 record set contains an alias record for app.anycompany.com that references the ALB in us-west-2 and uses a simple routing policy. The application is experiencing an increase in users from other locations in the world. These users are experiencing high latency.
Most of the new users are close to the ap-southeast-2 Region. The company deploys a copy of the application to ap-southeast-2. A SysOps administrator must implement a solution that automatically routes requests to the lowest latency endpoint for users without changing the URL.
Which solution will meet these requirements?

A. Add a new value to the existing alias record for app.anycompany.com with the DNS name of the new ALB in ap-southeast-2.
B. Change the existing alias record to use a geolocation routing policy. Create two geolocation records, one record that references each ALSelect the location that is closest to each Region.
C. Change the existing alias record to use a latency routing policy. Create two latency records, one record that references each ALB.
D. Change the existing alias record to use a multivalue routing policy Add the DNS name of each ALB to the record.

**Correct Answer: C**
**Section:**
**Explanation:**
To optimize the routing of requests to the application for users in different geographic locations, use Amazon Route 53's latency-based routing:
Latency Routing Policy: Modify the Route 53 DNS settings for app.anycompany.com by changing the routing policy to latency. This policy routes user requests to the server that has the lowest latency relative to the user's location.
Configure Latency Records: Create latency records for each ALB---one for the ALB in us-west-2 and another for the new ALB in ap-southeast-2. Route 53 will automatically direct traffic to the endpoint that provides the lowest latency connection to the user.
Seamless User Experience: By using latency routing, the URL remains the same (app.anycompany.com), but the backend service location varies based on which endpoint is likely to respond the fastest. This setup provides a seamless experience for users, reducing latency without requiring any changes to how they access the application.
This method leverages Route 53's advanced DNS capabilities to ensure optimal performance and user experience by dynamically routing traffic based on geographic latency considerations.

**QUESTION 23**
A company has applications that process transaction requests multiple times each minute. The applications write transaction data to a single Amazon RDS DB instance. As the company begins to process more transactions, the company becomes concerned that it has no failover solution in place for disaster recovery (DR). The company needs the DB instance to fail over automatically without losing any committed transactions.
Which solution will meet these requirements?

A. Create an RDS read replica in the same AWS Region. Configure an AWS Lambda function to promote the replica as the primary DB instance during a DR scenario.
B. Create an RDS read replica in a different AWS Region. Configure an AWS Lambda function to promote the replica as the primary DB instance during a DR scenario.
C. Modify the DB instance to be a Multi-AZ deployment.
D. Setup an Amazon CloudWatch alarm that monitors the DB instance memory utilization with a threshold greater than 90%. Invoke an AWS Lambda function to restart the DB instance.

**Correct Answer: C**
**Section:**
**Explanation:**
For an RDS instance that needs high availability and automatic failover capabilities, setting it up as a Multi-AZ deployment is the most effective solution:
Multi-AZ Deployment: This feature allows Amazon RDS to automatically provision and manage a synchronous standby replica of your database in a different Availability Zone (AZ). The primary DB instance and the standby

replica contain the same data, providing data redundancy and fail-safe mechanism.
Automatic Failover: In the event of a planned or unplanned outage of your primary DB instance (including DB instance failure, AZ failure, or network failure), RDS automatically fails over to the standby so that database operations can resume quickly without administrative intervention.
Data Integrity: This setup ensures no data loss for committed transactions, as the standby replica is always in sync with the primary.
By enabling Multi-AZ deployment, you ensure that your database environment has high availability and robustness, addressing both disaster recovery and operational continuity without losing any committed transactions.

**QUESTION 24**
A company is using AWS Certificate Manager (ACM) to manage public SSL/TLS certificates. A SysOps administrator needs to send an email notification when a certificate has less than 14 days until expiration.
Which solution will meet this requirement with the LEAST operational overhead?

A. Create an Amazon CloudWatch custom metric to monitor certificate expiration for all ACM certificates. Create an Amazon EventBridge rule that has an event source of a ws. cloud watch Configure the rule to send an event to a target Amazon Simple Notification Service (Amazon SNS) topic if the DaysToExpiry metric is less than 14. Subscribe the appropriate email addresses to the SNS topic.
B. Create an Amazon EventBridge rule that has an event source of aws.acm. Configure the rule to evaluate the DaysToExpiry melric for all ACM certificates. Configure the rule to send an event to a target Amazon Simple Notification Service (Amazon SNS) topic if DaysToExpiry is less than 14. Subscribe the appropriate email addresses to the SNS topic.
C. Create an Amazon CloudWatch dashboard that displays the DaysToExpiry metric for all ACM certificates. If DaysToExpiry is less than 14, send an email message to the appropriate email addresses. Send the email message by running a predefined CLI command to publish to an Amazon Simple Notification Service (Amazon SNS) topic.
D. Create an Amazon EventBridge rule that has an event source of aws.acm. Configure the rule to evaluate the DaysToExpiry metric for all ACM certificates. Configure a target SMS identity that uses a predefined email template. Configure the rule to send an event to the target SMS identity if DaysToExpiry is less than 14.

**Correct Answer: B**
**Section:**
**Explanation:**
To send an email notification when an ACM certificate has less than 14 days until expiration, the most operationally efficient solution is to create an Amazon EventBridge rule that directly monitors AWS ACM for certificate expiry details. EventBridge can capture aws.acm events related to certificate expiration and evaluate the DaysToExpiry metric directly. Configure the rule to send notifications to an Amazon SNS topic when DaysToExpiry is less than 14, and subscribe the necessary email addresses to this SNS topic. This approach avoids the need for custom metrics or manual monitoring, providing a streamlined and automatic notification system. Option B is the correct answer as it leverages native integration between ACM, EventBridge, and SNS for minimal operational overhead. More details on EventBridge rules can be found in AWS documentation EventBridge Rules.

**QUESTION 25**
A company wants to store sensitive financial data within Amazon S3 buckets. The company has a corporate policy that does not allow public read or write access to the buckets. A SysOps administrator must create a solution to automatically remove S3 permissions that allow public read or write access.
Which AWS service should the SysOps administrator use to meet these requirements in the MOST operationally efficient manner?

A. AWSConfig
B. AWS Security Hub
C. AWS Trusted Advisor
D. Amazon Inspector

**Correct Answer: A**
**Section:**
**Explanation:**
AWS Config is the best service to automatically manage and remediate S3 bucket permissions that violate corporate policies against public access. AWS Config continuously monitors and records AWS resource configurations and allows you to create rules that trigger automatic responses when public access configurations are detected. This approach is highly operationally efficient as it automates compliance and enforcement of security policies without manual intervention. Option A is correct. AWS Config can be used to assess, audit, and evaluate the configurations of AWS resources, including S3 buckets. Reference AWS Config.

**QUESTION 26**
A company is running a development application on an Amazon EC2 instance. The application uploads 500.000 files that are 1 GB in size into a large! Amazon S3 bucket that has default encryption enabled The EC2 instance is in the same AWS Region where the S3 bucket is deployed.
The company uses performance logging that is built into the application software. The logs show that the application is constantly waiting for the files to be written to the S3 bucket. A SysOps administrator needs to improve the application's throughput performance. The SysOps administrator validates that the networking on the EC2 instance is not constrained.

What should the SysOps administrator do to improve the S3 upload performance''

A. Enable S3 Transfer Acceleration on the S3 bucket.
B. Split the S3 write operations to use multiple bucket prefixes to write items in parallel.
C. Configure AWS PrivateLink for Amazon S3 Turn off encryption on the S3 bucket
D. Configure AWS Global Accelerator in the Region. Turn off encryption on the S3 bucket.

**Correct Answer: B**
**Section:**
**Explanation:**
Improve S3 Upload Performance:
Using multiple bucket prefixes can improve throughput by allowing parallel upload streams.
Steps:
Modify the application to write files to different prefixes in the S3 bucket.
Example: Instead of writing all files to s3://bucket-name/, write to s3://bucket-name/prefix1/, s3://bucket-name/prefix2/, etc.

**QUESTION 27**
A company that uses AWS Organizations recently implemented AWS Control Tower The company now needs to centralize identity management A SysOps administrator must federate AWS 1AM Identity Center with an external SAML 2.0 identity provider (IdP) to centrally manage access to all the company's accounts and cloud applications
Which prerequisites must the SysOps administrator have so that the SysOps administrator can connect to the external IdP? (Select TWO.)

A. A copy of the 1AM Identity Center SAML metadata
B. The IdP metadata, including the public X.509 certificate
C. The IP address of the IdP
D. Root access to the management account
E. Administrative permissions to the member accounts of the organization

**Correct Answer: A, B**
**Section:**
**Explanation:**
IAM Identity Center SAML Metadata:
This metadata is required to establish the trust relationship between AWS IAM Identity Center and the external SAML 2.0 identity provider.
Steps:
Download the IAM Identity Center SAML metadata from the AWS Management Console.
Provide this metadata to the external IdP.
IdP Metadata:
The metadata from the IdP, including the public X.509 certificate, is needed to configure the trust relationship.
Steps:
Obtain the IdP metadata, which includes the entity ID, endpoints, and X.509 certificate.
Configure the IAM Identity Center with this information.

**QUESTION 28**
A company is using an Amazon EC2 Auto Scaling group to support a workload A Sytfhe company now needs to centruito Scaling group is configured with two similar scaling policies dP) to centrally manage access to One scaling policy adds 5 instances when CPU utilization reaches 80%. The other sctrator can connect to the extemahen CPU utilization leaches 80%.
What will happen when CPU utilization reaches the 80% threshold?

A. Amazon EC2 Auto Scaling will add 5 instances
B. Amazon EC2 Auto Scaling will add 10 instances

C. Amazon EC2 Auto Scaling will add 15 instances.
D. The Auto Scaling group will not scale because of conflicting policies

**Correct Answer: B**
**Section:**
**Explanation:**
Scaling Policies in Auto Scaling:
When multiple scaling policies trigger at the same time, each policy is executed independently.
If both policies are set to add 5 instances when CPU utilization reaches 80%, they will both be executed when the threshold is met.
Therefore, the total number of instances added will be the sum of the instances specified in both policies.
In this case, 5 instances from one policy and 5 instances from the other policy will result in a total of 10 instances being added.
Steps to Configure and Verify Scaling Policies:
Go to the AWS Management Console.
Navigate to EC2 and select 'Auto Scaling Groups.'
Select your Auto Scaling group and review the scaling policies.
Ensure that both scaling policies are configured to trigger at 80% CPU utilization.
Monitor the Auto Scaling group's activity to verify the addition of instances when the CPU utilization threshold is reached.

**QUESTION 29**
A company hosts an application on Amazon EC2 instances The instances are in an Amazon EC2 Auto Scaling group that uses a launch template The amount of application traffic changes throughout the day. Scaling events happen frequently.
A SysOps administrator needs to help developers troubleshoot the application. When a scaling event removes an instance. EC2 Auto Scaling terminates the instance before the developers can log in to the instance to diagnose issues.
Which solution will prevent termination of the instance so that the developers can log in to the instance?

A. Ensure that the Delete on termination setting is turned off in the UserData section of the launch template
B. Update the Auto Scaling group by enabling instance scale-in protection for newly launched instances.
C. Use Amazon Inspector to configure a rules package to protect the instances from termination.
D. Use Amazon GuardDuty to configure rules to protect the instances from termination.

**Correct Answer: B**
**Section:**
**Explanation:**
Enabling Instance Scale-In Protection:
Instance scale-in protection prevents Auto Scaling from terminating specific instances.
Steps:
Go to the AWS Management Console.
Navigate to EC2 and select 'Auto Scaling Groups.'
Select your Auto Scaling group.
Go to the 'Instance management' tab.
Select the instances you want to protect and click 'Actions.'
Choose 'Enable scale-in protection.'
This ensures that instances are not terminated during troubleshooting.

**QUESTION 30**
A company has many accounts in an organization in AWS Organizations The company must automate resource provisioning from the organization's management account to the member accounts.
Which solution will meet this requirement?

A. Create an AWS CkHJdFormation change set Deploy the change set to all member accounts
B. Create an AWS CtoudFormation nested stack Deploy the nested stack to all member accounts.
C. Create an AWS CtoudFormation stack set Deploy the stack set to all member accounts.
D. Create an AWS Serverless Application Model (AWS SAM) template. Deploy the template to all member accounts.

**Correct Answer: C**
**Section:**
**Explanation:**
Using CloudFormation Stack Sets:
CloudFormation stack sets allow you to deploy CloudFormation stacks across multiple AWS accounts and regions.
Steps:
Go to the AWS Management Console.
Navigate to CloudFormation and select 'StackSets.'
Click on 'Create StackSet.'
Provide the template URL or upload a template file.
Configure the stack set options and specify the accounts and regions.
Deploy the stack set to the specified accounts and regions.

**QUESTION 31**
An AWS Cloud Formation 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 Cloud Format ton 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 live 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 Cloud Formation template definition of the RDS instance.
D. Create a new CloudFormallon template to perform backups of the RDS instance, and run this template before deleting the stack.

**Correct Answer: C**
**Section:**
**Explanation:**
Snapshot Deletion Policy:
The Snapshot Deletion Policy ensures that a snapshot is created when an RDS instance is deleted as part of a CloudFormation stack deletion.
Steps:
Update your CloudFormation template to include the DeletionPolicy attribute for the RDS instance resource.
Example template snippet:
Resources:
MyDBInstance:
Type: AWS::RDS::DBInstance
Properties:
# DB instance properties
DeletionPolicy: Snapshot
This configuration retains a snapshot of the RDS instance data when the stack is deleted.
Reference: AWS CloudFormation DeletionPolicy

**QUESTION 32**
A company wants to prohibit its developers from using a particular family of Amazon EC2 instances The company uses AWS Organizations and wants to apply the restriction across multiple accounts
What is the MOST operationally efficient way for the company lo apply service control policies (SCPs) to meet these requirements?

A. Add the accounts to an organizational unit (OUf Apply the SCPs to the OU.
B. Add the accounts to resource groups in AWS Resource Groups. Apply the SCPs to the resource groups.
C. Apply the SCPs to each developer account.
D. Enroll the accounts with AWS Control Tower. Apply the SCPs to the AWS Control Tower management account.

**Correct Answer: A**
**Section:**
**Explanation:**
Applying SCPs to an Organizational Unit:
Service Control Policies (SCPs) allow you to manage permissions for multiple AWS accounts within an organization.
Steps:
Go to the AWS Management Console.
Navigate to AWS Organizations.
Create an Organizational Unit (OU) if not already created.
Move the target accounts into the OU.
Create an SCP that denies the use of the specific EC2 instance family.
Attach the SCP to the OU.
This approach ensures that the policy is applied consistently across all accounts in the OU.

**QUESTION 33**
A company has an application that uses Amazon DynamoDB tables The tables are spread across AWS accounts and AWS Regions. The company uses AWS CloudFormation to deploy AWS resources.
A new team at the company is deleting unused AWS resources. The team accidentally deletes several production DynamoDB tables by running an AWS Lambda function that makes a DynamoDB DeleteTable API call. The table deletions cause an application outage
A SysOps administrator must implement a solution that minimizes the chance of accidental deletions of tables. The solution also must minimize data loss that results from accidental deletions.
Which combination of steps will meet these requirements? (Select TWO.)

A. Enable termination protection for the CloudFormation stacks that deploy the DynamoDB tables.
B. Enable deletion protection for the DynamoDB tables
C. Enable point-in-time recovery for (he DynamoDB tables. Restore the tables if they are accidentally deleted.
D. Schedule daily backups of the DynamoDB tables. Restore the tables if they are accidentally deleted.
E. Export the DynamoDB tables to Amazon S3 every day. Use Import from Amazon S3 to restore data for tables that are accidentally deleted

**Correct Answer: B, C**
**Section:**
**Explanation:**
Enable deletion protection for the DynamoDB tables:
Deletion protection is a feature that prevents accidental deletion of DynamoDB tables. When enabled, it requires an additional step to disable this protection before the table can be deleted.
Steps:
Go to the AWS Management Console.
Navigate to DynamoDB.
Select the table you want to protect.
Choose the 'Overview' tab.
Under 'Deletion protection,' click 'Enable deletion protection.'
Enable point-in-time recovery (PITR) for the DynamoDB tables:
PITR provides continuous backups of your DynamoDB tables. You can restore the table to any point in time within the last 35 days.
Steps:
Go to the AWS Management Console.
Navigate to DynamoDB.

Select the table you want to enable PITR for.
Choose the 'Backups' tab.
Click on 'Enable Point-in-Time Recovery.'
If a table is accidentally deleted, you can restore it using PITR.
Go to the DynamoDB console.
Select 'Backups' from the navigation pane.
Find the table backup and choose 'Restore.'

**QUESTION 34**
A company is running an application on a group of Amazon EC2 instances behind an Application Load Balancer The EC2 instances run across three Availability Zones The company needs to provide the customers with a maximum of two static IP addresses for their applications
How should a SysOps administrator meet these requirement?

A. Add AWS Global Accelerator in front of the Application Load Balancer
B. Add an internal Network Load Balancer behind the Application Load Balancer
C. Configure the Application Load Balancer in only two Availability Zones.
D. Create two Elastic IP addresses and assign them to the Application Load Balancer.

**Correct Answer: A**
**Section:**
**Explanation:**
AWS Global Accelerator:
AWS Global Accelerator is a service that improves the availability and performance of your applications with a global user base. It provides static IP addresses that act as a fixed entry point to your application endpoints (such as ALBs).
Steps:
Go to the AWS Management Console.
Navigate to Global Accelerator.
Click on 'Create accelerator.'
Configure the accelerator by providing a name and adding listeners.
Add your Application Load Balancer as an endpoint.
Allocate two static IP addresses.
This setup ensures that your application is accessible via two static IP addresses, fulfilling the requirement.

**QUESTION 35**
A company currently runs its infrastructure within a VPC in a single Availability Zone The VPC is connected to the company's on-premises data center through an AWS Site-to-SIte VPN connection attached to a virtual pnvate gateway. The on-premises route tables route all VPC networks to the VPN connection Communication between the two environments is working correctly. A SysOps administrator created new VPC subnets within a new Availability Zone, and deployed new resources within the subnets. However, communication cannot be established between the new resources and the on-premises environment.
Which steps should the SysOps administrator take to resolve the issue?

A. Add a route to the route tables of the new subnets that send on-premises traffic to the virtual private gateway.
B. Create a ticket with AWS Support to request adding Availability Zones to the Site-to-Site VPN route configuration.
C. Establish a new Site-to-Site VPN connection between a virtual private gateway attached to the new Availability Zone and the on-premises data center
D. Replace the Site-to-Site VPN connection with an AWS Direct Connect connection.

**Correct Answer: A**
**Section:**
**Explanation:**
Adding a Route to the Route Tables:
When new subnets are created, they need appropriate routing to ensure communication with on-premises networks.

Steps:
Go to the AWS Management Console.
Navigate to VPC.
Select the route table associated with the new subnets.
Choose 'Edit routes.'
Add a new route with the destination CIDR block of the on-premises network.
For the target, select the virtual private gateway (VGW).
This ensures that traffic destined for the on-premises network is routed correctly through the VPN connection.

**QUESTION 36**
A company deploys a new application on three Amazon EC2 instances across three Availability Zones The company uses a Network Load Balancer (NLB) to route traffic lo the EC2 instances. A SysOps administrator must implement a solution so that the EC2 instances allow traffic from only the NLB.
What should the SysOps administrator do to meet these requirements with the LEAST operational overhead?

A. Configure the security group that is associated with the EC2 instances to allow traffic from only the security group that is associated with the NLB.
B. Configure the security group that is associated with the EC2 instances to allow traffic from only the elastic network interfaces that are associated with the NLB.
C. Create a network ACL. Associate the network ACL with the application subnets. Configure the network ACL to allow inbound traffic from only the CIDR ranges of the NLB.
D. Use a third-party firewall solution that is installed on a separate EC2 instance. Configure a firewall rule that allows traffic to the application's EC2 instances from only the subnets where the NLB is deployed

**Correct Answer: A**
**Section:**
**Explanation:**
Configuring Security Groups:
Security groups act as virtual firewalls for your instances to control inbound and outbound traffic.
Steps:
Go to the AWS Management Console.
Navigate to EC2.
Select 'Security Groups' from the left-hand menu.
Find and select the security group associated with your EC2 instances.
Choose the 'Inbound rules' tab and click 'Edit inbound rules.'
Add a rule to allow traffic from the security group associated with the NLB.
Type: Custom TCP (or the specific port your application uses)
Source: Select 'Custom' and enter the ID of the NLB's security group.
This setup ensures that the EC2 instances accept traffic only from the NLB, enhancing security with minimal operational overhead.

**QUESTION 37**
A company has created an AWS CloudFormation template that consists of the AWS: EC2 Instance resource and a custom Cloud Formation resource The custom CloudFormation resource is an AWS Lambda function that attempts to run automation on the Amazon EC2 instance.
During testing, the Lambda function fails because the Lambda function tries to run 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.

**Correct Answer: A**
**Section:**
**Explanation:**

DependsOn Attribute in CloudFormation:
The DependsOn attribute in AWS CloudFormation ensures that one resource is created only after another resource has been successfully created. In this case, it ensures that the EC2 instance is fully launched before the custom resource (the Lambda function) is executed.
Steps:
Update the CloudFormation template to include the DependsOn attribute for the custom resource.
Ensure that the custom resource references the EC2 instance.
Resources:
MyEC2Instance:
Type: AWS::EC2::Instance
Properties:
# EC2 properties
MyCustomResource:
Type: Custom::MyCustomResource
DependsOn: MyEC2Instance
Properties:
ServiceToken: !GetAtt MyLambdaFunction.Arn
# Other properties

**QUESTION 38**
A company has scientists who upload large data objects to an Amazon S3 bucket The scientists upload the objects as multipart uploads. The multipart uploads often fail because of poor end-client connectivity.
The company wants to optimize storage costs that are associated with the data. A SysOps administrator must implement a solution that presents metrics for incomplete uploads. The solution also must automatically delete any incomplete uploads after 7 days.
Which solution will meet these requirements?

A. Review the Incomplete Multipart Upload Bytes metre in the S3 Storage Lens dashboard Create an S3 Lifecycle policy to automatically delete any incomplete multipart uploads after 7 days.
B. Implement S3 Intelligent-Tiering to move data into tower-cost storage classes after 7 days. Create an S3 Storage Lens policy to automatically delete any incomplete multipart uploads after 7 days.
C. Access the S3 console Review the Metrics tab to check the storage that incomplete multipart uploads are consuming. Create an AWS Lambda function to delete any incomplete multipart uploads after 7 days.
D. Use the S3 analytics storage class analysis tool to identify and measure incomplete multipart uploads. Configure an S3 bucket policy to enforce restrictions on multipart uploads to delete incomplete multipart uploads after 7 days.

**Correct Answer: A**
**Section:**
**Explanation:**
S3 Storage Lens and Lifecycle Policies:
Incomplete Multipart Upload Bytes Metric: This metric in S3 Storage Lens helps you identify the storage consumed by incomplete multipart uploads.
S3 Lifecycle Policies: Lifecycle policies allow you to automatically manage the lifecycle of objects, including deleting incomplete multipart uploads after a specified number of days.
Steps:
Go to the AWS Management Console.
Navigate to S3 and select the bucket.
Go to the 'Metrics' tab and view the 'Incomplete Multipart Upload Bytes' metric in the S3 Storage Lens dashboard.
To create a lifecycle policy:
Select the bucket.
Go to the 'Management' tab.
Under 'Lifecycle rules,' click 'Create lifecycle rule.'
Define a rule name.
Choose 'Multipart upload' and specify 'Delete incomplete multipart uploads' after 7 days.
Save the rule.
AWS S3 Storage Lens
AWS S3 Lifecycle Policies

**QUESTION 39**
An application uses an Amazon Aurora MySQL DB cluster that Includes one Aurora Replica The application's read performance degrades when there are more than 200 user connections. The number of user connections is approximately 180 on a consistent basis Occasionally, the number of user connections increases rapidly to more than 200
A SysOps administrator must implement a solution that will scale the application automatically as user demand increases or decreases.
Which solution will meet these requirements?

A. Modify the DB cluster by increasing the Aurora Replica instance size.
B. Modify the DB cluster by changing to serverless mode whenever the number of user connections exceeds 200.
C. Migrate to a new Aurora DB cluster that has multiple writer instances. Modify the application's database connection string.
D. Create an auto scaling policy that has a target value of 195 for the DatabaseConnections metric.

**Correct Answer: D**
**Section:**
**Explanation:**
Aurora Auto Scaling:
Aurora Auto Scaling adjusts the number of Aurora Replicas in response to changes in connectivity or workload.
Steps:
Go to the AWS Management Console.
Navigate to RDS and select the Aurora cluster.
Under 'Actions,' choose 'Add Aurora Replica' to initially add replicas if needed.
Go to the 'Auto Scaling' section and create an auto scaling policy.
Set the target value for the DatabaseConnections metric to 195.
Define the minimum and maximum number of replicas.
Save the configuration.
This ensures that the Aurora cluster scales automatically when the number of connections approaches the threshold, improving read performance.

**QUESTION 40**
A company runs a single-page web application on AWS The application uses Amazon CloudFront lo deliver static content from an Amazon S3 bucket origin The application also uses an Amazon Elastic Kubemetes Service (Amazon EKS) duster to serve API calls
Users sometimes report that the website is not operational, even when monitoring shows that the index page is reachable and that the EKS cluster is healthy. A SysOps administrator must Implement additional monitoring that can delect when the website is not operational before users report the problem.
Which solution will meet these requirements?

A. Create an Amazon CloudWatch Synthetics heartbeat monitor canary that points to the fully qualified domain name (FQDN) of the website.
B. Create an Amazon CloudWatch Synthetics API canary that monitors the availability of API endpoints from the EKS cluster.
C. Create an Amazon CloudWatch RUM app monitor that points to the fully qualified domain name (FQDN) of the website. Configure the app monitor to collect performance telemetry and JavaScript errors
D. Create an Amazon CloudWatch RUM app monitor that uses the API endpoints from the EKS cluster

**Correct Answer: A**
**Section:**
**Explanation:**
Amazon CloudWatch Synthetics:
CloudWatch Synthetics allows you to create canaries to monitor your endpoints and API calls, simulating user behavior to detect issues before users do.
Steps:
Go to the AWS Management Console.
Navigate to CloudWatch and select 'Synthetics.'
Click on 'Create canary.'
Choose 'Heartbeat monitoring' as the blueprint.
Configure the canary to point to the FQDN of the website.

Set the frequency and retention settings as per your requirement.
Create the canary.
This setup continuously checks the operational status of your website, alerting you if it becomes unreachable or has issues.

**QUESTION 41**
A company needs to monitor the disk utilization of Amazon Elastic Block Store (Amazon EBS) volumes The EBS volumes are attached to Amazon EC2 Linux Instances A SysOps administrator must set up an Amazon CloudWatch alarm that provides an alert when disk utilization increases to more than 80%.
Which combination of steps must the SysOps administrator lake lo meet these requirements? (Select THREE.)

A. Create an 1AM role that includes the Cloud Watch AgentServerPol icy AWS managed policy Attach me role to the instances
B. Create an 1AM role that includes the CloudWatchApplicationInsightsReadOnlyAccess AWS managed policy. Attach the role to the instances
C. Install and start the CloudWatch agent by using AWS Systems Manager or the command line
D. Install and start the CloudWatch agent by using an 1AM role. Attach the Cloud Watch AgentServerPolicy AWS managed policy to the role.
E. Configure a CloudWatch alarm to enter ALARM state when the disk_used_percent CloudWatch metric is greater than 80%.
F. Configure a CloudWatch alarm to enter ALARM state when the disk_used CloudWatch metric is greater than 80% or when the disk_free CloudWatch metric is less than 20%.

**Correct Answer: A, C, E**
**Section:**
**Explanation:**
Create an IAM role with the CloudWatchAgentServerPolicy:
This policy grants the necessary permissions for the CloudWatch agent to collect and send metrics.
Steps:
Go to the AWS Management Console.
Navigate to IAM and create a new role.
Choose 'EC2' as the trusted entity.
Attach the 'CloudWatchAgentServerPolicy' managed policy to the role.
Attach this IAM role to your EC2 instances.
Install and start the CloudWatch agent:
The CloudWatch agent must be installed and configured to collect disk utilization metrics.
Steps:
Use AWS Systems Manager or SSH to connect to your instances.
Install the CloudWatch agent using the following commands:
sudo yum install amazon-cloudwatch-agent
sudo /opt/aws/amazon-cloudwatch-agent/bin/amazon-cloudwatch-agent-config-wizard
sudo /opt/aws/amazon-cloudwatch-agent/bin/amazon-cloudwatch-agent-ctl -a fetch-config -m ec2 -c file:/path/to/your-config-file.json -s
Start the agent:
sudo systemctl start amazon-cloudwatch-agent
Configure a CloudWatch alarm:
Create an alarm based on the disk_used_percent metric.
Steps:
Go to the AWS Management Console.
Navigate to CloudWatch and select 'Alarms' from the left-hand menu.
Click on 'Create alarm.'
Select the disk_used_percent metric.
Set the threshold to 80% and configure the alarm actions (e.g., sending a notification).

**QUESTION 42**
A SysOps administrator is investigating a company's web application for performance problems The application runs on Amazon EC2 instances that are in an Auto Scaling group. The application receives large traffic increases at random times throughout the day. During periods of rapid traffic increases, the Auto Scaling group is not adding capacity fast enough. As a result, users are experiencing poor performance.

The company wants to minimize costs without adversely affecting the user experience when web traffic surges quickly. The company needs a solution that adds more capacity to me Auto Scaling group for larger traffic increases than for smaller traffic increases.
How should the SysOps administrator configure the Auto Scaling group to meet these requirements?

A. Create a simple scaling policy with settings to make larger adjustments in capacity when the system is under heavy load
B. Create a step scaling policy with settings to make larger adjustments in capacity when the system is under heavy load.
C. Create a target tracking scaling policy with settings to make larger adjustments in capacity when the system is under heavy load
D. Use Amazon EC2 Auto Scaling lifecycle hooks Adjust the Auto Scaling group's maximum number of instances after every scaling event

**Correct Answer: B**
**Section:**
**Explanation:**
Step Scaling Policy:
Step scaling policies allow you to define scaling actions based on different levels of CloudWatch alarms.
Steps:
Go to the AWS Management Console.
Navigate to EC2 Auto Scaling.
Select your Auto Scaling group.
Create or edit a scaling policy and choose 'Step scaling.'
Define different steps based on CloudWatch alarm thresholds (e.g., CPU usage or request count).
Configure larger adjustments for higher thresholds and smaller adjustments for lower thresholds.
Example Configuration:
For CPU > 80%, increase capacity by 4 instances.
For CPU > 60%, increase capacity by 2 instances.
For CPU > 40%, increase capacity by 1 instance.

**QUESTION 43**
A company runs an application on hundreds of Amazon EC2 instances in three Availability Zones The application calls a third-parly API over the public internet A SysOps administrator must provide the third party with a list of static IP addresses so that the third party can allow traffic from the application
Which solution will meet these requirements?

A. Add a NAT gateway in the public subnet of each Availability Zone. Make the NAT gateway the default route of all private subnets In those Availability Zones.
B. Allocate one Elastic IP address in each Availability Zone. Associate the Elastic IP address with all the instances in the Availability Zone
C. Place the instances behind a Network Load Balancer (NLB). Send the traffic to the interne! through the private IP address of the NLB
D. Update the main route table to send the traffic to the internet through an Elastic IP address that is assigned to each instance.

**Correct Answer: A**
**Section:**
**Explanation:**
NAT Gateway Setup:
A NAT gateway allows instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances.
Steps:
Go to the AWS Management Console.
Navigate to VPC and select 'NAT Gateways.'
Create a NAT gateway in the public subnet of each Availability Zone.
Allocate an Elastic IP address to each NAT gateway.
Update the route tables for the private subnets to route internet-bound traffic to the NAT gateways.

**QUESTION 44**

A company runs a high performance computing (HPC) application on an Amazon EC2 instance The company needs to scale this architecture to two or more EC2 instances. The EC2 instances wilt need to communicate with each other at high speeds with low latency to support the application.
The company wants to ensure that the network performance can support the required communication between the EC2 instances.
What should a SysOps administrator do to meet these requirements?

A. Create a cluster placement group. Back up the existing EC2 instance to an Amazon Machine Image (AMI). Restore the EC2 instance from the AMI into the placement group Launch the additional EC2 instances into the placement group
B. Back up the existing EC2 instance to an Amazon Machine Image (AMI). Create a launch template from the existing EC2 instance by specifying the AMI. Create an Auto Scaling group and configure the desired instance count.
C. Create a Network Load Balancer (NLB) and a target group. Launch the new EC2 instances and register them with the target group Register the existing EC2 instance with the target group. Pass all application traffic through the NLB.
D. Back up the existing EC2 Instance to an Amazon Machine Image (AMI). Create additional clones of the EC2 instance from the AMI in the same Availability Zone where the existing EC2 instance is located.

**Correct Answer: A**
**Section:**
**Explanation:**
Cluster Placement Group:
Cluster placement groups are used to ensure low-latency networking between EC2 instances. They place instances physically close to each other within the same Availability Zone.
Steps:
Go to the AWS Management Console.
Navigate to EC2 and select 'Placement Groups.'
Create a new cluster placement group.
Back up the existing EC2 instance to an AMI.
Launch new EC2 instances from the AMI into the cluster placement group.
Ensure all instances are in the same Availability Zone.

**QUESTION 45**
A company is uploading important files as objects to Amazon S3 The company needs to be informed if an object is corrupted during the upload
What should a SysOps administrator do to meet this requirement?

A. Pass the Content-Disposition value as a request body during the object upload.
B. Pass the Content-MD5 value as a request header during the object upload.
C. Pass x-amz-objecWock-mode as a request header during the object upload
D. Pass x-amz-server-side-encryption-customer-algorithm as a request body during the object upload.

**Correct Answer: B**
**Section:**
**Explanation:**
Content-MD5 Header:
The Content-MD5 header provides an MD5 checksum of the object being uploaded. Amazon S3 uses this checksum to verify the integrity of the object.
Steps:
When uploading an object to S3, calculate the MD5 checksum of the object.
Include the Content-MD5 header with the base64-encoded MD5 checksum value in the upload request.
This ensures that S3 can detect if the object is corrupted during the upload process.

**QUESTION 46**
A SysOps administrator needs to ensure that an Amazon RDS for PostgreSQL DB instance has available backups The DB instance has automated backups turned on with a backup retention period of 7 days. However, no automated backups for the DB instance have been created in the past month.
What could be the cause of the lack of automated backups?

A. The Amazon S3 bucket that stores the backups is full
B. The DB instance is in the STORAGE_FULL state
C. The DB instance is not configured for Multi-AZ.
D. The backup retention period must be 30 days.

**Correct Answer: B**
**Section:**
**Explanation:**
STORAGE_FULL State:
When an RDS instance is in the STORAGE_FULL state, automated backups cannot be performed because there is insufficient storage available.
Steps to Resolve:
Go to the AWS Management Console.
Navigate to RDS and select the DB instance.
Check the storage metrics to confirm the STORAGE_FULL state.
Increase the allocated storage for the DB instance to provide sufficient space for automated backups.

**QUESTION 47**
A company has a list of pre-appf oved Amazon Machine Images (AMIs) for developers lo use to launch Amazon EC2 instances However, developers are still launching EC2 instances from unapproved AMIs.
A SysOps administrator must implement a solution that automatically terminates any instances that are launched from unapproved AMIs.
Which solution will meet mis requirement?

A. Set up an AWS Config managed rule to check if instances are running from AMIs that are on the list of pre-approved AMIs. Configure an automatic remediation action so that an AWS Systems Manager Automation runbook terminates any instances that are noncompliant with the rule
B. Store the list of pre-approved AMIs in an Amazon DynamoDB global table that is replicated to all AWS Regions that the developers use. Create Regional EC2 launch templates. Configure the launch templates to check AMIs against the list and to terminate any instances that are not on the list
C. Select the Amazon CloudWatch metric that shows all running instances and the AMIs that the instances were launched from Create a CloudWatch alarm that terminates an instance if the metric shows the use of an unapproved AMI.
D. Create a custom Amazon Inspector finding to compare a running instance's AMI against the list of pre-approved AMIs Create an AWS Lambda function that terminates instances. Configure Amazon Inspector to report findings of unapproved AMIs to an Amazon Simple Queue Service (Amazon SQS) queue to invoke the Lambda function.

**Correct Answer: A**
**Section:**
**Explanation:**
AWS Config Managed Rule:
AWS Config can be used to assess, audit, and evaluate the configurations of AWS resources. The managed rule can check if instances are launched from approved AMIs.
Steps:
Go to the AWS Management Console.
Navigate to AWS Config.
Create a managed rule that checks for EC2 instances running approved AMIs.
Configure the rule to use a list of approved AMIs.
Automatic Remediation with Systems Manager Automation:
AWS Systems Manager Automation runbooks can automate the process of remediating non-compliant resources.
Steps:
Create a Systems Manager Automation runbook that terminates instances not running approved AMIs.
Attach the runbook to the AWS Config rule for automatic remediation.

**QUESTION 48**
Users of a company's internal web application recently experienced application performance issues for a brief period The application includes frontend web servers that run in an Amazon Elastic Kubernetes Service (Amazon

EKS) cluster The application also includes a bacKend Amazon Aurora PostgreSQL DB cluster that includes one DB instance.
A SysOps administrator determines that the source of the performance issues was high utilization of the DB cluster. The single writer instance experienced more than 90% utilization for 11 minutes The cause of the high utilization was an automated report that is scheduled to run one time each week
What should the SysOps administrator do to ensure that users do not experience performance Issues each week when the report runs?

A. Increase the size of the DB instance. Monitor the performance during the next scheduled run of the report
B. Add a reader instance. Change the database connection string of the report application to use the newly created reader instance.
C. Add another writer instance Change the database connection string of the report application to use the newly created writer instance.
D. Configure auto scaling for the DB cluster Set the minimum capacity units, maximum capacity units, and target utilization

**Correct Answer: A**
**Section:**
**Explanation:**
Increasing DB Instance Size:
Increasing the instance size provides more CPU and memory resources, which can help handle higher loads.
Steps:
Go to the AWS Management Console.
Navigate to RDS and select the DB instance.
Modify the instance to increase its size.
Apply the changes during the next maintenance window or immediately if it is a critical issue.
Monitoring Performance:
After resizing, monitor the instance during the next report run to ensure that it handles the load effectively.

**QUESTION 49**
A company is using AWS to deploy a critical application on a fleet of Amazon EC2 instances The company is rewriting the application because the application failed a security review The application will take 12 months to rewrite While this rewrite happens, the company needs to rotate IAM access keys that the application uses.
A SysOps administrator must implement an automated solution that finds and rotates IAM access Keys that are at least 30 days old. The solution must then continue to rotate the IAM access Keys every 30 days.
Which solution will meet this requirement with the MOST operational efficiency?

A. Use an AWS Config rule to identify IAM access Keys that are at least 30 days old. Configure AWS Config to invoKe an AWS Systems Manager Automation runbook to rotate the identified IAM access keys.
B. Use AWS Trusted Advisor to identify IAM access Keys that are at least 30 days old. Configure Trusted Advisor to invoke an AWS Systems Manager Automation runbook to rotate the identified IAM access keys
C. Create a script that checks the age of IAM access Keys and rotates them if they are at least 30 days old. Launch an EC2 instance. Schedule the script to run as a cron expression on the EC2 instance every day.
D. Create an AWS Lambda function that checks the age of IAM access keys and rotates them if they are at least 30 days old Use an Amazon EventBridge rule to invoke the Lambda function every time a new IAM access key is created.

**Correct Answer: D**
**Section:**
**Explanation:**
Lambda Function to Rotate IAM Access Keys:
A Lambda function can be used to automate the rotation of IAM access keys based on their age.
Steps:
Write a Lambda function that checks the age of IAM access keys.
The function should rotate keys that are at least 30 days old.
Deploy the Lambda function.
Amazon EventBridge Rule:
EventBridge can trigger the Lambda function periodically and when a new key is created.
Steps:
Create an EventBridge rule that triggers the Lambda function on a schedule (e.g., daily) and on IAM key creation events.

**QUESTION 50**
A company receives an alert from an Amazon CloudWatch alarm The alarm indicates that a web application that Is running on Amazon EC2 instances is not responding to requests The EC2 instances have a Red Hat Enterprise Linux operating system and are in an Auto Scaling group. The Auto Scaling group has a minimum capacity of 2 and a maximum capacity of 5.
An Investigation reveals that the web application is experiencing oul-of-memory errors. The company adds memory lo the web application and wants to track operating system memory utilization. A CloudWatch memory metric does not currently exist tor the EC2 Instances in the Auto Scaling group
What should a SysOps administrator do to provide a CloudWatch memory metric for the EC2 instances?

A. Use an Amazon Machine Image (AMI) that includes the CloudWatch agent.
B. Turn on CloudWatch detailed monitoring
C. Turn on Instance Metadata Service Version 2 (IMOSv2).
D. Use an Amazon Machine Image (AMI) that is based on Amazon Linux.

**Correct Answer: A**
**Section:**
**Explanation:**
Using an AMI with CloudWatch Agent:
The CloudWatch agent can collect memory utilization metrics and send them to CloudWatch.
Steps:
Create or use an existing AMI that includes the CloudWatch agent installed and configured.
Ensure the CloudWatch agent is configured to collect memory metrics.
Use this AMI for instances in the Auto Scaling group.

**QUESTION 51**
A company is using an Amazon CloudWatch alarm lo monitor the FreeLocalStorage metric for an Amazon Aurora PostgreSQL production database The alarm goes into ALARM state and indicates that the database is running low on temporary storage. A SysOps administrator discovers that a weekly report is using most of the temporary storage that is currently allocated.
What should the SysOps administrator do to solve this problem?

A. Turn on Aurora PostgreSQL query plan management.
B. Modify the configuration of the DB cluster to turn on storage auto scaling.
C. Add an Aurora read replica to the DB cluster. Modify the report lo use the new read replica.
D. Modify the DB instance class for each DB instance In the DB cluster to increase the instance size.

**Correct Answer: B**
**Section:**
**Explanation:**
Storage Auto Scaling:
Aurora storage auto scaling automatically increases the storage capacity of the database cluster when free storage space is running low.
Steps:
Go to the AWS Management Console.
Navigate to RDS and select your Aurora DB cluster.
Modify the DB cluster configuration to enable storage auto scaling.
Apply the changes.

**QUESTION 52**
A SysOps administrator is responsible for more than 50 Amazon EC2 instances mat are deployed in a single production AWS account The EC2 instances are running several different operating systems The company's standards require patching to be completed at least once a month.
The SysOps administrator wants to use AWS Systems Manager to reduce the number of hours the company spends on operating system patching each month.
Which combination of steps should the SysOps administrator take to meet these requirements? (Select THREE.)

A. Group similar EC2 instances together into resource groups by using AWS Resource Groups
B. Create a schedule in Systems Manager Patch Manager. Specify the appropriate resource group as the target
C. Specify Systems Manager Automation runbooks to patch the operating systems. Register the runbooks as tasks in the maintenance window. Specify the appropriate resource group as the target
D. Create a Systems Manager Automation runbook to monitor and control the state of the patches required. Apply the runbook to Systems Manager Patch Manager
E. Create a single Systems Manager maintenance window for each resource group.
F. Configure Systems Manager Fleet Manager to apply a Systems Manager Automation runbook to the appropriate resource group.

**Correct Answer: A, B, E**
**Section:**
**Explanation:**
Group EC2 Instances Using Resource Groups:
Resource groups help organize and manage AWS resources based on tags and other criteria.
Steps:
Go to the AWS Management Console.
Navigate to AWS Resource Groups.
Create resource groups for similar EC2 instances based on tags or other criteria.
Create a Schedule in Patch Manager:
AWS Systems Manager Patch Manager automates the process of patching managed instances.
Steps:
Go to the AWS Management Console.
Navigate to Systems Manager and select Patch Manager.
Create a patch baseline if not already created.
Create a schedule for patching and specify the resource group as the target.
Create Maintenance Windows for Resource Groups:
Maintenance windows define a period of time for performing administrative tasks on instances.
Steps:
Go to the AWS Management Console.
Navigate to Systems Manager and select Maintenance Windows.
Create a maintenance window for each resource group.
Specify tasks and targets (resource groups) for each maintenance window.

**QUESTION 53**
A company uses AWS Cloud Formation to deploy its infrastructure. The company recently retired an application. A cloud operations engineer initiates CloudFormation stack deletion, and the stack gets stuck in DELETE FAILED status.
A SysOps administrator discovers that the stack had deployed a security group. The security group is referenced by other security groups in the environment. The SysOps administrator needs to delete the stack without affecting other applications.
Which solution will meet these requirements m the MOST operationally efficient manner?

A. Create a new security group that has a different name Apply identical rules to the new security group. Replace all other security groups that reference the new security group. Delete the stack.
B. Create a CloudFormation change set to delete the security group. Deploy the change set.
C. Delete the stack again. Specify that the security group be retained.
D. Perform CloudFormation drift detection Delete the stack.

**Correct Answer: C**
**Section:**
**Explanation:**
Retain the Security Group:
When deleting a CloudFormation stack, you can specify resources to be retained instead of deleted.

Steps:
Go to the AWS Management Console.
Navigate to CloudFormation and select the stack.
Choose to delete the stack.
In the deletion options, specify that the security group should be retained.
This will delete the stack but keep the security group, ensuring no impact on other applications.

**QUESTION 54**
A company's architeclure team must receive immediate email notification whenever new Amazon EC2 Instances are launched In the company's main AWS production account
What should a SysOps administrator do to meet this requirement?

A. Create a user data script that sends an email message through a smarx 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 (he 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 SOS) 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 SOS 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.

**Correct Answer: B**
**Section:**
**Explanation:**
Create an SNS Topic and Subscription:
Amazon SNS allows you to send notifications to multiple endpoints.
Steps:
Go to the AWS Management Console.
Navigate to SNS and create a new topic.
Create a subscription for the topic using the email protocol.
Enter the architecture team's email address as the subscriber.
Create an EventBridge Rule:
Amazon EventBridge can monitor events and trigger actions.
Steps:
Go to the AWS Management Console.
Navigate to EventBridge.
Create a new rule that reacts to EC2 instance launch events.
Specify the SNS topic as the rule's target.

**QUESTION 55**
A SysOps administrator needs to update an AWS accoun1 name What should the SysOps administrator do to accomplish this goal?

A. Add the Administrator Access policy to the SysOps administrator's 1AM user.
B. Add the AWS_ConfigRole policy to the SysOps administrator's 1AM user.
C. Change the AWS account name through the AWS Trusted Advisor interface.
D. Sign in as the AWS account root user to make the change.

**Correct Answer: D**
**Section:**
**Explanation:**

Update AWS Account Name:
The AWS account name can only be changed by the root user of the account.
Steps:
Sign in to the AWS Management Console using the root user credentials.
Navigate to the 'My Account' page.
Update the account name field and save the changes.

**QUESTION 56**
A SysOps administrator needs to create a report that shows how many bytes are sent to and received from each target group member for an Application Load Balancer (ALB).
Which combination of steps should the SysOps administrator take to meet these requirements? (Select TWO.)

A. Enable access logging for the ALB. Save the logs to an Amazon S3 bucket.
B. Install the Amazon CloudWatch agent on the Instances in the target group.
C. Use Amazon Athena to query the ALB logs Query the table Use the received_bytes and senl_byt.es fields to calculate the total bytes grouped by the target:port field.
D. Use Amazon Athena to query the ALB logs Query the table. Use the received_bytes and sent_byt.es fields to calculate the total bytes grouped by the clientport field
E. Create an Amazon CloudWatch dashboard that shows the Sum statistic of the ProcessedBytes metric for the ALB.

**Correct Answer: A, C**
**Section:**
**Explanation:**
Enable Access Logging for the ALB:
Access logging provides detailed information about requests sent to your load balancer.
Steps:
Go to the AWS Management Console.
Navigate to EC2 and select 'Load Balancers.'
Select your Application Load Balancer.
Under the 'Attributes' tab, enable 'Access logs.'
Specify an S3 bucket where the logs will be saved.
Use Amazon Athena to Query the ALB Logs:
Athena allows you to run SQL queries on data stored in S3.
Steps:
Go to the AWS Management Console.
Navigate to Athena.
Create a table for the ALB logs using the appropriate schema.
Run queries to calculate the total bytes sent and received, grouped by the target field.
Example query:
SELECT target, SUM(received_bytes) as total_received, SUM(sent_bytes) as total_sent
FROM alb_logs
GROUP BY target, port

**QUESTION 57**
A SysOps administrator must ensure that all of a company's current and future Amazon S3 buckets have logging enabled If an S3 bucket does not have logging enabled an automated process must enable logging for the S3 bucket.
Which solution will meet these requirements?

A. Use AWS Trusted Advisor 10 perform a check for S3 buckets that do not have logging enabled Configure the check to enable logging for S3 buckets that do not have logging enabled.
B. Configure an S3 bucket policy that requires all current and future S3 buckets to have logging enabled
C. Use the s3-bucket-logging-enabled AWS Config managed rule. Add a remediation action that uses an AWS Lambda function to enable logging.

D. Use the s3-bucket-logging-enabled AWS Config managed rule. Add a remediation action that uses the AWS-ConfigureS3BucketLoggmg AWS Systems Manager Automation runbook to enable logging.

**Correct Answer: C, D**
**Section:**
**Explanation:**
AWS Config Managed Rule for S3 Logging:
The s3-bucket-logging-enabled AWS Config rule checks whether S3 buckets have logging enabled.
Steps:
Go to the AWS Management Console.
Navigate to AWS Config.
Create a rule using s3-bucket-logging-enabled.
Add a remediation action using an AWS Lambda function or Systems Manager Automation runbook.
Using AWS Lambda for Remediation:
Create a Lambda function that enables logging on S3 buckets.
Steps:
Write a Lambda function in Python or Node.js to enable logging.
Configure the function to trigger on non-compliant buckets.
Using AWS Systems Manager Automation:
The AWS-ConfigureS3BucketLogging runbook automates enabling logging.
Steps:
Go to the AWS Management Console.
Navigate to Systems Manager.
Create an Automation document or use the existing AWS-ConfigureS3BucketLogging runbook.
Configure the remediation action to use this runbook.

**QUESTION 58**
After creating a presigned URL for an S3 object, users can no longer access the file after a few days.

A. The presigned URL's expiration date and time have passed.
B. The SysOps administrator's access key is no longer valid.
C. The S3 bucket's Block Public Access settings are enabled.
D. The S3 object's ACL does not include READ access for the All Users group.
E. The S3 object's ACL does not include READ_ACP access for the All Users group.

**Correct Answer: A, B**
**Section:**
**Explanation:**
The presigned URL expiration is the most common reason for access issues after some time. Additionally, if the SysOps administrator's access key (used to generate the presigned URL) is invalid, the URL will no longer be usable. Block Public Access or ACL settings are irrelevant to presigned URLs.

**QUESTION 59**
The company needs to increase IOPS for two EC2 instances with gp2 volumes to support an upcoming promotion with higher I/O requirements.

A. Migrate the attached EBS volumes to Throughput Optimized HDD (st1) EBS volumes.
B. Configure Amazon ElastiCache integration on the EC2 instances.
C. Migrate the workload to two storage optimized EC2 instances.
D. Migrate the attached EBS volumes to General Purpose SSD (gp3) EBS volumes. Provision the appropriate IOPS.

**Correct Answer: D**
**Section:**
**Explanation:**
Migrating to gp3 volumes allows for customizable IOPS at a lower cost than gp2, meeting the requirement for higher IOPS during the promotion. Throughput Optimized HDD (st1) volumes do not support high IOPS, and ElastiCache does not address I/O for EBS volumes.

**QUESTION 60**
The SysOps administrator needs to create a key policy that grants data engineers least privilege access to decrypt and read data from an S3 bucket encrypted with KMS.

A. 'kms:ReEncrypt*', 'kms:GenerateDataKey*', 'kms:Encrypt', 'kms:DescribeKey'
B. 'kms:ListAliases', 'kms:GetKeyPolicy', 'kms:Describe*', 'kms:Decrypt'
C. 'kms:ListAliases', 'kms:DescribeKey', 'kms:Decrypt'
D. 'kms:Update*', 'kms:TagResource', 'kms:Revoke*', 'kms:Put*', 'kms:List*', 'kms:Get*', 'kms:Enable*', 'kms:Disable*', 'kms:Describe*', 'kms:Delete*', 'kms:Create*', kms:CancelKeyDeletion

**Correct Answer: C**
**Section:**
**Explanation:**
The least privilege required for reading encrypted data involves kms:Decrypt to decrypt, kms:DescribeKey to understand key properties, and kms:ListAliases if needed to identify the key alias.

**QUESTION 61**
The SysOps administrator must restart the web server if specific errors are detected in logs on EC2 instances behind a load balancer.

A. Install the Amazon CloudWatch agent on the EC2 instances.
B. Create an AWS CloudTrail metric filter for the web logs. Configure an alarm for the specific errors.
C. Create an Amazon CloudWatch metric filter for the web logs. Configure an alarm for the specific errors.
D. Publish alarm findings to Amazon Simple Email Service (Amazon SES). Invoke an AWS Lambda function to restart the web server software.
E. Create an Amazon EventBridge rule that responds to the alarm. Configure the rule to invoke an AWS Systems Manager Automation runbook to restart the web server software.
F. Create an Amazon Simple Notification Service (Amazon SNS) notification that responds to the alarm. Configure the notification to invoke an AWS Systems Manager Automation runbook to restart the web server software.

**Correct Answer: A, C, E**
**Section:**
**Explanation:**
Installing the CloudWatch agent enables log monitoring, and a CloudWatch metric filter allows alerting on specific errors. Using EventBridge to trigger a Systems Manager Automation runbook automates the restart of the web server, creating an efficient and automated solution.

**QUESTION 62**
The company requires a disaster recovery solution for an Aurora PostgreSQL database with a 20-second RPO.

A. Reconfigure the database to be an Aurora global database. Set the RPO to 20 seconds.
B. Reconfigure the database to be an Aurora Serverless v2 database with an Aurora Replica in a separate Availability Zone. Set the replica lag to 20 seconds.
C. Modify the database to use a Multi-AZ cluster that has two readable standby instances in separate Availability Zones. Add an Aurora Replica in a separate Availability Zone. Set the replica lag to 20 seconds.

**Correct Answer: A**
**Section:**
**Explanation:**
Aurora Global Databases are designed for cross-Region disaster recovery with very low RPO, meeting the 20-second requirement. Setting up Aurora as a global database with the correct configuration ensures low-latency replication and rapid failover, making it ideal for compliance with strict disaster recovery requirements.

**QUESTION 63**
The company needs a shared file solution for EC2 Windows instances in a Multi-AZ deployment that uses native Windows storage capabilities and maximizes consistency.

A. Create an Amazon FSx for Windows File Server Multi-AZ file system. Map file shares on the instances by using the file system's DNS name.
B. Grant the instances access to a shared Amazon S3 bucket. Use Windows Task Scheduler to synchronize the contents of the S3 bucket locally to each instance periodically.
C. Create an Amazon Elastic File System (Amazon EFS) file system that uses the EFS Standard storage class. Mount the file system to the instances by using the file system's DNS name and the EFS mount helper.
D. Create a new Amazon Elastic Block Store (Amazon EBS) Multi-Attach volume. Attach the EBS volume as an additional drive to each instance.

**Correct Answer: A**
**Section:**
**Explanation:**
Amazon FSx for Windows File Server provides a fully managed, highly available, and native Windows file system compatible with the SMB protocol, ideal for Windows workloads requiring shared access.
Multi-AZ File System: Ensures high availability across multiple Availability Zones.
Native Windows Capabilities: Allows instances to map file shares and access files using Windows storage features, offering strong consistency and performance for shared files.
Other options, like Amazon S3 and Amazon EFS, either lack native Windows integration or do not offer the desired consistency and high availability for shared file systems in a Windows environment.

**QUESTION 64**
To automatically reboot an EC2 instance when disk usage reaches 100%, a solution with minimal operational overhead is needed.

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 Simple Email Service (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.

**Correct Answer: A**
**Section:**
**Explanation:**
Using a CloudWatch alarm with an EventBridge rule provides an automated, direct way to reboot the EC2 instance without extra components like SES or Lambda. This is a straightforward approach with low operational overhead.

**QUESTION 65**
The SysOps administrator needs to prevent launching EC2 instances without a specific tag in the application OU.

A. Create an IAM group that has a policy allowing ec2:RunInstances when the CostCenter-Project tag is present. Place all IAM users in this group.
B. Create a service control policy (SCP) that denies ec2:RunInstances when the CostCenter-Project tag is missing. Attach the SCP to the application OU.
C. Create an IAM role with a policy that allows ec2:RunInstances when the CostCenter-Project tag is present. Attach the IAM role to users in the application OU accounts.
D. Create a service control policy (SCP) that denies ec2:RunInstances when the CostCenter-Project tag is missing. Attach the SCP to the root OU.

**Correct Answer: B**
**Section:**
**Explanation:**
An SCP applied to the application OU that denies ec2:RunInstances when the CostCenter-Project tag is missing ensures that all accounts in the OU adhere to the tagging policy. This approach is centralized and applies only to the intended OU.

**QUESTION 66**
A company has an AWS Config rule that identifies open SSH ports in security groups. The rule has an automatic remediation action to delete the SSH inbound rule for noncompliant security groups. However, business units require SSH access and can provide a list of trusted IPs to restrict access.

A. Create a new AWS Systems Manager Automation runbook that adds an IP set to the security group's inbound rule. Update the AWS Config rule to change the automatic remediation action to use the new runbook.
B. Create a new AWS Systems Manager Automation runbook that updates the security group's inbound rule with the IP addresses from the business units. Update the AWS Config rule to change the automatic remediation action to use the new runbook.
C. Create an AWS Lambda function that adds an IP set to the security group's inbound rule. Update the AWS Config rule to change the automatic remediation action to use the Lambda function.
D. Create an AWS Lambda function that updates the security group's inbound rule with the IP addresses from the business units. Update the AWS Config rule to change the automatic remediation action to use the Lambda function.

**Correct Answer: B**
**Section:**
**Explanation:**
The problem requires modifying the inbound SSH rule to restrict access to a list of trusted IPs instead of deleting it entirely. AWS Config rules can be configured with automatic remediation actions using either Systems Manager Automation runbooks or Lambda functions. However, AWS Systems Manager Automation runbooks are often more appropriate for managing infrastructure changes like security group modifications because they are reusable, parameterized, and easier to audit.
Create a Systems Manager Automation runbook: This runbook will contain steps to add or modify the existing security group rule, allowing SSH access only from the specified IP addresses.
Update the AWS Config rule: Modify the Config rule to call this new runbook for its automatic remediation. This will prevent deletion of the SSH rule and instead update it based on the IP list.

**QUESTION 67**
A company's application on EC2 instances relies on a Single-AZ RDS for MySQL DB instance. The SysOps administrator needs to ensure failover to minimize downtime.

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.

**Correct Answer: A**
**Section:**
**Explanation:**
To ensure high availability and failover for RDS, converting the instance to a Multi-AZ deployment is the best practice. Multi-AZ configurations provide automated standby in a different Availability Zone, automatically failing over to the standby in case of instance or Availability Zone issues.
Modify DB instance: AWS allows for seamless conversion of an existing Single-AZ RDS instance to a Multi-AZ deployment, making it more resilient to outages without requiring significant application changes.
Failover mechanism: In Multi-AZ, failover is managed automatically by AWS, minimizing application downtime.

**QUESTION 68**
To manage Auto Scaling group instances that have OS vulnerabilities, the SysOps administrator needs an automated patching solution.

A. Use AWS Systems Manager Patch Manager to patch the instances during a scheduled maintenance window. In the AWS-RunPatchBaseline document, ensure that the RebootOption parameter is set to RebootIfNeeded.
B. Use EC2 Image Builder pipelines on a schedule to create new Amazon Machine Images (AMIs) and new launch templates that reference the new AMIs. Use the instance refresh feature for EC2 Auto Scaling to replace instances.
C. Use AWS Config to scan for operating system vulnerabilities and to patch instances when the instance status changes to NON_COMPLIANT. Send an Amazon Simple Notification Service (Amazon SNS) notification to an operations team to reboot the instances during off-peak hours.
D. In the Auto Scaling launch template, provide an Amazon Machine Image (AMI) ID for an AWS-provided base image. Update the user data with a shell script to download and install patches.

**Correct Answer: A**
**Section:**
**Explanation:**
Using AWS Systems Manager Patch Manager with a maintenance window is a best practice for automating OS patch management across instances in an Auto Scaling group.
Patch Manager: Allows for scheduled patching according to maintenance windows, ensuring minimal impact on application uptime.
RebootOption parameter: Setting this to RebootIfNeeded ensures patches are applied fully when a reboot is required.

AWS-RunPatchBaseline: This document automates the patching process and can be customized based on compliance requirements.

**QUESTION 69**
The company is experiencing increased message load from the frontend to the backend, causing message loss due to backend capacity limitations.

A. Redevelop the backend application as a series of AWS Lambda functions.
B. Implement an Amazon Kinesis data stream to replace the backend application.
C. Implement an Application Load Balancer to distribute message traffic across the backend application instances.
D. Implement an Amazon Simple Queue Service (Amazon SQS) queue between the frontend and backend components.

**Correct Answer: D**
**Section:**
**Explanation:**
To handle the increased message load with minimal operational effort, implementing an Amazon Simple Queue Service (SQS) queue between the frontend and backend is an ideal solution. SQS decouples the frontend and backend by queuing messages, enabling the backend to process messages at its own pace without losing any.
SQS Queue: Acts as a buffer, ensuring messages are not lost if the backend application cannot immediately process them.
Decoupling: With SQS, the frontend can continue sending messages without concern for the backend's processing speed, providing a scalable solution with minimal management requirements.
Low Operational Overhead: SQS is fully managed, reducing the need for infrastructure management.

**QUESTION 70**
A company plans to launch a static website on its domain example com and subdomain www example.com using Amazon S3. How should the SysOps administrator meet this requirement?

A. Create one S3 bucket named example.com for both the domain and subdomain.
B. Create one S3 bucket with a wildcard named '.example.com tor both the domain and subdomain.
C. Create two S3 buckets named example.com and www.example.com. Configure the subdomain bucket to redirect requests to the domain bucket.
D. Create two S3 buckets named http//example.com and http//" exampte.com. Configure thewildcard (') bucket to redirect requests to the domain bucket.

**Correct Answer: C**
**Section:**

**QUESTION 71**
A SysOps administrator configuring AWS Client VPN to connect use's on a corporate network to AWS resources mat are running in a VPC According to compliance requirements, only traffic that is destined for the VPC can travel across the VPN tunnel.
How should the SysOps administrator configure Client VPN to meet these requirements?

A. Associate the Client VPN endpoint with a private subnet that has an internet route through a NAT gateway.
B. On the Client VPN endpoint, turns on the split-tunnel option.
C. On the Client VPN endpoint, specify DNS server IP addresses
D. Select a private certificate to use as the identity certificate tor the VPN client.

**Correct Answer: C**
**Section:**

**QUESTION 72**
A SysOps administrator has Nocked public access to all company Amazon S3 buckets. The SysOps administrator wants to be notified when an S3 bucket becomes publicly readable in the future. What is the MOST operationally efficient way to meet this requirement?

A. Create an AWS Lambda function that periodically checks the public access settings for each S3 bucket. Set up Amazon Simple Notification Service (Amazon SNS) to send notifications.

B. Create a cron script that uses the S3 API to check the public access settings for each S3 bucket. Set up Amazon Simple Notification Service (Amazon SNS) to send notifications
C. Enable S3 Event notified tons for each S3 bucket. Subscribe S3 Event Notifications to an Amazon Simple Notification Service (Amazon SNS) topic.
D. Enable the s3-bucket-public-read-prohibited managed rule in AWS Config. Subscribe the AWS Config rule to an Amazon Simple Notification Service (Amazon SNS) topic.

**Correct Answer: D**
**Section:**

**QUESTION 73**
A company wants to create an automated solution for all accounts managed by AWS Organizations to detect any worry groups that urn 0.0.0.0/0 as the source address for inbound traffic. The company also wants to automatically remediate any noncompliant security groups by restricting access to a specific CIDR block corresponds with the company's intranet.

A. Create an AWS Config rule to detect noncompliant security groups. Set up automatic remediation to change the 0.0.0.0/0 source address to the approved CIDK block.
B. Create an IAM policy to deny the creation of security groups that have 0.0.0.0/0 as the source address Attach this 1AM policy to every user in the company.
C. Create an AWS Lambda function to inspect now and existing security groups check for a noncompliant 0.0.0.0A) source address and change the source address to the approved CIDR block.
D. Create a service control policy (SCP) for the organizational unit (OU) to deny the creation of security groups that have the 0.0.0.0/0 source address. Set up automatic remediation to change Vie 0.0.0.0/0 source address to the approved CIDR block.

**Correct Answer: A**
**Section:**

**QUESTION 74**
A company runs an application on an Amazon EC2 instance A SysOps administrator creates an Auto Scaling group and an Application Load Balancer (ALB) to handle an increase in demand However, the EC2 instances are failing tie health check.
What should the SysOps administrator do to troubleshoot this issue?

A. Verity that the Auto Scaling group is configured to use all AWS Regions.
B. Verily that the application is running on the protocol and the port that the listens is expecting.
C. Verify the listener priority in the ALB Change the priority if necessary.
D. Verify the maximum number of instances in the Auto Scaling group Change the number if necessary

**Correct Answer: B**
**Section:**

**QUESTION 75**
A SysOps administrator needs to create alerts that are based on the read and write metrics of Amazon Elastic Block Store (Amazon EBS) volumes that are attached to an Amazon EC2 instance. The SysOps administrator creates and enables Amazon CloudWatch alarms for the DiskReadBytes metric and the DiskWriteBytes metric.
A custom monitoring tool that is installed on the EC2 instance with the same alarm configuration indicates that the volume metrics have exceeded the threshold. However, the CloudWatch alarms were not in ALARM state.
Which action will ensure that the CloudWatch alarms function correctly?

A. Install and configure the CloudWatch agent on the EC2 instance to capture the desired metrics.
B. Install and configure AWS Systems Manager Agent on the EC2 instance to capture the desired metrics.
C. Reconfigure the CloudWatch alarms to use the VolumeReadBytes metric and the VolumeWriteBytes metric for the EBS volumes.
D. Reconfigure the CloudWatch alarms to use the VolumeReadBytes metric and the VolumeWriteBytes metric for the EC2 instance.

**Correct Answer: A**
**Section:**

**QUESTION 76**

A company has created a NAT gateway in a public subnet in a VPC. The VPC also contains a private subnet that includes Amazon EC2 instances. The EC2 instances use the NAT gateway to access the internet to download patches and updates. The company has configured a VPC flow log for the elastic network interface of the NAT gateway. The company is publishing the output to Amazon CloudWatch Logs. A SysOps administrator must identify the top five internet destinations that the EC2 instances in the private subnet communicate with for downloads. What should the SysOps administrator do to meet this requirement in the MOST operationally efficient way?

A. Use AWS CloudTrail Insights events to identify the top five internet destinations.
B. Use Amazon CloudFront standard logs (access logs) to identify the top five internet destinations.
C. Use CloudWatch Logs Insights to identify the top five internet destinations.
D. Change the flow log to publish logs to Amazon S3. Use Amazon Athena to query the log files in Amazon S3.

**Correct Answer: C**
**Section:**

**QUESTION 77**
A company has an application that is deployed 10 two AWS Regions in an active-passive configuration. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) in each Region. The instances are in an Amazon EC2 Auto Scaling group in each Region. The application uses an Amazon Route 53 hosted zone (or DNS. A SysOps administrator needs to configure automatic failover to the secondary Region. What should the SysOps administrator do to meet these requirements?

A. Configure Route 53 alias records that point to each ALB. Choose a failover routing policy. Set Evaluate Target Health to Yes.
B. Configure CNAME records that point to each ALB. Choose a failover routing policy. Set Evaluate Target Health to Yes.
C. Configure Elastic Load Balancing (ELB) health checks for the Auto Scaling group. Add a target group to the ALB in the primary Region. Include the EC2 instances in the secondary Region as targets.
D. Configure EC2 health checks for the Auto Scaling group. Add a target group to the ALB in the primary Region. Include the EC2 instances in the secondary Region as targets.

**Correct Answer: A**
**Section:**

**QUESTION 78**
A company has a compliance requirement that no security groups can allow SSH ports to be open to all IP addresses. A SysOps administrator must implement a solution that will notify the company's SysOps team when a security group rule violates this requirement. The solution also must remediate the security group rule automatically. Which solution will meet these requirements?

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function when a security group changes. Configure the Lambda function to evaluate the security group for compliance, remove all inbound security group rules on all ports, and notify the SysOps team if the security group is noncompliant.
B. Create an AWS CloudTrail metric filter for security group changes. Create an Amazon CloudWatch alarm to notify the SysOps team through an Amazon Simple Notification Service (Amazon SNS) topic when (he metric is greater than 0. Subscribe an AWS Lambda function to the SNS topic to remediate
the security group rule by removing the rule.
C. Activate the AWS Config restricted-ssh managed rule. Add automatic remediation to the AWS Config rule by using the AWS Systems Manager Automation AWSDisablePublicAccessForSecurityGroup runbook. Create an Amazon EventBridge (Amazon CloudWatch
Events) rule to notify the SysOps team when the rule is noncompliant.
D. Create an AWS CloudTrail metric filter for security group changes. Create an Amazon CloudWatch alarm for when the metric is greater than 0. Add an AWS Systems Manager action to the CloudWatch alarm to suspend the security group by using the Systems Manager Automation AWSDisablePublicAccessForSecurityGroup runbook when the alarm is in ALARM state. Add an Amazon Simple Notification Service (Amazon SNS) topic as a second target to notify the SysOps team.

**Correct Answer: C**
**Section:**

**QUESTION 79**
A company has an application that customers use to search for records on a website. The application's data is stored in an Amazon Aurora DB cluster. The application's usage varies by season and by day of the week. The website's popularity is increasing, and the website is experiencing slower performance because of increased load on the DB cluster during periods of peak activity. The application logs show that the performance issues occur when users are searching for information. The same search is rarely performed multiple times.

A SysOps administrator must improve the performance of the platform by using a solution that maximizes resource efficiency. Which solution will meet these requirements?

A. Deploy an Amazon ElastiCache for Redis cluster in front of the DB cluster. Modify the application to check the cache before the application issues new queries to the database. Add the results of any queries to the cache.
B. Deploy an Aurora Replica for the DB cluster. Modify the application to use the reader endpoint for search operations. Use Aurora Auto Scaling to scale the number of replicas based on load. Most Voted
C. Use Provisioned IOPS on the storage volumes that support the DB cluster to improve performance sufficiently to support the peak load on the application.
D. Increase the instance size in the DB cluster to a size that is sufficient to support the peak load on the application. Use Aurora Auto Scaling to scale the instance size based on load.

**Correct Answer: B**
**Section:**
**Explanation:**
https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/aurora-replicasadding.html

**QUESTION 80**
A company hosts a web application on an Amazon EC2 instance. The web server logs are published to Amazon CloudWatch Logs. The log events have the same structure and include the HTTP response codes that are associated with the user requests. The company needs to monitor the number of times that the web server returns an HTTP 404 response. What is the MOST operationally efficient solution that meets these requirements?

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

**Correct Answer: A**
**Section:**
**Explanation:**
This is the most operationally efficient solution that meets the requirements, as it will allow the company to monitor the number of times that the web server returns an HTTP 404 response in realtime. The other solutions (creating a CloudWatch Logs subscription filter, an AWS Lambda function, or a script) will require additional steps and resources to monitor the number of times that the web server returns an HTTP 404 response.
A metric filter allows you to search for specific terms, phrases, or values in your log events, and then to create a metric based on the number of occurrences of those search terms. This allows you to create a CloudWatch Metric that can be used to create alarms and dashboards, which can be used to monitor the number of HTTP 404 responses returned by the web server.

**QUESTION 81**
A Sysops administrator has created an Amazon EC2 instance using an AWS CloudFormation template in the us-east-I Region. The administrator finds that this template has failed to create an EC2 instance in the us-west-2 Region. What is one cause for this failure?

A. Resource tags defined in the CloudFormation template are specific to the us-east-I Region.
B. The Amazon Machine Image (AMI) ID referenced in the CloudFormation template could not be found in the us-west-2 Region.
C. The cfn-init script did not run during resource provisioning in the us-west-2 Region.
D. The IAM user was not created in the specified Region.

**Correct Answer: B**
**Section:**
**Explanation:**
One possible cause for the failure of the CloudFormation template to create an EC2 instance in the us-west-2 Region is that the Amazon Machine Image (AMI) ID referenced in the template could not be found in the us-west-2 Region. This could be due to the fact that the AMI is not available in that region, or the credentials used to access the AMI were not configured properly. The other options (resource tags defined in the CloudFormation template are specific to the us-east-I Region, the cfninit script did not run during resource provisioning in the us-west-2 Region, and the IAM user was not created in the specified Region) are not valid causes for this failure.

**QUESTION 82**
A company plans to deploy a database on an Amazon Aurora MySQL DB cluster. The database will store data for a demonstration environment. The data must be reset on a daily basis. What is the MOST operationally efficient solution that meets these requirements?

A. Create a manual snapshot of the DB cluster after the data has been populated. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to restore the snapshot and then delete the previous DB cluster.
B. Enable the Backtrack feature during the creation of the DB cluster. Specify a target backtrack window of 48 hours. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to perform a backtrack operation.
C. Export a manual snapshot of the DB cluster to an Amazon S3 bucket after the data has been populated. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to restore the snapshot from Amazon S3.
D. Set the DB cluster backup retention period to 2 days. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to restore the DB cluster to a point in time and then delete the previous DB cluster.

**Correct Answer: D**
**Section:**
**Explanation:**
Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to restore the DB cluster to a point in time and then delete the previous DB cluster. This is the most operationally efficient solution that meets the requirements, as it will allow the company to reset the database on a daily basis without having to manually take and restore snapshots. The other solutions (creating a manual snapshot of the DB cluster, enabling the Backtrack feature, or exporting a manual snapshot of the DB cluster to Amazon S3) will require additional steps and resources to reset the database on a daily basis.

**QUESTION 83**
A company has a memory-intensive application that runs on a fleet of Amazon EC2 instances behind an Elastic Load Balancer (ELB). The instances run in an Auto Scaling group. A Sysops administrator must ensure that the application can scale based on the number of users that connect to the application. Which solution will meet these requirements?

A. Create a scaling policy that will scale the application based on the ActiveConnectionCount Amazon CloudWatch metric that is generated from the ELB.
B. Create a scaling policy that will scale the application based on the mem used Amazon CloudWatch metric that is generated from the ELB.
C. Create a scheduled scaling policy to increase the number of EC2 instances in the Auto Scaling group to support additional connections.
D. Create and deploy a script on the ELB to expose the number of connected users as a custom Amazon CloudWatch metric. Create a scaling policy that uses the metric.

**Correct Answer: A**
**Section:**
**Explanation:**

**QUESTION 84**
A company is using Amazon CloudFront to serve static content for its web application to its users.
The CloudFront distribution uses an existing on-premises website as a custom origin.
The company requires the use of TLS between CloudFront and the origin server. This configuration has worked as expected for several months. However, users are now experiencing HTTP 502 (Bad Gateway) errors when they view webpages that include content from the CloudFront distribution.
What should a SysOps administrator do to resolve this problem?

A. Examine the expiration date on the certificate on the origin site. Validate that the certificate has not expired. Replace the certificate if necessary.
B. Examine the hostname on the certificate on the origin site. Validate that the hostname matches one of the hostnames on the CloudFront distribution. Replace the certificate if necessary.
C. Examine the firewall rules that are associated with the origin server. Validate that port 443 is open for inbound traffic from the internet. Create an inbound rule if necessary.
D. Examine the network ACL rules that are associated with the CloudFront distribution. Validate that port 443 is open for outbound traffic to the origin server. Create an outbound rule if necessary.

**Correct Answer: A**
**Section:**
**Explanation:**
HTTP 502 errors from CloudFront can occur because of the following reasons:
There's an SSL negotiation failure because the origin is using SSL/TLS protocols and ciphers that aren't supported by CloudFront. There's an SSL negotiation failure because the SSL certificate on the origin is expired or invalid, or because the certificate chain is invalid. There's a host header mismatch in the SSL negotiation between your CloudFront distribution and the custom origin. The custom origin isn't responding on the ports specified in the origin settings of the CloudFront distribution. The custom origin is ending the connection to CloudFront too quickly.

https://aws.amazon.com/premiumsupport/knowledge-center/resolve-cloudfront-connection-error/

**QUESTION 85**
A company runs hundreds of Amazon EC2 instances in a single AWS Region. Each EC2 instance has two attached 1 GiB General Purpose SSD (gp2) Amazon Elastic Block Store (Amazon EBS) volumes. A critical workload is using all the available IOPS capacity on the EBS volumes.
According to company policy, the company cannot change instance types or EBS volume types without completing lengthy acceptance tests to validate that the company's applications will function properly. A SysOps administrator needs to increase the I/O performance of the EBS volumes as quickly as possible. Which action should the SysOps administrator take to meet these requirements?

A. Increase the size of the 1 GiB EBS volumes.
B. Add two additional elastic network interfaces on each EC2 instance.
C. Turn on Transfer Acceleration on the EBS volumes in the Region.
D. Add all the EC2 instances to a cluster placement group.

**Correct Answer: A**
**Section:**
**Explanation:**
Increasing the size of the 1 GiB EBS volumes will increase the IOPS capacity of the volumes, which will improve the I/O performance of the EBS volumes. This option does not require any changes to the instance types or EBS volume types, so it can be done quickly without the need for lengthy acceptance tests to validate that the company's applications will function properly. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/requesting-ebs-volumemodifications.html

**QUESTION 86**
A company is implementing a monitoring solution that is based on machine learning. The monitoring solution consumes Amazon EventBridge (Amazon CloudWatch Events) events that are generated by Amazon EC2 Auto Scaling. The monitoring solution provides detection of anomalous behavior such as unanticipated scaling events and is configured as an EventBridge (CloudWatch Events) API destination.
During initial testing, the company discovers that the monitoring solution is not receiving events. However, Amazon CloudWatch is showing that the EventBridge (CloudWatch Events) rule is being invoked. A SysOps administrator must implement a solution to retrieve client error details to help resolve this issue. Which solution will meet these requirements with the LEAST operational effort?

A. Create an EventBridge (CloudWatch Events) archive for the event pattern to replay the events.
Increase the logging on the monitoring solution. Use replay to invoke the monitoring solution.
Examine the error details.
B. Add an Amazon Simple Queue Service (Amazon SQS) standard queue as a dead-letter queue for the target. Process the messages in the dead-letter queue to retrieve error details.
C. Create a second EventBridge (CloudWatch Events) rule for the same event pattern to target an AWS Lambda function. Configure the Lambda function to invoke the monitoring solution and to record the results to Amazon CloudWatch Logs. Examine the errors in the logs.
D. Configure the EventBridge (CloudWatch Events) rule to send error messages to an Amazon Simple Notification Service (Amazon SNS) topic.

**Correct Answer: A**
**Section:**
**Explanation:**
"In EventBridge, you can create an archive of events so that you can easily replay them at a later time. For example, you might want to replay events to recover from errors or to validate new functionality in your application."
https://docs.aws.amazon.com/eventbridge/latest/userguide/ebarchive.html

**QUESTION 87**
A company is storing backups in an Amazon S3 bucket. The backups must not be deleted for at least 3 months after the backups are created. What should a SysOps administrator do to meet this requirement?

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.

**Correct Answer: B**

**Section:**
**Explanation:**
https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html Ingovernance mode, users can't overwrite or delete an object version or alter its lock settingsunless they have special permissions. With governance mode, you protect objects against beingdeleted by most users, but you can still grant some users permission to alter the retentionsettings or delete the object if necessary. You can also use governance mode to test retention-period settings before creating a compliance-mode retention period. In compliance mode, aprotected object version can't be overwritten or deleted by any user, including the root user inyour AWS account. When an object is locked in compliance mode, its retention mode can't bechanged, and its retention period can't be shortened. Compliance mode helps ensure that anobject version can't be overwritten or deleted for the duration of the retention period.

**QUESTION 88**
A SysOps administrator needs to track the costs of data transfer between AWS Regions. The SysOps administrator must implement a solution to send alerts to an email distribution list when transfer costs reach 75% of a specific threshold.
What should the SysOps administrator do to meet these requirements?

A. Create an AWS Cost and Usage Report. Analyze the results in Amazon Athena. Configure an alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic when costs reach 75% of the threshold. Subscribe the email distribution list to the topic.
B. Create an Amazon CloudWatch billing alarm to detect when costs reach 75% of the threshold.
Configure the alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the email distribution list to the topic.
C. Use AWS Budgets to create a cost budget for data transfer costs. Set an alert at 75% of the budgeted amount. Configure the budget to send a notification to the email distribution list when costs reach 75% of the threshold.
D. Set up a VPC flow log. Set up a subscription filter to an AWS Lambda function to analyze data transfer. Configure the Lambda function to send a notification to the email distribution list when costs reach 75% of the threshold.

**Correct Answer: B**
**Section:**
**Explanation:**
The reason is that it uses the Amazon CloudWatch billing alarm which is a built-in service specifically designed to monitor and alert on cost usage of your AWS account, which makes it a more suitable solution for this use case. The alarm can be configured to detect when costs reach 75% of the threshold and when it is triggered, it can publish a message to an Amazon Simple Notification Service (Amazon SNS) topic. The email distribution list can be subscribed to the topic, so that they will receive the alerts when costs reach 75% of the threshold. AWS Budgets allows you to track and manage your costs, but it doesn't specifically focus on data transfer costs between regions, and it might not provide as much granularity as CloudWatch Alarms.

**QUESTION 89**
A company needs to archive all audit logs for 10 years. The company must protect the logs from any future edits. Which solution will meet these requirements?

A. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume. Configure AWS Key Management Service (AWS KMS) encryption.
B. Store the data in an Amazon S3 Glacier vault. Configure a vault lock policy for write-once, readmany (WORM) access.
C. Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA). Configure server-side encryption.
D. Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA). Configure multi-factor authentication (MFA).

**Correct Answer: B**
**Section:**
**Explanation:**
To meet the requirements of the workload, a company should store the data in an Amazon S3 Glacier vault and configure a vault lock policy for write-once, read-many (WORM) access. This will ensure that the data is stored securely and cannot be edited in the future. The other solutions (storing the data in an Amazon Elastic Block Store (Amazon EBS) volume and configuring AWS Key Management Service (AWS KMS) encryption, storing the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA) and configuring server-side encryption, or storing the data in Amazon S3 StandardInfrequent Access (S3 Standard-IA) and configuring multi-factor authentication (MFA)) will not meet the requirements, as they do not provide a way to protect the audit logs from future edits. https://docs.aws.amazon.com/zh_tw/AmazonS3/latest/userguide/object-lock.html

**QUESTION 90**
A company's AWS Lambda function is experiencing performance issues. The Lambda function performs many CPU-intensive operations. The Lambda function is not running fast enough and is creating bottlenecks in the system.
What should a SysOps administrator do to resolve this issue?

A. In the CPU launch options for the Lambda function, activate hyperthreading.
B. Turn off the AWS managed encryption.
C. Increase the amount of memory for the Lambda function.
D. Load the required code into a custom layer.

**Correct Answer: C**
**Section:**
**Explanation:**
Increasing the amount of memory for the Lambda function will help to improve the performance of the function. This is because the Lambda function is CPU-intensive and increasing the memory will give it access to more CPU resources and help it run faster. The other options (activating hyperthreading in the CPU launch options for the Lambda function, turning off the AWS managed encryption, and loading the required code into a custom layer) will not help to improve the performance of the Lambda function and are not the correct solutions for this issue. https://docs.aws.amazon.com/lambda/latest/dg/configuration-functioncommon.html#configuration-memory-console

**QUESTION 91**
A company is attempting to manage its costs in the AWS Cloud. A SysOps administrator needs specific company-defined tags that are assigned to resources to appear on the billing report. What should the SysOps administrator do to meet this requirement?

A. Activate the tags as AWS generated cost allocation tags.
B. Activate the tags as user-defined cost allocation tags.
C. Create a new cost category. Select the account billing dimension.
D. Create a new AWS Cost and Usage Report. Include the resource IDs.

**Correct Answer: B**
**Section:**
**Explanation:**
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/custom-tags.html "User-defined tags are tags that you define, create, and apply to resources. After you have created and applied the user-defined tags, you can activate by using the Billing and Cost Management console for cost allocation tracking. " To meet this requirement, the SysOps administrator should activate the company-defined tags as user-defined cost allocation tags. This will ensure that the tags appear on the billing report and that the resources can be tracked with the specific tags. The other options (activating the tags as AWS generated cost allocation tags, creating a new cost category and selecting the account billing dimension, and creating a new AWS Cost and Usage Report and including the resource IDs) will not meet the requirements and are not the correct solutions for this issue.

**QUESTION 92**
A company's application currently uses an IAM role that allows all access to all AWS services. A SysOps administrator must ensure that the company's IAM policies allow only the permissions that the application requires. How can the SysOps administrator create a policy to meet this requirement?

A. Turn on AWS CloudTrail. Generate a policy by using AWS Security Hub.
B. Turn on Amazon EventBridge (Amazon CloudWatch Events). Generate a policy by using AWS Identity and Access Management Access Analyzer.
C. Use the AWS CLI to run the get-generated-policy command in AWS Identity and Access Management Access Analyzer.
D. Turn on AWS CloudTrail. Generate a policy by using AWS Identity and Access Management Access Analyzer.

**Correct Answer: D**
**Section:**
**Explanation:**
Generate a policy by using AWS Identity and Access Management Access Analyzer. AWS CloudTrail is a service that records all API calls made on your account. You can use this data to generate a policy with AWS Identity and Access Management Access Analyzer that only allows the permissions that the application requires. This will ensure that the application only has the necessary permissions and will protect the company from any unauthorized access.
https://docs.aws.amazon.com/IAM/latest/UserGuide/what-is-access-analyzer.html#what-is-accessanalyzer-policy-generation

**QUESTION 93**
A company updates its security policy to clarify cloud hosting arrangements for regulated workloads. Workloads that are identified as sensitive must run on hardware that is not shared with other customers or with other AWS accounts within the company. Which solution will ensure compliance with this policy?

A. Deploy workloads only to Dedicated Hosts.
B. Deploy workloads only to Dedicated Instances.
C. Deploy workloads only to Reserved Instances.
D. Place all instances in a dedicated placement group.

**Correct Answer: A**
**Section:**
**Explanation:**
Dedicated Hosts are physical servers that are dedicated to a single customer, ensuring that the customer's workloads are not shared with other customers or with other AWS accounts within the company. This will ensure that the company's security policy is followed and that sensitive workloads are running on hardware that is not shared with other customers or with other AWS accounts within the company.

**QUESTION 94**
A company needs to implement a managed file system to host Windows file shares for users on premises. Resources in the AWS Cloud also need access to the data on these file shares. A SysOps administrator needs to present the user file shares on premises and make the user file shares available on AWS with minimum latency. What should the SysOps administrator do to meet these requirements?

A. Set up an Amazon S3 File Gateway.
B. Set up an AWS Direct Connect connection.
C. Use AWS DataSync to automate data transfers between the existing file servers and AWS.
D. Set up an Amazon FSx File Gateway.



**Correct Answer: D**
**Section:**
**Explanation:**
Amazon FSx provides a fully managed file system that is optimized for Windows-based workloads and can be used to create file shares that can be accessed both on premises and in the AWS Cloud. The file shares that are created in Amazon FSx are highly available and can be accessed with low latency. Additionally, Amazon FSx supports Windows-based authentication, making it easy to integrate with existing Windows user accounts.

**QUESTION 95**
A company is hosting applications on Amazon EC2 instances. The company is hosting a database on an Amazon RDS for PostgreSQL DB instance. The company requires all connections to the DB instance to be encrypted. What should a SysOps administrator do to meet this requirement?

A. Allow SSL connections to the database by using an inbound security group rule.
B. Encrypt the database by using an AWS Key Management Service (AWS KMS) encryption key.
C. Enforce SSL connections to the database by using a custom parameter group.
D. Patch the database with SSL/TLS by using a custom PostgreSQL extension.

**Correct Answer: C**
**Section:**
**Explanation:**
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/PostgreSQL.Concepts.General.SSL.htm
Amazon RDS supports SSL/TLS encryption for connections to the database, and this can be enabled by creating a custom parameter group and setting the rds.force_ssl parameter to 1. This will ensure that all connections to the database are encrypted, protecting the data and maintaining compliance with the company's requirements.l

**QUESTION 96**
A company recently purchased Savings Plans. The company wants to receive email notification when the company's utilization drops below 90% for a given day. Which solution will meet this requirement?

A. Create an Amazon CloudWatch alarm to monitor the Savings Plan check in AWS Trusted Advisor.
Configure an Amazon Simple Queue Service (Amazon SQS) queue for email notification when the utilization drops below 90% for a given day.
B. Create an Amazon CloudWatch alarm to monitor the SavingsPlansUtilization metric under the AWS/SavingsPlans namespace in CloudWatch. Configure an Amazon Simple Queue Service (Amazon SQS) queue for email notification when the utilization drops below 90% for a given day.
C. Create a Savings Plans alert to monitor the daily utilization of the Savings Plans. Configure an Amazon Simple Notification Service (Amazon SNS) topic for email notification when the utilization drops below 90% for a given day.
D. Use AWS Budgets to create a Savings Plans budget to track the daily utilization of the Savings Plans. Configure an Amazon Simple Notification Service (Amazon SNS) topic for email notification when the utilization drops below 90% for a given day.

**Correct Answer: D**
**Section:**
**Explanation:**
AWS Budgets can be used to create a Savings Plans budget and track the daily utilization of the company's Savings Plans. By creating a budget, it will trigger an action when the utilization drops below 90%, which in this case will be to send an email notification via an Amazon SNS topic. This will ensure that the company is notified when their Savings Plans utilization drops below 90%, allowing them to take action if necessary.
Reference: [1]Â https://docs.aws.amazon.com/savingsplans/latest/userguide/sp-usingBudgets.html

**QUESTION 97**
A SysOps administrator receives an alert from Amazon GuardDuty about suspicious network activity on an Amazon EC2 instance. The GuardDuty finding lists a new external IP address as a traffic destination. The SysOps administrator does not recognize the external IP address. The SysOps administrator must block traffic to the external IP address that GuardDuty identified. Which solution will meet this requirement?

A. Create a new security group to block traffic to the external IP address. Assign the new security group to the EC2 instance.
B. Use VPC flow logs with Amazon Athena to block traffic to the external IP address.
C. Create a network ACL. Add an outbound deny rule for traffic to the external IP address.
D. Create a new security group to block traffic to the external IP address. Assign the new security group to the entire VPC.

**Correct Answer: C**
**Section:**
**Explanation:**
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-network-acls.html

**QUESTION 98**
A company's reporting job that used to run in 15 minutes is now taking an hour to run. An application generates the reports. The application runs on Amazon EC2 instances and extracts data from an Amazon RDS for MySQL database. A SysOps administrator checks the Amazon CloudWatch dashboard for the RDS instance and notices that the Read IOPS metrics are high, even when the reports are not running. The SysOps administrator needs to improve the performance and the availability of the RDS instance.
Which solution will meet these requirements?

A. Configure an Amazon ElastiCache cluster in front of the RDS instance. Update the reporting job to query the ElastiCache cluster.
B. Deploy an RDS read replica. Update the reporting job to query the reader endpoint.
C. Create an Amazon CloudFront distribution. Set the RDS instance as the origin. Update the reporting job to query the CloudFront distribution.
D. Increase the size of the RDS instance.

**Correct Answer: B**
**Section:**
**Explanation:**
Using an RDS read replica will improve the performance and availability of the RDS instance by offloading read queries to the replica. This will also ensure that the reporting job completes in a timely manner and does not affect the performance of other queries that might be running on the RDS instance. Additionally, updating the reporting job to query the reader endpoint will ensure that all read queries are directed to the read replica.
Reference: [1]Â https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

**QUESTION 99**
A company's SysOps administrator regularly checks the AWS Personal Health Dashboard in each of the company's accounts. The accounts are part of an organization in AWS Organizations. The company recently added 10 more accounts to the organization. The SysOps administrator must consolidate the alerts from each account's Personal Health Dashboard. Which solution will meet this requirement with the LEAST amount of effort?

A. Enable organizational view in AWS Health.
B. Configure the Personal 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 to write all events to an Amazon DynamoDB table.
D. Use the AWS Health API to write events to an Amazon DynamoDB table.

**Correct Answer: A**
**Section:**
**Explanation:**
Enabling the organizational view in AWS Health will allow the SysOps administrator to consolidate the alerts from each account's Personal Health Dashboard. It will also provide the administrator with a single view of all the accounts in the organization, allowing them to easily monitor the health of all the accounts in the organization.
Reference: [1]Â https://aws.amazon.com/premiumsupport/knowledge-center/organizational-viewhealth-dashboard/

**QUESTION 100**
A company runs an application on Amazon EC2 instances. The EC2 instances are in an Auto Scaling group and run behind an Application Load Balancer (ALB). The application experiences errors when total requests exceed 100 requests per second. A SysOps administrator must collect information about total requests for a 2-week period to determine when requests exceeded this threshold. What should the SysOps administrator do to collect this data?

A. Use the ALB's RequestCount metric. Configure a time range of 2 weeks and a period of 1 minute. Examine the chart to determine peak traffic times and volumes.
B. Use Amazon CloudWatch metric math to generate a sum of request counts for all the EC2 instances over a 2-week period. Sort by a 1-minute interval.
C. Create Amazon CloudWatch custom metrics on the EC2 launch configuration templates to create aggregated request metrics across all the EC2 instances.
D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule. Configure an EC2 event matching pattern that creates a metric that is based on EC2 requests. Display the data in a graph.

**Correct Answer: A**
**Section:**
**Explanation:**
Using the ALB's RequestCount metric will allow the SysOps administrator to collect information about total requests for a 2-week period and determine when requests exceeded the threshold of 100 requests per second. Configuring a time range of 2 weeks and a period of 1 minute will ensure that the data can be accurately examined to determine peak traffic times and volumes.

**QUESTION 101**
A company recently migrated its application to a VPC on AWS. An AWS Site-to-Site VPN connection connects the company's on-premises network to the VPC. The application retrieves customer data from another system that resides on premises. The application uses an on-premises DNS server to resolve domain records. After the migration, the application is not able to connect to the customer data because of name resolution errors.
Which solution will give the application the ability to resolve the internal domain names?

A. Launch EC2 instances in the VPC. On the EC2 instances, deploy a custom DNS forwarder that forwards all DNS requests to the on-premises DNS server. Create an Amazon Route 53 private hosted zone that uses the EC2 instances for name servers.
B. Create an Amazon Route 53 Resolver outbound endpoint. Configure the outbound endpoint to forward DNS queries against the on-premises domain to the on-premises DNS server.
C. Set up two AWS Direct Connect connections between the AWS environment and the on-premises network. Set up a link aggregation group (LAG) that includes the two connections. Change the VPC resolver address to point to the on-premises DNS server.
D. Create an Amazon Route 53 public hosted zone for the on-premises domain. Configure the network ACLs to forward DNS requests against the on-premises domain to the Route 53 public hosted zone.

**Correct Answer: B**
**Section:**
**Explanation:**
https://docs.aws.amazon.com/zh_tw/Route53/latest/DeveloperGuide/resolver-forwardingoutbound-queries.html

**QUESTION 102**
A SysOps administrator creates two VPCs, VPC1 and VPC2, in a company's AWS account The SysOps administrator deploys a Linux Amazon EC2 instance in VPC1 and deploys an Amazon RDS for MySQL DB instance in VPC2. The DB instance is deployed in a private subnet. An application that runs on the EC2 instance needs to connect to the database. What should the SysOps administrator do to give the EC2 instance the ability to connect to the database?

A. Enter the DB instance connection string into the VPC1 route table.
B. Configure VPC peering between the two VPCs.
C. Add the same IPv4 CIDR range for both VPCs.
D. Connect to the DB instance by using the DB instance's public IP address.

**Correct Answer: B**
**Section:**
**Explanation:**
VPC peering allows two VPCs to communicate with each other securely. By configuring VPC peering between the two VPCs, the SysOps administrator will be able to give the EC2 instance in VPC1 the ability to connect to the database in VPC2. Once the VPC peering is configured, the EC2 instance will be able to communicate with the database using the private IP address of the DB instance in the private subnet.

**QUESTION 103**
A company needs to take an inventory of applications that are running on multiple Amazon EC2 instances. The company has configured users and roles with the appropriate permissions for AWS Systems Manager. An updated version of Systems Manager Agent has been installed and is running \on every instance. While configuring an inventory collection, a SysOps administrator discovers that not all the instances in a single subnet are managed by Systems Manager.
What must the SysOps administrator do to fix this issue?

A. Ensure that all the EC2 instances have the correct tags for Systems Manager access.
B. Configure AWS Identity and Access Management Access Analyzer to determine and automatically remediate the issue.
C. Ensure that all the EC2 instances have an instance profile with Systems Manager access.
D. Configure Systems Manager to use an interface VPC endpoint.

**Correct Answer: C**
**Section:**
**Explanation:**
Ensuring that all the EC2 instances have an instance profile with Systems Manager access is the most effective way to fix this issue. Having an instance profile with Systems Manager access will allow the SysOps administrator to configure the inventory collection for all the instances in the subnet, regardless of whether or not they are managed by Systems Manager.

**QUESTION 104**
A company hosts an application on an Amazon EC2 instance in a single AWS Region. The application requires support for non-HTTP TCP traffic and HTTP traffic. The company wants to deliver content with low latency by leveraging the AWS network. The company also wants to implement an Auto Scaling group with an Elastic Load Balancer. How should a SysOps administrator meet these requirements?

A. Create an Auto Scaling group with an Application Load Balancer (ALB). Add an Amazon CloudFront distribution with the ALB as the origin.
B. Create an Auto Scaling group with an Application Load Balancer (ALB). Add an accelerator with AWS Global Accelerator with the ALB as an endpoint.
C. Create an Auto Scaling group with a Network Load Balancer (NLB). Add an Amazon CloudFront distribution with the NLB as the origin.
D. Create an Auto Scaling group with a Network Load Balancer (NLB). Add an accelerator with AWS Global Accelerator with the NLB as an endpoint.

**Correct Answer: D**
**Section:**
**Explanation:**

**QUESTION 105**

A SysOps administrator is managing a Memcached cluster in Amazon ElastiCache. The cluster has been heavily used recently, and the administrator wants to use a larger instance type with more memory.
What should the administrator use to make this change?

A. Use the ModifycacheCluster API and specify a new cacheNodeType.
B. Use the createcacheciuster API and specify a new cacheNodeType.
C. Use the Modi fyCacheParameterGcoup API and specify a new CacheNodeType.
D. Use the Rebootcacheclustcr API and specify a new CacheNodeType.

**Correct Answer: A**
**Section:**
**Explanation:**
To upgrade the instance type of a Memcached cluster in Amazon ElastiCache due to increased usage and the need for more memory:
ModifyCacheCluster API: Utilize the ModifyCacheCluster API call. This API allows you to change various settings of an existing cache cluster, including the instance type, which is referred to as cacheNodeType.
Instance Upgrade: Specify a new, larger cacheNodeType that provides more memory. This upgrade will involve a brief interruption as nodes are replaced with the larger type, but it is necessary to accommodate the increased load and memory requirements.
Cluster Availability: Ensure that the Memcached cluster is configured for minimal downtime during this change. The upgrade process is handled by ElastiCache, and the new nodes will join the cluster with more memory capacity.
This approach enables you to effectively scale up the resources available to your Memcached cluster, enhancing its performance and capacity to handle larger workloads.

**QUESTION 106**
A SysOps administrator is examining the following AWS CloudFormation template:

```
AWSTemplateFormatVersion: '2010-09-09'
Description: 'Creates an EC2 Instance'
Resources:
    EC2Instance:
      Type: AWS::EC2::Instance
      Properties:
        ImageId: ami-79fd7eee
        InstanceType: m5n.large
        SubnetId: subnet-1abc3d3fg
        PrivateDnsName: ip-10-24-34-0.ec2.internal
        Tags:
          - Key: Name
            Value: !Sub "${AWS::StackName} Instance"
```

Why will the stack creation fail?

A. The Outputs section of the Cloud Formation template was omitted.
B. The Parameters section of the CtoudFormation template was omitted.
C. The PnvateDnsName cannot be set from a CloudFormation template.
D. The VPC was not specified in the CloudFormation template.

**Correct Answer: C**
**Section:**
**Explanation:**
In AWS CloudFormation, the PrivateDnsName property of an EC2 instance cannot be directly set within the template. This property is automatically assigned by AWS when the instance is launched within a VPC and is associated with the private IP address of the instance. The attempt to explicitly set PrivateDnsName in a CloudFormation template will result in an error, causing the stack creation to fail. Therefore, option C is correct. For reference, the AWS documentation on EC2 instances in CloudFormation does not list PrivateDnsName as a configurable property AWS CloudFormation User Guide.

**QUESTION 107**

A SysOps administrator wants to securely share an object from a private Amazon S3 bucket with a group of users who do not have an AWS account. What is the MOST operationally efficient solution that will meet this requirement?

A. Attach an S3 bucket policy that only allows object downloads from the users' IP addresses.
B. Create an 1AM role that has access to the object. Instruct the users to assume the role.
C. Create an 1AM user that has access to the object. Share the credentials with the users.
D. Generate a presigned URL for the object. Share the URL with the users.

**Correct Answer: D**
**Section:**
**Explanation:**
The most operationally efficient and secure method to share an object from a private Amazon S3 bucket with users who do not have an AWS account is by generating a presigned URL. This URL grants temporary access to the object and can be limited by time, ensuring that users can only access the S3 object during a specified window. This does not require managing network configurations or sharing credentials, making it a secure and simple solution. Option D is therefore the correct answer. Reference to this method can be found in the AWS S3 documentation on presigned URLs Amazon S3 Presigned URLs.

**QUESTION 108**
A company's social media application has strict data residency requirements. The company wants to use Amazon Route 53 to provide the application with DNS services. A SysOps administrator must implement a solution that routes requests to a defined list of AWS Regions. The routing must be based on the user's location. Which solution will meet these requirements?

A. Configure a Route 53 latency routing policy.
B. Configure a Route 53 multivalue answer routing policy.
C. Configure a Route 53 geolocation routing policy.
D. Configure a Route 53 IP-based routing policy.

**Correct Answer: C**
**Section:**
**Explanation:**
For routing based on the user's geographic location to comply with data residency requirements, the best solution is to use Amazon Route 53 geolocation routing policy. This policy allows you to configure DNS responses based on the geographic location of the user, ensuring that requests are directed to specific AWS Regions that align with the company's data residency requirements. Option C is correct. The AWS Route 53 documentation provides details on implementing geolocation routing policies Amazon Route 53 Geolocation Routing.

**QUESTION 109**
A company runs its applications on a large number of Amazon EC2 instances. A SysOps administrator must implement a solution to notify the operations team whenever an EC2 instance slate changes.
What is the MOST operationally efficient solution that meets these requirements?

A. Create a script that captures instance state changes and publishes a notification to an Amazon Simple Notification Service (Amazon SNS) topic. Use AWS Systems Manager Run Command to run the script on all EC2 instances.
B. Create an Amazon EventBridge event rule that captures EC2 instance state changes. Set an Amazon Simple Notification Service (Amazon SNS) topic as the target.
C. Create an Amazon EventBridge event rule that captures EC2 instance state changes. Set as the target an AWS Lambda function that publishes a notification to an Amazon Simple Notification Service (Amazon SNS) topic.
D. Create an AWS Config custom rule that evaluates instance state changes with automatic remediation. Use the rule to invoke an AWS Lambda function that publishes a notification to an Amazon Simple Notification Service (Amazon SNS) topic.

**Correct Answer: B**
**Section:**
**Explanation:**
The most operationally efficient way to monitor state changes in EC2 instances and notify the operations team is by using Amazon EventBridge. EventBridge can be configured with a rule that listens for state change events from EC2 instances. These events can then be directed to an Amazon Simple Notification Service (Amazon SNS) topic, which will distribute the notification to the relevant parties. This solution does not require deploying additional scripts or functions, thereby enhancing operational efficiency. Option B is correct. For more details, see the Amazon EventBridge documentation Amazon EventBridge.

**QUESTION 110**
A company is running Amazon EC2 On-Demand Instances in an Auto Scaling group. The instances process messages from an Amazon Simple Queue Service (Amazon SQS) queue. The Auto Scaling group is set to scale based on the number of messages in the queue. Messages can take up to 12 hours to process completely. A SysOps administrator must ensure that instances are not interrupted during message processing.
What should the SysOps administrator do to meet these requirements?

A. Enable instance scale-in protection for the specific instance in the Auto Scaling group at the start of message processing by calling the Amazon EC2 Auto Scaling API from the processing script. Disable instance scale-in protection after message processing is complete by calling the Amazon EC2 Auto Scaling API from the processing script.
B. Set the Auto Scaling group's termination policy to OldestInstance.
C. Set the Auto Scaling group's termination policy to OldestLaunchConfiguration.
D. Suspend the Launch and Terminate scaling processes for the specific instance in the Auto Scaling group at the start of message processing by calling the Amazon EC2 Auto Scaling API from the processing script. Resume the scaling processes after message processing is complete by calling the Amazon EC2 Auto Scaling API from the processing script.

**Correct Answer: A**
**Section:**
**Explanation:**
# Enable instance scale-in protection for specific instance.
aws autoscaling set-instance-protection --instance-ids i-5f2e8a0d --auto-scaling-group-name my-asg --protected-from-scale-in
# Disable instance scale-in protection for the specified instance.
aws autoscaling set-instance-protection --instance-ids i-5f2e8a0d --auto-scaling-group-name my-asg --no-protected-from-scale-in
https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-instance-protection.html
To ensure that EC2 instances in an Auto Scaling group are not interrupted during message processing, the most effective method is to implement scale-in protection for the instances while they are actively processing messages. This can be done programmatically by modifying the Auto Scaling group's settings using the Amazon EC2 Auto Scaling API.
Starting Message Processing: When an instance begins processing a message, your application should make an API call to enable scale-in protection. This is done using the SetInstanceProtection action, setting the ProtectedFromScaleIn parameter to true for that specific instance.
Completing Message Processing: Once the message has been processed, another API call should be made to disable scale-in protection. This is done by calling the SetInstanceProtection action again, but this time setting the ProtectedFromScaleIn parameter to false.
This method ensures that while messages are being processed, the instances are not terminated by the Auto Scaling group regardless of any scale-in activities that might be triggered by other parameters like CPU utilization or a decrease in the number of messages in the queue.
AWS Documentation
Reference: You can refer to the AWS documentation on managing instance scale-in protection in Auto Scaling groups for more details: Instance Scale-In Protection.

**QUESTION 111**
A company is managing a website with a global user base hosted on Amazon EC2 with an Application Load Balancer (ALB). To reduce the load on the web servers, a SysOps administrator configures an Amazon CloudFront distribution with the ALB as the origin. After a week of monitoring the solution, the administrator notices that requests are still being served by the ALB and there is no change in the web server load.
What are possible causes for this problem? (Choose two.)

A. CloudFront does not have the ALB configured as the origin access identity.
B. The DNS is still pointing to the ALB instead of the CloudFront distribution.
C. The ALB security group is not permitting inbound traffic from CloudFront.
D. The default, minimum, and maximum Time to Live (TTL) are set to 0 seconds on the CloudFront distribution.
E. The target groups associated with the ALB are configured for sticky sessions.

**Correct Answer: B, D**
**Section:**
**Explanation:**
To effectively use Amazon CloudFront as a content delivery network for an application using an Application Load Balancer as the origin, several configuration steps need to be correctly implemented:
DNS Configuration: Ensure that the DNS records for the domain serving the content point to the CloudFront distribution's DNS name rather than directly to the ALB. If the DNS still points to the ALB, users' requests will bypass CloudFront, leading directly to the ALB and maintaining the existing load on your web servers.

TTL Settings: The Time to Live (TTL) settings in the CloudFront distribution dictate how long the content is cached in CloudFront edge locations before CloudFront fetches a fresh copy from the origin. If the TTL values are set to 0, it means that CloudFront does not cache the content at all, resulting in each user request being forwarded to the ALB, which does not reduce the load.
AWS Documentation
Reference: For more information on DNS and TTL configurations for CloudFront, you can refer to the following AWS documentation:
Configuring DNS
CloudFront TTL Settings.

**QUESTION 112**
A company has migrated its application to AWS. The company will host the application on Amazon EC2 instances of multiple instance families.
During initial testing, a SysOps administrator identifies performance issues on selected EC2 instances. The company has a strict budget allocation policy, so the
SysOps administrator must use the right resource types with the performance characteristics to match the workload.
What should the SysOps administrator do to meet this requirement?

A. Purchase regional Reserved Instances (RIs) for immediate cost savings. Review and take action on the EC2 rightsizing recommendations in Cost Explorer. Exchange the RIs for the optimal instance family after rightsizing.
B. Purchase zonal Reserved Instances (RIs) for the existing instances. Monitor the RI utilization in the AWS Billing and Cost Management console. Make adjustments to instance sizes to optimize utilization.
C. Review and take action on AWS Compute Optimizer recommendations. Purchase Compute Savings Plans to reduce the cost that is required to run the compute resources. Most Voted
D. Review resource utilization metrics in the AWS Cost and Usage Report. Rightsize the EC2 instances. Create On-Demand Capacity Reservations for the rightsized resources.

**Correct Answer: C**
**Section:**
**Explanation:**
When managing performance and cost for EC2 instances across different families, the following steps are recommended:
Utilize AWS Compute Optimizer: This service provides recommendations for EC2 instances based on historical usage patterns and existing configurations. It helps identify optimal EC2 instance types and sizes that could deliver better performance and cost savings for your specific workload.
Implement Compute Savings Plans: After determining the most suitable instance types and sizes through Compute Optimizer, purchasing Compute Savings Plans can offer significant cost savings. These savings plans apply to any instance family across any region, providing flexibility and cost efficiency without upfront commitment to specific instance types.
AWS Documentation
Reference: Further details can be found in the AWS documentation on Compute Optimizer and Compute Savings Plans:
AWS Compute Optimizer
AWS Compute Savings Plans.

**QUESTION 113**
A Sysops administrator wants to share a copy of a production database with a migration account. The production database is hosted on an Amazon RDS DB instance and is encrypted at rest with an AWS Key Management Service (AWS KMS) key that has an alias of
What must the Sysops administrator do to meet these requirements with the LEAST administrative overhead?

A. Take a snapshot of the RDS DB instance in the production account. Amend the KMS key policy of the production-rds-key KMS key to give access to the migration account's root user. Share the snapshot with the migration account.
B. Create an RDS read replica in the migration account. Configure the KMS key policy to replicate the production-rds-key KMS key to the migration account.
C. Take a snapshot of the RDS DB instance in the production account. Share the snapshot with the migration account. In the migration account, create a new KMS key that has an identical alias.
D. Use native database toolsets to export the RDS DB instance to Amazon S3. Create an S3 bucket and an S3 bucket policy for cross-account access between the production account and the migration account. Use native database toolsets to import the database from Amazon S3 to a new RDS DB instance.

**Correct Answer: A**
**Section:**
**Explanation:**
To share an encrypted Amazon RDS DB instance snapshot across accounts, the least administrative overhead involves directly managing permissions on the AWS KMS key and sharing the snapshot. Here's how to do it:
Take a Snapshot: Initiate a snapshot of your Amazon RDS DB instance in the production account. This captures the current state of the database.
Modify KMS Key Policy: Adjust the policy of the KMS key used for encryption (identified by the alias 'production-rds-key') to grant the kms:Decrypt permission to the migration account's root user. This step is crucial as it

allows the migration account to use the same encryption key to decrypt the snapshot.
Share the Snapshot: Share the newly created snapshot with the migration account using the RDS console or AWS CLI. The migration account will now be able to see and use this snapshot to create a new RDS instance.
AWS Documentation
Reference: You can refer to the AWS documentation on sharing encrypted snapshots: Sharing Encrypted Snapshots.

**QUESTION 114**
A company manages its production applications across several AWS accounts. The company hosts the production applications on Amazon EC2 instances that run Amazon Linux 2. The EC2 instances are spread across multiple VPCs. Each VPC uses its own Amazon Route 53 private hosted zone for private DNS.
A VPC from Account A needs to resolve private DNS records from a private hosted zone that is associated with a different VPC in Account B.
What should a SysOps administrator do to meet these requirements?

A. In Account A, create an AWS Systems Manager document that updates the /etc/resolv.conf file across all EC2 instances to point to the AWS provided default DNS resolver for the VPC in Account B.
B. In Account A, create an AWS CloudFormation template that associates the private hosted zone from Account B with the private hosted zone in Account A.
C. In Account A, use the AWS CLI to create a VPC association authorization. When the association is created, use the AWS CLI in Account B to associate the VPC from Account A with the private hosted zone in Account B.
D. In Account B, use the AWS CLI to create a VPC association authorization. When the association is created, use the AWS CLI in Account A to associate the VPC from Account B with the private hosted zone in Account A.

**Correct Answer: D**
**Section:**
**Explanation:**
To resolve DNS across VPCs in different accounts, you should:
Authorization: In Account B, initiate a VPC association authorization for the private hosted zone. This action allows another AWS account to associate a VPC with this hosted zone.
Association: In Account A, after receiving the authorization from Account B, associate its VPC with the private hosted zone that exists in Account B. This step will enable EC2 instances within the VPC in Account A to resolve DNS records hosted in Account B.
AWS Documentation
Reference: AWS provides detailed guidance on associating VPCs with private hosted zones across accounts in their documentation: Associating VPCs and Private Hosted Zones Across Accounts.

**QUESTION 115**
A SysOps administrator needs to implement a backup strategy for Amazon EC2 resources and Amazon RDS resources. The backup strategy must meet the following retention requirements:
* Daily backups: must be kept for 6 days
* Weekly backups: must be kept for 4 weeks:
* Monthly backups: must be kept for 11 months
* Yearly backups: must be kept for 7 years
Which backup strategy will meet these requirements with the LEAST administrative effort?

A. Use Amazon Data Lifecycle Manager to create an Amazon Elastic Block Store (Amazon EBS) snapshot policy. Create tags on each resource that needs to be backed up. Create multiple schedules according to the requirements within the policy. Set the appropriate frequency and retention period.
B. Use AWS Backup to create a new backup plan for each retention requirement with a backup frequency of daily, weekly, monthly, or yearly. Set the retention period to match the requirement. Create tags on each resource that needs to be backed up. Set up resource assignment by using the tags.
C. Create an AWS Lambda function. Program the Lambda function to use native tooling to take backups of file systems in Amazon EC2 and to make copies of databases in Amazon RDS. Create an Amazon EventBridge rule to invoke the Lambda function.
D. Use Amazon Data Lifecycle Manager to create an Amazon Elastic Block Store (Amazon EBS) snapshot policy. Create tags on each resource that needs to be backed up. Set up resource assignment by using the tags. Create multiple schedules according to the requirements within the policy. Set the appropriate frequency and retention period. In Amazon RDS, activate automated backups on the required DB instances.

**Correct Answer: B**
**Section:**
**Explanation:**
AWS Backup provides a centralized way to manage backups across AWS services. Here's how to implement the required backup strategy with minimal administrative effort:
Create Backup Plans: Set up different backup plans in AWS Backup, each configured for a specific backup frequency---daily, weekly, monthly, and yearly.
Set Retention Periods: For each backup plan, configure the retention settings to align with the required retention durations: 6 days, 4 weeks, 11 months, and 7 years respectively.

Tag Resources: Apply tags to each EC2 and RDS resource that needs to be backed up. This allows for the automated inclusion of these resources in the respective backup plans based on their tags.
Assign Resources to Backup Plans: Use the tags to define which resources are included in each backup plan, ensuring that all necessary resources are backed up according to the defined schedules and retention policies.
AWS Documentation
Reference: More details on setting up and managing AWS Backup can be found here: AWS Backup.

**QUESTION 116**
A company has multiple AWS accounts. The company uses AWS Organizations with an organizational unit (OU) for the production account and another OU for the development account. Corporate policies state that developers may use only approved AWS services in the production account.
What is the MOST operationally efficient solution to control the production account?

A. Create a customer managed policy in AWS Identity and Access Management (IAM). Apply the policy to all users within the production account.
B. Create a job function policy in AWS Identity and Access Management (IAM). Apply the policy to all users within the production OU.
C. Create a service control policy (SCP). Apply the SCP to the production OU.
D. Create an IAM policy. Apply the policy in Amazon API Gateway to restrict the production account.

**Correct Answer: C**
**Section:**

**QUESTION 117**
A company needs to upload gigabytes of files every day. The company need to achieve higher throughput and upload speeds to Amazon S3 Which action should a SysOps administrator take to meet this requirement?

A. Create an Amazon CloudFront distribution with the GET HTTP method allowed and the S3 bucketas an origin.
B. Create an Amazon ElastiCache duster and enable caching for the S3 bucket
C. Set up AWS Global Accelerator and configure it with the S3 bucket
D. Enable S3 Transfer Acceleration and use the acceleration endpoint when uploading files



**Correct Answer: D**
**Section:**
**Explanation:**
Enable Amazon S3 Transfer Acceleration Amazon S3 Transfer Acceleration can provide fast and secure transfers over long distances between your client and Amazon S3. Transfer Acceleration uses Amazon CloudFront's globally distributed edge locations.
https://aws.amazon.com/premiumsupport/knowledge-center/s3-upload-large-files/

**QUESTION 118**
A company requires that all IAM user accounts that have not been used for 90 days or more must have their access keys and passwords immediately disabled A SysOps administrator must automate the process of disabling unused keys using the MOST operationally efficient method.
How should the SysOps administrator implement this solution?

A. Create an AWS Step Functions workflow to identify IAM users that have not been active for 90 days Run an AWS Lambda function when a scheduled Amazon EventBridge (Amazon CloudWatch Events) rule is invoked to automatically remove the AWS access keys and passwords for these IAM users
B. Configure an AWS Config rule to identify IAM users that have not been active for 90 days Set up an automatic weekly batch process on an Amazon EC2 instance to disable the AWS access keys and passwords for these IAM users
C. Develop and run a Python script on an Amazon EC2 instance to programmatically identify IAM users that have not been active for 90 days Automatically delete these 1AM users
D. Set up an AWS Config managed rule to identify IAM users that have not been active for 90 days Set up an AWS Systems Manager automation runbook to disable the AWS access keys for these IAM users

**Correct Answer: D**
**Section:**

**QUESTION 119**
A company plans to run a public web application on Amazon EC2 instances behind an Elastic Load Balancer (ELB). The company's security team wants to protect the website by using AWS Certificate Manager (ACM) certificates The ELB must automatically redirect any HTTP requests to HTTPS Which solution will meet these requirements?

A. Create an Application Load Balancer that has one HTTPS listener on port 80 Attach an SSLTLScertificate to listener port 80 Create a rule to redirect requests from HTTP to HTTPS
B. Create an Application Load Balancer that has one HTTP listener on port 80 and one HTTPS protocollistener on port 443 Attach an SSL TLS certificate to listener port 443 Create a rule to redirect requestsfrom port 80 to port 443
C. Create an Application Load Balancer that has two TCP listeners on port 80 and port 443 Attach an SSLTLS certificate to listener port 443 Create a rule to redirect requests from port 80 to port 443
D. Create a Network Load Balancer that has two TCP listeners on port 80 and port 443 Attach an SSLTLS certificate to listener port 443 Create a rule to redirect requests from port 80 to port 443

**Correct Answer: B**
**Section:**

**QUESTION 120**
A company is planning to host its stateful web-based applications on AWS A SysOps administrator is using an Auto Scaling group of Amazon EC2 instances The web applications will run 24 hours a day 7 days a week throughout the year The company must be able to change the instance type within the same instance family later in the year based on the traffic and usage patterns Which EC2 instance purchasing option will meet these requirements MOST cost-effectively?

A. Convertible Reserved Instances
B. On-Demand instances
C. Spot instances
D. Standard Reserved instances

**Correct Answer: A**
**Section:**
**Explanation:**
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ri-convertible-exchange.html

**QUESTION 121**
A SysOps administrator is setting up a fleet of Amazon EC2 instances in an Auto Scaling group for an application. The fleet should have 50% CPU available at that times to accommodate bursts of traffic. The load will increase significantly between the hours of 09:00 and 17:00,7 days a week How should the SysOps administrator configure the scaling of the EC2 instances to meet these requirements?

A. Create a target tracking scaling policy that runs when the CPU utilization is higher than 90%
B. Create a target tracking scaling policy that runs when the CPU utilization is higher than 50%.
Create a scheduled scaling policy that ensures that the fleet is available at 09:00 Create a second scheduled scaling policy that scales in the fleet at 17:00
C. Set the Auto Scaling group to start with 2 instances by setting the desired instances maximum instances, and minimum instances to 2 Create a scheduled scaling policy that ensures that the fleet is available at 09:00
D. Create a scheduled scaling policy that ensures that the fleet is available at 09.00. Create a second scheduled scaling policy that scales in the fleet at 17:00

**Correct Answer: B**
**Section:**

**QUESTION 122**
A SysOps administrator receives an alert from Amazon GuardDuty about suspicious network activity on an Amazon FC2 instance. The GuardDuty finding lists a new external IP address as a traffic destination. The SysOps administrator does not recognize the external IP address. The SysOps administrator must block traffic to the external IP address that GuardDuty identified Which solution will meet this requirement?

A. Create a new security group to block traffic to the external IP address. Assign the new security group to the EC2 instance
B. Use VPC flow logs with Amazon Athena to block traffic to the external IP address
C. Create a network ACL Add an outbound deny rule tor traffic to the external IP address

D. Create a new security group to block traffic to the external IP address Assign the new security group to the entire VPC

**Correct Answer: A**
**Section:**

**QUESTION 123**
A SysOps administrator is designing a solution for an Amazon RDS for PostgreSQL DB instance.
Database credentials must be stored and rotated monthly. The applications that connect to the DB instance send write-intensive traffic with variable client connections that sometimes increase significantly in a short period of time. Which solution should a SysOps administrator choose to meet these requirements?

A. Configure AWS Key Management Service (AWS KMS) to automatically rotate the keys for the DB instance. Use RDS Proxy to handle the increases in database connections.
B. Configure AWS Key Management Service (AWS KMS) to automatically rotate the keys for the DB instance. Use RDS read replicas to handle the increases in database connections.
C. Configure AWS Secrets Manager to automatically rotate the credentials for the DB instance. Use RDS Proxy to handle the increases in database connections.
D. Configure AWS Secrets Manager to automatically rotate the credentials for the DB instance. Use RDS read replicas to handle the increases in database connections.

**Correct Answer: A**
**Section:**

**QUESTION 124**
An ecommerce company uses an Amazon ElastiCache for Memcached cluster for in-memory caching of popular product queries on the shopping site. When viewing recent Amazon CloudWatch metrics data for the ElastiCache cluster, the SysOps administrator notices a large number of evictions.
Which of the following actions will reduce these evictions? (Choose two.)

A. Add an additional node to the ElastiCache cluster.
B. Increase the ElastiCache time to live (TTL).
C. Increase the individual node size inside the ElastiCache cluster.
D. Put an Elastic Load Balancer in front of the ElastiCache cluster.
E. Use Amazon Simple Queue Service (Amazon SQS) to decouple the ElastiCache cluster.

**Correct Answer: A, C**
**Section:**
**Explanation:**
https://d1.awsstatic.com/training-and-certification/docs-sysops-associate/AWS-Certified-SysOps- Administrator-Associate_Sample-Questions_C02.pdf

**QUESTION 125**
A company is deploying a third-party unit testing solution that is delivered as an Amazon EC2 Amazon Machine Image (AMI). All system configuration data is stored in Amazon DynamoDB. The testing results are stored in Amazon S3. A minimum of three EC2 instances are required to operate the product. The company's testing team wants to use an additional three EC2 Instances when the Spot Instance prices are at a certain threshold. A SysOps administrator must Implement a highly available solution that provides this functionality.
Which solution will meet these requirements with the LEAST operational overhead?

A. Define an Amazon EC2 Auto Scaling group by using a launch configuration. Use the provided AMI In the launch configuration. Configure three On-Demand Instances and three Spot Instances. Configure a maximum Spot Instance price In the launch configuration.
B. Define an Amazon EC2 Auto Scaling group by using a launch template. Use the provided AMI in the launch template. Configure three On-Demand Instances and three Spot Instances. Configure a maximum Spot Instance price In the launch template.
C. Define two Amazon EC2 Auto Scaling groups by using launch configurations. Use the provided AMI in the launch configurations. Configure three On-Demand Instances for one Auto Scaling group. Configure three Spot Instances for the other Auto Scaling group. Configure a maximum Spot Instance price in the launch configuration for the Auto Scaling group that has Spot Instances.
D. Define two Amazon EC2 Auto Scaling groups by using launch templates. Use the provided AMI in the launch templates. Configure three On-Demand Instances for one Auto Scaling group. Configure three Spot Instances for the other Auto Scaling group. Configure a maximum Spot Instance price in the launch template for the Auto Scaling group that has Spot Instances.

**Correct Answer: A**
**Section:**
**Explanation:**
https://docs.aws.amazon.com/autoscaling/ec2/userguide/LaunchTemplates.html https://docs.aws.amazon.com/autoscaling/ec2/userguide/LaunchConfiguration.html

**QUESTION 126**
A company stores sensitive data in an Amazon S3 bucket. The company must log all access attempts to the S3 bucket. The company's risk team must receive immediate notification about any delete events. Which solution will meet these requirements?

A. Enable S3 server access logging for audit logs. Set up an Amazon Simple Notification Service (Amazon SNSJ notification for the S3 bucket. Select DeleteObject tor the event type for the alert system.
B. Enable S3 server access logging for audit logs. Launch an Amazon EC2 instance for the alert system. Run a cron job on the EC2 instance to download the access logs each day and to scan for a DeleteObject event.
C. Use Amazon CloudWatch Logs for audit logs. Use Amazon CloudWatch alarms with an Amazon Simple Notification Service (Amazon SNS) notification for the alert system.
D. Use Amazon CloudWatch Logs for audit logs. Launch an Amazon EC2 instance for The alert system. Run a cron job on the EC2 Instance each day to compare the list of the items with the list from the previous day. Configure the cron job to send a notification if an item is missing.

**Correct Answer: A**
**Section:**
**Explanation:**
To meet the requirements of logging all access attempts to the S3 bucket and receiving immediate notification about any delete events, the company can enable S3 server access logging and set up an Amazon Simple Notification Service (Amazon SNS) notification for the S3 bucket. The S3 server access logs will record all access attempts to the bucket, including delete events, and the SNS notification can be configured to send an alert when a DeleteObject event occurs.

**QUESTION 127**
A compliance learn requites all administrator passwords for Amazon RDS DB instances to be changed at least annually. Which solution meets this requirement in the MOST operationally efficient manner?

A. Store the database credentials in AWS Secrets Manager. Configure automatic rotation for the secret every 365 days.
B. Store the database credentials as a parameter In the RDS parameter group. Create a database trigger to rotate the password every 365 days.
C. Store the database credentials in a private Amazon S3 bucket. Schedule an AWS Lambda function to generate a new set of credentials every 365 days.
D. Store the database credentials in AWS Systems Manager Parameter Store as a secure string parameter. Configure automatic rotation for the parameter every 365 days.

**Correct Answer: A**
**Section:**

**QUESTION 128**
A company runs workloads on 90 Amazon EC2 instances in the eu-west-1 Region in an AWS account.
In 2 months, the company will migrate the workloads from eu-west-1 to the eu-west-3 Region.
The company needs to reduce the cost of the EC2 instances. The company is willing to make a 1-year commitment that will begin next week. The company must choose an EC2 Instance purchasing option that will provide discounts for the 90 EC2 Instances regardless of Region during the 1-year period.
Which solution will meet these requirements?

A. Purchase EC2 Standard Reserved Instances.
B. Purchase an EC2 Instance Savings Plan.
C. Purchase EC2 Convertible Reserved Instances.
D. Purchase a Compute Savings Plan.

**Correct Answer: B**
**Section:**

**QUESTION 129**
A company wants to archive sensitive data on Amazon S3 Glacier. The company's regulatory and compliance requirements do not allow any modifications to the data by any account. Which solution meets these requirements?

A. Attach a vault lock policy to an S3 Glacier vault that contains the archived data. Use the lock ID to validate the vault lock policy after 24 hours.
B. Attach a vault lock policy to an S3 Glacier vault that contains the archived data. Use the lock ID to validate the vault lock policy within 24 hours.
C. Configure S3 Object Lock in governance mode. Upload all files after 24 hours.
D. Configure S3 Object Lock in governance mode. Upload all files within 24 hours.

**Correct Answer: B**
**Section:**

**QUESTION 130**
A global company handles a large amount of personally identifiable information (Pll) through an internal web portal. The company's application runs in a corporate data center that is connected to AWS through an AWS Direct Connect connection. The application stores the Pll in Amazon S3.
According to a compliance requirement, traffic from the web portal to Amazon S3 must not travel across the internet. What should a SysOps administrator do to meet the compliance requirement?

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

**Correct Answer: A**
**Section:**
**Explanation:**
Using the interface endpoint, applications in your on-premises data center can easily query S3buckets over AWS Direct Connect or Site-to-Site VPN.https://aws.amazon.com/blogs/architecture/choosing-your-vpc-endpoint-strategy-for-amazon-s3/

**QUESTION 131**
A SysOps administrator recently configured Amazon S3 Cross-Region Replication on an S3 bucket Which of the following does this feature replicate to the destination S3 bucket by default?

A. Objects in the source S3 bucket for which the bucket owner does not have permissions
B. Objects that are stored in S3 Glacier
C. Objects that existed before replication was configured
D. Object metadata

**Correct Answer: B**
**Section:**

**QUESTION 132**
A company must migrate its applications to AWS The company is using Chef recipes for configuration management The company wants to continue to use the existing Chef recipes after the applications are migrated to AWS. What is the MOST operationally efficient solution that meets these requirements?

A. Use AWS Cloud Format ion to create an Amazon EC2 instance, install a Chef server, and add Chef recipes.
B. Use AWS CloudFormation to create a stack and add layers for Chef recipes.
C. Use AWS Elastic Beanstalk with the Docker platform to upload Chef recipes.
D. Use AWS OpsWorks to create a stack and add layers with Chef recipes.

**Correct Answer: D**
**Section:**

**QUESTION 133**
A company uses an Amazon CloudFront distribution to deliver its website. Traffic logs for the website must be centrally stored, and all data must be encrypted at rest. Which solution will meet these requirements?

A. Create an Amazon OpenSearch Service (Amazon Elasticsearch Service) domain with internet access and server-side encryption that uses the default AWS managed key. Configure CloudFront to use the Amazon OpenSearch Service (Amazon Elasticsearch Service) domain as a log destination.
B. Create an Amazon OpenSearch Service (Amazon Elasticsearch Service) domain with VPC access and server-side encryption that uses AES-256 Configure CloudFront to use the Amazon OpenSearch Service (Amazon Elasticsearch Service) domain as a log destination.
C. Create an Amazon S3 bucket that Is configured with default server-side encryption that uses AES- 256. Configure CloudFront to use the S3 bucket as a log destination.
D. Create an Amazon S3 bucket that is configured with no default encryption. Enable encryption in the CloudFront distribution, and use the S3 bucket as a log destination.

**Correct Answer: C**
**Section:**

**QUESTION 134**
A SysOps administrator is creating an Amazon EC2 Auto Scaling group in a new AWS account. After adding some instances, the SysOps administrator notices that the group has not reached the minimum number of instances. The SysOps administrator receives the following error message:

```
Launching a new EC2 instance. Status Reason: Your quota allows for 0 more running instance(s).
You requested at least 1. Launching EC2 instance failed.
```

Which action will resolve this issue?

A. Adjust the account spending limits for Amazon EC2 on the AWS Billing and Cost Management console
B. Modify the EC2 quota for that AWS Region in the EC2 Settings section of the EC2 console.
C. Request a quota Increase for the Instance type family by using Service Quotas on the AWS Management Console.
D. Use the Rebalance action In the Auto Scaling group on the AWS Management Console.

**Correct Answer: C**
**Section:**

**QUESTION 135**
A company needs to view a list of security groups that are open to the internet on port 3389.
What should a SysOps administrator do to meet this requirement?

A. Configure Amazon GuardDuty to scan security groups and report unrestricted access on port 3389.
B. Configure a service control policy (SCP) to identify security groups that allow unrestricted access on port 3389.
C. Use AWS Identity and Access Management Access Analyzer to find any instances that have unrestricted access on port 3389.
D. Use AWS Trusted Advisor to find security groups that allow unrestricted access on port 3389

**Correct Answer: D**
**Section:**

**QUESTION 136**
A company uses AWS Organizations to manage its AWS accounts. A SysOps administrator 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.

**Correct Answer: B**
**Section:**

**QUESTION 137**
A company uploaded its website files to an Amazon S3 bucket that has S3 Versioning enabled. The company uses an Amazon CloudFront distribution with the S3 bucket as the origin. The company recently modified the tiles, but the object names remained the same. Users report that old content is still appearing on the website.
How should a SysOps administrator remediate this issue?

A. Create a CloudFront invalidation, and add the path of the updated files.
B. Create a CloudFront signed URL to update each object immediately.
C. Configure an S3 origin access identity (OAI) to display only the updated files to users.
D. Disable S3 Versioning on the S3 bucket so that the updated files can replace the old files.

**Correct Answer: A**
**Section:**

**QUESTION 138**
A company uses AWS Organizations to manage multiple AWS accounts. The company's SysOps team has been using a manual process to create and manage 1AM roles. The team requires an automated solution to create and manage the necessary 1AM roles for multiple AWS accounts.
What is the MOST operationally efficient solution that meets these requirements?

A. Create AWS CloudFormation templates. Reuse the templates to create the necessary 1AM roles in each of the AWS accounts.
B. Use AWS Directory Service with AWS Organizations to automatically associate the necessary 1AM roles with Microsoft Active Directory users.
C. Use AWS Resource Access Manager with AWS Organizations to deploy and manage shared resources across the AWS accounts.
D. Use AWS CloudFormation StackSets with AWS Organizations to deploy and manage 1AM roles for the AWS accounts.

**Correct Answer: D**
**Section:**

**QUESTION 139**
A company's SysOps administrator attempts to restore an Amazon Elastic Block Store (Amazon EBS) snapshot. However, the snapshot is missing because another system administrator accidentally deleted the snapshot. The company needs the ability to recover snapshots for a specified period of time after snapshots are deleted.
Which solution will provide this functionality?

A. Turn on deletion protection on individual EBS snapshots that need to be kept.
B. Create an 1AM policy that denies the deletion of EBS snapshots by using a condition statement for the snapshot age Apply the policy to all users
C. Create a Recycle Bin retention rule for EBS snapshots for the desired retention period.
D. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function to copy EBS snapshots to Amazon S3 Glacier.

**Correct Answer: B**
**Section:**

**QUESTION 140**
A company is using Amazon Elastic Container Sen/ice (Amazon ECS) to run a containerized application on Amazon EC2 instances. A SysOps administrator needs to monitor only traffic flows between the ECS tasks. Which combination of steps should the SysOps administrator take to meet this requirement? (Select TWO.)

A. Configure Amazon CloudWatch Logs on the elastic network interface of each task.
B. Configure VPC Flow Logs on the elastic network interface of each task.

C. Specify the awsvpc network mode in the task definition.
D. Specify the bridge network mode in the task definition.
E. Specify the host network mode in the task definition.

**Correct Answer: B, C**
**Section:**
**Explanation:**
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-networking-awsvpc.html

**QUESTION 141**
A company runs a website from Sydney, Australi a. Users in the United States (US) and Europe are reporting that images and videos are taking a long time to load. However, local testing in Australia indicates no performance issues. The website has a large amount of static content in the form of images and videos that are stored m Amazon S3. Which solution will result In the MOST Improvement In the user experience for users In the US and Europe?

A. Configure AWS PrivateLink for Amazon S3.
B. Configure S3 Transfer Acceleration.
C. Create an Amazon CloudFront distribution. Distribute the static content to the CloudFront edge locations
D. Create an Amazon API Gateway API in each AWS Region. Cache the content locally.

**Correct Answer: D**
**Section:**

**QUESTION 142**
A SysOps administrator is using AWS Systems Manager Patch Manager to patch a fleet of Amazon EC2 instances. The SysOps administrator has configured a patch baseline and a maintenance window. The SysOps administrator also has used an instance tag to identify which instances to patch.
The SysOps administrator must give Systems Manager the ability to access the EC2 instances.
Which additional action must the SysOps administrator perform to meet this requirement?

A. Add an inbound rule to the instances' security group.
B. Attach an 1AM instance profile with access to Systems Manager to the instances.
C. Create a Systems Manager activation Then activate the fleet of instances.
D. Manually specify the instances to patch Instead of using tag-based selection.

**Correct Answer: A**
**Section:**

**QUESTION 143**
A company is expanding globally and needs to back up data on Amazon Elastic Block Store (Amazon EBS) volumes to a different AWS Region. Most of the EBS volumes that store the data are encrypted, but some of the EBS volumes are unencrypted. The company needs the backup data from all the EBS volumes to be encrypted.
Which solution will meet these requirements with the LEAST management overhead?

A. Configure a lifecycle policy in Amazon Data Lifecycle Manager (Amazon DLM) to create the EBS volume snapshots with cross-Region backups enabled. Encrypt the snapshot copies by using AWS Key Management Service (AWS KMS).
B. Create a point-in-time snapshot of the EBS volumes. When the snapshot status is COMPLETED, copy the snapshots to another Region and set the Encrypted parameter to False.
C. Create a point-in-time snapshot of the EBS volumes. Copy the snapshots to an Amazon S3 bucket that uses server-side encryption. Turn on S3 Cross-Region Replication on the S3 bucket.
D. Schedule an AWS Lambda function with the Python runtime. Configure the Lambda function to create the EBS volume snapshots, encrypt the unencrypted snapshots, and copy the snapshots to another Region.

**Correct Answer: A**

**Section:**
**Explanation:**
Encrypt the snapshot copies by using AWS Key Management Service (AWS KMS). This solution will allow the company to automatically create encrypted snapshots of the EBS volumes and copy them to different AWS Regions with minimal effort.

**QUESTION 144**
A company has an initiative to reduce costs associated with Amazon EC2 and AWS Lambd a. Which action should a SysOps administrator take to meet these requirements?

A. Analyze the AWS Cost and Usage Report by using Amazon Athena to identity cost savings.
B. Create an AWS Budgets alert to alarm when account spend reaches 80% of the budget.
C. Purchase Reserved Instances through the Amazon EC2 console.
D. Use AWS Compute Optimizer and take action on the provided recommendations.

**Correct Answer: D**
**Section:**

**QUESTION 145**
A company uses AWS Organizations. A SysOps administrator wants to use AWS Compute Optimizer and AWS tag policies in the management account to govern all member accounts in the billing family. The SysOps administrator navigates to the AWS Organizations console but cannot activate tag policies through the management account. What could be the reason for this issue?

A. All features have not been enabled in the organization.
B. Consolidated billing has not been enabled.
C. The member accounts do not have tags enabled for cost allocation.
D. The member accounts have not manually enabled trusted access for Compute Optimizer.

**Correct Answer: C**
**Section:**

**QUESTION 146**
A user working in the Amazon EC2 console increased the size of an Amazon Elastic Block Store
(Amazon EBS) volume attached to an Amazon EC2 Windows instance. The change is not reflected in the file system. What should a SysOps administrator do to resolve this issue?

A. Extend the file system with operating system-level tools to use the new storage capacity.
B. Reattach the EBS volume to the EC2 instance.
C. Reboot the EC2 instance that is attached to the EBS volume.
D. Take a snapshot of the EBS volume. Replace the original volume with a volume that is created from the snapshot.

**Correct Answer: B**
**Section:**

**QUESTION 147**
A SysOps administrator is reviewing AWS Trusted Advisor warnings and encounters a warning for an S3 bucket policy that has open access permissions. While discussing the issue with the bucket owner, the administrator realizes the S3 bucket is an origin for an Amazon CloudFront web distribution.
Which action should the administrator take to ensure that users access objects in Amazon S3 by using only CloudFront URLs?

A. Encrypt the S3 bucket content with Server-Side Encryption with Amazon S3-Managed Keys (SSES3).
B. Create an origin access identity and grant it permissions to read objects in the S3 bucket.
C. Assign an 1AM user to the CloudFront distribution and grant the user permissions in the S3 bucket policy.

D. Assign an 1AM role to the CloudFront distribution and grant the role permissions in the S3 bucket policy.

**Correct Answer: B**
**Section:**
**Explanation:**
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-contentrestricting-access-to-s3.html

**QUESTION 148**
A company is testing Amazon Elasticsearch Service (Amazon ES) as a solution for analyzing system logs from a fleet of Amazon EC2 instances. During the test phase, the domain operates on a singlenode cluster. A SysOps administrator needs to transition the test domain into a highly available production-grade deployment.
Which Amazon ES configuration should the SysOps administrator use to meet this requirement?

A. Use a cluster of four data nodes across two AWS Regions. Deploy four dedicated master nodes in each Region.
B. Use a cluster of six data nodes across three Availability Zones. Use three dedicated master nodes.
C. Use a cluster of six data nodes across three Availability Zones. Use six dedicated master nodes.
D. Use a cluster of eight data nodes across two Availability Zones. Deploy four master nodes in a failover AWS Region.

**Correct Answer: B**
**Section:**

**QUESTION 149**
A SysOps administrator is using AWS Compute Optimizer to get recommendations for a fleet of Amazon EC2 instances. After the analysis is complete, some of the EC2 instances are missing from the Compute Optimizer dashboard. What is the cause of this issue?

A. The missing instances do not have the Amazon CloudWatch agent installed.
B. Compute Optimizer does not support the instance types of the missing instances.
C. Compute Optimizer already considers the missing instances to be optimized.
D. The missing instances are running a Windows operating system.

**Correct Answer: A**
**Section:**

**QUESTION 150**
A company website contains a web tier and a database tier on AWS. The web tier consists of Amazon EC2 instances that run in an Auto Scaling group across two Availability Zones. The database tier runs on an Amazon RDS for MySQL Multi-AZ DB instance. The database subnet network ACLs are restricted to only the web subnets that need access to the database. The web subnets use the default network ACL with the default rules.
The company's operations team has added a third subnet to the Auto Scaling group configuration. After an Auto Scaling event occurs, some users report that they intermittently receive an error message. The error messages states that the server cannot connect to the database. The operations team has confirmed that the route tables are correct and that the required ports are open on all security groups. Which combination of actions should a SysOps administrator take so that the web servers can communicate with the DB instance? (Choose two.)

A. On the default ACL, create inbound Allow rules of type TCP with the ephemeral port range and the source as the database subnets.
B. On the default ACL. Create outbound Allow rules of type MySQL/Aurora (3306). Specify the destinations as the database subnets.
C. On the network ACLs for the database subnets, create an inbound Allow rule of type MySQL/Aurora (3306). Specify the source as the third web subnet.
D. On the network ACLs for the database subnets, create an outbound Allow rule of type TCP with the ephemeral port range and the destination as the third web subnet.
E. On the network ACLs for the database subnets, create an outbound Allow rule of type MySQL/Aurora (3306). Specify the destination as the third web subnet.

**Correct Answer: B, D**
**Section:**

**QUESTION 151**
An Amazon S3 Inventory report reveals that more than 1 million objects in an S3 bucket are not encrypted. These objects must be encrypted, and all future objects must be encrypted at the time they are written. Which combination of actions should a SysOps administrator take to meet these requirements? (Choose two.)

A. Create an AWS Config rule that runs evaluations against configuration changes to the S3 bucket. When an unencrypted object is found, run an AWS Systems Manager Automation document to encrypt the object in place.
B. Edit the properties of the S3 bucket to enable default server-side encryption.
C. Filter the S3 Inventory report by using S3 Select to find all objects that are not encrypted. Create an S3 Batch Operations job to copy each object in place with encryption enabled.
D. Filter the S3 Inventory report by using S3 Select to find all objects that are not encrypted. Send each object name as a message to an Amazon Simple Queue Service (Amazon SQS) queue. Use the SQS queue to invoke an AWS Lambda function to tag each object with a key of "Encryption" and a value of "SSE-KMS".
E. Use S3 Event Notifications to invoke an AWS Lambda function on all new object-created events for the S3 bucket. Configure the Lambda function to check whether the object is encrypted and to run an AWS Systems Manager Automation document to encrypt the object in place when an unencrypted object is found.

**Correct Answer: B, E**
**Section:**

**QUESTION 152**
A company hosts its website on Amazon EC2 instances behind an Application Load Balancer. The company manages its DNS with Amazon Route 53, and wants to point its domain's zone apex to the website. Which type of record should be used to meet these requirements?

A. An AAAA record for the domain's zone apex
B. An A record for the domain's zone apex
C. A CNAME record for the domain's zone apex
D. An alias record for the domain's zone apex



**Correct Answer: D**
**Section:**
**Explanation:**
Reference: https://aws.amazon.com/route53/faqs/

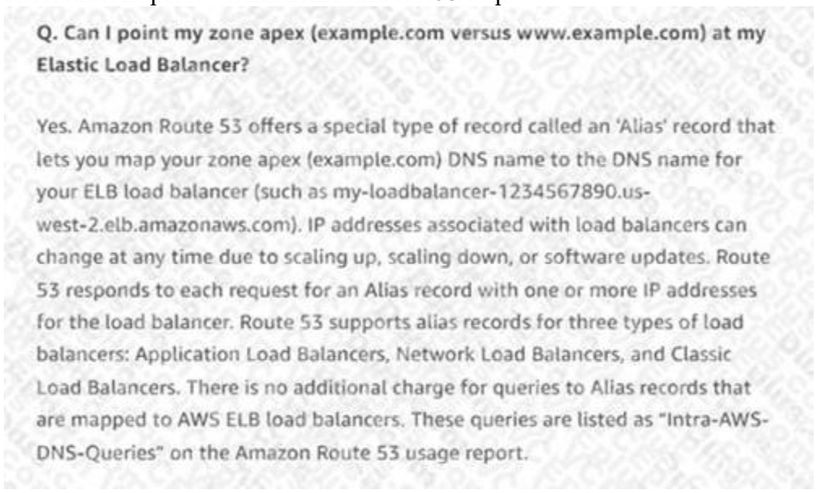


**QUESTION 153**
A company uses an Amazon Elastic File System (Amazon EFS) file system to share files across many Linux Amazon EC2 instances. A SysOps administrator notices that the file system's PercentIOLimit metric is consistently at 100% for 15 minutes or longer. The SysOps administrator also notices that the application that reads and writes to that file system is performing poorly. They application requires high throughput and IOPS while accessing the file system.
What should the SysOps administrator do to remediate the consistently high PercentIOLimit metric?

A. Create a new EFS file system that uses Max I/O performance mode. Use AWS DataSync to migrate data to the new EFS file system.
B. Create an EFS lifecycle policy to transition future files to the Infrequent Access (IA) storage class to improve performance. Use AWS DataSync to migrate existing data to IA storage.
C. Modify the existing EFS file system and activate Max I/O performance mode.
D. Modify the existing EFS file system and activate Provisioned Throughput mode.

**Correct Answer: A**
**Section:**

**QUESTION 154**
A company is migrating its production file server to AWS. All data that is stored on the file server must remain accessible if an Availability Zone becomes unavailable or when system maintenance is performed. Users must be able to interact with the file server through the SMB protocol. Users also must have the ability to manage file permissions by using Windows ACLs. Which solution will net these requirements?

A. Create a single AWS Storage Gateway file gateway.
B. Create an Amazon FSx for Windows File Server Multi-AZ file system.
C. Deploy two AWS Storage Gateway file gateways across two Availability Zones. Configure an Application Load Balancer in front of the file gateways.
D. Deploy two Amazon FSx for Windows File Server Single-AZ 2 file systems. Configure Microsoft Distributed File System Replication (DFSR).

**Correct Answer: B**
**Section:**
**Explanation:**
Reference: https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html

Amazon FSx for Windows File Server provides fully managed Microsoft Windows file servers, backed by a fully native Windows file system. FSx for Windows File Server has the features, performance, and compatibility to easily lift and shift enterprise applications to the AWS Cloud.

Amazon FSx supports a broad set of enterprise Windows workloads with fully managed file storage built on Microsoft Windows Server. Amazon FSx has native support for Windows file system features and for the industry-standard Server Message Block (SMB) protocol to access file storage over a network. Amazon FSx is optimized for enterprise applications in the AWS Cloud, with native Windows compatibility, enterprise performance and features, and consistent sub-millisecond latencies.

With file storage on Amazon FSx, the code, applications, and tools that Windows developers and administrators use today can continue to work unchanged. Windows applications and workloads ideal for Amazon FSx include business applications, home directories, web serving, content management, data analytics, software build setups, and media processing workloads.

**QUESTION 155**
A SysOps administrator is creating two AWS CloudFormation templates. The first template will create a VPC with associated resources, such as subnets, route tables, and an internet gateway. The second template will deploy application resources within the VPC that was created by the first template. The second template should refer to the resources created by the first template. How can this be accomplished with the LEAST amount of administrative effort?

A. Add an export field to the outputs of the first template and import the values in the second template.
B. Create a custom resource that queries the stack created by the first template and retrieves the required values.
C. Create a mapping in the first template that is referenced by the second template.
D. Input the names of resources in the first template and refer to those names in the second template as a parameter.

**Correct Answer: C**
**Section:**

**QUESTION 156**
A company is partnering with an external vendor to provide data processing services. For this integration, the vendor must host the company's data in an Amazon S3 bucket in the vendor's AWS account. The vendor is allowing the company to provide an AWS Key Management Service (AWS KMS) key to encrypt the company's data. The vendor has provided an IAM role Amazon Resources Name (ARN) to the company for this integration. What should a SysOps administrator do to configure this integration?

A. Create a new KMS key. Add the vendor's IAM role ARN to the KMS key policy. Provide the new KMS key ARN to the vendor.
B. Create a new KMS key. Create a new IAM key. Add the vendor's IAM role ARN to an inline policy that is attached to the IAM user. Provide the new IAM user ARN to the vendor.
C. Configure encryption using the KMS managed S3 key. Add the vendor's IAM role ARN to the KMS key policy. Provide the KMS managed S3 key ARN to the vendor.
D. Configure encryption using the KMS managed S3 key. Create an S3 bucket. Add the vendor's IAM role ARN to the S3 bucket policy. Provide the S3 bucket ARN to the vendor.

**Correct Answer: D**
**Section:**
**Explanation:**
Reference: https://bookdown.org/bingweiliu11/aws-tutorial-book/use-case.html

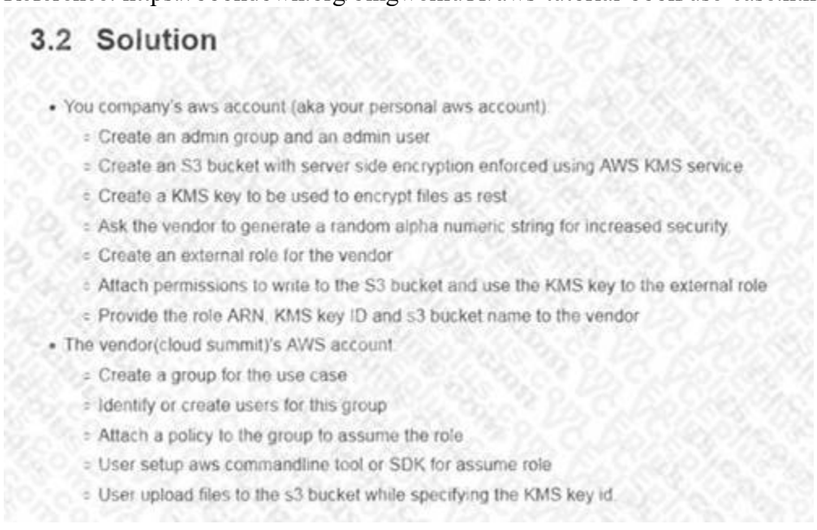


**QUESTION 157**
A company has an infernal web application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group in a single Availability Zone. A SysOps administrator must make the application highly available.
Which action should the SysOps administrator take to meet this requirement?

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

**Correct Answer: C**
**Section:**

**QUESTION 158**
A company uses AWS Organizations to host several applications across multiple AWS accounts. Several teams are responsible for building and maintaining the infrastructure of the application across the AWS accounts. A

SysOps administrator must implement a solution to ensure that user accounts and permissions are centrally managed. The solution must be integrated with the company's existing on-premises Active Directory environment. The SysOps administrator already has enabled AWS Single Sign-On (AWS SSO) and has set up an AWS Direct Connect connection. What is the MOST operationally efficient solution that meets these requirements?

A. Create a Simple AD domain, and establish a forest trust relationship with the on-premises Active Directory domain. Set the Simple AD domain as the identity source for AWS SSO. Create the required role-based permission sets. Assign each group of users to the AWS accounts that the group will manage.
B. Create an Active Directory domain controller on an Amazon EC2 instance that is joined to the on-premises Active Directory domain. Set the Active Directory domain controller as the identity source for AWS SSO. Create the required rolebased permission sets. Assign each group of users to the AWS accounts that the group will manage.
C. Create an AD Connector that is associated with the on-premises Active Directory domain. Set the AD Connector as the identity source for AWS SSO. Create the required role-based permission sets. Assign each group of users to the AWS accounts that the group will manage.
D. Use the built-in SSO directory as the identity source for AWS SSO. Copy the users and groups from the on-premises Active Directory domain. Create the required role-based permission sets. Assign each group of users to the AWS accounts that the group will manage.

**Correct Answer: C**
**Section:**
**Explanation:**
Reference: https://docs.aws.amazon.com/singlesignon/latest/userguide/connectonpremad.html


**QUESTION 159**
A company uses several large Chef recipes to automate the configuration of virtual machines (VMs) in its data center. A SysOps administrator is migrating this workload to Amazon EC2 Instances on AWS and must run the existing Chef recipes.
Which solution will meet these requirements MOST cost-effectively?

A. Create a Chef server that includes EC2 instances. Migrate the existing recipes. Modify the EC2 instance user data to connect to Chef.
B. Set up AWS OpsWorks for Chef Automate. Migrate the existing recipes. Modify the EC2 instance user data to connect to Chef.
C. Upload the existing recipes to Amazon S3. Run the recipes by using AWS Systems Manager State Manager.
D. Upload the existing recipes to the user data section during the creation of the EC2 instances.

**Correct Answer: B**
**Section:**

**QUESTION 160**
A company wants to be alerted through email when IAM CreateUser API calls are made within its AWS account. Which combination of actions should a SysOps administrator take to meet this requirement? (Choose two.)

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with AWS CloudTrail as the event source and IAM CreateUser as the specific API call for the event pattern.
B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with Amazon CloudSearch as the event source and IAM CreateUser as the specific API call for the event pattern.
C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with AWS IAM Access Analyzer as the event source and IAM CreateUser as the specific API call for the event pattern.
D. Use an Amazon Simple Notification Service (Amazon SNS) topic as an event target with an email subscription.
E. Use an Amazon Simple Email Service (Amazon SES) notification as an event target with an email subscription.

**Correct Answer: C, D**
**Section:**

**QUESTION 161**
A company needs to create a daily Amazon Machine Image (AMI) of an existing Amazon Linux EC2 instance that hosts the operating system, application, and database on multiple attached Amazon Elastic Block Store (Amazon EBS) volumes. File system integrity must be maintained.
Which solution will meet these requirements?

A. Create an AWS Lambda function to call the CreateImage API operation with the EC2 instance ID and the no-reboot parameter enabled. Create a daily scheduled Amazon EventBridge (Amazon CloudWatch Events) rule that invokes the function.
B. Create an AWS Lambda function to call the CreateImage API operation with the EC2 instance ID and the reboot parameter enabled. Create a daily scheduled Amazon EventBridge (Amazon CloudWatch Events) rule that invokes the function.
C. Use AWS Backup to create a backup plan with a backup rule that runs daily. Assign the resource ID of the EC2 instance with the no-reboot parameter enabled.
D. Use AWS Backup to create a backup plan with a backup rule that runs daily. Assign the resource ID of the EC2 instance with the reboot parameter enabled.

**Correct Answer: C**
**Section:**

**QUESTION 162**
A company is running a website on Amazon EC2 instances behind an Application Load Balancer (ALB). The company configured an Amazon CloudFront distribution and set the ALB as the origin. The company created an Amazon Route 53 CNAME record to send all traffic through the CloudFront distribution. As an unintended side effect, mobile users are now being served the desktop version of the website. Which action should a SysOps administrator take to resolve this issue?

A. Configure the CloudFront distribution behavior to forward the User-Agent header.
B. Configure the CloudFront distribution origin settings. Add a User-Agent header to the list of origin custom headers.
C. Enable IPv6 on the ALB. Update the CloudFront distribution origin settings to use the dualstack endpoint.
D. Enable IPv6 on the CloudFront distribution. Update the Route 53 record to use the dualstack endpoint.

**Correct Answer: C**
**Section:**
**Explanation:**
Reference: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-elb-load-balancer.html

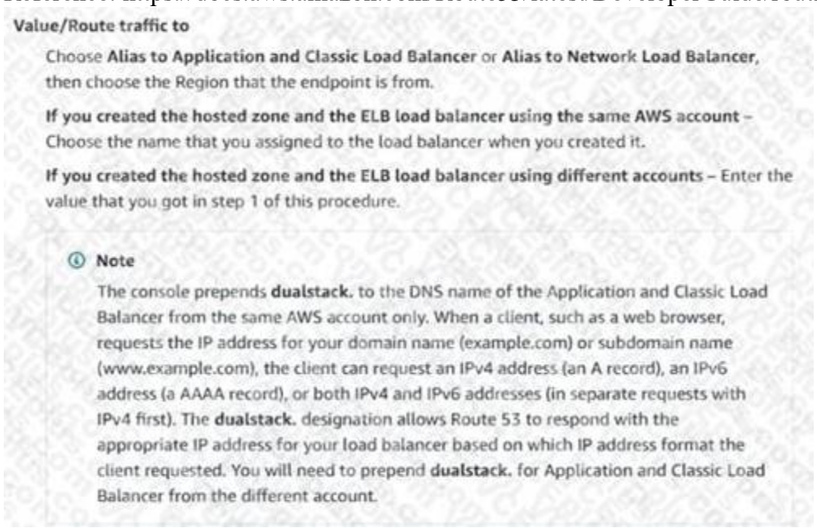

**QUESTION 163**
A company runs a web application on three Amazon EC2 instances behind an Application Load Balancer (ALB). The company notices that random periods of increased traffic cause a degradation in the application's performance. A SysOps administrator must scale the application to meet the increased traffic.
Which solution meets these requirements?

A. Create an Amazon CloudWatch alarm to monitor application latency and increase the size of each EC2 instance if the desired threshold is reached.
B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to monitor application latency and add an EC2 instance to the ALB if the desired threshold is reached.
C. Deploy the application to an Auto Scaling group of EC2 instances with a target tracking scaling policy. Attach the ALB to the Auto Scaling group.
D. Deploy the application to an Auto Scaling group of EC2 instances with a scheduled scaling policy. Attach the ALB to the Auto Scaling group.

**Correct Answer: C**
**Section:**
**Explanation:**
docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-target-tracking.html

**QUESTION 164**
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 SysOps administrator take to resolve this problem? (Choose 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.

**Correct Answer: C, E**
**Section:**

**QUESTION 165**
While setting up an AWS managed VPN connection, a SysOps administrator creates a customer gateway resource in AWS. The customer gateway device resides in a data center with a NAT gateway in front of it.
What address should be used to create the customer gateway resource?

A. The private IP address of the customer gateway device
B. The MAC address of the NAT device in front of the customer gateway device
C. The public IP address of the customer gateway device
D. The public IP address of the NAT device in front of the customer gateway device

**Correct Answer: D**
**Section:**
**Explanation:**
Reference: https://docs.aws.amazon.com/vpn/latest/s2svpn/cgw-options.html

**QUESTION 166**
A company is using an AWS KMS customer master key (CMK) with imported key material. The company references the CMK by its alias in the Java application to encrypt data. The CMK must be rotated every 6 months.
What is the process to rotate the key?

A. Enable automatic key rotation for the CMK, and specify a period of 6 months.

B. Create a new CMK with new imported material, and update the key alias to point to the new CMK.
C. Delete the current key material, and import new material into the existing CMK.
D. Import a copy of the existing key material into a new CMK as a backup, and set the rotation schedule for 6 months.

**Correct Answer: B**
**Section:**
**Explanation:**
Reference: https://aws.amazon.com/kms/faqs/

**QUESTION 167**
A company hosts an online shopping portal in the AWS Cloud. The portal provides HTTPS security by using a TLS certificateon an Elastic Load Balancer (ELB). Recently, the portal suffered an outage because the TLS certificate expired. A SysOpsadministrator must create a solution to automatically renew certificates to avoid this issue in the future. What is the MOST operationally efficient solution that meets these requirements?

A. Request a public certificate by using AWS Certificate Manager (ACM). Associate the certificate from ACM with the ELB. Write a scheduled AWS Lambda function to renew the certificate every 18 months.
B. Request a public certificate by using AWS Certificate Manager (ACM). Associate the certificate from ACM with the ELACM will automatically manage the renewal of the certificate.
C. Register a certificate with a third-party certificate authority (CA). Import this certificate into AWS Certificate Manager (ACM). Associate the certificate from ACM with the ELB. ACM will automatically manage the renewal of the certificate.
D. Register a certificate with a third-party certificate authority (CA). Configure the ELB to import the certificate directly from the CA. Set the certificate refresh cycle on the ELB to refresh when the certificate is within 3 months of the expiration date.

**Correct Answer: C**
**Section:**

**QUESTION 168**
A company is using an Amazon Aurora MySQL DB cluster that has point-in-time recovery, backtracking, and automatic backup enabled. A SysOps administrator needs to be able 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.

**Correct Answer: D**
**Section:**
**Explanation:**
Reference: https://aws.amazon.com/premiumsupport/knowledge-center/aurora-mysql-slow-snapshot-restore/

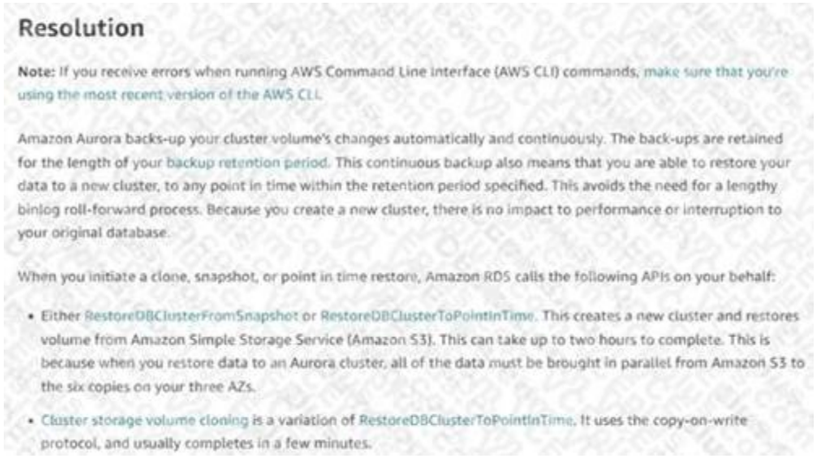

## Resolution

**Note:** If you receive errors when running AWS Command Line interface (AWS CLI) commands, make sure that you're using the most recent version of the AWS CLI.

Amazon Aurora backs-up your cluster volume's changes automatically and continuously. The back-ups are retained for the length of your backup retention period. This continuous backup also means that you are able to restore your data to a new cluster, to any point in time within the retention period specified. This avoids the need for a lengthy binlog roll-forward process. Because you create a new cluster, there is no impact to performance or interruption to your original database.

When you initiate a clone, snapshot, or point in time restore, Amazon RDS calls the following APIs on your behalf:

- Either RestoreDBClusterFromSnapshot or RestoreDBClusterToPointInTime. This creates a new cluster and restores volume from Amazon Simple Storage Service (Amazon S3). This can take up to two hours to complete. This is because when you restore data to an Aurora cluster, all of the data must be brought in parallel from Amazon S3 to the six copies on your three AZs.
- Cluster storage volume cloning is a variation of RestoreDBClusterToPointInTime. It uses the copy-on-write protocol, and usually completes in a few minutes.

**QUESTION 169**
A gaming application is deployed on four Amazon EC2 instances in a default VPC. The SysOps administrator has noticed consistently high latency in responses as data is transferred among the four instances. There is no way for the administrator to alter the application code.
The MOST effective way to reduce latency is to relaunch the EC2 instances in:

A. a dedicated VPC.
B. a single subnet inside the VPC.
C. a placement group.
D. a single Availability Zone.

**Correct Answer: C**
**Section:**

**QUESTION 170**
A new website will run on Amazon EC2 instances behind an Application Load Balancer. Amazon Route 53 will be used to manage DNS records. What type of record should be set in Route 53 to point the website's apex domain name (for example, "company.com") to the Application Load Balancer?

A. CNAME
B. SOA
C. TXT
D. ALIAS

**Correct Answer: D**
**Section:**
**Explanation:**
Reference: https://docs.aws.amazon.com/govcloud-us/latest/UserGuide/setting-up-route53-zoneapex-elb.html

Alias resource record sets are virtual records that work like CNAME records. But they differ from CNAME records in that they are not visible to resolvers. Resolvers only see the A record and the resulting IP address of the target record. As such, unlike CNAME records, alias resource record sets are available to configure a zone apex (also known as a root domain or naked domain) in a dynamic environment.

This section provides a solution for Route 53 zone apex alias support by setting up an Amazon CloudFront distribution between Route 53 and an AWS GovCloud (US) Elastic Load Balancing load balancer. The solution demonstrates how to configure Route 53 with a zone apex alias resource record set that maps to a CloudFront web distribution DNS name. The CloudFront distribution in turn points to the AWS GovCloud (US) load balancer DNS name as a custom origin.

An additional benefit of this approach is that CloudFront can help improve the performance of your website, including both static and dynamic content. For more information about CloudFront, see the CloudFront documentation.

**QUESTION 171**
A company has an existing web application that runs on two Amazon EC2 instances behind an Application Load Balancer (ALB) across two Availability Zones. The application uses an Amazon RDS MultiAZ DB Instance. Amazon Route 53 record sets route requests for dynamic content to the load balancer and requests for static content to an Amazon S3 bucket. Site visitors are reporting extremely long loading times. Which actions should be taken to improve the performance of the website? (Choose two.)

A. Add Amazon CloudFront caching for static content.
B. Change the load balancer listener from HTTPS to TCP.
C. Enable Amazon Route 53 latency-based routing.
D. Implement Amazon EC2 Auto Scaling for the web servers.
E. Move the static content from Amazon S3 to the web servers.

**Correct Answer: C, D**
**Section:**
**Explanation:**
Reference: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html
https://aws.amazon.com/ec2/autoscaling/

**QUESTION 172**
A SysOps administrator has launched a large general purpose Amazon EC2 instance to regularly process large data files. The instance has an attached 1 TB General Purpose SSD (gp2) Amazon Elastic Block Store (Amazon EBS) volume. The instance also is EBS-optimized. To save costs, the SysOps administrator stops the instance each evening and restarts the instance ach morning.
When data processing is active, Amazon CloudWatch metrics on the instance show a consistent 3,000 VolumeReadOps. The SysOps administrator must improve the I/O performance while ensuring data integrity.
Which action will meet these requirements?

A. Change the instance type to a large, burstable, general purpose instance.
B. Change the instance type to an extra large general purpose instance.
C. Increase the EBS volume to a 2 TB General Purpose SSD (gp2) volume.
D. Move the data that resides on the EBS volume to the instance store.

**Correct Answer: C**
**Section:**

**QUESTION 173**
A company uses Amazon Route 53 to manage the public DNS records for the domain example.com. The company deploys an Amazon CloudFront distribution to deliver static assets for a new corporate website. The company

wants to create a subdomain that is named "static" and must route traffic for the subdomain to the CloudFront distribution. How should a SysOps administrator create a new record for the subdomain in Route 53?

A. Create a CNAME record. Enter static.cloudfront.net as the record name. Enter the CloudFront distribution's public IP address as the value.
B. Create a CNAME record. Enter static.example.com as the record name. Enter the CloudFront distribution's private IP address as the value.
C. Create an A record. Enter static.cloudfront.net as the record name. Enter the CloudFront distribution's ID as an alias target.
D. Create an A record. Enter static.example.com as the record name. Enter the CloudFront distribution's domain name as an alias target.

**Correct Answer: D**
**Section:**

**QUESTION 174**
A manufacturing company uses an Amazon RDS DB instance to store inventory of all stock items. The company maintains several AWS Lambda functions that interact with the database to add, update, and delete items. The Lambda functions use hardcoded credentials to connect to the database.
A SysOps administrator must ensure that the database credentials are never stored in plaintext and that the password is rotated every 30 days. Which solution will meet these requirements in the MOST operationally efficient manner?

A. Store the database password as an environment variable for each Lambda function. Create a new Lambda function that is named PasswordRotate. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule the PasswordRotate function every 30 days to change the database password and update the environment variable for each Lambda function.
B. Use AWS Key Management Service (AWS KMS) to encrypt the database password and to store the encrypted password as an environment variable for each Lambda function. Grant each Lambda function access to the KMS key so that the database password can be decrypted when required. Create a new Lambda function that is named PasswordRotate to change the password every 30 days.
C. Use AWS Secrets Manager to store credentials for the database. Create a Secrets Manager secret and select the database so that Secrets Manager will use a Lambda function to update the database password automatically. Specify an automatic rotation schedule of 30 days. Update each Lambda function to access the database password from Secrets Manager.
D. Use AWS Systems Manager Parameter Store to create a secure string to store credentials for the database. Create a new Lambda function called PasswordRotate. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule the PasswordRotate function every 30 days to change the database password and to update the secret within Parameter Store. Update each Lambda function to access the database password from Parameter Store.

**Correct Answer: C**
**Section:**

**QUESTION 175**
The SysOps administrator needs to prevent any account within an AWS Organization from leaving the organization.

A. Create a service control policy (SCP) that denies the LeaveOrganization action. Apply the SCP to the root organizational unit (OU).
B. Create a service control policy (SCP) that denies the RemoveAccountFromOrganization action. Apply the SCP to the root organizational unit (OU).
C. Deploy an AWS Lambda function in each member account to remove any Organizations permissions when a user is created.
D. Turn on AWS Config. Set up the account-part-of-organizations managed rule. Configure the rule to run every hour.

**Correct Answer: A**
**Section:**
**Explanation:**
To prevent accounts from leaving an AWS Organization, an SCP that denies the LeaveOrganization action should be applied to the root organizational unit (OU).
Service Control Policy (SCP): By denying LeaveOrganization, member accounts are restricted from leaving the organization.
Root OU Application: Applying this policy at the root level ensures that no account in the organization can bypass it.
The RemoveAccountFromOrganization action pertains to removing an account by the organization's management account rather than preventing member accounts from leaving. AWS Config's account-part-of-organizations rule does not enforce this restriction but only monitors it.

**QUESTION 176**
The company's ecommerce website running on EC2 instances behind an ALB intermittently returns HTTP 500 errors. The Auto Scaling group is only using EC2 status checks.

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 the instances. Configure the agent to reboot the instances.

**Correct Answer: B**
**Section:**
**Explanation:**
Using ALB health checks in the Auto Scaling group will provide more accurate health monitoring and replace instances if they are unhealthy.
ALB Health Checks: Configure health checks based on HTTP response codes, which will detect application-level issues causing the HTTP 500 errors and replace instances as needed.
EC2 vs. ALB Health Checks: EC2 status checks only verify instance hardware and OS, not the application's responsiveness. By using ALB health checks, Auto Scaling can remove instances that are failing at the application level, thus preventing users from receiving 500 errors.
Replacing the ALB with a Network Load Balancer does not address HTTP 500 errors, and enabling session affinity does not resolve application health issues. The CloudWatch agent would not provide the required health check functionality for automated instance replacement.

**QUESTION 177**
The company needs a solution to provide failover for a Single-AZ RDS for MySQL DB instance to minimize application downtime.

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.

**Correct Answer: A**
**Section:**
**Explanation:**
Modifying the DB instance to a Multi-AZ deployment is the recommended solution for failover and high availability in RDS. Multi-AZ RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone, allowing automatic failover in case of an instance or Availability Zone failure.
Multi-AZ Deployment: Provides resilience with automatic failover and minimizes application downtime.
No Application Changes Needed: Multi-AZ is managed by AWS, so the failover is transparent to the application.
Read replicas and Auto Scaling groups do not provide automatic failover for write operations, and RDS Proxy only improves connection management rather than high availability.

**QUESTION 178**
The company's security team needs to consolidate Security Hub findings to reduce duplicate notifications for the same misconfigurations.

A. Turn on consolidated control findings in the Security Hub delegated administrator account.
B. Export the Security Hub findings. Consolidate the findings based on control ID. Visualize the findings in Amazon QuickSight.
C. Set up an AWS Config aggregator instead of Security Hub. Deploy a custom conformance pack by consolidating AWS Config rules.
D. Launch an Amazon EC2 instance in the organization's management account. Configure a custom script to assume a role in each linked account to extract and consolidate findings from the accounts.

**Correct Answer: A**
**Section:**
**Explanation:**
Enabling consolidated control findings in Security Hub reduces duplication by merging findings for similar controls across multiple standards. This reduces the operational burden of prioritizing remediation based on multiple copies of the same findings.
Consolidated Control Findings: Merges findings for controls across standards to avoid duplicates, providing a clearer view of misconfigurations without the need for additional infrastructure or manual processing.
Least Operational Overhead: This solution is managed within Security Hub without the need for external tools or manual exports.
Using AWS Config aggregators, QuickSight visualization, or custom EC2-based solutions would introduce additional complexity and overhead.

**QUESTION 179**
The SysOps administrator must dynamically reference the latest AMI ID from Systems Manager Parameter Store in CloudFormation templates for new AMI versions.

A. Create a new Systems Manager parameter to store the AMI value in the standard parameter tier.
B. Create a new Systems Manager parameter to store the AMI value in the advanced parameter tier.
C. Enable trusted access with Organizations.
D. Enable resource sharing with Organizations.
E. Create a resource share by using AWS Resource Access Manager (AWS RAM). Specify the new parameter as the resource. Specify the entire organization as the principal.
F. Create an Amazon EventBridge rule that invokes an AWS Lambda function when a new AMI is published. Program the Lambda function to assume an IAM role in all linked accounts and to update Parameter Store with the new AMI ID.

**Correct Answer: A, D, E**
**Section:**
**Explanation:**
To allow CloudFormation templates in all accounts within the organization to reference the latest AMI ID:
Parameter Store in Standard Tier: Storing the AMI ID in Systems Manager Parameter Store provides a central and easy-to-update source.
Enable Resource Sharing with Organizations: This allows the parameter to be shared across accounts in the organization.
Resource Share in AWS RAM: AWS Resource Access Manager (RAM) can be used to share the parameter with the entire organization, allowing other accounts to access the AMI ID.
Using the standard tier in Parameter Store is sufficient, and an EventBridge rule with Lambda for updating AMIs would add unnecessary complexity.

**QUESTION 180**
The company wants to improve the security and high availability of a two-tier web application that was rehosted to AWS, currently in a single Availability Zone.

A. Place the web-tier instances in an Auto Scaling group. Configure the Auto Scaling group to support a Multi-AZ deployment into private subnets that are behind an internet-facing Application Load Balancer.
B. Place the web-tier instances in an Auto Scaling group. Configure the Auto Scaling group in multiple AWS Regions. Deploy the EC2 instances into private subnets that are behind an internet-facing Application Load Balancer.
C. Launch an additional EC2 instance to host SQL Server. Place the new database EC2 instance in a second AWS Region. Enable replication between the two database EC2 instances.
D. Use AWS Database Migration Service (AWS DMS) to migrate the database EC2 instance to Amazon RDS for SQL Server with Multi-AZ Database Mirroring (DBM).
E. Use AWS Database Migration Service (AWS DMS) to migrate the database EC2 instance to Amazon DynamoDB.

**Correct Answer: A, D**
**Section:**
**Explanation:**
To improve security and availability, the best approach is to configure Multi-AZ for both the web and database tiers.
Multi-AZ Auto Scaling for Web Tier: Deploying the web-tier instances in an Auto Scaling group across multiple AZs with an internet-facing ALB provides high availability and fault tolerance.
RDS Multi-AZ for SQL Server: Migrating the SQL Server to RDS with Multi-AZ deployment ensures database redundancy and failover without additional management overhead.
Placing the web tier in multiple Regions would add unnecessary complexity, and migrating the database to DynamoDB is not suitable for applications requiring SQL Server's relational capabilities.

**QUESTION 181**
The company needs EC2 instances in the VPC to resolve DNS names for on-premises hosts using Direct Connect.

A. Create an Amazon Route 53 private hosted zone. Populate the zone with the hostnames and IP addresses of the hosts in the on-premises data center.
B. Create an Amazon Route 53 Resolver outbound endpoint. Add the IP addresses of an on-premises DNS server for the domain names that need to be forwarded.
C. Set up a forwarding rule for reverse DNS queries in Amazon Route 53 Resolver. Set the enableDnsHostnames attribute to true for the VPC.
D. Add the hostnames and IP addresses for the on-premises hosts to the /etc/hosts file of each EC2 instance.

**Correct Answer: B**
**Section:**

**Explanation:**
Using a Route 53 Resolver outbound endpoint allows DNS queries for on-premises hosts to be forwarded to the on-premises DNS server over the Direct Connect connection, minimizing maintenance and automating name resolution without the need for manual entry or file management.