

Exam Code: SOA-C02

Exam Name: AWS Certified SysOps Administrator - Associate



Exam A

QUESTION 1

A SysOps administrator needs to automate the invocation of an AWS Lambda function. The Lambda function must run at the end of each day to generate a report on data that is stored in an Amazon S3 bucket. What is the MOST operationally efficient solution that meets these requirements?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that has an event pattern for Amazon S3 and the Lambda function as a target.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that has a schedule and the Lambda function as a target.
- C. Create an S3 event notification to invoke the Lambda function whenever objects change in the S3 bucket.
- D. Deploy an Amazon EC2 instance with a cron job to invoke the Lambda function.

Correct Answer: C

Section:

QUESTION 2

A company deployed a new web application on multiple Amazon EC2 instances behind an Application Load Balancer (ALB). The EC2 instances run in an Auto Scaling group. Users report that they are frequently being prompted to log in.

What should a SysOps administrator do to resolve this issue?

- A. Configure an Amazon CloudFront distribution with the ALB as the origin.
- B. Enable sticky sessions (session affinity) for the target group of EC2 instances.
- C. Redeploy the EC2 instances in a spread placement group.
- D. Replace the ALB with a Network Load Balancer.



Correct Answer: C

Section:

Explanation:

QUESTION 3

A SysOps administrator manages the caching of an Amazon CloudFront distribution that serves pages of a website. The SysOps administrator needs to configure the distribution so that the TTL of individual pages can vary. The TTL of the individual pages must remain within the maximum TTL and the minimum TTL that are set for the distribution.

Which solution will meet these requirements?

- A. Create an AWS Lambda function that calls the CreateInvalidation API operation when a change in cache time is necessary.
- B. Add a Cache-Control: max-age directive to the object at the origin when content is being returned to CloudFront.
- C. Add a no-cache header through a Lambda@Edge function in response to the Viewer response.
- D. Add an Expires header through a CloudFront function in response to the Viewer response.

Correct Answer: B

Section:

Explanation:

To allow the TTL (Time to Live) of individual pages to vary while adhering to the maximum and minimum TTL settings configured for the Amazon CloudFront distribution, setting cache behaviors directly at the origin is most effective:

Use Cache-Control Headers: By configuring the Cache-Control: max-age directive in the HTTP headers of the objects served from the origin, you can specify how long an object should be cached by CloudFront before it is

considered stale.

Integration with CloudFront: When CloudFront receives a request for an object, it checks the cache-control header to determine the TTL for that specific object. This allows individual objects to have their own TTL settings, as long as they are within the globally set minimum and maximum TTL values for the distribution.

Operational Efficiency: This method does not require any additional AWS services or modifications to the distribution settings. It leverages HTTP standard practices, ensuring compatibility and ease of management.

Implementing the TTL management through cache-control headers at the origin provides precise control over caching behavior, aligning with varying content freshness requirements without complex configurations.

QUESTION 4

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 5

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 6

A company's SysOps administrator manages a fleet of hundreds of Amazon EC2 instances that run Windows-based workloads and Linux-based workloads. Each EC2 instance has a tag that identifies its operating system. All the EC2 instances run AWS Systems Manager Session Manager.

A zero-day vulnerability is reported, and no patches are available. The company's security team provides code for all the relevant operating systems to reduce the risk of the vulnerability. The SysOps administrator needs to implement the code on the EC2 instances and must provide a report that shows that the code has successfully run on all the instances.

What should the SysOps administrator do to meet these requirements as quickly as possible?

- A. Use Systems Manager Run Command. Choose either the AWS-RunShellScript document or the AWS-RunPowerShellScript document. Configure Run Command with the code from the security team. Specify the operating system tag in the Targets parameter. Run the command. Provide the command history's evidence to the security team.
- B. Create an AWS Lambda function that connects to the EC2 instances through Session Manager. Configure the Lambda function to identify the operating system, run the code from the security team, and return the results to an Amazon RDS DB instance. Query the DB instance for the results. Provide the results as evidence to the security team.
- C. Log on to each EC2 instance. Run the code from the security team on each EC2 instance. Copy and paste the results of each run into a single spreadsheet. Provide the spreadsheet as evidence to the security team.
- D. Update the launch templates of the EC2 instances to include the code from the security team in the user data. Relaunch the EC2 instances by using the updated launch templates. Retrieve the EC2 instance logs of each instance. Provide the EC2 instance logs as evidence to the security team.

Correct Answer: A

Section:

Explanation:

AWS Systems Manager Run Command provides an efficient method to execute administrative tasks on EC2 instances. This solution will minimize the time and complexity involved:

Select Document: Choose AWS-RunShellScript for Linux-based instances or AWS-RunPowerShellScript for Windows-based instances.

Configure Command: Enter the mitigation script provided by the security team into the command document.

Target Instances: Use the tagging system to target only the instances that match the specific OS as identified by their tags.

Execute Command: Run the command across the targeted instances.

Verification and Reporting: The command history in Systems Manager will serve as evidence of execution and success, which can be reported back to the security team.

AWS Documentation

Reference: More about Run Command can be found here: [AWS Systems Manager Run Command](#).

QUESTION 7

Accompany wants to monitor the number of Amazon EC2 instances that it is running. The company also wants to automate a service quota increase when the number of instances reaches a specific threshold.

Which solution meets these requirements?

- A. Create an Amazon CloudWatch alarm to monitor Service Quotas. Configure the alarm to invoke an AWS Lambda function to request a quota increase when the alarm reaches the threshold.
- B. Create an AWS Config rule to monitor Service Quotas. Call an AWS Lambda function to remediate the action and increase the quota.
- C. Create an Amazon CloudWatch alarm to monitor the AWS Health Dashboard. Configure the alarm to invoke an AWS Lambda function to request a quota increase when the alarm reaches the threshold.
- D. Create an Amazon CloudWatch alarm to monitor AWS Trusted Advisor service quotas. Configure the alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to increase the quota.

Correct Answer: A

Section:

Explanation:

This approach uses CloudWatch for monitoring and Lambda for automation, allowing for quick and efficient quota management:

Setup CloudWatch Alarm: Monitor the usage of EC2 instances against the service quota using CloudWatch.

Lambda Function: Write a Lambda function that triggers a quota increase request via the Service Quotas API when the threshold is met.

Integration: Configure the CloudWatch alarm to trigger this Lambda function when the instance count approaches the service quota.

AWS Documentation

Reference: Information on monitoring with CloudWatch and automating actions with Lambda can be found in these guides: Amazon CloudWatch Alarms, AWS Lambda.

QUESTION 8

A company uses AWS Organizations to manage its multi-account environment. The organization contains a dedicated account for security and a dedicated account for logging. A SysOps administrator needs to implement a centralized solution that provides alerts when a resource metric in any account crosses a standard defined threshold.

Which solution will meet these requirements?

- A. Deploy an AWS CloudFormation stack set to the accounts in the organization. Use a template that creates the required Amazon CloudWatch alarms and references an Amazon Simple Notification Service (Amazon SNS) topic in the logging account with publish permissions for all the accounts.
- B. Deploy an AWS CloudFormation stack in each account. Use the stack to deploy the required Amazon CloudWalch alarms and the required Amazon Simple Notification Service (Amazon SNS) topic.
- C. Deploy an AWS Lambda function on a cron job in each account. Configure the Lambda function to read resources that are in the account and to invoke an Amazon Simple Notification Service (Amazon SNS) topic if any metrics cross the defined threshold.
- D. Deploy an AWS CloudFormation change set to the organization. Use a template to create the required Amazon CloudWatch alarms and to send alerts to a verified Amazon Simple Email Service (Amazon SES) identity.

Correct Answer: A

Section:

Explanation:

Using AWS CloudFormation stack sets allows you to manage CloudWatch alarms across multiple accounts efficiently:

Create Stack Set: Use a CloudFormation template that defines the required CloudWatch alarms and configures them to publish alerts to an SNS topic.

Specify SNS Topic: Ensure the SNS topic is located in the logging account and has the necessary permissions set to receive publications from all accounts in the organization.

Deploy Across Organization: Implement the stack set across all accounts, ensuring centralized management and standardized deployment.

AWS Documentation

Reference: Learn more about deploying resources with CloudFormation StackSets: Working with AWS CloudFormation StackSets.

QUESTION 9

A company has developed a service that is deployed on a fleet of Linux-based Amazon EC2 instances that are in an Auto Scaling group. The service occasionally fails unexpectedly because of an error in the application code. The company's engineering team determines that resolving the underlying cause of the service failure could take several weeks.

A SysOps administrator needs to create a solution to automate recovery if the service crashes on any of the EC2 instances.

Which solutions will meet this requirement? (Select TWO.)

- A. Install the Amazon CloudWatch agent on the EC2 instances. Configure the CloudWatch agent to monitor the service. Set the CloudWatch action to restart if the service health check fails.
- B. Tag the EC2 instances. Create an AWS Lambda function that uses AWS Systems Manager Session Manager to log in to the tagged EC2 instances and restart the service. Schedule the Lambda function to run every 5 minutes.
- C. Tag the EC2 instances. Use AWS Systems Manager State Manager to create an association that uses the AWS-RunShellScript document. Configure the association command with a script that checks if the service is running and that starts the service if the service is not running. For targets, specify the EC2 instance tag. Schedule the association to run every 5 minutes.
- D. Update the EC2 user data that is specified in the Auto Scaling group's launch template to include a script that runs on a cron schedule every 5 minutes.
- E. Update the EC2 user data that is specified in the Auto Scaling group's launch template to ensure that the service runs during startup. Redeploy all the EC2 instances in the Auto Scaling group with the updated launch template.

Correct Answer: A, C

Section:**Explanation:**

The requirement is to automate recovery if the service crashes on any of the EC2 instances.

Option A: Install the Amazon CloudWatch agent on the EC2 instances. Configure the CloudWatch agent to monitor the service. Set the CloudWatch action to restart if the service health check fails. This is a valid solution because the CloudWatch agent can be configured to monitor the service and take action (restart the service) if the health check fails.

Option C: Tag the EC2 instances. Use AWS Systems Manager State Manager to create an association that uses the AWS-RunShellScript document. Configure the association command with a script that checks if the service is running and that starts the service if the service is not running. For targets, specify the EC2 instance tag. Schedule the association to run every 5 minutes. This is a valid solution because AWS Systems Manager State Manager can be used to maintain a consistent state of the EC2 instances. It can run a script to check if the service is running and start the service if it's not running.

Option B: Tag the EC2 instances. Create an AWS Lambda function that uses AWS Systems Manager Session Manager to log in to the tagged EC2 instances and restart the service. Schedule the Lambda function to run every 5 minutes. This is not a valid solution because AWS Lambda functions are not designed to log in to EC2 instances and restart services. They are used for running serverless applications.

Option D: Update the EC2 user data that is specified in the Auto Scaling group's launch template to include a script that runs on a cron schedule every 5 minutes. This is not a valid solution because user data scripts are run only during the launch of an EC2 instance. They are not designed to run on a schedule.

Option E: Update the EC2 user data that is specified in the Auto Scaling group's launch template to ensure that the service runs during startup. Redeploy all the EC2 instances in the Auto Scaling group with the updated launch template. This is not a valid solution because while user data can be used to ensure that the service runs during startup, it does not provide a solution for when the service crashes after the EC2 instance has started.

QUESTION 10

A SysOps administrator must analyze Amazon CloudWatch logs across 10 AWS Lambda functions for historical errors. The logs are in JSON format and are stored in Amazon S3. Errors sometimes do not appear in the same field, but all errors begin with the same string prefix.

What is the MOST operationally efficient way for the SysOps administrator to analyze the log files?

- A. Use S3 Select to write a query to search for errors. Run the query across all log groups of interest.
- B. Create an AWS Glue processing job to index the logs of interest. Run a query in Amazon Athena to search for errors.
- C. Use Amazon CloudWatch Logs Insights to write a query to search for errors. Run the query across all log groups of interest.
- D. Use Amazon CloudWatch Contributor Insights to create a rule. Apply the rule across all log groups of interest.

Correct Answer: C

Section:

QUESTION 11

A company hosts an internet web application on Amazon EC2 instances. The company is replacing the application with a new AWS Lambda function. During a transition period, the company must route some traffic to the legacy application and some traffic to the new Lambda function. The company needs to use the URL path of request to determine the routing.

Which solution will meet these requirements?

- A. Configure a Gateway Load Balancer to use the URL path to direct traffic to the legacy application and the new Lambda function.
- B. Configure a Network Load Balancer to use the URL path to direct traffic to the legacy application and the new Lambda function.
- C. Configure a Network Load Balancer to use a regular expression to match the URL path to direct traffic to the new Lambda function.
- D. Configure an Application Load Balancer to use the URL path to direct traffic to the legacy application and the new Lambda function.

Correct Answer: D

Section:

Explanation:

To route traffic based on the URL path during a transition period where both an EC2-based legacy application and a new AWS Lambda function are in use:

Use of Application Load Balancer (ALB): ALBs support advanced request routing based on the URL path, among other criteria. This capability allows the ALB to evaluate the URL path of incoming requests and route them appropriately to either the legacy EC2 instances or the Lambda function.

Path-Based Routing Rules: Configure the ALB with rules that specify which URL paths should be directed to the EC2 instances and which should be routed to the Lambda function. For example, requests to `/legacy/*` might go to the EC2 instances, while `/new/*` could be directed to the Lambda function.

Integration with Lambda: ALBs can directly invoke Lambda functions in response to HTTP requests, making them ideal for scenarios where both server-based and serverless components are used in tandem.

This setup not only facilitates a smooth transition by enabling simultaneous operation of both components but also leverages the native capabilities of ALBs to manage traffic based on application requirements effectively.

QUESTION 12

A SysOps administrator manages policies for many AWS member accounts in an AWS Organizations structure. Administrators on other teams have access to the account root user credentials of the member accounts. The SysOps administrator must prevent all teams, including their administrators, from using Amazon DynamoDB. The solution must not affect the ability of the teams to access other AWS services. Which solution will meet these requirements?

- A. In all member accounts, configure IAM policies that deny access to all DynamoDB resources for all users, including the root user.
- B. Create a service control policy (SCP) in the management account to deny all DynamoDB actions. Apply the SCP to the root of the organization
- C. In all member accounts, configure IAM policies that deny AmazonDynamoDBFullAccess to all users, including the root user.
- D. Remove the default service control policy (SCP) in the management account. Create a replacement SCP that includes a single statement that denies all DynamoDB actions.

Correct Answer: B

Section:

Explanation:

To prevent all teams within an AWS Organizations structure from using Amazon DynamoDB while allowing access to other AWS services, the most effective solution is to use a Service Control Policy (SCP). SCPs apply at the organization, organizational unit (OU), or account level and can override individual IAM policies, including the root user's permissions:

B: Create a service control policy (SCP) in the management account to deny all DynamoDB actions. Apply the SCP to the root of the organization. This policy will effectively block DynamoDB actions across all member accounts without affecting the ability to access other AWS services. SCPs are powerful tools for centrally managing permissions in AWS Organizations and can enforce policy compliance across all accounts. Further information on SCPs and their usage can be found in the AWS documentation on Service Control Policies [AWS Service Control Policies](#).

QUESTION 13

A company's SysOps administrator needs to change the AWS Support plan for one of the company's AWS accounts. The account has multi-factor authentication (MFA) activated, and the MFA device is lost. What should the SysOps administrator do to sign in?

- A. Sign in as a root user by using email and phone verification. Set up a new MFA device. Change the root user password.
- B. Sign in as an IAM user with administrator permissions. Resynchronize the MFA token by using the IAM console.
- C. Sign in as an IAM user with administrator permissions. Reset the MFA device for the root user by adding a new device.
- D. Use the forgot-password process to verify the email address. Set up a new password and MFA device.

Correct Answer: A

Section:

QUESTION 14

A SysOps administrator has created an AWS Service Catalog portfolio and has shared the portfolio with a second AWS account in the company. The second account is controlled by a different administrator. Which action will the administrator of the second account be able to perform?

- A. Add a product from the imported portfolio to a local portfolio.
- B. Add new products to the imported portfolio.
- C. Change the launch role for the products contained in the imported portfolio.
- D. Customize the products in the imported portfolio.

Correct Answer: A

Section:

QUESTION 15

A SysOps administrator wants to manage a web server application with AWS Elastic Beanstalk. The Elastic Beanstalk service must maintain full capacity for new deployments at all times. Which deployment policies satisfy this requirement? (Select TWO.)

- A. All at once

- B. Immutable
- C. Rebuild
- D. Rolling
- E. Rolling with additional batch

Correct Answer: B, E

Section:

Explanation:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html>

QUESTION 16

A company has a policy that requires all Amazon EC2 instances to have a specific set of tags. If an EC2 instance does not have the required tags, the noncompliant instance should be terminated. What is the MOST operationally efficient solution that meets these requirements?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to send all EC2 instance state changes to an AWS Lambda function to determine if each instance is compliant. Terminate any noncompliant instances.
- B. Create an IAM policy that enforces all EC2 instance tag requirements. If the required tags are not in place for an instance, the policy will terminate noncompliant instance.
- C. Create an AWS Lambda function to determine if each EC2 instance is compliant and terminate an instance if it is noncompliant. Schedule the Lambda function to invoke every 5 minutes.
- D. Create an AWS Config rule to check if the required tags are present. If an EC2 instance is noncompliant, invoke an AWS Systems Manager Automation document to terminate the instance.

Correct Answer: D

Section:

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-automation.html>

QUESTION 17

A SysOps administrator is evaluating Amazon Route 53 DNS options to address concerns about high availability for an on-premises website. The website consists of two servers: a primary active server and a secondary passive server. Route 53 should route traffic to the primary server if the associated health check returns 2xx or 3xx HTTP codes. All other traffic should be directed to the secondary passive server. The failover record type, set ID, and routing policy have been set appropriately for both primary and secondary servers. Which next step should be taken to configure Route 53?

- A. Create an A record for each server. Associate the records with the Route 53 HTTP health check.
- B. Create an A record for each server. Associate the records with the Route 53 TCP health check.
- C. Create an alias record for each server with evaluate target health set to yes. Associate the records with the Route 53 HTTP health check.
- D. Create an alias record for each server with evaluate target health set to yes. Associate the records with the Route 53 TCP health check.

Correct Answer: C

Section:

Explanation:

Reference: <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/health-checks-how-route-53-chooses-records.html>

QUESTION 18

A company is supporting a business-critical application that runs on Amazon EC2 instances. The application receives data from a service that runs in an on-premises data center. End users are reporting intermittent issues that are related to data refreshes. The issues are occurring because of fluctuations in available network bandwidth between AWS and the on-premises data center. A SysOps administrator must improve the user experience and the application's performance while minimizing changes to the application stack. Which solution will offer the MOST performance improvement while meeting these requirements?

- A. Migrate the service to AWS Implement auto scaling.
- B. Modify the service to use Amazon S3 Transfer Acceleration.
- C. Set up an AWS Direct Connect connection with the on-premises data center.

D. Use AWS Storage Gateway to move the data into AWS.

Correct Answer: B

Section:

Explanation:

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/transfer-acceleration.html>

QUESTION 19

A company uses AWS CloudFormation to deploy its application infrastructure. Recently, a user accidentally changed a property of a database in a CloudFormation template and performed a stack update that caused an interruption to the application. A SysOps administrator must determine how to modify the deployment process to allow the DevOps team to continue to deploy the infrastructure, but prevent against accidental modifications to specific resources. Which solution will meet these requirements?

- A. Set up an AWS Config rule to alert based on changes to any CloudFormation stack. An AWS Lambda function can then describe the stack to determine if any protected resources were modified and cancel the operation.
- B. Set up an Amazon CloudWatch Events event with a rule to trigger based on any CloudFormation API call. An AWS Lambda function can then describe the stack to determine if any protected resources were modified and cancel the operation.
- C. Launch the CloudFormation templates using a stack policy with an explicit allow for all resources and an explicit deny of the protected resources with an action of Update:*
- D. Attach an IAM policy to the DevOps team role that prevents a CloudFormation stack from updating, with a condition based on the specific Amazon Resource Names (ARNs) of the protected resources.

Correct Answer: C

Section:

Explanation:

Reference: <https://aws.amazon.com/blogs/devops/aws-cloudformation-security-best-practices/>

QUESTION 20

A software development company has multiple developers who work on the same product. Each developer must have their own development environments, and these development environments must be identical. Each development environment consists of Amazon EC2 instances and an Amazon RDS DB instance. The development environments should be created only when necessary, and they must be terminated each night to minimize costs. What is the MOST operationally efficient solution that meets these requirements?

- A. Provide developers with access to the same AWS CloudFormation template so that they can provision their development environment when necessary. Schedule a nightly cron job on each development instance to stop all running processes to reduce CPU utilization to nearly zero.
- B. Provide developers with access to the same AWS CloudFormation template so that they can provision their development environment when necessary. Schedule a nightly Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function to delete the AWS CloudFormation stacks.
- C. Provide developers with CLI commands so that they can provision their own development environment when necessary. Schedule a nightly Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function to terminate all EC2 instances and the DB instance.
- D. Provide developers with CLI commands so that they can provision their own development environment when necessary. Schedule a nightly Amazon EventBridge (Amazon CloudWatch Events) rule to cause AWS CloudFormation to delete all of the development environment resources.

Correct Answer: C

Section:

QUESTION 21

A SysOps administrator is optimizing the cost of a workload. The workload is running in multiple AWS Regions and is using AWS Lambda with Amazon EC2 On-Demand Instances for the compute. The overall usage is predictable. The amount of compute that is consumed in each Region varies, depending on the users' locations. Which approach should the SysOps administrator use to optimize this workload?

- A. Purchase Compute Savings Plans based on the usage during the past 30 days.
- B. Purchase Convertible Reserved Instances by calculating the usage baseline.
- C. Purchase EC2 Instance Savings Plans based on the usage during the past 30 days.
- D. Purchase Standard Reserved Instances by calculating the usage baseline.

Correct Answer: C

Section:

Explanation:

Reference: <https://pileuscloud.com/2019/11/14/aws-saving-plans-3-critical-things-to-know-before-buying-a-saving-plan/>

QUESTION 22

A SysOps administrator needs to give users the ability to upload objects to an Amazon S3 bucket. The SysOps administrator creates a presigned URL and provides the URL to a user, but the user cannot upload an object to the S3 bucket. The presigned URL has not expired, and no bucket policy is applied to the S3 bucket.

Which of the following could be the cause of this problem?

- A. The user has not properly configured the AWS CLI with their access key and secret access key.
- B. The SysOps administrator does not have the necessary permissions to upload the object to the S3 bucket.
- C. The SysOps administrator must apply a bucket policy to the S3 bucket to allow the user to upload the object.
- D. The object already has been uploaded through the use of the presigned URL, so the presigned URL is no longer valid.

Correct Answer: B

Section:

Explanation:

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies.html>

QUESTION 23

A company's IT department noticed an increase in the spend of their developer AWS account. There are over 50 developers using the account, and the finance team wants to determine the service costs incurred by each developer. What should a SysOps administrator do to collect this information? (Choose two.)

- A. Activate the createdBy tag in the account.
- B. Analyze the usage with Amazon CloudWatch dashboards.
- C. Analyze the usage with Cost Explorer.
- D. Configure AWS Trusted Advisor to track resource usage.
- E. Create a billing alarm in AWS Budgets.



Correct Answer: A, C

Section:

QUESTION 24

A SysOps administrator is maintaining a web application using an Amazon CloudFront web distribution, an Application Load Balancer (ALB), Amazon RDS, and Amazon EC2 in a VPC. All services have logging enabled. The administrator needs to investigate HTTP Layer 7 status codes from the web application. Which log sources contain the status codes? (Choose two.)

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

Correct Answer: C, D

Section:

QUESTION 25

A company uses an Amazon RDS DB instance for data storage for its web application. For disaster recovery purposes, a SysOps administrator has configured an AWS Lambda function that copies the daily DB snapshot to the

us-west-2- Region.

The SysOps administrator must provide a custom DNS name, myexampledb, for the DB instance so that the company's developers do not need to update the application code if the DB snapshot must be restored in another Region. The company hosts its corporate domain, example.com, on Amazon Route 53.

Which solution will meet these requirements?

- A. Create a Route 53 alias record that maps myexampledb.example.com to the DB instance domain name. Instruct the developers to refer to myexampledb.example.com in their application. After restoring the DB snapshot in us-west-2, update the alias record to point to the new DB instance domain name.
- B. Create a Route 53 CNAME record that maps myexampledb.example.com to the DB instance domain name. Instruct the developers to refer to myexampledb.example.com in their application. After restoring the DB snapshot in us-west-2, update the CNAME record to point to the new DB instance domain name.
- C. Locate the IP address of the DB instance. Create a Route 53 A record that maps myexampledb.example.com to the IP address. Instruct the developers to refer to myexampledb.example.com in their application. After restoring the DB snapshot in us-west-2, update the A record to point to the new DB instance IP address.
- D. Locate the IP address of the DB instance. Create a Route 53 alias record that maps myexampledb.example.com to the IP address. Instruct the developers to refer to myexampledb.example.com in their application. After restoring the DB snapshot in us-west-2, update the alias record to point to the new DB instance IP address.

Correct Answer: D

Section:

Explanation:

Reference: <https://www.amazonaws.cn/en/route53/faqs/>

QUESTION 26

A company has a new requirement stating that all resources in AWS must be tagged according to a set policy. Which AWS service should be used to enforce and continually identify all resources that are not in compliance with the policy?

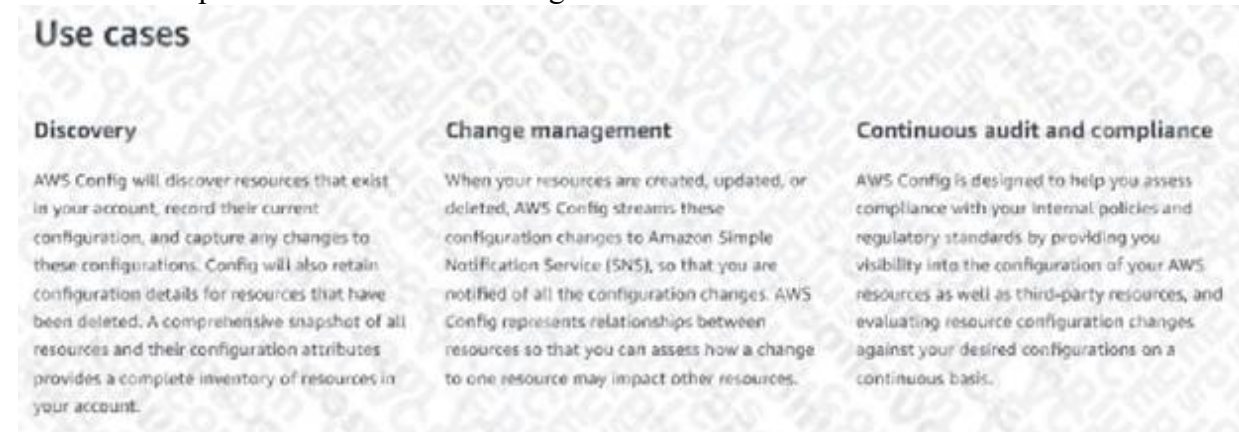
- A. AWS CloudTrail
- B. Amazon Inspector
- C. AWS Config
- D. AWS Systems Manager

Correct Answer: C

Section:

Explanation:

Reference: <https://aws.amazon.com/config/>



QUESTION 27

With the threat of ransomware viruses encrypting and holding company data hostage, which action should be taken to protect an Amazon S3 bucket?

- A. Deny Post, Put, and Delete on the bucket.
- B. Enable server-side encryption on the bucket.



- C. Enable Amazon S3 versioning on the bucket.
- D. Enable snapshots on the bucket.

Correct Answer: C

Section:

QUESTION 28

A company manages an application that uses Amazon ElastiCache for Redis with two extra-large nodes spread across two different Availability Zones. The company's IT team discovers that the ElastiCache for Redis cluster has 75% freeable memory. The application must maintain high availability. What is the MOST cost-effective way to resize the cluster?

- A. Decrease the number of nodes in the ElastiCache for Redis cluster from 2 to 1.
- B. Deploy a new ElastiCache for Redis cluster that uses large node types. Migrate the data from the original cluster to the new cluster. After the process is complete, shut down the original cluster.
- C. Deploy a new ElastiCache for Redis cluster that uses large node types. Take a backup from the original cluster, and restore the backup in the new cluster. After the process is complete, shut down the original cluster.
- D. Perform an online resizing for the ElastiCache for Redis cluster. Change the node types from extra-large nodes to large nodes.

Correct Answer: D

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/scaling-redis-cluster-modeenabled.html>

As demand on your clusters changes, you might decide to improve performance or reduce costs by changing the number of shards in your Redis (cluster mode enabled) cluster. We recommend using online horizontal scaling to do so, because it allows your cluster to continue serving requests during the scaling process. <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/redis-cluster-vertical-scalingscaling-down.html>

QUESTION 29

An existing, deployed solution uses Amazon EC2 instances with Amazon EBS General Purpose SSD volumes, an Amazon RDS PostgreSQL database, an Amazon EFS file system, and static objects stored in an Amazon S3 bucket. The Security team now mandates that at-rest encryption be turned on immediately for all aspects of the application, without creating new resources and without any downtime. To satisfy the requirements, which one of these services can the SysOps administrator enable at-rest encryption on?

- A. EBS General Purpose SSD volumes
- B. RDS PostgreSQL database
- C. Amazon EFS file systems
- D. S3 objects within a bucket

Correct Answer: B

Section:

QUESTION 30

A SysOps administrator notices a scale-up event for an Amazon EC2 Auto Scaling group. Amazon CloudWatch shows a spike in the RequestCount metric for the associated Application Load Balancer. The administrator would like to know the IP addresses for the source of the requests. Where can the administrator find this information?

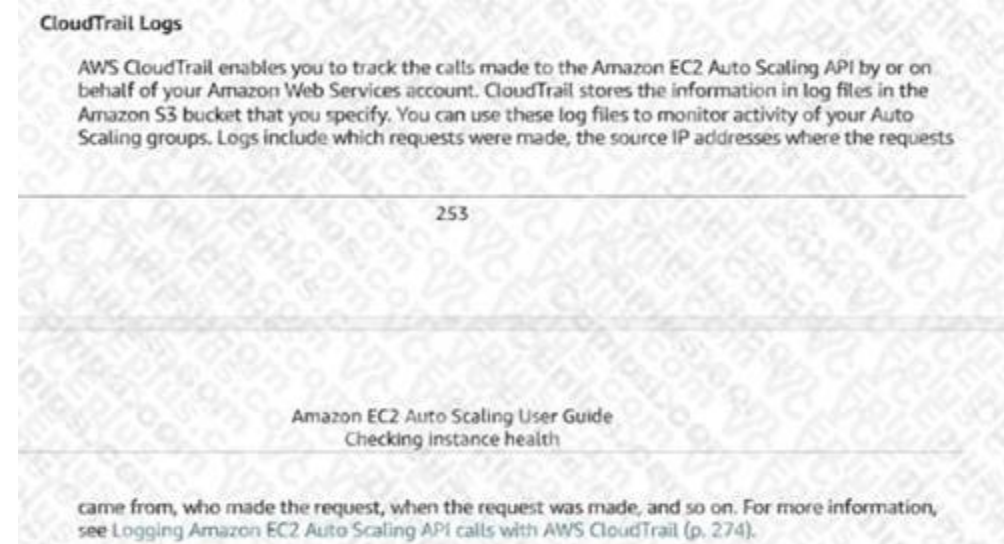
- A. Auto Scaling logs
- B. AWS CloudTrail logs
- C. EC2 instance logs
- D. Elastic Load Balancer access logs

Correct Answer: B

Section:

Explanation:

Reference: <https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-dg.pdf> page 253



QUESTION 31

A company is running an application on premises and wants to use AWS for data backup. All of the data must be available locally. The backup application can write only to block-based storage that is compatible with the Portable Operating System Interface (POSIX).

Which backup solution will meet these requirements?

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



Correct Answer: D

Section:

QUESTION 32

A company has an Amazon RDS DB instance. The company wants to implement a caching service while maintaining high availability. Which combination of actions will meet these requirements? (Choose two.)

- A. Add Auto Discovery to the data store.
- B. Create an Amazon ElastiCache for Memcached data store.
- C. Create an Amazon ElastiCache for Redis data store.
- D. Enable Multi-AZ for the data store.
- E. Enable Multi-threading for the data store.

Correct Answer: A, D

Section:

QUESTION 33

A SysOps administrator must create a solution that immediately notifies software developers if an AWS Lambda function experiences an error. Which solution will meet this requirement?

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic with an email subscription for each developer. Create an Amazon CloudWatch alarm by using the Errors metric and the Lambda function name as a dimension. Configure the alarm to send a notification to the SNS topic when the alarm state reaches ALARM.

- B. Create an Amazon Simple Notification Service (Amazon SNS) topic with a mobile subscription for each developer. Create an Amazon EventBridge (Amazon CloudWatch Events) alarm by using the LambdaError as the event pattern and the SNS topic name as a resource. Configure the alarm to send a notification to the SNS topic when the alarm state reaches ALARM.
- C. Verify each developer email address in Amazon Simple Email Service (Amazon SES). Create an Amazon CloudWatch rule by using the LambdaError metric and developer email addresses as dimensions. Configure the rule to send an email through Amazon SES when the rule state reaches ALARM.
- D. Verify each developer mobile phone in Amazon Simple Email Service (Amazon SES). Create an Amazon EventBridge (Amazon CloudWatch Events) rule by using Error as the event pattern and the Lambda function name as a resource. Configure the rule to send a push notification through Amazon SES when the rule state reaches ALARM.

Correct Answer: D

Section:

QUESTION 34

A large company is using AWS Organizations to manage its multi-account AWS environment. According to company policy, all users should have read-level access to a particular Amazon S3 bucket in a central account. The S3 bucket data should not be available outside the organization. A SysOps administrator must set up the permissions and add a bucket policy to the S3 bucket. Which parameters should be specified to accomplish this in the MOST efficient manner?

- A. Specify "*" as the principal and PrincipalOrgId as a condition.
- B. Specify all account numbers as the principal.
- C. Specify PrincipalOrgId as the principal.
- D. Specify the organization's master account as the principal.

Correct Answer: A

Section:

Explanation:

Reference: <https://aws.amazon.com/blogs/security/iam-share-aws-resources-groups-aws-accounts-aws-organizations/>



```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "TrainingDataS3ReadOnly",
      "Effect": "Allow",
      "Principal": "*",
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3:::training-data/*",
      "Condition": {
        "ForAnyValue:StringLike": {
          "aws:PrincipalOrgPaths": ["o-myorganization/*ou-machinelearn/*"]
        }
      }
    }
  ]
}

```

In the policy above, I assert that principals trying to read the contents of the training-data bucket must be either a member of the OU that corresponds to the ou-machinelearn ID I provided (my Machine Learning OU Identifier), or a member of any OUs that are children of it. For the aws:PrincipalOrgPaths value, I used two asterisk (*) wildcards. I used the first asterisk (*) between my organization ID and my OU ID because OU IDs are unique within my organization. This means specifying the full path is not necessary to select the OU I need. The second asterisk (*) at the end of the path, is used to specify that I want to allow all child OUs to be included in my string comparison. If I didn't want to include the child OUs, I could remove the wildcard character.

QUESTION 35

A company is planning to host an application on a set of Amazon EC2 instances that are distributed across multiple Availability Zones. The application must be able to scale to millions of requests each second. A SysOps administrator must design a solution to distribute the traffic to the EC2 instances. The solution must be optimized to handle sudden and volatile traffic patterns while using a single static IP address for each Availability Zone. Which solution will meet these requirements?

- A. Amazon Simple Queue Service (Amazon SQS) queue
- B. Application Load Balancer
- C. AWS Global Accelerator
- D. Network Load Balancer

Correct Answer: B

Section:

QUESTION 36

A company has a critical serverless application that uses multiple AWS Lambda functions. Each Lambda function generates 1 GB of log data daily in its own Amazon CloudWatch Logs log group. The company's security team asks for a count of application errors, grouped by type, across all of the log group.

What should a SysOps administrator do to meet this requirement?

- A. Perform a CloudWatch Logs Insights query that uses the stats command and count function.
- B. Perform a CloudWatch Logs search that uses the groupby keyword and count function.
- C. Perform an Amazon Athena query that uses the SELECT and GROUP BY keywords.
- D. Perform an Amazon RDS query that uses the SELECT and GROUP BY keywords.

Correct Answer: A

Section:

QUESTION 37

A company is running an application on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The EC2 instances are launched by an Auto Scaling group and are automatically registered in a target group. A SysOps administrator must set up a notification to alert application owners when targets fail health checks. What should the SysOps administrator do to meet these requirements?

- A. Create an Amazon CloudWatch alarm on the UnHealthyHostCount metric. Configure an action to send an Amazon Simple Notification Service (Amazon SNS) notification when the metric is greater than 0.
- B. Configure an Amazon EC2 Auto Scaling custom lifecycle action to send an Amazon Simple Notification Service (Amazon SNS) notification when an instance is in the Pending: Wait state.
- C. Update the Auto Scaling group. Configure an activity notification to send an Amazon Simple Notification Service (Amazon SNS) notification for the Unhealthy event type.
- D. Update the ALB health check to send an Amazon Simple Notification Service (Amazon SNS) notification when an instance is unhealthy.

Correct Answer: A

Section:

Explanation:

Reference: <https://aws.amazon.com/blogs/networking-and-content-delivery/identifying-unhealthy-targets-of-elastic-loadbalancer/>

QUESTION 38

A SysOps administrator is troubleshooting an AWS CloudFormation template whereby multiple Amazon EC2 instances are being created. The template is working in us-east-1, but it is failing in us-west-2 with the error code: AMI [ami-12345678] does not exist

How should the Administrator ensure that the AWS CloudFormation template is working in every region?

- A. Copy the source region's Amazon Machine Image (AMI) to the destination region and assign it the same ID.
- B. Edit the AWS CloudFormation template to specify the region code as part of the fully qualified AMI ID.
- C. Edit the AWS CloudFormation template to offer a drop-down list of all AMIs to the user by using the AWS::EC2::AMI::ImageID control.
- D. Modify the AWS CloudFormation template by including the AMI IDs in the "Mappings" section. Refer to the proper mapping within the template for the proper AMI ID.

Correct Answer: D

Section:

QUESTION 39

An organization created an Amazon Elastic File System (Amazon EFS) volume with a file system ID of fs-85ba41fc, and it is actively used by 10 Amazon EC2 hosts. The organization has become concerned that the file system is not encrypted. How can this be resolved?

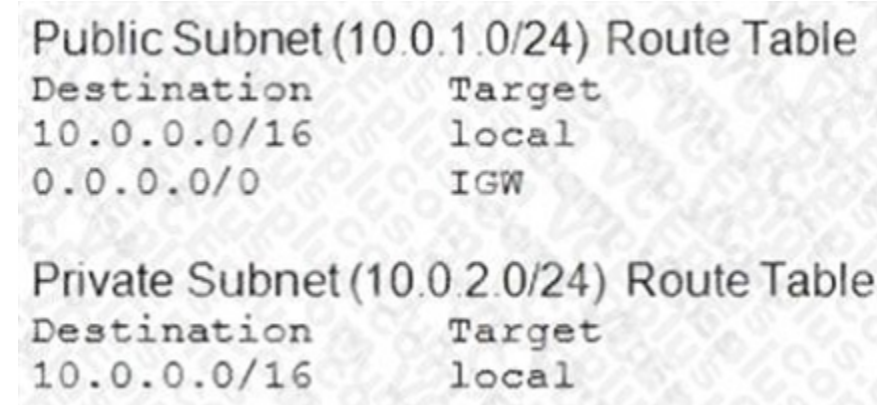
- A. Enable encryption on each host's connection to the Amazon EFS volume. Each connection must be recreated for encryption to take effect.
- B. Enable encryption on the existing EFS volume by using the AWS Command Line Interface.
- C. Enable encryption on each host's local drive. Restart each host to encrypt the drive.
- D. Enable encryption on a newly created volume and copy all data from the original volume. Reconnect each host to the new volume.

Correct Answer: D

Section:

QUESTION 40

A SysOps administrator is attempting to download patches from the internet into an instance in a private subnet. An internet gateway exists for the VPC, and a NAT gateway has been deployed on the public subnet; however, the instance has no internet connectivity. The resources deployed into the private subnet must be inaccessible directly from the public internet.



The image shows two screenshots of AWS Route Tables. The first is for the Public Subnet (10.0.1.0/24) and the second is for the Private Subnet (10.0.2.0/24). Both show a route for 10.0.0.0/16 to local and a route for 0.0.0.0/0 to IGW.

Public Subnet (10.0.1.0/24) Route Table	
Destination	Target
10.0.0.0/16	local
0.0.0.0/0	IGW

Private Subnet (10.0.2.0/24) Route Table	
Destination	Target
10.0.0.0/16	local



What should be added to the private subnet's route table in order to address this issue, given the information provided?

- A. 0.0.0.0/0 IGW
- B. 0.0.0.0/0 NAT
- C. 10.0.1.0/24 IGW
- D. 10.0.1.0/24 NAT

Correct Answer: B

Section:

Explanation:

Reference: https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Scenario2.html

QUESTION 41

An organization finds that a high number of gp2 Amazon EBS volumes are running out of space. Which solution will provide the LEAST disruption with MINIMAL effort?

- A. Create a snapshot and restore it to a larger gp2 volume.
- B. Create a RAID 0 with another new gp2 volume to increase capacity.
- C. Leverage the Elastic Volumes feature of EBS to increase gp2 volume size.
- D. Write a script to migrate data to a larger gp2 volume.

Correct Answer: C

Section:

Explanation:

Reference: <https://aws.amazon.com/ebs/features/>

Amazon EBS Elastic Volumes

Elastic Volumes is a feature that allows you to easily adapt your volumes as the needs of your applications change. Elastic Volumes allows you to dynamically increase capacity, tune performance, and change the type of any new or existing current generation volume with no downtime or performance impact. Easily right-size your deployment and adapt to performance changes.

Simply create a volume with the capacity and performance needed today knowing you have the ability to modify your volume configuration in the future, saving hours of planning cycles.

By using Amazon CloudWatch with AWS Lambda, you can automate volume changes to meet the changing needs of your applications.

The Elastic Volumes feature makes it easier to adapt your resources to changing application demands, giving you confidence that you can make modifications in the future as your business needs change.

QUESTION 42



A company hosts a website on multiple Amazon EC2 instances that run in an Auto Scaling group. Users are reporting slow responses during peak times between 6 PM and 11 PM every weekend. A SysOps administrator must implement a solution to improve performance during these peak times.



What is the MOST operationally efficient solution that meets these requirements?

- A. Create a scheduled Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function to increase the desired capacity before peak times.
- B. Configure a scheduled scaling action with a recurrence option to change the desired capacity before and after peak times.
- C. Create a target tracking scaling policy to add more instances when memory utilization is above 70%.
- D. Configure the cooldown period for the Auto Scaling group to modify desired capacity before and after peak times.

Correct Answer: B

Section:

QUESTION 43

A company's SysOps administrator has created an Amazon EC2 instance with custom software that will be used as a template for all new EC2 instances across multiple AWS accounts. The Amazon Elastic Block Store (Amazon EBS) volumes that are attached to the EC2 instance are encrypted with AWS managed keys. The SysOps administrator creates an Amazon Machine Image (AMI) of the custom EC2 instance and plans to share the AMI with the company's other AWS accounts. The company requires that all AMIs are encrypted with AWS Key Management Service (AWS KMS) keys and that only authorized AWS accounts can access the shared AMIs. Which solution will securely share the AMI with the other AWS accounts?

- A. In the account where the AMI was created, create a customer master key (CMK). Modify the key policy to provide kms:DescribeKey, kms:ReEncrypt*, kms:CreateGrant, and kms:Decrypt permissions to the AWS accounts that the AMI will be shared with. Modify the AMI permissions to specify the AWS account numbers that the AMI will be shared with.
- B. In the account where the AMI was created, create a customer master key (CMK). Modify the key policy to provide kms:DescribeKey, kms:ReEncrypt*, kms:CreateGrant, and kms:Decrypt permissions to the AWS accounts that the AMI will be shared with. Create a copy of the AMI, and specify the CMK. Modify the permissions on the copied AMI to specify the AWS account numbers that the AMI will be shared with.
- C. In the account where the AMI was created, create a customer master key (CMK). Modify the key policy to provide kms:DescribeKey, kms:ReEncrypt*, kms:CreateGrant, and kms:Decrypt permissions to the AWS accounts that the AMI will be shared with. Create a copy of the AMI, and specify the CMK. Modify the permissions on the copied AMI to make it public.
- D. In the account where the AMI was created, modify the key policy of the AWS managed key to provide kms:DescribeKey, kms:ReEncrypt*, kms:CreateGrant, and kms:Decrypt permissions to the AWS accounts that the AMI will be shared with. Modify the AMI permissions to specify the AWS account numbers that the AMI will be shared with.

Correct Answer: C

Section:

QUESTION 44

A company needs to restrict access to an Amazon S3 bucket to Amazon EC2 instances in a VPC only. All traffic must be over the AWS private network. What actions should the SysOps administrator take to meet these requirements?

- A. Create a VPC endpoint for the S3 bucket, and create an IAM policy that conditionally limits all S3 actions on the bucket to the VPC endpoint as the source.
- B. Create a VPC endpoint for the S3 bucket, and create an S3 bucket policy that conditionally limits all S3 actions on the bucket to the VPC endpoint as the source.
- C. Create a service-linked role for Amazon EC2 that allows the EC2 instances to interact directly with Amazon S3, and attach an IAM policy to the role that allows the EC2 instances full access to the S3 bucket.
- D. Create a NAT gateway in the VPC, and modify the VPC route table to route all traffic destined for Amazon S3 through the NAT gateway.

Correct Answer: B

Section:

QUESTION 45

A company has a VPC with public and private subnets. An Amazon EC2 based application resides in the private subnets and needs to process raw .csv files stored in an Amazon S3 bucket. A SysOps administrator has set up the correct IAM role with the required permissions for the application to access the S3 bucket, but the application is unable to communicate with the S3 bucket. Which action will solve this problem while adhering to least privilege access?

- A. Add a bucket policy to the S3 bucket permitting access from the IAM role.
- B. Attach an S3 gateway endpoint to the VPC. Configure the route table for the private subnet.

- C. Configure the route table to allow the instances on the private subnet access through the internet gateway.
- D. Create a NAT Gateway in a private subnet and configure the route table for the private subnets.

Correct Answer: C

Section:

QUESTION 46

A SysOps administrator noticed that a large number of Elastic IP addresses are being created on the company's AWS account, but they are not being associated with Amazon EC2 instance, and are incurring Elastic IP address charges in the monthly bill.

How can the administrator identify who is creating the Elastic IP addresses?

- A. Attach a cost-allocation tag to each requested Elastic IP address with the IAM user name of the developer who creates it.
- B. Query AWS CloudTrail logs by using Amazon Athena to search for Elastic IP address events.
- C. Create a CloudWatch alarm on the EIPCreated metric and send an Amazon SNS notification when the alarm triggers.
- D. Use Amazon Inspector to get a report of all Elastic IP addresses created in the last 30 days.

Correct Answer: A

Section:

QUESTION 47

A company uses an Amazon Simple Queue Service (Amazon SQS) standard queue with its application. The application sends messages to the queue with unique message bodies. The company decides to switch to an SQS FIFO queue. What must the company do to migrate to an SQS FIFO queue?

- A. Create a new SQS FIFO queue. Turn on content-based deduplication on the new FIFO queue. Update the application to include a message group ID in the messages.
- B. Create a new SQS FIFO queue. Update the application to include the DelaySeconds parameter in the messages.
- C. Modify the queue type from SQS standard to SQS FIFO. Turn off content-based deduplication on the queue. Update the application to include a message group ID in the messages.
- D. Modify the queue type from SQS standard to SQS FIFO. Update the application to send messages with identical message bodies and to include the DelaySeconds parameter in the messages.

Correct Answer: A

Section:

QUESTION 48

A database is running on an Amazon RDS Multi-AZ DB instance. A recent security audit found the database to be out of compliance because it was not encrypted. Which approach will resolve the encryption requirement?

- A. Log in to the RDS console and select the encryption box to encrypt the database.
- B. Create a new encrypted Amazon EBS volume and attach it to the instance.
- C. Encrypt the standby replica in the secondary Availability Zone and promote it to the primary instance.
- D. Take a snapshot of the RDS instance, copy and encrypt the snapshot and then restore to the new RDS instance.

Correct Answer: D

Section:

QUESTION 49

A SysOps administrator is tasked with deploying a company's infrastructure as code. The SysOps administrator wants to write a single template that can be reused for multiple environments. How should the SysOps administrator use AWS CloudFormation to create a solution?

- A. Use Amazon EC2 user data in a CloudFormation template.
- B. Use nested stacks to provision resources.

- C. Use parameters in a CloudFormation template
- D. Use stack policies to provision resources

Correct Answer: C

Section:

Explanation:

Reuse templates to replicate stacks in multiple environments After you have your stacks and resources set up, you can reuse your templates to replicate your infrastructure in multiple environments. For example, you can create environments for development, testing, and production so that you can test changes before implementing them into production. To make templates reusable, use the parameters, mappings, and conditions sections so that you can customize your stacks when you create them. For example, for your development environments, you can specify a lower-cost instance type compared to your production environment, but all other configurations and settings remain the same. <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html#reuse>

QUESTION 50

A company's web application is available through an Amazon CloudFront distribution and directly through an internet-facing Application Load Balancer (ALB) A SysOps administrator must make the application accessible only through the CloudFront distribution and not directly through the ALB. The SysOps administrator must make this change without changing the application code Which solution will meet these requirements?

- A. Modify the ALB type to internal Set the distribution's origin to the internal ALB domain name
- B. Create a Lambda@Edge function Configure the function to compare a custom header value in the request with a stored password and to forward the request to the origin in case of a match Associate the function with the distribution.
- C. Replace the ALB with a new internal ALB Set the distribution's origin to the internal ALB domain name Add a custom HTTP header to the origin settings for the distribution In the ALB listener add a rule to forward requests that contain the matching custom header and the header's value Add a default rule to return a fixed response code of 403.
- D. Add a custom HTTP header to the origin settings for the distribution in the ALB listener add a rule to forward requests that contain the matching custom header and the header's value Add a default rule to return a fixed response code of 403.

Correct Answer: D

Section:

Explanation:

To make the application accessible only through the CloudFront distribution and not directly through the Application Load Balancer (ALB), you can add a custom HTTP header to the origin settings for the CloudFront distribution. You can then create a rule in the ALB listener to forward requests that contain the matching custom header and its value to the origin. You can also add a default rule to the ALB listener to return a fixed response code of 403 for requests that do not contain the matching custom header. This will allow you to redirect all requests to the CloudFront distribution and block direct access to the application through the ALB. [https://docs.aws.amazon.com/AmazonCloudFront/latest/ DeveloperGuide/restrict-access-to-load-balancer.html](https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html)

QUESTION 51

A compliance team requires 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 52

A SysOps administrator is responsible for a large fleet of Amazon EC2 instances and must know whether any instances will be affected by upcoming hardware maintenance. Which option would provide this information with the LEAST administrative overhead?

- A. Deploy a third-party monitoring solution to provide real-time EC2 instance monitoring
- B. List any instances with failed system status checks using the AWS Management Console
- C. Monitor AWS CloudTrail for StopInstances API calls

D. Review the AWS Personal Health Dashboard

Correct Answer: D

Section:

QUESTION 53

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:

QUESTION 54

A SysOps administrator must configure a resilient tier of Amazon EC2 instances for a high performance computing (HPC) application. The HPC application requires minimum latency between nodes Which actions should the SysOps administrator take to meet these requirements? (Select TWO.)

- A. Create an Amazon Elastic File System (Amazon EFS) file system Mount the file system to the EC2 instances by using user data
- B. Create a Multi-AZ Network Load Balancer in front of the EC2 instances
- C. Place the EC2 instances in an Auto Scaling group within a single subnet
- D. Launch the EC2 instances into a cluster placement group
- E. Launch the EC2 instances into a partition placement group



Correct Answer: A, D

Section:

QUESTION 55

A SysOps administrator is unable to authenticate an AWS CLI call to an AWS service Which of the following is the cause of this issue?

- A. The IAM password is incorrect
- B. The server certificate is missing
- C. The SSH key pair is incorrect
- D. There is no access key

Correct Answer: C

Section:

QUESTION 56

A company is expanding its use of AWS services across its portfolios The company wants to provision AWS accounts for each team to ensure a separation of business processes for security compliance and billing Account creation and bootstrapping should be completed in a scalable and efficient way so new accounts are created with a defined baseline and governance guardrails in place A SysOps administrator needs to design a provisioning process that saves time and resources Which action should be taken to meet these requirements?

- A. Automate using AWS Elastic Beanstalk to provision the AWS accounts set up infrastructure and integrate with AWS Organizations
- B. Create bootstrapping scripts in AWS OpsWorks and combine them with AWS CloudFormation templates to provision accounts and infrastructure

- C. Use AWS Config to provision accounts and deploy instances using AWS Service Catalog
- D. Use AWS Control Tower to create a template in Account Factory and use the template to provision new accounts

Correct Answer: D

Section:

QUESTION 57

A SysOps administrator is unable to launch Amazon EC2 instances into a VPC because there are no available private IPv4 addresses in the VPC. Which combination of actions must the SysOps administrator take to launch the instances? (Select TWO.)

- A. Associate a secondary IPv4 CIDR block with the VPC
- B. Associate a primary IPv6 CIDR block with the VPC
- C. Create a new subnet for the VPC
- D. Modify the CIDR block of the VPC
- E. Modify the CIDR block of the subnet that is associated with the instances

Correct Answer: A, D

Section:

QUESTION 58

A SysOps administrator needs to develop a solution that provides email notification and inserts a record into a database every time a file is put into an Amazon S3 bucket. What is the MOST operationally efficient solution that meets these requirements?

- A. Set up an S3 event notification that targets an Amazon Simple Notification Service (Amazon SNS) topic. Create two subscriptions for the SNS topic. Use one subscription to send the email notification. Use the other subscription to invoke an AWS Lambda function that inserts the record into the database.
- B. Set up an Amazon CloudWatch alarm that enters ALARM state whenever an object is created in the S3 bucket. Configure the alarm to invoke an AWS Lambda function that sends the email notification and inserts the record into the database.
- C. Create an AWS Lambda function to send the email notification and insert the record into the database whenever a new object is detected in the S3 bucket. Invoke the function every minute with an Amazon EventBridge (Amazon CloudWatch Events) scheduled rule.
- D. Set up two S3 event notifications. Target a separate AWS Lambda function with each notification. Configure one function to send the email notification. Configure the other function to insert the record into the database.

Correct Answer: C

Section:

QUESTION 59

A company needs to upload gigabytes of files every day. The company needs 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 bucket as an origin.
- B. Create an Amazon ElastiCache cluster 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 60

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 IAM 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 61

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 SSL/TLS certificate 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 protocol listener on port 443. Attach an SSL/TLS certificate to listener port 443. Create a rule to redirect requests from port 80 to port 443.
- C. Create an Application Load Balancer that has two TCP listeners on port 80 and port 443. Attach an SSL/TLS 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 SSL/TLS certificate to listener port 443. Create a rule to redirect requests from port 80 to port 443.

Correct Answer: B

Section:

QUESTION 62

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 63

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 64

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: A

Section:

QUESTION 65

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 66

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 67

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 68

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 SNS) notification for the S3 bucket. Select DeleteObject for 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 69

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 70

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 71

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 72**

A global company handles a large amount of personally identifiable information (PII) 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 PII 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 S3 buckets over AWS Direct Connect or Site-to-Site VPN. <https://aws.amazon.com/blogs/architecture/choosing-your-vpc-endpoint-strategy-for-amazon-s3/>

QUESTION 73

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 74

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 CloudFormation 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 75

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 76

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 77

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 78

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 79

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 files, 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 80

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 IAM roles. The team requires an automated solution to create and manage the necessary IAM 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 IAM roles in each of the AWS accounts.
- B. Use AWS Directory Service with AWS Organizations to automatically associate the necessary IAM 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 IAM roles for the AWS accounts.

Correct Answer: D

Section:

QUESTION 81

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 IAM 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 82

A company is using Amazon Elastic Container Service (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 83

A company runs a website from Sydney, Australia. 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 in 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 84

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 IAM 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 85

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 86

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

- A. Analyze the AWS Cost and Usage Report by using Amazon Athena to identify 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 87

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 88

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 89

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 IAM user to the CloudFront distribution and grant the user permissions in the S3 bucket policy.
- D. Assign an IAM 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 90

A SysOps administrator is reviewing AWS Trusted Advisor recommendations. The SysOps administrator notices that all the application servers for a finance application are listed in the Low Utilization Amazon EC2 Instances check. The application runs on three instances across three Availability Zones. The SysOps administrator must reduce the cost of running the application without affecting the application's availability or design. Which solution will meet these requirements?

- A. Reduce the number of application servers.
- B. Apply rightsizing recommendations from AWS Cost Explorer to reduce the instance size.
- C. Provision an Application Load Balancer in front of the instances.
- D. Scale up the instance size of the application servers.

Correct Answer: C

Section:

QUESTION 91

A company is undergoing an external audit of its systems, which run wholly on AWS. A SysOps administrator must supply documentation of Payment Card Industry Data Security Standard (PCI DSS) compliance for the infrastructure managed by AWS.

Which set of action should the SysOps administrator take to meet this requirement?

- A. Download the applicable reports from the AWS Artifact portal and supply these to the auditors.
- B. Download complete copies of the AWS CloudTrail log files and supply these to the auditors.
- C. Download complete copies of the AWS CloudWatch logs and supply these to the auditors.
- D. Provide the auditors with administrative access to the production AWS account so that the auditors can determine compliance.

Correct Answer: A

Section:

QUESTION 92

A company wants to collect data from an application to use for analytics. For the first 90 days, the data will be infrequently accessed but must remain highly available. During this time, the company's analytics team requires access to the data in milliseconds. However, after 90 days, the company must retain the data for the long term at a lower cost. The retrieval time after 90 days must be less than 5 hours. Which solution will meet these requirements MOST cost-effectively?

- A. Store the data in S3 Standard-Infrequent Access (S3 Standard-IA) for the first 90 days. Set up an S3 Lifecycle rule to move the data to S3 Glacier Flexible Retrieval after 90 days.
- B. Store the data in S3 One Zone-Infrequent Access (S3 One Zone-IA) for the first 90 days. Set up an S3 Lifecycle rule to move the data to S3 Glacier Deep Archive after 90 days.
- C. Store the data in S3 Standard for the first 90 days. Set up an S3 Lifecycle rule to move the data to S3 Glacier Flexible Retrieval after 90 days.
- D. Store the data in S3 Standard for the first 90 days. Set up an S3 Lifecycle rule to move the data to S3 Glacier Deep Archive after 90 days.

Correct Answer: A

Section:

Explanation:

Glacier Deep Archive retrieval time more than 5 hours (it's 12 hours), so B&D out. S3 Standard IA is cheaper than S3 Standard. <https://aws.amazon.com/tw/s3/pricing/>

QUESTION 93

A company hosts several write-intensive applications. These applications use a MySQL database that runs on a single Amazon EC2 instance. The company asks a SysOps administrator to implement a highly available database solution that is ideal for multi-tenant workloads.

Which solution should the SysOps administrator implement to meet these requirements?

- A. Create a second EC2 instance for MySQL. Configure the second instance to be a read replica.
- B. Migrate the database to an Amazon Aurora DB cluster. Add an Aurora Replica.
- C. Migrate the database to an Amazon Aurora multi-master DB cluster.
- D. Migrate the database to an Amazon RDS for MySQL DB instance.

Correct Answer: C

Section:

QUESTION 94

A SysOps administrator created an AWS CloudFormation template that provisions Amazon EC2 instances, an Elastic Load Balancer (ELB), and an Amazon RDS DB instance. During stack creation, the creation of the EC2 instances and the creation of the ELB are successful. However, the creation of the DB instance fails.

What is the default behavior of CloudFormation in this scenario?

- A. CloudFormation will roll back the stack and delete the stack.
- B. CloudFormation will roll back the stack but will not delete the stack.
- C. CloudFormation will prompt the user to roll back the stack or continue.
- D. CloudFormation will successfully complete the stack but will report a failed status for the DB instance.

Correct Answer: C

Section:

QUESTION 95

A company needs to deploy a new workload on AWS. The company must encrypt all data at rest and must rotate the encryption keys once each year. The workload uses an Amazon RDS for MySQL Multi-AZ database for data storage. Which configuration approach will meet these requirements?

- A. Enable Transparent Data Encryption (TDE) in the MySQL configuration file. Manually rotate the key every 12 months.
- B. Enable RDS encryption on the database at creation time by using the AWS managed key for Amazon RDS.
- C. Create a new AWS Key Management Service (AWS KMS) customer managed key. Enable automatic key rotation. Enable RDS encryption on the database at creation time by using the KMS key.



- D. Create a new AWS Key Management Service (AWS KMS) customer managed key. Enable automatic key rotation. Enable encryption on the Amazon Elastic Block Store (Amazon EBS) volumes that are attached to the RDS DB instance.

Correct Answer: C

Section:

Explanation:

This configuration approach will meet the requirement of encrypting all data at rest and rotating the encryption keys once each year. By creating a new AWS KMS customer managed key and enabling automatic key rotation, the encryption keys will be rotated automatically every year. By enabling RDS encryption on the database at creation time using the KMS key, all data stored in the RDS for MySQL Multi-AZ database will be encrypted at rest. This approach provide more control over key management and rotation and provide additional security benefits

QUESTION 96

A SysOps administrator wants to upload a file that is 1 TB in size from on-premises to an Amazon S3 bucket using multipart uploads. What should the SysOps administrator do to meet this requirement?

- A. Upload the file using the S3 console.
- B. Use the s3api copy-object command.
- C. Use the s3api put-object command.
- D. Use the s3 cp command.

Correct Answer: D

Section:

Explanation:

It's a best practice to use aws s3 commands (such as aws s3 cp) for multipart uploads and downloads, because these aws s3 commands automatically perform multipart uploading and downloading based on the file size. By comparison, aws s3api commands, such as aws s3api create-multipart-upload, should be used only when aws s3 commands don't support a specific upload need, such as when the multipart upload involves multiple servers, a multipart upload is manually stopped and resumed later, or when the aws s3 command doesn't support a required request parameter.<https://aws.amazon.com/premiumsupport/knowledge-center/s3-multipart-upload-cli/>

QUESTION 97

A SysOps administrator is responsible for a company's security groups. The company wants to maintain a documented trail of any changes that are made to the security groups. The SysOps administrator must receive notification whenever the security groups change.

Which solution will meet these requirements?

- A. Set up Amazon Detective to record security group changes. Specify an Amazon CloudWatch Logs log group to store configuration history logs. Create an Amazon Simple Queue Service (Amazon SQS) queue for notifications about configuration changes. Subscribe the SysOps administrator's email address to the SQS queue.
- B. Set up AWS Systems Manager Change Manager to record security group changes. Specify an Amazon CloudWatch Logs log group to store configuration history logs. Create an Amazon Simple Notification Service (Amazon SNS) topic for notifications about configuration changes. Subscribe the SysOps administrator's email address to the SNS topic.
- C. Set up AWS Config to record security group changes. Specify an Amazon S3 bucket as the location for configuration snapshots and history files. Create an Amazon Simple Notification Service (Amazon SNS) topic for notifications about configuration changes. Subscribe the SysOps administrator's email address to the SNS topic.
- D. Set up Amazon Detective to record security group changes. Specify an Amazon S3 bucket as the location for configuration snapshots and history files. Create an Amazon Simple Notification Service (Amazon SNS) topic for notifications about configuration changes. Subscribe the SysOps administrator's email address to the SNS topic.

Correct Answer: D

Section:

QUESTION 98

A SysOps administrator created an Amazon VPC with an IPv6 CIDR block, which requires access to the internet. However, access from the internet towards the VPC is prohibited. After adding and configuring the required components to the VPC, the administrator is unable to connect to any of the domains that reside on the internet.

What additional route destination rule should the administrator add to the route tables?

- A. Route ::/0 traffic to a NAT gateway
- B. Route ::/0 traffic to an internet gateway

- C. Route 0.0.0.0/0 traffic to an egress-only internet gateway
- D. Route ::/0 traffic to an egress-only internet gateway

Correct Answer: D

Section:

Explanation:

<https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html>

QUESTION 99

A company is rolling out a new version of its website. Management wants to deploy the new website in a limited rollout to 20% of the company's customers. The company uses Amazon Route 53 for its website's DNS solution. Which configuration will meet these requirements?

- A. Create a failover routing policy. Within the policy, configure 80% of the website traffic to be sent to the original resource. Configure the remaining 20% of traffic as the failover record that points to the new resource.
- B. Create a multivalue answer routing policy. Within the policy, create 4 records with the name and IP address of the original resource. Configure 1 record with the name and IP address of the new resource.
- C. Create a latency-based routing policy. Within the policy, configure a record pointing to the original resource with a weight of 80. Configure a record pointing to the new resource with a weight of 20.
- D. Create a weighted routing policy. Within the policy, configure a weight of 80 for the record pointing to the original resource. Configure a weight of 20 for the record pointing to the new resource.

Correct Answer: C

Section:

QUESTION 100

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 101

A company has a simple web application that runs on a set of Amazon EC2 instances behind an Elastic Load Balancer in the eu-west-2 Region. Amazon Route 53 holds a DNS record for the application with a simple routing policy. Users from all over the world access the application through their web browsers. The company needs to create additional copies of the application in the us-east-1 Region and in the ap-south-1 Region. The company must direct users to the Region that provides the fastest response times when the users load the application. What should a SysOps administrator do to meet these requirements?

- A. In each new Region, create a new Elastic Load Balancer and a new set of EC2 Instances to run a copy of the application. Transition to a geolocation routing policy.
- B. In each new Region, create a copy of the application on new EC2 instances. Add these new EC2 instances to the Elastic Load Balancer in eu-west-2. Transition to a latency routing policy.
- C. In each new Region, create a copy of the application on new EC2 instances. Add these new EC2 instances to the Elastic Load Balancer in eu-west-2. Transition to a multivalue routing policy.
- D. In each new Region, create a new Elastic Load Balancer and a new set of EC2 instances to run a copy of the application. Transition to a latency routing policy.

Correct Answer: B

Section:

QUESTION 102

A company is managing many accounts by using a single organization in AWS Organizations. The organization has all features enabled. The company wants to turn on AWS Config in all the accounts of the organization and in all AWS Regions.

What should a Sysops administrator do to meet these requirements in the MOST operationally efficient way?

- A. Use AVVS CloudFormation StackSets to deploy stack instances that turn on AWS Config in all accounts and in all Regions.
- B. Use AWS CloudFormation StackSets to deploy stack policies that turn on AWS Config in all accounts and in all Regions.
- C. Use service control policies (SCPs) to configure AWS Config in all accounts and in all Regions.
- D. Create a script that uses the AWS CLI to turn on AWS Config in all accounts in the organization. Run the script from the organization's management account.

Correct Answer: C

Section:

QUESTION 103

A company updates its security policy to prohibit the public exposure of any data in Amazon S3 buckets in the company's account. What should a SysOps administrator do to meet this requirement?

- A. Turn on S3 Block Public Access from the account level.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to enforce that all S3 objects are private.
- C. Use Amazon Inspector to search for S3 buckets and to automatically reset S3 ACLs if any public S3 buckets are found.
- D. Use S3 Object Lambda to examine S3 ACLs and to change any public S3 ACLs to private.

Correct Answer: A

Section:

Explanation:

Using Amazon S3 Block Public Access as a centralized way to limit public access. Block Public Access settings override bucket policies and object permissions. Be sure to enable Block Public Access for all accounts and buckets that you don't want publicly accessible.

<https://aws.amazon.com/premiumsupport/knowledge-center/secure-s3resources/#:~:text=Using%20Amazon%20S3%20Block%20Public,don't%20want%20publicly%20accessible.>

QUESTION 104

A SysOps administrator needs to secure the credentials for an Amazon RDS database that is created by an AWS CloudFormation template. The solution must encrypt the credentials and must support automatic rotation. Which solution will meet these requirements?

- A. Create an AWS::SecretsManager::Secret resource in the CloudFormation template. Reference the credentials in the AWS::RDS::DBInstance resource by using the resolve:secretsmanager dynamic reference.
- B. Create an AWS::SecretsManager::Secret resource in the CloudFormation template. Reference the credentials in the AWS::RDS::DBInstance resource by using the resolve:ssm-secure dynamic reference.
- C. Create an AWS::SSM::Parameter resource in the CloudFormation template. Reference the credentials in the AWS::RDS::DBInstance resource by using the resolve:ssm dynamic reference.
- D. Create parameters for the database credentials in the CloudFormation template. Use the Ref intrinsic function to provide the credentials to the AWS::RDS::DBInstance resource.

Correct Answer: A

Section:

QUESTION 105

A company has two VPC networks named VPC A and VPC B. The VPC A CIDR block is 10.0.0.0/16 and the VPC B CIDR block is 172.31.0.0/16. The company wants to establish a VPC peering connection named pcx-12345 between both VPCs.

Which rules should appear in the route table of VPC A after configuration? (Select TWO.)

- A. Destination: 10.0.0.0/16, Target: Local
- B. Destination: 172.31.0.0/16, Target: Local
- C. Destination: 10.0.0.0/16, Target: pcx-12345
- D. Destination: 172.31.0.0/16, Target: pcx-12345

E. Destination: 10.0.0.0/16. Target: 172.31.0.0/16

Correct Answer: A, D

Section:

Explanation:

<https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-routing.html>

QUESTION 106

A company has a public website that recently experienced problems. Some links led to missing webpages, and other links rendered incorrect webpages. The application infrastructure was running properly, and all the provisioned resources were healthy. Application logs and dashboards did not show any errors, and no monitoring alarms were raised. Systems administrators were not aware of any problems until end users reported the issues. The company needs to proactively monitor the website for such issues in the future and must implement a solution as soon as possible. Which solution will meet these requirements with the LEAST operational overhead?

- A. Rewrite the application to surface a custom error to the application log when issues occur. Automatically parse logs for errors. Create an Amazon CloudWatch alarm to provide alerts when issues are detected.
- B. Create an AWS Lambda function to test the website. Configure the Lambda function to emit an Amazon CloudWatch custom metric when errors are detected. Configure a CloudWatch alarm to provide alerts when issues are detected.
- C. Create an Amazon CloudWatch Synthetics canary. Use the CloudWatch Synthetics Recorder plugin to generate the script for the canary run. Configure the canary in line with requirements. Create an alarm to provide alerts when issues are detected.

Correct Answer: A

Section:

QUESTION 107

A company hosts a database on an Amazon RDS Multi-AZ DB instance. The database is not encrypted.

The company's new security policy requires all AWS resources to be encrypted at rest and in transit. What should a SysOps administrator do to encrypt the database?

- A. Configure encryption on the existing DB instance.
- B. Take a snapshot of the DB instance. Encrypt the snapshot. Restore the snapshot to the same DB instance.
- C. Encrypt the standby replica in a secondary Availability Zone. Promote the standby replica to the primary DB instance.
- D. Take a snapshot of the DB instance. Copy and encrypt the snapshot. Create a new DB instance by restoring the encrypted copy.

Correct Answer: B

Section:

QUESTION 108

A company uses an Amazon S3 bucket to store data files. The S3 bucket contains hundreds of objects. The company needs to replace a tag on all the objects in the S3 bucket with another tag. What is the MOST operationally efficient way to meet this requirement?

- A. Use S3 Batch Operations. Specify the operation to replace all object tags.
- B. Use the AWS CLI to get the tags for each object. Save the tags in a list. Use S3 Batch Operations. Specify the operation to delete all object tags. Use the AWS CLI and the list to retag the objects.
- C. Use the AWS CLI to get the tags for each object. Save the tags in a list. Use the AWS CLI and the list to remove the object tags. Use the AWS CLI and the list to retag the objects.
- D. Use the AWS CLI to copy the objects to another S3 bucket. Add the new tag to the copied objects. Delete the original objects.

Correct Answer: A

Section:

Explanation:

Ref. <https://aws.amazon.com/es/blogs/storage/adding-and-removing-object-tags-with-s3-batch-operations/>

QUESTION 109

A company runs several workloads on AWS. The company identifies five AWS Trusted Advisor service quota metrics to monitor in a specific AWS Region. The company wants to receive email notification each time resource usage exceeds 60% of one of the service quotas.

Which solution will meet these requirements?

- A. Create five Amazon CloudWatch alarms, one for each Trusted Advisor service quota metric.
Configure an Amazon Simple Notification Service (Amazon SNS) topic for email notification each time that usage exceeds 60% of one of the service quotas.
- B. Create five Amazon CloudWatch alarms, one for each Trusted Advisor service quota metric.
Configure an Amazon Simple Queue Service (Amazon SQS) queue for email notification each time that usage exceeds 60% of one of the service quotas.
- C. Use the AWS Service Health Dashboard to monitor each Trusted Advisor service quota metric.
Configure an Amazon Simple Queue Service (Amazon SQS) queue for email notification each time that usage exceeds 60% of one of the service quotas.
- D. Use the AWS Service Health Dashboard to monitor each Trusted Advisor service quota metric.
Configure an Amazon Simple Notification Service (Amazon SNS) topic for email notification each time that usage exceeds 60% of one of the service quotas.

Correct Answer: A

Section:

QUESTION 110

A company runs its entire suite of applications on Amazon EC2 instances. The company plans to move the applications to containers and AWS Fargate. Within 6 months, the company plans to retire its EC2 instances and use only Fargate. The company has been able to estimate its future Fargate costs.

A SysOps administrator needs to choose a purchasing option to help the company minimize costs.

The SysOps administrator must maximize any discounts that are available and must ensure that there are no unused reservations. Which purchasing option will meet these requirements?

- A. Compute Savings Plans for 1 year with the No Upfront payment option
- B. Compute Savings Plans for 1 year with the Partial Upfront payment option
- C. EC2 Instance Savings Plans for 1 year with the All Upfront payment option
- D. EC2 Reserved Instances for 1 year with the Partial Upfront payment option



Correct Answer: C

Section:

QUESTION 111

A company creates a new member account by using AWS Organizations. A SysOps administrator needs to add AWS Business Support to the new account. Which combination of steps must the SysOps administrator take to meet this requirement? (Select TWO.)

- A. Sign in to the new account by using IAM credentials. Change the support plan.
- B. Sign in to the new account by using root user credentials. Change the support plan.
- C. Use the AWS Support API to change the support plan.
- D. Reset the password of the account root user.
- E. Create an IAM user that has administrator privileges in the new account.

Correct Answer: B, E

Section:

Explanation:

The best combination of steps to meet this requirement is to sign in to the new account by using root user credentials and change the support plan, and to create an IAM user that has administrator privileges in the new account. Signing in to the new account by using root user credentials will allow the SysOps administrator to access the account and change the support plan to AWS Business Support. Additionally, creating an IAM user that has administrator privileges in the new account will ensure that the SysOps administrator has the necessary access to manage the account and make changes to the support plan if necessary. Reference:[1]

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html#orgs_manage_accounts_access_signin-root

QUESTION 112

A company needs to automatically monitor an AWS account for potential unauthorized AWS Management Console logins from multiple geographic locations. Which solution will meet this requirement?

- A. Configure Amazon Cognito to detect any compromised IAM credentials.
- B. Set up Amazon Inspector. Scan and monitor resources for unauthorized logins.
- C. Set up AWS Config. Add the iam-policy-blacklisted-check managed rule to the account.
- D. Configure Amazon GuardDuty to monitor the UnauthorizedAccess:IAMUser/ConsoleLoginSuccess finding.

Correct Answer: D

Section:

QUESTION 113

Application A runs on Amazon EC2 instances behind a Network Load Balancer (NLB). The EC2 instances are in an Auto Scaling group and are in the same subnet that is associated with the NLB. Other applications from an on-premises environment cannot communicate with Application A on port 8080.

To troubleshoot the issue, a SysOps administrator analyzes the flow logs. The flow logs include the following records:

```
2 123456789010 eni-1235b8ca123456789 192.168.0.13 172.31.16.139 59003 8080 1 4 336 1432917027 1432917142 ACCEPT OK
2 123456789010 eni-1235b8ca123456789 172.31.16.139 192.168.0.13 8080 59003 1 4 336 1432917094 1432917142 REJECT OK
```

What is the reason for the rejected traffic?

- A. The security group of the EC2 instances has no Allow rule for the traffic from the NLB.
- B. The security group of the NLB has no Allow rule for the traffic from the on-premises environment.
- C. The ACL of the on-premises environment does not allow traffic to the AWS environment.
- D. The network ACL that is associated with the subnet does not allow outbound traffic for the ephemeral port range.

Correct Answer: A

Section:



QUESTION 114

A company's SysOps administrator deploys a public Network Load Balancer (NLB) in front of the company's web application. The web application does not use any Elastic IP addresses. Users must access the web application by using the company's domain name. The SysOps administrator needs to configure Amazon Route 53 to route traffic to the NLB. Which solution will meet these requirements MOST cost-effectively?

- A. Create a Route 53 AAAA record for the NLB.
- B. Create a Route 53 alias record for the NLB.
- C. Create a Route 53 CAA record for the NLB.
- D. Create a Route 53 CNAME record for the NLB.

Correct Answer: B

Section:

QUESTION 115

A company has an application that is deployed in 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 116

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 the 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 117

SIMULATION

You need to update an existing AWS CloudFormation stack. If needed, a copy of the CloudFormation template is available in an Amazon S3 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 EC2 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 CFServiceRole 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 address range 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 CFServiceRole 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 118

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 tag all the untagged 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 119

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 120

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 kubectl command line tool.

Which of the following must be configured on the Sysops administrator's machine so that kubectl can 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 121

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 122

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 123

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 124

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 125

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 126

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-1 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 127

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 128

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 129

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 130

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 131

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 132

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> In governance mode, users can't overwrite or delete an object version or alter its lock settings unless they have special permissions. With governance mode, you protect objects against being deleted by most users, but you can still grant some users permission to alter the retention settings 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, a protected object version can't be overwritten or deleted by any user, including the root user in your AWS account. When an object is locked in compliance mode, its retention mode can't be changed, and its retention period can't be shortened. Compliance mode helps ensure that an object version can't be overwritten or deleted for the duration of the retention period.

QUESTION 133

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 134

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, read-many (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 Standard-Infrequent 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 135

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 136

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 137

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 138

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 139

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 140

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.

QUESTION 141

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] <https://docs.aws.amazon.com/savingsplans/latest/userguide/sp-usingBudgets.html>

QUESTION 142

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 143

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] https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

QUESTION 144

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] <https://aws.amazon.com/premiumsupport/knowledge-center/organizational-viewhealth-dashboard/>

QUESTION 145

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 146

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 147

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 148

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 149

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 150

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 createCacheCluster API and specify a new cacheNodeType.
- C. Use the ModifyCacheParameterGroup API and specify a new CacheNodeType.
- D. Use the RebootCacheCluster 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 151

A SysOps administrator is examining the following AWS CloudFormation template:

```
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 CloudFormation template was omitted.
- C. The PrivateDnsName 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 152

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 IAM role that has access to the object. Instruct the users to assume the role.
- C. Create an IAM 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 153

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 154

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 state 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 155

A company needs to deploy instances of an application and associated infrastructure to multiple AWS Regions. The company wants to use a single AWS CloudFormation template to achieve this goal. The company uses AWS Organizations and wants to administer and run this template from a central administration account.

What should a SysOps administrator do to meet these requirements?

- A. Create a CloudFormation template that is stored in Amazon S3. Configure Cross-Region Replication (CRR) on the S3 bucket. Reference the required accounts and remote Regions in the input template parameters.
- B. In the central administration account, create a CloudFormation primary template that loads CloudFormation nested stacks from Amazon S3 buckets in the target Regions.
- C. Create CloudFormation nested stacks by using a primary template in the central administration account. Configure the required accounts and Regions for deployment of the nested stacks.
- D. Create a CloudFormation stack set that includes service-managed permissions. Deploy the stack set into the required accounts and Regions from the central administration account.

Correct Answer: D

Section:

Explanation:

AWS CloudFormation StackSets extends the functionality of stacks by enabling you to create, update, or delete stacks across multiple accounts and regions with a single operation. Using a stack set, the SysOps administrator can manage deployments across different regions and accounts within AWS Organizations efficiently.

Setting up StackSets: First, define your CloudFormation template that describes all the resources that need to be deployed across the regions. Store this template in an S3 bucket accessible by the central administration account.

Service-Managed Permissions: When creating a stack set, select the option for service-managed permissions if you are using AWS Organizations. This allows AWS CloudFormation to automatically set up the necessary permissions in the target accounts.

Deploying the Stack Set: From the central administration account, create the stack set and specify the target accounts and regions. CloudFormation will then ensure that the resources defined in the template are instantiated in each of the specified regions and accounts.

This method simplifies management and ensures consistency of infrastructure across multiple regions and accounts, leveraging the organizational units in AWS Organizations for centralized governance.

QUESTION 156

A company decides to stop non-production Amazon EC2 instances during the EC2 instances. The company's IT manager must receive notification in near real time whenever an EC2 instance that has an environment type tag value of non-production is started during the night.

Which solution will meet this requirement with the MOST operational efficiency?

- A. Configure an AWS Lambda function with an SMTP client library. Subscribe the Lambda function to the AWS Health Dashboard to receive notification whenever an EC2 instance is in the running state. Configure the Lambda function to use Amazon Pinpoint to send email notifications to the IT manager. Deploy a second Lambda function to throttle calls from the first Lambda function during the daytime.
- B. Deploy an AWS Lambda function that queries the Amazon EC2 API to determine the state of each EC2 instance. Use the EC2 instance scheduler to configure the Lambda function to run every minute during the night and to send an email notification to the IT manager for each non-production EC2 instance that is in the running state.
- C. Create an Amazon EventBridge rule that includes the EC2 Instance State-change Notification event type. Filter the event to capture only the running state. Create an AWS Lambda function as a target of the rule. Configure the Lambda function to check the current time and the EC2 instances' tags to determine the environment type. Create an Amazon Simple Notification Service (Amazon SNS) topic as a target of the Lambda function for notifications. Subscribe the IT manager's email address to the SNS topic.
- D. Store the EC2 instance metadata, including the environment type, in an Amazon DynamoDB table. Deploy a custom application to an EC2 instance. Configure the custom application to poll the DynamoDB data every minute during the night and to query the Amazon EC2 API to determine the state of each instance. Additionally, configure the custom application to send an email notification to the IT manager for each non-production EC2 instance that is in the running state.

Correct Answer: C

Section:**Explanation:**

The requirement is to monitor and notify whenever a non-production EC2 instance is started during the night. Amazon EventBridge offers a robust solution by triggering workflows in response to events.

Setting up Amazon EventBridge: Create an EventBridge rule that listens for the 'EC2 Instance State-change Notification' event. Configure the rule to trigger only when instances transition to the 'running' state.

Lambda Function: Attach a Lambda function as the target of the EventBridge rule. This function will execute when an EC2 instance starts. Inside the Lambda function, implement logic to check the current time and confirm it is during the night hours. Additionally, the function will check the instance's tags to verify if it's labeled as 'non-production'.

Notification via Amazon SNS: If the conditions are met (non-production and nighttime), the Lambda function publishes a message to an Amazon SNS topic specifically set up for this alert. The IT manager is subscribed to this topic, enabling them to receive an email notification almost instantaneously when the event occurs.

This solution is operationally efficient as it leverages serverless components that are inherently scalable and cost-effective, providing real-time monitoring and notifications without the need for continuous polling or complex infrastructure.

QUESTION 157

A SysOps administrator must configure Amazon S3 to host a simple nonproduction webpage. The SysOps administrator has created an empty S3 bucket from the AWS Management Console. The S3 bucket has the default configuration in place.

Which combination of actions should the SysOps administrator take to complete this process? (Choose two.)

- A. Configure the S3 bucket by using the 'Redirect requests for an object' functionality to point to the bucket root URL.
- B. Turn off the 'Block all public access' setting. Allow public access by using a bucket ACL that contains <Permission>WEBSITE</Permission>.
- C. Turn off the 'Block all public access' setting. Allow public access by using a bucket ACL that allows access to the AuthenticatedUsers grantee.
- D. Turn off the 'Block all public access' setting. Set a bucket policy that allows 'Principal': the s3:GetObject action.
- E. Create an index.html document. Configure static website hosting, and upload the index document to the S3 bucket.

Correct Answer: D, E

Section:**Explanation:**

To host a static website on Amazon S3, the SysOps administrator needs to configure the bucket for public access and set up the static website hosting. Here's how to complete this process:

Turn off 'Block all public access': Amazon S3 buckets have 'Block all public access' settings enabled by default for security. Since the webpage needs to be accessible publicly, this setting must be disabled. This step is crucial to allow public read access to the web content.

Set a bucket policy: After disabling 'Block all public access,' set a bucket policy that explicitly allows public read access to the S3 bucket. This policy should allow the s3:GetObject action for everyone, which can be set by specifying 'Principal': '*'. This policy ensures that anyone can view the webpage but does not grant permissions to modify or delete the content.

Create an index.html document and configure static website hosting: The next step is to create an index.html file, which will serve as the entry point of the website. After creating this file, upload it to the bucket. Then, configure the bucket for static website hosting through the S3 management console. This setting enables the S3 bucket to serve the webpage directly from the index.html file.

Combining these actions, the S3 bucket will be properly configured to host and serve the static website with minimal operational overhead and maximum accessibility.

QUESTION 158

A company is experiencing issues with legacy software running on Amazon EC2 instances. Errors occur when the total CPU utilization on the EC2 instances exceeds 80%. A short-term solution is required while the software is being rewritten. A SysOps administrator is tasked with creating a solution to restart the instances when the CPU utilization rises above 80%.

Which solution meets these requirements with the LEAST operational overhead?

- A. Write a script that monitors the CPU utilization of the EC2 instances and reboots the instances when utilization exceeds 80%. Run the script as a cron job.
- B. Add an Amazon CloudWatch alarm for CPU utilization and configure the alarm action to reboot the EC2 instances.
- C. Create an Amazon EventBridge rule using the predefined patterns for CPU utilization of the EC2 instances. When utilization exceeds 80%, invoke an AWS Lambda function to restart the instances.
- D. Add an Amazon CloudWatch alarm for CPU utilization and configure an AWS Systems Manager Automation runbook to reboot the EC2 instances when utilization exceeds 80%.

Correct Answer: B

Section:**Explanation:**

The simplest and most efficient solution to ensure that EC2 instances are restarted when CPU utilization exceeds 80% is to use Amazon CloudWatch alarms:

Create a CloudWatch Alarm: Navigate to the CloudWatch dashboard in the AWS Management Console and create a new alarm. Set the alarm to monitor the CPU utilization metric of the EC2 instances.

Set the Alarm Condition: Configure the alarm to trigger when the CPU utilization exceeds 80%. You can specify this threshold in the alarm settings.

Configure Alarm Actions: In the actions settings of the alarm, select the option to reboot the instance. This action ensures that the instance is automatically restarted whenever the alarm condition is met, without the need for manual intervention or additional scripts.

This method leverages AWS's native capabilities, minimizing operational overhead and eliminating the need for external tools or custom scripts.

QUESTION 159

ASysOps administrator configures an application to run on Amazon EC2 instances behind an Application Load Balancer (ALB) in a simple scaling Auto Scaling group with the default settings. The Auto Scaling group is configured to use the RequestCountPerTarget metric for scaling. The SysOps administrator notices that the RequestCountPerTarget metric exceeded the specified limit twice in 180 seconds.

How will the number of EC2 instances in this Auto Scaling group be affected in this scenario?

- A. The Auto Scaling group will launch an additional EC2 instance every time the RequestCountPerTarget metric exceeds the predefined limit.
- B. The Auto Scaling group will launch one EC2 instance and will wait for the default cooldown period before launching another instance.
- C. The Auto Scaling group will send an alert to the ALB to rebalance the traffic and not add new EC2 instances until the load is normalized.
- D. The Auto Scaling group will try to distribute the traffic among all EC2 instances before launching another instance.

Correct Answer: B

Section:

Explanation:

When using the RequestCountPerTarget metric for scaling in an Auto Scaling group, the behavior of instance scaling follows specific rules set by Auto Scaling policies and cooldown periods:

Scaling Trigger: The Auto Scaling group triggers a scaling action whenever the RequestCountPerTarget exceeds the predefined limit set in the scaling policy.

Cooldown Period: After launching an EC2 instance due to a scaling action, the Auto Scaling group enters a cooldown period. During this period, despite further breaches of the threshold, no additional instances will be launched. This is designed to give the newly launched instance time to start and begin handling traffic, preventing the Auto Scaling group from launching too many instances too quickly.

This mechanism helps maintain efficient use of resources by adapting to changes in load while avoiding rapid, unnecessary scaling actions.

QUESTION 160

A company hosts a web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses Amazon Route 53 to route traffic.

The company also has a static website that is configured in an Amazon S3 bucket.

A SysOps administrator must use the static website as a backup to the web application. The failover to the static website must be fully automated.

Which combination of actions will meet these requirements? (Choose two.)

- A. Create a primary failover routing policy record. Configure the value to be the ALB.
- B. Create an AWS Lambda function to switch from the primary website to the secondary website when the health check fails.
- C. Create a primary failover routing policy record. Configure the value to be the ALB. Associate the record with a Route 53 health check.
- D. Create a secondary failover routing policy record. Configure the value to be the static website. Associate the record with a Route 53 health check.
- E. Create a secondary failover routing policy record. Configure the value to be the static website.

Correct Answer: C, E

Section:

QUESTION 161

A company has an application that uses a scheduled AWS Lambda function to retrieve datasets from external sources over the internet. The function is not associated with a VPC. The company is modifying the application to store the information that the Lambda function retrieves on an Amazon RDS DB instance in a private subnet. The VPC has two public subnets and two private subnets.

A SysOps administrator must deploy a solution that allows the Lambda function to access the new database and continue to access the internet.

Which solution meets these requirements?

- A. Create a new Lambda function with VPC access and an Elastic IP address. Attach the function to public subnets in two Availability Zones. Associate a security group with the Elastic IP address. Configure the security group outbound rules to allow Lambda to access the required resources.

- B. Create a new Lambda function with VPC access and two public IP addresses. Attach the function to public subnets in the same Availability Zones that the database uses. Associate a security group with the function. Configure the security group inbound rules to allow Lambda to access the required resources.
- C. Reconfigure the Lambda function for VPC access. Add NAT gateways to the public subnets in the VPC. Add route table entries in the private subnets to route through the NAT gateways to the internet. Attach the function to the private subnets that support the database. Associate a security group with the function. Configure the security group outbound rules to allow Lambda to access the internet.
- D. Reconfigure the Lambda function for VPC access. Attach the function to the private subnets. Add route table entries in the private subnets to route through the internet gateway to the internet. Associate a security group with the subnets. Configure the security group inbound rules to allow Lambda to access the required resources through the internet gateway.

Correct Answer: C

Section:

QUESTION 162

A company is running distributed computing software to manage a fleet of 20 Amazon EC2 instances for calculations. The fleet includes 2 control nodes and 18 task nodes to run the calculations. Control nodes can automatically start the task nodes.

Currently, all the nodes run on demand. The control nodes must be available 24 hours a day, 7 days a week. The task nodes run for 4 hours each day. A SysOps administrator needs to optimize the cost of this solution. Which combination of actions will meet these requirements? (Choose two.)

- A. Purchase EC2 Instance Savings Plans for the control nodes.
- B. Use Dedicated Hosts for the control nodes.
- C. Use Reserved Instances for the task nodes.
- D. Use Spot Instances for the control nodes. Use On-Demand Instances if there is no Spot availability.
- E. Use Spot Instances for the task nodes. Use On-Demand Instances if there is no Spot availability.

Correct Answer: A, E

Section:

Explanation:

To optimize the cost of a computing environment consisting of control nodes that are always on and task nodes that operate for a limited number of hours each day, consider the following strategies:

Purchase EC2 Instance Savings Plans for the Control Nodes: Since the control nodes are required to be operational 24/7, purchasing EC2 Instance Savings Plans is a cost-effective choice. These plans provide a lower price compared to on-demand instances, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a one or three-year period.

Use Spot Instances for the Task Nodes: Given that task nodes are used for a shorter duration (4 hours a day) and presumably can tolerate interruptions, using Spot Instances can significantly reduce costs. Spot Instances offer unused EC2 capacity at a fraction of the regular price, which can lead to substantial cost savings. Additionally, configure the system to fall back to On-Demand Instances during periods when Spot Instances are not available to ensure availability.

This combination leverages cost savings for continuous use and flexible, lower-cost options for intermittent use, optimizing overall operational costs efficiently.

QUESTION 163

A company has a secure website running on Amazon EC2 instances behind an Application Load Balancer (ALB). An SSL certificate from AWS Certificate Manager (ACM) is used on the ALB. Users with legacy web browsers are experiencing issues with the website.

How should the SysOps administrator resolve these issues in the MOST operationally efficient manner?

- A. Create a new SSL certificate in ACM and install the new certificate on the ALB to support legacy web browsers.
- B. Create a second ALB and install a custom SSL certificate with a different domain name on the second ALB to support legacy web browsers.
- C. Remove the ALB from the configuration and install a custom SSL certificate on each web server.
- D. Update the SSL negotiation configuration of the ALB with a security policy that contains ciphers for legacy web browsers.

Correct Answer: D

Section:

Explanation:

The issues experienced by users with legacy browsers typically stem from the SSL/TLS ciphers that are supported or enforced by the ALB. Modern security policies may exclude older ciphers that are necessary for compatibility with older browsers. Here's how to resolve it:

Access the ALB Settings: Go to the AWS Management Console, navigate to the ALB settings, and locate the SSL negotiation configurations.

Modify Security Policy: Update the SSL/TLS security policy on the ALB to include ciphers that are compatible with legacy browsers. AWS provides predefined security policies, and some of these policies are designed to support older ciphers while still maintaining a level of security that complies with general best practices.

Apply Changes: Once the security policy is updated, the ALB will start using this new configuration, which should resolve compatibility issues with legacy browsers without needing to replace the SSL certificate or alter the infrastructure.

This solution maintains the operational efficiency of the setup and avoids the need for additional resources like a second ALB or new certificates.

QUESTION 164

A Sysops administrator launches an Amazon EC2 instance from a Windows Amazon Machine Image (AMI). The EC2 instance includes additional Amazon Elastic Block Store (Amazon EBS) volumes. When the instance is launched, none of the additional Amazon Elastic Block Store (Amazon EBS) volumes are initialized and ready for use through a drive letter. The SysOps administrator needs to automate the EBS volume initialization. Which solution will meet these requirements in the MOST operationally efficient way?

- A. Create an Amazon EventBridge rule. Configure an AWS Systems Manager Automation runbook as a target of the EventBridge rule to initialize the disks after an EC2 instance launch event.
- B. Create an Amazon EventBridge rule. Configure an AWS Lambda function as a target of the EventBridge rule to initialize the drives after the AMI is launched.
- C. Create an AWS Config rule to automatically initialize the EBS volumes on Windows EC2 instances.
- D. Add the secondary volume configuration to the DriveLetterMappingConfig.json file. Configure the InitializeDisks.ps1 Windows PowerShell script to run at launch. Create a new AMI from the running EC2 instance.

Correct Answer: D

Section:

Explanation:

To automate the initialization of additional EBS volumes on Windows EC2 instances, the most effective approach is to integrate initialization scripts within the instance so that they execute upon startup:

Configure Initialization Script: Use a Windows PowerShell script (InitializeDisks.ps1) to initialize and format the additional EBS volumes. The script can assign drive letters based on configurations specified in DriveLetterMappingConfig.json.

Automate at Launch: Ensure that the PowerShell script runs automatically upon instance startup. This can be configured through Windows Task Scheduler or by setting it up in the startup folder.

Create a Custom AMI: Once the instance is configured with the script and successfully initializes the disks on startup, create a new AMI from this setup. This AMI can then be used to launch new instances that will automatically initialize their additional EBS volumes with no manual intervention required.

This method leverages native Windows tools and AWS capabilities to automate EBS volume initialization, enhancing operational efficiency without additional external dependencies.

QUESTION 165

A company has a cluster of Linux Amazon EC2 Spot Instances that read many files from and write many files to attached Amazon Elastic Block Store (Amazon EBS) volumes. The EC2 instances are frequently started and stopped. As part of the process when an EC2 instance starts, an EBS volume is restored from a snapshot.

EBS volumes that are restored from snapshots are experiencing initial performance that is lower than expected. The company's workload needs almost all the provisioned IOPS on the attached EBS volumes. The EC2 instances are unable to support the workload when the performance of the EBS volumes is too low. A SysOps administrator must implement a solution to ensure that the EBS volumes provide the expected performance when they are restored from snapshots.

Which solution will meet these requirements?

- A. Configure fast snapshot restore (FSR) on the snapshots that are used.
- B. Restore each snapshot onto an unencrypted EBS volume. Encrypt the EBS volume when the performance stabilizes.
- C. Format the EBS volumes as XFS file systems before restoring the snapshots.
- D. Increase the Linux read-ahead buffer to 1 MiB.

Correct Answer: A

Section:

Explanation:

For EBS volumes restored from snapshots to immediately achieve the required IOPS performance, Fast Snapshot Restore (FSR) can be utilized:

Enable FSR: Fast Snapshot Restore can be enabled on specific snapshots. This feature pre-warms the EBS volume created from a snapshot to its full performance level immediately after it is provisioned.

Operational Impact: By enabling FSR, any EBS volume created from these enabled snapshots will provide the provisioned IOPS performance right from the start, eliminating the performance lag that typically occurs as the data is lazily loaded from S3.

Cost Considerations: While FSR increases costs due to the pre-warming of data, it is justified by the need for immediate high performance, especially in environments where EBS volume responsiveness is critical to application

performance.

This solution directly addresses the challenge of initial performance degradation and ensures that the EBS volumes can handle the required workload immediately upon restoration from a snapshot.

QUESTION 166

A SysOps administrator launches an Amazon EC2 instance in a private subnet of a VPC. When the SysOps administrator attempts a curl command from the command line of the EC2 instance, the SysOps administrator cannot connect to <https://www.example.com>.

What should the SysOps administrator do to resolve this issue?

- A. Ensure that there is an outbound security group for port 443 to 0.0.0.0/0.
- B. Ensure that there is an inbound security group for port 443 from 0.0.0.0/0.
- C. Ensure that there is an outbound network ACL for ephemeral ports 1024-65535 to 0.0.0.0/0.
- D. Ensure that there is an outbound network ACL for port 80 to 0.0.0.0/0.

Correct Answer: A

Section:

Explanation:

To resolve the issue of the EC2 instance in a private subnet not being able to connect to external websites via HTTPS (port 443), it is necessary to adjust the security group settings:

Outbound Security Group Rules: Verify that the security group associated with the EC2 instance allows outbound traffic on port 443 to any destination (0.0.0.0/0). This rule is crucial because it enables the instance to initiate HTTPS connections to external websites.

Network ACLs: While the primary concern here is the security group, ensure also that the Network Access Control List (ACL) associated with the subnet permits outbound HTTPS traffic. However, the ACLs by default allow all outbound traffic unless specifically restricted.

Internet Connectivity: Since the instance is in a private subnet, ensure that it has a route to the internet through a NAT Gateway or NAT Instance located in a public subnet. Without this, the instance won't be able to reach external networks even if the security groups and ACLs are correctly configured.

By ensuring that the security group permits outbound HTTPS traffic, you address the most common configuration oversight that would prevent such connectivity.

QUESTION 167

A company has several business units that want to use Amazon EC2. The company wants to require all business units to provision their EC2 instances by using only approved EC2 instance configurations.

What should a SysOps administrator do to implement this requirement?

- A. Create an EC2 instance launch configuration. Allow the business units to launch EC2 instances by specifying this launch configuration in the AWS Management Console.
- B. Develop an IAM policy that limits the business units to provision EC2 instances only. Instruct the business units to launch instances by using an AWS CloudFormation template.
- C. Publish a product and launch constraint role for EC2 instances by using AWS Service Catalog. Allow the business units to perform actions in AWS Service Catalog only.
- D. Share an AWS CloudFormation template with the business units. Instruct the business units to pass a role to AWS CloudFormation to allow the service to manage EC2 instances.

Correct Answer: C

Section:

Explanation:

To enforce the use of approved EC2 instance configurations across different business units efficiently:

AWS Service Catalog: Utilize AWS Service Catalog to manage and govern commonly deployed IT services. Create a catalog of pre-approved products (in this case, EC2 instance configurations).

Publish Products: Define and publish EC2 instance configurations as products within the Service Catalog. These products will incorporate all the necessary and approved configurations, options, and software.

Launch Constraints: Assign launch constraints to these products, ensuring that users can only launch EC2 instances as defined by the pre-approved configurations.

Control Access: Grant business units access only to the Service Catalog for provisioning EC2 instances. This ensures they use only those configurations that comply with company policies and standards.

This approach not only standardizes resource deployment but also simplifies management and enhances compliance across the organization.

QUESTION 168

A company is supposed to receive a data file every hour in an Amazon S3 bucket. An S3 event notification invokes an AWS Lambda function each time a file arrives. The function processes the data for use by an application.

The application team notices that sometimes the file does not arrive. The application team wants to receive a notification whenever the file does not arrive.

What is the MOST operationally efficient solution that meets these requirements?

- A. Add an S3 Lifecycle rule on the S3 bucket with a scope that is limited to objects that were created in the last hour. Configure another S3 event notification to be invoked by the lifecycle transition when the number of objects transitioned is zero. Publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to notify the application team.
- B. Configure another S3 event notification to invoke a Lambda function that posts a message to an Amazon Simple Queue Service (Amazon SQS) queue. Create an Amazon CloudWatch alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to notify the application team when the ApproximateAgeOfOldestMessage metric of the queue is greater than 1 hour.
- C. Create an Amazon CloudWatch alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to alert the application team when the Invocations metric of the Lambda function is zero for an hour. Configure the alarm to treat missing data as breaching.
- D. Create a new Lambda function to get the timestamp of the newest file in the S3 bucket. If the timestamp is more than 1 hour ago, publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to notify the application team. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the new function hourly.

Correct Answer: C

Section:

QUESTION 169

A global company operates out of five AWS Regions. A SysOps administrator wants to identify all the company's tagged and untagged Amazon EC2 instances.

The company requires the output to display the instance ID and tags.

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

- A. Create a tag-based resource group in AWS Resource Groups.
- B. Use AWS Trusted Advisor. Export the EC2 On-Demand Instances check results from Trusted Advisor.
- C. Use Cost Explorer. Choose a service type of EC2-Instances, and group by Resource.
- D. Use Tag Editor in AWS Resource Groups. Select all Regions, and choose a resource type of AWS::EC2::Instance.

Correct Answer: D

Section:

Explanation:

To identify both tagged and untagged EC2 instances across multiple AWS Regions efficiently:

AWS Tag Editor: Tag Editor allows you to search for resources across your AWS account by tags, including both tagged and untagged resources.

Search Setup: In the Tag Editor, select all the Regions where the company operates. Specify the resource type as AWS::EC2::Instance to focus the search on EC2 instances.

View and Export Data: Execute the search to view all EC2 instances, along with their associated tags and instance IDs. This data can be exported for further analysis or reporting.

Using the Tag Editor is an operationally efficient way to quickly get a comprehensive view of resource tagging across multiple Regions, aiding in compliance and resource management tasks.

QUESTION 170

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

Which additional actions should the administrator take to control access? (Choose two.)

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

Correct Answer: A, E

Section:

Explanation:

To control access to Amazon EC2 instances using AWS Systems Manager Session Manager based on specific tags:

Attach an IAM Policy to Users or Groups: Create and attach an Identity and Access Management (IAM) policy to the IAM users or groups who need access to the EC2 instances. This policy should specify the permissions required to use Session Manager to start sessions with the instances.

Create an IAM Policy with Tag-Based Conditions: Create an IAM policy that includes a condition element to allow access to EC2 instances based on specific tags. This policy can be designed to grant the ssm:StartSession



permission only for instances that match certain tags, as defined in the condition block of the IAM policy. Here is a sample condition block that could be used:

```
'Condition': {  
'StringEquals': {  
'ec2:ResourceTag/YourTagName': 'YourTagValue'  
}  
}
```

This ensures that only authorized users can initiate sessions with instances that have the specified tags, enhancing security and operational management.

By implementing these policies, you ensure that only the appropriate personnel have the controlled access required, based on the specific business needs and security guidelines.

QUESTION 171

An application runs on Amazon EC2 instances in an Auto Scaling group. Following the deployment of a new feature on the EC2 instances, some instances were marked as unhealthy and then replaced by the Auto Scaling group. The EC2 instances terminated before a SysOps administrator could determine the cause of the health status changes. To troubleshoot this issue, the SysOps administrator wants to ensure that an AWS Lambda function is invoked in this situation.

How should the SysOps administrator meet these requirements?

- A. Activate the instance scale-in protection setting for the Auto Scaling group. Invoke the Lambda function through Amazon EventBridge (Amazon CloudWatch Events).
- B. Activate the instance scale-in protection setting for the Auto Scaling group. Invoke the Lambda function through Amazon Route 53.
- C. Add a lifecycle hook to the Auto Scaling group to invoke the Lambda function through Amazon EventBridge (Amazon CloudWatch Events).
- D. Add a lifecycle hook to the Auto Scaling group to invoke the Lambda function through Amazon Route 53.

Correct Answer: C

Section:

Explanation:

To enable troubleshooting of EC2 instances marked as unhealthy before they are terminated by the Auto Scaling group, you can use lifecycle hooks:

Add a Lifecycle Hook: Configure a lifecycle hook in the Auto Scaling group. This hook will hold the instance in a 'wait' state either when it launches or terminates (in this case, when it's about to be terminated due to health check failure).

Integration with Amazon EventBridge (CloudWatch Events): Set up the lifecycle hook to send an event to EventBridge (formerly CloudWatch Events) when an instance is in the termination lifecycle state.

Invoke Lambda Function: Configure EventBridge to trigger an AWS Lambda function when it receives the termination lifecycle event from the Auto Scaling group. This Lambda function can then perform necessary diagnostics, logging, or data capture activities on the instance before it's terminated.

This configuration allows the SysOps administrator to perform necessary investigations on why instances were marked unhealthy before they are automatically replaced, offering a chance to diagnose and potentially correct underlying issues.

QUESTION 172

A company hosts an internal application on Amazon EC2 On-Demand Instances behind an Application Load Balancer (ALB). The instances are in an Amazon EC2 Auto Scaling group. Employees use the application to provide product prices to potential customers. The Auto Scaling group is configured with a dynamic scaling policy and tracks average CPU utilization of the instances.

Employees have noticed that sometimes the application becomes slow or unresponsive. A SysOps administrator finds that some instances are experiencing a high CPU load. The Auto Scaling group cannot scale out because the company is reaching the EC2 instance service quota.

The SysOps administrator needs to implement a solution that provides a notification when the company reaches 70% or more of the EC2 instance service quota.

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

- A. Create an AWS Lambda function that lists the EC2 instances, counts the EC2 instances, and compares the total number against the applied quota value by using the Service Quotas API. Configure the Lambda function to publish an Amazon Simple Notification Service (Amazon SNS) notification if the quota utilization is equal to or greater than 70%. Create an Amazon EventBridge rule to invoke the Lambda function.
- B. Create an AWS Lambda function that lists the EC2 instances, counts the EC2 instances, and compares the total number against the applied quota value by using the Amazon CloudWatch Metrics API. Configure the Lambda function to publish an Amazon Simple Notification Service (Amazon SNS) notification if the quota utilization is equal to or greater than 70%. Create an Amazon EventBridge rule to invoke the Lambda function.
- C. Use the Service Quotas console to create an Amazon CloudWatch alarm for the EC2 instances. Configure the alarm with quota utilization equal to or greater than 70%. Configure the alarm to publish an Amazon Simple Notification Service (Amazon SNS) notification when the alarm enters ALARM state.
- D. Create an Amazon CloudWatch alarm. Configure the alarm with a threshold of 70% for the CPUUtilization metric for the EC2 instances. Configure the alarm to publish an Amazon Simple Notification Service (Amazon SNS) notification when the alarm enters ALARM state.

Correct Answer: C

Section:

Explanation:

To monitor and receive alerts when the EC2 instance service quota usage reaches 70% or more:

Service Quotas Console: Navigate to the Service Quotas console within AWS and identify the specific quota for EC2 instances.

Create a CloudWatch Alarm: Directly from the Service Quotas console, set up a CloudWatch alarm for the EC2 instance quota metric. Configure the alarm to trigger when the quota utilization reaches or exceeds 70%.

Notification Setup: Link this alarm to an Amazon SNS topic that will send a notification to relevant stakeholders or systems when the quota usage threshold is breached.

This method provides an automated, straightforward way to monitor resource limits and ensures that stakeholders are promptly notified, enabling them to take proactive measures to manage the quota and prevent service disruption.

QUESTION 173

A company wants to track its expenditures for Amazon EC2 and Amazon RDS within AWS. The company decides to implement more rigorous tagging requirements for resources in its AWS accounts. A SysOps administrator needs to identify all noncompliant resources.

What is the MOST operationally efficient solution that meets these requirements?

- A. Create a rule in Amazon EventBridge (Amazon CloudWatch Events) that invokes a custom AWS Lambda function that will evaluate all created or updated resources for the specified tags.
- B. Create a rule in AWS Config that invokes a custom AWS Lambda function that will evaluate all resources for the specified tags.
- C. Create a rule in AWS Config with the required-tags managed rule to evaluate all resources for the specified tags.
- D. Create a rule in Amazon EventBridge (Amazon CloudWatch Events) with a managed rule to evaluate all created or updated resources for the specified tags.

Correct Answer: C

Section:

Explanation:

To efficiently monitor and identify noncompliant resources in terms of tagging within AWS, using AWS Config with a managed rule for required tagging is most appropriate:

AWS Config Setup: Configure AWS Config to monitor and record configurations of AWS resources within your environment.

Managed Rule for Required Tags: Utilize the 'required-tags' managed rule in AWS Config, which checks whether your resources have the specific tags you define as mandatory. This rule can be customized to specify which tags are required and can automatically evaluate all existing and new resources in your environment.

Compliance Reporting: AWS Config provides detailed compliance reporting that helps you identify resources that do not meet the tagging requirements, facilitating easy remediation.

This approach leverages AWS Config's capabilities for continuous monitoring and evaluation without needing to write custom code or manage additional services, providing an operationally efficient solution for compliance management.

QUESTION 174

A user is connected to an Amazon EC2 instance in a private subnet. The user is unable to access the internet from the instance by using the following curl command: curl http://www.example.com.

A SysOps administrator reviews the VPC configuration and learns the following information:

- * The private subnet has a route to a NAT gateway for CIDR 0.0.0.0/0
- * The outbound security group for the EC2 instance contains one rule: outbound for port 443 to CIDR 0.0.0.0/0
- * The inbound security group for the EC2 instance allows ports 22 and 443 from the user's IP address.
- * The inbound network ACL for the subnet allows port 22 and port range 1024-65535 from CIDR 0.0.0.0/0

Which action will allow the user to complete the curl request successfully?

- A. Add an additional inbound network ACL rule for port 80 to CIDR 0.0.0.0/0.
- B. Add an additional inbound security group rule for port 80 to CIDR 0.0.0.0/0.
- C. Add an additional outbound security group rule for port 80 to CIDR 0.0.0.0/0.
- D. Add an additional outbound security group rule for port 80 to the user's IP address.

Correct Answer: C

Section:

Explanation:

Since the EC2 instance is attempting to access the internet using HTTP (port 80) but is configured only to allow HTTPS (port 443) traffic, the security group needs adjustment:

Security Group Configuration: The outbound rules of the security group associated with the EC2 instance must allow traffic over HTTP. Add an outbound rule that enables port 80 to destination 0.0.0.0/0. This rule will allow the instance to send HTTP requests to any IP address on the internet.

Test Connectivity: After updating the security group, test the connectivity using the curl command again to ensure the configuration allows internet access via HTTP.

This change is necessary because the existing security group configuration does not permit outbound HTTP traffic, which is essential for accessing websites using HTTP.

QUESTION 175

A SysOps administrator needs to configure the Amazon Route 53 hosted zone for example.com and www.example.com to point to an Application Load Balancer (ALB). Which combination of actions should the SysOps administrator take to meet these requirements? (Select TWO.)

- A. Configure an A record for example.com to point to the IP address of the ALB.
- B. Configure an A record for www.example.com to point to the IP address of the ALB.
- C. Configure an alias record for example.com to point to the CNAME of the ALB.
- D. Configure an alias record for www.example.com to point to the Route 53 example.com record.
- E. Configure a CNAME record for example.com to point to the CNAME of the ALB.

Correct Answer: C, D

Section:

Explanation:

You are correct that an A record typically points to an IP address. However, in the case of an Application Load Balancer (ALB), you cannot use an A record with an IP address because the IP addresses of an ALB can change over time. Instead, you can use an alias record to point to the DNS name of the ALB. An alias record is a Route 53 extension to DNS that allows you to route traffic to selected AWS resources, such as an ALB, by using a friendly DNS name, such as example.com, instead of the resource's IP address or DNS name.

QUESTION 176

A SysOps administrator deployed a three-tier web application to a OA environment and is now evaluating the high availability of the application. The SysOps administrator notices that, when they simulate an unavailable Availability Zone, the application fails to respond. The application stores data in Amazon RDS and Amazon DynamoDB. How should the SysOps administrator resolve this issue?

- A. Add additional subnets to the RDS instance subnet group.
- B. Add an Elastic Load Balancer in front of the RDS instance.
- C. Distribute the data in DynamoDB across Availability Zones.
- D. Enable Multi-AZ for the RDS instance.

Correct Answer: D

Section:

Explanation:

To improve the high availability of an application that utilizes Amazon RDS and experiences failure when an Availability Zone becomes unavailable:

Multi-AZ Deployment for RDS: Enable Multi-AZ deployments for your Amazon RDS instance. This setting ensures that RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone.

Automatic Failover: In the event of a primary RDS instance failure, RDS will automatically failover to the standby so that database operations can resume quickly with minimal disruption.

High Availability Configuration: This configuration not only enhances the robustness of the database component but also ensures that the application remains operational even if one Availability Zone is experiencing issues.

Enabling Multi-AZ for RDS is crucial for maintaining high availability and ensuring that the application remains resilient in the face of AZ disruptions.

QUESTION 177

A company wants to reduce costs for jobs that can be completed at any time. The jobs currently run by using multiple Amazon EC2 On-Demand Instances, and the jobs take slightly less than 2 hours to complete. If a job fails for any reason, it must be restarted from the beginning.

Which solution will meet these requirements MOST cost-effectively?

- A. Purchase Reserved Instances for the jobs.
- B. Submit a request for a one-time Spot Instance for the jobs.

- C. Submit a request for Spot Instances with a defined duration for the jobs.
- D. Use a mixture of On-Demand Instances and Spot Instances for the jobs.

Correct Answer: C

Section:

Explanation:

To reduce costs effectively for jobs that are flexible in their scheduling and have a clear, predictable runtime:

Spot Instances with Defined Duration (Spot Blocks): Spot Instances offer significant discounts compared to On-Demand pricing. For workloads like the described jobs that have a predictable duration (slightly less than 2 hours), requesting Spot Instances with a defined duration (also known as Spot Blocks) is ideal. This option allows you to request Spot Instances guaranteed to not be terminated by AWS due to price changes for the duration specified.

Cost Efficiency: This method ensures that the instances will run for the duration required to complete the jobs without interruption, unless AWS experiences an exceptional shortage of capacity. The cost savings compared to On-Demand Instances can be substantial, especially for regular, predictable workloads.

Risk Mitigation: Although Spot Instances can be interrupted, using them with a defined duration reduces the risk of interruptions within the set time frame, making them suitable for jobs that can tolerate a restart in rare cases of interruption after the block time expires.

This strategy combines cost savings with the performance requirements of the jobs, making it an optimal choice for tasks that are not time-critical but need completion within a predictable timeframe.

QUESTION 178

A company runs a worker process on three Amazon EC2 instances. The instances are in an Auto Scaling group that is configured to use a simple scaling policy. The instances process messages from an Amazon Simple Queue Service (Amazon SQS) queue.

Random periods of increased messages are causing a decrease in the performance of the worker process. A SysOps administrator must scale the instances to accommodate the increased number of messages.

Which solution will meet these requirements?

- A. Use CloudWatch to create a metric math expression to calculate the approximate age of the oldest message in the SQS queue. Create a target tracking scaling policy for the metric math expression to modify the Auto Scaling group.
- B. Use CloudWatch to create a metric math expression to calculate the approximate number of messages visible in the SQS queue for each instance. Create a target tracking scaling policy for the metric math expression to modify the Auto Scaling group.
- C. Create an Application Load Balancer (ALB). Attach the ALB to the Auto Scaling group. Create a target tracking scaling policy for the ALBRequestCountPerTarget metric to modify the Auto Scaling group.
- D. Create an Application Load Balancer (ALB). Attach the ALB to the Auto Scaling group. Create a scheduled scaling policy for the Auto Scaling group.

Correct Answer: B

Section:

Explanation:

To manage scaling of EC2 instances in response to variable SQS message loads effectively:

Monitor SQS Queue Size: Utilize Amazon CloudWatch to monitor the number of visible messages in the SQS queue. This metric gives an indication of the workload that needs to be processed by the worker instances.

Metric Math Expression: Create a CloudWatch metric math expression that calculates the approximate number of messages visible per instance. This provides a more precise scaling metric, ensuring that each instance in the Auto Scaling group has a manageable load.

Target Tracking Scaling Policy: Implement a target tracking scaling policy based on this metric math expression. Configure the Auto Scaling group to automatically adjust its size to maintain a target value for the average number of SQS messages per instance. This approach ensures that the EC2 instances scale up during high traffic periods and scale down when the message load decreases.

This solution optimizes resource utilization and cost while maintaining performance by ensuring that the worker processes are neither overwhelmed nor idle.

QUESTION 179

A company's security policy states that connecting to Amazon EC2 instances is not permitted through SSH and RDP. If access is required, authorized staff can connect to instances by using AWS Systems Manager Session Manager.

Users report that they are unable to connect to one specific Amazon EC2 instance that is running Ubuntu and has AWS Systems Manager Agent (SSM Agent) pre-installed. These users are able to use Session Manager to connect to other instances in the same subnet, and they are in a 1AM group that has Session Manager permission for all instances.

What should a SysOps administrator do to resolve this issue?

- A. Add an inbound rule for port 22 in the security group associated with the Ubuntu instance.
- B. Assign the AmazonSSMManagedInstanceCore managed policy to the EC2 instance profile for the Ubuntu instance.

- C. Configure the SSM Agent to log in with a user name of 'ubuntu'.
- D. Generate a new key pair, configure Session Manager to use this new key pair, and provide the private key to the users.

Correct Answer: B

Section:

Explanation:

If users are unable to connect to a specific Ubuntu EC2 instance using AWS Systems Manager Session Manager while other instances are accessible, the issue is likely due to IAM permissions:

Instance Profile Permissions: Ensure that the EC2 instance has the necessary IAM permissions to interact with Systems Manager. The AmazonSSMManagedInstanceCore managed policy includes permissions required for the SSM Agent on the instance to communicate with the AWS Systems Manager service.

Attach Managed Policy: Attach the AmazonSSMManagedInstanceCore policy to the IAM role that is associated with the Ubuntu instance's instance profile. This step is crucial as it authorizes the instance to use Systems Manager services, including Session Manager.

Verify Configuration and Connectivity: After updating the instance profile, verify that users can connect via Session Manager. This solution does not require any changes to network security settings like security groups.

By ensuring that the instance has the appropriate IAM permissions, you resolve issues related to access control and Systems Manager functionality, allowing authorized personnel to connect securely without using SSH or RDP.

QUESTION 180

A fleet of servers must send local logs to Amazon CloudWatch. How should the servers be configured to meet this requirement?

- A. Configure AWS Config to forward events to CloudWatch.
- B. Configure a Simple Network Management Protocol (SNMP) agent to forward events to CloudWatch.
- C. Install and configure the unified CloudWatch agent.
- D. Install and configure the Amazon Inspector agent.

Correct Answer: C

Section:

Explanation:

To send local logs from a fleet of servers to Amazon CloudWatch:

Install the Unified CloudWatch Agent: The unified CloudWatch agent is capable of collecting both logs and metrics from servers. This agent supports various operating systems and can be configured according to specific logging requirements.

Configuration of the Agent: The agent's configuration involves specifying which log files to monitor and how they should be processed. This configuration can be done manually or through the AWS Systems Manager for automated deployment across multiple servers.

Send Logs to CloudWatch: Once configured and running, the CloudWatch agent will continuously monitor the specified log files and send the log data to Amazon CloudWatch Logs, allowing you to view, search, and set alarms on log data.

This setup ensures a robust and scalable way to manage log data across a fleet of servers, leveraging AWS native services for seamless integration and management.

QUESTION 181

A company has 50 AWS accounts and wants to create an identical Amazon VPC in each account. Any changes the company makes to the VPCs in the future must be implemented on every VPC.

What is the MOST operationally efficient method to deploy and update the VPCs in each account?

- A. Create an AWS Cloud Formation template that defines the VPC. Sign in to the AWS Management Console under each account. Create a stack from the template.
- B. Create a shell script that configures the VPC using the AWS CLI. Provide a list of accounts to the shell script from a text file. Create the VPC in every account in the list.
- C. Create an AWS Lambda function that configures the VPC. Store the account information in Amazon DynamoDB. Grant Lambda access to the DynamoDB table. Create the VPC in every account in the list.
- D. Create an AWS Cloud Formation template that defines the VPC. Create an AWS CloudFormation StackSet based on the template. Deploy the template to all accounts using the stack set.

Correct Answer: D

Section:

Explanation:

To deploy and manage an identical Amazon VPC configuration across multiple AWS accounts efficiently:



AWS CloudFormation Template: Create a CloudFormation template that defines the VPC configuration. This template should include all necessary resources like subnets, route tables, internet gateways, etc.
Use CloudFormation StackSets: Utilize AWS CloudFormation StackSets to manage the deployment of the VPC template across the 50 AWS accounts. StackSets allow you to specify management and target accounts, automate deployments, and ensure consistency across all accounts.
Updating VPCs: When updates are required, modify the CloudFormation template and update the stack set. This will automatically apply the changes to all VPCs in the target accounts, ensuring uniformity and reducing operational overhead.
This method provides a centralized, consistent, and scalable way to manage resources across multiple AWS accounts, greatly simplifying the administration and ensuring compliance with organizational standards.

QUESTION 182

A company runs a web application on three Amazon EC2 instances behind an Application Load Balancer (ALB). Web traffic increases significantly during the same 9-hour period every day and causes a decrease in the application's performance. A SysOps administrator must scale the application ahead of the changes in demand to accommodate the increased traffic. Which solution will meet these requirements?

- A. Create an Amazon CloudWatch alarm to monitor application latency. Configure an alarm action to increase the size of each EC2 instance if the latency threshold is reached.
- B. Create an Amazon EventBridge rule to monitor application latency. Configure the rule to add an EC2 instance to the ALB if the latency threshold is reached
- C. Deploy the application to an EC2 Auto Scaling group that uses a target tracking scaling policy. Attach the ALB to the Auto Scaling group.
- D. Deploy the application to an EC2 Auto Scaling group that uses a scheduled scaling policy. Attach the ALB to the Auto Scaling group.

Correct Answer: D

Section:

Explanation:

For predictable, significant traffic increases during a specific time period every day:

EC2 Auto Scaling Group: Set up an Auto Scaling group for the EC2 instances running the web application. This group automatically adjusts the number of instances based on policies defined.

Scheduled Scaling Policy: Use a scheduled scaling policy to pre-emptively increase the number of instances before the expected increase in traffic each day. Scheduled scaling allows you to specify the scaling actions to occur at specific times, based on known or expected demand patterns.

Attach to ALB: Ensure the Auto Scaling group is attached to the Application Load Balancer, which will distribute incoming traffic across the dynamically adjusted pool of EC2 instances. This approach ensures that the application scales up resources ahead of the expected load, maintaining performance and user experience without manual intervention.

QUESTION 183

A company hosts a production MySQL database on an Amazon Aurora single-node DB cluster. The database is queried heavily for reporting purposes. The DB cluster is experiencing periods of performance degradation because of high CPU utilization and maximum connections errors. A SysOps administrator needs to improve the stability of the database. Which solution will meet these requirements?

- A. Create an Aurora Replica node. Create an Auto Scaling policy to scale replicas based on CPU utilization. Ensure that all reporting requests use the read-only connection string.
- B. Create a second Aurora MySQL single-node DB cluster in a second Availability Zone. Ensure that all reporting requests use the connection string for this additional node.
- C. Create an AWS Lambda function that caches reporting requests. Ensure that all reporting requests call the Lambda function.
- D. Create a multi-node Amazon ElastiCache cluster. Ensure that all reporting requests use the ElastiCache cluster. Use the database if the data is not in the cache.

Correct Answer: A

Section:

Explanation:

To alleviate performance degradation on a heavily queried Amazon Aurora DB cluster:

A: Create an Aurora Replica node and implement an Auto Scaling policy based on CPU utilization. Ensure all reporting requests use the read-only connection string to redirect read queries to the replica. This setup alleviates the load on the primary DB instance by balancing read traffic, which can significantly improve stability during periods of high demand. Aurora Replicas are ideal for scaling read operations and can improve the performance of the primary instance by offloading read requests. More details on Aurora Replicas and their benefits can be found in the AWS documentation on Aurora Replicas Amazon Aurora Replicas.

QUESTION 184

A SysOps administrator is responsible for managing a fleet of Amazon EC2 instances. These EC2 instances upload build artifacts to a third-party service. The third-party service recently implemented a strict IP allow list that requires all build uploads to come from a single IP address. What change should the systems administrator make to the existing build fleet to comply with this new requirement?

- A. Move all of the EC2 instances behind a NAT gateway and provide the gateway IP address to the service.
- B. Move all of the EC2 instances behind an internet gateway and provide the gateway IP address to the service.
- C. Move all of the EC2 instances into a single Availability Zone and provide the Availability Zone IP address to the service.
- D. Move all of the EC2 instances to a peered VPC and provide the VPC IP address to the service.

Correct Answer: A

Section:

Explanation:

To ensure all EC2 instances upload build artifacts through a single IP address:

A: Move all of the EC2 instances behind a NAT gateway. Provide the IP address of the NAT gateway to the third-party service for the allow list. A NAT gateway enables instances in a private subnet to connect to services outside AWS (such as a third-party service) but prevents the internet from initiating connections with those instances. Using a NAT gateway standardizes all outgoing traffic to use a single IP address. More information on NAT gateways can be found in AWS documentation NAT Gateways.

QUESTION 185

A SysOps administrator needs to monitor a process that runs on Linux Amazon EC2 instances. If the process stops, the process must restart automatically. The Amazon CloudWatch agent is already installed on all the EC2 Instances.

Which solution will meet these requirements?

- A. Add a procstat monitoring configuration to the CloudWatch agent for the process. Create an Amazon EventBridge event rule that initiates an AWS Systems Manager Automation runbook to restart the process after the process stops.
- B. Add a StatsD monitoring configuration to the CloudWatch agent for the process. Create a CloudWatch alarm that initiates an AWS Systems Manager Automation runbook to restart the process after the process stops.
- C. Add a StatsD monitoring configuration to the CloudWatch agent for the process. Create an Amazon EventBridge event rule that initiates an AWS Systems Manager Automation runbook to restart the process after the process stops.
- D. Add a procstat monitoring configuration to the CloudWatch agent for the process. Create a CloudWatch alarm that initiates an AWS Systems Manager Automation runbook to restart the process after the process stops.

Correct Answer: A

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch-Agent-procstat-process-metrics.html>

QUESTION 186

A SysOps administrator is preparing to deploy an application to Amazon EC2 instances that are in an Auto Scaling group. The application requires dependencies to be installed. Application updates are Issued weekly. The SysOps administrator needs to implement a solution to incorporate the application updates on a regular basis. The solution also must conduct a vulnerability scan during Amazon Machine Image (AMI) creation. What is the MOST operationally efficient solution that meets these requirements?

- A. Create a script that uses Packer. Schedule a cron job to run the script.
- B. Install the application and its dependencies on an EC2 instance. Create an AMI of the H2 instance.
- C. Use EC2 Image Builder with a custom recipe to install the application and its dependencies.
- D. Invoke the EC2 CreateImage API operation by using an Amazon EventBridge scheduled rule.

Correct Answer: C

Section:

Explanation:

To efficiently manage application deployments and updates on Amazon EC2 instances within an Auto Scaling group, along with ensuring security through vulnerability scans:

EC2 Image Builder: This AWS service automates the creation, management, and deployment of customized, secure, and up-to-date 'golden' server images. By using EC2 Image Builder, you can automate the installation of

software, patches, and security configurations.

Custom Recipes: Define a custom recipe in EC2 Image Builder that includes steps to install the application and its dependencies. Additionally, configure the recipe to perform vulnerability scans as part of the image creation process.

Automated Pipeline: Set up an Image Builder pipeline that triggers on a regular schedule (e.g., weekly) to incorporate the latest application updates and security patches into the AMI. The new AMIs can then be automatically used by the Auto Scaling group to launch updated and secure instances.

This solution not only streamlines the management of application deployments and updates but also ensures that all instances launched by the Auto Scaling group meet the latest security and compliance standards, minimizing operational overhead and enhancing security.

QUESTION 187

A company has a public web application that experiences rapid traffic increases after advertisements appear on local television. The application runs on Amazon EC2 instances that are in an Auto Scaling group. The Auto Scaling group is not keeping up with the traffic surges after an advertisement runs. The company often needs to scale out to 100 EC2 instances during the traffic surges.

The instance startup times are lengthy because of a boot process that creates machine-specific data caches that are unique to each instance. The exact timing of when the advertisements will appear on television is not known. A SysOps administrator must implement a solution so that the application can function properly during the traffic surges.

Which solution will meet these requirements?

- A. Create a warm pool. Keep enough instances in the Stopped state to meet the increased demand.
- B. Start 100 instances. Allow the boot process to finish running. Store this data on the instance store volume before stopping the instances.
- C. Increase the value of the instance warmup time in the scaling policy.
- D. Use predictive scaling for the Auto Scaling group.

Correct Answer: A

Section:

Explanation:

To address the issue of slow startup times during unexpected traffic surges, a warm pool for the Auto Scaling group is an effective solution:

Warm Pool Concept: A warm pool allows you to maintain a set of pre-initialized or partially initialized EC2 instances that are not actively serving traffic but can be quickly brought online when needed.

Management of Instances: Instances in the warm pool can be kept in a stopped state and then started much more quickly than launching new instances, as the machine-specific data caches are already created.

Scalability and Responsiveness: During a surge in traffic, especially unpredictable ones like those triggered by advertisements, instances from the warm pool can be rapidly activated to handle the increased load, ensuring that the application remains responsive without the typical delays associated with boot processes.

This method significantly reduces the time to scale out by utilizing pre-warmed instances, enhancing the application's ability to cope with sudden and substantial increases in traffic.

QUESTION 188

A company is running production workloads that use a Multi-AZ deployment of an Amazon RDS for MySQL db.m6g.xlarge (general purpose) standard DB instance. Users report that they are frequently encountering a 'too many connections' error. A SysOps administrator observes that the number of connections on the database is high.

The SysOps administrator needs to resolve this issue while keeping code changes to a minimum.

Which solution will meet these requirements MOST cost-effectively?

- A. Modify the RDS for MySQL DB instance to a larger instance size.
- B. Migrate the RDS for MySQL DB instance to Amazon DynamoDB.
- C. Configure RDS Proxy. Modify the application configuration file to use the RDS Proxy endpoint.
- D. Modify the RDS for MySQL DB instance to a memory optimized DB instance.

Correct Answer: C

Section:

Explanation:

For the issue of 'too many connections' on a MySQL database, using RDS Proxy offers a streamlined solution:

RDS Proxy Setup: RDS Proxy sits between your application and the database. It pools and efficiently manages database connections, which reduces the number of direct connections to the database.

Connection Management: By handling connection pooling, RDS Proxy can help mitigate issues related to connection overhead and limits, such as the 'too many connections' error, by allowing the database to serve more requests from a smaller and more stable number of connections.

Minimal Code Changes: Integrating RDS Proxy requires changes only to the database connection settings in the application's configuration files to point to the RDS Proxy endpoint instead of directly to the database. This

minimizes the amount of code change needed and leverages RDS Proxy to handle connection scaling and management more efficiently. This approach enhances database performance and scalability by efficiently managing connections without the need for significant application changes or database resizing.

QUESTION 189

A development team created and deployed a new AWS Lambda function 15 minutes ago. Although the function was invoked many times, Amazon CloudWatch Logs are not showing any log messages. What is one cause of this?

- A. The developers did not enable log messages for this Lambda function.
- B. The Lambda function's role does not include permissions to create CloudWatch Logs items.
- C. The Lambda function raises an exception before the first log statement has been reached.
- D. The Lambda function creates local log files that have to be shipped to CloudWatch Logs first before becoming visible.

Correct Answer: B

Section:

Explanation:

If AWS Lambda function logs are not appearing in Amazon CloudWatch, it is typically due to insufficient permissions:

IAM Role Permissions: The execution role assigned to the Lambda function must have the necessary permissions to interact with CloudWatch Logs. This includes actions like `logs:CreateLogGroup`, `logs:CreateLogStream`, and `logs:PutLogEvents`.

Check and Update Role: Verify that the IAM role used by the Lambda function includes a policy granting these permissions. If not, update the role to include these permissions.

Log Group and Stream: With the appropriate permissions, the Lambda function will be able to create or use a log group and stream in CloudWatch Logs and publish log messages accordingly.

Ensuring the Lambda function has the correct permissions is essential for diagnostics and monitoring, allowing log data to be captured and reviewed in CloudWatch Logs.

QUESTION 190

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

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

- A. Designate a central security account as the AWS Security Hub administrator account. Create a script that sends an invitation from the Security Hub administrator account and accepts the invitation from the member account. Run the script every time a new account is created. Configure Security Hub to run the CIS AWS Foundations Benchmark scans.
- B. Run the CIS AWS Foundations Benchmark across all accounts by using Amazon Inspector.
- C. Designate a central security account as the Amazon GuardDuty administrator account. Create a script that sends an invitation from the GuardDuty administrator account and accepts the invitation from the member account. Run the script every time a new account is created. Configure GuardDuty to run the CIS AWS Foundations Benchmark scans.
- D. Designate an AWS Security Hub administrator account. Configure new accounts in the organization to automatically become member accounts. Enable CIS AWS Foundations Benchmark scans.

Correct Answer: D

Section:

Explanation:

To ensure comprehensive and automated security scanning across multiple AWS accounts:

Security Hub Administrator Account: Designate one account within AWS Organizations as the Security Hub administrator account. This centralizes security findings management.

Automate Account Association: Configure Security Hub to automatically associate new accounts in the organization as member accounts. This ensures all new and existing accounts are continuously monitored under the same security policies.

Enable CIS Benchmark Scans: Within Security Hub, enable the CIS AWS Foundations Benchmark standard. This automatically scans all member accounts against this set of security best practices and compliance standards. This configuration provides an operationally efficient and scalable way to manage security and compliance across an extensive AWS environment, leveraging the native integration of AWS services.

QUESTION 191

A SysOps administrator needs to configure an Amazon S3 bucket to host a web application. The SysOps administrator has created the S3 bucket and has copied the static files for the web application to the S3 bucket. The company has a policy that all S3 buckets must not be public.

What should the SysOps administrator do to meet these requirements?

- A. Create an Amazon CloudFront distribution. Configure the S3 bucket as an origin with an origin access identity (OAI). Give the OAI the s3:GetObject permission in the S3 bucket policy.
- B. Configure static website hosting in the S3 bucket. Use Amazon Route 53 to create a DNS CNAME to point to the S3 website endpoint.
- C. Create an Application Load Balancer (ALB). Change the protocol to HTTPS in the ALB listener configuration. Forward the traffic to the S3 bucket.
- D. Create an accelerator in AWS Global Accelerator. Set up a listener configuration for port 443. Set the endpoint type to forward the traffic to the S3 bucket.

Correct Answer: A

Section:

Explanation:

To host a web application in an S3 bucket while adhering to the policy that prohibits public S3 buckets:

Amazon CloudFront: Set up a CloudFront distribution and designate the S3 bucket as its origin. This allows the web application to be served via CloudFront, which can handle web traffic at scale and provide additional features such as HTTPS delivery.

Origin Access Identity (OAI): Create an OAI for the CloudFront distribution and configure the S3 bucket policy to grant the s3:GetObject permission to the OAI. This allows only CloudFront to access the content in the S3 bucket, keeping the bucket private from direct public access.

Security and Performance: This configuration ensures that the web application is only accessible through CloudFront, enhancing security and performance. It also complies with the company's policy against public S3 buckets by controlling access strictly through CloudFront.

This method leverages CloudFront's capabilities to securely serve web applications from S3, maintaining privacy and compliance with organizational policies.

QUESTION 192

A company recently deployed an application in production. The production environment currently runs on a single Amazon EC2 instance that hosts the application's web application and a MariaDB database. Company policy states that all IT production environments must be highly available.

What should a SysOps administrator do to meet this requirement?

- A. Migrate the database from the EC2 instance to an Amazon RDS for MariaDB Multi-AZ DB instance. Run the application on EC2 instances that are in an Auto Scaling group that extends across multiple Availability Zones. Place the EC2 instances behind a load balancer.
- B. Migrate the database from the EC2 instance to an Amazon RDS for MariaDB Multi-AZ DB instance. Use AWS Application Migration Service to convert the application into an AWS Lambda function. Specify the Multi-AZ option for the Lambda function.
- C. Copy the database to a different EC2 instance in a different Availability Zone. Use AWS Backup to create Amazon Machine Images (AMIs) of the application EC2 instance and the database EC2 instance. Create an AWS Lambda function that performs health checks every minute. In case of failure, configure the Lambda function to launch a new EC2 instance from the AMIs that AWS Backup created.
- D. Migrate the database to a different EC2 instance. Place the application EC2 instance in an Auto Scaling group that extends across multiple Availability Zones. Create an Amazon Machine Image (AMI) from the database EC2 instance. Use the AMI to launch a second database EC2 instance in a different Availability Zone. Put the second database EC2 instance in the stopped state. Use the second database EC2 instance as a standby.

Correct Answer: A

Section:

Explanation:

To make the production environment highly available in accordance with company policy:

Database Migration: Move the MariaDB database from a single EC2 instance to Amazon RDS for MariaDB configured for Multi-AZ. This setup ensures high availability of the database with synchronous replication to a standby instance in a different Availability Zone.

Application Scalability: Deploy the application on EC2 instances within an Auto Scaling group. Configure the Auto Scaling group to operate across multiple Availability Zones to ensure that the application remains available even if one zone becomes unavailable.

Load Balancing: Place the EC2 instances behind an Elastic Load Balancer (ELB). The load balancer will distribute incoming application traffic across the multiple, geographically dispersed EC2 instances, further enhancing the availability and fault tolerance of the application.

This solution leverages AWS managed services to increase the reliability and availability of both the application and database layers, adhering to best practices for deploying critical production environments on AWS.

QUESTION 193

A SysOps administrator maintains the security and compliance of a company's AWS account. To ensure the company's Amazon EC2 instances are following company policy, a SysOps administrator wants to terminate any EC2 instance that do not contain a department tag. Noncompliant resources must be terminated in near real time.

Which solution will meet these requirements?

- A. Create an AWS Config rule with the required-tags managed rule to identify noncompliant resources. Configure automatic remediation to run the AWS-TerminateEC2Instance automation runbook to terminate noncompliant resources.
- B. Create a new Amazon EventBridge rule to monitor when new EC2 instances are created. Send the event to an Simple Notification Service (Amazon SNS) topic for automatic remediation.
- C. Ensure all users who can create EC2 instances also have the permissions to use the ec2:CreateTags and ec2:DescribeTags actions. Change the instance's shutdown behavior to terminate.
- D. Ensure AWS Systems Manager Compliance is configured to manage the EC2 instances. Call the AWS-StopEC2Instances automation runbook to stop noncompliant resources.

Correct Answer: A

Section:

Explanation:

To enforce compliance with tagging policies in real-time:

AWS Config Setup: Implement an AWS Config rule to continuously monitor and evaluate EC2 instances for compliance with the tagging requirements. The required-tags managed rule can be configured to specifically check for the presence of a 'department' tag.

Automatic Remediation: Configure AWS Config to automatically execute the AWS-TerminateEC2Instance Systems Manager Automation document as a remediation action. This runbook will terminate any EC2 instance identified as noncompliant due to missing required tags.

Operational Efficiency: This setup allows for the enforcement of company tagging policies automatically and in near real-time, reducing the manual overhead of monitoring and ensuring compliance.

This method provides an efficient and effective solution to ensure that all EC2 instances meet the company's tagging requirements and that any noncompliant instances are dealt with promptly.

QUESTION 194

A company has deployed an application on AWS. The application runs on a fleet of Linux Amazon EC2 instances that are in an Auto Scaling group. The Auto Scaling group is configured to use launch templates. The launch templates launch Amazon Elastic Block Store (Amazon EBS) backed EC2 instances that use General Purpose SSD (gp3) EBS volumes for primary storage.

A SysOps administrator needs to implement a solution to ensure that all the EC2 instances can share the same underlying files. The solution also must ensure that the data is consistent.

Which solution will meet these requirements?

- A. Create an Amazon Elastic File System (Amazon EFS) file system. Create a new launch template version that includes user data that mounts the EFS file system. Update the Auto Scaling group to use the new launch template version to cycle in newer EC2 instances and to terminate the older EC2 instances.
- B. Enable Multi-Attach on the EBS volumes. Create a new launch template version that includes user data that mounts the EBS volume. Update the Auto Scaling group to use the new template version to cycle in newer EC2 instances and to terminate the older EC2 instances.
- C. Create a cron job that synchronizes the data between the EBS volumes for all the EC2 instances in the Auto Scaling group. Create a lifecycle hook during instance launch to configure the cron job on all the EC2 instances. Rotate out the older EC2 instances.
- D. Create a new launch template version that creates an Amazon Elastic File System (Amazon EFS) file system. Update the Auto Scaling group to use the new template version to cycle in newer EC2 instances and to terminate the older EC2 instances.

Correct Answer: A

Section:

Explanation:

The requirement to share the same underlying files among EC2 instances with data consistency is best met by using Amazon Elastic File System (EFS), which supports concurrent access from multiple EC2 instances. A new launch template version should include user data scripts that mount the EFS file system on each instance launched by the Auto Scaling group. Older instances can be cycled out to ensure all instances use the new configuration. Option A is correct and provides the necessary solution while ensuring data consistency and availability. For implementation guidance, refer to the AWS documentation on integrating EFS with EC2 Amazon EFS Integration with EC2.

QUESTION 195

A SysOps administrator is re-architecting an application. The SysOps administrator has moved the database from a public subnet, where the database used a public endpoint, into a private subnet to restrict access from the public network. After this change, an AWS Lambda function that requires read access to the database cannot connect to the database. The SysOps administrator must resolve this issue without compromising security.

Which solution meets these requirements?

- A. Create an AWS PrivateLink interface endpoint for the Lambda function. Connect to the database using its private endpoint.
- B. Connect the Lambda function to the database VPC. Connect to the database using its private endpoint.
- C. Attach an IAM role to the Lambda function with read permissions to the database.

D. Move the database to a public subnet. Use security groups for secure access.

Correct Answer: B

Section:

Explanation:

To resolve the issue of an AWS Lambda function unable to connect to a database that has been moved to a private subnet, the Lambda function needs to be connected to the same VPC as the database. This is done by configuring the Lambda function with VPC access. This involves specifying the VPC, subnets, and security groups for the Lambda function so that it can communicate with the database using its private endpoint. Option B is correct as it directly addresses the issue without compromising security. AWS documentation on configuring VPC access for Lambda provides guidance on this setup [Configuring VPC Access for Lambda](#).

QUESTION 196

A company is running Amazon RDS for PostgreSQL Multi-AZ DB clusters. The company uses an AWS Cloud Formation template to create the databases individually with a default size of 100 GB. The company creates the databases every Monday and deletes the databases every Friday.

Occasionally, the databases run low on disk space and initiate an Amazon CloudWatch alarm. A SysOps administrator must prevent the databases from running low on disk space in the future.

Which solution will meet these requirements with the FEWEST changes to the application?

- A. Modify the CloudFormation template to use Amazon Aurora PostgreSQL as the DB engine.
- B. Modify the CloudFormation template to use Amazon DynamoDB as the database. Activate storage auto scaling during creation of the tables
- C. Modify the Cloud Formation template to activate storage auto scaling on the existing DB instances.
- D. Create a CloudWatch alarm to monitor DB instance storage space. Configure the alarm to invoke the VACUUM command.

Correct Answer: C

Section:

Explanation:

To prevent Amazon RDS for PostgreSQL Multi-AZ DB instances from running low on disk space, enabling storage auto-scaling is the most straightforward solution. This feature automatically adjusts the storage capacity of the DB instance when it approaches its limit, thus preventing the database from running out of space and triggering CloudWatch alarms. Option C is the least intrusive and most effective solution as it only requires a modification to the existing CloudFormation template to enable auto-scaling on storage. For reference, see [AWS documentation on managing RDS storage automatically](#) [Managing RDS Storage Automatically](#).

QUESTION 197

A SysOps administrator manages a company's Amazon S3 buckets. The SysOps administrator has identified 5 GB of incomplete multipart uploads in an S3 bucket in the company's AWS account. The SysOps administrator needs to reduce the number of incomplete multipart upload objects in the S3 bucket.

Which solution will meet this requirement?

- A. Create an S3 Lifecycle rule on the S3 bucket to delete expired markers or incomplete multipart uploads
- B. Require users that perform uploads of files into Amazon S3 to use the S3 TransferUtility.
- C. Enable S3 Versioning on the S3 bucket that contains the incomplete multipart uploads.
- D. Create an S3 Object Lambda Access Point to delete incomplete multipart uploads.

Correct Answer: A

Section:

Explanation:

To manage incomplete multipart uploads in an Amazon S3 bucket, creating an S3 Lifecycle rule to specifically address these uploads is the most effective method. The rule can be configured to automatically delete expired multipart upload parts, which helps in cleaning up unused data and reducing storage costs. Option A is correct as it directly addresses the requirement to manage incomplete uploads effectively. Reference on setting up S3 Lifecycle policies can be found here [Amazon S3 Lifecycle](#).

QUESTION 198

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 199

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 out-of-memory errors. The company adds memory to the web application and wants to track operating system memory utilization. A CloudWatch memory metric does not currently exist for 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 200

A company is using an Amazon CloudWatch alarm to 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 to 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 201

A SysOps administrator is responsible for more than 50 Amazon EC2 instances that 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 202

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 in 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.

