Exam Questions DOP-C02

AWS Certified DevOps Engineer - Professional

https://www.2passeasy.com/dumps/DOP-C02/

NEW QUESTION 1
A company's DevOps engineer uses AWS Systems Manager to perform maintenance tasks during maintenance windows. The company has a few Amazon EC2 instances that require a restart after notifications from AWS Health. The DevOps engineer needs to implement an automated solution to remediate these notifications. The DevOps engineer creates an Amazon EventBridge rule.
How should the DevOps engineer configure the EventBridge rule to meet these requirements?

A. Configure an event source of AWS Health, a service of EC2. and an event type that indicates instance maintenanc
B. Target a Systems Manager document to restart the EC2 instance.
C. Configure an event source of Systems Manager and an event type that indicates a maintenance windo
D. Target a Systems Manager document to restart the EC2 instance.
E. Configure an event source of AWS Health, a service of EC2, and an event type that indicates instance maintenanc
F. Target a newly created AWS Lambda function thatregisters an automation task to restart the EC2 instance during a maintenance window.
G. Configure an event source of EC2 and an event type that indicates instance maintenanc
H. Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.

Answer: C

Explanation:
 AWS Health provides real-time events and information related to your AWS infrastructure. It can be integrated with Amazon EventBridge to act upon the health events automatically. If the maintenance notification from AWS Health indicates that an EC2 instance requires a restart, you can set up an EventBridge rule to respond to such events. In this case, the target of this rule would be a Lambda function that would trigger a Systems Manager automation to restart the EC2 instance during a maintenance window. Remember, AWS Health is the source of the events (not EC2 or Systems Manager), and AWS Lambda can be used to execute complex remediation tasks, such as scheduling maintenance tasks via Systems Manager.
The following are the steps involved in configuring the EventBridge rule to meet these requirements:
? Configure an event source of AWS Health, a service of EC2, and an event type that indicates instance maintenance.
? Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.
The AWS Lambda function will be triggered by the event from AWS Health. The function will then register an automation task to restart the EC2 instance during the next maintenance window.

NEW QUESTION 2
A company is examining its disaster recovery capability and wants the ability to switch over its daily operations to a secondary AWS Region. The company uses AWS CodeCommit as a source control tool in the primary Region.
A DevOps engineer must provide the capability for the company to develop code in the secondary Region. If the company needs to use the secondary Region, developers can add an additional remote URL to their local Git configuration.
Which solution will meet these requirements?

A. Create a CodeCommit repository in the secondary Regio
B. Create an AWS CodeBuild project to perform a Git mirror operation of the primary Region's CodeCommit repository to the secondary Region's CodeCommit repositor
C. Create an AWS Lambda function that invokes the CodeBuild projec
D. Create an Amazon EventBridge rule that reacts to merge events in the primary Region's CodeCommit repositor
E. Configure the EventBridge rule to invoke the Lambda function.
F. Create an Amazon S3 bucket in the secondary Regio
G. Create an AWS Fargate task to perform a Git mirror operation of the primary Region's CodeCommit repository and copy the result to the S3 bucke
H. Create an AWS Lambda function that initiates the Fargate tas
I. Create an Amazon EventBridge rule that reacts to merge events in the CodeCommitrepositor
J. Configure the EventBridge rule to invoke the Lambda function.
K. Create an AWS CodeArtifact repository in the secondary Regio
L. Create an AWS CodePipeline pipeline that uses the primary Region's CodeCommit repository for the source actio
M. Create a Cross-Region stage in the pipeline that packages the CodeCommit repository contents and stores the contents in the CodeArtifact repository when a pull request is merged into the CodeCommit repository.
N. Create an AWS Cloud9 environment and a CodeCommit repository in the secondary Regio
O. Configure the primary Region's CodeCommit repository as a remote repository in the AWS Cloud9 environmen
P. Connect the secondary Region's CodeCommit repository to the AWS Cloud9 environment.

Answer: A

Explanation:
 The best solution to meet the disaster recovery capability and allow developers to switch over to a secondary AWS Region for code development is option A. This involves creating a CodeCommit repository in the secondary Region and setting up an AWS CodeBuild project to perform a Git mirror operation of the primary Region's CodeCommit repository to the secondary Region's repository. An AWS Lambda function is then created to invoke the CodeBuild project. Additionally, an Amazon EventBridge rule is configured to react to merge events in the primary Region's CodeCommit repository and invoke the Lambda function12. This setup ensures that the secondary Region's repository is always up-to-date with the primary repository, allowing for a seamless transition in case of a disaster recovery event1.
References:
? AWS CodeCommit User Guide on resilience and disaster recovery1.
? AWS Documentation on monitoring CodeCommit events in Amazon EventBridge and Amazon CloudWatch Events2.

NEW QUESTION 3
A DevOps engineer needs to apply a core set of security controls to an existing set of AWS accounts. The accounts are in an organization in AWS Organizations. Individual teams will administer individual accounts by using the AdministratorAccess AWS managed policy. For all accounts. AWS CloudTrail and AWS Config must be turned on in all available AWS Regions. Individual account administrators must not be able to edit or delete any of the baseline resources. However, individual account administrators must be able to edit or delete their own CloudTrail trails and AWS Config rules.
Which solution will meet these requirements in the MOST operationally efficient way?

A. Create an AWS CloudFormation template that defines the standard account resource
B. Deploy the template to all accounts from the organization's management account by using CloudFormation StackSet
C. Set the stack policy to deny Update:Delete actions.
D. Enable AWS Control Towe

E. Enroll the existing accounts in AWS Control Towe
F. Grant the individual account administrators access to CloudTrail and AWS Config.
G. Designate an AWS Config management accoun
H. Create AWS Config recorders in all accounts by using AWS CloudFormation StackSet
I. Deploy AWS Config rules to the organization by using the AWS Config management accoun
J. Create a CloudTrail organization trail in the organization's management accoun
K. Deny modification or deletion of the AWS Config recorders by using an SCP.
L. Create an AWS CloudFormation template that defines the standard account resource
M. Deploy the template to all accounts from the organization's management account by using Cloud Formation StackSets Create an SCP that prevents updates or deletions to CloudTrail resources or AWS Config resources unless the principal is an administrator of the organization's management account.

Answer: D

NEW QUESTION 4
A company is implementing a well-architected design for its globally accessible API stack. The design needs to ensure both high reliability and fast response times for users located in North America and Europe.
The API stack contains the following three tiers: Amazon API Gateway
AWS Lambda Amazon DynamoDB
Which solution will meet the requirements?

A. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using health check
B. Configure the APIs to forward requests to a Lambda function in that Regio
C. Configure the Lambda functions to retrieve and update the data in a DynamoDB table in the same Region as the Lambda function.
D. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using latency-based routing and health check
E. Configure the APIs to forward requests to a Lambda function in that Regio
F. Configure the Lambda functions to retrieve and update the data in a DynamoDB global table.
G. Configure Amazon Route 53 to point to API Gateway in North America, create a disaster recovery API in Europe, and configure both APIs to forward requests to the Lambda functions in that Regio
H. Retrieve the data from a DynamoDB global tabl
I. Deploy a Lambda function to check the North America API health every 5 minute
J. In the event of a failure, update Route 53 to point to the disaster recovery API.
K. Configure Amazon Route 53 to point to API Gateway API in North America using latency-based routin
L. Configure the API to forward requests to the Lambda function in the Region nearest to the use
M. Configure the Lambda function to retrieve and update the data in a DynamoDB table.

Answer: B

NEW QUESTION 5
A company hosts its staging website using an Amazon EC2 instance backed with Amazon EBS storage. The company wants to recover quickly with minimal data losses in the event of network connectivity issues or power failures on the EC2 instance.
Which solution will meet these requirements?

A. Add the instance to an EC2 Auto Scaling group with the minimum, maximum, and desired capacity set to 1.
B. Add the instance to an EC2 Auto Scaling group with a lifecycle hook to detach the EBS volume when the EC2 instance shuts down or terminates.
C. Create an Amazon CloudWatch alarm for the StatusCheckFailed System metric and select the EC2 action to recover the instance.
D. Create an Amazon CloudWatch alarm for the StatusCheckFailed Instance metric and select the EC2 action to reboot the instance.

Answer: C

Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance- recover.html

NEW QUESTION 6
A company has a data ingestion application that runs across multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to monitor the application and consolidate access to the application. Currently the company is running the application on Amazon EC2 instances from several Auto Scaling groups. The EC2 instances have no access to the internet because the data is sensitive Engineers have deployed the necessary VPC endpoints. The EC2 instances run a custom AMI that is built specifically tor the application.
To maintain and troubleshoot the application, system administrators need the ability to log in to the EC2 instances. This access must be automated and controlled centrally. The company's security team must receive a notification whenever the instances are accessed.
Which solution will meet these requirements?

A. Create an Amazon EventBridge rule to send notifications to the security team whenever a user logs in to an EC2 instance Use EC2 Instance Connect to log in to the instance
B. Deploy Auto Scaling groups by using AWS Cloud Formation Use the cfn-init helper script to deploy appropriate VPC routes for external access Rebuild the custom AMI so that the custom AMI includes AWS Systems Manager Agent.
C. Deploy a NAT gateway and a bastion host that has internet access Create a security group that allows incoming traffic on all the EC2 instances from the bastion host Install AWS Systems Manager Agent on all the EC2 instances Use Auto Scaling group lifecycle hooks for monitoring and auditing access Use Systems Manager Session Manager to log into the instances Send logs to a log group m Amazon CloudWatch Log
D. Export data to Amazon S3 for auditing Send notifications to the security team by using S3 event notifications.
E. Use EC2 Image Builder to rebuild the custom AMI Include the most recent version of AWS Systems Manager Agent in the Image Configure the Auto Scaling group to attach the AmazonSSMManagedinstanceCore role to all the EC2 instances Use Systems Manager Session Manager to log in to the instances Enable logging of session details to Amazon S3 Create an S3 event notification for new file uploads to send a message to the security team through an Amazon Simple Notification Service (Amazon SNS) topic.
F. Use AWS Systems Manager Automation to build Systems Manager Agent into the custom AMI Configure AWS Configure to attach an SCP to the root organization account to allow the EC2 instances to connect to Systems Manager Use Systems Manager Session Manager to log in to the instances Enable logging of session details to Amazon S3 Create an S3 event notification for new file uploads to send a message to the security team through an Amazon Simple Notification Service (Amazon SNS) topic.

Answer: C

Explanation:
Even if AmazonSSMManagedInstanceCore is a managed policy and not an IAM role I will go with C because this policy is to be attached to an IAM role for EC2 to access System Manager.

NEW QUESTION 7
A company has a new AWS account that teams will use to deploy various applications. The teams will create many Amazon S3 buckets for application- specific purposes and to store AWS CloudTrail logs. The company has enabled Amazon Macie for the account.
A DevOps engineer needs to optimize the Macie costs for the account without compromising the account's functionality.
Which solutions will meet these requirements? (Select TWO.)

A. Exclude S3 buckets that contain CloudTrail logs from automated discovery.
B. Exclude S3 buckets that have public read access from automated discovery.
C. Configure scheduled daily discovery jobs for all S3 buckets in the account.
D. Configure discovery jobs to include S3 objects based on the last modified criterion.
E. Configure discovery jobs to include S3 objects that are tagged as production only.

Answer: AD

Explanation:
To optimize the Macie costs for the account without compromising the account's functionality, the DevOps engineer needs to exclude S3 buckets that do not contain sensitive data from automated discovery. S3 buckets that contain CloudTrail logs are unlikely to have sensitive data, and Macie charges for scanning and monitoring data in S3 buckets. Therefore, excluding S3 buckets that contain CloudTrail logs from automated discovery can reduce Macie costs. Similarly, configuring discovery jobs to include S3 objects based on the last modified criterion can also reduce Macie costs, as it will only scan and monitor new or updated objects, rather than all objects in the bucket.

NEW QUESTION 8
A company has an application and a CI/CD pipeline. The CI/CD pipeline consists of an AWS CodePipeline pipeline and an AWS CodeBuild project. The CodeBuild project runs tests against the application as part of the build process and outputs a test report. The company must keep the test reports for 90 days.
Which solution will meet these requirements?

A. Add a new stage in the CodePipeline pipeline after the stage that contains the CodeBuild projec
B. Create an Amazon S3 bucket to store the report
C. Configure an S3 deploy action type in the new CodePipeline stage with the appropriate path and format for the reports.
D. Add a report group in the CodeBuild project buildspec file with the appropriate path and format for the report
E. Create an Amazon S3 bucket to store the report
F. Configure an Amazon EventBridge rule that invokes an AWS Lambda function to copy the reports to the S3 bucket when a build is complete
G. Create an S3 Lifecycle rule to expire the objects after 90 days.
H. Add a new stage in the CodePipeline pipelin
I. Configure a test action type with the appropriate path and format for the report
J. Configure the report expiration time to be 90 days in the CodeBuild project buildspec file.
K. Add a report group in the CodeBuild project buildspec file with the appropriate path and format for the report
L. Create an Amazon S3 bucket to store the report
M. Configure the report group as an artifact in the CodeBuild project buildspec fil
N. Configure the S3 bucket as the artifact destinatio
O. Set the object expiration to 90 days.

Answer: B

Explanation:
The correct solution is to add a report group in the AWS CodeBuild project buildspec file with the appropriate path and format for the reports. Then, create an Amazon S3 bucket to store the reports. You should configure an Amazon EventBridge rule that invokes an AWS Lambda function to copy the reports to the S3 bucket when a build is completed. Finally, create an S3 Lifecycle rule to expire the objects after 90 days. This approach allows for the automated transfer of reports to long-term storage and ensures
they are retained for the required duration without manual intervention1. References:
? AWS CodeBuild User Guide on test reporting1.
? AWS CodeBuild User Guide on working with report groups2.
? AWS Documentation on using AWS CodePipeline with AWS CodeBuild3.

NEW QUESTION 9
A company recently launched multiple applications that use Application Load Balancers. Application response time often slows down when the applications experience problems A DevOps engineer needs to Implement a monitoring solution that alerts the company when the applications begin to perform slowly The DevOps engineer creates an Amazon Simple Notification Semce (Amazon SNS) topic and subscribe the company's email address to the topic
What should the DevOps engineer do next to meet the requirements?

A. Create an Amazon EventBridge rule that invokes an AWS Lambda function to query the applications on a 5-minute interval Configure the Lambda function to publish a notification to the SNS topic when the applications return errors.
B. Create an Amazon CloudWatch Synthetics canary that runs a custom script to query the applications on a 5-minute interva
C. Configure the canary to use the SNS topic when the applications return errors.
D. Create an Amazon CloudWatch alarm that uses the AWS/AppljcabonELB namespace RequestCountPerTarget metric Configure the CloudWatch alarm to send a notification when the number of connections becomes greater than the configured number of threads that the application supports Configure the CloudWatch alarm to use the SNS topic.
E. Create an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric Configure the CloudWatch alarm to send a notification when the average response time becomes greater than the longest response time that the application supports Configure the CloudWatch alarm to use the SNS topic

Answer: B

Explanation:
? Option A is incorrect because creating an Amazon EventBridge rule that invokes an AWS Lambda function to query the applications on a 5-minute interval is not a valid solution. EventBridge rules can only trigger Lambda functions based on events, not on time intervals. Moreover, querying the applications on a 5-minute

interval might incur unnecessary costs and network overhead, and might not detect performance issues in real time.
? Option B is correct because creating an Amazon CloudWatch Synthetics canary that runs a custom script to query the applications on a 5-minute interval is a valid solution. CloudWatch Synthetics canaries are configurable scripts that monitor endpoints and APIs by simulating customer behavior. Canaries can run as often as once per minute, and can measure the latency and availability of the applications. Canaries can also send notifications to an Amazon SNS topic when they detect errors or performance issues1.
? Option C is incorrect because creating an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric is not a valid solution. The RequestCountPerTarget metric measures the number of requests completed or connections made per target in a target group2. This metric does not reflect the application response time, which is the requirement. Moreover, configuring the CloudWatch alarm to send a notification when the number of connections becomes greater than the configured number of threads that the application supports is not a valid way to measure the application performance, as it depends on the application design and implementation.
? Option D is incorrect because creating an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric is not a valid solution, for the same reason as option C. The RequestCountPerTarget metric does not reflect the application response time, which is the requirement. Moreover, configuring the CloudWatch alarm to send a notification when the average response time becomes greater than the longest response time that the application supports is not a valid way to measure the application performance, as it does not account for variability or outliers in the response time distribution.
References:
? 1: Using synthetic monitoring
? 2: Application Load Balancer metrics

NEW QUESTION 10
A company has an organization in AWS Organizations. The organization includes workload accounts that contain enterprise applications. The company centrally manages users from an operations account. No users can be created in the workload accounts. The company recently added an operations team and must provide the operations team members with administrator access to each workload account.
Which combination of actions will provide this access? (Choose three.)

A. Create a SysAdmin role in the operations accoun
B. Attach the AdministratorAccess policy to the rol
C. Modify the trust relationship to allow the sts:AssumeRole action from the workload accounts.
D. Create a SysAdmin role in each workload accoun
E. Attach the AdministratorAccess policy to the rol
F. Modify the trust relationship to allow the sts:AssumeRole action from the operations account.
G. Create an Amazon Cognito identity pool in the operations accoun
H. Attach the SysAdmin role as an authenticated role.
I. In the operations account, create an IAM user for each operations team member.
J. In the operations account, create an IAM user group that is named SysAdmin
K. Add an IAM policy that allows the sts:AssumeRole action for the SysAdmin role in each workload accoun
L. Add all operations team members to the group.
M. Create an Amazon Cognito user pool in the operations accoun
N. Create an Amazon Cognito user for each operations team member.

Answer: BDE

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account- with-roles.html

NEW QUESTION 10
A company uses an Amazon API Gateway regional REST API to host its application API. The REST API has a custom domain. The REST API's default endpoint is deactivated.
The company's internal teams consume the API. The company wants to use mutual TLS between the API and the internal teams as an additional layer of authentication.
Which combination of steps will meet these requirements? (Select TWO.)

A. Use AWS Certificate Manager (ACM) to create a private certificate authority (CA). Provision a client certificate that is signed by the private CA.
B. Provision a client certificate that is signed by a public certificate authority (CA). Import the certificate into AWS Certificate Manager (ACM).
C. Upload the provisioned client certificate to an Amazon S3 bucke
D. Configure the API Gateway mutual TLS to use the client certificate that is stored in the S3 bucket as the trust store.
E. Upload the provisioned client certificate private key to an Amazon S3 bucke
F. Configure the API Gateway mutual TLS to use the private key that is stored in the S3 bucket as the trust store.
G. Upload the root private certificate authority (CA) certificate to an Amazon S3 bucke
H. Configure the API Gateway mutual TLS to use the private CA certificate that is stored in the S3 bucket as the trust store.

Answer: AE

Explanation:
Mutual TLS (mTLS) authentication requires two-way authentication between the client and the server. For Amazon API Gateway, you can enable mTLS for a custom domain name, which requires clients to present X.509 certificates to verify their identity to access your API. To set up mTLS, you would typically use AWS Certificate Manager (ACM) to create a private certificate authority (CA) and provision a client certificate signed by this private
CA. The root CA certificate is then uploaded to an Amazon S3 bucket and configured in API Gateway as the trust store12.
References:
? Introducing mutual TLS authentication for Amazon API Gateway1.
? Configuring mutual TLS authentication for a REST API2.
? AWS Private Certificate Authority details3.
? AWS Certificate Manager Private Certificate Authority updates4.

NEW QUESTION 12
A company has multiple accounts in an organization in AWS Organizations. The company's SecOps team needs to receive an Amazon Simple Notification Service (Amazon SNS) notification if any account in the organization turns off the Block Public Access feature on an Amazon S3 bucket. A DevOps engineer must implement this change without affecting the operation of any AWS accounts. The implementation must ensure that individual member accounts in the organization cannot turn off the notification.
Which solution will meet these requirements?

A. Designate an account to be the delegated Amazon GuardDuty administrator accoun
B. Turn on GuardDuty for all accounts across the organizatio
C. In the GuardDuty administrator account, create an SNS topi
D. Subscribe the SecOps team's email address to the SNS topi
E. In the same account, create an Amazon EventBridge rule that uses an event pattern for GuardDuty findings and a target of the SNS topic.
F. Create an AWS CloudFormation template that creates an SNS topic and subscribes the SecOps team's email address to the SNS topi
G. In the template, include an Amazon EventBridge rule that uses an event pattern of CloudTrail activity for s3:PutBucketPublicAccessBlock and a target of the SNS topi
H. Deploy the stack to every account in the organization by using CloudFormation StackSets.
I. Turn on AWS Config across the organizatio
J. In the delegated administrator account, create an SNS topi
K. Subscribe the SecOps team's email address to the SNS topi
L. Deploy a conformance pack that uses the s3-bucket-level-public-access-prohibited AWS Config managed rule in each account and uses an AWS Systems Manager document to publish an event to the SNS topic to notify the SecOps team.
M. Turn on Amazon Inspector across the organizatio
N. In the Amazon Inspector delegated administrator account, create an SNS topi
O. Subscribe the SecOps team's email address to the SNS topi
P. In the same account, create an Amazon EventBridge rule that uses an event pattern for public network exposure of the S3 bucket and publishes an event to the SNS topic to notify the SecOps team.

Answer: C

Explanation:
 Amazon GuardDuty is primarily on threat detection and response, not configuration monitoring A conformance pack is a collection of AWS Config rules and remediation actions that can be easily deployed as a single entity in an account and a Region or across an organization in AWS Organizations.
https://docs.aws.amazon.com/config/latest/developerguide/conformance-packs.html https://docs.aws.amazon.com/config/latest/developerguide/s3-account-level-public-access- blocks.html

NEW QUESTION 17
A company has chosen AWS to host a new application. The company needs to implement a multi-account strategy. A DevOps engineer creates a new AWS account and an organization in AWS Organizations. The DevOps engineer also creates the OU structure for the organization and sets up a landing zone by using AWS Control Tower.
The DevOps engineer must implement a solution that automatically deploys resources for new accounts that users create through AWS Control Tower Account Factory. When a user creates a new account, the solution must apply AWS CloudFormation templates and SCPs that are customized for the OU or the account to automatically deploy all the resources that are attached to the account. All the OUs are enrolled in AWS Control Tower.
Which solution will meet these requirements in the MOST automated way?

A. Use AWS Service Catalog with AWS Control Towe
B. Create portfolios and products in AWS Service Catalo
C. Grant granular permissions to provision these resource
D. Deploy SCPs by using the AWS CLI and JSON documents.
E. Deploy CloudFormation stack sets by using the required template
F. Enable automatic deploymen
G. Deploy stack instances to the required account
H. Deploy a CloudFormationstack set to the organization's management account to deploy SCPs.
I. Create an Amazon EventBridge rule to detect the CreateManagedAccount even
J. Configure AWS Service Catalog as the target to deploy resources to any new account
K. Deploy SCPs by using the AWS CLI and JSON documents.
L. Deploy the Customizations for AWS Control Tower (CfCT) solutio
M. Use an AWS CodeCommit repository as the sourc
N. In the repository, create a custom package that includes the CloudFormation templates and the SCP JSON documents.

Answer: D

Explanation:
 The CfCT solution is designed for the exact purpose stated in the question. It extends the capabilities of AWS Control Tower by providing you with a way to automate resource provisioning and apply custom configurations across all AWS accounts created in the Control Tower environment. This enables the company to implement additional account customizations when new accounts are provisioned via the Control Tower Account Factory. The CloudFormation templates and SCPs can be added to a CodeCommit repository and will be automatically deployed to new accounts when they are created. This provides a highly automated solution that does not require manual intervention to deploy resources and SCPs to new accounts.

NEW QUESTION 21
A company's application is currently deployed to a single AWS Region. Recently, the company opened a new office on a different continent. The users in the new office are experiencing high latency. The company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and uses Amazon DynamoDB as the database layer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. A DevOps engineer is tasked with minimizing application response times and improving availability for users in both Regions.
Which combination of actions should be taken to address the latency issues? (Choose three.)

A. Create a new DynamoDB table in the new Region with cross-Region replication enabled.
B. Create new ALB and Auto Scaling group global resources and configure the new ALB to direct traffic to the new Auto Scaling group.
C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group.
D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB.
E. Create Amazon Route 53 aliases, health checks, and failover routing policies to route to the ALB.
F. Convert the DynamoDB table to a global table.

Answer: CDF

Explanation:
 C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group. This will allow users in the new Region to access the application with lower latency by reducing the network hops between the user and the application servers.
* D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB. This will enable Route 53 to route user traffic to the

nearest healthy ALB, based on the latency between the user and the ALBs.
* F. Convert the DynamoDB table to a global table. This will enable reads and writes to the table in both Regions with low latency, improving the overall response time of the application

NEW QUESTION 23
A DevOps team manages an API running on-premises that serves as a backend for an Amazon API Gateway endpoint. Customers have been complaining about high response latencies, which the development team has verified using the API Gateway latency metrics in Amazon CloudWatch. To identify the cause, the team needs to collect relevant data without introducing additional latency.
Which actions should be taken to accomplish this? (Choose two.)

A. Install the CloudWatch agent server side and configure the agent to upload relevant logs to CloudWatch.
B. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and upload those segments to X-Ray during each request.
C. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and use the X-Ray daemon to upload segments to X-Ray.
D. Modify the on-premises application to send log information back to API Gateway with each request.
E. Modify the on-premises application to calculate and upload statistical data relevant to the API service requests to CloudWatch metrics.

Answer: AC

Explanation:
 https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/install-CloudWatch-Agent-on-premise.html
https://docs.aws.amazon.com/xray/latest/devguide/xray-api-sendingdata.html

NEW QUESTION 27
A company has multiple AWS accounts. The company uses AWS IAM Identity Center (AWS Single Sign-On) that is integrated with AWS Toolkit for Microsoft Azure DevOps. The attributes for access control feature is enabled in IAM Identity Center.
The attribute mapping list contains two entries. The department key is mapped to
${path:enterprise.department}. The costCenter key is mapped to
${path:enterprise.costCenter}.
All existing Amazon EC2 instances have a department tag that corresponds to three company departments (d1, d2, d3). A DevOps engineer must create policies based on the matching attributes. The policies must minimize administrative effort and must grant each Azure AD user access to only the EC2 instances that are tagged with the user's respective department name.
Which condition key should the DevOps engineer include in the custom permissions policies to meet these requirements?

A.
```
"Condition": {
    "ForAllValues:StringEquals": {
        "aws:TagKeys": ["department"]
    )
}
```

B.
```
"Condition": {
    "StringEquals": {
        "aws:PrincipalTag/department": "$(aws:ResourceTag/department)"
    )
}
```

C.
```
"Condition": {
    "StringEquals": {
        "ec2:ResourceTag/department": "$(aws:PrincipalTag/department)"
    )
}
```

D.
```
"Condition": {
    "ForAllValues:StringEquals": {
        "ec2:ResourceTag/department": ["d1", "d2", "d3"]
    )
}
```

A.

Answer: C

Explanation:
https://docs.aws.amazon.com/singlesignon/latest/userguide/configure- abac.html

NEW QUESTION 32
A company requires its internal business teams to launch resources through pre-approved AWS CloudFormation templates only. The security team requires automated monitoring when resources drift from their expected state.
Which strategy should be used to meet these requirements?

A. Allow users to deploy CloudFormation stacks using a CloudFormation service role onl
B. Use CloudFormation drift detection to detect when resources have drifted from their expected state.
C. Allow users to deploy CloudFormation stacks using a CloudFormation service role onl
D. Use AWS Config rules to detect when resources have drifted from their expected state.
E. Allow users to deploy CloudFormation stacks using AWS Service Catalog onl
F. Enforce the use of a launch constrain
G. Use AWS Config rules to detect when resources have drifted from their expected state.
H. Allow users to deploy CloudFormation stacks using AWS Service Catalog onl
I. Enforce the use of a template constrain
J. Use Amazon EventBridge notifications to detect when resources have drifted from their expected state.

Answer: C

Explanation:
The correct answer is C. Allowing users to deploy CloudFormation stacks using AWS Service Catalog only and enforcing the use of a launch constraint is the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only. AWS Service Catalog is a service that enables organizations to create and manage catalogs of IT services that are approved for use on AWS. A launch constraint is a rule that specifies the role that AWS Service Catalog assumes when launching a product.
By using a launch constraint, the DevOps engineer can control the permissions that the users have when launching a product. Using AWS Config rules to detect when resources have drifted from their expected state is the best way to automate the monitoring of the resources. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config rules are custom or managed rules that AWS Config uses to evaluate whether your AWS resources comply with your desired configurations. By using AWS Config rules, the DevOps engineer can track the changes in the resources and identify any non-compliant resources.
Option A is incorrect because allowing users to deploy CloudFormation stacks using a CloudFormation service role only is not the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only. A CloudFormation service role is an IAM role that CloudFormation assumes to create, update, or delete the stack resources. By using a CloudFormation service role, the DevOps engineer can control the permissions that CloudFormation has when acting on the resources, but not the permissions that the users have when launching a stack. Therefore, option A does not prevent the users from launching resources that are not approved by the company. Using CloudFormation drift detection to detect when resources have drifted from their expected state is a valid way to monitor the resources, but it is not as automated and scalable as using AWS Config rules. CloudFormation drift detection is a feature that enables you to detect whether a stack's actual configuration differs, or has drifted, from its expected configuration. To use this feature, the DevOps engineer would need to manually initiate a drift detection operation on the stack or the stack resources, and then view the drift status and details in the CloudFormation console or API.
Option B is incorrect because allowing users to deploy CloudFormation stacks using a CloudFormation service role only is not the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only, as explained in option A. Using AWS Config rules to detect when resources have drifted from their expected state is a valid way to monitor the resources, as explained in option C. Option D is incorrect because enforcing the use of a template constraint is not the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only. A template constraint is a rule that defines the values or properties that users can specify when launching a product. By using a template constraint, the DevOps engineer can control the parameters that the users can provide when launching a product, but not the permissions that the users have when launching a product. Therefore, option D does not prevent the users from launching resources that are not approved by the company. Using Amazon EventBridge notifications to detect when resources have drifted from their expected state is a less reliable and consistent solution than using AWS Config rules. Amazon EventBridge is a service that enables you to connect your applications with data from a variety of sources. Amazon EventBridge can deliver a stream of real-time data from event sources, such as AWS services, and route
that data to targets, such as AWS Lambda functions. However, to use this solution, the DevOps engineer would need to configure the event source, the event bus, the event rule, and the event target for each resource type that needs to be monitored, which is more complex and error-prone than using AWS Config rules.

NEW QUESTION 35
A company's application development team uses Linux-based Amazon EC2 instances as bastion hosts. Inbound SSH access to the bastion hosts is restricted to specific IP addresses, as defined in the associated security groups. The company's security team wants to receive a notification if the security group rules are modified to allow SSH access from any IP address.
What should a DevOps engineer do to meet this requirement?

A. Create an Amazon EventBridge rule with a source of aws.cloudtrail and the event name AuthorizeSecurityGroupIngres
B. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.
C. Enable Amazon GuardDuty and check the findings for security groups in AWS Security Hu
D. Configure an Amazon EventBridge rule with a custom pattern that matches GuardDuty events with an output of NON_COMPLIAN
E. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.
F. Create an AWS Config rule by using the restricted-ssh managed rule to check whether security groups disallow unrestricted incoming SSH traffi
G. Configure automatic remediation to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.
H. Enable Amazon Inspecto
I. Include the Common Vulnerabilities and Exposures-1.1 rules package to check the security groups that are associated with the bastion host
J. Configure Amazon Inspector to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/monitor-security-group-changes-ec2/

NEW QUESTION 37
A company has multiple development groups working in a single shared AWS account. The Senior Manager of the groups wants to be alerted via a third-party API call when the creation of resources approaches the service limits for the account.
Which solution will accomplish this with the LEAST amount of development effort?

A. Create an Amazon CloudWatch Event rule that runs periodically and targets an AWS Lambda functio
B. Within the Lambda function, evaluate the current state of the AWS environment and compare deployed resource values to resource limits on the accoun
C. Notify the Senior Manager if the account is approaching a service limit.

D. Deploy an AWS Lambda function that refreshes AWS Trusted Advisor checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodicall

E. Create another CloudWatch Events rule with an event pattern matching Trusted Advisor events and a target Lambda functio

F. In the target Lambda function, notify the Senior Manager.

G. Deploy an AWS Lambda function that refreshes AWS Personal Health Dashboard checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodicall

H. Create another CloudWatch Events rule with an event pattern matching Personal Health Dashboard events and a target Lambda functio

I. In the target Lambda function, notify the Senior Manager.

J. Add an AWS Config custom rule that runs periodically, checks the AWS service limit status, and streams notifications to an Amazon SNS topi

K. Deploy an AWS Lambda function that notifies the Senior Manager, and subscribe the Lambda function to the SNS topic.

Answer: B

Explanation:
To meet the requirements, the company needs to create a solution that alerts the Senior Manager when the creation of resources approaches the service limits for the account with the least amount of development effort. The company can use AWS Trusted Advisor, which is a service that provides best practice recommendations for cost optimization, performance, security, and service limits. The company can deploy an AWS Lambda function that refreshes Trusted Advisor checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically. This will ensure that Trusted Advisor checks are up to date and reflect the current state of the account. The company can then create another CloudWatch Events rule with an event pattern matching Trusted Advisor events and a target Lambda function. The event pattern can filter for events related to service limit checks and their status. The target Lambda function can notify the Senior Manager via a third-party API call if the event indicates that the account is approaching or exceeding a service limit.

NEW QUESTION 41

A company provides an application to customers. The application has an Amazon API Gateway REST API that invokes an AWS Lambda function. On initialization, the Lambda function loads a large amount of data from an Amazon DynamoDB table. The data load process results in long cold-start times of 8-10 seconds. The DynamoDB table has DynamoDB Accelerator (DAX) configured.

Customers report that the application intermittently takes a long time to respond to requests. The application receives thousands of requests throughout the day. In the middle of the day, the application experiences 10 times more requests than at any other time of the day. Near the end of the day, the application's request volume decreases to 10% of its normal total.

A DevOps engineer needs to reduce the latency of the Lambda function at all times of the day.

Which solution will meet these requirements?

A. Configure provisioned concurrency on the Lambda function with a concurrency value of 1. Delete the DAX cluster for the DynamoDB table.

B. Configure reserved concurrency on the Lambda function with a concurrency value of 0.

C. Configure provisioned concurrency on the Lambda functio

D. Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.

E. Configure reserved concurrency on the Lambda functio

F. Configure AWS Application Auto Scaling on the API Gateway API with a reserved concurrency maximum value of 100.

Answer: C

Explanation:
The following are the steps that the DevOps engineer should take to reduce the latency of the Lambda function at all times of the day:
? Configure provisioned concurrency on the Lambda function.
? Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.
The provisioned concurrency setting ensures that there is always a minimum number of Lambda function instances available to handle requests. The Application Auto Scaling setting will automatically scale the number of Lambda function instances up or down based on the demand for the application.
This solution will ensure that the Lambda function is able to handle the increased load during the middle of the day, while also keeping the cold-start latency low.
The following are the reasons why the other options are not correct:
? Option A is incorrect because it will not reduce the cold-start latency of the Lambda function.
? Option B is incorrect because it will not scale the number of Lambda function instances up or down based on demand.
? Option D is incorrect because it will only configure reserved concurrency on the API Gateway API, which will not affect the Lambda function.

NEW QUESTION 43

A media company has several thousand Amazon EC2 instances in an AWS account. The company is using Slack and a shared email inbox for team communications and important updates. A DevOps engineer needs to send all AWS-scheduled EC2 maintenance notifications to the Slack channel and the shared inbox. The solution must include the instances' Name and Owner tags.

Which solution will meet these requirements?

A. Integrate AWS Trusted Advisor with AWS Config Configure a custom AWS Config rule to invoke an AWS Lambda function to publish notifications to an Amazon Simple Notification Service (Amazon SNS) topic Subscribe a Slack channel endpoint and the shared inbox to the topic.

B. Use Amazon EventBridge to monitor for AWS Health Events Configure the maintenance events to target an Amazon Simple Notification Service (Amazon SNS) topic Subscribe an AWS Lambda function to the SNS topic to send notifications to the Slack channel and the shared inbox.

C. Create an AWS Lambda function that sends EC2 maintenance notifications to the Slack channel and the shared inbox Monitor EC2 health events by using Amazon CloudWatch metrics Configure a CloudWatch alarm that invokes the Lambda function when a maintenance notification is received.

D. Configure AWS Support integration with AWS CloudTrail Create a CloudTrail lookup event to invoke an AWS Lambda function to pass EC2 maintenance notifications to Amazon Simple Notification Service (Amazon SNS) Configure Amazon SNS to target the Slack channel and the shared inbox.

Answer: B

Explanation:
https://docs.aws.amazon.com/health/latest/ug/cloudwatch-events-health.html

NEW QUESTION 46

A DevOps engineer needs to configure a blue green deployment for an existing three-tier application. The application runs on Amazon EC2 instances and uses an Amazon RDS database The EC2 instances run behind an Application Load Balancer (ALB) and are in an Auto Scaling group.

The DevOps engineer has created a launch template and an Auto Scaling group for the blue environment. The DevOps engineer also has created a launch template and an Auto Scaling group for the green environment. Each Auto Scaling group deploys to a matching blue or green target group. The target group also specifies which software blue or green gets loaded on the EC2 instances. The ALB can be configured to send traffic to the blue environments target group or the green environments target group. An Amazon Route 53 record for www example com points to the ALB.

The deployment must move traffic all at once between the software on the blue environment's EC2 instances to the newly deployed software on the green environments EC2 instances

What should the DevOps engineer do to meet these requirements?

A. Start a rolling restart to the Auto Scaling group tor the green environment to deploy the new software on the green environment's EC2 instances When the rolling restart is complete, use an AWS CLI command to update the ALB to send traffic to the green environment's target group.
B. Use an AWS CLI command to update the ALB to send traffic to the green environment's target grou
C. Then start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances.
D. Update the launch template to deploy the green environment's software on the blue environment's EC2 instances Keep the target groups and Auto Scaling groups unchanged in both environments Perform a rolling restart of the blue environment's EC2 instances.
E. Start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances When the rolling restart is complete, update the Route 53 DNS to point to the green environments endpoint on the ALB.

Answer: A

Explanation:
 This solution will meet the requirements because it will use a rolling restart to gradually replace the EC2 instances in the green environment with new instances that have the new software version installed. A rolling restart is a process that terminates and launches instances in batches, ensuring that there is always a minimum number of healthy instances in service. This way, the green environment can be updated without affecting the availability or performance of the application. When the rolling restart is complete, the DevOps engineer can use an AWS CLI command to modify the listener rules of the ALB and change the default action to forward traffic to the green environment's target group. This will switch the traffic from the blue environment to the green environment all at once, as required by the question.

NEW QUESTION 49
A company's security policies require the use of security hardened AMIS in production environments. A DevOps engineer has used EC2 Image Builder to create a pipeline that builds the AMIs on a recurring schedule.
The DevOps engineer needs to update the launch templates of the companys Auto Scaling groups. The Auto Scaling groups must use the newest AMIS during the launch of Amazon EC2 instances.
Which solution will meet these requirements with the MOST operational efficiency?

A. Configure an Amazon EventBridge rule to receive new AMI events from Image Builde
B. Target an AWS Systems Manager Run Command document that updates the launch templates of the Auto Scaling groups with the newest AMI ID.
C. Configure an Amazon EventBridge rule to receive new AMI events from Image Builde
D. Target an AWS Lambda function that updates the launch templates of the Auto Scaling groups with the newest AMI ID.
E. Configure the launch template to use a value from AWS Systems Manager Parameter Store for the AMI I
F. Configure the Image Builder pipeline to update the Parameter Store value with the newest AMI ID.
G. Configure the Image Builder distribution settings to update the launch templates with the newest AMI I
H. Configure the Auto Scaling groups to use the newest version of the launch template.

Answer: C

Explanation:
? The most operationally efficient solution is to use AWS Systems Manager Parameter Store1 to store the AMI ID and reference it in the launch template2. This way, the launch template does not need to be updated every time a new AMI is created by Image Builder. Instead, the Image Builder pipeline can update the Parameter Store value with the newest AMI ID3, and the Auto Scaling group can launch instances using the latest value from Parameter Store.
? The other solutions require updating the launch template or creating a new version of it every time a new AMI is created, which adds complexity and overhead. Additionally, using EventBridge rules and Lambda functions or Run Command documents introduces additional dependencies and potential points of failure.
References: 1: AWS Systems Manager Parameter Store 2: Using AWS Systems Manager parameters instead of AMI IDs in launch templates 3: Update an SSM parameter with
Image Builder

NEW QUESTION 53
A company deploys updates to its Amazon API Gateway API several times a week by using an AWS CodePipeline pipeline. As part of the update process the company exports the JavaScript SDK for the API from the API. Gateway console and uploads the SDK to an Amazon S3 bucket
The company has configured an Amazon CloudFront distribution that uses the S3 bucket as an origin Web client then download the SDK by using the CloudFront distribution's endpoint. A DevOps engineer needs to implement a solution to make the new SDK available automatically during new API deployments.
Which solution will meet these requirements?

A. Create a CodePipeline action immediately after the deployment stage of the AP
B. Configure the action to invoke an AWS Lambda functio
C. Configure the Lambda function to download the SDK from API Gateway, upload the SDK to the S3 bucket and create a CloudFront invalidation for the SDK path.
D. Create a CodePipeline action immediately after the deployment stage of the API Configure the action to use the CodePipelme integration with AP
E. Gateway to export the SDK to Amazon S3 Create another action that uses the CodePipeline integration with Amazon S3 to invalidate the cache for the SDK path.
F. Create an Amazon EventBridge rule that reacts to UpdateStage events from aws apigateway Configure the rule to invoke an AWS Lambda function to download the SDK from API Gateway upload the SDK to the S3 bucket and call the CloudFront API to create an invalidation for the SDK path.
G. Create an Amazon EventBridge rule that reacts to Creat
H. Deployment events from aws apigatewa
I. Configure the rule to invoke an AWS Lambda function to download the SDK from AP
J. Gateway upload the SDK to the S3 bucket and call the S3 API to invalidate the cache for the SDK path.

Answer: A

Explanation:
 This solution would allow the company to automate the process of updating the SDK and making it available to web clients. By adding a CodePipeline action immediately after the deployment stage of the API, the Lambda function will be invoked automatically each time the API is updated. The Lambda function should be able to download the new SDK from API Gateway, upload it to the S3 bucket and also create a CloudFront invalidation for the SDK path so that the latest version of the SDK is available for the web clients. This is the most straight forward solution and it will meet the requirements.

NEW QUESTION 55
A company that runs many workloads on AWS has an Amazon EBS spend that has increased over time. The DevOps team notices there are many unattached EBS volumes. Although there are workloads where volumes are detached, volumes over 14 days old are stale and no longer needed. A DevOps engineer has been tasked with creating automation that deletes unattached EBS volumes that have been unattached for 14 days.
Which solution will accomplish this?

A. Configure the AWS Config ec2-volume-inuse-check managed rule with a configuration changes trigger type and an Amazon EC2 volume resource targe
B. Create a new Amazon CloudWatch Events rule scheduled to execute an AWS Lambda function in 14 days to delete the specified EBS volume.
C. Use Amazon EC2 and Amazon Data Lifecycle Manager to configure a volume lifecycle polic
D. Set the interval period for unattached EBS volumes to 14 days and set the retention rule to delet
E. Set the policy target volumes as *.
F. Create an Amazon CloudWatch Events rule to execute an AWS Lambda function dail
G. The Lambda function should find unattached EBS volumes and tag them with the current date, and delete unattached volumes that have tags with dates that are more than 14 days old.
H. Use AWS Trusted Advisor to detect EBS volumes that have been detached for more than 14 day
I. Execute an AWS Lambda function that creates a snapshot and then deletes the EBS volume.

Answer: C

Explanation:
 The requirement is to create automation that deletes unattached EBS volumes that have been unattached for 14 days. To do this, the DevOps engineer needs to use the following steps:
? Create an Amazon CloudWatch Events rule to execute an AWS Lambda function
daily. CloudWatch Events is a service that enables event-driven architectures by delivering events from various sources to targets. Lambda is a service that lets you
run code without provisioning or managing servers. By creating a CloudWatch Events rule that executes a Lambda function daily, the DevOps engineer can schedule a recurring task to check and delete unattached EBS volumes.
? The Lambda function should find unattached EBS volumes and tag them with the
current date, and delete unattached volumes that have tags with dates that are more than 14 days old. The Lambda function can use the EC2 API to list and filter unattached EBS volumes based on their state and tags. The function can then tag each unattached volume with the current date using the create-tags command. The function can also compare the tag value with the current date and delete any unattached volume that has been tagged more than 14 days ago using the delete- volume command.

NEW QUESTION 60
A company needs to ensure that flow logs remain configured for all existing and new VPCs in its AWS account. The company uses an AWS CloudFormation stack to manage its VPCs. The company needs a solution that will work for any VPCs that any IAM user creates.
Which solution will meet these requirements?

A. Add the resource to the CloudFormation stack that creates the VPCs.
B. Create an organization in AWS Organization
C. Add the company's AWS account to the organizatio
D. Create an SCP to prevent users from modifying VPC flow logs.
E. Turn on AWS Confi
F. Create an AWS Config rule to check whether VPC flow logs are turned o
G. Configure automatic remediation to turn on VPC flow logs.
H. Create an IAM policy to deny the use of API calls for VPC flow log
I. Attach the IAM policy to all IAM users.

Answer: C

Explanation:
To meet the requirements of ensuring that flow logs remain configured for all existing and new VPCs in the AWS account, the company should use AWS Config and automatic remediation. AWS Config is a service that enables customers to assess, audit, and evaluate the configurations of their AWS resources. AWS Config continuously monitors and records the configuration changes of the AWS resources and evaluates them against desired configurations. Customers can use AWS Config rules to define the desired configuration state of their AWS resources and trigger actions when a resource configuration violates a rule.
One of the AWS Config rules that customers can use is vpc-flow-logs-enabled, which checks whether VPC flow logs are enabled for all VPCs in an AWS account. Customers can also configure automatic remediation for this rule, which means that AWS Config will automatically enable VPC flow logs for any VPCs that do not have them enabled. Customers can specify the destination (CloudWatch Logs or S3) and the traffic type (all, accept, or reject) for the flow logs as remediation parameters. By using AWS Config and automatic remediation, the company can ensure that flow logs remain configured for all existing and new VPCs in its AWS account, regardless of who creates them or how they are created.
The other options are not correct because they do not meet the requirements or follow best practices. Adding the resource to the CloudFormation stack that creates the VPCs is not a sufficient solution because it will only work for VPCs that are created by using the CloudFormation stack. It will not work for VPCs that are created by using other methods, such as the console or the API. Creating an organization in AWS Organizations and creating an SCP to prevent users from modifying VPC flow logs is not a good solution because it will not ensure that flow logs are enabled for all VPCs in the first place. It will only prevent users from disabling or changing flow logs after they are enabled. Creating an IAM policy to deny the use of API calls for VPC flow logs and attaching it to all IAM users is not a valid solution because it will prevent users from enabling or disabling flow logs at all.
It will also not work for VPCs that are created by using other methods, such as the console or CloudFormation.
References:
? 1: AWS::EC2::FlowLog - AWS CloudFormation
? 2: Amazon VPC Flow Logs extends CloudFormation Support to custom format subscriptions, 1-minute aggregation intervals and tagging
? 3: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud
? : About AWS Config - AWS Config
? : vpc-flow-logs-enabled - AWS Config
? : Remediate Noncompliant Resources with AWS Config Rules - AWS Config

NEW QUESTION 63
A company hosts a security auditing application in an AWS account. The auditing application uses an IAM role to access other AWS accounts. All the accounts are in the same organization in AWS Organizations.
A recent security audit revealed that users in the audited AWS accounts could modify or delete the auditing application's IAM role. The company needs to prevent any modification to the auditing application's IAM role by any entity other than a trusted administrator IAM role.
Which solution will meet these requirements?

A. Create an SCP that includes a Deny statement for changes to the auditing application's IAM rol
B. Include a condition that allows the trusted administrator IAM role to make change
C. Attach the SCP to the root of the organization.
D. Create an SCP that includes an Allow statement for changes to the auditing application's IAM role by the trusted administrator IAM rol
E. Include a Deny statement for changes by all other IAM principal
F. Attach the SCP to the IAM service in each AWS account where the auditing application has an IAM role.
G. Create an IAM permissions boundary that includes a Deny statement for changes to the auditing application's IAM rol
H. Include a condition that allows the trusted administrator IAM role to make change
I. Attach the permissions boundary to the audited AWS accounts.
J. Create an IAM permissions boundary that includes a Deny statement for changes to the auditing application's IAM rol
K. Include a condition that allows the trusted administrator IAM role to make change
L. Attach the permissions boundary to the auditing application's IAM role in the AWS accounts.

Answer: A

Explanation:
 https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps. html?icmpid=docs_orgs_console
SCPs (Service Control Policies) are the best way to restrict permissions at the organizational level, which in this case would be used to restrict modifications to the IAM role used by the auditing application, while still allowing trusted administrators to make changes to it. Options C and D are not as effective because IAM permission boundaries are applied to IAM entities (users, groups, and roles), not the account itself, and must be applied to all IAM entities in the account.

NEW QUESTION 64
A company is performing vulnerability scanning for all Amazon EC2 instances across many accounts. The accounts are in an organization in AWS Organizations.
Each account's VPCs are attached to a shared transit gateway. The VPCs send traffic to the internet through a central egress VPC. The company has enabled Amazon Inspector in a delegated administrator account and has enabled scanning for all member accounts.
A DevOps engineer discovers that some EC2 instances are listed in the "not scanning" tab in Amazon Inspector.
Which combination of actions should the DevOps engineer take to resolve this issue? (Choose three.)

A. Verify that AWS Systems Manager Agent is installed and is running on the EC2 instances that Amazon Inspector is not scanning.
B. Associate the target EC2 instances with security groups that allow outbound communication on port 443 to the AWS Systems Manager service endpoint.
C. Grant inspector: StartAssessmentRun permissions to the IAM role that the DevOps engineer is using.
D. Configure EC2 Instance Connect for the EC2 instances that Amazon Inspector is not scanning.
E. Associate the target EC2 instances with instance profiles that grant permissions to communicate with AWS Systems Manager.
F. Create a managed-instance activatio
G. Use the Activation Code and the Activation ID to register the EC2 instances.

Answer: ABE

Explanation:
 https://docs.aws.amazon.com/inspector/latest/user/scanning-ec2.html

NEW QUESTION 65
A company uses an organization in AWS Organizations that has all features enabled. The company uses AWS Backup in a primary account and uses an AWS Key Management Service (AWS KMS) key to encrypt the backups.
The company needs to automate a cross-account backup of the resources that AWS Backup backs up in the primary account. The company configures cross-account backup in the Organizations management account. The company creates a new AWS account in the organization and configures an AWS Backup backup vault in the new account. The company creates a KMS key in the new account to encrypt the backups. Finally, the company configures a new backup plan in the primary account. The destination for the new backup plan is the backup vault in the new account.
When the AWS Backup job in the primary account is invoked, the job creates backups in the primary account. However, the backups are not copied to the new account's backup vault.
Which combination of steps must the company take so that backups can be copied to the new account's backup vault? (Select TWO.)

A. Edit the backup vault access policy in the new account to allow access to the primary account.
B. Edit the backup vault access policy in the primary account to allow access to the new account.
C. Edit the backup vault access policy in the primary account to allow access to the KMS key in the new account.
D. Edit the key policy of the KMS key in the primary account to share the key with the new account.
E. Edit the key policy of the KMS key in the new account to share the key with the primary account.

Answer: AE

Explanation:
To enable cross-account backup, the company needs to grant permissions to both the backup vault and the KMS key in the destination account. The backup vault access policy in the destination account must allow the primary account to copy backups into the vault. The key policy of the KMS key in the destination account must allow the primary account to use the key to encrypt and decrypt the backups. These steps are described in the AWS documentation12. Therefore, the correct answer is A and E.
References:
? 1: Creating backup copies across AWS accounts - AWS Backup
? 2: Using AWS Backup with AWS Organizations - AWS Backup

NEW QUESTION 67
A company is hosting a web application in an AWS Region. For disaster recovery purposes, a second region is being used as a standby. Disaster recovery requirements state that session data must be replicated between regions in near-real time and 1% of requests should route to the secondary region to continuously verify system functionality. Additionally, if there is a disruption in service in the main region, traffic should be automatically routed to the secondary region, and the secondary region must be able to
scale up to handle all traffic.
How should a DevOps engineer meet these requirements?

A. In both regions, deploy the application on AWS Elastic Beanstalk and use Amazon DynamoDB global tables for session dat
B. Use an Amazon Route 53 weighted routing policy with health checks to distribute the traffic across the regions.
C. In both regions, launch the application in Auto Scaling groups and use DynamoDB for session dat

D. Use a Route 53 failover routing policy with health checks to distribute the traffic across the regions.
E. In both regions, deploy the application in AWS Lambda, exposed by Amazon API Gateway, and use Amazon RDS for PostgreSQL with cross-region replication for session dat
F. Deploy the web application with client-side logic to call the API Gateway directly.
G. In both regions, launch the application in Auto Scaling groups and use DynamoDB global tables for session dat
H. Enable an Amazon CloudFront weighted distribution across region
I. Point the Amazon Route 53 DNS record at the CloudFront distribution.

Answer: D

NEW QUESTION 68
A company uses AWS CodeArtifact to centrally store Python packages. The CodeArtifact repository is configured with the following repository policy.

```json
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "codeartifact:DescribePackageVersion",
                "codeartifact:DescribeRepository",
                "codeartifact:GetPackageVersionReadme",
                "codeartifact:GetRepositoryEndpoint",
                "codeartifact:ListPackageVersionAssets",
                "codeartifact:ListPackageVersionDependencies",
                "codeartifact:ListPackageVersions",
                "codeartifact:ListPackages",
                "codeartifact:ReadFromRepository"
            ],
            "Effect": "Allow",
            "Resource": "*",
            "Principal": "*",
            "Condition": {
                "StringEquals": {
                    "aws:PrincipalOrgID": [
                        "o-xxxxxxxxxxx"
                    ]
                }
            }
        }
    ]
}
```

A development team is building a new project in an account that is in an organization in AWS Organizations. The development team wants to use a Python library that has already been stored in the CodeArtifact repository in the organization. The development team uses AWS CodePipeline and AWS CodeBuild to build the new application. The CodeBuild job that the development team uses to build the application is configured to run in a VPC Because of compliance requirements the VPC has no internet connectivity.
The development team creates the VPC endpoints for CodeArtifact and updates the CodeBuild buildspec yaml file. However, the development team cannot download the Python library from the repository.
Which combination of steps should a DevOps engineer take so that the development team can use Code Artifact? (Select TWO.)

A. Create an Amazon S3 gateway endpoint Update the route tables for the subnets thatare running the CodeBuild job.
B. Update the repository policy's Principal statement to include the ARN of the role that the CodeBuild project uses.
C. Share the CodeArtifact repository with the organization by using AWS Resource Access Manager (AWS RAM).
D. Update the role that the CodeBuild project uses so that the role has sufficient permissions to use the CodeArtifact repository.
E. Specify the account that hosts the repository as the delegated administrator for CodeArtifact in the organization.

Answer: AD

Explanation:
"AWS CodeArtifact operates in multiple Availability Zones and stores artifact data and metadata in Amazon S3 and Amazon DynamoDB. Your encrypted data is redundantly stored across multiple facilities and multiple devices in each facility, making it highly available and highly durable."
https://aws.amazon.com/codeartifact/features/ With no internet connectivity, a gateway endpoint becomes necessary to access S3.

NEW QUESTION 73
A company manages multiple AWS accounts in AWS Organizations. The company's security policy states that AWS account root user credentials for member accounts must not be used. The company monitors access to the root user credentials.
A recent alert shows that the root user in a member account launched an Amazon EC2 instance. A DevOps engineer must create an SCP at the organization's root level that will prevent the root user in member accounts from making any AWS service API calls.
Which SCP will meet these requirements?
A)

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "*",
            "Resource": "*",
            "Condition": {
                "StringNotLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
            }
        }
    ]
}
```

B)
```
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Deny",
            "Action": "*",
            "Resource": "*",
            "Principal": { "AWS": "arn:aws:iam::*:root" }
        }
    ]
}
```

C)
```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Deny",
            "Action": "*",
            "Resource": "*",
            "Condition": {
                "StringLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
            }
        }
    ]
}
```

D)
```
    {
        "Version": "2012-10-17",
        "Statement": [
            {
                "Effect": "Allow",
                "Action": "*",
                "Resource": "*",
                "Principal": "root"
            }
        ]
    }
```

A. Option A
B. Option B
C. Option C
D. Option D

Answer: D

NEW QUESTION 76
An ecommerce company has chosen AWS to host its new platform. The company's DevOps team has started building an AWS Control Tower landing zone. The DevOps team has set the identity store within AWS IAM Identity Center (AWS Single Sign-On) to external identity provider (IdP) and has configured SAML 2.0. The DevOps team wants a robust permission model that applies the principle of least privilege. The model must allow the team to build and manage only the team's own resources.
Which combination of steps will meet these requirements? (Choose three.)

A. Create IAM policies that include the required permission
B. Include the aws:PrincipalTag condition key.
C. Create permission set
D. Attach an inline policy that includes the required permissions and uses the aws:PrincipalTag condition key to scope the permissions.
E. Create a group in the Id
F. Place users in the grou
G. Assign the group to accounts and the permission sets in IAM Identity Center.
H. Create a group in the Id
I. Place users in the grou
J. Assign the group to OUs and IAM policies.
K. Enable attributes for access control in IAM Identity Cente
L. Apply tags to user
M. Map the tags as key-value pairs.
N. Enable attributes for access control in IAM Identity Cente
O. Map attributes from the IdP as key-value pairs.

Answer: BCF

Explanation:
Using the principalTag in the Permission Set inline policy a logged in user belonging to a specific AD group in the IDP can be permitted access to perform operations on certain resources if their group matches the group used in the PrincipleTag. Basically you are narrowing the scope of privileges assigned via Permission policies conditionally based on whether the logged in user belongs to a specific AD Group in IDP. The mapping of the AD group to the request attributes can be done using SSO attributes where we can pass other attributes like the SAML token as well.
https://docs.aws.amazon.com/singlesignon/latest/userguide/abac.html

NEW QUESTION 77
An online retail company based in the United States plans to expand its operations to Europe and Asia in the next six months. Its product currently runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. All data is stored in an Amazon Aurora database instance.
When the product is deployed in multiple regions, the company wants a single product catalog across all regions, but for compliance purposes, its customer information and purchases must be kept in each region.
How should the company meet these requirements with the LEAST amount of application changes?

A. Use Amazon Redshift for the product catalog and Amazon DynamoDB tables for the customer information and purchases.
B. Use Amazon DynamoDB global tables for the product catalog and regional tables for the customer information and purchases.
C. Use Aurora with read replicas for the product catalog and additional local Aurora instances in each region for the customer information and purchases.
D. Use Aurora for the product catalog and Amazon DynamoDB global tables for the customer information and purchases.

Answer: C

NEW QUESTION 80
A DevOps engineer is architecting a continuous development strategy for a company's software as a service (SaaS) web application running on AWS. For application and security reasons users subscribing to this application are distributed across multiple. Application Load Balancers (ALBs) each of which has a dedicated Auto Scaling group and fleet of Amazon EC2 instances The application does not require a build stage and when it is committed to AWS CodeCommit, the application must trigger a simultaneous deployment to all ALBs Auto Scaling groups and EC2 fleets.
Which architecture will meet these requirements with the LEAST amount of configuration?

A. Create a single AWS CodePipeline pipeline that deploys the application in parallel using unique AWS CodeDeploy applications and deployment groups created for each ALB-Auto Scaling group pair.
B. Create a single AWS CodePipeline pipeline that deploys the application using a single AWS CodeDeploy application and single deployment group.
C. Create a single AWS CodePipeline pipeline that deploys the application in parallel using a single AWS CodeDeploy application and unique deployment group for each ALB-Auto Scaling group pair.
D. Create an AWS CodePipeline pipeline for each ALB-Auto Scaling group pair that deploys the application using an AWS CodeDeploy application and deployment group created for the same ALB-Auto Scaling group pair.

Answer: C

Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment- groups.html

NEW QUESTION 81
A company uses AWS Secrets Manager to store a set of sensitive API keys that an AWS Lambda function uses. When the Lambda function is invoked, the Lambda function retrieves the API keys and makes an API call to an external service. The Secrets Manager secret is encrypted with the default AWS Key Management Service (AWS KMS) key.
A DevOps engineer needs to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege.
Which combination of steps will meet these requirements? (Select TWO.)

A. Update the default KMS key for Secrets Manager to allow only the Lambda function's execution role to decrypt.
B. Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function's execution role to decryp
C. Update Secrets Manager to use the new customer managed key.
D. Create a KMS customer managed key that trusts Secrets Manager and allows the account's :root principal to decryp
E. Update Secrets Manager to use the new customer managed key.

F. Ensure that the Lambda function's execution role has the KMS permissions scoped on the resource leve
G. Configure the permissions so that the KMS key can encrypt the Secrets Manager secret.
H. Remove all KMS permissions from the Lambda function's execution role.

Answer: BD

Explanation:
 The requirement is to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege, which means granting the minimum permissions necessary to perform a task.
To do this, the DevOps engineer needs to use the following steps:
? Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function's execution role to decrypt. A customer managed key is a symmetric encryption key that is fully managed by the customer. The customer can define the key policy, which specifies who can use and manage the key. By creating a customer managed key, the DevOps engineer can restrict the decryption permission to only the Lambda function's execution role, and prevent other principals from accessing the secret values. The customer managed key also needs to trust Secrets Manager, which means allowing Secrets Manager to use the key to encrypt and decrypt secrets on behalf of the customer.
? Update Secrets Manager to use the new customer managed key. Secrets Manager allows customers to choose which KMS key to use for encrypting each secret. By default, Secrets Manager uses the default KMS key for Secrets Manager, which is a service-managed key that is shared by all customers in the same AWS Region. By updating Secrets Manager to use the new customer managed key, the DevOps engineer can ensure that only the Lambda function's execution role can decrypt the secret values using that key.
? Ensure that the Lambda function's execution role has the KMS permissions scoped on the resource level. The Lambda function's execution role is an IAM role that grants permissions to the Lambda function to access AWS services and resources. The role needs to have KMS permissions to use the customer managed key for decryption. However, to apply the principle of least privilege, the role should have the permissions scoped on the resource level, which means specifying the ARN of the customer managed key as a condition in the IAM policy statement. This way, the role can only use that specific key and not any other KMS keys in the account.

NEW QUESTION 85
A company updated the AWS Cloud Formation template for a critical business application. The stack update process failed due to an error in the updated template and AWS CloudFormation automatically began the stack rollback process Later a DevOps engineer discovered that the application was still unavailable and that the stack was in the UPDATE_ROLLBACK_FAILED state.
Which combination of actions should the DevOps engineer perform so that the stack rollback can complete successfully? (Select TWO.)

A. Attach the AWSC loud Formation FullAccess IAM policy to the AWS CtoudFormation role.
B. Automatically recover the stack resources by using AWS CloudFormation drift detection.
C. Issue a ContinueUpdateRollback command from the AWS CloudFormation console or the AWS CLI.
D. Manually adjust the resources to match the expectations of the stack.
E. Update the existing AWS CloudFormation stack by using the original template.

Answer: CD

Explanation:
 https://docs.aws.amazon.com/cli/latest/reference/cloudformation/continue- update-rollback.html For a specified stack that is in the UPDATE_ROLLBACK_FAILED state, continues rolling it back to the UPDATE_ROLLBACK_COMPLETE state. Depending on the cause of the failure, you can manually fix the error and continue the rollback. By continuing the rollback, you can return your stack to a working state (the UPDATE_ROLLBACK_COMPLETE state), and then try to update the stack again.

NEW QUESTION 89
A company uses AWS Organizations to manage multiple accounts. Information security policies require that all unencrypted Amazon EBS volumes be marked as non-compliant. A DevOps engineer needs to automatically deploy the solution and ensure that this compliance check is always present.
Which solution will accomplish this?

A. Create an AWS CloudFormation template that defines an AWS Inspector rule to check whether EBS encryption is enable
B. Save the template to an Amazon S3 bucket that has been shared with all accounts within the compan
C. Update the account creation script pointing to the CloudFormation template in Amazon S3.
D. Create an AWS Config organizational rule to check whether EBS encryption is enabled and deploy the rule using the AWS CL
E. Create and apply an SCP to prohibit stopping and deleting AWS Config across the organization.
F. Create an SCP in Organization
G. Set the policy to prevent the launch of Amazon EC2 instances without encryption on the EBS volumes using a conditional expressio
H. Apply the SCP to all AWS account
I. Use Amazon Athena to analyze the AWS CloudTrail output, looking for events that deny an ec2: RunInstances action.
J. Deploy an IAM role to all accounts from a single trusted accoun
K. Build a pipeline withAWS CodePipeline with a stage in AWS Lambda to assume the IAM role, and list all EBS volumes in the accoun
L. Publish a report to Amazon S3.

Answer: B

Explanation:
 https://docs.aws.amazon.com/config/latest/developerguide/ec2-ebs-encryption-by-default.html

NEW QUESTION 92
A DevOps engineer is building a multistage pipeline with AWS CodePipeline to build, verify, stage, test, and deploy an application. A manual approval stage is required between the test stage and the deploy stage. The development team uses a custom chat tool with webhook support that requires near-real-time notifications.
How should the DevOps engineer configure status updates for pipeline activity and approval requests to post to the chat tool?

A. Create an Amazon CloudWatch Logs subscription that filters on CodePipeline Pipeline Execution State Chang
B. Publish subscription events to an Amazon Simple Notification Service (Amazon SNS) topi
C. Subscribe the chat webhook URL to the SNS topic, and complete the subscription validation.
D. Create an AWS Lambda function that is invoked by AWS CloudTrail event
E. When a CodePipeline Pipeline Execution State Change event is detected, send the event details to the chat webhook URL.
F. Create an Amazon EventBridge rule that filters on CodePipeline Pipeline Execution State Chang

G. Publish the events to an Amazon Simple Notification Service (Amazon SNS) topi
H. Create an AWS Lambda function that sends event details to the chat webhook UR
I. Subscribe the function to the SNS topic.
J. Modify the pipeline code to send the event details to the chat webhook URL at the end of each stag
K. Parameterize the URL so that each pipeline can send to a different URL based on the pipeline environment.

Answer: C

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/sns-lambda-webhooks-chime-slack-teams/

NEW QUESTION 95
A company is divided into teams Each team has an AWS account and all the accounts are in an organization in AWS Organizations. Each team must retain full administrative rights to its AWS account. Each team also must be allowed to access only AWS services that the company approves for use AWS services must gam approval through a request and approval process.
How should a DevOps engineer configure the accounts to meet these requirements?

A. Use AWS CloudFormation StackSets to provision IAM policies in each account to deny access to restricted AWS service
B. In each account configure AWS Config rules that ensure that the policies are attached to IAM principals in the account.
C. Use AWS Control Tower to provision the accounts into OUs within the organization Configure AWS Control Tower to enable AWS IAM identity Center (AWS Single Sign-On). Configure 1AM Identity Center to provide administrative access Include deny policies on user roles for restricted AWS services.
D. Place all the accounts under a new top-level OU within the organization Create an SCP that denies access to restricted AWS services Attach the SCP to the OU.
E. Create an SCP that allows access to only approved AWS service
F. Attach the SCP to the root OU of the organizatio
G. Remove the FullAWSAccess SCP from the root OU of the organization.

Answer: C

Explanation:
https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html A managed prefix list is a set of one or more CIDR blocks. You can use prefix lists to make it easier to configure and maintain your security groups and route tables. https://docs.aws.amazon.com/vpc/latest/userguide/sharing-managed-prefix-lists.html With AWS Resource Access Manager (AWS RAM), the owner of a prefix list can share a prefix list with the following: Specific AWS accounts inside or outside of its organization in AWS Organizations An organizational unit inside its organization in AWS Organizations An entire organization in AWS Organizations

NEW QUESTION 98
A company manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The EC2 instances run in an Auto Scaling group across multiple Availability Zones. The application uses an Amazon RDS for MySQL DB instance to store the data. The company has configured Amazon Route 53 with an alias record that points to the ALB.
A new company guideline requires a geographically isolated disaster recovery (DR> site with an RTO of 4 hours and an RPO of 15 minutes.
Which DR strategy will meet these requirements with the LEAST change to the application stack?

A. Launch a replica environment of everything except Amazon RDS in a different Availability Zone Create an RDS read replica in the new Availability Zone: and configure the new stack to point to the local RDS DB instanc
B. Add the new stack to the Route 53 record set by using a hearth check to configure a failover routing policy.
C. Launch a replica environment of everything except Amazon RDS in a different AW
D. Region Create an RDS read replica in the new Region and configure the new stack to point to the local RDS DB instanc
E. Add the new stack to the Route 53 record set by using a health check to configure a latency routing policy.
F. Launch a replica environment of everything except Amazon RDS ma different AWS Regio
G. In the event of an outage copy and restore the latest RDS snapshot from the primar
H. Region to the DR Region Adjust the Route 53 record set to point to the ALB in the DR Region.
I. Launch a replica environment of everything except Amazon RDS in a different AWS Regio
J. Create an RDS read replica in the new Region and configure the new environment to point to the local RDS DB instanc
K. Add the new stack to the Route 53 record set by using a health check to configure a failover routing polic
L. In the event of an outage promote the read replica to primary.

Answer: D

NEW QUESTION 99
A company's security team requires that all external Application Load Balancers (ALBs) and Amazon API Gateway APIs are associated with AWS WAF web ACLs. The company
has hundreds of AWS accounts, all of which are included in a single organization in AWS Organizations. The company has configured AWS Config for the organization. During an audit, the company finds some externally facing ALBs that are not associated with AWS WAF web ACLs.
Which combination of steps should a DevOps engineer take to prevent future violations? (Choose two.)

A. Delegate AWS Firewall Manager to a security account.
B. Delegate Amazon GuardDuty to a security account.
C. Create an AWS Firewall Manager policy to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.
D. Create an Amazon GuardDuty policy to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.
E. Configure an AWS Config managed rule to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.

Answer: AC

Explanation:
If instead you want to automatically apply the policy to existing in-scope resources, choose Auto remediate any noncompliant resources. This option creates a web ACL in each applicable account within the AWS organization and associates the web ACL with the resources in the accounts. When you choose Auto remediate any noncompliant resources, you can also choose to remove existing web ACL associations from in-scope resources, for the web ACLs that aren't managed by another active Firewall Manager policy. If you choose this option, Firewall Manager first associates the policy's web ACL with the resources, and then removes the prior associations. If a resource has an association with another web ACL that's managed by a different active Firewall Manager policy, this choice doesn't affect that association.

NEW QUESTION 101

A company plans to use Amazon CloudWatch to monitor its Amazon EC2 instances. The company needs to stop EC2 instances when the average of the NetworkPacketsIn metric is less than 5 for at least 3 hours in a 12-hour time window. The company must evaluate the metric every hour. The EC2 instances must continue to run if there is missing data for the NetworkPacketsIn metric during the evaluation period.

A DevOps engineer creates a CloudWatch alarm for the NetworkPacketsIn metric. The DevOps engineer configures a threshold value of 5 and an evaluation period of 1 hour.

Which set of additional actions should the DevOps engineer take to meet these requirements?

A. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as breaching the threshold
B. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.
C. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as not breaching the threshol
D. Add an EC2 action to stop the instance when the alarm enters the ALARM state.
E. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as breaching the threshold
F. Add an EC2 action to stop the instance when the alarm enters the ALARM state.
G. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as not breaching the threshol
H. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.

Answer: B

Explanation:
To meet the requirements, the DevOps engineer needs to configure the CloudWatch alarm to stop the EC2 instances when the average of the NetworkPacketsIn metric is less than 5 for at least 3 hours in a 12-hour time window. This means that the alarm should trigger when 3 out of 12 datapoints are below the threshold of 5. The alarm should also treat missing data as not breaching the threshold, so that the EC2 instances continue to run if there is no data for the metric during the evaluation period. The DevOps engineer can add an EC2 action to stop the instance when the alarm enters the ALARM state, which is a built-in action type for CloudWatch alarms.

NEW QUESTION 104

A company runs an application on Amazon EC2 instances. The company uses a series of AWS CloudFormation stacks to define the application resources. A developer performs updates by building and testing the application on a laptop and then uploading the build output and CloudFormation stack templates to Amazon S3. The developer's peers review the changes before the developer performs the CloudFormation stack update and installs a new version of the application onto the EC2 instances.

The deployment process is prone to errors and is time-consuming when the developer updates each EC2 instance with the new application. The company wants to automate as much of the application deployment process as possible while retaining a final manual approval step before the modification of the application or resources.

The company already has moved the source code for the application and the CloudFormation templates to AWS CodeCommit. The company also has created an AWS CodeBuild project to build and test the application.

Which combination of steps will meet the company's requirements? (Choose two.)

A. Create an application group and a deployment group in AWS CodeDeplo
B. Install the CodeDeploy agent on the EC2 instances.
C. Create an application revision and a deployment group in AWS CodeDeplo
D. Create an environment in CodeDeplo
E. Register the EC2 instances to the CodeDeploy environment.
F. Use AWS CodePipeline to invoke the CodeBuild job, run the CloudFormation update, and pause for a manual approval ste
G. After approval, start the AWS CodeDeploy deployment.
H. Use AWS CodePipeline to invoke the CodeBuild job, create CloudFormation change sets for each of the application stacks, and pause for a manual approval ste
I. After approval, run the CloudFormation change sets and start the AWS CodeDeploy deployment.
J. Use AWS CodePipeline to invoke the CodeBuild job, create CloudFormation change sets for each of the application stacks, and pause for a manual approval ste
K. After approval, start the AWS CodeDeploy deployment.

Answer: AD

Explanation:
 A- https://docs.aws.amazon.com/codedeploy/latest/userguide/codedeploy- agent.html D - This option correctly utilizes AWS CodePipeline to invoke the CodeBuild job and create CloudFormation change sets. It adds a manual approval step before executing the change sets and starting the AWS CodeDeploy deployment. This ensures that the deployment process is automated while retaining the final manual approval step.

NEW QUESTION 109

A DevOps engineer is planning to deploy a Ruby-based application to production. The application needs to interact with an Amazon RDS for MySQL database and should have automatic scaling and high availability. The stored data in the database is critical and should persist regardless of the state of the application stack. The DevOps engineer needs to set up an automated deployment strategy for the application with automatic rollbacks. The solution also must alert the application team when a deployment fails.

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

A. Deploy the application on AWS Elastic Beanstal
B. Deploy an Amazon RDS for MySQL DB instance as part of the Elastic Beanstalk configuration.
C. Deploy the application on AWS Elastic Beanstal
D. Deploy a separate Amazon RDS for MySQL DB instance outside of Elastic Beanstalk.
E. Configure a notification email address that alerts the application team in the AWS Elastic Beanstalk configuration.
F. Configure an Amazon EventBridge rule to monitor AWS Health event
G. Use an Amazon Simple Notification Service (Amazon SNS) topic as a target to alert the application team.
H. Use the immutable deployment method to deploy new application versions.
I. Use the rolling deployment method to deploy new application versions.

Answer: BDE

Explanation:
For deploying a Ruby-based application with requirements for interaction with an Amazon RDS for MySQL database, automatic scaling, high availability, and data

persistence, the following steps will meet the requirements:
? B. Deploy the application on AWS Elastic Beanstalk. Deploy a separate Amazon
RDS for MySQL DB instance outside of Elastic Beanstalk. This approach ensures that the database persists independently of the Elastic Beanstalk environment, which can be torn down and recreated without affecting the database123.
? E. Use the immutable deployment method to deploy new application
versions. Immutable deployments provide a zero-downtime deployment method that ensures that if any part of the deployment process fails, the environment is rolled back to the original state automatically4.
? D. Configure an Amazon EventBridge rule to monitor AWS Health events. Use an
Amazon Simple Notification Service (Amazon SNS) topic as a target to alert the application team. This setup allows for automated monitoring and alerting of the application team in case of deployment failures or other health events56.
References:
? AWS Elastic Beanstalk documentation on deploying Ruby applications1.
? AWS documentation on application auto-scaling7.
? AWS documentation on automated deployment strategies with automatic rollbacks and alerts456.

NEW QUESTION 110
A company builds a container image in an AWS CodeBuild project by running Docker commands. After the container image is built, the CodeBuild project uploads the container image to an Amazon S3 bucket. The CodeBuild project has an 1AM service role that has permissions to access the S3 bucket.
A DevOps engineer needs to replace the S3 bucket with an Amazon Elastic Container Registry (Amazon ECR) repository to store the container images. The DevOps engineer creates an ECR private image repository in the same AWS Region of the CodeBuild project. The DevOps engineer adjusts the 1AM service role with the permissions that are necessary to work with the new ECR repository. The DevOps engineer also places new repository information into the docker build command and the docker push command that are used in the buildspec.yml file.
When the CodeBuild project runs a build job, the job fails when the job tries to access the ECR repository.
Which solution will resolve the issue of failed access to the ECR repository?

A. Update the buildspec.yml file to log in to the ECR repository by using the aws ecr get- login-password AWS CLI command to obtain an authentication toke
B. Update the docker login command to use the authentication token to access the ECR repository.
C. Add an environment variable of type SECRETS_MANAGER to the CodeBuild projec
D. In the environment variable, include the ARN of the CodeBuild project's IAM service rol
E. Update the buildspec.yml file to use the new environment variable to log in with the docker login command to access the ECR repository.
F. Update the ECR repository to be a public image repositor
G. Add an ECR repository policy that allows the 1AM service role to have access.
H. Update the buildspec.yml file to use the AWS CLI to assume the 1AM service role for ECR operation
I. Add an ECR repository policy that allows the 1AM service role to have access.

Answer: A

Explanation:
 (A) When Docker communicates with an Amazon Elastic Container Registry (ECR) repository, it requires authentication. You can authenticate your Docker client to the Amazon ECR registry with the help of the AWS CLI (Command Line Interface). Specifically, you can use the "aws ecr get-login-password" command to get an authorization token and then use Docker's "docker login" command with that token to authenticate to the registry. You would need to perform these steps in your buildspec.yml file before attempting to push or pull images from/to the ECR repository.

NEW QUESTION 115
A company uses AWS and has a VPC that contains critical compute infrastructure with predictable traffic patterns. The company has configured VPC flow logs that are published to a log group in Amazon CloudWatch Logs.
The company's DevOps team needs to configure a monitoring solution for the VPC flow logs to identify anomalies in network traffic to the VPC over time. If the monitoring solution detects an anomaly, the company needs the ability to initiate a response to the anomaly.
How should the DevOps team configure the monitoring solution to meet these requirements?

A. Create an Amazon Kinesis data strea
B. Subscribe the log group to the data strea
C. Configure Amazon Kinesis Data Analytics to detect log anomalies in the data strea
D. Create anAWS Lambda function to use as the output of the data strea
E. Configure the Lambda function to write to the default Amazon EventBridge event bus in the event of an anomaly finding.
F. Create an Amazon Kinesis Data Firehose delivery stream that delivers events to an Amazon S3 bucke
G. Subscribe the log group to the delivery strea
H. Configure Amazon Lookout for Metrics to monitor the data in the S3 bucket for anomalie
I. Create an AWS Lambda function to run in response to Lookout for Metrics anomaly finding
J. Configure the Lambda function to publish to the default Amazon EventBridge event bus.
K. Create an AWS Lambda function to detect anomalie
L. Configure the Lambda function to publish an event to the default Amazon EventBridge event bus if the Lambda function detects an anomal
M. Subscribe the Lambda function to the log group.
N. Create an Amazon Kinesis data strea
O. Subscribe the log group to the data strea
P. Create an AWS Lambda function to detect log anomalie
Q. Configure the Lambda function to write to the default Amazon EventBridge event bus if the Lambda function detects an anomal
R. Set the Lambda function as the processor for the data stream.

Answer: D

Explanation:
To meet the requirements, the DevOps team needs to configure a monitoring solution for the VPC flow logs that can detect anomalies in network traffic over time and initiate a response to the anomaly. The DevOps team can use Amazon Kinesis Data Streams to ingest and process streaming data from CloudWatch Logs. The DevOps team can subscribe the log group to a Kinesis data stream, which will deliver log events from CloudWatch Logs to Kinesis Data Streams in near real-time. The DevOps team can then create an AWS Lambda function to detect log anomalies using machine learning or statistical methods. The Lambda function can be set as a processor for the data stream, which means that it will process each record from the stream before sending it to downstream applications or destinations. The Lambda function can also write to the default Amazon EventBridge event bus if it detects an anomaly, which will allow other AWS services or custom applications to respond to the anomaly event.

NEW QUESTION 119

A company needs a strategy for failover and disaster recovery of its data and application. The application uses a MySQL database and Amazon EC2 instances. The company requires a maximum RPO of 2 hours and a maximum RTO of 10 minutes for its data and application at all times.

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

A. Create an Amazon Aurora Single-AZ cluster in multiple AWS Regions as the data stor
B. Use Aurora's automatic recovery capabilities in the event of a disaster.
C. Create an Amazon Aurora global database in two AWS Regions as the data stor
D. In the event of a failure, promote the secondary Region to the primary for the applicatio
E. Update the application to use the Aurora cluster endpoint in the secondary Region.
F. Create an Amazon Aurora cluster in multiple AWS Regions as the data stor
G. Use a Network Load Balancer to balance the database traffic in different Regions.
H. Set up the application in two AWS Region
I. Use Amazon Route 53 failover routing that points to Application Load Balancers in both Region
J. Use health checks and Auto Scaling groups in each Region.
K. Set up the application in two AWS Region
L. Configure AWS Global Accelerator to point to Application Load Balancers (ALBs) in both Region
M. Add both ALBs to a single endpoint grou
N. Use health checks and Auto Scaling groups in each Region.

Answer: BE

Explanation:

To meet the requirements of failover and disaster recovery, the company should use the following deployment strategies:

? Create an Amazon Aurora global database in two AWS Regions as the data store.

In the event of a failure, promote the secondary Region to the primary for the application. Update the application to use the Aurora cluster endpoint in the secondary Region. This strategy can provide a low RPO and RTO for the data, as Aurora global database replicates data with minimal latency across Regions and allows fast and easy failover12. The company can use the Amazon Aurora cluster endpoint to connect to the current primary DB cluster without needing to change any application code1.

? Set up the application in two AWS Regions. Configure AWS Global Accelerator to

point to Application Load Balancers (ALBs) in both Regions. Add both ALBs to a single endpoint group. Use health checks and Auto Scaling groups in each Region. This strategy can provide high availability and performance for the application, as AWS Global Accelerator uses the AWS global network to route traffic to the closest healthy endpoint3. The company can also use static IP addresses that are assigned by Global Accelerator as a fixed entry point for their application1. By using health checks and Auto Scaling groups, the company can ensure that their application can scale up or down based on demand and handle any instance failures4.

The other options are incorrect because:

? Creating an Amazon Aurora Single-AZ cluster in multiple AWS Regions as the data store would not provide a fast failover or disaster recovery solution, as the company would need to manually restore data from backups or snapshots in another Region in case of a failure.

? Creating an Amazon Aurora cluster in multiple AWS Regions as the data store and using a Network Load Balancer to balance the database traffic in different Regions would not work, as Network Load Balancers do not support cross-Region routing. Moreover, this strategy would not provide a consistent view of the data across Regions, as Aurora clusters do not replicate data automatically between Regions unless they are part of a global database.

? Setting up the application in two AWS Regions and using Amazon Route 53 failover routing that points to Application Load Balancers in both Regions would not provide a low RTO, as Route 53 failover routing relies on DNS resolution, which can take time to propagate changes across different DNS servers and clients. Moreover, this strategy would not provide deterministic routing, as Route 53 failover routing depends on DNS caching behavior, which can vary depending on different factors.

NEW QUESTION 120

A company is testing a web application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company uses a blue green deployment process with immutable instances when deploying new software.

During testing users are being automatically logged out of the application at random times. Testers also report that when a new version of the application is deployed all users are logged out. The development team needs a solution to ensure users remain logged m across scaling events and application deployments.

What is the MOST operationally efficient way to ensure users remain logged in?

A. Enable smart sessions on the load balancer and modify the application to check tor an existing session.
B. Enable session sharing on the toad balancer and modify the application to read from the session store.
C. Store user session information in an Amazon S3 bucket and modify the application to read session information from the bucket.
D. Modify the application to store user session information in an Amazon ElastiCache cluster.

Answer: D

Explanation:

https://aws.amazon.com/caching/session-management/

NEW QUESTION 125

A company is developing an application that will generate log events. The log events consist of five distinct metrics every one tenth of a second and produce a large amount of data The company needs to configure the application to write the logs to Amazon Time stream The company will configure a daily query against the Timestream table.

Which combination of steps will meet these requirements with the FASTEST query performance? (Select THREE.)

A. Use batch writes to write multiple log events in a Single write operation
B. Write each log event as a single write operation
C. Treat each log as a single-measure record
D. Treat each log as a multi-measure record
E. Configure the memory store retention period to be longer than the magnetic store retention period
F. Configure the memory store retention period to be shorter than the magnetic store retention period

Answer: ADF

Explanation:

A comprehensive and detailed explanation is:

? Option A is correct because using batch writes to write multiple log events in a single write operation is a recommended practice for optimizing the performance

and cost of data ingestion in Timestream. Batch writes can reduce the number of network round trips and API calls, and can also take advantage of parallel processing by Timestream. Batch writes can also improve the compression ratio of data in the memory store and the magnetic store, which can reduce the storage costs and improve the query performance1.

? Option B is incorrect because writing each log event as a single write operation is not a recommended practice for optimizing the performance and cost of data ingestion in Timestream. Writing each log event as a single write operation would increase the number of network round trips and API calls, and would also reduce the compression ratio of data in the memory store and the magnetic store. This would increase the storage costs and degrade the query performance1.

? Option C is incorrect because treating each log as a single-measure record is not a recommended practice for optimizing the query performance in Timestream. Treating each log as a single-measure record would result in creating multiple records for each timestamp, which would increase the storage size and the query latency. Moreover, treating each log as a single-measure record would require using joins to query multiple measures for the same timestamp, which would add complexity and overhead to the query processing2.

? Option D is correct because treating each log as a multi-measure record is a recommended practice for optimizing the query performance in Timestream. Treating each log as a multi-measure record would result in creating a single record for each timestamp, which would reduce the storage size and the query latency. Moreover, treating each log as a multi-measure record would allow querying multiple measures for the same timestamp without using joins, which would simplify and speed up the query processing2.

? Option E is incorrect because configuring the memory store retention period to be longer than the magnetic store retention period is not a valid option in Timestream. The memory store retention period must always be shorter than or equal to the magnetic store retention period. This ensures that data is moved from the memory store to the magnetic store before it expires out of the memory store3.

? Option F is correct because configuring the memory store retention period to be shorter than the magnetic store retention period is a valid option in Timestream. The memory store retention period determines how long data is kept in the memory store, which is optimized for fast point-in-time queries. The magnetic store retention period determines how long data is kept in the magnetic store, which is optimized for fast analytical queries. By configuring these retention periods appropriately, you can balance your storage costs and query performance according to your application needs3.

References:
? 1: Batch writes
? 2: Multi-measure records vs. single-measure records
? 3: Storage

NEW QUESTION 128
A company has an application that runs on a fleet of Amazon EC2 instances. The application requires frequent restarts. The application logs contain error messages when a restart is required. The application logs are published to a log group in Amazon CloudWatch Logs.
An Amazon CloudWatch alarm notifies an application engineer through an Amazon Simple Notification Service (Amazon SNS) topic when the logs contain a large number of restart- related error messages. The application engineer manually restarts the application on the instances after the application engineer receives a notification from the SNS topic.
A DevOps engineer needs to implement a solution to automate the application restart on the instances without restarting the instances.
Which solution will meet these requirements in the MOST operationally efficient manner?

A. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instance
B. Configure the SNS topic to invoke the runbook.
C. Create an AWS Lambda function that restarts the application on the instance
D. Configure the Lambda function as an event destination of the SNS topic.
E. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instance
F. Create an AWS Lambda function to invoke the runboo
G. Configure the Lambda function as an event destination of the SNS topic.
H. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instance
I. Configure an Amazon EventBridge rule that reacts when the CloudWatch alarm enters ALARM stat
J. Specify the runbook as a target of the rule.

Answer: D

Explanation:
This solution meets the requirements in the most operationally efficient manner by automating the application restart process on the instances without restarting them. When the CloudWatch alarm enters the ALARM state, the EventBridge rule is triggered, which in turn invokes the Systems Manager Automation runbook that contains the script to restart the application on the instances.

NEW QUESTION 131
A company has an AWS CodePipeline pipeline that is configured with an Amazon S3 bucket in the eu-west-1 Region. The pipeline deploys an AWS Lambda application to the same Region. The pipeline consists of an AWS CodeBuild project build action and an AWS CloudFormation deploy action.
The CodeBuild project uses the aws cloudformation package AWS CLI command to build an artifact that contains the Lambda function code's .zip file and the CloudFormation template. The CloudFormation deploy action references the CloudFormation template from the output artifact of the CodeBuild project's build action.
The company wants to also deploy the Lambda application to the us-east-1 Region by using the pipeline in eu-west-1. A DevOps engineer has already updated the CodeBuild project to use the aws cloudformation package command to produce an additional output artifact for us-east-1.
Which combination of additional steps should the DevOps engineer take to meet these requirements? (Choose two.)

A. Modify the CloudFormation template to include a parameter for the Lambda function code's zip file locatio
B. Create a new CloudFormation deploy action for us-east-1 in the pipelin
C. Configure the new deploy action to pass in the us-east-1 artifact location as a parameter override.
D. Create a new CloudFormation deploy action for us-east-1 in the pipelin
E. Configure the new deploy action to use the CloudFormation template from the us-east-1 output artifact.
F. Create an S3 bucket in us-east-1. Configure the S3 bucket policy to allow CodePipeline to have read and write access.
G. Create an S3 bucket in us-east-1. Configure S3 Cross-Region Replication (CRR) from the S3 bucket in eu-west-1 to the S3 bucket in us-east-1.
H. Modify the pipeline to include the S3 bucket for us-east-1 as an artifact stor
I. Create a new CloudFormation deploy action for us-east-1 in the pipelin
J. Configure the new deploy action to use the CloudFormation template from the us-east-1 output artifact.

Answer: AB

Explanation:
A. The CloudFormation template should be modified to include a parameter that indicates the location of the .zip file containing the Lambda function's code. This allows the CloudFormation deploy action to use the correct artifact depending on the region. This is critical because Lambda functions need to reference their code artifacts from the same region they are being deployed in. B. You would also need to create a new CloudFormation deploy action for the us-east-1 Region within the pipeline. This action should be configured to use the CloudFormation template from the artifact that was specifically created for us- east-1.

NEW QUESTION 134

A company recently created a new AWS Control Tower landing zone in a new organization in AWS Organizations. The landing zone must be able to demonstrate compliance with the Center tor Internet Security (CIS) Benchmarks tor AWS Foundations.

The company's security team wants to use AWS Security Hub to view compliance across all accounts Only the security team can be allowed to view aggregated Security Hub Findings. In addition specific users must be able to view findings from their own accounts within the organization All accounts must be enrolled m Security Hub after the accounts are created.

Which combination of steps will meet these requirements in the MOST automated way? (Select THREE.)

A. Turn on trusted access for Security Hub in the organization's management accoun
B. Create a new security account by using AWS Control Tower Configure the new security account as the delegated administrator account for Security Hu
C. In the new security account provid
D. Security Hub with the CIS Benchmarks for AWS Foundations standards.
E. Turn on trusted access for Security Hub in the organ ration's management accoun
F. From the management account, provide Security Hub with the CIS Benchmarks for AWS Foundations standards.
G. Create an AWS IAM identity Center (AWS Single Sign-On) permission set that includes the required permissions Use the CreateAccountAssignment API operation to associate the security team users with the permission set and with the delegated security account.
H. Create an SCP that explicitly denies any user who is not on the security team from accessing Security Hub.
I. In Security Hub, turn on automatic enablement.
J. In the organization's management account create an Amazon EventBridge rule that reacts to the CreateManagedAccount event Create an AWS Lambda function that uses the Security Hub CreateMembers API operation to add new accounts to Security Hu
K. Configure the EventBridge rule to invoke the Lambda function.

Answer: ACE

Explanation:
https://docs.aws.amazon.com/securityhub/latest/userguide/accounts-orgs- auto-enable.html

NEW QUESTION 135

A company uses AWS CloudFormation stacks to deploy updates to its application. The stacks consist of different resources. The resources include AWS Auto Scaling groups, Amazon EC2 instances, Application Load Balancers (ALBs), and other resources that are necessary to launch and maintain independent stacks. Changes to application resources outside of CloudFormation stack updates are not allowed.

The company recently attempted to update the application stack by using the AWS CLI. The stack failed to update and produced the following error message:
"ERROR: both the deployment and the CloudFormation stack rollback failed. The deployment failed because the following resource(s) failed to update: [AutoScalingGroup]."

The stack remains in a status of UPDATE_ROLLBACK_FAILED. * Which solution will resolve this issue?

A. Update the subnet mappings that are configured for the ALB
B. Run the aws cloudformation update-stack-set AWS CLI command.
C. Update the 1AM role by providing the necessary permissions to update the stac
D. Run the aws cloudformation continue-update-rollback AWS CLI command.
E. Submit a request for a quota increase for the number of EC2 instances for the accoun
F. Run the aws cloudformation cancel-update-stack AWS CLI command.
G. Delete the Auto Scaling group resourc
H. Run the aws cloudformation rollback-stack AWS CLI command.

Answer: B

Explanation:
https://repost.aws/knowledge-center/cloudformation-update-rollback-failed If your stack is stuck in the UPDATE_ROLLBACK_FAILED state after a failed update, then the only actions that you can perform on the stack are the ContinueUpdateRollback or DeleteStack operations.

NEW QUESTION 137
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DOP-C02 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DOP-C02 Product From:

https://www.2passeasy.com/dumps/DOP-C02/

Money Back Guarantee

DOP-C02 Practice Exam Features:

* DOP-C02 Questions and Answers Updated Frequently

* DOP-C02 Practice Questions Verified by Expert Senior Certified Staff

* DOP-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* DOP-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year