Amazon-Web-Services

Exam Questions DOP-C02

AWS Certified DevOps Engineer - Professional

NEW QUESTION 1

A DevOps engineer manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an EC2 Auto Scaling group across multiple Availability Zones. The engineer needs to implement a deployment strategy that:
Launches a second fleet of instances with the same capacity as the original fleet. Maintains the original fleet unchanged while the second fleet is launched.
Transitions traffic to the second fleet when the second fleet is fully deployed. Terminates the original fleet automatically 1 hour after transition.
Which solution will satisfy these requirements?

A. Use an AWS CloudFormation template with a retention policy for the ALB set to 1 hou
B. Update the Amazon Route 53 record to reflect the new ALB.
C. Use two AWS Elastic Beanstalk environments to perform a blue/green deployment from the original environment to the new on
D. Create an application version lifecycle policy to terminate the original environment in 1 hour.
E. Use AWS CodeDeploy with a deployment group configured with a blue/green deployment configuration Select the option Terminate the original instances in the deployment group with a waiting period of 1 hour.
F. Use AWS Elastic Beanstalk with the configuration set to Immutabl
G. Create an .ebextension using the Resources key that sets the deletion policy of the ALB to 1 hour, and deploy the application.

Answer: C

Explanation:
https://docs.aws.amazon.com/codedeploy/latest/APIReference/API_BlueInstanceTerminationOption.html The original revision termination settings are configured to wait 1 hour after traffic has been rerouted before
terminating the blue task set. https://docs.aws.amazon.com/AmazonECS/latest/developerguide/deployment-type-bluegreen.html

NEW QUESTION 2

An application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). A DevOps engineer is using AWS CodeDeploy to release a new version. The deployment fails during the AllowTraffic lifecycle event, but a cause for the failure is not indicated in the deployment logs.
What would cause this?

A. The appspe
B. yml file contains an invalid script that runs in the AllowTraffic lifecycle hook.
C. The user who initiated the deployment does not have the necessary permissions to interact with the ALB.
D. The health checks specified for the ALB target group are misconfigured.
E. The CodeDeploy agent was not installed in the EC2 instances that are pad of the ALB target group.

Answer: C

Explanation:
This failure is typically due to incorrectly configured health checks in Elastic Load Balancing for the Classic Load Balancer, Application Load Balancer, or Network Load Balancer used to manage traffic for the deployment group. To resolve the issue, review and correct any errors in the health check configuration for the load balancer.
https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting-deployments.html#troubleshooting-d

NEW QUESTION 3

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"
    )
}
```

Answer: C

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

NEW QUESTION 4
A company uses AWS CodeArtifact to centrally store Python packages. The CodeArtifact repository is configured with the following repository policy.

```
{
    "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 that are 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 5
A company has a mobile application that makes HTTP API calls to an Application Load Balancer (ALB). The ALB routes requests to an AWS Lambda function.
Many different versions of the application are in use at any given time, including versions that are in testing by a subset of users. The version of the application is defined in the user-agent header that is sent with all requests to the API.
After a series of recent changes to the API, the company has observed issues with the application. The company needs to gather a metric for each API operation by response code for each version of the application that is in use. A DevOps engineer has modified the Lambda function to extract the API operation name, version information from the user-agent header and response code.
Which additional set of actions should the DevOps engineer take to gather the required metrics?

A. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log grou
B. Configure a CloudWatch Logs metric filter that increments a metric for each API operation nam

C. Specify response code and application version as dimensions for the metric.
D. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log grou
E. Configure a CloudWatch Logs Insights query topopulate CloudWatch metrics from the log line
F. Specify response code and application version as dimensions for the metric.
G. Configure the ALB access logs to write to an Amazon CloudWatch Logs log grou
H. Modify the Lambda function to respond to the ALB with the API operation name, response code, and version number as response metadat
I. Configure a CloudWatch Logs metric filter that increments a metric for each API operation nam
J. Specify response code and application version as dimensions for the metric.
K. Configure AWS X-Ray integration on the Lambda functio
L. Modify the Lambda function to create an X-Ray subsegment with the API operation name, response code, and version numbe
M. Configure X-Ray insights to extract an aggregated metric for each API operation name and to publish the metric to Amazon CloudWatc
N. Specify response code and application version as dimensions for the metric.

Answer: A

Explanation:
"Note that the metric filter is different from a log insights query, where the experience is interactive and provides immediate search results for the user to investigate. No automatic action can be invoked from an insights query. Metric filters, on the other hand, will generate metric data in the form of a time series. This lets you create alarms that integrate into your ITSM processes, execute AWS Lambda functions, or even create anomaly detection models."
https://aws.amazon.com/blogs/mt/quantify-custom-application-metrics-with-amazon-cloudwatch-logs-and-metri

NEW QUESTION 6
A development team uses AWS CodeCommit for version control for applications. The development team uses AWS CodePipeline, AWS CodeBuild. and AWS CodeDeploy for CI/CD infrastructure. In CodeCommit, the development team recently merged pull requests that did not pass long-running tests in the code base. The development team needed to perform rollbacks to branches in the codebase, resulting in lost time and wasted effort.
A DevOps engineer must automate testing of pull requests in CodeCommit to ensure that reviewers more easily see the results of automated tests as part of the pull request review.
What should the DevOps engineer do to meet this requirement?

A. Create an Amazon EventBridge rule that reacts to the pullRequestStatusChanged even
B. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the applicatio
C. Program the Lambda function to post the CodeBuild badge as a comment on the pull request so that developers will see the badge in their code review.
D. Create an Amazon EventBridge rule that reacts to the pullRequestCreated even
E. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the applicatio
F. Program the Lambda function to post the CodeBuild test results as a comment on the pull request when the test results are complete.
G. Create an Amazon EventBridge rule that reacts to pullRequestCreated and pullRequestSourceBranchUpdated event
H. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the applicatio
I. Program the Lambda function to post the CodeBuild badge as a comment on the pull request so that developers will see the badge in their code review.
J. Create an Amazon EventBridge rule that reacts to the pullRequestStatusChanged even
K. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the applicatio
L. Program the Lambda function to post the CodeBuild test results as a comment on the pullrequest when the test results are complete.

Answer: B

Explanation:
https://aws.amazon.com/es/blogs/devops/complete-ci-cd-with-aws-codecommit-aws-codebuild-aws-codedeploy

NEW QUESTION 7
A DevOps engineer has automated a web service deployment by using AWS CodePipeline with the following steps:
1) An AWS CodeBuild project compiles the deployment artifact and runs unit tests.
2) An AWS CodeDeploy deployment group deploys the web service to Amazon EC2 instances in the staging environment.
3) A CodeDeploy deployment group deploys the web service to EC2 instances in the production environment. The quality assurance (QA) team requests permission to inspect the build artifact before the deployment to the
production environment occurs. The QA team wants to run an internal penetration testing tool to conduct
manual tests. The tool will be invoked by a REST API call.
Which combination of actions should the DevOps engineer take to fulfill this request? (Choose two.)

A. Insert a manual approval action between the test actions and deployment actions of the pipeline.
B. Modify the buildspec.yml file for the compilation stage to require manual approval before completion.
C. Update the CodeDeploy deployment groups so that they require manual approval to proceed.
D. Update the pipeline to directly call the REST API for the penetration testing tool.
E. Update the pipeline to invoke an AWS Lambda function that calls the REST API for the penetration testing tool.

Answer: AE

NEW QUESTION 8
A company requires an RPO of 2 hours and an RTO of 10 minutes for its data and application at all times. An application uses a MySQL database and Amazon EC2 web servers. The development team needs a strategy for failover and disaster recovery.
Which combination of deployment strategies will meet these requirements? (Select TWO.)

A. Create an Amazon Aurora cluster in one Availability Zone across multiple Regions as the data store Use Aurora's automatic recovery capabilities in the event of a disaster
B. Create an Amazon Aurora global database in two Regions as the data stor
C. In the event of a failure promote the secondary Region as the primary for the application.
D. Create an Amazon Aurora multi-master cluster across multiple Regions as the data stor
E. Use a Network Load Balancer to balance the database traffic in different Regions.
F. Set up the application in two Regions and use Amazon Route 53 failover-based routing that points to the Application Load Balancers in both Region
G. Use hearth checks to determine the availability in a givenRegio
H. Use Auto Scaling groups in each Region to adjust capacity based on demand.
I. Set up the application m two Regions and use a multi-Region Auto Scaling group behind Application Load Balancers to manage the capacity based on deman

J. In the event of a disaster adjust the Auto Scaling group's desired instance count to increase baseline capacity in the failover Region.

Answer: BD

NEW QUESTION 9

A large enterprise is deploying a web application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon RDS for Oracle DB instance and Amazon DynamoDB. There are separate environments tor development testing and production.

What is the MOST secure and flexible way to obtain password credentials during deployment?

A. Retrieve an access key from an AWS Systems Manager securestring parameter to access AWS services.Retrieve the database credentials from a Systems Manager SecureString parameter.
B. Launch the EC2 instances with an EC2 1AM role to access AWS services Retrieve the database credentials from AWS Secrets Manager.
C. Retrieve an access key from an AWS Systems Manager plaintext parameter to access AWS services.Retrieve the database credentials from a Systems Manager SecureString parameter.
D. Launch the EC2 instances with an EC2 1AM role to access AWS services Store the database passwords in an encrypted config file with the application artifacts.

Answer: B

Explanation:
AWS Secrets Manager is a secrets management service that helps you protect access to your applications, services, and IT resources. This service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Using Secrets Manager, you can secure and manage secrets used to access resources in the AWS Cloud, on third-party services, and on-premises. SSM parameter store and AWS Secret manager are both a secure option. However, Secrets manager is more flexible and has more options like password generation. Reference: https://www.1strategy.com/blog/2019/02/28/aws-parameter-store-vs-aws-secrets-manager/

NEW QUESTION 10

A company is adopting AWS CodeDeploy to automate its application deployments for a Java-Apache Tomcat application with an Apache Webserver. The development team started with a proof of concept, created a deployment group for a developer environment, and performed functional tests within the application. After completion, the team will create additional deployment groups for staging and production.

The current log level is configured within the Apache settings, but the team wants to change this configuration dynamically when the deployment occurs, so that they can set different log level configurations depending on the deployment group without having a different application revision for each group.

How can these requirements be met with the LEAST management overhead and without requiring different script versions for each deployment group?

A. Tag the Amazon EC2 instances depending on the deployment grou
B. Then place a script into the application revision that calls the metadata service and the EC2 API to identify which deployment group the instance is part o
C. Use this information to configure the log level setting
D. Reference the script as part of the AfterInstall lifecycle hook in the appspec.yml file.
E. Create a script that uses the CodeDeploy environment variable DEPLOYMENT_GROUP_ NAME to identify which deployment group the instance is part o
F. Use this information to configure the log level setting
G. Reference this script as part of the BeforeInstall lifecycle hook in the appspec.yml file.
H. Create a CodeDeploy custom environment variable for each environmen
I. Then place a script into the application revision that checks this environment variable to identify which deployment group the instance is part o
J. Use this information to configure the log level setting
K. Reference this script as part of the ValidateService lifecycle hook in the appspec.yml file.
L. Create a script that uses the CodeDeploy environment variable DEPLOYMENT_GROUP_ID to identify which deployment group the instance is part of to configure the log level setting
M. Reference this script as part of the Install lifecycle hook in the appspec.yml file.

Answer: B

Explanation:
The following are the steps that the company can take to change the log level dynamically when the deployment occurs:
➢ Create a script that uses the CodeDeploy environment variable DEPLOYMENT_GROUP_NAME to identify which deployment group the instance is part of.
➢ Use this information to configure the log level settings.
➢ Reference this script as part of the BeforeInstall lifecycle hook in the appspec.yml file.
The DEPLOYMENT_GROUP_NAME environment variable is automatically set by CodeDeploy when the deployment is triggered. This means that the script does not need to call the metadata service or the EC2 API to identify the deployment group.
This solution is the least complex and requires the least management overhead. It also does not require different script versions for each deployment group.
The following are the reasons why the other options are not correct:
➢ Option A is incorrect because it would require tagging the Amazon EC2 instances, which would be a manual and time-consuming process.
➢ Option C is incorrect because it would require creating a custom environment variable for each
environment. This would be a complex and error-prone process.
➢ Option D is incorrect because it would use the DEPLOYMENT_GROUP_ID environment variable.
However, this variable is not automatically set by CodeDeploy, so the script would need to call the metadata service or the EC2 API to get the deployment group ID. This would add complexity and overhead to the solution.

NEW QUESTION 10

A company uses Amazon S3 to store proprietary information. The development team creates buckets for new projects on a daily basis. The security team wants to ensure that all existing and future buckets have encryption logging and versioning enabled. Additionally, no buckets should ever be publicly read or write accessible.

What should a DevOps engineer do to meet these requirements?

A. Enable AWS CloudTrail and configure automatic remediation using AWS Lambda.
B. Enable AWS Conflg rules and configure automatic remediation using AWS Systems Manager documents.
C. Enable AWS Trusted Advisor and configure automatic remediation using Amazon EventBridge.
D. Enable AWS Systems Manager and configure automatic remediation using Systems Manager documents.

Answer: B

Explanation:
https://aws.amazon.com/blogs/mt/aws-config-auto-remediation-s3-compliance/ https://aws.amazon.com/blogs/aws/aws-config-rules-dynamic-compliance-checking-for-cloud-resources/

NEW QUESTION 13
A DevOps team is merging code revisions for an application that uses an Amazon RDS Multi-AZ DB cluster for its production database. The DevOps team uses continuous integration to periodically verify that the application works. The DevOps team needs to test the changes before the changes are deployed to the production database.
Which solution will meet these requirements'?

A. Use a buildspec file in AWS CodeBuild to restore the DB cluster from a snapshot of the production database run integration tests, and drop the restored database after verification.
B. Deploy the application to productio
C. Configure an audit log of data control language (DCL) operations to capture database activities to perform if verification fails.
D. Create a snapshot of the DB duster before deploying the application Use the Update requires Replacement property on the DB instance in AWS CloudFormation to deploy the application and apply the changes.
E. Ensure that the DB cluster is a Multi-AZ deploymen
F. Deploy the application with the update
G. Fail over to the standby instance if verification fails.

Answer: A

Explanation:
This solution will meet the requirements because it will create a temporary copy of the production database using a snapshot, run the integration tests on the copy, and delete the copy after the tests are done. This way, the production database will not be affected by the code revisions, and the DevOps team can test the changes before deploying them to production. A buildspec file is a YAML file that contains the commands and settings that CodeBuild uses to run a build1. The buildspec file can specify the steps to restore the DB cluster from a snapshot, run the integration tests, and drop the restored database2

NEW QUESTION 16
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 17
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 19

A Company uses AWS CodeCommit for source code control. Developers apply their changes to various feature branches and create pull requests to move those changes to the main branch when the changes are ready for production.
The developers should not be able to push changes directly to the main branch. The company applied the AWSCodeCommitPowerUser managed policy to the developers' IAM role, and now these developers can push changes to the main branch directly on every repository in the AWS account.
What should the company do to restrict the developers' ability to push changes to the main branch directly?

A. Create an additional policy to include a Deny rule for the GitPush and PutFile action
B. Include a restriction for the specific restriction for the specific repositories in the policy repositories in the policy statement with a condition that references the main branch.A Create an additional policy to include a Deny rule for the GitPush and PutFile actions Include a restriction for the specific repositories in the policy statement with a condition that references the main branch
C. Remove the IAM policy, and add an AWSCodeCommitReadOnly managed polic
D. Add an Allow rule for the GitPush and PutFile actions for the specific repositories in the policy statement with a condition that references the mam branch.
E. Modify the IAM policy Include a Deny rule for the GitPush and PutFile actions for the specific repositories in the policy statement with a condition that references the main branch.
F. Create an additional policy to include an Allow rule for the GitPush and PutFile action
G. Include a restriction for the specific repositories in the policy statement with a condition that references the feature branches.

Answer: A

Explanation:
By default, the AWSCodeCommitPowerUser managed policy allows users to push changes to any branch in any repository in the AWS account. To restrict the developers' ability to push changes to the main branch directly, an additional policy is needed that explicitly denies these actions for the main branch.
The Deny rule should be included in a policy statement that targets the specific repositories and includes a condition that references the main branch. The policy statement should look something like this:
{
"Effect": "Deny", "Action": ["codecommit:GitPush", "codecommit:PutFile"
],
"Resource": "arn:aws:codecommit:<region>:<account-id>:<repository-name>", "Condition": {
"StringEqualsIfExists": { "codecommit:References": ["refs/heads/main"
]
}
}

NEW QUESTION 21
A company's developers use Amazon EC2 instances as remote workstations. The company is concerned that users can create or modify EC2 security groups to allow unrestricted inbound access.
A DevOps engineer needs to develop a solution to detect when users create unrestricted security group rules. The solution must detect changes to security group rules in near real time, remove unrestricted rules, and send email notifications to the security team. The DevOps engineer has created an AWS Lambda function that checks for security group ID from input, removes rules that grant unrestricted access, and sends notifications through Amazon Simple Notification Service (Amazon SNS).
What should the DevOps engineer do next to meet the requirements?

A. Configure the Lambda function to be invoked by the SNS topi
B. Create an AWS CloudTrail subscription for the SNS topi
C. Configure a subscription filter for security group modification events.
D. Create an Amazon EventBridge scheduled rule to invoke the Lambda functio
E. Define a schedule pattern that runs the Lambda function every hour.
F. Create an Amazon EventBridge event rule that has the default event bus as the sourc
G. Define the rule's event pattern to match EC2 security group creation and modification event
H. Configure the rule to invoke the Lambda function.
I. Create an Amazon EventBridge custom event bus that subscribes to events from all AWS services.Configure the Lambda function to be invoked by the custom event bus.

Answer: C

Explanation:
To meet the requirements, the DevOps engineer should create an Amazon EventBridge event rule that has the default event bus as the source. The rule's event pattern should match EC2 security group creation and modification events, and it should be configured to invoke the Lambda function. This solution will allow for near real-time detection of security group rule changes and will trigger the Lambda function to remove any unrestricted rules and send email notifications to the security team.
https://repost.aws/knowledge-center/monitor-security-group-changes-ec2

NEW QUESTION 25
A video-sharing company stores its videos in Amazon S3. The company has observed a sudden increase in video access requests, but the company does not know which videos are most popular. The company needs to identify the general access pattern for the video files. This pattern includes the number of users who access a certain file on a given day, as well as the numb A DevOps engineer manages a large commercial website that runs on Amazon EC2 The website uses Amazon Kinesis Data Streams to collect and process web togs The DevOps engineer manages the Kinesis consumer application, which also runs on Amazon EC2 Sudden increases of data cause the Kinesis consumer application to (all behind and the Kinesis data streams drop records before the records can be processed The DevOps engineer must implement a solution to improve stream handling
Which solution meets these requirements with the MOST operational efficiency'' er of pull requests for certain files.
How can the company meet these requirements with the LEAST amount of effort?

A. Activate S3 server access loggin
B. Import the access logs into an Amazon Aurora databas
C. Use an Aurora SQL query to analyze the access patterns.
D. Activate S3 server access loggin
E. Use Amazon Athena to create an external table with the log file
F. Use Athena to create a SQL query to analyze the access patterns.
G. Invoke an AWS Lambda function for every S3 object access even
H. Configure the Lambda function to write the file access information, such as use
I. S3 bucket, and file key, to an Amazon Aurora databas

J. Use an Aurora SQL query to analyze the access patterns.
K. Record an Amazon CloudWatch Logs log message for every S3 object access even
L. Configure a CloudWatch Logs log stream to write the file access information, such as user, S3 bucket, and file key, to an Amazon Kinesis Data Analytics for SQL applicatio
M. Perform a sliding window analysis.

Answer: B

Explanation:
Activating S3 server access logging and using Amazon Athena to create an external table with the log files is the easiest and most cost-effective way to analyze access patterns. This option requires minimal setup and allows for quick analysis of the access patterns with SQL queries. Additionally, Amazon Athena scales automatically to match the query load, so there is no need for additional infrastructure provisioning or management.

NEW QUESTION 27
A company wants to ensure that their EC2 instances are secure. They want to be notified if any new vulnerabilities are discovered on their instances and they also want an audit trail of all login activities on the instances.
Which solution will meet these requirements'?

A. Use AWS Systems Manager to detect vulnerabilities on the EC2 instances Install the Amazon Kinesis Agent to capture system logs and deliver them to Amazon S3.
B. Use AWS Systems Manager to detect vulnerabilities on the EC2 instances Install the Systems Manager Agent to capture system logs and view login activity in the CloudTrail console.
C. Configure Amazon CloudWatch to detect vulnerabilities on the EC2 instances Install the AWS Config daemon to capture system logs and view them in the AWS Config console.
D. Configure Amazon Inspector to detect vulnerabilities on the EC2 instances Install the Amazon CloudWatch Agent to capture system logs and record them via Amazon CloudWatch Logs.

Answer: D

Explanation:
This solution will meet the requirements because it will use Amazon Inspector to scan the EC2 instances for any new vulnerabilities and generate findings that can be viewed in the Inspector console or sent as notifications via Amazon Simple Notification Service (SNS). It will also use the Amazon CloudWatch Agent to collect and send system logs from the EC2 instances to Amazon CloudWatch Logs, where they can be
stored, searched, and analyzed. The system logs can provide an audit trail of all login activities on the instances, as well as other useful information such as performance metrics, errors, and events.
https://docs.aws.amazon.com/inspector/latest/user/what-is-inspector.html

NEW QUESTION 28
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 29
A development team manually builds an artifact locally and then places it in an Amazon S3 bucket. The application has a local cache that must be cleared when a deployment occurs. The team runs a command to do this downloads the artifact from Amazon S3 and unzips the artifact to complete the deployment.
A DevOps team wants to migrate to a CI/CD process and build in checks to stop and roll back the deployment when a failure occurs. This requires the team to track the progression of the deployment.
Which combination of actions will accomplish this? (Select THREE)

A. Allow developers to check the code into a code repository Using Amazon EventBridge on every pull into the mam branch invoke an AWS Lambda function to build the artifact and store it in Amazon S3.
B. Create a custom script to clear the cache Specify the script in the BeforeInstall lifecycle hook in the AppSpec file.
C. Create user data for each Amazon EC2 instance that contains the clear cache script Once deployed test the application If it is not successful deploy it again.
D. Set up AWS CodePipeline to deploy the application Allow developers to check the code into a code repository as a source tor the pipeline.
E. Use AWS CodeBuild to build the artifact and place it in Amazon S3 Use AWS CodeDeploy to deploy the artifact to Amazon EC2 instances.
F. Use AWS Systems Manager to fetch the artifact from Amazon S3 and deploy it to all the instances.

Answer: BDE

NEW QUESTION 31

A company requires that its internally facing web application be highly available. The architecture is made up of one Amazon EC2 web server instance and one NAT instance that provides outbound internet access for updates and accessing public data.
Which combination of architecture adjustments should the company implement to achieve high availability? (Choose two.)

A. Add the NAT instance to an EC2 Auto Scaling group that spans multiple Availability Zone
B. Update the route tables.
C. Create additional EC2 instances spanning multiple Availability Zone
D. Add an Application Load Balancer to split the load between them.
E. Configure an Application Load Balancer in front of the EC2 instanc
F. Configure Amazon CloudWatch alarms to recover the EC2 instance upon host failure.
G. Replace the NAT instance with a NAT gateway in each Availability Zon
H. Update the route tables.
I. Replace the NAT instance with a NAT gateway that spans multiple Availability Zone
J. Update the route tables.

Answer: BD

Explanation:
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html

NEW QUESTION 32
A company is implementing an Amazon Elastic Container Service (Amazon ECS) cluster to run its workload. The company architecture will run multiple ECS services on the cluster. The architecture includes an Application Load Balancer on the front end and uses multiple target groups to route traffic.
A DevOps engineer must collect application and access logs. The DevOps engineer then needs to send the logs to an Amazon S3 bucket for near-real-time analysis.
Which combination of steps must the DevOps engineer take to meet these requirements? (Choose three.)

A. Download the Amazon CloudWatch Logs container instance from AW
B. Configure this instance as a tas
C. Update the application service definitions to include the logging task.
D. Install the Amazon CloudWatch Logs agent on the ECS instance
E. Change the logging driver in the ECS task definition to awslogs.
F. Use Amazon EventBridge to schedule an AWS Lambda function that will run every 60 seconds and will run the Amazon CloudWatch Logs create-export-task comman
G. Then point the output to the logging S3 bucket.
H. Activate access logging on the AL
I. Then point the ALB directly to the logging S3 bucket.
J. Activate access logging on the target groups that the ECS services us
K. Then send the logs directly to the logging S3 bucket.
L. Create an Amazon Kinesis Data Firehose delivery stream that has a destination of the logging S3 bucket.Then create an Amazon CloudWatch Logs subscription filter for Kinesis Data Firehose.

Answer: BDF

Explanation:
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-logging-monitoring.html

NEW QUESTION 34
A company has developed an AWS Lambda function that handles orders received through an API. The company is using AWS CodeDeploy to deploy the Lambda function as the final stage of a CI/CD pipeline.
A DevOps engineer has noticed there are intermittent failures of the ordering API for a few seconds after deployment. After some investigation the DevOps engineer believes the failures are due to database changes not having fully propagated before the Lambda function is invoked
How should the DevOps engineer overcome this?

A. Add a BeforeAllowTraffic hook to the AppSpec file that tests and waits for any necessary database changes before traffic can flow to the new version of the Lambda function.
B. Add an AfterAllowTraffic hook to the AppSpec file that forces traffic to wait for any pending database changes before allowing the new version of the Lambda function to respond.
C. Add a BeforeAllowTraffic hook to the AppSpec file that tests and waits for any necessary database changes before deploying the new version of the Lambda function.
D. Add a validateService hook to the AppSpec file that inspects incoming traffic and rejects the payload if dependent services such as the database are not yet ready.

Answer: A

Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html#appspec

NEW QUESTION 39
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 accounts.Use Amazon Athena to analyze the AWS CloudTrail output, looking for events that deny an ec2:RunInstances action.
I. Deploy an IAM role to all accounts from a single trusted accoun
J. Build a pipeline with AWS CodePipeline with a stage in AWS Lambda to assume the IAM role, and list all EBS volumes in the accoun
K. 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 44
A DevOps engineer has implemented a CI/CO pipeline to deploy an AWS Cloud Format ion template that provisions a web application. The web application consists of an Application Load Balancer (ALB) a target group, a launch template that uses an Amazon Linux 2 AMI an Auto Scaling group of Amazon EC2 instances, a security group and an Amazon RDS for MySQL database The launch template includes user data that specifies a script to install and start the application.
The initial deployment of the application was successful. The DevOps engineer made changes to update the version of the application with the user data. The CI/CD pipeline has deployed a new version of the template However, the health checks on the ALB are now failing The health checks have marked all targets as unhealthy.
During investigation the DevOps engineer notices that the Cloud Formation stack has a status of UPDATE_COMPLETE. However, when the DevOps engineer connects to one of the EC2 instances and checks /varar/log messages, the DevOps engineer notices that the Apache web server failed to start successfully because of a configuration error
How can the DevOps engineer ensure that the CloudFormation deployment will fail if the user data fails to successfully finish running?

A. Use the cfn-signal helper script to signal success or failure to CloudFormation Use the WaitOnResourceSignals update policy within the CloudFormation template Set an appropriate timeout for the update policy.
B. Create an Amazon CloudWatch alarm for the UnhealthyHostCount metri
C. Include an appropriate alarm threshold for the target group Create an Amazon Simple Notification Service (Amazon SNS) topic as the target to signal success or failure to CloudFormation
D. Create a lifecycle hook on the Auto Scaling group by using the AWS AutoScaling LifecycleHook resource Create an Amazon Simple Notification Service (Amazon SNS) topic as the target to signal success or failure to CloudFormation Set an appropriate timeout on the lifecycle hook.
E. Use the Amazon CloudWatch agent to stream the cloud-init logs Create a subscription filter that includes an AWS Lambda function with an appropriate invocation timeout Configure the Lambda function to use the SignalResource API operation to signal success or failure to CloudFormation.

Answer: A

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-updatepolicy.html

NEW QUESTION 45
An application running on a set of Amazon EC2 instances in an Auto Scaling group requires a configuration file to operate. The instances are created and maintained with AWS CloudFormation. A DevOps engineer wants the instances to have the latest configuration file when launched and wants changes to the configuration file to be reflected on all the instances with a minimal delay when the CloudFormation template is updated. Company policy requires that application configuration files be maintained along with AWS infrastructure configuration files m source control.
Which solution will accomplish this?

A. In the CloudFormaiion template add an AWS Config rul
B. Place the configuration file content in the rule's InputParameters property and set the Scope property to the EC2 Auto Scaling grou
C. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
D. In the CloudFormation template add an EC2 launch template resourc
E. Place the configuration file content in the launch templat
F. Configure the cfn-mit script to run when the instance is launched and configure the cfn-hup script to poll for updates to the configuration.
G. In the CloudFormation template add an EC2 launch template resourc
H. Place the configuration file content in the launch templat
I. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
J. In the CloudFormation template add CloudFormation imt metadat
K. Place the configuration file content m the metadat
L. Configure the cfn-init script to run when the instance is launched and configure thecfn-hup script to poll for updates to the configuration.

Answer: D

Explanation:
Use the AWS::CloudFormation::Init type to include metadata on an Amazon EC2 instance for the cfn-init helper script. If your template calls the cfn-init script, the script looks for resource metadata rooted in the AWS::CloudFormation::Init metadata key. Reference:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-init.html

NEW QUESTION 47
A company has an on-premises application that is written in Go. A DevOps engineer must move the application to AWS. The company's development team wants to enable blue/green deployments and perform A/B testing.
Which solution will meet these requirements?

A. Deploy the application on an Amazon EC2 instance, and create an AMI of the instanc
B. Use the AMI to create an automatic scaling launch configuration that is used in an Auto Scaling grou
C. Use Elastic Load Balancing to distribute traffi
D. When changes are made to the application, a new AMI will be created, which will initiate an EC2 instance refresh.
E. Use Amazon Lightsail to deploy the applicatio
F. Store the application in a zipped format in an Amazon S3 bucke
G. Use this zipped version to deploy new versions of the application to Lightsai
H. Use Lightsail deployment options to manage the deployment.
I. Use AWS CodeArtifact to store the application cod
J. Use AWS CodeDeploy to deploy the application to a fleet of Amazon EC2 instance

K. Use Elastic Load Balancing to distribute the traffic to the EC2 instance
L. When making changes to the application, upload a new version to CodeArtifact and create a new CodeDeploy deployment.
M. Use AWS Elastic Beanstalk to host the applicatio
N. Store a zipped version of the application in Amazon S3. Use that location to deploy new versions of the applicatio
O. Use Elastic Beanstalk to manage the deployment options.

Answer: D

Explanation:
https://aws.amazon.com/quickstart/architecture/blue-green-deployment/

NEW QUESTION 51
A company's DevOps engineer is creating an AWS Lambda function to process notifications from an Amazon Simple Notification Service (Amazon SNS) topic. The Lambda function will process the notification messages and will write the contents of the notification messages to an Amazon RDS Multi-AZ DB instance.
During testing a database administrator accidentally shut down the DB instance. While the database was down the company lost several of the SNS notification messages that were delivered during that time.
The DevOps engineer needs to prevent the loss of notification messages in the future Which solutions will meet this requirement? (Select TWO.)

A. Replace the RDS Multi-AZ DB instance with an Amazon DynamoDB table.
B. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination of the Lambda function.
C. Configure an Amazon Simple Queue Service (Amazon SQS> dead-letter queue for the SNS topic.
D. Subscribe an Amazon Simple Queue Service (Amazon SQS) queue to the SNS topic Configure the Lambda function to process messages from the SQS queue.
E. Replace the SNS topic with an Amazon EventBridge event bus Configure an EventBridge rule on the new event bus to invoke the Lambda function for each event.

Answer: CD

Explanation:
These solutions will meet the requirement because they will prevent the loss of notification messages in the future. An Amazon SQS queue is a service that provides a reliable, scalable, and secure message queue for asynchronous communication between distributed components. You can use an SQS queue to buffer messages from an SNS topic and ensure that they are delivered and processed by a Lambda function, even if the function or the database is temporarily unavailable.
Option C will configure an SQS dead-letter queue for the SNS topic. A dead-letter queue is a queue that receives messages that could not be delivered to any subscriber after a specified number of retries. You can use a dead-letter queue to store and analyze failed messages, or to reprocess them later. This way, you can avoid losing messages that could not be delivered to the Lambda function due to network errors, throttling, or other issues.
Option D will subscribe an SQS queue to the SNS topic and configure the Lambda function to process messages from the SQS queue. This will decouple the SNS topic from the Lambda function and provide more flexibility and control over the message delivery and processing. You can use an SQS queue to store messages from the SNS topic until they are ready to be processed by the Lambda function, and also to retry processing in case of failures. This way, you can avoid losing messages that could not be processed by the Lambda function due to database errors, timeouts, or other issues.

NEW QUESTION 55
A company wants to use a grid system for a proprietary enterprise m-memory data store on top of AWS. This system can run in multiple server nodes in any Linux-based distribution. The system must be able to reconfigure the entire cluster every time a node is added or removed. When adding or removing nodes an /etc./cluster/nodes config file must be updated listing the IP addresses of the current node members of that cluster.
The company wants to automate the task of adding new nodes to a cluster. What can a DevOps engineer do to meet these requirements?

A. Use AWS OpsWorks Stacks to layer the server nodes of that cluste
B. Create a Chef recipe that populates the content of the 'etc./cluster/nodes config file and restarts the service by using the current members of the laye
C. Assign that recipe to the Configure lifecycle event.
D. Put the file nodes config in version contro
E. Create an AWS CodeDeploy deployment configuration and deployment group based on an Amazon EC2 tag value for thecluster node
F. When adding a new node to the cluster update the file with all tagged instances and make a commit in version contro
G. Deploy the new file and restart the services.
H. Create an Amazon S3 bucket and upload a version of the /etc./cluster/nodes config file Create a crontab script that will poll for that S3 file and download it frequentl
I. Use a process manager such as Monit or system, to restart the cluster services when it detects that the new file was modifie
J. When adding a node to the cluster edit the file's most recent members Upload the new file to the S3 bucket.
K. Create a user data script that lists all members of the current security group of the cluster and automatically updates the /etc/cluster/. nodes confi
L. Tile whenever a new instance is added to the cluster.

Answer: A

Explanation:
You can run custom recipes manually, but the best approach is usually to have AWS OpsWorks Stacks run them automatically. Every layer has a set of built-in recipes assigned each of five lifecycle events—Setup, Configure, Deploy, Undeploy, and Shutdown. Each time an event occurs for an instance, AWS OpsWorks Stacks runs the associated recipes for each of the instance's layers, which handle the corresponding tasks. For example, when an instance finishes booting, AWS OpsWorks Stacks triggers a Setup event. This event runs the associated layer's Setup recipes, which typically handle tasks such as installing and configuring packages

NEW QUESTION 58
A company is using AWS CodePipeline to automate its release pipeline. AWS CodeDeploy is being used in the pipeline to deploy an application to Amazon Elastic Container Service (Amazon ECS) using the blue/green deployment model. The company wants to implement scripts to test the green version of the application before shifting traffic. These scripts will complete in 5 minutes or less. If errors are discovered during these tests, the application must be rolled back.
Which strategy will meet these requirements?

A. Add a stage to the CodePipeline pipeline between the source and deploy stage
B. Use AWS CodeBuild to create a runtime environment and build commands in the buildspec file to invoke test script
C. If errors are found, use the aws deploy stop-deployment command to stop the deployment.
D. Add a stage to the CodePipeline pipeline between the source and deploy stage

E. Use this stage to invoke an AWS Lambda function that will run the test script
F. If errors are found, use the aws deploystop-deployment command to stop the deployment.
G. Add a hooks section to the CodeDeploy AppSpec fil
H. Use the AfterAllowTestTraffic lifecycle event to invoke an AWS Lambda function to run the test script
I. If errors are found, exit the Lambda function with an error to initiate rollback.
J. Add a hooks section to the CodeDeploy AppSpec fil
K. Use the AfterAllowTraffic lifecycle event to invoke the test script
L. If errors are found, use the aws deploy stop-deployment CLI command to stop the deployment.

Answer: C

Explanation:

https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html

NEW QUESTION 63

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 window.Target a Systems Manager document to restart the EC2 instance.
D. Configure an event source of AWS Health, a service of EC2, and an event type that indicates instance maintenanc
E. Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.
F. Configure an event source of EC2 and an event type that indicates instance maintenanc
G. 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 67

A production account has a requirement that any Amazon EC2 instance that has been logged in to manually must be terminated within 24 hours. All applications in the production account are using Auto Scaling groups with the Amazon CloudWatch Logs agent configured.
How can this process be automated?

A. Create a CloudWatch Logs subscription to an AWS Step Functions applicatio
B. Configure an AWS Lambda function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissione
C. Create an Amazon EventBridge rule to invoke a second Lambda function once a day that will terminate all instances with this tag.
D. Create an Amazon CloudWatch alarm that will be invoked by the login even
E. Send the notification to an Amazon Simple Notification Service (Amazon SNS) topic that the operations team is subscribed to, and have them terminate the EC2 instance within 24 hours.
F. Create an Amazon CloudWatch alarm that will be invoked by the login even
G. Configure the alarm to send to an Amazon Simple Queue Service (Amazon SQS) queu
H. Use a group of worker instances to process messages from the queue, which then schedules an Amazon EventBridge rule to be invoked.
I. Create a CloudWatch Logs subscription to an AWS Lambda functio
J. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned.Create an Amazon EventBridge rule to invoke a daily Lambda function that terminates all instances with this tag.

Answer: D

Explanation:

"You can use subscriptions to get access to a real-time feed of log events from CloudWatch Logs and have it delivered to other services such as an Amazon Kinesis stream, an Amazon Kinesis Data Firehose stream, or AWS Lambda for custom processing, analysis, or loading to other systems. When log events are sent to the receiving service, they are Base64 encoded and compressed with the gzip format." See
https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/Subscriptions.html

NEW QUESTION 69

A company is hosting a static website from an Amazon S3 bucket. The website is available to customers at example.com. The company uses an Amazon Route 53 weighted routing policy with a TTL of 1 day. The company has decided to replace the existing static website with a dynamic web application. The dynamic web application uses an Application Load Balancer (ALB) in front of a fleet of Amazon EC2 instances.
On the day of production launch to customers, the company creates an additional Route 53 weighted DNS record entry that points to the ALB with a weight of 255 and a TTL of 1 hour. Two days later, a DevOps engineer notices that the previous static website is displayed sometimes when customers navigate to example.com.
How can the DevOps engineer ensure that the company serves only dynamic content for example.com?

A. Delete all objects, including previous versions, from the S3 bucket that contains the static website content.

B. Update the weighted DNS record entry that points to the S3 bucke
C. Apply a weight of 0. Specify the domain reset option to propagate changes immediately.
D. Configure webpage redirect requests on the S3 bucket with a hostname that redirects to the ALB.
E. Remove the weighted DNS record entry that points to the S3 bucket from the example.com hosted zone.Wait for DNS propagation to become complete.

Answer: D

NEW QUESTION 71
An AWS CodePipeline pipeline has implemented a code release process. The pipeline is integrated with AWS CodeDeploy to deploy versions of an application to multiple Amazon EC2 instances for each CodePipeline stage.
During a recent deployment the pipeline failed due to a CodeDeploy issue. The DevOps team wants to improve monitoring and notifications during deployment to decrease resolution times.
What should the DevOps engineer do to create notifications. When issues are discovered?

A. Implement Amazon CloudWatch Logs for CodePipeline and CodeDeploy create an AWS Config rule to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.
B. Implement Amazon EventBridge for CodePipeline and CodeDeploy create an AWS Lambda function to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.
C. Implement AWS CloudTrail to record CodePipeline and CodeDeploy API call information create an AWS Lambda function to evaluate code deployment issues and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.
D. Implement Amazon EventBridge for CodePipeline and CodeDeploy create an Amazo
E. Inspector assessment target to evaluate code deployment issues and create an Amazon Simpl
F. Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.

Answer: B

Explanation:
AWS CloudWatch Events can be used to monitor events across different AWS resources, and a CloudWatch Event Rule can be created to trigger an AWS Lambda function when a deployment issue is detected in the pipeline. The Lambda function can then evaluate the issue and send a notification to the appropriate stakeholders through an Amazon SNS topic. This approach allows for real-time notifications and faster resolution times.

NEW QUESTION 75
A company has a guideline that every Amazon EC2 instance must be launched from an AMI that the company's security team produces Every month the security team sends an email message with the latest approved AMIs to all the development teams.
The development teams use AWS CloudFormation to deploy their applications. When developers launch a new service they have to search their email for the latest AMIs that the security department sent. A DevOps engineer wants to automate the process that the security team uses to provide the AMI IDs to the development teams.
What is the MOST scalable solution that meets these requirements?

A. Direct the security team to use CloudFormation to create new versions of the AMIs and to list! the AMI ARNs in an encrypted Amazon S3 object as part of the stack's Outputs Section Instruct the developers to use a cross-stack reference to load the encrypted S3 object and obtain the most recent AMI ARNs.
B. Direct the security team to use a CloudFormation stack to create an AWS CodePipeline pipeline that builds new AMIs and places the latest AMI ARNs in an encrypted Amazon S3 object as part of the pipeline output Instruct the developers to use a cross-stack reference within their own CloudFormation template to obtain the S3 object location and the most recent AMI ARNs.
C. Direct the security team to use Amazon EC2 Image Builder to create new AMIs and to place the AMI ARNs as parameters in AWS Systems Manager Parameter Store Instruct the developers to specify a parameter of type SSM in their CloudFormation stack to obtain the most recent AMI ARNs from Parameter Store.
D. Direct the security team to use Amazon EC2 Image Builder to create new AMIs and to create an Amazon Simple Notification Service (Amazon SNS) topic so that every development team can receive notification
E. When the development teams receive a notification instruct them to write an AWS Lambda function that will update their CloudFormation stack with the most recent AMI ARNs.

Answer: C

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/dynamic-references.html

NEW QUESTION 76
An IT team has built an AWS CloudFormation template so others in the company can quickly and reliably deploy and terminate an application. The template creates an Amazon EC2 instance with a user data script to install the application and an Amazon S3 bucket that the application uses to serve static webpages while it is running.
All resources should be removed when the CloudFormation stack is deleted. However, the team observes that CloudFormation reports an error during stack deletion, and the S3 bucket created by the stack is not deleted.
How can the team resolve the error in the MOST efficient manner to ensure that all resources are deleted without errors?

A. Add a DelelionPolicy attribute to the S3 bucket resource, with the value Delete forcing the bucket to be removed when the stack is deleted.
B. Add a custom resource with an AWS Lambda function with the DependsOn attribute specifying the S3bucket, and an IAM rol
C. Write the Lambda function to delete all objects from the bucket when RequestType is Delete.
D. Identify the resource that was not delete
E. Manually empty the S3 bucket and then delete it.
F. Replace the EC2 and S3 bucket resources with a single AWS OpsWorks Stacks resourc
G. Define a custom recipe for the stack to create and delete the EC2 instance and the S3 bucket.

Answer: B

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-s3-custom-resources/

NEW QUESTION 78

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: A

NEW QUESTION 80
A company wants to use AWS CloudFormation for infrastructure deployment. The company has strict tagging and resource requirements and wants to limit the deployment to two Regions. Developers will need to deploy multiple versions of the same application.
Which solution ensures resources are deployed in accordance with company policy?

A. Create AWS Trusted Advisor checks to find and remediate unapproved CloudFormation StackSets.
B. Create a Cloud Formation drift detection operation to find and remediate unapproved CloudFormation StackSets.
C. Create CloudFormation StackSets with approved CloudFormation templates.
D. Create AWS Service Catalog products with approved CloudFormation templates.

Answer: D

Explanation:
service catalog uses stacksets and can enforce tag and restrict resources AWS Customer case with tag enforcement
https://aws.amazon.com/ko/blogs/apn/enforce-centralized-tag-compliance-using-aws-service-catalog-amazon-dy And Youtube video showing how to restrict resources per user with portfolio https://www.youtube.com/watch?v=LzvhTcqqyog

NEW QUESTION 82
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 .z ip 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 85
The security team depends on AWS CloudTrail to detect sensitive security issues in the company's AWS account. The DevOps engineer needs a solution to auto-remediate CloudTrail being turned off in an AWS account.
What solution ensures the LEAST amount of downtime for the CloudTrail log deliveries?

A. Create an Amazon EventBridge rule for the CloudTrail StopLogging even
B. Create an AWS Lambda (unction that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was calle
C. Add the Lambda function ARN as a target to the EventBridge rule.
D. Deploy the AWS-managed CloudTrail-enabled AWS Config rule set with a periodic interval to 1 hour.Create an Amazon EventBridge rule tor AWS Config rules compliance chang
E. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLoggmg was calle
F. Add the Lambda function ARN as a target to the EventBridge rule.
G. Create an Amazon EventBridge rule for a scheduled event every 5 minute
H. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on a CloudTrail trail in the AWS accoun
I. Add the Lambda function ARN as a target to the EventBridge rule.

J. Launch a t2 nano instance with a script running every 5 minutes that uses the AWS SDK to query CloudTrail in the current accoun
K. If the CloudTrail trail is disabled have the script re-enable the trail.

Answer: A

Explanation:
https://aws.amazon.com/blogs/mt/monitor-changes-and-auto-enable-logging-in-aws-cloudtrail/

NEW QUESTION 90
A DevOps engineer is building a continuous deployment pipeline for a serverless application that uses AWS Lambda functions. The company wants to reduce the customer impact of an unsuccessful deployment. The company also wants to monitor for issues.
Which deploy stage configuration will meet these requirements?

A. Use an AWS Serverless Application Model (AWS SAM) template to define the serverless application.Use AWS CodeDeploy to deploy the Lambda functions with the Canary10Percent15Minutes Deployment Preference Typ
B. Use Amazon CloudWatch alarms to monitor the health of the functions.
C. Use AWS CloudFormation to publish a new stack update, and include Amazon CloudWatch alarms on all resource
D. Set up an AWS CodePipeline approval action for a developer to verify and approve the AWS CloudFormation change set.
E. Use AWS CloudFormation to publish a new version on every stack update, and include Amazon CloudWatch alarms on all resource
F. Use the RoutingConfig property of the AWS::Lambda::Alias resource to update the traffic routing during the stack update.
G. Use AWS CodeBuild to add sample event payloads for testing to the Lambda function
H. Publish a new version of the functions, and include Amazon CloudWatch alarm
I. Update the production alias to point to the new versio
J. Configure rollbacks to occur when an alarm is in the ALARM state.

Answer: D

Explanation:
Use routing configuration on an alias to send a portion of traffic to a second function version. For example, you can reduce the risk of deploying a new version by configuring the alias to send most of the traffic to the existing version, and only a small percentage of traffic to the new version.
https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html
The following are the steps involved in the deploy stage configuration that will meet the requirements:
➢ Use AWS CodeBuild to add sample event payloads for testing to the Lambda functions.
➢ Publish a new version of the functions, and include Amazon CloudWatch alarms.
➢ Update the production alias to point to the new version.
➢ Configure rollbacks to occur when an alarm is in the ALARM state.
This configuration will help to reduce the customer impact of an unsuccessful deployment by deploying the new version of the functions to a staging environment first. This will allow the DevOps engineer to test the new version of the functions before deploying it to production.
The configuration will also help to monitor for issues by including Amazon CloudWatch alarms. These alarms will alert the DevOps engineer if there are any problems with the new version of the functions.

NEW QUESTION 94
A company manages an application that stores logs in Amazon CloudWatch Logs. The company wants to archive the logs to an Amazon S3 bucket Logs are rarely accessed after 90 days and must be retained tor 10 years.
Which combination of steps should a DevOps engineer take to meet these requirements? (Select TWO.)

A. Configure a CloudWatch Logs subscription filter to use AWS Glue to transfer all logs to an S3 bucket.
B. Configure a CloudWatch Logs subscription filter to use Amazon Kinesis Data Firehose to stream all logs to an S3 bucket.
C. Configure a CloudWatch Logs subscription fitter to stream all logs to an S3 bucket.
D. Configure the S3 bucket lifecycle policy to transition logs to S3 Glacier after 90 days and to expire logs after 3.650 days.
E. Configure the S3 bucket lifecycle policy to transition logs to Reduced Redundancy after 90 days and to expire logs after 3.650 days.

Answer: BD

Explanation:
https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html

NEW QUESTION 98
A company has many AWS accounts. During AWS account creation the company uses automation to create an Amazon CloudWatch Logs log group in every AWS Region that the company operates in. The automaton configures new resources in the accounts to publish logs to the provisioned log groups in their Region. The company has created a logging account to centralize the logging from all the other accounts. A DevOps engineer needs to aggregate the log groups from all the accounts to an existing Amazon S3 bucket in the logging account.
Which solution will meet these requirements in the MOST operationally efficient manner?

A. In the logging account create a CloudWatch Logs destination with a destination polic
B. For each new account subscribe the CloudWatch Logs log groups to th
C. Destination Configure a single Amazon Kinesis data stream and a single Amazon Kinesis Data Firehose delivery stream to deliver the logs from the CloudWatch Logs destination to the S3 bucket.
D. In the logging account create a CloudWatch Logs destination with a destination policy for each Region.For each new account subscribe the CloudWatch Logs log groups to the destinatio
E. Configure a single Amazon Kinesis data stream and a single Amazon Kinesis Data Firehose delivery stream to deliver the logs from all the CloudWatch Logs destinations to the S3 bucket.
F. In the logging account create a CloudWatch Logs destination with a destination policy for each Region.For each new account subscribe the CloudWatch Logs log groups to the destination Configure an Amazon Kinesis data stream and an Amazon Kinesis Data Firehose delivery stream for each Region to deliver the logs from the CloudWatch Logs destinations to the S3 bucket.
G. In the logging account create a CloudWatch Logs destination with a destination polic
H. For each new account subscribe the CloudWatch Logs log groups to the destinatio
I. Configure a single Amazon Kinesis data stream to deliver the logs from the CloudWatch Logs destination to the S3 bucket.

Answer: C

Explanation:
This solution will meet the requirements in the most operationally efficient manner because it will use CloudWatch Logs destination to aggregate the log groups from all the accounts to a single S3 bucket in the logging account. However, unlike option A, this solution will create a CloudWatch Logs destination for each region, instead of a single destination for all regions. This will improve the performance and reliability of the log delivery, as it will avoid cross-region data transfer and latency issues. Moreover, this solution will use an Amazon Kinesis data stream and an Amazon Kinesis Data Firehose delivery stream for each region, instead of a single stream for all regions. This will also improve the scalability and throughput of the log delivery, as it will avoid bottlenecks and throttling issues that may occur with a single stream.

NEW QUESTION 103
A company deploys its corporate infrastructure on AWS across multiple AWS Regions and Availability Zones. The infrastructure is deployed on Amazon EC2 instances and connects with AWS IoT Greengrass devices. The company deploys additional resources on on-premises servers that are located in the corporate headquarters.
The company wants to reduce the overhead involved in maintaining and updating its resources. The company's DevOps team plans to use AWS Systems Manager to implement automated management and application of patches. The DevOps team confirms that Systems Manager is available in the Regions that the resources are deployed m Systems Manager also is available in a Region near the corporate headquarters.
Which combination of steps must the DevOps team take to implement automated patch and configuration management across the company's EC2 instances IoT devices and on-premises infrastructure? (Select THREE.)

A. Apply tags lo all the EC2 instance
B. AWS IoT Greengrass devices, and on-premises server
C. Use Systems Manager Session Manager to push patches to all the tagged devices.
D. Use Systems Manager Run Command to schedule patching for the EC2 instances AWS IoT Greengrass devices and on-premises servers.
E. Use Systems Manager Patch Manager to schedule patching loT the EC2 instances AWS IoT Greengrass devices and on-premises servers as a Systems Manager maintenance window task.
F. Configure Amazon EventBridge to monitor Systems Manager Patch Manager for updates to patch baseline
G. Associate Systems Manager Run Command with the event lo initiate a patch action for all EC2 instances AWS IoT Greengrass devices and on-premises servers.
H. Create an IAM instance profile for Systems Manager Attach the instance profile to all the EC2 instances in the AWS accoun
I. For the AWS IoT Greengrass devices and on-premises servers create an IAM service role for Systems Manager.
J. Generate a managed-instance activation Use the Activation Code and Activation ID to install Systems Manager Agent (SSM Agent) on each server in the on-premises environment Update the AWS IoT Greengrass IAM token exchange role Use the role to deploy SSM Agent on all the loT devices.

Answer: CEF

Explanation:
https://aws.amazon.com/blogs/mt/how-to-centrally-manage-aws-iot-greengrass-devices-using-aws-systems-man

NEW QUESTION 105
A development team is using AWS CodeCommit to version control application code and AWS CodePipeline to orchestrate software deployments. The team has decided to use a remote main branch as the trigger for the pipeline to integrate code changes. A developer has pushed code changes to the CodeCommit repository, but noticed that the pipeline had no reaction, even after 10 minutes.
Which of the following actions should be taken to troubleshoot this issue?

A. Check that an Amazon EventBridge rule has been created for the main branch to trigger the pipeline.
B. Check that the CodePipeline service role has permission to access the CodeCommit repository.
C. Check that the developer's IAM role has permission to push to the CodeCommit repository.
D. Check to see if the pipeline failed to start because of CodeCommit errors in Amazon CloudWatch Logs.

Answer: A

Explanation:
When you create a pipeline from CodePipeline during the step-by-step it creates a CloudWatch Event rule for a given branch and repo
like this:
```
{
"source": [ "aws.codecommit"
],
"detail-type": [
"CodeCommit Repository State Change"
],
"resources": [
"arn:aws:codecommit:us-east-1:xxxxx:repo-name"
],
"detail": {
"event": [ "referenceCreated", "referenceUpdated"
],
"referenceType": [ "branch"
],
"referenceName": [ "master"
]
}
}
```
https://docs.aws.amazon.com/codepipeline/latest/userguide/pipelines-trigger-source-repo-changes-console.html

NEW QUESTION 109
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 113
A space exploration company receives telemetry data from multiple satellites. Small packets of data are received through Amazon API Gateway and are placed directly into an Amazon Simple Queue Service (Amazon SQS) standard queue. A custom application is subscribed to the queue and transforms the data into a standard format.
Because of inconsistencies in the data that the satellites produce, the application is occasionally unable to transform the data. In these cases, the messages remain in the SQS queue. A DevOps engineer must develop a solution that retains the failed messages and makes them available to scientists for review and future processing.
Which solution will meet these requirements?

A. Configure AWS Lambda to poll the SQS queue and invoke a Lambda function to check whether the queue messages are vali
B. If validation fails, send a copy of the data that is not valid to an Amazon S3 bucket so that the scientists can review and correct the dat
C. When the data is corrected, amend the message in the SQS queue by using a replay Lambda function with the corrected data.
D. Convert the SQS standard queue to an SQS FIFO queu
E. Configure AWS Lambda to poll the SQS queue every 10 minutes by using an Amazon EventBridge schedul
F. Invoke the Lambda function to identify any messages with a SentTimestamp value that is older than 5 minutes, push the data to the same location as the application's output location, and remove the messages from the queue.
G. Create an SQS dead-letter queu
H. Modify the existing queue by including a redrive policy that sets the Maximum Receives setting to 1 and sets the dead-letter queue ARN to the ARN of the newly created queu
I. Instruct the scientists to use the dead-letter queue to review the data that is not vali
J. Reprocess this data at a later time.
K. Configure API Gateway to send messages to different SQS virtual queues that are named for each of the satellite
L. Update the application to use a new virtual queue for any data that it cannot transform, and send the message to the new virtual queu
M. Instruct the scientists to use the virtual queue to review the data that is not vali
N. Reprocess this data at a later time.

Answer: C

Explanation:
Create an SQS dead-letter queue. Modify the existing queue by including a redrive policy that sets the Maximum Receives setting to 1 and sets the dead-letter queue ARN to the ARN of the newly created queue. Instruct the scientists to use the dead-letter queue to review the data that is not valid. Reprocess this data at a later time.

NEW QUESTION 114
A company is using an Amazon Aurora cluster as the data store for its application. The Aurora cluster is configured with a single DB instance. The application performs read and write operations on the database by using the cluster's instance endpoint.
The company has scheduled an update to be applied to the cluster during an upcoming maintenance window. The cluster must remain available with the least possible interruption during the maintenance window.
What should a DevOps engineer do to meet these requirements?

A. Add a reader instance to the Aurora cluste
B. Update the application to use the Aurora cluster endpoint for write operation
C. Update the Aurora cluster's reader endpoint for reads.
D. Add a reader instance to the Aurora cluste
E. Create a custom ANY endpoint for the cluste
F. Update the application to use the Aurora cluster's custom ANY endpoint for read and write operations.
G. Turn on the Multi-AZ option on the Aurora cluste
H. Update the application to use the Aurora cluster endpoint for write operation
I. Update the Aurora cluster's reader endpoint for reads.
J. Turn on the Multi-AZ option on the Aurora cluste
K. Create a custom ANY endpoint for the cluster.Update the application to use the Aurora cluster's custom ANY endpoint for read and write operations.

Answer: C

Explanation:
To meet the requirements, the DevOps engineer should do the following:
➢ Turn on the Multi-AZ option on the Aurora cluster.
➢ Update the application to use the Aurora cluster endpoint for write operations.

> Update the Aurora cluster's reader endpoint for reads.

Turning on the Multi-AZ option will create a replica of the database in a different Availability Zone. This will ensure that the database remains available even if one of the Availability Zones is unavailable.

Updating the application to use the Aurora cluster endpoint for write operations will ensure that all writes are sent to both the primary and replica databases. This will ensure that the data is always consistent.

Updating the Aurora cluster's reader endpoint for reads will allow the application to read data from the replica database. This will improve the performance of the application during the maintenance window.

NEW QUESTION 117
A DevOps engineer is designing an application that integrates with a legacy REST API. The application has an AWS Lambda function that reads records from an Amazon Kinesis data stream. The Lambda function sends the records to the legacy REST API.

Approximately 10% of the records that the Lambda function sends from the Kinesis data stream have data errors and must be processed manually. The Lambda function event source configuration has an Amazon Simple Queue Service (Amazon SQS) dead-letter queue as an on-failure destination. The DevOps engineer has configured the Lambda function to process records in batches and has implemented retries in case of failure.

During testing the DevOps engineer notices that the dead-letter queue contains many records that have no data errors and that already have been processed by the legacy REST API. The DevOps engineer needs to configure the Lambda function's event source options to reduce the number of errorless records that are sent to the dead-letter queue.

Which solution will meet these requirements?

A. Increase the retry attempts
B. Configure the setting to split the batch when an error occurs
C. Increase the concurrent batches per shard
D. Decrease the maximum age of record

Answer: B

Explanation:
This solution will meet the requirements because it will reduce the number of errorless records that are sent to the dead-letter queue. When you configure the setting to split the batch when an error occurs, Lambda will retry only the records that caused the error, instead of retrying the entire batch. This way, the records that have no data errors and have already been processed by the legacy REST API will not be retried and sent to the dead-letter queue unnecessarily.
https://docs.aws.amazon.com/lambda/latest/dg/with-kinesis.html

NEW QUESTION 118
......

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

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

100% Actual & Verified — Instant Download, Please Click
[Order The DOP-C02 Practice Test Here](https://www.certshared.com/exam/DOP-C02/)