Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

NEW QUESTION 1
A company runs an application on one Amazon EC2 instance. Application metadata is stored in Amazon S3 and must be retrieved if the instance is restarted. The instance must restart or relaunch automatically if the instance becomes unresponsive.
Which solution will meet these requirements?

A. Create an Amazon CloudWatch alarm for the StatusCheckFailed metri
B. Use the recover action to stop and start the instanc
C. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
D. Configure AWS OpsWorks, and use the auto healing feature to stop and start the instanc
E. Use a lifecycle event in OpsWorks to pull the metadata from Amazon S3 and update it on the instance.
F. Use EC2 Auto Recovery to automatically stop and start the instance in case of a failur
G. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
H. Use AWS CloudFormation to create an EC2 instance that includes the UserData property for the EC2 resourc
I. Add a command in UserData to retrieve the application metadata from Amazon S3.

Answer: B

Explanation:
https://aws.amazon.com/blogs/mt/how-to-set-up-aws-opsworks-stacks-auto-healing-notifications-in-amazon-cloudwatch-events/

NEW QUESTION 2
A DevOps engineer is building an application that uses an AWS Lambda function to query an Amazon Aurora MySQL DB cluster. The Lambda function performs only read queries. Amazon EventBridge events invoke the Lambda function.
As more events invoke the Lambda function each second, the database's latency increases and the database's throughput decreases. The DevOps engineer needs to improve the performance of the application.
Which combination of steps will meet these requirements? (Select THREE.)

A. Use Amazon RDS Proxy to create a prox
B. Connect the proxy to the Aurora cluster reader endpoin
C. Set a maximum connections percentage on the proxy.
D. Implement database connection pooling inside the Lambda cod
E. Set a maximum number of connections on the database connection pool.
F. Implement the database connection opening outside the Lambda event handler code.
G. Implement the database connection opening and closing inside the Lambda event handler code.
H. Connect to the proxy endpoint from the Lambda function.
I. Connect to the Aurora cluster endpoint from the Lambda function.

Answer: ACE

Explanation:
To improve the performance of the application, the DevOps engineer should use Amazon RDS Proxy, implement the database connection opening outside the Lambda event handler code, and connect to the proxy endpoint from the Lambda function. References:
? Amazon RDS Proxy is a fully managed, highly available database proxy for Amazon Relational Database Service (RDS) that makes applications more scalable, more resilient to database failures, and more secure1. By using Amazon RDS Proxy, the DevOps engineer can reduce the overhead of opening and closing connections to the database, which can improve latency and throughput2.
? The DevOps engineer should connect the proxy to the Aurora cluster reader
endpoint, which allows read-only connections to one of the Aurora Replicas in the DB cluster3. This can help balance the load across multiple read replicas and improve performance for read-intensive workloads4.
? The DevOps engineer should implement the database connection opening outside the Lambda event handler code, which means using a global variable to store the database connection object5. This can enable connection reuse across multiple invocations of the Lambda function, which can reduce latency and improve performance.
? The DevOps engineer should connect to the proxy endpoint from the Lambda function, which is a unique URL that represents the proxy. This can allow the Lambda function to access the database through the proxy, which can provide benefits such as connection pooling, load balancing, failover handling, and enhanced security.
? The other options are incorrect because:

NEW QUESTION 3
A company's DevOps engineer uses AWS Systems Manager to perform maintenance tasks during maintenance windows. The company has a few Amazon EC2 instances that require a restart after notifications from AWS Health. The DevOps engineer needs to implement an automated solution to remediate these notifications. The DevOps engineer creates an Amazon EventBridge rule.
How should the DevOps engineer configure the EventBridge rule to meet these requirements?

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

Answer: C

Explanation:
AWS Health provides real-time events and information related to your AWS infrastructure. It can be integrated with Amazon EventBridge to act upon the health events automatically. If the maintenance notification from AWS Health indicates that an EC2 instance requires a restart, you can set up an EventBridge rule to respond to such events. In this case, the target of this rule would be a Lambda function that would trigger a Systems Manager automation to restart the EC2 instance during a maintenance window. Remember, AWS Health is the source of the events (not EC2 or Systems Manager), and AWS Lambda can be used to execute complex remediation tasks, such as scheduling maintenance tasks via Systems Manager.

The following are the steps involved in configuring the EventBridge rule to meet these requirements:
? Configure an event source of AWS Health, a service of EC2, and an event type that indicates instance maintenance.
? Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.
The AWS Lambda function will be triggered by the event from AWS Health. The function will then register an automation task to restart the EC2 instance during the next maintenance window.

NEW QUESTION 4
A company is building a new pipeline by using AWS CodePipeline and AWS CodeBuild in a build account. The pipeline consists of two stages. The first stage is a CodeBuild job to build and package an AWS Lambda function. The second stage consists of deployment actions that operate on two different AWS accounts a development environment account and a production environment account. The deployment stages use the AWS Cloud Format ion action that CodePipeline invokes to deploy the infrastructure that the Lambda function requires.
A DevOps engineer creates the CodePipeline pipeline and configures the pipeline to encrypt build artifacts by using the AWS Key Management Service (AWS KMS) AWS managed key for Amazon S3 (the aws/s3 key). The artifacts are stored in an S3 bucket When the pipeline runs, the Cloud Formation actions fail with an access denied error.
Which combination of actions must the DevOps engineer perform to resolve this error? (Select TWO.)

A. Create an S3 bucket in each AWS account for the artifacts Allow the pipeline to write to the S3 bucket
B. Create a CodePipeline S3 action to copy the artifacts to the S3 bucket in each AWS account Update the CloudFormation actions to reference the artifacts S3 bucket in the production account.
C. Create a customer managed KMS key Configure the KMS key policy to allow the IAM roles used by the CloudFormation action to perform decrypt operations Modify the pipeline to use the customer managed KMS key to encrypt artifacts.
D. Create an AWS managed KMS key Configure the KMS key policy to allow the development account and the production account to perform decrypt operation
E. Modify the pipeline to use the KMS key to encrypt artifacts.
F. In the development account and in the production account create an IAM role for CodePipelin
G. Configure the roles with permissions to perform CloudFormation operations and with permissions to retrieve and decrypt objects from the artifacts S3 bucke
H. In the CodePipeline account configure the CodePipeline CloudFormation action to use the roles.
I. In the development account and in the production account create an IAM role for CodePipeline Configure the roles with permissions to perform CloudFormation operations and with permissions to retrieve and decrypt objects from the artifacts S3bucke
J. In the CodePipelme account modify the artifacts S3 bucket policy to allow the roles access Configure the CodePipeline CloudFormation action to use the roles.

Answer: BE

NEW QUESTION 5
A company is migrating its on-premises Windows applications and Linux applications to AWS. The company will use automation to launch Amazon EC2 instances to mirror the on- premises configurations. The migrated applications require access to shared storage that uses SMB for Windows and NFS for Linux.
The company is also creating a pilot light disaster recovery (DR) environment in another AWS Region. The company will use automation to launch and configure the EC2 instances in the DR Region. The company needs to replicate the storage to the DR Region.
Which storage solution will meet these requirements?

A. Use Amazon S3 for the application storag
B. Create an S3 bucket in the primary Region and an S3 bucket in the DR Regio
C. Configure S3 Cross-Region Replication (CRR) from the primary Region to the DR Region.
D. Use Amazon Elastic Block Store (Amazon EBS) for the application storag
E. Create a backup plan in AWS Backup that creates snapshots of the EBS volumes that are in the primary Region and replicates the snapshots to the DR Region.
F. Use a Volume Gateway in AWS Storage Gateway for the application storag
G. Configure Cross-Region Replication (CRR) of the Volume Gateway from the primary Region to the DR Region.
H. Use Amazon FSx for NetApp ONTAP for the application storag
I. Create an FSx for ONTAP instance in the DR Regio
J. Configure NetApp SnapMirror replication from the primary Region to the DR Region.

Answer: D

Explanation:
To meet the requirements of migrating its on-premises Windows and Linux applications to AWS and creating a pilot light DR environment in another AWS Region, the company should use Amazon FSx for NetApp ONTAP for the application storage. Amazon FSx for NetApp ONTAP is a fully managed service that provides highly reliable, scalable, high- performing, and feature-rich file storage built on NetApp's popular ONTAP file system. FSx for ONTAP supports multiple protocols, including SMB for Windows and NFS for Linux, so the company can access the shared storage from both types of applications. FSx for ONTAP also supports NetApp SnapMirror replication, which enables the company to replicate the storage to the DR Region. NetApp SnapMirror replication is efficient, secure, and incremental, and it preserves the data deduplication and compression benefits of FSx for ONTAP. The company can use automation to launch and configure the EC2 instances in the DR Region and then use NetApp SnapMirror to restore the data from the primary Region.
The other options are not correct because they do not meet the requirements or follow best practices. Using Amazon S3 for the application storage is not a good option because S3 is an object storage service that does not support SMB or NFS protocols natively. The company would need to use additional services or software to mount S3 buckets as file systems, which would add complexity and cost. Using Amazon EBS for the application storage is also not a good option because EBS is a block storage service that does not support SMB or NFS protocols natively. The company would need to set up and manage file servers on EC2 instances to provide shared access to the EBS volumes, which would add overhead and maintenance. Using a Volume Gateway in AWS Storage Gateway for the application storage is not a valid option because Volume Gateway does not support SMB protocol. Volume Gateway only supports iSCSI protocol, which means that only Linux applications can access the shared storage.
References:
? 1: What is Amazon FSx for NetApp ONTAP? - FSx for ONTAP
? 2: Amazon FSx for NetApp ONTAP
? 3: Amazon FSx for NetApp ONTAP | NetApp
? 4: AWS Announces General Availability of Amazon FSx for NetApp ONTAP
? : Replicating Data with NetApp SnapMirror - FSx for ONTAP
? : What Is Amazon S3? - Amazon Simple Storage Service
? : What Is Amazon Elastic Block Store (Amazon EBS)? - Amazon Elastic Compute Cloud
? : What Is AWS Storage Gateway? - AWS Storage Gateway

NEW QUESTION 6
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 7
A company is developing an application that will generate log events. The log events consist of five distinct metrics every one tenth of a second and produce a large amount of data The company needs to configure the application to write the logs to Amazon Time stream The company will configure a daily query against the Timestream table.
Which combination of steps will meet these requirements with the FASTEST query performance? (Select THREE.)

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

Answer: ADF

Explanation:
A comprehensive and detailed explanation is:
? Option A is correct because using batch writes to write multiple log events in a single write operation is a recommended practice for optimizing the performance and cost of data ingestion in Timestream. Batch writes can reduce the number of network round trips and API calls, and can also take advantage of parallel processing by Timestream. Batch writes can also improve the compression ratio of data in the memory store and the magnetic store, which can reduce the storage costs and improve the query performance1.
? Option B is incorrect because writing each log event as a single write operation is not a recommended practice for optimizing the performance and cost of data ingestion in Timestream. Writing each log event as a single write operation would increase the number of network round trips and API calls, and would also reduce the compression ratio of data in the memory store and the magnetic store. This would increase the storage costs and degrade the query performance1.
? Option C is incorrect because treating each log as a single-measure record is not a recommended practice for optimizing the query performance in Timestream. Treating each log as a single-measure record would result in creating multiple records for each timestamp, which would increase the storage size and the query latency. Moreover, treating each log as a single-measure record would require using joins to query multiple measures for the same timestamp, which would add complexity and overhead to the query processing2.
? Option D is correct because treating each log as a multi-measure record is a recommended practice for optimizing the query performance in Timestream. Treating each log as a multi-measure record would result in creating a single record for each timestamp, which would reduce the storage size and the query latency. Moreover, treating each log as a multi-measure record would allow querying multiple measures for the same timestamp without using joins, which would simplify and speed up the query processing2.
? Option E is incorrect because configuring the memory store retention period to be longer than the magnetic store retention period is not a valid option in Timestream. The memory store retention period must always be shorter than or equal to the magnetic store retention period. This ensures that data is moved from the memory store to the magnetic store before it expires out of the memory store3.
? Option F is correct because configuring the memory store retention period to be shorter than the magnetic store retention period is a valid option in Timestream. The memory store retention period determines how long data is kept in the memory store, which is optimized for fast point-in-time queries. The magnetic store retention period determines how long data is kept in the magnetic store, which is optimized for fast analytical queries. By configuring these retention periods appropriately, you can balance your storage costs and query performance according to your application needs3.
References:
? 1: Batch writes
? 2: Multi-measure records vs. single-measure records
? 3: Storage

NEW QUESTION 8
A development team uses AWS CodeCommit, AWS CodePipeline, and AWS CodeBuild to develop and deploy an application. Changes to the code are submitted by pull requests. The development team reviews and merges the pull requests, and then the pipeline builds
and tests the application.
Over time, the number of pull requests has increased. The pipeline is frequently blocked because of failing tests. To prevent this blockage, the development team wants to run the unit and integration tests on each pull request before it is merged.
Which solution will meet these requirements?

A. Create a CodeBuild project to run the unit and integration test
B. Create a CodeCommit approval rule templat
C. Configure the template to require the successful invocation of the CodeBuild projec
D. Attach the approval rule to the project's CodeCommit repository.
E. Create an Amazon EventBridge rule to match pullRequestCreated events from CodeCommit Create a CodeBuild project to run the unit and integration test

F. Configure the CodeBuild project as a target of the EventBridge rule that includes a custom event payload with the CodeCommit repository and branch information from the event.
G. Create an Amazon EventBridge rule to match pullRequestCreated events from CodeCommi
H. Modify the existing CodePipeline pipeline to not run the deploy steps if the build is started from a pull reques
I. Configure the EventBridge rule to run the pipeline with a custom payload that contains the CodeCommit repository and branch information from the event.
J. Create a CodeBuild project to run the unit and integration test
K. Create a CodeCommit notification rule that matches when a pull request is created or update
L. Configure the notification rule to invoke the CodeBuild project.

Answer: B

Explanation:
CodeCommit generates events in CloudWatch, CloudWatch triggers the CodeBuild https://aws.amazon.com/es/blogs/devops/complete-ci-cd-with-aws-codecommit-aws-codebuild-aws-codedeploy-and-aws-codepipeline/

NEW QUESTION 9
A company is launching an application that stores raw data in an Amazon S3 bucket. Three applications need to access the data to generate reports. The data must be redacted differently for each application before
the applications can access the data.
Which solution will meet these requirements?

A. Create an S3 bucket for each applicatio
B. Configure S3 Same-Region Replication (SRR) from the raw data's S3 bucket to each application's S3 bucke
C. Configure each application to consume data from its own S3 bucket.
D. Create an Amazon Kinesis data strea
E. Create an AWS Lambda function that isinvoked by object creation events in the raw data's S3 bucke
F. Program the Lambda function to redact data for each applicatio
G. Publish the data on the Kinesis data strea
H. Configure each application to consume data from the Kinesis data stream.
I. For each application, create an S3 access point that uses the raw data's S3 bucket as the destinatio
J. Create an AWS Lambda function that is invoked by object creation events in the raw data's S3 bucke
K. Program the Lambda function to redact data for each applicatio
L. Store the data in each application's S3 access poin
M. Configure each application to consume data from its own S3 access point.
N. Create an S3 access point that uses the raw data's S3 bucket as the destinatio
O. For each application, create an S3 Object Lambda access point that uses the S3 access poin
P. Configure the AWS Lambda function for each S3 Object Lambda access point to redact data when objects are retrieve
Q. Configure each application to consume data from its own S3 Object Lambda access point.

Answer: D

Explanation:
? The best solution is to use S3 Object Lambda1, which allows you to add your own code to S3 GET, LIST, and HEAD requests to modify and process data as it is returned to an application2. This way, you can redact the data differently for each application without creating and storing multiple copies of the data or running proxies.
? The other solutions are less efficient or scalable because they require replicating the data to multiple buckets, streaming the data through Kinesis, or storing the data in S3 access points.
References: 1: Amazon S3 Features | Object Lambda | AWS 2: Transforming objects with S3 Object Lambda - Amazon Simple Storage Service

NEW QUESTION 10
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 pull request when the test results are complete.

Answer: C

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

NEW QUESTION 10
A company builds a container image in an AWS CodeBuild project by running Docker commands. After the container image is built, the CodeBuild project uploads the container image to an Amazon S3 bucket. The CodeBuild project has an 1AM service role that has permissions to access the S3 bucket.
A DevOps engineer needs to replace the S3 bucket with an Amazon Elastic Container Registry (Amazon ECR) repository to store the container images. The

DevOps engineer creates an ECR private image repository in the same AWS Region of the CodeBuild project. The DevOps engineer adjusts the 1AM service role with the permissions that are necessary to work with the new ECR repository. The DevOps engineer also places new repository information into the docker build command and the docker push command that are used in the buildspec.yml file.
When the CodeBuild project runs a build job, the job fails when the job tries to access the ECR repository.
Which solution will resolve the issue of failed access to the ECR repository?

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

Answer: A

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

NEW QUESTION 14
A DevOps engineer is planning to deploy a Ruby-based application to production. The application needs to interact with an Amazon RDS for MySQL database and should have automatic scaling and high availability. The stored data in the database is critical and should persist regardless of the state of the application stack.
The DevOps engineer needs to set up an automated deployment strategy for the application with automatic rollbacks. The solution also must alert the application team when a deployment fails.
Which combination of steps will meet these requirements? (Select THREE.)

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

Answer: BDE

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

NEW QUESTION 16
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 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 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 22
A company's production environment uses an AWS CodeDeploy blue/green deployment to deploy an application. The deployment incudes Amazon EC2 Auto Scaling groups that launch instances that run Amazon Linux 2.
A working appspec. ymi file exists in the code repository and contains the following text.

```
version: 0.0
os: linux
files:
  - source: /
    destination: /var/www/html/application
```

A DevOps engineer needs to ensure that a script downloads and installs a license file onto the instances before the replacement instances start to handle request traffic. The DevOps engineer adds a hooks section to the appspec. yml file.
Which hook should the DevOps engineer use to run the script that downloads and installs the license file?

A. AfterBlockTraffic
B. BeforeBlockTraffic
C. BeforeInstall
D. Down load Bundle

Answer: C

Explanation:
This hook runs before the new application version is installed on the replacement instances. This is the best place to run the script because it ensures that the license file is downloaded and installed before the replacement instances start to handle request traffic. If you use any other hook, you may encounter errors or inconsistencies in your application.

NEW QUESTION 23
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?

A. Modify the Kinesis consumer application to store the logs durably in Amazon S3 Use Amazon EMR to process the data directly on Amazon S3 to derive customer insights Store the results in Amazon S3.
B. Horizontally scale the Kinesis consumer application by adding more EC2 instances based on the Amazon CloudWatch GetRecords IteratorAgeMilliseconds metric Increase the retention period of the Kinesis data streams.
C. Convert the Kinesis consumer application to run as an AWS Lambda functio
D. Configure the Kinesis data streams as the event source for the Lambda function to process the data streams
E. Increase the number of shards in the Kinesis data streams to increase the overall throughput so that the consumer application processes the data faster.

Answer: B

Explanation:
https://docs.aws.amazon.com/streams/latest/dev/monitoring-with- cloudwatch.html
GetRecords.IteratorAgeMilliseconds - The age of the last record in all GetRecords calls made against a Kinesis stream, measured over the specified time period. Age is the difference between the current time and when the last record of the GetRecords call was written to the stream. The Minimum and Maximum statistics can be used to track the progress of Kinesis consumer applications. A value of zero indicates that the records being read are completely caught up.

NEW QUESTION 28
A company is using an AWS CodeBuild project to build and package an application. The packages are copied to a shared Amazon S3 bucket before being deployed across multiple AWS accounts.
The buildspec.yml file contains the following:

```
version: 0.2
phases:
  build:
    commands:
      - go build -o myapp
  post_build:
    commands:
      - aws s3 cp --acl authenticated-read myapp s3://artifacts/
```

The DevOps engineer has noticed that anybody with an AWS account is able to download the artifacts.
What steps should the DevOps engineer take to stop this?

A. Modify the post_build command to use --acl public-read and configure a bucket policy that grants read access to the relevant AWS accounts only.
B. Configure a default ACL for the S3 bucket that defines the set of authenticated users as the relevant AWS accounts only and grants read-only access.
C. Create an S3 bucket policy that grants read access to the relevant AWS accounts and denies read access to the principal “*”.
D. Modify the post_build command to remove --acl authenticated-read and configure a bucket policy that allows read access to the relevant AWS accounts only.

Answer: D

Explanation:
When setting the flag authenticated-read in the command line, the owner gets FULL_CONTROL. The AuthenticatedUsers group (Anyone with an AWS account) gets READ access. Reference: https://docs.aws.amazon.com/AmazonS3/latest/userguide/acl-overview.html

NEW QUESTION 31
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 bucke
M. 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 33
A company manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The EC2 instances run in an Auto Scaling group across multiple Availability Zones. The application uses an Amazon RDS for MySQL DB instance to store the data. The company has configured Amazon Route 53 with an alias record that points to the ALB.
A new company guideline requires a geographically isolated disaster recovery (DR> site with an RTO of 4 hours and an RPO of 15 minutes.
Which DR strategy will meet these requirements with the LEAST change to the application stack?

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

Answer: D

NEW QUESTION 37
A company needs to implement failover for its application. The application includes an Amazon CloudFront distribution and a public Application Load Balancer (ALB) in an AWS Region. The company has configured the ALB as the default origin for the distribution.
After some recent application outages, the company wants a zero-second RTO. The company deploys the application to a secondary Region in a warm standby configuration. A DevOps engineer needs to automate the failover of the application to the secondary Region so that HTTP GET requests meet the desired RTO.
Which solution will meet these requirements?

A. Create a second CloudFront distribution that has the secondary ALB as the default origi
B. Create Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both CloudFront distribution
C. Update the application to use the new record set.
D. Create a new origin on the distribution for the secondary AL
E. Create a new origin grou
F. Set the original ALB as the primary origi
G. Configure the origin group to fail over for HTTP 5xx status code
H. Update the default behavior to use the origin group.
I. Create Amazon Route 53 alias records that have a failover policy and Evaluate TargetHealth set to Yes for both ALB
J. Set the TTL of both records to 0. Update the distribution's origin to use the new record set.
K. Create a CloudFront function that detects HTTP 5xx status code
L. Configure the function to return a 307 Temporary Redirect error response to the secondary ALB if the function detects 5xx status code
M. Update the distribution's default behavior to send origin responses to the function.

Answer: B

Explanation:
The best solution to implement failover for the application is to use CloudFront origin groups. Origin groups allow CloudFront to automatically switch to a secondary origin when the primary origin is unavailable or returns specific HTTP status codes that indicate a failure1. This way, CloudFront can serve the requests from the secondary ALB in the secondary Region without any delay or redirection. To set up origin groups, the DevOps engineer needs to create a new origin on the distribution for the secondary ALB, create a new origin group with the original ALB as the primary origin and the secondary ALB as the secondary origin, and configure the origin group to fail over for HTTP 5xx status
codes. Then, the DevOps engineer needs to update the default behavior to use the origin group instead of the single origin2.
The other options are not as effective or efficient as the solution in option B. Option A is not suitable because creating a second CloudFront distribution will

increase the complexity and cost of the application. Moreover, using Route 53 alias records with a failover policy will introduce some delay in detecting and switching to the secondary CloudFront distribution, which may not meet the zero-second RTO requirement. Option C is not feasible because CloudFront does not support using Route 53 alias records as origins3. Option D is not advisable because using a CloudFront function to redirect the requests to the secondary ALB will add an extra round-trip and latency to the failover process, which may also not meet the zero-second RTO requirement.
References:
? 1: Optimizing high availability with CloudFront origin failover - Amazon CloudFront
? 2: Creating an origin group - Amazon CloudFront
? 3: Values That You Specify When You Create or Update a Web Distribution - Amazon CloudFront

NEW QUESTION 40
A company uses AWS Storage Gateway in file gateway mode in front of an Amazon S3 bucket that is used by multiple resources. In the morning when business begins, users do not see the objects processed by a third party the previous evening. When a DevOps engineer looks directly at the S3 bucket, the data is there, but it is missing in Storage Gateway.
Which solution ensures that all the updated third-party files are available in the morning?

A. Configure a nightly Amazon EventBridge event to invoke an AWS Lambda function to run the RefreshCache command for Storage Gateway.
B. Instruct the third party to put data into the S3 bucket using AWS Transfer for SFTP.
C. Modify Storage Gateway to run in volume gateway mode.
D. Use S3 Same-Region Replication to replicate any changes made directly in the S3 bucket to Storage Gateway.

Answer: A

Explanation:
https://docs.aws.amazon.com/storagegateway/latest/APIReference/API_RefreshCache.ht ml " It only updates the cached inventory to reflect changes in the inventory of the objects in the S3 bucket. This operation is only supported in the S3 File Gateway types."

NEW QUESTION 42
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 45
A business has an application that consists of five independent AWS Lambda functions.
The DevOps engineer has built a CI/CD pipeline using AWS CodePipeline and AWS CodeBuild that builds tests packages and deploys each Lambda function in sequence. The pipeline uses an Amazon EventBridge rule to ensure the pipeline starts as quickly as possible after a change is made to the application source code.
After working with the pipeline for a few months the DevOps engineer has noticed the pipeline takes too long to complete.
What should the DevOps engineer implement to BEST improve the speed of the pipeline?

A. Modify the CodeBuild projects within the pipeline to use a compute type with more available network throughput.
B. Create a custom CodeBuild execution environment that includes a symmetricmultiprocessing configuration to run the builds in parallel.
C. Modify the CodePipeline configuration to run actions for each Lambda function in parallel by specifying the same runorder.
D. Modify each CodeBuild protect to run within a VPC and use dedicated instances to increase throughput.

Answer: C

Explanation:
https://docs.aws.amazon.com/codepipeline/latest/userguide/reference- pipeline-structure.html
AWS doc: "To specify parallel actions, use the same integer for each action you want to run in parallel. For example, if you want three actions to run in sequence in a stage, you would give the first action the runOrder value of 1, the second action the runOrder value of 2, and the third the runOrder value of 3. However, if you want the second and third actions to run in parallel, you would give the first action the runOrder value of 1 and both the second and third actions the runOrder value of 2."

NEW QUESTION 46
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 hou
E. Create an Amazon EventBridge rule tor AWS Config rules compliance chang
F. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLoggmg was calle
G. Add the Lambda function ARN as a target to the EventBridge rule.
H. Create an Amazon EventBridge rule for a scheduled event every 5 minute
I. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on a CloudTrail trail in the AWS accoun
J. Add the Lambda function ARN as a target to the EventBridge rule.
K. Launch a t2 nano instance with a script running every 5 minutes that uses the AWS SDK to query CloudTrail in the current accoun
L. 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 48
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-dynamodb-aws-lambda- and-amazon-cloudwatch-events/ And Youtube video showing how to restrict resources per user with portfolio https://www.youtube.com/watch?v=LzvhTcqqyog

NEW QUESTION 52
A company uses AWS Organizations to manage multiple accounts. Information security policies require that all unencrypted Amazon EBS volumes be marked as non-compliant. A DevOps engineer needs to automatically deploy the solution and ensure that this compliance check is always present.
Which solution will accomplish this?

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

Answer: B

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

NEW QUESTION 56
A company updated the AWS Cloud Formation template for a critical business application. The stack update process failed due to an error in the updated template and AWS CloudFormation automatically began the stack rollback process Later a DevOps engineer discovered that the application was still unavailable and that the stack was in the UPDATE_ROLLBACK_FAILED state.
Which combination of actions should the DevOps engineer perform so that the stack rollback can complete successfully? (Select TWO.)

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

Answer: CD

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

NEW QUESTION 61
A developer is maintaining a fleet of 50 Amazon EC2 Linux servers. The servers are part of an Amazon EC2 Auto Scaling group, and also use Elastic Load Balancing for load balancing.

Occasionally, some application servers are being terminated after failing ELB HTTP health checks. The developer would like to perform a root cause analysis on the issue, but before being able to access application logs, the server is terminated.
How can log collection be automated?

A. Use Auto Scaling lifecycle hooks to put instances in a Pending:Wait stat
B. Create an Amazon CloudWatch alarm for EC2 Instance Terminate Successful and trigger an AWS Lambda function that invokes an SSM Run Command script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
C. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait stat
D. Create an AWS Config rule for EC2 Instance-terminate Lifecycle Action and trigger a step function that invokes a script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
E. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait stat
F. Create an Amazon CloudWatch subscription filter for EC2 Instance Terminate Successful and trigger a CloudWatch agent that invokes a script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
G. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait stat
H. Create an Amazon EventBridge rule for EC2 Instance-terminate Lifecycle Action and trigger an AWS Lambda function that invokes an SSM Run Command script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.

Answer: D

Explanation:
https://blog.fourninecloud.com/auto-scaling-lifecycle-hooks-to-export-server-logs-when-instance-terminating-58e06d7c0d6a

NEW QUESTION 65
An online retail company based in the United States plans to expand its operations to Europe and Asia in the next six months. Its product currently runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. All data is stored in an Amazon Aurora database instance.
When the product is deployed in multiple regions, the company wants a single product catalog across all regions, but for compliance purposes, its customer information and purchases must be kept in each region.
How should the company meet these requirements with the LEAST amount of application changes?

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

Answer: C

NEW QUESTION 69
A company builds a container image in an AWS CodeBuild project by running Docker commands. After the container image is built, the CodeBuild project uploads the container image to an Amazon S3 bucket. The CodeBuild project has an IAM service role that has permissions to access the S3 bucket.
A DevOps engineer needs to replace the S3 bucket with an Amazon Elastic Container Registry (Amazon ECR) repository to store the container images. The DevOps engineer creates an ECR private image repository in the same AWS Region of the CodeBuild project. The DevOps engineer adjusts the IAM service role with the permissions that are necessary to work with the new ECR repository. The DevOps engineer also places
new repository information into the docker build command and the docker push command that are used in the buildspec.yml file.
When the CodeBuild project runs a build job, the job fails when the job tries to access the ECR repository.
Which solution will resolve the issue of failed access to the ECR repository?

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

Answer: A

Explanation:
Update the buildspec.yml file to log in to the ECR repository by using the aws ecr get-login- password AWS CLI command to obtain an authentica-tion token.
Update the docker login command to use the authentication token to access the ECR repository.
This is the correct solution. The aws ecr get-login-password AWS CLI command retrieves and displays an authentication token that can be used to log in to an ECR repository. The docker login command can use this token as a password to authenticate with the ECR repository. This way, the CodeBuild project can push and pull images from the ECR repository without any errors. For more information, see Using Amazon ECR with the AWS CLI and get-login-password.

NEW QUESTION 73
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 78
A DevOps engineer has implemented a Cl/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 79
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-xxxxxxxxxx"
                    ]
                }
            }
        }
    ]
}
```

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

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

Answer: AD

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

NEW QUESTION 84
A highly regulated company has a policy that DevOps engineers should not log in to their Amazon EC2 instances except in emergencies. It a DevOps engineer does log in the security team must be notified within 15 minutes of the occurrence.
Which solution will meet these requirements'?

A. Install the Amazon Inspector agent on each EC2 instance Subscribe to Amazon EventBridge notifications Invoke an AWS Lambda function to check if a message is about user logins If it is send a notification to the security team using Amazon SNS.
B. Install the Amazon CloudWatch agent on each EC2 instance Configure the agent to push all logs to Amazon CloudWatch Logs and set up a CloudWatch metric filter that searches for user login
C. If a login is found send a notification to the security team using Amazon SNS.
D. Set up AWS CloudTrail with Amazon CloudWatch Log
E. Subscribe CloudWatch Logs to Amazon Kinesis Attach AWS Lambda to Kinesis to parse and determine if a log contains a user login If it does, send a notification to the security team using Amazon SNS.
F. Set up a script on each Amazon EC2 instance to push all logs to Amazon S3 Set up an S3 event to invoke an AWS Lambda function which invokes an Amazon Athena query to ru
G. The Athena query checks tor logins and sends the output to the security team using Amazon SNS.

Answer: B

Explanation:
https://aws.amazon.com/blogs/security/how-to-monitor-and-visualize-failed-ssh-access-attempts-to-amazon-ec2-linux-instances/

NEW QUESTION 85
A company is hosting a web application in an AWS Region. For disaster recovery purposes, a second region is being used as a standby. Disaster recovery requirements state that session data must be replicated between regions in near-real time and 1% of requests should route to the secondary region to continuously verify system functionality. Additionally, if there is a disruption in service in the main region, traffic should be automatically routed to the secondary region, and the secondary region must be able to
scale up to handle all traffic.
How should a DevOps engineer meet these requirements?

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

Answer: D

NEW QUESTION 86
A company uses an organization in AWS Organizations that has all features enabled. The company uses AWS Backup in a primary account and uses an AWS Key Management Service (AWS KMS) key to encrypt the backups.
The company needs to automate a cross-account backup of the resources that AWS Backup backs up in the primary account. The company configures cross-account backup in the Organizations management account. The company creates a new AWS account in the organization and configures an AWS Backup backup vault in the new account. The company creates a KMS key in the new account to encrypt the backups. Finally, the company configures a new backup plan in the primary account. The destination for the new backup plan is the backup vault in the new account.
When the AWS Backup job in the primary account is invoked, the job creates backups in the primary account. However, the backups are not copied to the new account's backup vault.
Which combination of steps must the company take so that backups can be copied to the new account's backup vault? (Select TWO.)

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

Answer: AE

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

NEW QUESTION 91

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

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

Answer: ABE

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

NEW QUESTION 96
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 97
A company needs to ensure that flow logs remain configured for all existing and new VPCs in its AWS account. The company uses an AWS CloudFormation stack to manage its VPCs. The company needs a solution that will work for any VPCs that any IAM user creates.
Which solution will meet these requirements?

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

Answer: C

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

? 2: Amazon VPC Flow Logs extends CloudFormation Support to custom format subscriptions, 1-minute aggregation intervals and tagging
? 3: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud
? : About AWS Config - AWS Config
? : vpc-flow-logs-enabled - AWS Config
? : Remediate Noncompliant Resources with AWS Config Rules - AWS Config

NEW QUESTION 98
A company that runs many workloads on AWS has an Amazon EBS spend that has increased over time. The DevOps team notices there are many unattached EBS volumes. Although there are workloads where volumes are detached, volumes over 14 days old are stale and no longer needed. A DevOps engineer has been tasked with creating automation that deletes unattached EBS volumes that have been unattached for 14 days.
Which solution will accomplish this?

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

Answer: C

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

NEW QUESTION 100
A DevOps engineer is working on a data archival project that requires the migration of on- premises data to an Amazon S3 bucket. The DevOps engineer develops a script that incrementally archives on-premises data that is older than 1 month to Amazon S3. Data that is transferred to Amazon S3 is deleted from the on-premises location The script uses the S3 PutObject operation.
During a code review the DevOps engineer notices that the script does not verity whether the data was successfully copied to Amazon S3. The DevOps engineer must update the script to ensure that data is not corrupted during transmission. The script must use MD5 checksums to verify data integrity before the on-premises data is deleted.
Which solutions for the script will meet these requirements'? (Select TWO.)

A. Check the returned response for the Versioned Compare the returned Versioned against the MD5 checksum.
B. Include the MD5 checksum within the Content-MD5 paramete
C. Check the operationcall's return status to find out if an error was returned.
D. Include the checksum digest within the tagging parameter as a URL query parameter.
E. Check the returned response for the ETa
F. Compare the returned ETag against the MD5 checksum.
G. Include the checksum digest within the Metadata parameter as a name-value pair After upload use the S3 HeadObject operation to retrieve metadata from the object.

Answer: BD

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/checking-object- integrity.html

NEW QUESTION 105
A company has an application that is using a MySQL-compatible Amazon Aurora Multi-AZ DB cluster as the database. A cross-Region read replica has been created for disaster recovery purposes. A DevOps engineer wants to automate the promotion of the replica so it becomes the primary database instance in the event of a failure.
Which solution will accomplish this?

A. Configure a latency-based Amazon Route 53 CNAME with health checks so it points to both the primary and replica endpoint
B. Subscribe an Amazon SNS topic to Amazon RDS failure notifications from AWS CloudTrail and use that topic to invoke an AWS Lambda function that will promote the replica instance as the primary.
C. Create an Aurora custom endpoint to point to the primary database instanc
D. Configure the application to use this endpoin
E. Configure AWS CloudTrail to run an AWS Lambda function to promote the replica instance and modify the custom endpoint to point to the newly promoted instance.
F. Create an AWS Lambda function to modify the application's AWS CloudFormation template to promote the replica, apply the template to update the stack, and point the application to the newly promoted instanc
G. Create an Amazon CloudWatch alarm to invoke this Lambda function after the failure event occurs.
H. Store the Aurora endpoint in AWS Systems Manager Parameter Stor
I. Create an Amazon EventBridge event that detects the database failure and runs an AWS Lambda function to promote the replica instance and update the

endpoint URL stored in AWS Systems Manager Parameter Stor
J. Code the application to reload the endpoint from Parameter Store if a database connection fails.

Answer: D

Explanation:
EventBridge is needed to detect the database failure. Lambda is needed to promote the replica as it's in another Region (manual promotion, otherwise). Storing and updating the endpoint in Parameter store is important in updating the application. Look at High Availability section of Aurora FAQ: https://aws.amazon.com/rds/aurora/faqs/

NEW QUESTION 110
A company uses AWS Organizations to manage its AWS accounts. The company has a root OU that has a child OU. The root OU has an SCP that allows all actions on all resources. The child OU has an SCP that allows all actions for Amazon DynamoDB and AWS Lambda, and denies all other actions.
The company has an AWS account that is named vendor-data in the child OU. A DevOps engineer has an 1AM user that is attached to the AdministratorAccess 1AM policy in the vendor-data account. The DevOps engineer attempts to launch an Amazon EC2 instance in the vendor-data account but receives an access denied error.
Which change should the DevOps engineer make to launch the EC2 instance in the vendor-data account?

A. Attach the AmazonEC2FullAccess 1AM policy to the 1AM user.
B. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the vendor-data account.
C. Update the SCP in the child OU to allow all actions for Amazon EC2.
D. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the root OU.

Answer: C

Explanation:
The correct answer is C. Updating the SCP in the child OU to allow all actions for Amazon EC2 will enable the DevOps engineer to launch the EC2 instance in the vendor-data account. SCPs are applied to OUs and accounts in a hierarchical manner, meaning that the SCPs attached to the parent OU are inherited by the child OU and accounts. Therefore, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. By adding EC2 to the allowed actions in the child OU's SCP, the DevOps engineer can access EC2 resources in the vendor-data account.
Option A is incorrect because attaching the AmazonEC2FullAccess IAM policy to the IAM user will not grant the user access to EC2 resources. IAM policies are evaluated after SCPs, so even if the IAM policy allows EC2 actions, the SCP will still deny them.
Option B is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the vendor-data account will not work. SCPs are not cumulative, meaning that only one SCP is applied to an account at a time. The SCP attached to the account will be the SCP attached to the OU that contains the account. Therefore, option B will not change the SCP that is applied to the vendor-data account.
Option D is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the root OU will not work. As explained earlier, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. Therefore, option D will not affect the SCP that is applied to the vendor-data account.

NEW QUESTION 115
A DevOps engineer is working on a project that is hosted on Amazon Linux and has failed a security review. The DevOps manager has been asked to review the company buildspec.yaml die for an AWS CodeBuild project and provide recommendations. The buildspec. yaml file is configured as follows:

```
env:
  variables:
    AWS_ACCESS_KEY_ID: AKIAJF7BRFWJBA4GHXNA
    AWS_SECRET_ACCESS_KEY: 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
    AWS_DEFAULT_REGION: us-east-1
    DB_PASSWORD: cuj5RptFa3va
phases:
  build:
    commands:
      - aws s3 cp s3://db-deploy-bucket/my.cnf.template /tmp/my.cnf
      - sed -i '' s/DB_PW/${DB_PASSWORD}/ /tmp/my.cnf
      - aws s3 cp s3://db-deploy-bucket/instance.key /tmp/instance.key
      - chmod 600 /tmp/instance.key
      - scp -i /tmp/instance.key /tmp/my.cnf root@10.25.15.23:/etc/my.cnf
      - ssh -i /tmp/instance.key root@10.25.15.23 /etc/init.d/mysqld restart
```

What changes should be recommended to comply with AWS security best practices? (Select THREE.)

A. Add a post-build command to remove the temporary files from the container before termination to ensure they cannot be seen by other CodeBuild users.
B. Update the CodeBuild project role with the necessary permissions and then remove the AWS credentials from the environment variable.
C. Store the db_password as a SecureString value in AWS Systems Manager Parameter Store and then remove the db_password from the environment variables.
D. Move the environment variables to the 'db.-deploy-bucket 'Amazon S3 bucket, add a prebuild stage to download then export the variables.
E. Use AWS Systems Manager run command versus sec and ssh commands directly to the instance.

Answer: BCE

Explanation:
B. Update the CodeBuild project role with the necessary permissions and then remove the AWS credentials from the environment variable. C. Store the DB_PASSWORD as a SecureString value in AWS Systems Manager Parameter Store and then remove the DB_PASSWORD from the environment variables. E. Use AWS Systems Manager run command versus scp and ssh commands directly to the instance.

NEW QUESTION 116
A company's security policies require the use of security hardened AMIS in production environments. A DevOps engineer has used EC2 Image Builder to create a pipeline that builds the AMIs on a recurring schedule.
The DevOps engineer needs to update the launch templates of the companys Auto Scaling groups. The Auto Scaling groups must use the newest AMIS during the launch of Amazon EC2 instances.
Which solution will meet these requirements with the MOST operational efficiency?

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

Answer: C

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

NEW QUESTION 120
A DevOps engineer needs to configure a blue green deployment for an existing three-tier application. The application runs on Amazon EC2 instances and uses an Amazon RDS database The EC2 instances run behind an Application Load Balancer (ALB) and are in an Auto Scaling group.
The DevOps engineer has created a launch template and an Auto Scaling group for the blue environment. The DevOps engineer also has created a launch template and an Auto Scaling group for the green environment. Each Auto Scaling group deploys to a matching blue or green target group. The target group also specifies which software blue or green gets loaded on the EC2 instances. The ALB can be configured to send traffic to the blue environments target group or the green environments target group. An Amazon Route 53 record for www example com points to the ALB.
The deployment must move traffic all at once between the software on the blue environment's EC2 instances to the newly deployed software on the green environments EC2 instances
What should the DevOps engineer do to meet these requirements?

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

Answer: A

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

NEW QUESTION 125
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 S3 bucket, 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 126
An ecommerce company is receiving reports that its order history page is experiencing delays in reflecting the processing status of orders. The order processing system consists of an AWS Lambda function that uses reserved concurrency. The Lambda function processes order messages from an Amazon Simple Queue Service (Amazon SQS) queue and inserts processed orders into an Amazon DynamoDB table. The DynamoDB table has auto scaling enabled for read and write capacity.
Which actions should a DevOps engineer take to resolve this delay? (Choose two.)

A. Check the ApproximateAgeOfOldestMessage metric for the SQS queu

B. Increase the Lambda function concurrency limit.
C. Check the ApproximateAgeOfOldestMessage metnc for the SQS queue Configure a redrive policy on the SQS queue.
D. Check the NumberOfMessagesSent metric for the SQS queu
E. Increase the SQS queue visibility timeout.
F. Check the WriteThrottleEvents metric for the DynamoDB tabl
G. Increase the maximum write capacity units (WCUs) for the table's scaling policy.
H. Check the Throttles metric for the Lambda functio
I. Increase the Lambda function timeout.

Answer: AD

Explanation:
A: If the ApproximateAgeOfOldestMessages indicate that orders are remaining in the SQS queue for longer than expected, the reserved concurrency limit may be set too small to keep up with the number of orders entering the queue and is being throttled. D: The DynamoDB table is using Auto Scaling. With Auto Scaling, you create a scaling policy that specifies whether you want to scale read capacity or write capacity (or both), and the minimum and maximum provisioned capacity unit settings for the table. The ThottledWriteRequests metric will indicate if there is a throttling issue on the DynamoDB table, which can be resolved by increasing the maximum write capacity units for the table's Auto Scaling policy. https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html

NEW QUESTION 129
A company requires its developers to tag all Amazon Elastic Block Store (Amazon EBS) volumes in an account to indicate a desired backup frequency. This requirement Includes EBS volumes that do not require backups. The company uses custom tags named Backup_Frequency that have values of none, dally, or weekly that correspond to the desired backup frequency. An audit finds that developers are occasionally not tagging the EBS volumes.
A DevOps engineer needs to ensure that all EBS volumes always have the Backup_Frequency tag so that the company can perform backups at least weekly unless a different value is specified.
Which solution will meet these requirements?

A. Set up AWS Config in the accoun
B. Create a custom rule that returns a compliance failure for all Amazon EC2 resources that do not have a Backup Frequency tag applied.Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
C. Set up AWS Config in the accoun
D. Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applie
E. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
F. Turn on AWS CloudTrail in the accoun
G. Create an Amazon EventBridge rule that reacts to EBS CreateVolume event
H. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekl
I. Specify the runbook as the target of the rule.
J. Turn on AWS CloudTrail in the accoun
K. Create an Amazon EventBridge rule that reacts to EBS CreateVolume events or EBS ModifyVolume event
L. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekl
M. Specify the runbook as the target of the rule.

Answer: B

Explanation:
The following are the steps that the DevOps engineer should take to ensure that all EBS volumes always have the Backup_Frequency tag so that the company can perform backups at least weekly unless a different value is specified:
? Set up AWS Config in the account.
? Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applied.
? Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
The managed rule AWS::Config::EBSVolumesWithoutBackupTag will return a compliance failure for any EBS volume that does not have the Backup_Frequency tag applied. The remediation action will then use the Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly to the EBS volume.

NEW QUESTION 131
A media company has several thousand Amazon EC2 instances in an AWS account. The company is using Slack and a shared email inbox for team communications and important updates. A DevOps engineer needs to send all AWS-scheduled EC2 maintenance notifications to the Slack channel and the shared inbox. The solution must include the instances' Name and Owner tags.
Which solution will meet these requirements?

A. Integrate AWS Trusted Advisor with AWS Config Configure a custom AWS Config rule to invoke an AWS Lambda function to publish notifications to an Amazon Simple Notification Service (Amazon SNS) topic Subscribe a Slack channel endpoint and the shared inbox to the topic.
B. Use Amazon EventBridge to monitor for AWS Health Events Configure the maintenance events to target an Amazon Simple Notification Service (Amazon SNS) topic Subscribe an AWS Lambda function to the SNS topic to send notifications to the Slack channel and the shared inbox.
C. Create an AWS Lambda function that sends EC2 maintenance notifications to the Slack channel and the shared inbox Monitor EC2 health events by using Amazon CloudWatch metrics Configure a CloudWatch alarm that invokes the Lambda function when a maintenance notification is received.
D. Configure AWS Support integration with AWS CloudTrail Create a CloudTrail lookup event to invoke an AWS Lambda function to pass EC2 maintenance notifications to Amazon Simple Notification Service (Amazon SNS) Configure Amazon SNS to target the Slack channel and the shared inbox.

Answer: B

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

NEW QUESTION 133
A company has 20 service learns Each service team is responsible for its own microservice. Each service team uses a separate AWS account for its microservice and a VPC with the 192 168 0 0/22 CIDR block. The company manages the AWS accounts with AWS Organizations.
Each service team hosts its microservice on multiple Amazon EC2 instances behind an Application Load Balancer. The microservices communicate with each other across the public internet. The company's security team has issued a new guideline that all communication between microservices must use HTTPS over private network connections and cannot traverse the public internet.
A DevOps engineer must implement a solution that fulfills these obligations and minimizes the number of changes for each service team.

Which solution will meet these requirements?

A. Create a new AWS account in AWS Organizations Create a VPC in this account and use AWS Resource Access Manager to share the private subnets of this VPC with the organization Instruct the service teams to launch a ne
B. Network Load Balancer (NLB) and EC2 instances that use the shared private subnets Use the NLB DNS names for communication between microservices.
C. Create a Network Load Balancer (NLB) in each of the microservice VPCs Use AWS PrivateLink to create VPC endpoints in each AWS account for the NLBs Create subscriptions to each VPC endpoint in each of the other AWS accounts Use the VPC endpoint DNS names for communication between microservices.
D. Create a Network Load Balancer (NLB) in each of the microservice VPCs Create VPC peering connections between each of the microservice VPCs Update the route tables for each VPC to use the peering links Use the NLB DNS names for communication between microservices.
E. Create a new AWS account in AWS Organizations Create a transit gateway in this account and use AWS Resource Access Manager to share the transit gateway with the organizatio
F. In each of the microservice VPC
G. create a transit gateway attachment to the shared transit gateway Update the route tables of each VPC to use the transit gateway Create a Network Load Balancer (NLB) in each of the microservice VPCs Use the NLB DNS names for communication between microservices.

Answer: B

Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-ranges/ Private link is the best option because Transit Gateway doesn't support overlapping CIDR ranges.

NEW QUESTION 137
A company needs to implement failover for its application. The application includes an Amazon CloudFront distribution and a public Application Load Balancer (ALB) in an AWS Region. The company has configured the ALB as the default origin for the distribution.
After some recent application outages, the company wants a zero-second RTO. The company deploys the application to a secondary Region in a warm standby configuration. A DevOps engineer needs to automate the failover of the application to the secondary Region so that HTTP GET requests meet the desired R TO. Which solution will meet these requirements?

A. Create a second CloudFront distribution that has the secondary ALB as the default origi
B. Create Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both CloudFront distribution
C. Update the application to use the new record set.
D. Create a new origin on the distribution for the secondary AL
E. Create a new origin grou
F. Set the original ALB as the primary origi
G. Configure the origin group to fail over for HTTP 5xx status code
H. Update the default behavior to use the origin group.
I. Create Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both ALB
J. Set the TTL of both records to
K. Update the distribution's origin to use the new record set.
L. Create a CloudFront function that detects HTTP 5xx status code
M. Configure the function to return a 307 Temporary Redirect error response to the secondary ALB if the function detects 5xx status code
N. Update the distribution's default behavior to send origin responses to the function.

Answer: B

Explanation:
To implement failover for the application to the secondary Region so that HTTP GET requests meet the desired RTO, the DevOps engineer should use the following solution:
? Create a new origin on the distribution for the secondary ALB. A CloudFront origin
is the source of the content that CloudFront delivers to viewers. By creating a new origin for the secondary ALB, the DevOps engineer can configure CloudFront to route traffic to the secondary Region when the primary Region is unavailable1
? Create a new origin group. Set the original ALB as the primary origin. Configure
the origin group to fail over for HTTP 5xx status codes. An origin group is a logical grouping of two origins: a primary origin and a secondary origin. By creating an origin group, the DevOps engineer can specify which origin CloudFront should use as a fallback when the primary origin fails. The DevOps engineer can also define which HTTP status codes should trigger a failover from the primary origin to the secondary origin. By setting the original ALB as the primary origin and configuring the origin group to fail over for HTTP 5xx status codes, the DevOps engineer can ensure that CloudFront will switch to the secondary ALB when the primary ALB returns server errors2
? Update the default behavior to use the origin group. A behavior is a set of rules
that CloudFront applies when it receives requests for specific URLs or file types. The default behavior applies to all requests that do not match any other behaviors. By updating the default behavior to use the origin group, the DevOps engineer can enable failover routing for all requests that are sent to the distribution3
This solution will meet the requirements because it will automate the failover of the
application to the secondary Region with zero-second RTO. When CloudFront receives an HTTP GET request, it will first try to route it to the primary ALB in the primary Region. If the primary ALB is healthy and returns a successful response, CloudFront will deliver it to the viewer. If the primary ALB is unhealthy or returns an HTTP 5xx status code, CloudFront will automatically route the request to the secondary ALB in the secondary Region and deliver its response to the viewer.
The other options are not correct because they either do not provide zero-second RTO or do not work as expected. Creating a second CloudFront distribution that has the secondary ALB as the default origin and creating Amazon Route 53 alias records that have a failover policy is not a good option because it will introduce additional latency and complexity to the solution. Route 53 health checks and DNS propagation can take several minutes or longer, which means that viewers might experience delays or errors when accessing the application during a failover event. Creating Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both ALBs and setting the TTL of both records to O is not a valid option because it will not work with CloudFront distributions. Route 53 does not support health checks for alias records that point to CloudFront distributions, so it cannot detect if an ALB behind a distribution is healthy or not. Creating a CloudFront function that detects HTTP 5xx status codes and returns a 307 Temporary Redirect error response to the secondary ALB is not a valid option because it will not provide zero-second RTO. A 307 Temporary Redirect error response tells viewers to retry their requests with a different URL, which means that viewers will have to make an additional request and wait for another response from CloudFront before reaching the secondary ALB.
References:
? 1: Adding, Editing, and Deleting Origins - Amazon CloudFront
? 2: Configuring Origin Failover - Amazon CloudFront
? 3: Creating or Updating a Cache Behavior - Amazon CloudFront

NEW QUESTION 138
A company provides an application to customers. The application has an Amazon API Gateway REST API that invokes an AWS Lambda function. On initialization,

the Lambda function loads a large amount of data from an Amazon DynamoDB table. The data load process results in long cold-start times of 8-10 seconds. The DynamoDB table has DynamoDB Accelerator (DAX) configured.
Customers report that the application intermittently takes a long time to respond to requests. The application receives thousands of requests throughout the day. In the middle of the day, the application experiences 10 times more requests than at any other time of the day. Near the end of the day, the application's request volume decreases to 10% of its normal total.
A DevOps engineer needs to reduce the latency of the Lambda function at all times of the day.
Which solution will meet these requirements?

A. Configure provisioned concurrency on the Lambda function with a concurrency value of 1. Delete the DAX cluster for the DynamoDB table.
B. Configure reserved concurrency on the Lambda function with a concurrency value of 0.
C. Configure provisioned concurrency on the Lambda functio
D. Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.
E. Configure reserved concurrency on the Lambda functio
F. Configure AWS Application Auto Scaling on the API Gateway API with a reserved concurrency maximum value of 100.

Answer: C

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

NEW QUESTION 140
A company's development team uses AVMS Cloud Formation to deploy its application resources The team must use for an changes to the environment The team cannot use AWS Management Console or the AWS CLI to make manual changes directly.
The team uses a developer IAM role to access the environment The role is configured with the Admnistratoraccess managed policy. The company has created a new Cloudformationdeployment IAM role that has the following policy.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "elasticloadbalancing:*",
                "lambda:*",
                "dynamodb:*"
            ],
            "Resource": "*"
        }
    ]
}
```

The company wants ensure that only CloudFormation can use the new role. The development team cannot make any manual changes to the deployed resources.
Which combination of steps meet these requirements? (Select THREE.)

A. Remove the AdministratorAccess polic
B. Assign the ReadOnlyAccess managed IAM policy to the developer rol
C. Instruct the developers to use the CloudFormationDeployment role as a CloudFormation service role when the developers deploy new stacks.
D. Update the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDepoyment role.
E. Configure the IAM to be to get and pass the CloudFormationDeployment role if cloudformation actions for resources,
F. Update the trust Of the CloudFormationDepoyment role to anow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeR01e action
G. Remove me Administratoraccess polic
H. Assign the ReadOnly/Access managed IAM policy to the developer role Instruct the developers to assume the CloudFormatondeployment role when the developers new stacks
I. Add an IAM policy to CloudFormationDeplyment to allow cloudformation * on an Add a policy that allows the iam.PassR01e action for ARN of if iam PassedT0Service equal cloudformation.amazonaws.com

Answer: ADF

Explanation:
A comprehensive and detailed explanation is:
? Option A is correct because removing the AdministratorAccess policy and assigning the ReadOnlyAccess managed IAM policy to the developer role is a valid way to prevent the developers from making any manual changes to the deployed resources. The AdministratorAccess policy grants full access to all AWS resources and actions, which is not necessary for the developers. The ReadOnlyAccess policy grants read-only access to most AWS resources and actions, which is sufficient for the developers to view the status of their stacks. Instructing the developers to use the CloudFormationDeployment role as a CloudFormation service role when they deploy new stacks is also a valid way to ensure that only CloudFormation can use the new role. A CloudFormation service role is an IAM role that

allows CloudFormation to make calls to resources in a stack on behalf of the user1. The user can specify a service role when they create or update a stack, and CloudFormation will use that role's credentials for all operations that are performed on that stack1.
? Option B is incorrect because updating the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDeployment role is not a valid solution. This would allow the developers to manually assume the CloudFormationDeployment role and perform actions on the deployed resources, which is not what the company wants. The trust of CloudFormationDeployment role should only allow the cloudformation.amazonaws.com AWS principal to assume the role, as in option D.
? Option C is incorrect because configuring the IAM user to be able to get and pass the CloudFormationDeployment role if cloudformation actions for resources is not a valid solution. This would allow the developers to manually pass the CloudFormationDeployment role to other services or resources, which is not what the company wants. The IAM user should only be able to pass the CloudFormationDeployment role as a service role when they create or update a stack with CloudFormation, as in option A.
? Option D is correct because updating the trust of CloudFormationDeployment role
to allow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeRole action is a valid solution. This allows CloudFormation to assume the CloudFormationDeployment role and access resources in other services on behalf of the user2. The trust policy of an IAM role defines which entities can assume the role2. By specifying cloudformation.amazonaws.com as the principal, you grant permission only to CloudFormation to assume this role.
? Option E is incorrect because instructing the developers to assume the
CloudFormationDeployment role when they deploy new stacks is not a valid solution. This would allow the developers to manually assume the CloudFormationDeployment role and perform actions on the deployed resources, which is not what the company wants. The developers should only use the CloudFormationDeployment role as a service role when they deploy new stacks with CloudFormation, as in option A.
? Option F is correct because adding an IAM policy to CloudFormationDeployment
that allows cloudformation:* on all resources and adding a policy that allows the iam:PassRole action for ARN of CloudFormationDeployment if iam:PassedToService equals cloudformation.amazonaws.com are valid solutions. The first policy grants permission for CloudFormationDeployment to perform any action with any resource using cloudformation.amazonaws.com as a service principal3. The second policy grants permission for passing this role only if it is passed by cloudformation.amazonaws.com as a service principal4. This ensures that only CloudFormation can use this role.
References:
? 1: AWS CloudFormation service roles
? 2: How to use trust policies with IAM roles
? 3: AWS::IAM::Policy
? 4: IAM: Pass an IAM role to a specific AWS service

NEW QUESTION 144
.....

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

AWS-Certified-DevOps-Engineer-Professional Practice Exam Features:

* AWS-Certified-DevOps-Engineer-Professional Questions and Answers Updated Frequently

* AWS-Certified-DevOps-Engineer-Professional Practice Questions Verified by Expert Senior Certified Staff

* AWS-Certified-DevOps-Engineer-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Certified-DevOps-Engineer-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year