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 company has an application that runs on AWS Lambda and sends logs to Amazon CloudWatch Logs. An Amazon Kinesis data stream is subscribed to the log groups in CloudWatch Logs. A single consumer Lambda function processes the logs from the data stream and stores the logs in an Amazon S3 bucket.
The company's DevOps team has noticed high latency during the processing and ingestion of some logs.
Which combination of steps will reduce the latency? (Select THREE.)

A. Create a data stream consumer with enhanced fan-ou
B. Set the Lambda function that processes the logs as the consumer.
C. Increase the ParallelizationFactor setting in the Lambda event source mapping.
D. Configure reserved concurrency for the Lambda function that processes the logs.
E. Increase the batch size in the Kinesis data stream.
F. Turn off the ReportBatchItemFailures setting in the Lambda event source mapping.
G. Increase the number of shards in the Kinesis data stream.

Answer: ABC

Explanation:
The latency in processing and ingesting logs can be caused by several factors, such as the throughput of the Kinesis data stream, the concurrency of the Lambda function, and the configuration of the event source mapping. To reduce the latency, the following steps can be taken:
? Create a data stream consumer with enhanced fan-out. Set the Lambda function that processes the logs as the consumer. This will allow the Lambda function to receive records from the data stream with dedicated throughput of up to 2 MB per second per shard, independent of other consumers1. This will reduce the contention and delay in accessing the data stream.
? Increase the ParallelizationFactor setting in the Lambda event source mapping. This will allow the Lambda service to invoke more instances of the function concurrently to process the records from the data stream2. This will increase the processing capacity and reduce the backlog of records in the data stream.
? Configure reserved concurrency for the Lambda function that processes the logs. This will ensure that the function has enough concurrency available to handle the increased load from the data stream3. This will prevent the function from being throttled by the account-level concurrency limit.
The other options are not effective or may have negative impacts on the latency. Option D is not suitable because increasing the batch size in the Kinesis data stream will increase the amount of data that the Lambda function has to process in each invocation, which may increase the execution time and latency4. Option E is not advisable because turning off the ReportBatchItemFailures setting in the Lambda event source mapping will prevent the Lambda service from retrying the failed records, which may result in data loss. Option F is not necessary because increasing the number of shards in the Kinesis data stream will increase the throughput of the data stream, but it will not affect the processing speed of the Lambda function, which is the bottleneck in this scenario.
References:
? 1: Using AWS Lambda with Amazon Kinesis Data Streams - AWS Lambda
? 2: AWS Lambda event source mappings - AWS Lambda
? 3: Managing concurrency for a Lambda function - AWS Lambda
? 4: AWS Lambda function scaling - AWS Lambda
? : AWS Lambda event source mappings - AWS Lambda
? : Scaling Amazon Kinesis Data Streams with AWS CloudFormation - Amazon Kinesis Data Streams

NEW QUESTION 3
A company deploys a web application on Amazon EC2 instances that are behind an
Application Load Balancer (ALB). The company stores the application code in an AWS CodeCommit repository. When code is merged to the main branch, an AWS Lambda function invokes an AWS CodeBuild project. The CodeBuild project packages the code, stores the packaged code in AWS CodeArtifact, and invokes AWS Systems Manager Run Command to deploy the packaged code to the EC2 instances.
Previous deployments have resulted in defects, EC2 instances that are not running the latest version of the packaged code, and inconsistencies between instances.
Which combination of actions should a DevOps engineer take to implement a more reliable deployment solution? (Select TWO.)

A. Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provide
B. Configure pipeline stages that run the CodeBuild project in parallel to build and test the applicatio
C. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action.
D. Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provide
E. Create separate pipeline stages that run a CodeBuild project to build and then test the applicatio
F. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action.
G. Create an AWS CodeDeploy application and a deployment group to deploy the packaged code to the EC2 instance
H. Configure the ALB for the deployment group.
I. Create individual Lambda functions that use AWS CodeDeploy instead of Systems Manager to run build, test, and deploy actions.
J. Create an Amazon S3 bucke
K. Modify the CodeBuild project to store the packages in the S3 bucket instead of in CodeArtifac

L. Use deploy actions in CodeDeploy to deploy the artifact to the EC2 instances.

Answer: AC

Explanation:
To implement a more reliable deployment solution, a DevOps engineer should take the following actions:
? Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provider. Configure pipeline stages that run the CodeBuild project in parallel to build and test the application. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action. This action will improve the deployment reliability by automating the entire process from code commit to deployment, reducing human errors and inconsistencies. By running the build and test stages in parallel, the pipeline can also speed up the delivery time and provide faster feedback. By using CodeDeploy as the deployment action, the pipeline can leverage the features of CodeDeploy, such as traffic shifting, health checks, rollback, and deployment configuration123
? Create an AWS CodeDeploy application and a deployment group to deploy the packaged code to the EC2 instances. Configure the ALB for the deployment group. This action will improve the deployment reliability by using CodeDeploy to orchestrate the deployment across multiple EC2 instances behind an ALB. CodeDeploy can perform blue/green deployments or in-place deployments with traffic shifting, which can minimize downtime and reduce risks. CodeDeploy can also monitor the health of the instances during and after the deployment, and automatically roll back if any issues are detected. By configuring the ALB for the deployment group, CodeDeploy can register and deregister instances from the load balancer as needed, ensuring that only healthy instances receive traffic45
The other options are not correct because they do not improve the deployment reliability or follow best practices. Creating separate pipeline stages that run a CodeBuild project to build and then test the application is not a good option because it will increase the pipeline execution time and delay the feedback loop. Creating individual Lambda functions that use CodeDeploy instead of Systems Manager to run build, test, and deploy actions is not a valid option because it will add unnecessary complexity and cost to the solution. Lambda functions are not designed for long-running tasks such as building or deploying applications. Creating an Amazon S3 bucket and modifying the CodeBuild project to store the packages in the S3 bucket instead of in CodeArtifact is not a necessary option because it will not affect the deployment reliability. CodeArtifact is a secure, scalable, and cost- effective package management service that can store and share software packages for application development67
References:
? 1: What is AWS CodePipeline? - AWS CodePipeline
? 2: Create a pipeline in AWS CodePipeline - AWS CodePipeline
? 3: Deploy an application with AWS CodeDeploy - AWS CodePipeline
? 4: What is AWS CodeDeploy? - AWS CodeDeploy
? 5: Configure an Application Load Balancer for your blue/green deployments - AWS CodeDeploy
? 6: What is AWS Lambda? - AWS Lambda
? 7: What is AWS CodeArtifact? - AWS CodeArtifact

NEW QUESTION 4
A company uses Amazon S3 to store proprietary information. The development team creates buckets for new projects on a daily basis. The security team wants to ensure that all existing and future buckets have encryption logging and versioning enabled. Additionally, no buckets should ever be publicly read or write accessible.
What should a DevOps engineer do to meet these requirements?

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

Answer: B

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

NEW QUESTION 5
A company uses a single AWS account lo test applications on Amazon EC2 instances. The company has turned on AWS Config in the AWS account and has activated the restricted- ssh AWS Config managed rule.
The company needs an automated monitoring solution that will provide a customized notification in real time if any security group in the account is not compliant with the restricted-ssh rule. The customized notification must contain the name and ID of the noncompliant security group.
A DevOps engineer creates an Amazon Simple Notification Service (Amazon SNS) topic in the account and subscribes the appropriate personnel to the topic.
What should me DevOps engineer do next to meet these requirements?

A. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT tor the restricted-ssh rul
B. Configure an input transformer for the EventBridge rule Configure the EventBridge rule to publish a notification to the SNS topic.
C. Configure AWS Config to send all evaluation results for the restricted-ssh rule to the SNS topi
D. Configure a filter policy on the SNS topic to send only notifications that contain the text of NON_COMPLIANT in the notification to subscribers.
E. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule Configure the EventBridge rule to invoke AWS Systems Manager Run Command on the SNS topic to customize a notification and to publish the notification to the SNS topic
F. Create an Amazon EventBridge rule that matches all AWS Config evaluation results of NON_COMPLIANT Configure an input transformer for the restricted-ssh rule Configure the EventBridge rule to publish a notification to the SNS topic.

Answer: A

Explanation:
Create an Amazon EventBridge (Amazon CloudWatch Events) rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule. Configure an input transformer for the EventBridge (CloudWatch Events) rule. Configure the EventBridge (CloudWatch Events) rule to publish a notification to the SNS topic. This approach uses Amazon EventBridge (previously known as Amazon CloudWatch Events) to filter AWS Config evaluation results based on the restricted-ssh rule and its compliance status (NON_COMPLIANT). An input transformer can be used to customize the information contained in the notification, such as the name and ID of the noncompliant security group. The EventBridge (CloudWatch Events) rule can then be configured to publish a notification to the SNS topic, which will notify the appropriate personnel in real-time.

NEW QUESTION 6
A company is examining its disaster recovery capability and wants the ability to switch over its daily operations to a secondary AWS Region. The company uses AWS CodeCommit as a source control tool in the primary Region.
A DevOps engineer must provide the capability for the company to develop code in the secondary Region. If the company needs to use the secondary Region,

developers can add an additional remote URL to their local Git configuration.
Which solution will meet these requirements?

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

Answer: A

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

NEW QUESTION 7
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 8

A company is using an organization in AWS Organizations to manage multiple AWS accounts. The company's development team wants to use AWS Lambda functions to meet resiliency requirements and is rewriting all applications to work with Lambda functions that are deployed in a VPC. The development team is using Amazon Elastic Pile System (Amazon EFS) as shared storage in Account A in the organization.
The company wants to continue to use Amazon EPS with Lambda Company policy requires all serverless projects to be deployed in Account B.
A DevOps engineer needs to reconfigure an existing EFS file system to allow Lambda functions to access the data through an existing EPS access point.
Which combination of steps should the DevOps engineer take to meet these requirements? (Select THREE.)

A. Update the EFS file system policy to provide Account B with access to mount and write to the EFS file system in Account A.
B. Create SCPs to set permission guardrails with fine-grained control for Amazon EFS.
C. Create a new EFS file system in Account B Use AWS Database Migration Service (AWS DMS) to keep data from Account A and Account B synchronized.
D. Update the Lambda execution roles with permission to access the VPC and the EFS file system.
E. Create a VPC peering connection to connect Account A to Account B.
F. Configure the Lambda functions in Account B to assume an existing IAM role in Account A.

Answer: AEF

Explanation:
A Lambda function in one account can mount a file system in a different account. For this scenario, you configure VPC peering between the function VPC and the file system VPC. https://docs.aws.amazon.com/lambda/latest/dg/services-efs.html https://aws.amazon.com/ru/blogs/storage/mount-amazon-efs-file-systems-cross-account- from-amazon-eks/
* 1. Need to update the file system policy on EFS to allow mounting the file system into Account B.
File System Policy
$ cat file-system-policy.json
{
"Statement": [
{
"Effect": "Allow", "Action": [
"elasticfilesystem:ClientMount", "elasticfilesystem:ClientWrite"
],
"Principal": {
"AWS": "arn:aws:iam::<aws-account-id-A>:root" # Replace with AWS account ID of EKS cluster
}
}
]
}
* 2. Need VPC peering between Account A and Account B as the pre-requisite
* 3. Need to assume cross-account IAM role to describe the mounts so that a specific mount can be chosen.

NEW QUESTION 9
A development team manually builds an artifact locally and then places it in an Amazon S3 bucket. The application has a local cache that must be cleared when a deployment occurs. The team runs a command to do this downloads the artifact from Amazon S3 and unzips the artifact to complete the deployment.
A DevOps team wants to migrate to a CI/CD process and build in checks to stop and roll back the deployment when a failure occurs. This requires the team to track the progression of the deployment.
Which combination of actions will accomplish this? (Select THREE)

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

Answer: BDE

NEW QUESTION 10
A company runs an application with an Amazon EC2 and on-premises configuration. A DevOps engineer needs to standardize patching across both environments. Company policy dictates that patching only happens during non-business hours.
Which combination of actions will meet these requirements? (Choose three.)

A. Add the physical machines into AWS Systems Manager using Systems Manager Hybrid Activations.
B. Attach an IAM role to the EC2 instances, allowing them to be managed by AWS Systems Manager.
C. Create IAM access keys for the on-premises machines to interact with AWS Systems Manager.
D. Run an AWS Systems Manager Automation document to patch the systems every hour.
E. Use Amazon EventBridge scheduled events to schedule a patch window.
F. Use AWS Systems Manager Maintenance Windows to schedule a patch window.

Answer: ABF

Explanation:
https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-managed-instance-activation.html

NEW QUESTION 10
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 15
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 16
A company requires that its internally facing web application be highly available. The architecture is made up of one Amazon EC2 web server instance and one NAT instance that provides outbound internet access for updates and accessing public data.
Which combination of architecture adjustments should the company implement to achieve high availability? (Choose two.)

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

Answer: BD

Explanation:

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

NEW QUESTION 17
A company wants to ensure that their EC2 instances are secure. They want to be notified if any new vulnerabilities are discovered on their instances and they also want an audit trail of all login activities on the instances.
Which solution will meet these requirements'?

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

Answer: D

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

NEW QUESTION 18
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 19
A DevOps engineer used an AWS Cloud Formation custom resource to set up AD Connector. The AWS Lambda function ran and created AD Connector, but Cloud Formation is not transitioning from CREATE_IN_PROGRESS to CREATE_COMPLETE.
Which action should the engineer take to resolve this issue?

A. Ensure the Lambda function code has exited successfully.
B. Ensure the Lambda function code returns a response to the pre-signed URL.
C. Ensure the Lambda function IAM role has cloudformation UpdateStack permissions for the stack ARN.
D. Ensure the Lambda function IAM role has ds ConnectDirectory permissions for the AWS account.

Answer: B

Explanation:

Reference: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/crpg-ref- responses.html

NEW QUESTION 24
A company uses a series of individual Amazon Cloud Formation templates to deploy its multi-Region Applications. These templates must be deployed in a specific order. The company is making more changes to the templates than previously expected and wants to deploy new templates more efficiently. Additionally, the data engineering team must be notified of all changes to the templates.

What should the company do to accomplish these goals?

A. Create an AWS Lambda function to deploy the Cloud Formation templates m the required order Use stack policies to alert the data engineering team.
B. Host the Cloud Formation templates in Amazon S3 Use Amazon S3 events to directly trigger CloudFormation updates and Amazon SNS notifications.
C. Implement CloudFormation StackSets and use drift detection to trigger update alerts to the data engineering team.
D. Leverage CloudFormation nested stacks and stack sets (or deployments Use Amazon SNS to notify the data engineering team.

Answer: D

Explanation:
This solution will meet the requirements because it will use CloudFormation nested stacks and stack sets to deploy the templates more efficiently and consistently across multiple regions. Nested stacks allow the company to separate out common components and reuse templates, while stack sets allow the company to create stacks in multiple accounts and regions with a single template. The company can also use Amazon SNS to send notifications to the data engineering team whenever a change is made to the templates or the stacks. Amazon SNS is a service that allows you to publish messages to subscribers, such as email addresses, phone numbers, or other AWS services. By using Amazon SNS, the company can ensure that the data engineering team is aware of all changes to the templates and can take appropriate actions if needed. What is Amazon SNS? - Amazon Simple Notification Service

NEW QUESTION 27
A company has a legacy application A DevOps engineer needs to automate the process of building the deployable artifact for the legacy application. The solution must store the deployable artifact in an existing Amazon S3 bucket for future deployments to reference
Which solution will meet these requirements in the MOST operationally efficient way?

A. Create a custom Docker image that contains all the dependencies tor the legacy application Store the custom Docker image in a new Amazon Elastic Container Registry (Amazon ECR) repository Configure a new AWS CodeBuild project to use the custom Docker image to build the deployable artifact and to save the artifact to the S3 bucket.
B. Launch a new Amazon EC2 instance Install all the dependencies (or the legacy application on the EC2 instance Use the EC2 instance to build the deployable artifact and to save the artifact to the S3 bucket.
C. Create a custom EC2 Image Builder image Install all the dependencies for the legacy application on the image Launch a new Amazon EC2 instance from the image Use the new EC2 instance to build the deployable artifact and to save the artifact to the S3 bucket.
D. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster with an AWS Fargate profile that runs in multiple Availability Zones Create a custom Docker image that contains all the dependencies for the legacy application Store the custom Docker image in a new Amazon Elastic Container Registry (Amazon ECR) repository Use the custom Docker image inside the EKS cluster to build the deployable artifact and to save the artifact to the S3 bucket.

Answer: A

Explanation:
This approach is the most operationally efficient because it leverages the benefits of containerization, such as isolation and reproducibility, as well as AWS managed services. AWS CodeBuild is a fully managed build service that can compile your source code, run tests, and produce deployable software packages. By using a custom Docker image that includes all dependencies, you can ensure that the environment in which your code is built is consistent. Using Amazon ECR to store Docker images lets you easily deploy the images to any environment. Also, you can directly upload the build artifacts to Amazon S3 from AWS CodeBuild, which is beneficial for version control and archival purposes.

NEW QUESTION 31
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 36
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 39
A company has configured an Amazon S3 event source on an AWS Lambda function The company needs the Lambda function to run when a new object is created or an existing object IS modified In a particular S3 bucket The Lambda function will use the S3 bucket name and the S3 object key of the incoming event to read the contents of the created or modified S3 object The Lambda function will parse the contents and save the parsed contents to an Amazon DynamoDB table. The Lambda function's execution role has permissions to read from the S3 bucket and to write to the DynamoDB table, During testing, a DevOps engineer discovers that the Lambda
function does not run when objects are added to the S3 bucket or when existing objects are modified.
Which solution will resolve this problem?

A. Increase the memory of the Lambda function to give the function the ability to process large files from the S3 bucket.
B. Create a resource policy on the Lambda function to grant Amazon S3 the permission to invoke the Lambda function for the S3 bucket
C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an OnFailure destination for the Lambda function
D. Provision space in the /tmp folder of the Lambda function to give the function the ability to process large files from the S3 bucket

Answer: B

Explanation:
? Option A is incorrect because increasing the memory of the Lambda function does not address the root cause of the problem, which is that the Lambda function is not triggered by the S3 event source. Increasing the memory of the Lambda function might improve its performance or reduce its execution time, but it does not affect its invocation. Moreover, increasing the memory of the Lambda function might incur higher costs, as Lambda charges based on the amount of memory allocated to the function.
? Option B is correct because creating a resource policy on the Lambda function to grant Amazon S3 the permission to invoke the Lambda function for the S3 bucket is a necessary step to configure an S3 event source. A resource policy is a JSON document that defines who can access a Lambda resource and under what conditions. By granting Amazon S3 permission to invoke the Lambda function, the company ensures that the Lambda function runs when a new object is created or an existing object is modified in the S3 bucket1.
? Option C is incorrect because configuring an Amazon Simple Queue Service (Amazon SQS) queue as an On-Failure destination for the Lambda function does not help with triggering the Lambda function. An On-Failure destination is a feature that allows Lambda to send events to another service, such as SQS or Amazon Simple Notification Service (Amazon SNS), when a function invocation fails. However, this feature only applies to asynchronous invocations, and S3 event sources use synchronous invocations. Therefore, configuring an SQS queue as an On-Failure destination would have no effect on the problem.
? Option D is incorrect because provisioning space in the /tmp folder of the Lambda function does not address the root cause of the problem, which is that the Lambda function is not triggered by the S3 event source. Provisioning space in the /tmp folder of the Lambda function might help with processing large files from the S3 bucket, as it provides temporary storage for up to 512 MB of data. However, it does not affect the invocation of the Lambda function.
References:
? Using AWS Lambda with Amazon S3
? Lambda resource access permissions
? AWS Lambda destinations
? [AWS Lambda file system]

NEW QUESTION 44
A company has multiple development teams in different business units that work in a shared single AWS account All Amazon EC2 resources that are created in the account must include tags that specify who created the resources. The tagging must occur within the first hour of resource creation.
A DevOps engineer needs to add tags to the created resources that Include the user ID that created the resource and the cost center ID The DevOps engineer configures an AWS Lambda function With the cost center mappings to tag the resources. The DevOps engineer also sets up AWS CloudTrail in the AWS account. An Amazon S3 bucket stores the CloudTrail event logs
Which solution will meet the tagging requirements?

A. Create an S3 event notification on the S3 bucket to invoke the Lambda function for s3.ObJectTagging:Put event
B. Enable bucket versioning on the S3 bucket.
C. Enable server access logging on the S3 bucke
D. Create an S3 event notification on the S3 bucket for s3. ObjectTaggIng.• events

E. Create a recurring hourly Amazon EventBridge scheduled rule that invokes the Larnbda functio
F. Modify the Lambda function to read the logs from the S3 bucket
G. Create an Amazon EventBridge rule that uses Amazon EC2 as the event sourc
H. Configure the rule to match events delivered by CloudTrai
I. Configure the rule to target the Lambda function

Answer: D

Explanation:
? Option A is incorrect because S3 event notifications do not support s3.ObjectTagging:Put events. S3 event notifications only support events related to object creation, deletion, replication, and restore. Moreover, enabling bucket versioning on the S3 bucket is not relevant to the tagging requirements, as it only keeps multiple versions of objects in the bucket.
? Option B is incorrect because enabling server access logging on the S3 bucket does not help with tagging the resources. Server access logging only records requests for access to the bucket or its objects. It does not capture the user ID or the cost center ID of the resources. Furthermore, creating an S3 event notification on the S3 bucket for s3.ObjectTagging:Put events is not possible, as explained in option A.
? Option C is incorrect because creating a recurring hourly Amazon EventBridge scheduled rule that invokes the Lambda function is not efficient or timely. The Lambda function would have to read the logs from the S3 bucket every hour and tag the resources accordingly, which could incur unnecessary costs and delays. A better solution would be to trigger the Lambda function as soon as a resource is created, rather than waiting for an hourly schedule.
? Option D is correct because creating an Amazon EventBridge rule that uses Amazon EC2 as the event source and matches events delivered by CloudTrail is a valid way to tag the resources. CloudTrail records all API calls made to AWS services, including EC2, and delivers them as events to EventBridge. The EventBridge rule can filter the events based on the user ID and the resource type, and then target the Lambda function to tag the resources with the cost center ID. This solution meets the tagging requirements in a timely and efficient manner.
References:
? S3 event notifications
? Server access logging
? Amazon EventBridge rules
? AWS CloudTrail

NEW QUESTION 46
A company has developed a serverless web application that is hosted on AWS. The application consists of Amazon S3. Amazon API Gateway, several AWS Lambda functions, and an Amazon RDS for MySQL database. The company is using AWS CodeCommit to store the source code. The source code is a combination of AWS Serverless Application Model (AWS SAM) templates and Python code.
A security audit and penetration test reveal that user names and passwords for authentication to the database are hardcoded within CodeCommit repositories. A DevOps engineer must implement a solution to automatically detect and prevent hardcoded secrets.
What is the MOST secure solution that meets these requirements?

A. Enable Amazon CodeGuru Profile
B. Decorate the handler function with@with_lambda_profiler(). Manually review the recommendation repor
C. Write the secret to AWS Systems Manager Parameter Store as a secure strin
D. Update the SAM templates and the Python code to pull the secret from Parameter Store.
E. Associate the CodeCommit repository with Amazon CodeGuru Reviewe
F. Manually check the code review for any recommendation
G. Choose the option to protect the secre
H. Update the SAM templates and the Python code to pull the secret from AWS Secrets Manager.
I. Enable Amazon CodeGuru Profile
J. Decorate the handler function with@with_lambda_profiler(). Manually review the recommendation repor
K. Choose the option to protect the secre
L. Update the SAM templates and the Python code to pull the secret from AWS Secrets Manager.
M. Associate the CodeCommit repository with Amazon CodeGuru Reviewe
N. Manually check the code review for any recommendation
O. Write the secret to AWS Systems Manager Parameter Store as a strin
P. Update the SAM templates and the Python code to pull the secret from Parameter Store.

Answer: B

Explanation:
https://docs.aws.amazon.com/codecommit/latest/userguide/how-to-amazon-codeguru-reviewer.html

NEW QUESTION 50
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 53
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 55
A company that uses electronic health records is running a fleet of Amazon EC2 instances with an Amazon Linux operating system. As part of patient privacy requirements, the company must ensure continuous compliance for patches for operating system and applications running on the EC2 instances.
How can the deployments of the operating system and application patches be automated using a default and custom repository?

A. Use AWS Systems Manager to create a new patch baseline including the custom repositor
B. Run the AWS-RunPatchBaseline document using the run command to verify and install patches.
C. Use AWS Direct Connect to integrate the corporate repository and deploy the patchesusing Amazon CloudWatch scheduled events, then use the CloudWatch dashboard to create reports.
D. Use yum-config-manager to add the custom repository under /etc/yum.repos.d and run yum-config-manager-enable to activate the repository.
E. Use AWS Systems Manager to create a new patch baseline including the corporate repositor
F. Run the AWS-AmazonLinuxDefaultPatchBaseline document using the run command to verify and install patches.

Answer: A

Explanation:
https://docs.aws.amazon.com/systems-manager/latest/userguide/patch-manager-how-it-works-alt-source-repository.html

NEW QUESTION 58
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 59
A company is using AWS to run digital workloads. Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations.
The company wants to enforce security standards across the entire organization. To avoid noncompliance because of security misconfiguration, the company has enforced the use of AWS CloudFormation. A production support team can modify resources in the production environment by using the AWS Management Console to troubleshoot and resolve application-related issues.
A DevOps engineer must implement a solution to identify in near real time any AWS
service misconfiguration that results in noncompliance. The solution must automatically remediate the issue within 15 minutes of identification. The solution also must track noncompliant resources and events in a centralized dashboard with accurate timestamps.
Which solution will meet these requirements with the LEAST development overhead?

A. Use CloudFormation drift detection to identify noncompliant resource
B. Use drift detection events from CloudFormation to invoke an AWS Lambda function for remediatio
C. Configure the Lambda function to publish logs to an Amazon CloudWatch Logs log grou
D. Configure an Amazon CloudWatch dashboard to use the log group for tracking.
E. Turn on AWS CloudTrail in the AWS account
F. Analyze CloudTrail logs by using Amazon Athena to identify noncompliant resource
G. Use AWS Step Functions to track query results on Athena for drift detection and to invoke an AWS Lambda function for remediatio
H. For tracking, set up an Amazon QuickSight dashboard that uses Athena as the data source.
I. Turn on the configuration recorder in AWS Config in all the AWS accounts to identify noncompliant resource
J. Enable AWS Security Hub with the ~no-enable-default-standards option in all the AWS account
K. Set up AWS Config managed rules and custom rule
L. Set up automatic remediation by using AWS Config conformance pack
M. For tracking, set up a dashboard on Security Hub in a designated Security Hub administrator account.
N. Turn on AWS CloudTrail in the AWS account
O. Analyze CloudTrail logs by using Amazon CloudWatch Logs to identify noncompliant resource
P. Use CloudWatch Logs filters for drift detectio
Q. Use Amazon EventBridge to invoke the Lambda function for remediatio
R. Stream filtered CloudWatch logs to Amazon OpenSearch Servic
S. Set up a dashboard on OpenSearch Service for tracking.

Answer: C

Explanation:
The best solution is to use AWS Config and AWS Security Hub to identify and remediate noncompliant resources across multiple AWS accounts. AWS Config enables continuous monitoring of the configuration of AWS resources and evaluates them against desired configurations. AWS Config can also automatically remediate noncompliant resources by using conformance packs, which are a collection of AWS Config rules and remediation actions that can be deployed as a single entity. AWS Security Hub provides a comprehensive view of the security posture of AWS accounts and resources. AWS Security Hub can aggregate and normalize the findings from AWS Config and other AWS services, as well as from partner solutions. AWS Security Hub can also be used to create a dashboard for tracking noncompliant resources and events in a centralized location.
The other options are not optimal because they either require more development overhead, do not provide near real time detection and remediation, or do not provide a centralized dashboard for tracking.
Option A is not optimal because CloudFormation drift detection is not a near real time solution. Drift detection has to be manually initiated on each stack or resource, or scheduled using a cron expression. Drift detection also does not provide remediation
actions, so a custom Lambda function has to be developed and invoked. CloudWatch Logs and dashboard can be used for tracking, but they do not provide a comprehensive view of the security posture of the AWS accounts and resources.
Option B is not optimal because CloudTrail logs analysis using Athena is not a near real time solution. Athena queries have to be manually run or scheduled using a cron expression. Athena also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. Step Functions can be used to orchestrate the query and remediation workflow, but it adds more complexity and cost. QuickSight dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources.
Option D is not optimal because CloudTrail logs analysis using CloudWatch Logs is not a near real time solution. CloudWatch Logs filters have to be manually created or updated for each resource type and configuration change. CloudWatch Logs also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. EventBridge can be used to trigger the Lambda function, but it adds more complexity and cost. OpenSearch Service dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources. References:
? AWS Config conformance packs
? Introducing AWS Config conformance packs
? Managing conformance packs across all accounts in your organization

NEW QUESTION 62
A company uses AWS Directory Service for Microsoft Active Directory as its identity provider (IdP). The company requires all infrastructure to be defined and deployed by AWS CloudFormation.
A DevOps engineer needs to create a fleet of Windows-based Amazon EC2 instances to host an application. The DevOps engineer has created a CloudFormation template that contains an EC2 launch template, IAM role, EC2 security group, and EC2 Auto Scaling group. The DevOps engineer must implement a solution that joins all EC2 instances to the domain of the AWS Managed Microsoft AD directory.
Which solution will meet these requirements with the MOST operational efficiency?

A. In the CloudFormation template, create an AWS::SSM::Document resource that joins the EC2 instance to the AWS Managed Microsoft AD domain by using the parameters for the existing director
B. Update the launch template to include the SSMAssociation property to use the new SSM documen
C. Attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use.
D. In the CloudFormation template, update the launch template to include specific tags that propagate on launc

E. Create an AWS::SSM::Association resource to associate the AWS- JoinDirectoryServiceDomain Automation runbook with the EC2 instances that have the specified tag
F. Define the required parameters to join the AWS Managed Microsoft AD director
G. Attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use.
H. Store the existing AWS Managed Microsoft AD domain connection details in AWS Secrets Manage
I. In the CloudFormation template, create an AWS::SSM::Association resource to associate the AWS-CreateManagedWindowsInstanceWithApproval Automation runbook with the EC2 Auto Scaling grou
J. Pass the ARNs for the parameters from Secrets Manager to join the domai
K. Attach the AmazonSSMDirectoryServiceAccess and SecretsManagerReadWrite AWS managed policies to the IAM role that the EC2 instances use.
L. Store the existing AWS Managed Microsoft AD domain administrator credentials in AWS Secrets Manage
M. In the CloudFormation template, update the EC2 launch template to include user dat
N. Configure the user data to pull the administrator credentials from Secrets Manager and to join the AWS Managed Microsoft AD domai
O. Attach the AmazonSSMManagedInstanceCore and SecretsManagerReadWrite AWS managed policies to the IAM role that the EC2 instances use.

Answer: B

Explanation:
To meet the requirements, the DevOps engineer needs to create a solution that joins all EC2 instances to the domain of the AWS Managed Microsoft AD directory with the most operational efficiency. The DevOps engineer can use AWS Systems Manager Automation to automate the domain join process using an existing runbook called AWS- JoinDirectoryServiceDomain. This runbook can join Windows instances to an AWS Managed Microsoft AD or Simple AD directory by using PowerShell commands. The DevOps engineer can create an AWS::SSM::Association resource in the CloudFormation template to associate the runbook with the EC2 instances that have specific tags. The tags can be defined in the launch template and propagated on launch to the EC2 instances. The DevOps engineer can also define the required parameters for the runbook, such as the directory ID, directory name, and organizational unit. The DevOps engineer can attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use. These policies grant the necessary permissions for Systems Manager and Directory Service operations.

NEW QUESTION 64
A company has developed an AWS Lambda function that handles orders received through an API. The company is using AWS CodeDeploy to deploy the Lambda function as the final stage of a CI/CD pipeline.
A DevOps engineer has noticed there are intermittent failures of the ordering API for a few seconds after deployment. After some investigation the DevOps engineer believes the failures are due to database changes not having fully propagated before the Lambda function is invoked
How should the DevOps engineer overcome this?

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

Answer: A

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

NEW QUESTION 69
A company's application uses a fleet of Amazon EC2 On-Demand Instances to analyze and process data. The EC2 instances are in an Auto Scaling group. The Auto Scaling group is a target group for an Application Load Balancer (ALB). The application analyzes critical data that cannot tolerate interruption. The application also analyzes noncritical data that can withstand interruption.
The critical data analysis requires quick scalability in response to real-time application demand. The noncritical data analysis involves memory consumption. A DevOps engineer must implement a solution that reduces scale-out latency for the critical data. The solution also must process the noncritical data.
Which combination of steps will meet these requirements? (Select TWO.)

A. For the critical data, modify the existing Auto Scaling grou
B. Create a warm pool instance in the stopped stat
C. Define the warm pool siz
D. Create a new version of the launch template that has detailed monitoring enable
E. use Spot Instances.
F. For the critical data, modify the existing Auto Scaling grou
G. Create a warm pool instance in the stopped stat
H. Define the warm pool siz
I. Create a new version of the launch template that has detailed monitoring enable
J. Use On-Demand Instances.
K. For the critical dat
L. modify the existing Auto Scaling grou
M. Create a lifecycle hook to ensure that bootstrap scripts are completed successfull
N. Ensure that the application on the instances is ready to accept traffic before the instances are registere
O. Create a new version of the launch template that has detailed monitoring enabled.
P. For the noncritical data, create a second Auto Scaling group that uses a launch templat
Q. Configure the launch template to install the unified Amazon CloudWatch agent and to configure the CloudWatch agent with a custom memory utilization metri
R. Use Spot Instance
S. Add the new Auto Scaling group as the target group for the AL
T. Modify the application to use two target groups for critical data and noncritical data.
. For the noncritical data, create a second Auto Scaling grou
. Choose the predefined memory utilization metric type for the target tracking scaling polic
. Use Spot Instance
. Add the new Auto Scaling group as the target group for the AL
. Modify the application to use two target groups for critical data and noncritical data.

Answer: BD

Explanation:
? For the critical data, using a warm pool1 can reduce the scale-out latency by having pre-initialized EC2 instances ready to serve the application traffic. Using On-Demand Instances can ensure that the instances are always available and not interrupted by Spot interruptions2.
? For the noncritical data, using a second Auto Scaling group with Spot Instances can reduce the cost and leverage the unused capacity of EC23. Using a launch template with the CloudWatch agent4 can enable the collection of memory utilization metrics, which can be used to scale the group based on the memory demand. Adding the second group as a target group for the ALB and modifying the application to use two target groups can enable routing the traffic based on the data type.
References: 1: Warm pools for Amazon EC2 Auto Scaling 2: Amazon EC2 On-Demand Capacity Reservations 3: Amazon EC2 Spot Instances 4: Metrics collected by the CloudWatch agent

NEW QUESTION 74
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 78
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 83
A global company manages multiple AWS accounts by using AWS Control Tower. The company hosts internal applications and public applications.
Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations.
One of the AWS Control Tower member accounts serves as a centralized DevOps account with CI/CD pipelines that application teams use to deploy applications to their respective target AWS accounts. An 1AM role for deployment exists in the centralized DevOps account.
An application team is attempting to deploy its application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster in an application AWS account. An 1AM role for deployment exists in the application AWS account. The deployment is through an AWS CodeBuild project that is set up in the centralized DevOps account. The CodeBuild project uses an 1AM service role for CodeBuild. The deployment is failing with an Unauthorized error during attempts to connect to the cross-account EKS cluster from CodeBuild.
Which solution will resolve this error?

A. Configure the application account's deployment 1AM role to have a trust relationship with the centralized DevOps accoun
B. Configure the trust relationship to allow the sts:AssumeRole actio
C. Configure the application account's deployment 1AM role to have the required access to the EKS cluste
D. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.
E. Configure the centralized DevOps account's deployment I AM role to have a trust relationship with the application accoun
F. Configure the trust relationship to allow the sts:AssumeRole actio
G. Configure the centralized DevOps account's deployment 1AM role to allow the required access to CodeBuild.
H. Configure the centralized DevOps account's deployment 1AM role to have a trust relationship with the application accoun
I. Configure the trust relationship to allow the sts:AssumeRoleWithSAML actio
J. Configure the centralized DevOps account's deployment 1AM role to allow the required access to CodeBuild.
K. Configure the application account's deployment 1AM role to have a trust relationship with the AWS Control Tower management accoun
L. Configure the trust relationship to allow the sts:AssumeRole actio
M. Configure the application account's deployment 1AM role to have the required access to the EKS cluste

N. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.

Answer: A

Explanation:
In the source AWS account, the IAM role used by the CI/CD pipeline should have permissions to access the source code repository, build artifacts, and any other resources required for the build process. In the destination AWS accounts, the IAM role used for deployment should have permissions to access the AWS resources required for deploying the application, such as EC2 instances, RDS databases, S3 buckets, etc. The exact permissions required will depend on the specific resources being used by the application. the IAM role used for deployment in the destination accounts should also have permissions to assume the IAM role for deployment in the centralized DevOps account. This is typically done using an IAM role trust policy that allows the destination account to assume the DevOps account role.

NEW QUESTION 86
A company must encrypt all AMIs that the company shares across accounts. A DevOps engineer has access to a source account where an unencrypted custom AMI has been built. The DevOps engineer also has access to a target account where an Amazon EC2 Auto Scaling group will launch EC2 instances from the AMI. The DevOps engineer must share the AMI with the target account.
The company has created an AWS Key Management Service (AWS KMS) key in the source account.
Which additional steps should the DevOps engineer perform to meet the requirements? (Choose three.)

A. In the source account, copy the unencrypted AMI to an encrypted AM
B. Specify the KMS key in the copy action.
C. In the source account, copy the unencrypted AMI to an encrypted AM
D. Specify the default Amazon Elastic Block Store (Amazon EBS) encryption key in the copy action.
E. In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.
F. In the source account, modify the key policy to give the target account permissions to create a gran
G. In the target account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role.
H. In the source account, share the unencrypted AMI with the target account.
I. In the source account, share the encrypted AMI with the target account.

Answer: ADF

Explanation:
The Auto Scaling group service-linked role must have a specific grant in the source account in order to decrypt the encrypted AMI. This is because the service-linked role does not have permissions to assume the default IAM role in the source account. The following steps are required to meet the requirements:
? In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the KMS key in the copy action.
? In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.
? In the source account, share the encrypted AMI with the target account.
? In the target account, attach the KMS grant to the Auto Scaling group service- linked role.
The first three steps are the same as the steps that I described earlier. The fourth step is required to grant the Auto Scaling group service-linked role permissions to decrypt the AMI
in the target account.

NEW QUESTION 88
An ecommerce company has chosen AWS to host its new platform. The company's DevOps team has started building an AWS Control Tower landing zone. The DevOps team has set the identity store within AWS IAM Identity Center (AWS Single Sign-On) to external identity provider (IdP) and has configured SAML 2.0. The DevOps team wants a robust permission model that applies the principle of least privilege. The model must allow the team to build and manage only the team's own resources.
Which combination of steps will meet these requirements? (Choose three.)

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

Answer: BCF

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

NEW QUESTION 92
A DevOps engineer has implemented a CI/CO pipeline to deploy an AWS Cloud Format ion template that provisions a web application. The web application consists of an Application Load Balancer (ALB) a target group, a launch template that uses an Amazon Linux 2 AMI an Auto Scaling group of Amazon EC2 instances, a security group and an Amazon RDS for MySQL database The launch template includes user data that specifies a script to install and start the application.
The initial deployment of the application was successful. The DevOps engineer made changes to update the version of the application with the user data. The

CI/CD pipeline has deployed a new version of the template However, the health checks on the ALB are now failing The health checks have marked all targets as unhealthy.
During investigation the DevOps engineer notices that the Cloud Formation stack has a status of UPDATE_COMPLETE. However, when the DevOps engineer connects to one of the EC2 instances and checks /varar/log messages, the DevOps engineer notices that the Apache web server failed to start successfully because of a configuration error
How can the DevOps engineer ensure that the CloudFormation deployment will fail if the user data fails to successfully finish running?

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

Answer: A

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

NEW QUESTION 93
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 96

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 100
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 102
A rapidly growing company wants to scale for developer demand for AWS development environments. Development environments are created manually in the AWS Management Console. The networking team uses AWS CloudFormation to manage the networking infrastructure, exporting stack output values for the Amazon VPC and all subnets. The development environments have common standards, such as Application Load Balancers, Amazon EC2 Auto Scaling groups, security groups, and Amazon DynamoDB tables.
To keep up with demand, the DevOps engineer wants to automate the creation of development environments. Because the infrastructure required to support the application is expected to grow, there must be a way to easily update the deployed infrastructure. CloudFormation will be used to create a template for the development environments.
Which approach will meet these requirements and quickly provide consistent AWS environments for developers?

A. Use Fn::ImportValue intrinsic functions in the Resources section of the template to retrieve Virtual Private Cloud (VPC) and subnet value
B. Use CloudFormation StackSets for the development environments, using the Count input parameter to indicate the number of environments neede
C. Use the UpdateStackSet command to update existing development environments.
D. Use nested stacks to define common infrastructure component
E. To access the exported values, use TemplateURL to reference the networking team's templat
F. To retrieve Virtual Private Cloud (VPC) and subnet values, use Fn::ImportValue intrinsic functions in the Parameters section of the root templat
G. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.
H. Use nested stacks to define common infrastructure component
I. Use Fn::ImportValue intrinsic functions with the resources of the nested stack to retrieve Virtual Private Cloud (VPC) and subnet value
J. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.
K. Use Fn::ImportValue intrinsic functions in the Parameters section of the root template to retrieve Virtual Private Cloud (VPC) and subnet value
L. Define the development resources in the order they need to be created in the CloudFormation nested stack
M. Use the CreateChangeSe
N. and ExecuteChangeSet commands to update existing development environments.

Answer: C

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function- reference-importvalue.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-importvalue.html CF of network exports the VPC, subnet or needed information CF of application imports the above information to its stack and UpdateChangeSet/ ExecuteChangeSet

NEW QUESTION 105
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 110
A company is deploying a new application that uses Amazon EC2 instances. The company needs a solution to query application logs and AWS account API activity Which solution will meet these requirements?

A. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon CloudWatch Logs Configure AWS CloudTrail to deliver the API logs to Amazon S3 Use CloudWatch to query both sets of logs.
B. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon CloudWatch Logs Configure AWS CloudTrail to deliver the API logs to CloudWatch Logs Use CloudWatch Logs Insights to query both sets of logs.
C. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon Kinesis Configure AWS CloudTrail to deliver the API logs to Kinesis Use Kinesis to load the data into Amazon Redshift Use Amazon Redshift to query both sets of logs.
D. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon S3 Use AWS CloudTrail to deliver the API togs to Amazon S3 Use Amazon Athena to query both sets of logs in Amazon S3.

Answer: D

Explanation:
This solution will meet the requirements because it will use Amazon S3 as a common data lake for both the application logs and the API logs. Amazon S3 is a service that provides scalable, durable, and secure object storage for any type of data. You can use the Amazon CloudWatch agent to send logs from your EC2 instances to S3 buckets, and use AWS CloudTrail to deliver the API logs to S3 buckets as well. You can also use Amazon Athena to query both sets of logs in S3 using standard SQL, without loading or transforming them. Athena is a serverless interactive query service that allows you to analyze data in S3 using a variety of data formats, such as JSON, CSV, Parquet, and ORC.

NEW QUESTION 111
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 112
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 113
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 118
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 119
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 122
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 127
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 130
A company has many applications. Different teams in the company developed the applications by using multiple languages and frameworks. The applications run on premises and on different servers with different operating systems. Each team has its own release protocol and process. The company wants to reduce the complexity of the release and maintenance of these applications.
The company is migrating its technology stacks, including these applications, to AWS. The
company wants centralized control of source code, a consistent and automatic delivery pipeline, and as few maintenance tasks as possible on the underlying infrastructure.
What should a DevOps engineer do to meet these requirements?

A. Create one AWS CodeCommit repository for all application
B. Put each application's code in a different branc
C. Merge the branches, and use AWS CodeBuild to build the application
D. Use AWS CodeDeploy to deploy the applications to one centralized application server.
E. Create one AWS CodeCommit repository for each of the application
F. Use AWS CodeBuild to build the applications one at a tim
G. Use AWS CodeDeploy to deploy the applications to one centralized application server.
H. Create one AWS CodeCommit repository for each of the application
I. Use AWS CodeBuild to build the applications one at a time and to create one AMI for each serve
J. Use AWS CloudFormation StackSets to automatically provision and decommission Amazon EC2 fleets by using these AMIs.
K. Create one AWS CodeCommit repository for each of the application
L. Use AWS CodeBuild to build one Docker image for each application in Amazon Elastic Container Registry (Amazon ECR). Use AWS CodeDeploy to deploy the applications to Amazon Elastic Container Service (Amazon ECS) on infrastructure that AWS Fargate manages.

Answer: D

Explanation:
because of "as few maintenance tasks as possible on the underlying infrastructure". Fargate does that better than "one centralized application server"

NEW QUESTION 132
A company has multiple development groups working in a single shared AWS account. The Senior Manager of the groups wants to be alerted via a third-party API call when the creation of resources approaches the service limits for the account.
Which solution will accomplish this with the LEAST amount of development effort?

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

periodicall
E. Create another CloudWatch Events rule with an event pattern matching Trusted Advisor events and a target Lambda functio
F. In the target Lambda function, notify the Senior Manager.
G. Deploy an AWS Lambda function that refreshes AWS Personal Health Dashboard checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodicall
H. Create another CloudWatch Events rule with an event pattern matching Personal Health Dashboard events and a target Lambda functio
I. In the target Lambda function, notify the Senior Manager.
J. Add an AWS Config custom rule that runs periodically, checks the AWS service limit status, and streams notifications to an Amazon SNS topi
K. Deploy an AWS Lambda function that notifies the Senior Manager, and subscribe the Lambda function to the SNS topic.

Answer: B

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

NEW QUESTION 136
.....

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