

# Amazon

## Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional



### NEW QUESTION 1

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-out
- 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 consumers<sup>1</sup>. 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 stream<sup>2</sup>. 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 stream<sup>3</sup>. 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 latency<sup>4</sup>. 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 2

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 provider
- B. Configure pipeline stages that run the CodeBuild project in parallel to build and test the application
- 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 provider
- E. Create separate pipeline stages that run a CodeBuild project to build and then test the application
- 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 instances
- 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 bucket
- K. Modify the CodeBuild project to store the packages in the S3 bucket instead of in CodeArtifact
- 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 configuration<sup>123</sup>

? 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 traffic<sup>45</sup>

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 development<sup>67</sup>

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 3

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 Config 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 4

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 storage
- B. Create an S3 bucket in the primary Region and an S3 bucket in the DR Region
- 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 storage
- 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 storage
- 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 storage
- I. Create an FSx for ONTAP instance in the DR Region
- 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 5

A company is using an Amazon Aurora cluster as the data store for its application. The Aurora cluster is configured with a single DB instance. The application performs read and write operations on the database by using the cluster's instance endpoint.

The company has scheduled an update to be applied to the cluster during an upcoming maintenance window. The cluster must remain available with the least possible interruption during the maintenance window.

What should a DevOps engineer do to meet these requirements?

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

**Answer: C**

**Explanation:**

To meet the requirements, the DevOps engineer should do the following:

- ? Turn on the Multi-AZ option on the Aurora cluster.
- ? Update the application to use the Aurora cluster endpoint for write operations.
- ? Update the Aurora cluster's reader endpoint for reads.

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

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

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

**NEW QUESTION 6**

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 performance<sup>1</sup>.

? 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 performance<sup>1</sup>.

? 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 processing<sup>2</sup>.

? 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 processing<sup>2</sup>.

? 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 store<sup>3</sup>.

? 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 needs<sup>3</sup>.

References:

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

**NEW QUESTION 7**

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 instance

- 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 Zone
- 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 8**

A development team uses AWS CodeCommit, AWS CodePipeline, and AWS CodeBuild to develop and deploy an application. Changes to the code are submitted by pull requests. The development team reviews and merges the pull requests, and then the pipeline builds and tests the application.

Over time, the number of pull requests has increased. The pipeline is frequently blocked because of failing tests. To prevent this blockage, the development team wants to run the unit and integration tests on each pull request before it is merged.

Which solution will meet these requirements?

- A. Create a CodeBuild project to run the unit and integration test
- B. Create a CodeCommit approval rule template
- C. Configure the template to require the successful invocation of the CodeBuild project
- D. Attach the approval rule to the project's CodeCommit repository.
- E. Create an Amazon EventBridge rule to match pullRequestCreated events from CodeCommit Create a CodeBuild project to run the unit and integration test
- F. Configure the CodeBuild project as a target of the EventBridge rule that includes a custom event payload with the CodeCommit repository and branch information from the event.
- G. Create an Amazon EventBridge rule to match pullRequestCreated events from CodeCommit
- H. Modify the existing CodePipeline pipeline to not run the deploy steps if the build is started from a pull request
- I. Configure the EventBridge rule to run the pipeline with a custom payload that contains the CodeCommit repository and branch information from the event.
- J. Create a CodeBuild project to run the unit and integration test
- K. Create a CodeCommit notification rule that matches when a pull request is created or updated
- L. Configure the notification rule to invoke the CodeBuild project.

**Answer:** B

**Explanation:**

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

**NEW QUESTION 9**

A company's DevOps engineer is working in a multi-account environment. The company uses AWS Transit Gateway to route all outbound traffic through a network operations account. In the network operations account all account traffic passes through a firewall appliance for inspection before the traffic goes to an internet gateway.

The firewall appliance sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO. The security team wants to receive an alert if any CRITICAL events occur.

What should the DevOps engineer do to meet these requirements?

- A. Create an Amazon CloudWatch Synthetics canary to monitor the firewall status
- B. If the firewall reaches a CRITICAL state or logs a CRITICAL event use a CloudWatch alarm to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic Subscribe the security team's email address to the topic.
- C. Create an Amazon CloudWatch metric filter by using a search for CRITICAL events Publish a custom metric for the findings
- D. Use a CloudWatch alarm based on the custom metric to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic
- E. Subscribe the security team's email address to the topic.
- F. Enable Amazon GuardDuty in the network operations account
- G. Configure GuardDuty to monitor flow logs Create an Amazon EventBridge event rule that is invoked by GuardDuty events that are CRITICAL Define an Amazon Simple Notification Service (Amazon SNS) topic as a target Subscribe the security team's email address to the topic.
- H. Use AWS Firewall Manager to apply consistent policies across all accounts
- I. Create an Amazon EventBridge event rule that is invoked by Firewall Manager events that are CRITICAL Define an Amazon Simple Notification Service (Amazon SNS) topic as a target Subscribe the security team's email address to the topic.
- J. EventBridge event rule that is invoked by Firewall Manager events that are CRITICAL Define an Amazon Simple Notification Service (Amazon SNS) topic as a target Subscribe the security team's email address to the topic.

**Answer:** B

**Explanation:**

"The firewall appliance sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO"

**NEW QUESTION 10**

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 application
- B. Configure S3 Same-Region Replication (SRR) from the raw data's S3 bucket to each application's S3 bucket
- C. Configure each application to consume data from its own S3 bucket.
- D. Create an Amazon Kinesis data stream
- E. Create an AWS Lambda function that is invoked by object creation events in the raw data's S3 bucket
- F. Program the Lambda function to redact data for each application
- G. Publish the data on the Kinesis data stream
- 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 destination
- J. Create an AWS Lambda function that is invoked by object creation events in the raw data's S3 bucket
- K. Program the Lambda function to redact data for each application
- L. Store the data in each application's S3 access point
- 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 destination
- O. For each application, create an S3 Object Lambda access point that uses the S3 access point
- P. Configure the AWS Lambda function for each S3 Object Lambda access point to redact data when objects are retrieved
- 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 Lambda<sup>1</sup>, 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 application<sup>2</sup>. This way, you can redact the data differently for each application without creating and storing multiple copies of the data or running proxies.

? The other solutions are less efficient or scalable because they require replicating the data to multiple buckets, streaming the data through Kinesis, or storing the data in S3 access points.

References: 1: Amazon S3 Features | Object Lambda | AWS 2: Transforming objects with S3 Object Lambda - Amazon Simple Storage Service

**NEW QUESTION 10**

A company is testing a web application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company uses a blue green deployment process with immutable instances when deploying new software. During testing users are being automatically logged out of the application at random times. Testers also report that when a new version of the application is deployed all users are logged out. The development team needs a solution to ensure users remain logged in across scaling events and application deployments. What is the MOST operationally efficient way to ensure users remain logged in?

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

**Answer: D**

**Explanation:**

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

**NEW QUESTION 13**

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

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

**Answer: BE**

**Explanation:**

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

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

In the event of a failure, promote the secondary Region to the primary for the application. Update the application to use the Aurora cluster endpoint in the secondary Region. This strategy can provide a low RPO and RTO for the data, as Aurora global database replicates data with minimal latency across Regions and allows fast and easy failover<sup>12</sup>. The company can use the Amazon Aurora cluster endpoint to connect to the current primary DB cluster without needing to change any application code<sup>1</sup>.

? Set up the application in two AWS Regions. Configure AWS Global Accelerator to point to Application Load Balancers (ALBs) in both Regions. Add both ALBs to a single endpoint group. Use health checks and Auto Scaling groups in each Region. This strategy can provide high availability and performance for the application, as AWS Global Accelerator uses the AWS global network to route traffic to the closest healthy endpoint<sup>3</sup>. The company can also use static IP addresses that are assigned by Global Accelerator as a fixed entry point for their application<sup>1</sup>. By using health checks and Auto Scaling groups, the company can ensure that their application can scale up or down based on demand and handle any instance failures<sup>4</sup>.

The other options are incorrect because:

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

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

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

different factors.

#### NEW QUESTION 16

A company uses AWS and has a VPC that contains critical compute infrastructure with predictable traffic patterns. The company has configured VPC flow logs that are published to a log group in Amazon CloudWatch Logs.

The company's DevOps team needs to configure a monitoring solution for the VPC flow logs to identify anomalies in network traffic to the VPC over time. If the monitoring solution detects an anomaly, the company needs the ability to initiate a response to the anomaly.

How should the DevOps team configure the monitoring solution to meet these requirements?

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

**Answer: D**

#### Explanation:

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

#### NEW QUESTION 21

A company builds a container image in an AWS CodeBuild project by running Docker commands. After the container image is built, the CodeBuild project uploads the container image to an Amazon S3 bucket. The CodeBuild project has an IAM service role that has permissions to access the S3 bucket.

A DevOps engineer needs to replace the S3 bucket with an Amazon Elastic Container Registry (Amazon ECR) repository to store the container images. The DevOps engineer creates an ECR private image repository in the same AWS Region of the CodeBuild project. The DevOps engineer adjusts the IAM service role with the permissions that are necessary to work with the new ECR repository. The DevOps engineer also places new repository information into the docker build command and the docker push command that are used in the buildspec.yml file.

When the CodeBuild project runs a build job, the job fails when the job tries to access the ECR repository.

Which solution will resolve the issue of failed access to the ECR repository?

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

**Answer: A**

#### Explanation:

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

#### NEW QUESTION 24

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 Beanstalk
- B. Deploy an Amazon RDS for MySQL DB instance as part of the Elastic Beanstalk configuration.
- C. Deploy the application on AWS Elastic Beanstalk
- 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 database<sup>123</sup>.
  - ? 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 automatically<sup>4</sup>.
  - ? 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 events<sup>56</sup>.
- References:
- ? AWS Elastic Beanstalk documentation on deploying Ruby applications<sup>1</sup>.
  - ? AWS documentation on application auto-scaling<sup>7</sup>.
  - ? AWS documentation on automated deployment strategies with automatic rollbacks and alerts<sup>456</sup>.

**NEW QUESTION 28**

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 valid
- 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 data
- 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 queue
- E. Configure AWS Lambda to poll the SQS queue every 10 minutes by using an Amazon EventBridge schedule
- 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 queue
- 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 queue
- I. Instruct the scientists to use the dead-letter queue to review the data that is not valid
- 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 satellites
- 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 queue
- M. Instruct the scientists to use the virtual queue to review the data that is not valid
- 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 30**

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 bucket
- 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 Lambda function
- 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 source
- H. Configure the rule to match events delivered by CloudTrail
- 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 34

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 report
- C. Write the secret to AWS Systems Manager Parameter Store as a secure string
- D. Update the SAM templates and the Python code to pull the secret from Parameter Store.
- E. Associate the CodeCommit repository with Amazon CodeGuru Reviewer
- F. Manually check the code review for any recommendation
- G. Choose the option to protect the secret
- 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 report
- K. Choose the option to protect the secret
- 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 Reviewer
- N. Manually check the code review for any recommendation
- O. Write the secret to AWS Systems Manager Parameter Store as a string
- 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 38

A company's production environment uses an AWS CodeDeploy blue/green deployment to deploy an application. The deployment includes Amazon EC2 Auto Scaling groups that launch instances that run Amazon Linux 2.

A working `appspec.yml` 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. `DownloadBundle`

**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 43

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 repository
- 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 patches using 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 repository
- 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 45**

A DevOps engineer is setting up a container-based architecture. The engineer has decided to use AWS CloudFormation to automatically provision an Amazon ECS cluster and an Amazon EC2 Auto Scaling group to launch the EC2 container instances. After successfully creating the CloudFormation stack, the engineer noticed that, even though the ECS cluster and the EC2 instances were created successfully and the stack finished the creation, the EC2 instances were associating with a different cluster.

How should the DevOps engineer update the CloudFormation template to resolve this issue?

- A. Reference the EC2 instances in the AWS: ECS: Cluster resource and reference the ECS cluster in the AWS: ECS: Service resource.
- B. Reference the ECS cluster in the AWS: AutoScaling: LaunchConfiguration resource of the UserData property.
- C. Reference the ECS cluster in the AWS:EC2: Instance resource of the UserData property.
- D. Reference the ECS cluster in the AWS: CloudFormation: CustomResource resource to trigger an AWS Lambda function that registers the EC2 instances with the appropriate ECS cluster.

**Answer:** B

**Explanation:**

The UserData property of the AWS: AutoScaling: LaunchConfiguration resource can be used to specify a script that runs when the EC2 instances are launched. This script can include the ECS cluster name as an environment variable for the ECS agent running on the EC2 instances. This way, the EC2 instances will register with the correct ECS cluster. Option A is incorrect because the AWS: ECS: Cluster resource does not have a property to reference the EC2 instances. Option C is incorrect because the EC2 instances are launched by the Auto Scaling group, not by the AWS: EC2: Instance resource. Option D is incorrect because using a custom resource and a Lambda function is unnecessary and overly complex for this scenario. References: AWS::AutoScaling::LaunchConfiguration, Amazon ECS Container Agent Configuration

**NEW QUESTION 48**

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

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

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 51**

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 55**

A company is implementing an Amazon Elastic Container Service (Amazon ECS) cluster to run its workload. The company architecture will run multiple ECS services on the cluster. The architecture includes an Application Load Balancer on the front end and uses multiple target groups to route traffic. A DevOps engineer must collect application and access logs. The DevOps engineer then needs to send the logs to an Amazon S3 bucket for near-real-time analysis.

Which combination of steps must the DevOps engineer take to meet these requirements? (Choose three.)

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

**Answer:** BDF

**Explanation:**

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

**NEW QUESTION 60**

A company manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The EC2 instances run in an Auto Scaling group across multiple Availability Zones. The application uses an Amazon RDS for MySQL DB instance to store the data. The company has configured Amazon Route 53 with an alias record that points to the ALB.

A new company guideline requires a geographically isolated disaster recovery (DR) site with an RTO of 4 hours and an RPO of 15 minutes.

Which DR strategy will meet these requirements with the LEAST change to the application stack?

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

**Answer:** D

**NEW QUESTION 64**

A company builds an application that uses an Application Load Balancer in front of Amazon EC2 instances that are in an Auto Scaling group. The application is stateless. The Auto Scaling group uses a custom AMI that is fully prebuilt. The EC2 instances do not have a custom bootstrapping process.

The AMI that the Auto Scaling group uses was recently deleted. The Auto Scaling group's scaling activities show failures because the AMI ID does not exist.

Which combination of steps should a DevOps engineer take to meet these requirements? (Select THREE.)

- A. Create a new launch template that uses the new AMI.
- B. Update the Auto Scaling group to use the new launch template.
- C. Reduce the Auto Scaling group's desired capacity to 0.
- D. Increase the Auto Scaling group's desired capacity by 1.
- E. Create a new AMI from a running EC2 instance in the Auto Scaling group.
- F. Create a new AMI by copying the most recent public AMI of the operating system that the EC2 instances use.

**Answer:** ABF

**Explanation:**

To restore the functionality of the Auto Scaling group after the AMI was deleted, the DevOps engineer needs to create a new AMI and update the Auto Scaling group to use it. The DevOps engineer can create a new AMI by copying the most recent public AMI of the operating system that the EC2 instances use. This will ensure that the new AMI has the same operating system as the custom AMI that was deleted. The DevOps engineer can then create a new launch template that

uses the new AMI and update the Auto Scaling group to use the new launch template. This will allow the Auto Scaling group to launch new instances with the new AMI.

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

**Answer:** BD

#### Explanation:

? For the critical data, using a warm pool<sup>1</sup> 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 interruptions<sup>2</sup>.

? For the noncritical data, using a second Auto Scaling group with Spot Instances can reduce the cost and leverage the unused capacity of EC2<sup>3</sup>. Using a launch template with the CloudWatch agent<sup>4</sup> 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 70

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

Which solution ensures that all the updated third-party files are available in the morning?

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

**Answer:** A

#### Explanation:

[https://docs.aws.amazon.com/storagegateway/latest/APIReference/API\\_RefreshCache.html](https://docs.aws.amazon.com/storagegateway/latest/APIReference/API_RefreshCache.html) " It only updates the cached inventory to reflect changes in the inventory of the objects in the S3 bucket. This operation is only supported in the S3 File Gateway types."

#### NEW QUESTION 71

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 gain 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 services
- 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 IAM 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 services
- F. Attach the SCP to the root OU of the organization

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 72**

A business has an application that consists of five independent AWS Lambda functions.

The DevOps engineer has built a CI/CD pipeline using AWS CodePipeline and AWS CodeBuild that builds tests packages and deploys each Lambda function in sequence. The pipeline uses an Amazon EventBridge rule to ensure the pipeline starts as quickly as possible after a change is made to the application source code.

After working with the pipeline for a few months the DevOps engineer has noticed the pipeline takes too long to complete.

What should the DevOps engineer implement to BEST improve the speed of the pipeline?

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

**Answer: C**

**Explanation:**

<https://docs.aws.amazon.com/codepipeline/latest/userguide/reference-pipeline-structure.html>

AWS doc: "To specify parallel actions, use the same integer for each action you want to run in parallel. For example, if you want three actions to run in sequence in a stage, you would give the first action the runOrder value of 1, the second action the runOrder value of 2, and the third the runOrder value of 3. However, if you want the second and third actions to run in parallel, you would give the first action the runOrder value of 1 and both the second and third actions the runOrder value of 2."

**NEW QUESTION 76**

A company has an application that includes AWS Lambda functions. The Lambda functions run Python code that is stored in an AWS CodeCommit repository. The company has recently experienced failures in the production environment because of an error in the Python code. An engineer has written unit tests for the Lambda functions to help avoid releasing any future defects into the production environment.

The company's DevOps team needs to implement a solution to integrate the unit tests into an existing AWS CodePipeline pipeline. The solution must produce reports about the unit tests for the company to view.

Which solution will meet these requirements?

- A. Associate the CodeCommit repository with Amazon CodeGuru Reviewer
- B. Create a new AWS CodeBuild project
- C. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project
- D. Create a buildspec.yml file in the CodeCommit repository
- E. In the buildspec.yml file, define the actions to run a CodeGuru review.
- F. Create a new AWS CodeBuild project
- G. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project
- H. Create a CodeBuild report group
- I. Create a buildspec.yml file in the CodeCommit repository
- J. In the buildspec.yml file, define the actions to run the unit tests with an output of JUNITXML in the build phase section. Configure the test reports to be uploaded to the new CodeBuild report group.
- K. Create a new AWS CodeArtifact repository
- L. Create a new AWS CodeBuild project
- M. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project
- N. Create an appspec.yml file in the original CodeCommit repository
- O. In the appspec.yml file, define the actions to run the unit tests with an output of CUCUMBERJSON in the build phase section
- P. Configure the test reports to be sent to the new CodeArtifact repository.
- Q. Create a new AWS CodeBuild project
- R. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project
- S. Create a new Amazon S3 bucket
- T. Create a buildspec.yml file in the CodeCommit repository
- . In the buildspec.yml file, define the actions to run the unit tests with an output of HTML in the phases section
- . In the reports section, upload the test reports to the S3 bucket.

**Answer: B**

**Explanation:**

The correct answer is B. Creating a new AWS CodeBuild project and configuring a test stage in the AWS CodePipeline pipeline that uses the new CodeBuild project is the best way to integrate the unit tests into the existing pipeline. Creating a CodeBuild report group and uploading the test reports to the new CodeBuild report group will produce reports about the unit tests for the company to view. Using JUNITXML as the output format for the unit tests is supported by CodeBuild and will generate a valid report. Option A is incorrect because Amazon CodeGuru Reviewer is a service that provides automated code reviews and recommendations for improving code quality and performance. It is not a tool for running unit tests or producing test reports. Therefore, option A will not meet the requirements.

Option C is incorrect because AWS CodeArtifact is a service that provides secure, scalable, and cost-effective artifact management for software development. It is not a tool for running unit tests or producing test reports. Moreover, option C uses CUCUMBERJSON as the output format for the unit tests, which is not supported by CodeBuild and will not generate a valid report.

Option D is incorrect because uploading the test reports to an Amazon S3 bucket is not the best way to produce reports about the unit tests for the company to view. CodeBuild has a built-in feature to create and manage test reports, which is more convenient and efficient than using S3. Furthermore, option D uses HTML as the output format for the unit tests, which is not supported by CodeBuild and will not generate a valid report.

#### NEW QUESTION 81

The security team depends on AWS CloudTrail to detect sensitive security issues in the company's AWS account. The DevOps engineer needs a solution to auto-remediate CloudTrail being turned off in an AWS account.

What solution ensures the LEAST amount of downtime for the CloudTrail log deliveries?

- A. Create an Amazon EventBridge rule for the CloudTrail StopLogging even
- B. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called
- C. Add the Lambda function ARN as a target to the EventBridge rule.
- D. Deploy the AWS-managed CloudTrail-enabled AWS Config rule set with a periodic interval to 1 hour
- E. Create an Amazon EventBridge rule for AWS Config rules compliance changes
- F. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called
- G. Add the Lambda function ARN as a target to the EventBridge rule.
- H. Create an Amazon EventBridge rule for a scheduled event every 5 minutes
- I. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on a CloudTrail trail in the AWS account
- J. Add the Lambda function ARN as a target to the EventBridge rule.
- K. Launch a t2 nano instance with a script running every 5 minutes that uses the AWS SDK to query CloudTrail in the current account
- L. If the CloudTrail trail is disabled have the script re-enable the trail.

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 86

A DevOps engineer manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an EC2 Auto Scaling group across multiple Availability Zones. The engineer needs to implement a deployment strategy that:

Launches a second fleet of instances with the same capacity as the original fleet. Maintains the original fleet unchanged while the second fleet is launched.

Transitions traffic to the second fleet when the second fleet is fully deployed. Terminates the original fleet automatically 1 hour after transition.

Which solution will satisfy these requirements?

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

**Answer:** C

#### Explanation:

[https://docs.aws.amazon.com/codedeploy/latest/APIReference/API\\_BlueInstanceTerminationOption.html](https://docs.aws.amazon.com/codedeploy/latest/APIReference/API_BlueInstanceTerminationOption.html)

The original revision termination settings are configured to wait 1 hour after traffic has been rerouted before terminating the blue task set.

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/deployment-type-bluegreen.html>

#### NEW QUESTION 87

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

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

How can log collection be automated?

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

**Answer:** D

#### Explanation:

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

#### NEW QUESTION 91

An online retail company based in the United States plans to expand its operations to Europe and Asia in the next six months. Its product currently runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. All data is stored in an Amazon Aurora database instance.

When the product is deployed in multiple regions, the company wants a single product catalog across all regions, but for compliance purposes, its customer information and purchases must be kept in each region.

How should the company meet these requirements with the LEAST amount of application changes?

- A. Use Amazon Redshift for the product catalog and Amazon DynamoDB tables for the customer information and purchases.

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

**Answer: C**

#### NEW QUESTION 93

AnyCompany is using AWS Organizations to create and manage multiple AWS accounts. AnyCompany recently acquired a smaller company, Example Corp. During the acquisition process, Example Corp's single AWS account joined AnyCompany's management account through an Organizations invitation. AnyCompany moved the new member account under an OU that is dedicated to Example Corp. AnyCompany's DevOps engineer has an IAM user that assumes a role that is named OrganizationAccountAccessRole to access member accounts. This role is configured with a full access policy. When the DevOps engineer tries to use the AWS Management Console to assume the role in Example Corp's new member account, the DevOps engineer receives the following error message: "Invalid information in one or more fields. Check your information or contact your administrator." Which solution will give the DevOps engineer access to the new member account?

- A. In the management account, grant the DevOps engineer's IAM user permission to assume the OrganizationAccountAccessRole IAM role in the new member account.
- B. In the management account, create a new SCP. In the SCP, grant the DevOps engineer's IAM user full access to all resources in the new member account.
- C. Attach the SCP to the OU that contains the new member account.
- D. In the new member account, create a new IAM role that is named OrganizationAccountAccessRole.
- E. Attach the AdministratorAccess AWS managed policy to the role.
- F. In the role's trust policy, grant the management account permission to assume the role.
- G. In the new member account, edit the trust policy for the OrganizationAccountAccessRole IAM role.
- H. Grant the management account permission to assume the role.

**Answer: C**

#### Explanation:

The problem is that the DevOps engineer cannot assume the OrganizationAccountAccessRole IAM role in the new member account that joined AnyCompany's management account through an Organizations invitation. The solution is to create a new IAM role with the same name and trust policy in the new member account.

? Option A is incorrect, as it does not address the root cause of the error. The DevOps engineer's IAM user already has permission to assume the OrganizationAccountAccessRole IAM role in any member account, as this is the default role name that AWS Organizations creates when a new account joins an organization. The error occurs because the new member account does not have this role, as it was not created by AWS Organizations.

? Option B is incorrect, as it does not address the root cause of the error. An SCP is a policy that defines the maximum permissions for account members of an organization or organizational unit (OU). An SCP does not grant permissions to IAM users or roles, but rather limits the permissions that identity-based policies or resource-based policies grant to them. An SCP also does not affect how IAM roles are assumed by other principals.

? Option C is correct, as it addresses the root cause of the error. By creating a new IAM role with the same name and trust policy as the OrganizationAccountAccessRole IAM role in the new member account, the DevOps engineer can assume this role and access the account. The new role should have the AdministratorAccess AWS managed policy attached, which grants full access to all AWS resources in the account. The trust policy should allow the management account to assume the role, which can be done by specifying the management account ID as a principal in the policy statement.

? Option D is incorrect, as it assumes that the new member account already has the OrganizationAccountAccessRole IAM role, which is not true. The new member account does not have this role, as it was not created by AWS Organizations. Editing the trust policy of a non-existent role will not solve the problem.

#### NEW QUESTION 94

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

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

**Answer: A**

#### Explanation:

Update the buildspec.yml file to log in to the ECR repository by using the aws ecr get-login-password AWS CLI command to obtain an authentication token. Update the docker login command to use the authentication token to access the ECR repository.

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

#### NEW QUESTION 95

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

- A. Configure a CloudWatch Logs subscription filter to use AWS Glue to transfer all logs to an S3 bucket.

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

**Answer:** BD

**Explanation:**

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

**NEW QUESTION 99**

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

The company wants to automate the task of adding new nodes to a cluster. What can a DevOps engineer do to meet these requirements?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 101**

A company manages multiple AWS accounts in AWS Organizations. The company's security policy states that AWS account root user credentials for member accounts must not be used. The company monitors access to the root user credentials.

A recent alert shows that the root user in a member account launched an Amazon EC2 instance. A DevOps engineer must create an SCP at the organization's root level that will prevent the root user in member accounts from making any AWS service API calls.

Which SCP will meet these requirements?

A)

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

B)

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

C)

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

D)

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

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

**Answer: D**

**NEW QUESTION 104**

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-xxxxxxxxxxxx"
          ]
        }
      }
    }
  ]
}
```

A development team is building a new project in an account that is in an organization in AWS Organizations. The development team wants to use a Python library that has already been stored in the CodeArtifact repository in the organization. The development team uses AWS CodePipeline and AWS CodeBuild to build the new application. The CodeBuild job that the development team uses to build the application is configured to run in a VPC. Because of compliance requirements the VPC has no internet connectivity.

The development team creates the VPC endpoints for CodeArtifact and updates the CodeBuild buildspec yaml file. However, the development team cannot download the Python library from the repository.

Which combination of steps should a DevOps engineer take so that the development team can use Code Artifact? (Select TWO.)

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

**Answer:** AD

**Explanation:**

"AWS CodeArtifact operates in multiple Availability Zones and stores artifact data and metadata in Amazon S3 and Amazon DynamoDB. Your encrypted data is redundantly stored across multiple facilities and multiple devices in each facility, making it highly available and highly durable."

<https://aws.amazon.com/codeartifact/features/> With no internet connectivity, a gateway endpoint becomes necessary to access S3.

**NEW QUESTION 108**

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

How should a DevOps engineer meet these requirements?

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

**Answer:** D

**NEW QUESTION 111**

A company hosts a security auditing application in an AWS account. The auditing application uses an IAM role to access other AWS accounts. All the accounts are in the same organization in AWS Organizations.

A recent security audit revealed that users in the audited AWS accounts could modify or delete the auditing application's IAM role. The company needs to prevent

any modification to the auditing application's IAM role by any entity other than a trusted administrator IAM role. Which solution will meet these requirements?

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

**Answer: A**

**Explanation:**

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

**NEW QUESTION 112**

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

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

**Answer: D**

**Explanation:**

Use routing configuration on an alias to send a portion of traffic to a second function version. For example, you can reduce the risk of deploying a new version by configuring the alias to send most of the traffic to the existing version, and only a small percentage of traffic to the new version.

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html>

The following are the steps involved in the deploy stage configuration that will meet the requirements:

- ? Use AWS CodeBuild to add sample event payloads for testing to the Lambda functions.
- ? Publish a new version of the functions, and include Amazon CloudWatch alarms.
- ? Update the production alias to point to the new version.
- ? Configure rollbacks to occur when an alarm is in the ALARM state.

This configuration will help to reduce the customer impact of an unsuccessful deployment

by deploying the new version of the functions to a staging environment first. This will allow the DevOps engineer to test the new version of the functions before deploying it to production.

The configuration will also help to monitor for issues by including Amazon CloudWatch alarms. These alarms will alert the DevOps engineer if there are any problems with the new version of the functions.

**NEW QUESTION 114**

A company has a guideline that every Amazon EC2 instance must be launched from an AMI that the company's security team produces. Every month the security team sends an email message with the latest approved AMIs to all the development teams.

The development teams use AWS CloudFormation to deploy their applications. When developers launch a new service they have to search their email for the latest AMIs that the security department sent. A DevOps engineer wants to automate the process that the security team uses to provide the AMI IDs to the development teams.

What is the MOST scalable solution that meets these requirements?

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 117**

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 organization
- D. Create an SCP to prevent users from modifying VPC flow logs.
- E. Turn on AWS Config
- F. Create an AWS Config rule to check whether VPC flow logs are turned on
- 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 logs
- 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 122**

A company has its AWS accounts in an organization in AWS Organizations. AWS Config is manually configured in each AWS account. The company needs to implement a solution to centrally configure AWS Config for all accounts in the organization. The solution also must record resource changes to a central account. Which combination of actions should a DevOps engineer perform to meet these requirements? (Choose two.)

- A. Configure a delegated administrator account for AWS Config
- B. Enable trusted access for AWS Config in the organization.
- C. Configure a delegated administrator account for AWS Config
- D. Create a service-linked role for AWS Config in the organization's management account.
- E. Create an AWS CloudFormation template to create an AWS Config aggregator
- F. Configure a CloudFormation stack set to deploy the template to all accounts in the organization.
- G. Create an AWS Config organization aggregator in the organization's management account
- H. Configure data collection from all AWS accounts in the organization and from all AWS Regions.
- I. Create an AWS Config organization aggregator in the delegated administrator account
- J. Configure data collection from all AWS accounts in the organization and from all AWS Regions.

**Answer: AE**

**Explanation:**

<https://aws.amazon.com/blogs/mt/org-aggregator-delegated-admin/> <https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-config.html>

**NEW QUESTION 125**

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

Which change should the DevOps engineer make to launch the EC2 instance in the `vendor-data` account?

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

**Answer:** C

**Explanation:**

The correct answer is C. Updating the SCP in the child OU to allow all actions for Amazon EC2 will enable the DevOps engineer to launch the EC2 instance in the vendor-data account. SCPs are applied to OUs and accounts in a hierarchical manner, meaning that the SCPs attached to the parent OU are inherited by the child OU and accounts. Therefore, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. By adding EC2 to the allowed actions in the child OU's SCP, the DevOps engineer can access EC2 resources in the vendor-data account.

Option A is incorrect because attaching the AmazonEC2FullAccess IAM policy to the IAM user will not grant the user access to EC2 resources. IAM policies are evaluated after SCPs, so even if the IAM policy allows EC2 actions, the SCP will still deny them.

Option B is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the vendor-data account will not work. SCPs are not cumulative, meaning that only one SCP is applied to an account at a time. The SCP attached to the account will be the SCP attached to the OU that contains the account. Therefore, option B will not change the SCP that is applied to the vendor-data account.

Option D is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the root OU will not work. As explained earlier, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. Therefore, option D will not affect the SCP that is applied to the vendor-data account.

**NEW QUESTION 129**

A company deploys updates to its Amazon API Gateway API several times a week by using an AWS CodePipeline pipeline. As part of the update process the company exports the JavaScript SDK for the API from the API Gateway console and uploads the SDK to an Amazon S3 bucket

The company has configured an Amazon CloudFront distribution that uses the S3 bucket as an origin Web client then download the SDK by using the CloudFront distribution's endpoint. A DevOps engineer needs to implement a solution to make the new SDK available automatically during new API deployments.

Which solution will meet these requirements?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 131**

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 133**

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 account
- 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 account
- D. Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applied
- 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 account
- 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 weekly
- I. Specify the runbook as the target of the rule.
- J. Turn on AWS CloudTrail in the account
- 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 weekly
- 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 135**

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 branch
- 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 time
- 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 server
- 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 139**

A company's development team uses AWS CloudFormation to deploy its application resources. The team must use for any changes to the environment. The team cannot use the AWS Management Console or the AWS CLI to make manual changes directly.

The team uses a developer IAM role to access the environment. The role is configured with the AdministratorAccess managed policy. The company has created a new CloudFormationDeployment IAM role that has the following policy.

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

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

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

**Answer:** ADF

**Explanation:**

A comprehensive and detailed explanation is:

? Option A is correct because removing the AdministratorAccess policy and assigning the ReadOnlyAccess managed IAM policy to the developer role is a valid way to prevent the developers from making any manual changes to the deployed resources. The AdministratorAccess policy grants full access to all AWS resources and actions, which is not necessary for the developers. The ReadOnlyAccess policy grants read-only access to most AWS resources and actions, which is sufficient for the developers to view the status of their stacks. Instructing the developers to use the CloudFormationDeployment role as a CloudFormation service role when they deploy new stacks is also a valid way to ensure that only CloudFormation can use the new role. A CloudFormation service role is an IAM role that allows CloudFormation to make calls to resources in a stack on behalf of the user1. The user can specify a service role when they create or update a stack, and CloudFormation will use that role's credentials for all operations that are performed on that stack1.

? Option B is incorrect because updating the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDeployment role is not a valid solution. This would allow the developers to manually assume the CloudFormationDeployment role and perform actions on the deployed resources, which is not what the company wants. The trust of CloudFormationDeployment role should only allow the cloudformation.amazonaws.com AWS principal to assume the role, as in option D.

? Option C is incorrect because configuring the IAM user to be able to get and pass the CloudFormationDeployment role if cloudformation actions for resources is not a valid solution. This would allow the developers to manually pass the CloudFormationDeployment role to other services or resources, which is not what the company wants. The IAM user should only be able to pass the CloudFormationDeployment role as a service role when they create or update a stack with CloudFormation, as in option A.

? Option D is correct because updating the trust of CloudFormationDeployment role to allow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeRole action is a valid solution. This allows CloudFormation to assume the CloudFormationDeployment role and access resources in other services on behalf of the user2. The trust policy of an IAM role defines which entities can assume the role2. By specifying cloudformation.amazonaws.com as the principal, you grant permission only to CloudFormation to assume this role.

? Option E is incorrect because instructing the developers to assume the CloudFormationDeployment role when they deploy new stacks is not a valid solution. This would allow the developers to manually assume the CloudFormationDeployment role and perform actions on the deployed resources, which is not what the company wants. The developers should only use the CloudFormationDeployment role as a service role when they deploy new stacks with CloudFormation, as in option A.

? Option F is correct because adding an IAM policy to CloudFormationDeployment that allows cloudformation:\* on all resources and adding a policy that allows the iam:PassRole action for ARN of CloudFormationDeployment if iam:PassedToService equals cloudformation.amazonaws.com are valid solutions. The first policy grants permission for CloudFormationDeployment to perform any action with any resource using cloudformation.amazonaws.com as a service principal3. The second policy grants permission for passing this role only if it is passed by cloudformation.amazonaws.com as a service principal4. This ensures that only CloudFormation can use this role.

References:

- ? 1: AWS CloudFormation service roles
- ? 2: How to use trust policies with IAM roles
- ? 3: AWS::IAM::Policy
- ? 4: IAM: Pass an IAM role to a specific AWS service

**NEW QUESTION 143**

.....

## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions 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 First Try
- \* AWS-Certified-DevOps-Engineer-Professional Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
[Order The AWS-Certified-DevOps-Engineer-Professional Practice Test Here](#)