

# Amazon

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

Amazon AWS Certified DevOps Engineer Professional



### NEW QUESTION 1

An application runs on Amazon EC2 instances behind an Application Load Balancer. Amazon RDS MySQL is used on the backend. The instances run in an Auto Scaling group across multiple Availability Zones. The Application Load Balancer health check ensures the web servers are operating and able to make read/write SQL connections. Amazon Route 53 provides DNS functionality with a record pointing to the Application Load Balancer. A new policy requires a geographically isolated disaster recovery site with an RTO of 4 hours and an RPO of 15 minutes.

Which disaster recovery strategy will require the LEAST amount of changes to the application stack?

- A. Launch a replica stack of everything except RDS in a different Availability Zone
- B. Create an RDS read-only replica in a new Availability Zone and configure the new stack to point to the local RDS instance
- C. Add the new stack to the Route 53 record set with a failover routing policy.
- D. Launch a replica stack of everything except RDS in a different region
- E. Create an RDS read-only replica in a new region and configure the new stack to point to the local RDS instance
- F. Add the new stack to the Route 53 record set with a latency routing policy.
- G. Launch a replica stack of everything except RDS in a different region
- H. Upon failure, copy the snapshot over from the primary region to the disaster recovery region
- I. Adjust the Amazon Route 53 record set to point to the disaster recovery region's Application Load Balancer.
- J. Launch a replica stack of everything except RDS in a different region
- K. Create an RDS read-only replica in a new region and configure the new stack to point to the local RDS instance
- L. Add the new stack to the Amazon Route 53 record set with a failover routing policy

**Answer: D**

### NEW QUESTION 2

A company wants to use AWS development tools to replace its current bash deployment scripts. The company currently deploys a LAMP application to a group of Amazon EC2 instances behind an Application Load Balancer (ALB). During the deployments, the company unit tests the committed application, stops and starts services, unregisters and re-registers instances with the load balancer, and updates file permissions. The company wants to maintain the same deployment functionality through the shift to using AWS services.

Which solution will meet these requirements?

- A. Use AWS CodeBuild to test the application
- B. Use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart services, and deregister and register instances with the ALB. Use the appspec.yml file to update file permissions without a custom script.
- C. Use AWS CodePipeline to move the application from the AWS CodeCommit repository to AWS CodeDeploy
- D. Use CodeDeploy's deployment group to test the application, unregister and reregister instances with the ALB, and restart service
- E. and restart service
- F. Use the appspec.yml file to update file permissions without a custom script.
- G. Use AWS CodePipeline to move the application source code from the AWS CodeCommit repository to AWS CodeDeploy
- H. Use CodeDeploy to test the application
- I. Use CodeDeploy's appspec.yml file to restart services and update permissions without a custom script
- J. Use AWS CodeBuild to unregister and re-register instances with the ALB.
- K. Use AWS CodePipeline to trigger AWS CodeBuild to test the application. Use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart service
- L. Unregister and re-register the instances in the AWS CodeDeploy deployment group with the ALB
- M. Update the appspec.yml file to update file permissions without a custom script.

**Answer: D**

### NEW QUESTION 3

A DevOps Engineer has a single Amazon DynamoDB table that received shipping orders and tracks inventory. The Engineer has three AWS Lambda functions reading from a DynamoDB stream on that table. The Lambda functions perform various functions such as doing an item count, moving items to Amazon Kinesis Data Firehose, monitoring inventory levels, and creating vendor orders when parts are low.

While reviewing logs, the Engineer notices the Lambda functions occasionally fail under increased load, receiving a stream throttling error.

Which is the MOST cost-effective solution that requires the LEAST amount of operational management?

- A. Use AWS Glue integration to ingest the DynamoDB stream, then migrate the Lambda code to an AWS Fargate task.
- B. Use Amazon Kinesis streams instead of DynamoDB streams, then use Kinesis analytics to trigger the Lambda functions.
- C. Create a fourth Lambda function and configure it to be the only Lambda reading from the stream
- D. Then use this Lambda function to pass the payload to the other three Lambda functions.
- E. Have the Lambda functions query the table directly and disable DynamoDB stream
- F. Then have the Lambda functions query from a global secondary index.

**Answer: C**

### NEW QUESTION 4

You have deployed an application to AWS which makes use of Autoscaling to launch new instances. You now want to change the instance type for the new instances. Which of the following is one of the action items to achieve this deployment?

- A. Use Elastic Beanstalk to deploy the new application with the new instance type
- B. Use CloudFormation to deploy the new application with the new instance type
- C. Create a new launch configuration with the new instance type
- D. Create new EC2 instances with the new instance type and attach it to the Autoscaling Group

**Answer: C**

#### Explanation:

The ideal way is to create a new launch configuration, attach it to the existing Auto Scaling group, and terminate the running instances.

Option A is invalid because Elastic Beanstalk cannot launch new instances on demand. Since the current scenario requires Autoscaling, this is not the ideal option

Option B is invalid because this will be a maintenance overhead, since you just have an Autoscaling Group.

There is no need to create a whole CloudFormation

template for this.

Option D is invalid because Autoscaling Group will still launch EC2 instances with the older launch configuration

For more information on Autoscaling Launch configuration, please refer to the below document link: from AWS

> [http://docs.aws.amazon.com/autoscaling/latest/userguide/l\\_launchConfiguration.html](http://docs.aws.amazon.com/autoscaling/latest/userguide/l_launchConfiguration.html)

**NEW QUESTION 5**

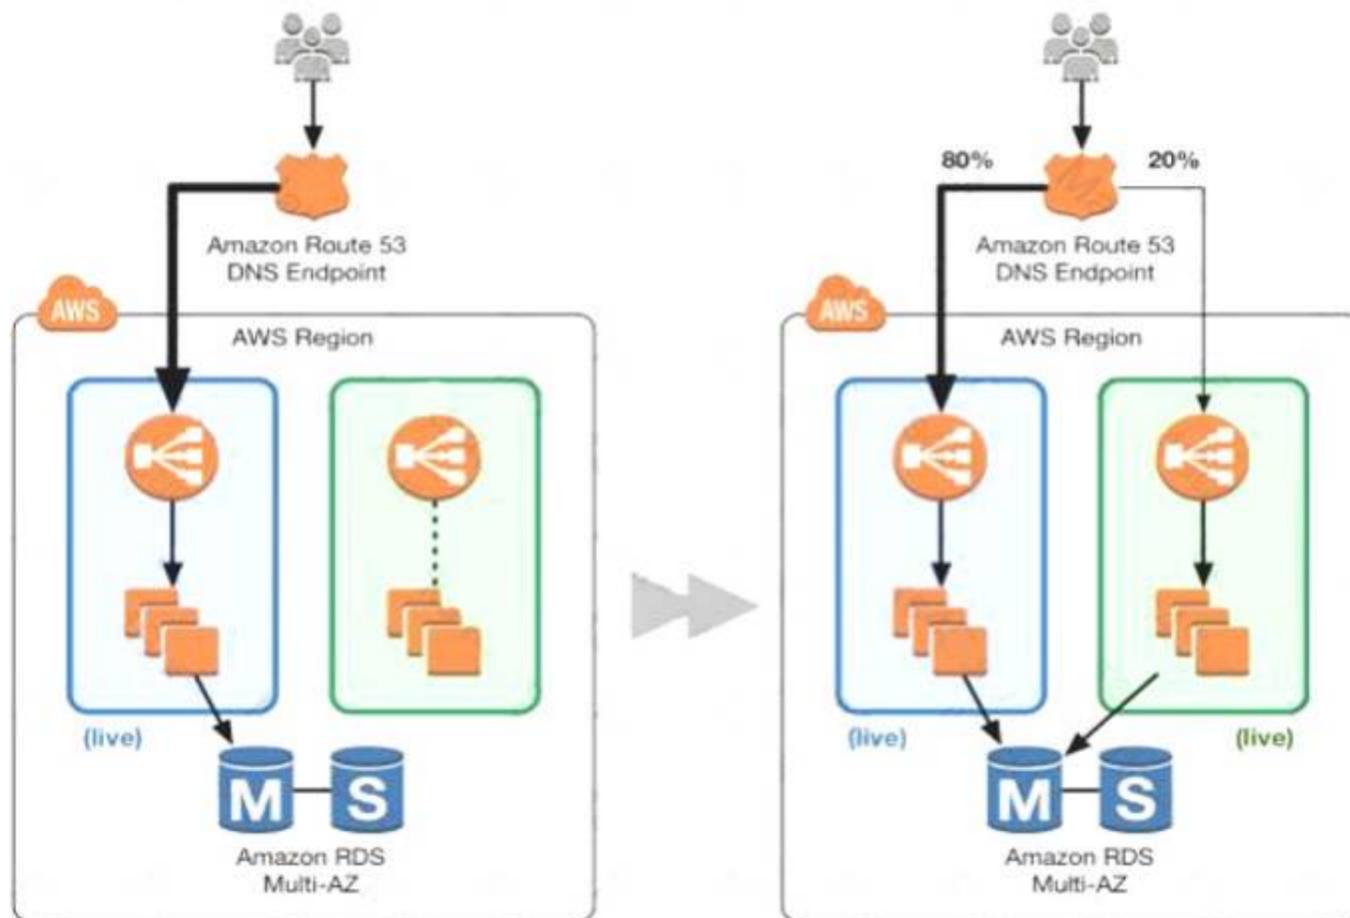
Your application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment?

- A. Set up Amazon Route 53 health checks to fail over from any Amazon EC2 instance that is currently being deployed to.
- B. Using AWS CloudFormation, create a test stack for validating the code, and then deploy the code to each production Amazon EC2 instance.
- C. Create a new load balancer with new Amazon EC2 instances, carry out the deployment, and then switch DNS over to the new load balancer using Amazon Route 53 after testing.
- D. Launch more Amazon EC2 instances to ensure high availability, de-register each Amazon EC2 instance from the load balancer, upgrade it, and test it, and then register it again with the load balancer.

**Answer: C**

**Explanation:**

The below diagram shows how this can be done C:\Users\wk\Desktop\mudassar\Untitled.jpg



- 1) First create a new ELB which will be used to point to the new production changes.
- 2) Use the Weighted Route policy for Route53 to distribute the traffic to the 2 ELB's based on a 80-20% traffic scenario. This is the normal case, the % can be changed based on the requirement.
- 3) Finally when all changes have been tested, Route53 can be set to 100% for the new ELB.

Option A is incorrect because this is a failover scenario and cannot be used for Blue green deployments. In Blue Green deployments, you need to have 2 environments running side by side.

Option B is incorrect, because you need to have a production stack with the changes which will run side by side.

Option D is incorrect because this is not a blue green deployment scenario. You cannot control which users will go to the new EC2 instances.

For more information on blue green deployments, please refer to the below document link: from AWS >

[https://dOawsstatic.com/whitepapers/AWS\\_Blue\\_Green\\_Deployments.pdf](https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf)

**NEW QUESTION 6**

A company recently launched an application that is more popular than expected. The company wants to ensure the application can scale to meet increasing demands and provide reliability using multiple Availability Zones (AZs). The application runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). A DevOps engineer has created an Auto Scaling group across multiple AZs for the application. Instances launched in the newly added AZs are not receiving any traffic for the application.

What is likely causing this issue?

- A. Auto Scaling groups can create new instances in a single AZ only.
- B. The EC2 instances have not been manually associated to the ALB.
- C. The ALB should be replaced with a Network Load Balancer (NLB).
- D. The new AZ has not been added to the ALB.

**Answer: A**

**NEW QUESTION 7**

A DevOps Engineer is leading the implementation for automating patching of Windows-based workstations in a hybrid cloud environment by using AWS Systems Manager (SSM).

What steps should the Engineer follow to set up Systems Manager to automate patching in this environment? (Select TWO.)

- A. Create multiple IAM service roles for Systems Manager so that the ssm.amazonaws.com service can execute the AssumeRole operation on every instance.
- B. Register the role on a per-resource level to enable the creation of a service token.
- C. Perform managed-instance activation with the newly created service role attached to each managed instance.
- D. Create an IAM service role for Systems Manager so that the ssm.amazonaws.com service can execute the AssumeRole operation.
- E. Register the role to enable the creation of a service token.
- F. Perform managed-instance activation with the newly created service role.
- G. Using previously obtained activation codes and activation IDs, download and install the SSM Agent on the hybrid servers, and register the servers or virtual machines on the Systems Manager service.
- H. Hybrid instances will show with an "mi-" prefix in the SSM console.
- I. Using previously obtained activation codes and activation IDs, download and install the SSM Agent on the hybrid servers, and register the servers or virtual machines on the Systems Manager service.
- J. Hybrid instances will show with an "i-" prefix in the SSM console as if they were provisioned as a regular Amazon EC2 instance.
- K. Run AWS Config to create a list of instances that are unpatched and not compliant.
- L. Create an instance scheduler job, and through an AWS Lambda function, perform the instance patching to bring them up to compliance.

**Answer:** BC

**Explanation:**

<https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-managed-instance-activation.html>  
<https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-install-managed-win.html>

**NEW QUESTION 8**

A company wants to migrate a legacy application to AWS and develop a deployment pipeline that uses AWS services only. A DevOps engineer is migrating all of the application code from a Git repository to AWS CodeCommit while preserving the history of the repository. The DevOps engineer has set all the permissions within CodeCommit, installed the Git client and the AWS CLI on a local computer, and is ready to migrate the repository. Which actions will follow?

- A. Create the CodeCommit repository using the AWS CLI.
- B. Clone the Git repository directly to CodeCommit using the AWS CLI.
- C. Validate that the files were migrated, and publish the CodeCommit repository.
- D. Create the CodeCommit repository using the AWS Management Console.
- E. Clone both the Git and CodeCommit repositories to the local computer.
- F. Copy the files from the Git repository to the CodeCommit repository on the local computer.
- G. Commit the CodeCommit repository.
- H. Validate that the files were migrated, and share the CodeCommit repository.
- I. Create the CodeCommit repository using the AWS Management Console.
- J. Use the console to clone the Git repository into the CodeCommit repository.
- K. Validate that the files were migrated, and publish the CodeCommit repository.
- L. Create the CodeCommit repository using the AWS Management Console or the AWS CLI.
- M. Clone the Git repository with a mirror argument to the local computer and push the repository to CodeCommit.
- N. Validate that the files were migrated, and share the CodeCommit repository.

**Answer:** D

**NEW QUESTION 9**

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 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

**NEW QUESTION 10**

A company has an application that has predictable peak traffic times. The company wants the application instances to scale up only during the peak times. The application stores state in Amazon DynamoDB. The application environment uses a standard Node.js application stack and custom Chef recipes stored in a private Git repository. Which solution is MOST cost-effective and requires the LEAST amount of management overhead when performing rolling updates of the application environment?

- A. Create a custom AMI with the Node.js environment and application stack using Chef recipe.
- B. Use the AMI in an Auto Scaling group and set up scheduled scaling for the required times, then set up an Amazon EC2 IAM role that provides permission to access DynamoDB.
- C. Create a Docker file that uses the Chef recipes for the application environment based on an official Node.js Docker image.
- D. Create an Amazon ECS cluster and a service for the application environment, then create a task based on this Docker image.
- E. Use scheduled scaling to scale the containers at the appropriate times and attach a task-level IAM role that provides permission to access DynamoDB.
- F. Configure AWS OpsWorks stacks and use custom Chef cookbook.
- G. Add the Git repository information where the custom recipes are stored, and add a layer in OpsWorks for the Node.js application server.
- H. Then configure the custom recipe to deploy the application in the deploy step.
- I. Configure time-based instances and attach an Amazon EC2 IAM role that provides permission to access DynamoDB.
- J. Configure AWS OpsWorks stacks and push the custom recipes to an Amazon S3 bucket and configure custom recipes to point to the S3 bucket.
- K. Then add an application layer type for a standard Node.js application server and configure the custom recipe to deploy the application in the deploy step from the S3 bucket.
- L. Configure time-based instances and attach an Amazon EC2 IAM role that provides permission to access DynamoDB.

**Answer:** D

#### NEW QUESTION 10

A company indexes all of its Amazon CloudWatch Logs on Amazon ES and uses Kibana to view a dashboard for actionable insight. The company wants to restrict user access to Kibana by user

Which actions can a DevOps Engineer take to meet this requirement? (Select TWO.)

- A. Create a proxy server with user authentication in an Auto Scaling group and restrict access of the Amazon ES endpoint to an Auto Scaling group tag
- B. Create a proxy server with user authentication and an Elastic IP address and restrict access of the Amazon ES endpoint to the IP address
- C. Create a proxy server with AWS IAM user and restrict access of the Amazon ES endpoint to the IAM user
- D. Use AWS SSO to offer user name and password protection for Kibana
- E. Use Amazon Cognito to offer user name and password protection for Kibana

**Answer: BE**

#### NEW QUESTION 12

A Development team creates a build project in AWS CodeBuild. The build project invokes automated tests of modules that access AWS services.

Which of the following will enable the tests to run the MOST securely?

- A. Generate credentials for an IAM user with a policy attached to allow the actions on AWS service
- B. Store credentials as encrypted environment variables for the build project
- C. As part of the build script, obtain the credentials to run the integration tests.
- D. Have CodeBuild run only the integration tests as a build job on a Jenkins server
- E. Create a role that has a policy attached to allow the actions on AWS service
- F. Generate credentials for an IAM user that is allowed to assume the role
- G. Configure the credentials as secrets in Jenkins, and allow the build job to use them to run the integration tests.
- H. Create a service role in IAM to be assumed by CodeBuild with a policy attached to allow the actions on AWS service
- I. Configure the build project to use the role created.
- J. Use AWS managed credential
- K. Encrypt the credentials with AWS KMS
- L. As part of the build script, decrypt with AWS KMS and use these credentials to run the integration tests.

**Answer: B**

#### NEW QUESTION 13

A company is running an application on Amazon EC2 instances in an Auto Scaling group. Recently, an issue occurred that prevented EC2 instances from launching successfully, and it took several hours for the support team to discover the issue. The support team wants to be notified by email whenever an EC2 instance does not start successfully.

Which action will accomplish this?

- A. Add a health check to the Auto Scaling group to invoke an AWS Lambda function whenever an instance status is impaired.
- B. Configure the Auto Scaling group to send a notification to an Amazon SNS topic whenever a failed instance launch occurs.
- C. Create an Amazon CloudWatch alarm that invokes an AWS Lambda function when a failed AttachInstances Auto Scaling API call is made.
- D. Create a status check alarm on Amazon EC2 to send a notification to an Amazon SNS topic whenever a status check fail occurs.

**Answer: B**

#### NEW QUESTION 16

A development team is using AWS CodeCommit to version control application code and AWS CodePipeline to orchestrate software deployments. The team has decided to use a remote master branch as the trigger (or the pipeline to integrate code changes. A developer has pushed code changes to the CodeCommit repository, but noticed that the pipeline had no reaction, even after 10 minutes.

Which of the following actions should be taken to troubleshoot this issue?

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

**Answer: C**

#### NEW QUESTION 17

A DevOps engineer needs to back up sensitive Amazon S3 objects that are stored within an S3 bucket with a private bucket policy using S3 cross-Region replication functionality. The objects need to be copied to a target bucket in a different AWS Region and account.

Which combination of actions should be performed to enable this replication? (Select THREE.)

- A. Create a replication IAM role in the source account.
- B. Create a replication IAM role in the target account.
- C. Add statements to the source bucket policy allowing the replication IAM role to replicate objects
- D. Add statements to the target bucket policy allowing the replication IAM role to replicate objects.
- E. Create a replication rule in the source bucket to enable the replication.
- F. Create a replication rule in the target bucket to enable the replication

**Answer: BE**

#### NEW QUESTION 19

The Development team at an online retailer has moved to Business support and want to take advantage of the AWS Health Dashboard and the AWS Health API to automate remediation actions for issues with the health of AWS resources. The first use case is to respond to AWS detecting an IAM access key that is listed on a public code repository site. The automated response will be to delete the IAM access key and send a notification to the Security team.

How should this be achieved?

- A. Create an AWS Lambda function to delete the IAM access ke
- B. Send AWS CloudTrail logs to AWS CloudWatch log
- C. Create a CloudWatch Logs metric filter for the AWS\_RISK\_CREDENTIALS\_EXPOSED event with two actions: first, run the Lambda function; second, use Amazon SNS to send a notification to the Security team.
- D. Create an AWS Lambda function to delete the IAM access ke
- E. Create an AWS Config rule for changes to aws.health and the AWS\_RISK\_CREDENTIALS\_EXPOSED event with two actions: first, run the Lambda function; second, use Amazon SNS to send a notification to the Security team.
- F. Use AWS Step Functions to create a function to delete the IAM access key, and then use Amazon SNS to send a notification to the Security tea
- G. Create an AWS Personal Health Dashboard rule for the AWS\_RISK\_CREDENTIALS\_EXPOSED event; set the target of the Personal Health Dashboard rule to Step Functions.
- H. Use AWS Step Functions to create a function to delete the IAM access key, and then use Amazon SNS to send a notification to the Security tea
- I. Create an Amazon CloudWatch Events rule with an aws.health event source and the AWS\_RISK\_CREDENTIALS\_EXPOSED event, set the target of the CloudWatch Events rule to Step Functions.

**Answer: A**

#### NEW QUESTION 21

A company has built a web service that runs on Amazon EC2 instances behind an Application Load Balancer (ALB) the company has deployed the application in us-east-1 Amazon Route 53 provides an external DNS that routes traffic from example.com to the application, created with appropriate health checks. The company has deployed a second environment for the application in eu-west-1 the company wants traffic to be routed to whichever environment results in the best response time for each user. If there is an outage in one Region, traffic should be directed to the other environment.

Which configuration will achieve this requirements?

- A. •A subdomain us example com with weighted routing the US ALB with weight 2 and the EU ALB with weight 1•Another subdomain eu.example.com with weighted routing the EU ALB with weight 2 and the US ALU with weight 1•Geolocation routing records for example.com North America aliased to us example.com and Europe aliased to eu.example.com
- B. •A subdomain us example com with latency-based routing the US ALB as the first target and the EU ALB as the second target. •Another subdomain eu.example.com with latency-based routing
- C. The EU ALB as the first target and the US ALB as the second target. •Failover routing records for example.com aliased to us.example.com as the first target and eu.example.com as the second target.
- D. •A subdomain us.example.com with failover routing the US ALB as primary and the EU ALB as secondary •Another subdomain eu.example.com with failover routing the EU ALB as primary and the US ALB as secondary •Latency-based routing records for example.com that are aliased to us example.com and eu.example.com
- E. •A subdomain us.example.com with multivalued answer routing
- F. the US ALB as first and the EU ALB as second •Another subdomain eu.example.com with failover routing the EU ALB as first and the US ALB as second •Failover routing records for example.com that are aliased to us.example.com and eu.example.com

**Answer: B**

#### NEW QUESTION 23

A company requires its internal business teams to launch resources through pre-approved AWS CloudFormation templates only. The security team requires automated monitoring when resources drift from their expected state.

Which strategy should be used to meet these requirements?

- A. Allow users to deploy Cloud Formation stacks using a CloudFormation service role onl
- B. Use CloudFormation drift detection to detect when resources have drifted from their expected state.
- C. Allow users to deploy CloudFormation stacks using a CloudFormation service role onl
- D. Use AWS Config rules to detect when resources have drifted from their expected state.
- E. Allow users to deploy CloudFormation stacks using AWS Service Catalog only Enforce the use of a launch constraint Use AWS Config rules to detect when resources have drifted from their expected state.
- F. Allow users to deploy CloudFormation stacks using AWS Service Catalog only Enforce the use of a template constraint Use Amazon EventBridge (Amazon CloudWatch Events) notifications to detect when resources have drifted from their expected state.

**Answer: B**

#### NEW QUESTION 27

A DevOps Engineer administers an application that manages video files for a video production company. The application runs on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. Data is stored in an Amazon RDS PostgreSQL Multi-AZ DB instance, and the video files are stored in an Amazon S3 bucket. On a typical day, 50 GB of new video are added to the S3 bucket. The Engineer must implement a multi-region disaster recovery plan with the least data loss and the lowest recovery times. The current application infrastructure is already described using AWS CloudFormation.

Which deployment option should the Engineer choose to meet the uptime and recovery objectives for the system?

- A. Launch the application from the CloudFormation template in the second region, which sets the capacity of the Auto Scaling group to 1. Create an Amazon RDS read replica in the second regio
- B. In the second region, enable cross-region replication between the original S3 bucket and a new S3 bucke
- C. To fail over, promote the read replica as maste
- D. Update the CloudFormation stack and increase the capacity of the Auto Scaling group.
- E. Launch the application from the CloudFormation template in the second region, which sets the capacity of the Auto Scaling group to 1. Create a scheduled task to take daily Amazon RDS cross-region snapshots to the second regio
- F. In the second region, enable cross-region replication between the original S3 bucket and Amazon Glacie
- G. In a disaster, launch a new application stack in the second region and restore the database from the most recent snapshot.
- H. Launch the application from the CloudFormation template in the second region which sets the capacity of the Auto Scaling group to 1. Use Amazon CloudWatch Events to schedule a nightly task to take a snapshot of the database, copy the snapshot to the second region, and replace the DB instance in the second region from the snapsho
- I. In the second region, enable cross-region replication between the original S3 bucket and a new S3 bucke
- J. To fail over, increase the capacity of the Auto Scaling group.
- K. Use Amazon CloudWatch Events to schedule a nightly task to take a snapshot of the database and copy the snapshot to the second regio
- L. Create an AWS Lambda function that copies each object to a new S3 bucket in the second region in response to S3 event notification
- M. In the second region, launch the application from the CloudFormation template and restore the database from the most recent snapshot.

**Answer:** A

#### NEW QUESTION 30

A DevOps team manages an API running on-premises that serves as a backend for an Amazon API Gateway endpoint. Customers have been complaining about high response latencies, which the development team has verified using the API Gateway latency metrics in Amazon CloudWatch. To identify the cause, the team needs to collect relevant data without introducing additional latency.

Which actions should be taken to accomplish this? (Select TWO.)

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

**Answer:** AC

#### NEW QUESTION 35

A company has multiple child accounts that are part of an organization in AWS Organizations. The security team needs to review every Amazon EC2 security group and their inbound and outbound rules. The security team wants to programmatically retrieve this information from the child accounts using an AWS Lambda function in the master account of the organization.

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

- A. Create a trust relationship that allows users in the child accounts to assume the master account IAM role.
- B. Create a trust relationship that allows users in the master account to assume the IAM roles of the child accounts.
- C. Create an IAM role in each child account that has access to the AmazonEC2ReadOnlyAccess managed policy.
- D. Create an IAM role in each child account to allow the sts:AssumeRole action against the master account IAM role's ARN.
- E. Create an IAM role in the master account that allows the sts:AssumeRole action against the child account IAM role's ARN.
- F. Create an IAM role in the master account that has access to the AmazonEC2ReadOnlyAccess managed policy.

**Answer:** ADF

#### NEW QUESTION 38

A company runs a production application workload in a single AWS account that uses Amazon Route 53, AWS Elastic Beanstalk, and Amazon RDS. In the event of a security incident, the Security team wants the application workload to fail over to a new AWS account. The Security team also wants to block all access to the original account immediately, with no access to any AWS resources in the original AWS account, during forensic analysis.

What is the most cost-effective way to prepare to fail over to the second account prior to a security incident?

- A. Migrate the Amazon Route 53 configuration to a dedicated AWS account
- B. Mirror the Elastic Beanstalk configuration in a different account
- C. Enable RDS Database Read Replicas in a different account.
- D. Migrate the Amazon Route 53 configuration to a dedicated AWS account
- E. Save/copy the Elastic Beanstalk configuration files in a different AWS account
- F. Copy snapshots of the RDS Database to a different account.
- G. Save/copy the Amazon Route 53 configurations for use in a different AWS account after an incident. Save/copy Elastic Beanstalk configuration files to a different account
- H. Enable the RDS database read replica in a different account.
- I. Save/copy the Amazon Route 53 configurations for use in a different AWS account after an incident. Mirror the configuration of Elastic Beanstalk in a different account
- J. Copy snapshots of the RDS database to a different account.

**Answer:** B

#### NEW QUESTION 40

A government agency has multiple AWS accounts, many of which store sensitive citizen information. A Security team wants to detect anomalous account and network activities (such as SSH brute force attacks) in any account and centralize that information in a dedicated security account. Event information should be stored in an Amazon S3 bucket in the security account, which is monitored by the department's Security Information and Event Manager (SIEM) system.

How can this be accomplished?

- A. Enable Amazon Macie in every account
- B. Configure the security account as the Macie Administrator for every member account using invitation/acceptance
- C. Create an Amazon CloudWatch Events rule in the security account to send all findings to Amazon Kinesis Data Firehose, which should push the findings to the S3 bucket.
- D. Enable Amazon Macie in the security account only
- E. Configure the security account as the Macie Administrator for every member account using invitation/acceptance
- F. Create an Amazon CloudWatch Events rule in the security account to send all findings to Amazon Kinesis Data Stream
- G. Write an application using KCL to read data from the Kinesis Data Streams and write to the S3 bucket.
- H. Enable Amazon GuardDuty in every account
- I. Configure the security account as the GuardDuty Administrator for every member account using invitation/acceptance
- J. Create an Amazon CloudWatch rule in the security account to send all findings to Amazon Kinesis Data Firehose, which will push the findings to the S3 bucket.
- K. Enable Amazon GuardDuty in the security account only
- L. Configure the security account as the GuardDuty Administrator for every member account using invitation/acceptance
- M. Create an Amazon CloudWatch rule in the security account to send all findings to Amazon Kinesis Data Stream
- N. Write an application using KCL to read data from Kinesis Data Streams and write to the S3 bucket.

**Answer:** C

#### Explanation:

<https://aws.amazon.com/blogs/security/how-to-manage-amazon-guardduty-security-findings-across-multiple-acc>

#### NEW QUESTION 42

A company has several AWS accounts. The accounts are shared and used across multiple teams globally, primarily for Amazon EC2 instances. Each EC2 instance has tags for team, environment, and cost center to ensure accurate cost allocations. How should a DevOps Engineer help the teams audit their costs and automate infrastructure cost optimization across multiple shared environments and accounts?

- A. Set up a scheduled script on the EC2 instances to report utilization and store the instances in an Amazon DynamoDB tabl
- B. Create a dashboard in Amazon QuickSight with DynamoDB as the source data to find underutilized instance
- C. Set up triggers from Amazon QuickSight in AWS Lambda to reduce underutilized instances.
- D. Create a separate Amazon CloudWatch dashboard for EC2 instance tags based on cost center, environment, and team, and publish the instance tags out using unique links for each tea
- E. For each team, set up a CloudWatch Events rule with the CloudWatch dashboard as the source, and set up a trigger to initiate an AWS Lambda function to reduce underutilized instances.
- F. Create an Amazon CloudWatch Events rule with AWS Trusted Advisor as the source for low utilization EC2 instance
- G. Trigger an AWS Lambda function that filters out reported data based on tags for each team, environment, and cost center, and store the Lambda function in Amazon S3. Set up a second trigger to initiate a Lambda function to reduce underutilized instances.
- H. Use AWS Systems Manager to track instance utilization and report underutilized instances to Amazon CloudWate
- I. Filter data in CloudWatch based on tags for team, environment, and cost cente
- J. Set up triggers from CloudWatch into AWS Lambda to reduce underutilized instances

**Answer: C**

#### Explanation:

<https://github.com/aws/Trusted-Advisor-Tools/tree/master/LowUtilizationEC2Instances> <https://docs.aws.amazon.com/quicksight/latest/user/supported-data-sources.html>

#### NEW QUESTION 44

A DevOps engineer has been tasked with ensuring that all Amazon S3 buckets, except for those with the word "public" in the name, allow access only to authorized users utilizing S3 bucket policies. The security team wants to be notified when a bucket is created without the proper policy and for the policy to be automatically updated. Which solutions will meet these requirements?

- A. Create a custom AWS Config rule that will trigger an AWS Lambda function when an S3 bucket is created or update
- B. Use the Lambda function to look for S3 buckets that should be private, but that do not have a bucket policy that enforces privac
- C. When such a bucket is found, invoke a remediation action and use Amazon SNS to notify the security team.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that triggers when an S3 bucket is create
- E. Use an AWS Lambda function to determine whether the bucket should be privat
- F. If the bucket should be private, update the PublicAccessBlock configuratio
- G. Configure a second EventBridge (CloudWatch Events) rule to notify the security team using Amazon SNS when PutBucketPolicy is called.
- H. Create an Amazon S3 event notification that triggers when an S3 bucket is created that does not have the word "public" in the nam
- I. Define an AWS Lambda function as a target for this notification and use the function to apply a new default policy to the S3 bucke
- J. Create an additional notification with the same filter and use Amazon SNS to send an email to the security team.
- K. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that triggers when a new object is created in a bucket that does not have the word "public" in the nam
- L. Target and use an AWS Lambda function to update the PublicAccessBlock configuratio
- M. Create an additional notification with the same filter and use Amazon SNS to send an email to the security team.

**Answer: A**

#### NEW QUESTION 45

A company has a legacy application running on AWS. The application can only run on one Amazon EC2 instance at a time. Application metadata is stored in Amazon S3 and must be retrieved if the instance is restarted. The instance should be automatically restarted or relaunched if performance degrades. Which solution will satisfy those requirements?

- A. Create an Amazon Cloud Watch alarm to monitor the EC2 instanc
- B. When the StatusCheckFailed system alarm is triggered use the recover action to stop and start the instanc
- C. Use a trigger in Amazon S3 to push the metadata to the instance when it is back up and running
- D. Use the auto healing feature in AWS OpsWorks to stop and start the EC2 instanc
- E. Use a lifecycle event in OpsWorks to pull the data from Amazon S3 and update it on the instance.
- F. Use the Auto Recovery feature in Amazon EC2 to automatically stop and start the EC2 instance in case of a failure .Use a trigger in Amazon S3 to push the metadata to the instance when it is back up and running
- G. Use AWS CloudFormation to create an EC2 instance that includes the user-data property for the EC2resourc
- H. Add a command in user data to retrieve the application metadata from Amazon S3.

**Answer: C**

#### NEW QUESTION 47

A company's web application will be migrated to AWS. The application is designed so that there is no server-side code required. As part of the migration, the company would like to improve the security of the application by adding HTTP response headers, following the Open Web Application Security Project (OWASP) secure headers recommendations. How can this solution be implemented to meet the security requirements using best practices?

- A. Use an Amazon S3 bucket configured for website hosting, then set up server access logging on the S3 bucket to track user activit
- B. Then configure the static website hosting and execute a scheduled AWS Lambda function to verify, and if missing, add security headers to the metadata.
- C. Use an Amazon S3 bucket configured for website hosting, then set up server access logging on the S3 bucket to track user activit
- D. Configure the static website hosting to return the required security headers.
- E. Use an Amazon S3 bucket configured for website hostin
- F. Create an Amazon CloudFront distribution that refers to this S3 bucket, with the origin response event set to trigger a Lambda@Edge Node.js function to add in the security headers.
- G. set an Amazon S3 bucket configured for website hostin
- H. Create an Amazon CloudFront distribution that refers to this S3 bucke

I. Set "Cache Based on Selected Request Headers" to "Whitelist," and add the security headers into the whitelist.

**Answer:** C

**Explanation:**

<https://aws.amazon.com/blogs/networking-and-content-delivery/adding-http-security-headers-using-lambdaedge>

#### NEW QUESTION 49

An ecommerce company uses a large number of Amazon EBS backed Amazon EC2 instances. To decrease manual work across all the instances, a DevOps engineer is tasked with automating restart actions when EC2 instance retirement events are scheduled.

How can this be accomplished?

- A. Create a scheduled Amazon CloudWatch Events rule to execute an AWS Systems Manager automation document that checks if any EC2 instances are scheduled for retirement once a week
- B. If the instance is scheduled for retirement, the automation document will hibernate the instance.
- C. Enable EC2 Auto Recovery on all of the instance
- D. Create an AWS Config rule to limit the recovery to occur during a maintenance window only.
- E. Reboot all EC2 instances during an approved maintenance window that is outside of standard business hour
- F. Set up Amazon CloudWatch alarms to send a notification in case any instance is failing EC2 instance status checks.
- G. Set up an AWS Health Amazon CloudWatch Events rule to execute AWS Systems Manager automation documents that stop and start the EC2 instance when a retirement scheduled event occurs.

**Answer:** D

#### NEW QUESTION 54

A DevOps Engineer must implement monitoring for a workload running on Amazon EC2 and Amazon RDS MySQL. The monitoring must include:

Application logs and operating system metrics for the Amazon EC2 instances Database logs and operating system metrics for the Amazon RDS database Which steps should the Engineer take?

- A. Install an Amazon CloudWatch agent on the EC2 and RDS instance
- B. Configure the agent to send the operating system metrics and application and database logs to CloudWatch.
- C. Install an Amazon CloudWatch agent on the EC2 instance, and configure the agent to send the application logs and operating system metrics to CloudWatch
- D. Enable RDS Enhanced Monitoring, and modify the RDS instance to publish database logs to CloudWatch Logs.
- E. Install an Amazon CloudWatch Logs agent on the EC2 instance and configure it to send application logs to CloudWatch.
- F. Set up scheduled tasks on the EC2 and RDS instances to put operating system metrics and application and database logs into an Amazon S3 bucket
- G. Set up an event on the bucket to invoke an AWS Lambda function to monitor for errors each time an object is put into the bucket.

**Answer:** B

#### NEW QUESTION 57

An application's users are encountering bugs immediately after Amazon API Gateway deployments. The development team deploys once or twice a day and uses a blue/green deployment strategy with custom health checks and automated rollbacks. The team wants to limit the number of users affected by deployment bugs and receive notifications when rollbacks are needed.

Which combination of steps should a DevOps engineer use to meet these requests? (Select TWO.)

- A. Implement a blue/green strategy using path mappings.
- B. Implement a canary deployment strategy.
- C. Implement a rolling deployment strategy using multiple stages.
- D. Use Amazon CloudWatch alarms to notify the development team.
- E. Use Amazon CloudWatch Events to notify the development team.

**Answer:** BD

#### NEW QUESTION 62

A Development team is adding a new country to an e-commerce application. This addition requires that new application features be added to the shipping component of the application. The team has not decided if all new features should be added, as some will take approximately six weeks to build. While the final decision on the shipping component features is being made, other team members are continuing to work on other features of the application.

Based on this situation, how should the application feature deployments be managed?

- A. Add the code updates as commits to the release branch
- B. The team can delay the deployment until all features are ready.
- C. Add the code updates as commits to a feature branch
- D. Merge the commits to a release branch as features are ready.
- E. Add the code updates as a single commit when a feature is ready
- F. Tag this commit with "new-country."
- G. Create a new repository named "new-country". Commit all the code changes to the new repository.

**Answer:** B

#### NEW QUESTION 64

A DevOps Engineer is developing a deployment strategy that will allow for data-driven decisions before a feature is fully approved for general availability. The current deployment process uses AWS CloudFormation and blue/green-style deployments. The development team has decided that customers should be randomly assigned to groups, rather than using a set percentage, and redirects should be avoided.

What process should be followed to implement the new deployment strategy?

- A. Configure Amazon Route 53 weighted records for the blue and green stacks, with 50% of traffic configured to route to each stack.
- B. Configure Amazon CloudFront with an AWS Lambda@Edge function to set a cookie when CloudFront receives a request
- C. Assign the user to a version A or B, and configure the web server to redirect to version A or B.

- D. Configure Amazon CloudFront with an AWS Lambda@Edge function to set a cookie when CloudFront receives a request
- E. Assign the user to a version A or B, then return the corresponding version to the viewer.
- F. Configure Amazon Route 53 with an AWS Lambda function to set a cookie when Amazon CloudFront receives a request
- G. Assign the user to version A or B, then return the corresponding version to the viewer.

**Answer: C**

**Explanation:**

[https://docs.aws.amazon.com/zh\\_cn/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html](https://docs.aws.amazon.com/zh_cn/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html)

**NEW QUESTION 65**

A DevOps Engineer has several legacy applications that all generate different log formats. The Engineer must standardize the formats before writing them to Amazon S3 for querying and analysis.

How can this requirement be met at the LOWEST cost?

- A. Have the application send its logs to an Amazon EMR cluster and normalize the logs before sending them to Amazon S3
- B. Have the application send its logs to Amazon QuickSight then use the Amazon QuickSight SPICE engine to normalize the logs Do the analysis directly from Amazon QuickSight.
- C. Keep the logs in Amazon S3 and use Amazon Redshift Spectrum to normalize the logs in place
- D. Use Amazon Kinesis Agent on each server to upload the logs and have Amazon Kinesis Data Firehose use an AWS Lambda function to normalize the logs before writing them to Amazon S3

**Answer: D**

**NEW QUESTION 67**

An e-commerce company is running a web application in an AWS Elastic Beanstalk environment. In recent months, the average load of the Amazon EC2 instances has been increased to handle more traffic.

The company would like to improve the scalability and resilience of the environment. The Development team has been asked to decouple long-running tasks from the environment if the tasks can be executed asynchronously. Examples of these tasks include confirmation emails when users are registered to the platform, and processing images or videos. Also, some of the periodic tasks that are currently running within the web server should be offloaded.

What is the most time-efficient and integrated way to achieve this?

- A. Create an Amazon SQS queue and send the tasks that should be decoupled from the Elastic Beanstalk web server environment to the SQS queue
- B. Create a fleet of EC2 instances under an Auto Scaling group
- C. Use an AMI that contains the application to process the asynchronous tasks, configure the application to listen for messages within the SQS queue, and create periodic tasks by placing those into the cron in the operating system
- D. Create an environment variable within the Elastic Beanstalk environment with a value pointing to the SQS queue endpoint.
- E. Create a second Elastic Beanstalk worker tier environment and deploy the application to process the asynchronous tasks there
- F. Send the tasks that should be decoupled from the original Elastic Beanstalk web server environment to the auto-generated Amazon SQS queue by the Elastic Beanstalk worker environment
- G. Place a cron.yaml file within the root of the application source bundle for the worker environment periodic task
- H. Use environment links to link the web server environment with the worker environment.
- I. Create a second Elastic Beanstalk web server tier environment and deploy the application to process the asynchronous task
- J. Send the tasks that should be decoupled from the original Elastic Beanstalk web server to the auto-generated Amazon SQS queue by the Elastic Beanstalk web server tier environment
- K. Place a cron.yaml file within the root of the application source bundle for the second web server tier environment with the necessary periodic task
- L. Use environment links to link both web server environments.
- M. Create an Amazon SQS queue and send the tasks that should be decoupled from the Elastic Beanstalk web server environment to the SQS queue
- N. Create a fleet of EC2 instances under an Auto Scaling group
- O. Install and configure the application to listen for messages within the SQS queue from UserData and create periodic tasks by placing those into the cron in the operating system
- P. Create an environment variable within the Elastic Beanstalk web server environment with a value pointing to the SQS queue endpoint.

**Answer: C**

**NEW QUESTION 69**

An IT department manages a portfolio with Windows and Linux (Amazon and Red Hat Enterprise Linux) servers both on-premises and on AWS. An audit reveals that there is no process for updating OS and core application patches, and that the servers have inconsistent patch levels.

Which of the following provides the MOST reliable and consistent mechanism for updating and maintaining all servers at the recent OS and core application patch levels?

- A. Install AWS Systems Manager agent on all on-premises and AWS server
- B. Create Systems Manager Resource Group
- C. Use Systems Manager Patch Manager with a preconfigured patch baseline to run scheduled patch updates during maintenance windows.
- D. Install the AWS OpsWorks agent on all on-premises and AWS server
- E. Create an OpsWorks stack with separate layers for each operating system, and get a recipe from the Chef supermarket to run the patch commands for each layer during maintenance windows.
- F. Use a shell script to install the latest OS patches on the Linux servers using yum and schedule it to run automatically using cron
- G. Use Windows Update to automatically patch Windows servers.
- H. Use AWS Systems Manager Parameter Store to securely store credentials for each Linux and Windows server
- I. Create Systems Manager Resource Group
- J. Use the Systems Manager Run Command to remotely deploy patch updates using the credentials in Systems Manager Parameter Store

**Answer: A**

**Explanation:**

1- <https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-patch-patchgroups.html> 2- <https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-patch.html>

#### NEW QUESTION 74

A DevOps engineer has automated a web service deployment using AWS CodePipeline with the following steps:

- An AWS CodeBuild project compiles the deployment artifact and runs unit tests.
- An AWS CodeDeploy deployment group deploys the web service to Amazon EC2 instances in the staging environment.
- A CodeDeploy deployment group deploys the web service to EC2 instances in the production environment. The quality assurance (QA) team has asked for permission to inspect the build artifact before the deployment to the production environment occurs. The QA team wants to run an internal automated penetration testing tool (invoked using a REST API call) to run some manual tests.

Which combination of actions will fulfill this request? (Select TWO.)

- A. Insert a manual approval action between the test and deployment actions of the pipeline.
- B. Modify the buildspec.yml file for the compilation stage to require manual approval before completion.
- C. Update the CodeDeploy deployment group so it requires manual approval to proceed.
- D. Update the pipeline to directly trigger the REST API for the automated penetration testing tool.
- E. Update the pipeline to invoke a Lambda function that triggers the REST API for the automated penetration testing tool.

**Answer: BD**

#### NEW QUESTION 79

A company's legacy application uses IAM user credentials to access resources in the company's AWS Organizations organization. A DevOps engineer needs to ensure new IAM users cannot be created unless the employee creating the IAM user is on an exception list.

Which solution will meet these requirements?

- A. Attach an Organizations SCP with an explicit deny for all iam:CreateAccessKey actions with a condition that excludes StringNotEquals for aws:username with a value of the exception list.
- B. Attach an Organizations SCP with an explicit deny for all iam:CreateUser actions with a condition that includes StringEquals for aws:username with a value of the exception list.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with a pattern that matches the iam:CreateAccessKey action with an AWS Lambda function target.
- D. The function will check the user name against an exception list.
- E. If the user is not in the exception list, the function will delete the user.
- F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with a pattern that matches the iam:CreateUser action with an AWS Lambda function target.
- G. The function will check the user name and account against an exception list.
- H. If the user is not in the exception list, the function will delete the user.

**Answer: B**

#### NEW QUESTION 81

Which Auto Scaling process would be helpful when testing new instances before sending traffic to them, while still keeping them in your Auto Scaling Group?

- A. Suspend the process AZ Rebalance
- B. Suspend the process Health Check
- C. Suspend the process Replace Unhealthy
- D. Suspend the process AddToLoadBalancer

**Answer: D**

#### Explanation:

If you suspend AddToLoadBalancer, Auto Scaling launches the instances but does not add them to the load balancer or target group. If you resume the AddToLoadBalancer process, Auto Scaling resumes adding instances to the load balancer or target group when they are launched. However, Auto Scaling does not add the instances that were launched while this process was suspended. You must register those instances manually.

Option A is invalid because this just balances the number of EC2 instances in the group across the Availability Zones in the region.

Option B is invalid because this just checks the health of the instances. Auto Scaling marks an instance as unhealthy if Amazon EC2 or Elastic Load Balancing tells

Auto Scaling that the instance is unhealthy.

Option C is invalid because this process just terminates instances that are marked as unhealthy and later creates new instances to replace them.

For more information on process suspension, please refer to the below document link: from AWS

> <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html>

#### NEW QUESTION 83

A development team wants to deploy an application using AWS CloudFormation stacks, but the developer IAM role does not currently have the required permissions to provision the resources specified in the CloudFormation template. A DevOps engineer is tasked with allowing developers to deploy the stacks while following the principle of least privilege.

Which solution will meet these requirements?

- A. Create an IAM policy that allows developers to provision the required resource.
- B. Attach the policy to the developer role.
- C. Create an IAM policy that allows full access to CloudFormation.
- D. Attach the policy to the developer role.
- E. Create a new IAM role with the required permissions to use as a CloudFormation service role.
- F. Grant the developer role a cloudformation:\* action.
- G. Create a new IAM role with the required permissions to use as a CloudFormation service role.
- H. Grant the developer role the iam:PassRole permission.

**Answer: C**

#### Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-service-role.html>

#### NEW QUESTION 85

A company wants to use AWS Systems Manager documents to bootstrap physical laptops for developers. The bootstrap code is stored in GitHub. A DevOps engineer has already created a Systems Manager activation, installed the Systems Manager agent with the registration code, and installed an activation ID on all the laptops.

Which set of steps should be taken next?

- A. Configure the Systems Manager document to use the AWS-RunShellScript command to copy the files from GitHub to Amazon S3, then use the aws-downloadContent plugin with a source Type of S3.
- B. Configure the Systems Manager document to use the aws-configurePackage plugin with an install action and point to the Git repository.
- C. Configure the Systems Manager document to use the aws-downloadContent plugin with a sourceType of GitHub and sourceInfo with the repository details.
- D. Configure the Systems Manager document to use the aws:softwareInventory plugin and run the script from the Git repository.

**Answer: D**

#### NEW QUESTION 89

A company mandates the creation of capture logs for everything running in its AWS account. The account has multiple VPCs with Amazon EC2 instances, Application Load Balancers, Amazon RDS MySQL databases, and AWS WAF rules configured. The logs must be protected from deletion. A daily visual analysis of log anomalies from the previous day is required.

Which combination of actions should a DevOps Engineer take to accomplish this? (Choose three.)

- A. Configure an AWS Lambda function to send all CloudWatch logs to an Amazon S3 bucket
- B. Create a dashboard report in Amazon QuickSight.
- C. Configure AWS CloudTrail to send all logs to Amazon Inspector
- D. Create a dashboard report in Amazon QuickSight.
- E. Configure Amazon S3 MFA Delete on the logging Amazon S3 bucket.
- F. Configure an Amazon S3 object lock legal hold on the logging Amazon S3 bucket.
- G. Configure AWS Artifact to send all logs to the logging Amazon S3 bucket
- H. Create a dashboard report in Amazon QuickSight.
- I. Deploy an Amazon CloudWatch agent to all Amazon EC2 instances.

**Answer: ADF**

#### NEW QUESTION 90

The management team at a company with a large on-premises OpenStack environment wants to move non-production workloads to AWS. An AWS Direct Connect connection has been provisioned and configured to connect the environments. Due to contractual obligations, the production workloads must remain on-premises, and will be moved to AWS after the next contract negotiation. The company follows Center for Internet Security (CIS) standards for hardening images; this configuration was developed using the company's configuration management system.

Which solution will automatically create an identical image in the AWS environment without significant overhead?

- A. Write an AWS CloudFormation template that will create an Amazon EC2 instance
- B. Use cloud-init to install the configuration management agent, use cfn-wait to wait for configuration management to successfully apply, and use an AWS Lambda-backed custom resource to create the AMI.
- C. Log in to the console, launch an Amazon EC2 instance, and install the configuration management agent. When changes are applied through the configuration management system, log in to the console and create a new AMI from the instance.
- D. Create a new AWS OpsWorks layer and mirror the image hardening standard
- E. Use this layer as the baseline for all AWS workloads.
- F. When a change is made in the configuration management system, a job in Jenkins is triggered to use the VM Import command to create an Amazon EC2 instance in the Amazon VPC
- G. Use lifecycle hooks to launch an AWS Lambda function to create the AMI.

**Answer: D**

#### Explanation:

<https://www.brad-x.com/2015/10/01/importing-an-openstack-vm-into-amazon-ec2/> <https://aws.amazon.com/ec2/vm-import/>

#### NEW QUESTION 93

A company wants to use a grid system for a proprietary enterprise in-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
- I. Create a crontab script that will poll for that S3 file and download it frequently
- J. Use a process manager, such as Monit or systemd, to restart the cluster services when it detects that the new file was modified
- K. When adding a node to the cluster, edit the file's most recent member
- L. Upload the new file to the S3 bucket.
- M. 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 file whenever a new instance is added to the cluster

**Answer: A**

#### Explanation:

<https://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-events.html>

#### NEW QUESTION 96

A startup company is developing a web application on AWS. It plans to use Amazon RDS for persistence and deploy the application to Amazon EC2 with an Auto Scaling group. The company would also like to separate the environments for development, testing, and production. What is the MOST secure and flexible approach to manage the application configuration?

- A. Create a property file to include the configuration and the encrypted password
- B. Check in the property file to the source repository, package the property file with the application, and deploy the application
- C. Create an environment tag for the EC2 instances and tag the instances respectively
- D. The application will extract the necessary property values based on the environment tag.
- E. Create a property file for each environment to include the environment-specific configuration and an encrypted password
- F. Check in the property files to the source repository
- G. During deployment, use only the environment-specific property file with the application
- H. The application will read the needed property values from the deployed property file.
- I. Create a property file for each environment to include the environment-specific configuration
- J. Create a private Amazon S3 bucket and save the property files in the bucket
- K. Save the passwords in the bucket with AWS KMS encryption
- L. During deployment, the application will read the needed property values from the environment-specific property file in the S3 bucket.
- M. Create a property file for each environment to include the environment-specific configuration
- N. Create a private Amazon S3 bucket and save the property files in the bucket
- O. Save the encrypted passwords in the AWS Systems Manager Parameter Store
- P. Create an environment tag for the EC2 instances and tag the instances respectively
- Q. The application will read the needed property values from the environment-specific property file in the S3 bucket and the parameter store.

**Answer: D**

#### NEW QUESTION 100

During metric analysis, your team has determined that the company's website during peak hours is experiencing response times higher than anticipated. You currently rely on Auto Scaling to make sure that you are scaling your environment during peak windows. How can you improve your Auto Scaling policy to reduce this high response time? Choose 2 answers.

- A. Push custom metrics to CloudWatch to monitor your CPU and network bandwidth from your servers, which will allow your Auto Scaling policy to have better fine-grain insight.
- B. Increase your Auto Scaling group's number of max servers.
- C. Create a script that runs and monitors your servers; when it detects an anomaly in load, it posts to an Amazon SNS topic that triggers Elastic Load Balancing to add more servers to the load balancer.
- D. Push custom metrics to CloudWatch for your application that include more detailed information about your web application, such as how many requests it is handling and how many are waiting to be processed.

**Answer: BD**

#### Explanation:

Option B makes sense because maybe the max servers is low hence the application cannot handle the peak load.

Option D helps in ensuring Autoscaling can scale the group on the right metrics.

For more information on Autoscaling health checks, please refer to the below document link: from AWS

> <http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html>

#### NEW QUESTION 105

A company wants to automatically re-create its infrastructure using AWS CloudFormation as part of the company's quality assurance (QA) pipeline. For each QA run, a new VPC must be created in a single account, resources must be deployed into the VPC, and tests must be run against this new infrastructure. The company policy states that all VPCs must be peered with a central management VPC to allow centralized logging. The company has existing CloudFormation templates to deploy its VPC and associated resources.

Which combination of steps will achieve the goal in a way that is automated and repeatable? (Choose two.)

- A. Create an AWS Lambda function that is invoked by an Amazon CloudWatch Events rule when a CreateVpcPeeringConnection API call is made
- B. The Lambda function should check the source of the peering request, accept the request, and update the route tables for the management VPC to allow traffic to go over the peering connection.
- C. In the CloudFormation template: Invoke a custom resource to generate unique VPC CIDR ranges for the VPC and subnets. Create a peering connection to the management VPC. Update route tables to allow traffic to the management VPC.
- D. In the CloudFormation template: Use the Fn::Cidr function to allocate an unused CIDR range for the VPC and subnets. Create a peering connection to the management VPC. Update route tables to allow traffic to the management VPC.
- E. Modify the CloudFormation template to include a mappings object that includes a list of /16 CIDR ranges for each account where the stack will be deployed.
- F. Use CloudFormation StackSets to deploy the VPC and associated resources to multiple AWS accounts using a custom resource to allocate unique CIDR range
- G. Create peering connections from each VPC to the central management VPC and accept those connections in the management VPC.

**Answer: BD**

#### NEW QUESTION 108

You have decided that you need to change the instance type of your production instances which are running as part of an AutoScaling group. The entire architecture is deployed using CloudFormation Template. You currently have 4 instances in Production. You cannot have any interruption in service and need to ensure 2 instances are always running during the update? Which of the options below listed can be used for this?

- A. AutoScalingRollingUpdate
- B. AutoScalingScheduledAction
- C. AutoScalingReplacingUpdate
- D. AutoScalingIntegrationUpdate

**Answer: A**

**Explanation:**

The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto Scaling group resource is updated when an update to the Cloud Formation stack occurs. A common approach to updating an Auto Scaling group is to perform a rolling update, which is done by specifying the AutoScalingRollingUpdate policy. This retains the same Auto Scaling group and replaces old instances with new ones, according to the parameters specified. For more information on Autoscaling updates, please refer to the below link:

➤ <https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/>

**NEW QUESTION 111**

A DevOps Engineer is using AWS CodeDeploy across a fleet of Amazon EC2 instances in an EC2 Auto Scaling group. The associated CodeDeploy deployment group, which is integrated with EC2 Auto Scaling, is configured to perform in-place deployments with CodeDeployDefault.OneAtATime. During an ongoing new deployment, the Engineer discovers that, although the overall deployment finished successfully, two out of five instances have the previous application revision deployed. The other three instances have the newest application revision.

What is likely causing this issue?

- A. The two affected instances failed to fetch the new deployment.
- B. A failed AfterInstall lifecycle event hook caused the CodeDeploy agent to roll back to the previous version on the affected instances.
- C. The CodeDeploy agent was not installed in two affected instances.
- D. EC2 Auto Scaling launched two new instances while the new deployment had not yet finished, causing the previous version to be deployed on the affected instances.

**Answer: D**

**NEW QUESTION 116**

A company develops and maintains a web application using Amazon EC2 instances and an Amazon RDS for SQL Server DB instance in a single Availability Zone. The resources need to run only when new deployments are being tested using AWS CodePipeline. Testing occurs one or more times a week and each test takes 2-3 hours to run. A DevOps engineer wants a solution that does not change the architecture components.

Which solution will meet these requirements in the MOST cost-effective manner?

- A. Convert the RDS database to an Amazon Aurora Serverless database. Use an AWS Lambda function to start and stop the EC2 instances before and after tests.
- B. Put the EC2 instances into an Auto Scaling group.
- C. Schedule scaling to run at the start of the deployment tests.
- D. Replace the EC2 instances with EC2 Spot Instances and the RDS database with an RDS Reserved Instance.
- E. Subscribe Amazon CloudWatch Events to CodePipeline to trigger AWS Systems Manager Automation documents that start and stop all EC2 and RDS instances before and after deployment tests.

**Answer: A**

**NEW QUESTION 120**

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

Which solution will accomplish this?

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

**Answer: C**

**NEW QUESTION 121**

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? (Select TWO.)

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

**Answer: AD**

**NEW QUESTION 123**

A company is using Amazon EC2 for various workloads. Company policy requires that instances be managed centrally to standardize configurations. These configurations include standard logging, metrics, security assessments, and weekly patching.

How can the company meet these requirements? (Select THREE.)

- A. Use AWS Config to ensure all EC2 instances are managed by Amazon Inspector.
- B. Use AWS Config to ensure all EC2 instances are managed by AWS Systems Manager.
- C. Use AWS Systems Manager to install and manage Amazon Inspector, Systems Manager Patch Manager, and the Amazon CloudWatch agent on all instances.

- D. Use Amazon Inspector to install and manage AWS Systems Manager, Systems Manager Patch Manager, and the Amazon CloudWatch agent on all instances.
- E. Use AWS Systems Manager maintenance windows with Systems Manager Run Command to schedule Systems Manager Patch Manager task
- F. Use the Amazon CloudWatch agent to schedule Amazon Inspector assessment runs.
- G. Use AWS Systems Manager maintenance windows with Systems Manager Run Command to schedule Systems Manager Patch Manager task
- H. Use Amazon CloudWatch Events to schedule Amazon Inspector assessment runs.

**Answer:** BDE

#### NEW QUESTION 128

A DevOps Engineer is reviewing a system that uses Amazon EC2 instances in an Auto Scaling group. This system uses a configuration management tool that runs locally on each EC2 instance. Because of the volatility of the application load, new instances must be fully functional within 3 minutes of entering a running state. Current setup tasks include:

- \* Installing the configuration management agent – 2 minutes
- \* Installing the application framework – 15 minutes
- \* Copying configuration data from Amazon S3 – 2 minutes
- \* Running the configuration management agent to configure instances – 1 minute
- \* Deploying the application code from Amazon S3 – 2 minutes

How should the Engineer set up system so it meets the launch time requirement?

- A. Trigger an AWS Lambda function from an Amazon CloudWatch Events rule when a new EC2 instance launches
- B. Have the function install the configuration management agent and the application framework, pull configuration data from Amazon S3, run the agent to configure the instance, and deploy the application from S3.
- C. Write a bootstrap script to install the configuration management agent, install and the application framework, pull configuration data from Amazon S3, run the agent to configure the instance, and deploy the application from S3.
- D. Build a custom AMI that includes the configuration management agent and application framework. Write a bootstrap script to pull configuration data from Amazon S3, run the agent to configure the instance, and deploy the application from S3.
- E. Build a custom AMI that includes the configuration management agent, application framework, and configuration data
- F. Write a bootstrap script to run the agent to configure the instance and deploy the application from Amazon S3.

**Answer:** D

#### NEW QUESTION 132

A software company wants to automate the build process for a project where the code is stored in GitHub. When the repository is updated, source code should be compiled, tested, and pushed to Amazon S3.

Which combination of steps would address these requirements? (Select THREE.)

- A. Add a buildspec.yml file to the source code with build instructions.
- B. Configure a GitHub webhook to trigger a build every time a code change is pushed to the repository.
- C. Create an AWS CodeBuild project with GitHub as the source repository.
- D. Create an AWS CodeDeploy application with the Amazon EC2/On-Premises compute platform.
- E. Create an AWS OpsWorks deployment with the install dependencies command.
- F. Provision an Amazon EC2 instance to perform the build.

**Answer:** ACD

#### NEW QUESTION 135

A DevOps Engineer uses Docker container technology to build an image-analysis application. The application often sees spikes in traffic. The Engineer must automatically scale the application in response to customer demand while maintaining cost effectiveness and minimizing any impact on availability.

What will allow the FASTEST response to spikes in traffic while fulfilling the other requirements?

- A. Create an Amazon ECS cluster with the container instances in an Auto Scaling group
- B. Configure the ECS service to use Service Auto Scaling
- C. Set up Amazon CloudWatch alarms to scale the ECS service and cluster.
- D. Deploy containers on an AWS Elastic Beanstalk Multicontainer Docker environment
- E. Configure Elastic Beanstalk to automatically scale the environment based on Amazon CloudWatch metrics.
- F. Create an Amazon ECS cluster using Spot instance
- G. Configure the ECS service to use Service Auto Scaling
- H. Set up Amazon CloudWatch alarms to scale the ECS service and cluster.
- I. Deploy containers on Amazon EC2 instance
- J. Deploy a container scheduler to schedule containers onto EC2 instance
- K. Configure EC2 Auto Scaling for EC2 instances based on available Amazon CloudWatch metrics.

**Answer:** A

#### Explanation:

<https://aws.amazon.com/blogs/compute/automatic-scaling-with-amazon-ecs/>

#### NEW QUESTION 140

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 CloudWatch Events rule to ensure the pipeline execution 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 symmetric multiprocessing configuration to run the builds in parallel.
- C. Modify the CodePipeline configuration to execute 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://aws.amazon.com/blogs/devops/using-aws-codepipeline-aws-codebuild-and-aws-lambda-for-serverless-au>  
<https://docs.aws.amazon.com/codepipeline/latest/userguide/reference-pipeline-structure.html>

**NEW QUESTION 141**

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

- A. Configure a latency-based Amazon Route 53 CNAME with health checks so it points to both the primary and replica endpoints Subscribe an Amazon SNS topic to Amazon RDS failure notifications from AWS CloudTrail and use that topic to trigger an AWS Lambda function that will promote the replica instance as the master.
- B. Create an Aurora custom endpoint to point to the primary database instance Configure the application to use this endpoint Configure AWS CloudTrail to run an AWS Lambda function to promote the replica instance and modify the custom endpoint to point to the newly promoted instance.
- C. Create an AWS Lambda function to modify the application's AWS CloudFormation template to promote the replica, apply the template to update the stack, and point the application to the newly promoted instance Create an Amazon CloudWatch alarm to trigger this Lambda function after the failure event occurs
- D. Store the Aurora endpoint in AWS Systems Manager Parameter Store Create an Amazon EventBridge (Amazon CloudWatch Events) event that detects the database failure and runs an AWS Lambda function to promote the replica instance and update the endpoint URL stored in AWS Systems Manager Parameter Store Code the application to reload the endpoint from Parameter Store if a database connection fails.

**Answer: A**

**NEW QUESTION 146**

After presenting a working proof of concept for a new application that uses AWS API Gateway, a Developer must set up a team development environment for the project. Due to a tight timeline, the Developer wants to minimize time spent on infrastructure setup, and would like to reuse the code repository created for the proof of concept. Currently, all source code is stored in AWS CodeCommit.  
 Company policy mandates having alpha, beta, and production stages with separate Jenkins servers to build code and run tests for every stage. The Development Manager must have the ability to block code propagation between admins at any time. The Security team wants to make sure that users will not be able to modify the environment without permission.  
 How can this be accomplished?

- A. Create API Gateway alpha, beta, and production stage
- B. Create a CodeCommit trigger to deploy code to the different stages using an AWS Lambda function.
- C. Create API Gateway alpha, beta, and production stage
- D. Create an AWS CodePipeline that pulls code from the CodeCommit repository
- E. Create CodePipeline actions to deploy code to the API Gateway stages.
- F. Create Jenkins servers for the alpha, beta, and production stages on Amazon EC2 instance
- G. Create multiple CodeCommit triggers to deploy code to different stages using an AWS Lambda function.
- H. Create an AWS CodePipeline pipeline that pulls code from the CodeCommit repository
- I. Create alpha, beta, and production stages with Jenkins servers on CodePipeline.

**Answer: D**

**NEW QUESTION 148**

A company runs a database on a single Amazon EC2 instance in a development environment. The data is stored on separate Amazon EBS volumes that are attached to the EC2 instance. An Amazon Route 53 A record has been created and configured to point to the EC2 instance. The company would like to automate the recovery of the database instance when an instance or Availability Zone (AZ) fails. The company also wants to keep its costs low. The RTO is 4 hours and RPO is 12 hours.  
 Which solution should a DevOps Engineer implement to meet these requirements?

- A. Run the database in an Auto Scaling group with a minimum and maximum instance count of 1 in multiple AZ
- B. Add a lifecycle hook to the Auto Scaling group and define an Amazon CloudWatch Events rule that is triggered when a lifecycle event occur
- C. Have the CloudWatch Events rule invoke an AWS Lambda function to detach or attach the Amazon EBS data volumes from the EC2 instance based on the event
- D. Configure the EC2 instance UserData to mount the data volumes (retry on failure with a short delay), then start the database and update the Route 53 record.
- E. Run the database on two separate EC2 instances in different AZs with one active and the other as a standby
- F. Attach the data volumes to the active instance
- G. Configure an Amazon CloudWatch Events rule to invoke an AWS Lambda function on EC2 instance termination
- H. The Lambda function launches a replacement EC2 instance
- I. If the terminated instance was the active node, then the function attaches the data volumes to the standby node
- J. Start the database and update the Route 53 record.
- K. Run the database in an Auto Scaling group with a minimum and maximum instance count of 1 in multiple AZ
- L. Create an AWS Lambda function that is triggered by a scheduled Amazon CloudWatch Events rule every 4 hours to take a snapshot of the data volume and apply it to a standby
- M. Have the instance UserData get the latest snapshot, create a new volume from it, and attach and mount the volume
- N. Then start the database and update the Route 53 record.
- O. Run the database on two separate EC2 instances in different AZ
- P. Configure one of the instances as a master and the other as a standby
- Q. Set up replication between the master and standby instance
- R. Point the Route 53 record to the master
- S. Configure an Amazon CloudWatch Events rule to invoke an AWS Lambda function upon the EC2 instance termination
- T. The Lambda function launches a replacement EC2 instance
- . If the terminated instance was the active node, the function promotes the standby to master and points the Route 53 record to it.

**Answer: D**

#### NEW QUESTION 149

A company is adopting AWS CodeDeploy to automate its application deployments for a Java-Apache Tomcat application with an Apache webserver. The Development team started with a proof of concept, created a deployment group for a developer environment, and performed functional tests within the application. After completion, the team will create additional deployment groups for staging and production. The current log level is configured within the Apache settings, but the team wants to change this configuration dynamically when the deployment occurs, so that they can set different log level configurations depending on the deployment group without having a different application revision for each group. How can these requirements be met with the LEAST management overhead and without requiring different script versions for each deployment group?

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 151

A company is implementing AWS CodePipeline to automate its testing process. The company wants to be notified when the execution state fails and used the following custom event pattern in Amazon CloudWatch:

```
{
  "source": [
    "aws.codepipeline"
  ],
  "detail-type": [
    "CodePipeline Action Execution State Change"
  ],
  "detail": {
    "state": [
      "FAILED"
    ]
  },
  "type": {
    "category": ["Approval"]
  }
}
```

Which type of events will match this event pattern?

- A. Failed deploy and build actions across all the pipelines.
- B. All rejected or failed approval actions across all the pipelines.
- C. All the events across all pipelines.
- D. Approval actions across all the pipelines.

**Answer:** B

#### Explanation:

<https://docs.aws.amazon.com/codepipeline/latest/userguide/detect-state-changes-cloudwatch-events.html>

#### NEW QUESTION 155

A company needs to introduce automatic DNS failover for a distributed web application to a disaster recovery or standby installation. The DevOps Engineer plans to configure Amazon Route 53 to provide DNS routing to alternate endpoint in the event of an application failure.

What steps should the Engineer take to accomplish this? (Select TWO.)

- A. Create Amazon Route 53 health checks for each endpoint that cannot be entered as alias record
- B. Ensure firewall and routing rules allow Amazon Route 53 to send requests to the endpoints that are specified in the health checks.
- C. Create alias records that route traffic to AWS resources and set the value of the Evaluate Target Health option to Yes, then create all the non-alias records.
- D. Create a governing Amazon Route 53 record set, set it to failover, and associate it with the primary and secondary Amazon Route 53 record sets to distribute

traffic to healthy DNS entries.

E. Create an Amazon CloudWatch alarm to monitor the primary Amazon Route 53 DNS entr

F. Then create an associated AWS Lambda function to execute the failover API call to Route 53 to the secondary DNS entry.

G. Map the primary and secondary Amazon Route 53 record sets to an Amazon CloudFront distribution using primary and secondary origins.

**Answer:** AC

#### NEW QUESTION 160

A web application has been deployed using an AWS Elastic Beanstalk application. The Application Developers are concerned that they are seeing high latency in two different areas of the application: HTTP client requests to a third-party API, MySQL client library queries to an Amazon RDS database. A DevOps Engineer must gather trace data to diagnose the issues. Which steps will gather the trace information with the LEAST amount of changes and performance impacts to the application?

A. Add additional logging to the application code.

B. Use the Amazon CloudWatch agent to stream the application logs into Amazon Elasticsearch Service.

C. Query the log data in Amazon ES.

D. Instrument the application to use the AWS X-Ray SDK.

E. Post trace data to an Amazon Elasticsearch Service cluster.

F. Query the trace data for calls to the HTTP client and the MySQL client.

G. On the AWS Elastic Beanstalk management page for the application, enable the AWS X-Ray daemon. View the trace data in the X-Ray console.

H. Instrument the application using the AWS X-Ray SDK.

I. On the AWS Elastic Beanstalk management page for the application, enable the X-Ray daemon.

J. View the trace data in the X-Ray console.

**Answer:** C

#### NEW QUESTION 161

A company runs several applications across multiple AWS accounts in an organization in AWS Organizations. Some of the resources are not tagged properly, and the company's finance team cannot determine which costs are associated with which applications. A DevOps engineer must remediate this issue and prevent this issue from happening in the future.

Which combination of actions should the DevOps engineer take to meet these requirements? (Select TWO.)

A. Activate the user-defined cost allocation tags in each AWS account.

B. Create and attach an SCP that requires a specific tag.

C. Define each line of business (LOB) in AWS Budget.

D. Assign the required tag to each resource.

E. Scan all accounts with Tag Editor.

F. Assign the required tag to each resource.

G. Use the budget report to find untagged resources.

H. Assign the required tag to each resource.

**Answer:** CD

#### NEW QUESTION 163

An Application team is refactoring one of its internal tools to run in AWS instead of on-premises hardware. All of the code is currently written in Python and is standalone. There is also no external state store or relational database to be queried.

Which deployment pipeline incurs the LEAST amount of changes between development and production?

A. Developers should use Docker for local development.

B. Use AWS SSM to import these containers as AMIs for Amazon EC2 whenever dependencies are updated.

C. Use AWS CodePipeline to test new code changes against the Auto Scaling group.

D. Developers should use their native Python environment.

E. When dependencies are changed and a new container is ready, use AWS CodePipeline and AWS CodeBuild to perform functional tests and then upload the new container to the Amazon EC2 instance.

F. Use AWS CloudFormation with the custom container to deploy the new Amazon ECS.

G. Developers should use their native Python environment.

H. When dependencies are changed and a new code is ready, use AWS CodePipeline and AWS CodeBuild to perform functional tests and then upload the new container to the Amazon EC2 instance.

I. Use CodePipeline and CodeBuild with the custom container to test new code changes inside AWS Elastic Beanstalk.

**Answer:** A

#### NEW QUESTION 165

A company hosts its staging website using an Amazon EC2 instance backed with Amazon EBS storage. The company wants to recover quickly with minimal data losses in the event of network connectivity issues or power failures on the EC2 instance.

Which solution will meet these requirements?

A. Add the instance to an EC2 Auto Scaling group with the minimum, maximum, and desired capacity set to 1.

B. Add the instance to an EC2 Auto Scaling group with a lifecycle hook to detach the EBS volume when the EC2 instance shuts down or terminates.

C. Create an Amazon CloudWatch alarm for the StatusCheckFailed\_System metric and select the EC2 action to recover the instance.

D. Create an Amazon CloudWatch alarm for the StatusCheckFailed\_Instance metric and select the EC2 action to reboot the instance.

**Answer:** A

#### NEW QUESTION 166

A company has microservices running in AWS Lambda that read data from Amazon DynamoDB. The Lambda code is manually deployed by Developers after successful testing. The company now needs the tests and deployments to be automated and run in the cloud. Additionally, traffic to the new versions of each microservice should be incrementally shifted over time after deployment.

What solution meets all the requirements, ensuring the MOST developer velocity?

- A. Create an AWS CodePipeline configuration and set up a post-commit hook to trigger the pipeline after tests have passed
- B. Use AWS CodeDeploy and create a Canary deployment configuration that specifies the percentage of traffic and interval.
- C. Create an AWS CodeBuild configuration that triggers when the test code is pushed
- D. Use AWS CloudFormation to trigger an AWS CodePipeline configuration that deploys the new Lambda versions and specifies the traffic shift percentage and interval.
- E. Create an AWS CodePipeline configuration and set up the source code step to trigger when code is pushed
- F. Set up the build step to use AWS CodeBuild to run the test
- G. Set up an AWS CodeDeploy configuration to deploy, then select the CodeDeployDefault.LambdaLinear10PercentEvery3Minutes option.
- H. Use the AWS CLI to set up a post-commit hook that uploads the code to an Amazon S3 bucket after tests have passed
- I. Set up an S3 event trigger that runs a Lambda function that deploys the new version
- J. Use an interval in the Lambda function to deploy the code over time at the required percentage.

**Answer: C**

**Explanation:**

<https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations.html>

#### NEW QUESTION 169

A company has established tagging and configuration standards for its infrastructure resources running on AWS. A DevOps Engineer is developing a design that will provide a near-real-time dashboard of the compliance posture with the ability to highlight violations. Which approach meets the stated requirements?

- A. Define the resource configurations in AWS Service Catalog, and monitor the AWS Service Catalog compliance and violations in Amazon CloudWatch
- B. Then, set up and share a live CloudWatch dashboard
- C. Set up Amazon SNS notifications for violations and corrections.
- D. Use AWS Config to record configuration changes and output the data to an Amazon S3 bucket
- E. Create an Amazon QuickSight analysis of the dataset, and use the information on dashboards and mobile devices.
- F. Create a resource group that displays resources with the specified tags and those without tag
- G. Use the AWS Management Console to view compliant and non-compliant resources.
- H. Define the compliance and tagging requirements in Amazon Inspector
- I. Output the results to Amazon CloudWatch Log
- J. Build a metric filter to isolate the monitored elements of interest and present the data in a CloudWatch dashboard.

**Answer: B**

**Explanation:**

<https://aws.amazon.com/about-aws/whats-new/2019/03/aws-config-now-supports-tagging-of-aws-config-resources/>

#### NEW QUESTION 171

After conducting a disaster recovery exercise, an Enterprise Architect discovers that a large team of Database and Storage Administrators need more than seven hours of manual effort to make a flagship application's database functional in a different AWS Region. The Architect also discovers that the recovered database is often missing as much as two hours of data transactions.

Which solution provides improved RTO and RPO in a cross-region failover scenario?

- A. Deploy an Amazon RDS Multi-AZ instance backed by a multi-region Amazon ElastiCache
- B. Configure the RDS option group to enable multi-region availability for native automation of cross-region recovery and continuous data replication
- C. Create an Amazon SNS topic subscribed to RDS-impacted events to send emails to the Database Administration team when significant query latency is detected in a single Availability Zone.
- D. Use Amazon SNS topics to receive published messages from Amazon RDS availability and backup events
- E. Use AWS Lambda for three separate functions with calls to Amazon RDS to snapshot a database instance, create a cross-region snapshot copy, and restore an instance from a snapshot
- F. Use a scheduled Amazon CloudWatch Events rule at a frequency matching the RPO to trigger the Lambda function to snapshot a database instance
- G. Trigger the Lambda function to create a cross-region snapshot copy when the SNS topic for backup events receives a new message
- H. Configure the Lambda function to restore an instance from a snapshot to trigger sending new messages published to the availability SNS topic.
- I. Create a scheduled Amazon CloudWatch Events rule to make a call to Amazon RDS to create a snapshot from a database instance and specify a frequency to match the RPO
- J. Create an AWS Step Functions task to call Amazon RDS to perform a cross-region snapshot copy into the failover region, and configure the state machine to execute the task when the RDS snapshot creation state is complete
- K. Create an SNS topic subscribed to RDS availability events, and push these messages to an Amazon SQS queue located in the failover region
- L. Configure an Auto Scaling group of worker nodes to poll the queue for new messages and make a call to Amazon RDS to restore a database from a snapshot after a checksum on the cross-region copied snapshot returns valid.
- M. Use Amazon RDS scheduled instance lifecycle events to create a snapshot and specify a frequency to match the RPO
- N. Use Amazon RDS scheduled instance lifecycle event configuration to perform a cross-region snapshot copy into the failover region upon SnapshotCreateComplete event
- O. Configure Amazon CloudWatch to alert when the CloudWatch RDS namespace CPUUtilization metric for the database instance falls to 0% and make a call to Amazon RDS to restore the database snapshot in the failover region.

**Answer: B**

**Explanation:**

<https://aws.amazon.com/blogs/database/cross-region-automatic-disaster-recovery-on-amazon-rds-for-oracle-data/>

#### NEW QUESTION 175

A company has a hybrid architecture solution in which some legacy systems remain on-premises, while a specific cluster of servers is moved to AWS. The company cannot reconfigure the legacy systems, so the cluster nodes must have a fixed hostname and local IP address for each server that is part of the cluster. The DevOps Engineer must automate the configuration for a six-node cluster with high availability across three Availability Zones (AZs), placing two elastic network interfaces in a specific subnet for each AZ. Each node's hostname and local IP address should remain the same between reboots or instance failures. Which solution involves the LEAST amount of effort to automate this task?

- A. Create an AWS Elastic Beanstalk application and a specific environment for each server of the cluster. For each environment, give the hostname, elastic network interface, and AZ as input parameter
- B. Use the local health agent to name the instance and attach a specific elastic network interface based on the current environment.
- C. Create a reusable AWS CloudFormation template to manage an Amazon EC2 Auto Scaling group with a minimum size of 1 and a maximum size of 1. Give the hostname, elastic network interface, and AZ as stack parameter
- D. Use those parameters to set up an EC2 instance with EC2 Auto Scaling and a user data script to attach to the specific elastic network interface
- E. Use CloudFormation nested stacks to nest the template six times for a total of six nodes needed for the cluster, and deploy using the master template.
- F. Create an Amazon DynamoDB table with the list of hostnames, subnets, and elastic network interfaces to be used
- G. Create a single AWS CloudFormation template to manage an Auto Scaling group with a minimum size of 6 and a maximum size of 6. Create a programmatic solution that is installed in each instance that will lock/release the assignment of each hostname and local IP address, depending on the subnet in which a new instance will be launched.
- H. Create a reusable AWS CLI script to launch each instance individually, which will name the instance, place it in a specific AZ, and attach a specific elastic network interface
- I. Monitor the instances and in the event of failure, replace the missing instance manually by running the script again.

**Answer: B**

#### **NEW QUESTION 177**

A DevOps Engineer has been asked by the Security team to ensure that AWS CloudTrail files are not tampered with after being created. Currently, there is a process with multiple trails, using AWS IAM to restrict access to specific trails. The Security team wants to ensure they can trace the integrity of each file and make sure there has been no tampering.

Which option will require the LEAST effort to implement and ensure the legitimacy of the file while allowing the Security team to prove the authenticity of the logs?

- A. Create an Amazon CloudWatch Events rule that triggers an AWS Lambda function when a new file is delivered
- B. Configure the Lambda function to perform an MD5 hash check on the file, store the name and location of the file, and post the returned hash to an Amazon DynamoDB table
- C. The Security team can use the values stored in DynamoDB to verify the file authenticity.
- D. Enable the CloudTrail file integrity feature on an Amazon S3 bucket
- E. Create an IAM policy that grants the Security team access to the file integrity logs stored in the S3 bucket.
- F. Enable the CloudTrail file integrity feature on the trail
- G. Use the digest file created by CloudTrail to verify the integrity of the delivered CloudTrail files.
- H. Create an AWS Lambda function that is triggered each time a new file is delivered to the CloudTrail bucket
- I. Configure the Lambda function to execute an MD5 hash check on the file, and store the result on a tag in an Amazon S3 object
- J. The Security team can use the information on the tag to verify the integrity of the file.

**Answer: C**

#### **Explanation:**

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

#### **NEW QUESTION 178**

.....

## **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](#)