

## Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional

<https://www.2passeasy.com/dumps/DOP-C01/>



#### NEW QUESTION 1

A company used AWS CloudFormation to deploy a three-tier web application that stores data in an Amazon RDS MySQL Multi-AZ DB instance. A DevOps Engineer must upgrade the RDS instance to the latest major version of MySQL while incurring minimal downtime. How should the Engineer upgrade the instance while minimizing downtime?

- A. Update the EngineVersion property of the AWS::RDS::DBInstance resource type in the CloudFormation template to the latest desired version
- B. Launch a second stack and make the new RDS instance a read replica.
- C. Update the DBEngineVersion property of the AWS::RDS::DBInstance resource type in the CloudFormation template to the latest desired version
- D. Perform an Update Stack operation
- E. Create a new RDS Read Replicas resource with the same properties as the instance to be upgraded
- F. Perform a second Update Stack operation.
- G. Update the DBEngineVersion property of the AWS::RDS::DBInstance resource type in the CloudFormation template to the latest desired version
- H. Create a new RDS Read Replicas resource with the same properties as the instance to be upgraded
- I. Perform an Update Stack operation.
- J. Update the EngineVersion property of the AWS::RDS::DBInstance resource type in the CloudFormation template to the latest version, and perform an operation
- K. Update Stack

**Answer:** A

#### NEW QUESTION 2

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 3

A company has a web application that uses an Amazon DynamoDB table in a single AWS Region to store user information. To support an increasingly global user base, the application must run in a secondary Region and allow users to connect to their closest Region and fail over to the secondary Region. Which approach should be used to ensure the deployment meets these requirements?

- A. Configure DynamoDB streams to copy data between Regions, deploy the web stack in both Regions, and configure Amazon Route 53 to use a geoproximity routing policy with health checks.
- B. Convert the DynamoDB table to a global table, deploy the web stack in both Regions, and configure Amazon Route 53 to use a geoproximity routing policy with health checks.
- C. Define DynamoDB cross-region backups to copy data to the secondary Region, deploy the web stack in both Regions, and configure Amazon Route 53 to use a latency-based routing policy with health checks.
- D. Use DynamoDB Accelerator to copy data to the secondary Region, deploy the web stack in both Regions, and configure Amazon Route 53 to use a failover routing policy.

**Answer:** A

#### NEW QUESTION 4

A company wants to implement a CI/CD pipeline for building and testing its mobile apps. A DevOps Engineer has been given the following requirements: Use AWS CodePipeline to orchestrate the workflow. Test the application on real devices. Trigger a notification. Stage the application binary on a production bucket in a different account. Make the application binary publicly accessible. Which sequence of actions should the Engineer perform in the pipeline to meet the requirements?

- A. Use AWS CodeCommit as the code source and AWS CodeDeploy to compile and package the application
- B. Use CodeDeploy to deploy the application binary to an AWS Lambda function for testing
- C. Use a third-party library on AWS Lambda to simulate the device platform
- D. Allow a Lambda role to upload to the production Amazon S3 bucket
- E. Make the binary publicly accessible
- F. Trigger notifications using Amazon SNS.
- G. Use GitHub as the code source and AWS Lambda to compile and package the application
- H. Use another Lambda function to run unit tests and deliver the application binary to a development bucket
- I. Use the binary from the development bucket and install the application on a personal device for testing
- J. Deliver the binary to the production bucket after approval
- K. Trigger notifications using Amazon SNS.
- L. Use an Amazon S3 bucket as the code source and AWS CodeBuild to compile and package the application
- M. Use AWS CodeDeploy to deploy the application binary to a device farm for testing
- N. Deliver the binary to the production S3 bucket
- O. Use an S3 bucket policy to allow public read on the production S3 bucket
- P. Trigger notifications using an Amazon CloudWatch Events rule with Amazon SNS.

Q. Use AWS CodeCommit as the code source and AWS CodeBuild to compile and package the applicatio  
R. Invoke an AWS Lambda function that uploads the application binary to a device farm for testin  
S. Deliver the binary to the production Amazon S3 bucke  
T. Use an S3 bucket policy to allow public read on the production S3 bucke  
. Trigger notifications by using an Amazon CloudWatch Events rule.

**Answer:** D

#### NEW QUESTION 5

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

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

**Answer:** C

#### NEW QUESTION 6

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 developmen
- B. Use AWS SMS to import these containers as AMIs for Amazon EC2 whenever dependencies are update
- C. Use AWS CodePipeline to test new code changes against the Auto Scaling group.
- D. Developers should use their native Python environmen
- 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 EC
- F. Use AWS CloudFormation with the custom container to deploy the new Amazon ECS.
- G. Developers should use their native Python environmen
- 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 EC
- I. Use CodePipeline and CodeBuild with the custom container to test new code changes inside AWS Elastic Beanstalk

**Answer:** A

#### NEW QUESTION 7

According to Information Security Policy, changes to the contents of objects inside production Amazon S3 bucket that contain encrypted secrets should only be made by a trusted group of administrators. How should a DevOps Engineer create real-time, automated checks to meet this requirement?

- A. Create an AWS Lambda function that is triggered by Amazon S3 data events for object changes and that also checks the IAM user's membership in an administrator's IAM role.
- B. Create a periodic AWS Config rule to query Amazon S3 Logs for changes and to check the IAM user's membership in an administrator's IAM role.
- C. Create a metrics filter for Amazon CloudWatch logs to check for Amazon S3 bucket-level permission changes and to check the IAM user's membership in an administrator's IAM role.
- D. Create a periodic AWS Config rule to query AWS CloudTrail logs for changes to the Amazon S3 bucket-level permissions and to check the IAM user's membership in an administrator's IAM role.

**Answer:** A

#### NEW QUESTION 8

A DevOps Engineer needs to deploy a scalable three-tier Node.js application in AWS. The application must have zero downtime during deployments and be able to roll back to previous versions. Other applications will also connect to the same MySQL backend database.

The CIO has provided the following guidance for logging:

- \*Centrally view all current web access server logs.
- \*Search and filter web and application logs in near-real time.
- \*Retain log data for three months.

How should these requirements be met?

- A. Deploy the application using AWS Elastic Beanstal
- B. Configure the environment type for Elastic Load Balancing and Auto Scalin
- C. Create an Amazon RDS MySQL instance inside the Elastic Beanstalk stac
- D. Configure the Elastic Beanstalk log options to stream logs to Amazon CloudWatch Log
- E. Set retention to 90 days.
- F. Deploy the application on Amazon EC2. Configure Elastic Load Balancing and Auto Scalin
- G. Use an Amazon RDS MySQL instance for the database tie
- H. Configure the application to store log files in Amazon S3. Use Amazon EMR to search and filter the dat
- I. Set an Amazon S3 lifecycle rule to expire objects after 90 days.
- J. Deploy the application using AWS Elastic Beanstal
- K. Configure the environment type for Elastic Load Balancing and Auto Scalin
- L. Create the Amazon RDS MySQL instance outside the Elastic Beanstalk stac
- M. Configure the Elastic Beanstalk log options to stream logs to Amazon CloudWatch Log
- N. Set retention to 90 days.
- O. Deploy the application on Amazon EC2. Configure Elastic Load Balancing and Auto Scalin
- P. Use an Amazon RDS MySQL instance for the database tie

Q. Configure the application to load streaming log data using Amazon Kinesis Data Firehouse into Amazon E  
R. Delete and create a new Amazon ES domain every 90 days.

**Answer: B**

**Explanation:**

<https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-plan-debugging.html>

#### NEW QUESTION 9

An IT team has built an AWS CloudFormation template so others in the company can quickly and reliably deploy and terminate an application. The template creates an Amazon EC2 instance with a user data script to install the application and an Amazon S3 bucket that the application uses to serve static webpages while it is running.

All resources should be removed when the CloudFormation stack is deleted. However, the team observes that CloudFormation reports an error during stack deletion, and the S3 bucket created by the stack is not deleted.

How can the team resolve the error in the MOST efficient manner to ensure that all resources are deleted without errors?

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

**Answer: C**

#### NEW QUESTION 10

A DevOps Engineer discovered a sudden spike in a website's page load times and found that a recent deployment occurred. A brief diff of the related commit shows that the URL for an external API call was altered and the connecting port changed from 80 to 443. The external API has been verified and works outside the application. The application logs show that the connection is now timing out, resulting in multiple retries and eventual failure of the call.

Which debug steps should the Engineer take to determine the root cause of the issue?

- A. Check the VPC Flow Logs looking for denies originating from Amazon EC2 instances that are part of the web Auto Scaling group
- B. Check the ingress security group rules and routing rules for the VPC.
- C. Check the existing egress security group rules and network ACLs for the VP
- D. Also check the application logs being written to Amazon CloudWatch Logs for debug information.
- E. Check the egress security group rules and network ACLs for the VP
- F. Also check the VPC flow logs looking for accepts originating from the web Auto Scaling group.
- G. Check the application logs being written to Amazon CloudWatch Logs for debug information
- H. Check the ingress security group rules and routing rules for the VPC.

**Answer: C**

#### NEW QUESTION 10

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 entry
- 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 13

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 log
- C. Do the analysis directly from Amazon QuickSight.
- D. Keep the logs in Amazon S3 and use Amazon Redshift Spectrum to normalize the logs in place.
- E. 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.

**Answer: D**

#### NEW QUESTION 14

A company recently migrated its legacy application from on-premises to AWS. The application is hosted on Amazon EC2 instances behind an Application Load Balancer, which is behind Amazon API Gateway. The company wants to ensure users experience minimal disruptions during any deployment of a new version of the application. The company also wants to ensure it can quickly roll back updates if there is an issue.

Which solution will meet these requirements with MINIMAL changes to the application?



- A. Introduce changes as a separate environment parallel to the existing on
- B. Configure API Gateway to use a canary release deployment to send a small subset of user traffic to the new environment.
- C. Introduce changes as a separate environment parallel to the existing on
- D. Update the application's DNS alias records to point to the new environment.
- E. Introduce changes as a separate target group behind the existing Application Load Balance
- F. Configure API Gateway to route user traffic to the new target group in steps.
- G. Introduce changes as a separate target group behind the existing Application Load Balance
- H. Configure API Gateway to route all traffic to the Application Load Balancer, which then sends the traffic to the new target group.

**Answer:** A

#### NEW QUESTION 18

A company hosts parts of a Python-based application using AWS Elastic Beanstalk. An Elastic Beanstalk CLI is being used to create and update the environments. The Operations team detected an increase in requests in one of the Elastic Beanstalk environments that caused downtime overnight. The team noted that the policy used for AWS Auto Scaling is NetworkOut. Based on load testing metrics, the team determined that the application needs to scale CPU utilization to improve the resilience of the environments. The team wants to implement this across all environments automatically. Following AWS recommendations, how should this automation be implemented?

- A. Using ebextensions, place a command within the container\_commands key to perform an API call to modify the scaling metric to CPUUtilization for the Auto Scaling configuration
- B. Use leader\_only to execute this command in only the first instance launched within the environment.
- C. Using ebextensions, create a custom resource that modifies the AWSEBAutoScalingScaleUpPolicy and AWSEBAutoScalingScaleDownPolicy resources to use CPUUtilization as a metric to scale for the Auto Scaling group.
- D. Using ebextensions, configure the option setting MeasureName to CPUUtilization within the aws:autoscaling:trigger namespace.
- E. Using ebextensions, place a script within the files key and place it in /opt/elasticbeanstalk/hooks/appdeploy/pre to perform an API call to modify the scaling metric to CPUUtilization for the Auto Scaling configuration
- F. Use leader\_only to place this script in only the first instance launched within the environment.

**Answer:** C

#### NEW QUESTION 21

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

**Answer:** AC

#### NEW QUESTION 22

A company is creating a software solution that executes a specific parallel-processing mechanism. The software can scale to tens of servers in some special scenarios. This solution uses a proprietary library that is license-based, requiring that each individual server have a single, dedicated license installed. The company has 200 licenses and is planning to run 200 server nodes concurrently at most.

The company has requested the following features:

"¢ A mechanism to automate the use of the licenses at scale. "¢ Creation of a dashboard to use in the future to verify which licenses are available at any moment. What is the MOST effective way to accomplish these requirements?

- A. Upload the licenses to a private Amazon S3 bucket
- B. Create an AWS CloudFormation template with a Mappings section for the license
- C. In the template, create an Auto Scaling group to launch the server
- D. In the user data script, acquire an available license from the Mappings section
- E. Create an Auto Scaling lifecycle hook, then use it to update the mapping after the instance is terminated.
- F. Upload the licenses to an Amazon DynamoDB table
- G. Create an AWS CloudFormation template that uses an Auto Scaling group to launch the server
- H. In the user data script, acquire an available license from the DynamoDB table
- I. Create an Auto Scaling lifecycle hook, then use it to update the mapping after the instance is terminated.
- J. Upload the licenses to a private Amazon S3 bucket
- K. Populate an Amazon SQS queue with the list of licenses stored in S3. Create an AWS CloudFormation template that uses an Auto Scaling group to launch the server
- L. In the user data script acquire an available license from SQS
- M. Create an Auto Scaling lifecycle hook, then use it to put the license back in SQS after the instance is terminated.
- N. Upload the licenses to an Amazon DynamoDB table
- O. Create an AWS CLI script to launch the servers by using the parameter --count, with min:max instances to launch
- P. In the user data script, acquire an available license from the DynamoDB table
- Q. Monitor each instance and, in case of failure, replace the instance, then manually update the DynamoDB table.

**Answer:** D

#### NEW QUESTION 24

A DevOps Engineer just joined a new company that is already running workloads on Amazon EC2 instances. AWS has been adopted incrementally with no central governance. The Engineer must now assess how well the existing deployments comply with the following requirements:

\*EC2 instances are running only approved AMIs.

\*Amazon EBS volumes are encrypted.

\*EC2 instances have an Owner tag.

\*Root login over SSH is disabled on EC2 instances.

Which services should the Engineer use to perform this assessment with the LEAST amount of effort? (Select TWO.)

- A. AWS Config
- B. Amazon GuardDuty
- C. AWS System Manager
- D. AWS Directory Service
- E. Amazon Inspector

**Answer:** AE

**Explanation:**

[https://docs.aws.amazon.com/ja\\_jp/inspector/latest/userguide/inspector\\_security-best-practices.html](https://docs.aws.amazon.com/ja_jp/inspector/latest/userguide/inspector_security-best-practices.html)

#### NEW QUESTION 25

A company's application is currently deployed to a single AWS Region. Recently, the company opened a new office on a different continent. The users in the new office are experiencing high latency. The company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and uses Amazon DynamoDB as the database layer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. A DevOps Engineer is tasked with minimizing application response times and improving availability for users in both Regions.

Which combination of actions should be taken to address the latency issues? (Choose three.)

- A. Create a new DynamoDB table in the new Region with cross-Region replication enabled.
- B. Create new ALB and Auto Scaling group global resources and configure the new ALB to direct traffic to the new Auto Scaling group.
- C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group.
- D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB.
- E. Create Amazon Route 53 aliases, health checks, and failover routing policies to route to the ALB.
- F. Convert the DynamoDB table to a global table.

**Answer:** CDF

#### NEW QUESTION 30

A company runs an application with an Amazon EC2 and on-premises configuration. A DevOps engineer needs to standardize patching across both environments. Company policy dictates that patching only happens during non-business hours.

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

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

**Answer:** ABF

#### NEW QUESTION 35

The Deployment team has grown substantially in recent months and so has the number of projects that use separate code repositories. The current process involves configuring AWS CodePipeline manually, and there have been service limit alerts for the count of Amazon S3 buckets.

Which pipeline option will reduce S3 bucket sprawl alerts?

- A. Combine the multiple separate code repositories into a single one, and deploy using a global AWS CodePipeline that has logic for each project.
- B. Create new pipelines by using the AWS API or AWS CLI, and configure them to use a single global S3 bucket with separate prefixes for each project.
- C. Create a new pipeline in a different region for each project to bypass the service limits for S3 buckets in a single region.
- D. Create a new pipeline and for S3 bucket for each project by using the AWS API or AWS CLI to bypass the service limits for S3 buckets in a single account

**Answer:** A

#### NEW QUESTION 36

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 41

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 46**

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 49**

A company is using AWS Organizations and wants to implement a governance strategy with the following requirements:

- AWS resource access is restricted to the same two Regions for all accounts.
- AWS services are limited to a specific group of authorized services for all accounts.
- Authentication is provided by Active Directory.
- Access permissions are organized by job function and are identical in each account. Which solution will meet these requirements?

- A. Establish an organizational unit (OU) with group policies in the master account to restrict Regions and authorized services
- B. Use AWS CloudFormation StackSets to provision roles with permissions for each job function, including an IAM trust policy for IAM identity provider authentication in each account.
- C. Establish a permission boundary in the master account to restrict Regions and authorized services
- D. Use AWS CloudFormation StackSet to provision roles with permissions for each job function, including an IAM trust policy for IAM identity provider authentication in each account.
- E. Establish a service control policy in the master account to restrict Regions and authorized services
- F. Use AWS Resource Access Manager to share master account roles with permissions for each job function, including AWS SSO for authentication in each account.
- G. Establish a service control policy in the master account to restrict Regions and authorized services
- H. Use CloudFormation StackSet to provision roles with permissions for each job function, including an IAM trust policy for IAM identity provider authentication in each account.

**Answer:** D

**NEW QUESTION 52**

A company has developed a Node.js web application which provides REST services to store and retrieve time series data. The web application is built by the Development team on company laptops, tested locally, and manually deployed to a single on-premises server, which accesses a local MySQL database. The company is starting a trial in two weeks, during which the application will undergo frequent updates based on customer feedback. The following requirements must be met:

\*The team must be able to reliably build, test, and deploy new updates on a daily basis, without downtime or degraded performance.

\*The application must be able to scale to meet an unpredictable number of concurrent users during the trial. Which action will allow the team to quickly meet these



objectives?

- A. Create two Amazon Lightsail virtual private servers for Node.js; one for test and one for production. Build the Node.js application using existing process and upload it to the new Lightsail test server using the AWS CL
- B. Test the application, and if it passes all tests, upload it to the production serve
- C. During the trial, monitor the production server usage, and if needed, increase performance by upgrading the instance type.
- D. Develop an AWS CloudFormation template to create an Application Load Balancer and two Amazon EC2 instances with Amazon EBS (SSD) volumes in an Auto Scaling group with rolling updates enable
- E. Use AWS CodeBuild to build and test the Node.js application and store it in an Amazon S3 bucke
- F. Use user- data scripts to install the application and the MySQL database on each EC2 instanc
- G. Update the stack to deploy new application versions.
- H. Configure AWS Elastic Beanstalk to automatically build the application using AWS CodeBuild and to deploy it to a test environment that is configured to support auto scalin
- I. Create a second Elastic Beanstalk environment for productio
- J. Use Amazon RDS to store dat
- K. When new versions of the applications have passed all tests, use Elastic Beanstalk "'swap cname' to promote the test environment to production.
- L. Modify the application to use Amazon DynamoDB instead of a local MySQL databas
- M. Use AWS OpsWorks to create a stack for the application with a DynamoDB layer, an Application Load Balancer layer, and an Amazon EC2 instance laye
- N. Use a Chef recipe to build the application and a Chef recipe to deploy the application to the EC2 instance laye
- O. Use custom health checks to run unit tests on each instance with rollback on failure.

**Answer: C**

#### NEW QUESTION 56

A company uses AWS KMS with CMKs and manual key rotation to meet regulatory compliance requirements. The security team wants to be notified when any keys have not been rotated after 90 days.  
Which solution will accomplish this?

- A. Configure AWS KMS to publish to an Amazon SNS topic when keys are more than 90 days old.
- B. Configure an Amazon CloudWatch Events event to launch an AWS Lambda function to call the AWS Trusted Advisor API and publish to an Amazon SNS topic
- C. Develop an AWS Config custom rule that publishes to an Amazon SNS topic when keys are more than 90 days old
- D. Configure AWS Security Hub to publish to an Amazon SNS topic when keys are more than 90 days old.

**Answer: C**

#### NEW QUESTION 58

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 branc
- B. The team can delay the deployment until all features are ready.
- C. Add the code updates as commits to a feature branc
- 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 read
- 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 63

You have an application running a specific process that is critical to the application's functionality, and have added the health check process to your Auto Scaling Group. The instances are showing healthy but the application itself is not working as it should. What could be the issue with the health check, since it is still showing the instances as healthy.

- A. You do not have the time range in the health check properly configured
- B. It is not possible for a health check to monitor a process that involves the application
- C. The health check is not configured properly
- D. The health check is not checking the application process

**Answer: D**

#### Explanation:

If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use this information. For example, if you determine that an instance is not functioning as expected, you can set the health status of the instance to Unhealthy. The next time that Auto Scaling performs a health check on the instance, it will determine that the instance is unhealthy and then launch a replacement instance  
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 67

A company has multiple development teams sharing one AWS account. The development team's manager wants to be able to automatically stop Amazon EC2 instances and receive notifications if resources are idle and not tagged as production resources  
Which solution will meet these requirements?

- A. Use a scheduled Amazon CloudWatch Events rule to filter for Amazon EC2 instance status checks and identify idle EC2 instance
- B. Use the CloudWatch Events rule to target an AWS Lambda function to stop non-production instances and send notifications.
- C. Use a scheduled Amazon CloudWatch Events rule to filter AWS Systems Manager events and identify idle EC2 instances and resource



- D. Use the CloudWatch Events rule to target an AWS Lambda function to stop non-production instances and send notifications.
- E. Use a scheduled Amazon CloudWatch Events rule to target a custom AWS Lambda function that runs AWS Trusted Advisor checks Create a second CloudWatch Events rule to filter events from Trusted Advisor to trigger a Lambda function to stop idle non-production instances and send notifications
- F. Use a scheduled Amazon CloudWatch Events rule to target Amazon Inspector events for idle EC2 instances Use the CloudWatch Events rule to target the AWS Lambda function to stop non-production instances and send notifications

**Answer:** A

#### NEW QUESTION 69

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

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

How can log collection be automated?

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

**Answer:** D

#### Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html>

#### NEW QUESTION 73

After a recent audit, a company decided to implement a new disaster recovery strategy for its Amazon S3 data and its MySQL database running on Amazon EC2. Management wants the ability to recover to a secondary AWS Region with an RPO under 5 seconds and a RTO under 1 minute.

Which actions will meet the requirements while MINIMIZING operational overhead? (Select TWO.)

- A. Modify the application to write to both Regions at the same time when uploading objects to Amazon S3
- B. Migrate the database to an Amazon Aurora multi-master in the primary and secondary Regions.
- C. Migrate the database to Amazon RDS with a read replica in the secondary Region
- D. Migrate to Amazon Aurora Global Database.
- E. Set up S3 cross-Region replication with a replication SLA for the S3 buckets where objects are being put.

**Answer:** AE

#### NEW QUESTION 75

A company is using Docker containers for an application deployment and wants to move its application to AWS. The company currently manages its own clusters on premises to manage the deployment of these containers. It wants to deploy its application to a managed service in AWS and wants the entire flow of the deployment process to be automated. In addition, the company has the following requirements:

Focus first on the development workload. The environment must be easy to manage.

Deployment should be repeatable and reusable for new environments. Store the code in a GitHub repository.

Which solution will meet these requirements?

- A. Set up an Amazon ECS environment
- B. Use AWS CodePipeline to create a pipeline that is triggered on a commit to the GitHub repository
- C. Use AWS CodeBuild to create the container images and AWS CodeDeploy to publish the container image to the ECS environment.
- D. Use AWS CodePipeline that triggers on a commit from the GitHub repository, build the container images with AWS CodeBuild, and publish the container images to Amazon EC
- E. In the final stage, use AWS CloudFormation to create an Amazon ECS environment that gets the container images from the ECR repository.
- F. Create a Kubernetes Cluster on Amazon EC2. Use AWS CodePipeline to create a pipeline that is triggered when the code is committed to the repository
- G. Create the container images with a Jenkins server on EC2 and store them in the Docker Hub
- H. Use AWS Lambda from the pipeline to trigger the deployment to the Kubernetes Cluster.
- I. Set up an Amazon ECS environment
- J. Use AWS CodePipeline to create a pipeline that is triggered on a commit to the GitHub repository
- K. Use AWS CodeBuild to create the container and store it in the Docker Hub
- L. Use an AWS Lambda function to trigger a deployment and pull the new container image from the Docker Hub.

**Answer:** A

#### NEW QUESTION 80

A DevOps Engineer is asked to implement a strategy for deploying updates to a web application with zero downtime. The application infrastructure is defined in AWS CloudFormation and is made up of an Amazon Route 53 record, an Application Load Balancer, Amazon EC2 instances in an EC2 Auto Scaling group, and Amazon DynamoDB tables. To avoid downtime, there must be an active instance serving the application at all times.

Which strategies will ensure the deployment happens with zero downtime? (Select TWO.)

- A. In the CloudFormation template, modify the AWS::AutoScaling::DeploymentUpdates resource and add an UpdatePolicy attribute to define the required elements for a deployment with zero downtime.

- B. Add a new Application Load Balancer and Auto Scaling group to the CloudFormation template
- C. Deploy new changes to the inactive Auto Scaling group
- D. Use Route 53 to change the active Application Load Balancer.
- E. Add a new Application Load Balancer and Auto Scaling group to the CloudFormation template
- F. Modify the AWS::AutoScaling::AutoScalingGroup resource and add an UpdatePolicy attribute to perform rolling updates.
- G. In the CloudFormation template, modify the UpdatePolicy attribute for the CloudFormation stack and specify the Auto Scaling group that will be updated
- H. Configure MinSuccessfulInstancesPercent and PauseTime to ensure the deployment happens with zero downtime.

**Answer:** AC

#### NEW QUESTION 83

A security review has identified that an AWS CodeBuild project is downloading a database population script from an Amazon S3 bucket using an unauthenticated request. The security team does not allow unauthenticated requests to S3 buckets for this project. How can this issue be corrected in the MOST secure manner?

- A. Add the bucket name to the AllowedBuckets section of the CodeBuild project setting
- B. Update the build spec to use the AWS CLI to download the database population script.
- C. Modify the S3 bucket settings to enable HTTPS basic authentication and specify a token
- D. Update the build spec to use cURL to pass the token and download the database population script.
- E. Remove unauthenticated access from the S3 bucket with a bucket policy
- F. Modify the service role for the CodeBuild project to include Amazon S3 access
- G. Use the AWS CLI to download the database population script.
- H. Remove unauthenticated access from the S3 bucket with a bucket policy
- I. Use the AWS CLI to download the database population script using an IAM access key and a secret access key.

**Answer:** C

#### NEW QUESTION 84

A DevOps Engineer must automate a weekly process of identifying unnecessary permissions on a per-user basis, across all users in an AWS account. This process should evaluate the permissions currently granted to each user by examining the user's attached IAM access policies compared to the permissions the user has actually used in the past 90 days. Any differences in the comparison would indicate that the user has more permissions than are required. A report of the deltas should be sent to the Information Security team for further review and IAM user access policy revisions, as required. Which solution is fully automated and will produce the MOST detailed deltas report?

- A. Create an AWS Lambda function that calls the IAM Access Advisor API to pull service permissions granted on a user-by-user basis for all users in the AWS account
- B. Ensure that Access Advisor is configured with a tracking period of 90 days
- C. Invoke the Lambda function using an Amazon CloudWatch Events rule on a weekly schedule
- D. For each record, by user, by service, if the Access Advisor Last Accessed field indicates a day count instead of "Not accessed in the tracking period," this indicates a delta compared to what is in the user's currently attached access policy
- E. After Lambda has iterated through all users in the AWS account, configure it to generate a report and send the report using Amazon SES.
- F. Configure an AWS CloudTrail trail that spans all AWS Regions and all read/write events, and point this trail to an Amazon S3 bucket
- G. Create an Amazon Athena table and specify the S3 bucket ARN in the CREATE TABLE query
- H. Create an AWS Lambda function that accesses the Athena table using the SDK, which performs a SELECT, ensuring that the WHERE clause includes userIdentity, eventName, and eventTime
- I. Compare the results against the user's currently attached IAM access policies to determine any delta
- J. Configure an Amazon CloudWatch Events schedule to automate this process to run once a week
- K. Configure Amazon SES to send a consolidated report to the Information Security team.
- L. Configure VPC Flow Logs on all subnets across all VPCs in all regions to capture user traffic across the entire account
- M. Ensure that all logs are being sent to a centralized Amazon S3 bucket, so all flow logs can be consolidated and aggregated
- N. Create an AWS Lambda function that is triggered once a week by an Amazon CloudWatch Events schedule
- O. Ensure that the Lambda function parses the flow log files for the following information: IAM user ID, subnet ID, VPC ID, Allow/Reject status per API call, and service name
- P. Then have the function determine the deltas on a user-by-user basis
- Q. Configure the Lambda function to send the consolidated report using Amazon SES.
- R. Create an Amazon ES cluster and note its endpoint URL, which will be provided as an environment variable into a Lambda function
- S. Configure an Amazon S3 event on a AWS CloudTrail trail destination S3 bucket and ensure that the event is configured to send to a Lambda function
- T. Create the Lambda function to consume the events, parse the input from JSON, and transform it to an Amazon ES document format
- U. POST the documents to the Amazon ES cluster's endpoint by way of the passed-in environment variable
- V. Make sure that the proper indexing exists in Amazon ES and use Apache Lucene queries to parse the permissions on a user-by-user basis
- W. Export the deltas into a report and have Amazon ES send the reports to the Information Security team using Amazon SES every week.

**Answer:** C

#### NEW QUESTION 88

You are responsible for your company's large multi-tiered Windows-based web application running on Amazon EC2 instances situated behind a load balancer. While reviewing metrics, you've started noticing an upwards trend for slow customer page load time. Your manager has asked you to come up with a solution to ensure that customer load time is not affected by too many requests per second. Which technique would you use to solve this issue?

- A. Re-deploy your infrastructure using an AWS CloudFormation template
- B. Configure Elastic Load Balancing health checks to initiate a new AWS CloudFormation stack when health checks return failed.
- C. Re-deploy your infrastructure using an AWS CloudFormation template
- D. Spin up a second AWS CloudFormation stack
- E. Configure Elastic Load Balancing SpillOver functionality to spill over any slow connections to the second AWS CloudFormation stack.
- F. Re-deploy your infrastructure using AWS CloudFormation, Elastic Beanstalk, and Auto Scaling
- G. Set up your Auto Scaling group policies to scale based on the number of requests per second as well as the current customer load time
- H. Re-deploy your application using an Auto Scaling template
- I. Configure the Auto Scaling template to spin up a new Elastic Beanstalk application when the customer load time surpasses your threshold.

**Answer:** C

**Explanation:**

Auto Scaling helps you ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application. You create collections of EC2 instances, called Auto Scaling groups. You can specify the minimum number of instances in each Auto Scaling group, and Auto Scaling ensures that your group never goes below this size. You can specify the maximum number of instances in each Auto Scaling group, and Auto Scaling ensures that your group never goes above this size. If you specify the desired capacity, either when you create the group or at any time thereafter. Auto Scaling ensures that your group has this many instances. If you specify scaling policies, then Auto Scaling can launch or terminate instances as demand on your application increases or decreases. Option A and B are invalid because Autoscaling is required to solve the issue to ensure the application can handle high traffic loads. Option D is invalid because there is no Autoscaling template. For more information on Autoscaling, please refer to the below document link: from AWS  
➤ <http://docs.aws.amazon.com/autoscaling/latest/userguide/WhatIsAutoScaling.html>

**NEW QUESTION 91**

A Development team is building more than 40 applications. Each app is a three-tiered web application based on an ELB Application Load Balancer, Amazon EC2, and Amazon RDS. Because the applications will be used internally, the Security team wants to allow access to the 40 applications only from the corporate network and block access from external IP addresses. The corporate network reaches the internet through proxy servers. The proxy servers have 12 proxy IP addresses that are being changed one or two times per month. The Network Infrastructure team manages the proxy servers; they upload the file that contains the latest proxy IP addresses into an Amazon S3 bucket. The DevOps Engineer must build a solution to ensure that the applications are accessible from the corporate network. Which solution achieves these requirements with MINIMAL impact to application development, MINIMAL operational effort, and the LOWEST infrastructure cost?

- A. Implement an AWS Lambda function to read the list of proxy IP addresses from the S3 object and to update the ELB security groups to allow HTTPS only from the given IP addresses
- B. Configure the S3 bucket to invoke the Lambda function when the object is updated
- C. Save the IP address list to the S3 bucket when they are changed.
- D. Ensure that all the applications are hosted in the same Virtual Private Cloud (VPC). Otherwise, consolidate the applications into a single VPC
- E. Establish an AWS Direct Connect connection with an active/standby configuration
- F. Change the ELB security groups to allow only inbound HTTPS connections from the corporate network IP addresses.
- G. Implement a Python script with the AWS SDK for Python (Boto), which downloads the S3 object that contains the proxy IP addresses, scans the ELB security groups, and updates them to allow only HTTPS inbound from the given IP addresses
- H. Launch an EC2 instance and store the script in the instance
- I. Use a cron job to execute the script daily.
- J. Enable ELB security groups to allow HTTPS inbound access from the Internet
- K. Use Amazon Cognito to integrate the company's Active Directory as the identity provider
- L. Change the 40 applications to integrate with Amazon Cognito so that only company employees can log into the applications
- M. Save the user access logs to Amazon CloudWatch Logs to record user access activities

**Answer:** A

**NEW QUESTION 96**

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-accounts/>

**NEW QUESTION 99**

A company has developed an AWS Lambda function that handles orders received through an API. The company is using AWS CodeDeploy to deploy the Lambda function as the final stage of a CI/CD pipeline. A DevOps Engineer has noticed there are intermittent failures of the ordering API for a few seconds after deployment. After some investigation, the DevOps Engineer believes the failures are due to database changes. The CloudFormation stack for the application lambda function begins executing. How should the DevOps Engineer overcome this?

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



ready

**Answer:** B

#### NEW QUESTION 104

A company plans to stop using Amazon EC2 key pairs for SSH access, and instead plans to use AWS Systems Manager Session Manager. To further enhance security, access to Session Manager must take place over a private network only.

Which combinations of actions will accomplish this? (Select TWO.)

- A. Allow inbound access to TCP port 22 in all associated EC2 security groups from the VPC CIDR range.
- B. Attach an IAM policy with the necessary Systems Manager permissions to the existing IAM instance profile.
- C. Create a VPC endpoint for Systems Manager in the desired Region.
- D. Deploy a new EC2 instance that will act as a bastion host to the rest of the EC2 instance fleet.
- E. Remove any default routes in the associated route tables.

**Answer:** BC

#### NEW QUESTION 108

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 CloudFormation stacks using a CloudFormation service role only
- 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 only
- 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 111

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 ALB 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 the EU ALB as the first target and the US ALB as the second target.
- 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 multivalue answer routing 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
- 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 115

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 key
- 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 key
- 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 team
- 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 team
- 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 119

A DevOps engineer is assisting with a multi-Region disaster recovery solution for a new application. The application consists of Amazon EC2 instances running in an Auto Scaling group and an Amazon Aurora MySQL DB cluster. The application must be available with an RTO of 120 minutes and an RPO of 60 minutes. What is the MOST cost-effective way to meet these requirements?

- A. Launch an Aurora DB cluster as an Aurora Replica in a different Region
- B. Create an AWS CloudFormation template for all compute resources and create a stack in two Region
- C. Write a script that promotes the Aurora Replica to the primary instance in the event of a failure.
- D. Launch an Aurora DB cluster as an Aurora Replica in a different Region and configure automatic cross-Region failover
- E. Create an AWS CloudFormation template that includes an Auto Scaling group, and create a stack in two Region
- F. Write a script that updates the CloudFormation stack in the disaster recovery Region to increase the number of instances.
- G. Use AWS Lambda to create and copy a snapshot of the Aurora DB cluster to the destination Region hourly
- H. Create an AWS CloudFormation template that includes an Auto Scaling group, and create a stack in two Region
- I. Restore the Aurora DB cluster from a snapshot and update the Auto Scaling group to start launching instances.
- J. Configure Amazon DynamoDB cross-Region replication
- K. Create an AWS CloudFormation template that includes an Auto Scaling group, and create a stack in two Region
- L. Write a script that will update the CloudFormation stack in the disaster recovery Region and promote the DynamoDB replica to the primary instance in the event of a failure.

**Answer:** D

#### NEW QUESTION 123

An application has microservices spread across different AWS accounts and is integrated with an on-premises legacy system for some of its functionality. Because of the segmented architecture and missing logs, every time the application experiences issues, it is taking too long to gather the logs to identify the issues. A DevOps Engineer must fix the log aggregation process and provide a way to centrally analyze the logs. Which is the MOST efficient and cost-effective solution?

- A. Collect system logs and application logs by using the Amazon CloudWatch Logs agent
- B. Use the Amazon S3 API to export on-premises logs, and store the logs in an S3 bucket in a central account
- C. Build an Amazon EMR cluster to reduce the logs and derive the root cause.
- D. Collect system logs and application logs by using the Amazon CloudWatch Logs agent
- E. Use the Amazon S3 API to import on-premises log
- F. Store all logs in S3 buckets in individual account
- G. Use Amazon Macie to write a query to search for the required specific event-related data point.
- H. Collect system logs and application logs using the Amazon CloudWatch Logs agent
- I. Install the CloudWatch Logs agent on the on-premises server
- J. Transfer all logs from AWS to the on-premises data center
- K. Use an Amazon Elasticsearch Logstash Kibana stack to analyze logs on premises.
- L. Collect system logs and application logs by using the Amazon CloudWatch Logs agent
- M. Install a CloudWatch Logs agent for on-premises resource
- N. Store all logs in an S3 bucket in a central account
- O. Set up an Amazon S3 trigger and an AWS Lambda function to analyze incoming logs and automatically identify anomalies
- P. Use Amazon Athena to run ad hoc queries on the logs in the central account.

**Answer:** D

#### NEW QUESTION 124

A company wants to use Amazon DynamoDB for maintaining metadata on its forums. See the sample data set in the image below.

#### Thread

ForumName	Subject	LastPostDateTime	Thread
"S3"	"aaa"	"2015-03-15:17:24:31"	12
"S3"	"bbb"	"2015-01-22:23:18:01"	3
"S3"	"ccc"	"2015-02-31:13:14:21"	4
"S3"	"ddd"	"2015-01-03:09:21:11"	9
"EC2"	"yyy"	"2015-02-12:11:07:56"	18
"EC2"	"zzz"	"2015-01-18:07:33:42"	0
"RDS"	"ttt"	"2015-01-19:01:13:24"	3
"RDS"	"sss"	"2015-03-11:06:53:00"	11
"RDS"	"ttt"	"2015-10-22:12:19:44"	5

A DevOps Engineer is required to define the table schema with the partition key, the sort key, the local secondary index, projected attributes, and fetch operations. The schema should support the following example searches using the least provisioned read capacity units to minimize cost.

- Search within ForumName for items where the subject starts with "a".
- Search forums within the given LastPostDateTime time frame.
- Return the thread value where LastPostDateTime is within the last three months. Which schema meets the requirements?

- A. Use Subject as the primary key and ForumName as the sort ke
- B. Have LSI with LastPostDateTime as the sort key and fetch operations for thread.
- C. Use ForumName as the primary key and Subject as the sort ke
- D. Have LSI with LastPostDateTime as the sort key and the projected attribute thread.
- E. Use ForumName as the primary key and Subject as the sort ke
- F. Have LSI with Thread as the sort key and the projected attribute LastPostDateTime.
- G. Use Subject as the primary key and ForumName as the sort ke
- H. Have LSI with Thread as the sort key and fetch operations for LastPostDateTime.

**Answer:** B

**Explanation:**

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/LSI.html>

#### NEW QUESTION 125

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 projec
- 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 serve
- 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 rol
- 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 KM
- L. As part of the build script, decrypt with AWS KMS and use these credentials to run the integration tests.

**Answer:** B

#### NEW QUESTION 129

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 cluser.

**Answer:** D

#### NEW QUESTION 134

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

- A. Allow developers to check the code into a code repositor
- B. Using Amazon CloudWatch Events, on every pull into master, trigger an AWS Lambda function to build the artifact and store it in Amazon
- C. Create a custom script to clear the cach
- D. Specify the script in the BeforeInstall lifecycle hook in the AppSpec file.
- E. Create user data for each Amazon EC2 instance that contains the clear cache scrip
- F. Once deployed, test the applicatio
- G. If it is not successful, deploy it again.
- H. Set up AWS CodePipeline to deploy the applicatio
- I. Allow developers to check the code into a code repository as a source for the pipeline.
- J. Use AWS CodeBuild to build the artifact and place it in Amazon S3. Use AWS CodeDeploy to deploy the artifact to Amazon EC2 instances.
- K. Use AWS Systems Manager to fetch the artifact from Amazon S3 and deploy it to all the instances.

**Answer:** ADE

#### NEW QUESTION 139

A company is reviewing its IAM policies. One policy written by the DevOps Engineer has been flagged as too permissive. The policy is used by an AWS Lambda function that issues a stop command to Amazon EC2 instances tagged with Environment: Nonproduction over the weekend. The current policy is:



```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*"
    }
  ]
}
```

What changes should the Engineer make to achieve a policy of least permission? (Select THREE.)

A)

Add the following conditional expression:

```
"Condition": {
  "StringEquals": {
    "aws:principaltype": "lambda.amazonaws.com"
  }
}
```

B)

Change "Resource": "\*" to "Resource":  
 "arn:aws:ec2:\*:\*:instance/\*"

C)

Add the following conditional expression:

```
"Condition": {
  "StringNotEquals": {
    "ec2:ResourceTag/Environment": "Production"
  }
}
```

D)

Add the following conditional expression:

```
"Condition": {
  "StringEquals": {
    "ec2:ResourceTag/Environment": "NonProduction"
  }
}
```

E)

Change "Action": "ec2:\*" to "Action": "ec2:StopInstances"

F)

Add the following conditional expression:

```
"Condition" : {
  "DateGreaterThan" : {
    "aws:CurrentTime" : "${aws:DateTime:Friday}"
  },
  "DateLessThan": {
    "aws:CurrentTime" : "${aws:DateTime:Monday}"
  }
}
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E
- F. Option F

**Answer:** BDE

**Explanation:**

[https://docs.aws.amazon.com/ja\\_jp/IAM/latest/UserGuide/reference\\_policies\\_variables.html](https://docs.aws.amazon.com/ja_jp/IAM/latest/UserGuide/reference_policies_variables.html) <https://aws.amazon.com/jp/premiumsupport/knowledge-center/restrict-ec2-iam/>

#### NEW QUESTION 143

A government agency is storing highly confidential files in an encrypted Amazon S3 bucket. The agency has configured federated access and has allowed only a particular on-premises Active Directory user group to access this bucket. The agency wants to maintain audit records and automatically detect and revert any accidental changes administrators make to the IAM policies used for providing this restricted federated access. Which of the following options provide the FASTEST way to meet these requirements?

- A. Configure an Amazon CloudWatch Events Event Bus on an AWS CloudTrail API for triggering the AWS Lambda function that detects and reverts the change.
- B. Configure an AWS Config rule to detect the configuration change and execute an AWS Lambda function to revert the change.
- C. Schedule an AWS Lambda function that will scan the IAM policy attached to the federated access role for detecting and reverting any changes.
- D. Restrict administrators in the on-premises Active Directory from changing the IAM policies

**Answer:** B

#### Explanation:

<https://www.puresec.io/blog/aws-security-best-practices-config-rules-lambda-security> "Cloudwatch Event Bus" are used for -> "Sending and Receiving Events Between AWS Accounts"

<https://aws.amazon.com/about-aws/whats-new/2017/06/cloudwatch-events-adds-cross-account-event-delivery-s>

<https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config-rules.html>

#### NEW QUESTION 147

A company is required to collect user consent to a privacy agreement. An application is deployed in six AWS Regions with two in North America, two in Europe, and two in Asia with a user base of 20-30 million users. The company needs to read and write data related to each user's response, and ensure the responses are available in all six Regions. What solution will satisfy these requirements while MINIMIZING latency?

- A. Implement Amazon Aurora Global Database in each of the six Regions.
- B. Implement Amazon DocumentDB (with MongoDB compatibility) in each of the six Regions.
- C. Implement Amazon DynamoDB global tables in each of the six Regions.
- D. Implement Amazon ElastiCache for Redis replication group in each of the six Regions.

**Answer:** C

#### NEW QUESTION 150

Your company has multiple applications running on AWS. Your company wants to develop a tool that notifies on-call teams immediately via email when an alarm is triggered in your environment. You have multiple on-call teams that work different shifts, and the tool should handle notifying the correct teams at the correct times. How should you implement this solution?

- A. Create an Amazon SNS topic and an Amazon SQS queue
- B. Configure the Amazon SQS queue as a subscriber to the Amazon SNS topic. Configure CloudWatch alarms to notify this topic when an alarm is triggered
- C. Create an Amazon EC2 Auto Scaling group with both minimum and desired Instances configured to 0. Worker nodes in this group spawn when messages are added to the queue
- D. Workers then use Amazon Simple Email Service to send messages to your on-call teams.
- E. Create an Amazon SNS topic and configure your on-call team email addresses as subscriber
- F. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to this new topic
- G. Notifications will be sent to on-call users when a CloudWatch alarm is triggered.
- H. Create an Amazon SNS topic and configure your on-call team email addresses as subscriber
- I. Create a secondary Amazon SNS topic for alarms and configure your CloudWatch alarms to notify this topic when triggered
- J. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggered
- K. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to the first topic so that on-call engineers receive alerts.
- L. Create an Amazon SNS topic for each on-call group, and configure each of these with the team member emails as subscriber
- M. Create another Amazon SNS topic and configure your CloudWatch alarms to notify this topic when triggered
- N. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggered
- O. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to the correct team topic when on shift.

**Answer:** D

#### Explanation:

Option D fulfills all the requirements

1) First is to create a SNS topic for each group so that the required members get the email addresses.

2) Ensure the application uses the HTTPS endpoint and the SDK to publish messages Option A is invalid because the SQS service is not required.

Option B and C are incorrect. As per the requirement we need to provide notification to only those on-call teams who are working in that particular shift when an alarm is triggered. It need not have to be send to all the on-call teams of the company. With Option B & C, since we are not configuring the SNS topic for each on-call team the notifications will be send to all the on-call teams. Hence these 2 options are invalid. For more information on setting up notifications, please refer to the below document link: from AWS

➤ [http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US\\_SetupSNS.html](http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html)

#### NEW QUESTION 155

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 160

A global company with distributed Development teams built a web application using a microservices architecture running on Amazon ECS. Each application service is independent and runs as a service in the ECS cluster. The container build files and source code reside in a private GitHub source code repository. Separate ECS clusters exist for development, testing, and production environments.

Developers are required to push features to branches in the GitHub repository and then merge the changes into an environment-specific branch (development, test, or production). This merge needs to trigger an automated pipeline to run a build and a deployment to the appropriate ECS cluster.

What should the DevOps Engineer recommend as an automated solution to these requirements?

- A. Create an AWS CloudFormation stack for the ECS cluster and AWS CodePipeline service
- B. Store the container build files in an Amazon S3 bucket
- C. Use a post-commit hook to trigger a CloudFormation stack update that deploys the ECS cluster
- D. Add a task in the ECS cluster to build and push images to Amazon ECR, based on the container build files in S3.
- E. Create a separate pipeline in AWS CodePipeline for each environment
- F. Trigger each pipeline based on commits to the corresponding environment branch in GitHub
- G. Add a build stage to launch AWS CodeBuild to create the container image from the build file and push it to Amazon ECR
- H. Then add another stage to update the Amazon ECS task and service definitions in the appropriate cluster for that environment.
- I. Create a pipeline in AWS CodePipeline
- J. Configure it to be triggered by commits to the master branch in GitHub
- K. Add a stage to use the Git commit message to determine which environment the commit should be applied to, then call the create-image Amazon ECR command to build the image, passing it to the container build file
- L. Then add a stage to update the ECS task and service definitions in the appropriate cluster for that environment.
- M. Create a new repository in AWS CodeCommit
- N. Configure a scheduled project in AWS CodeBuild to synchronize the GitHub repository to the new CodeCommit repository
- O. Create a separate pipeline for each environment triggered by changes to the CodeCommit repository
- P. Add a stage using AWS Lambda to build the container image and push to Amazon ECR
- Q. Then add another stage to update the ECS task and service definitions in the appropriate cluster for that environment.

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 161

A highly regulated company has a policy that DevOps Engineers should not log in to their Amazon EC2 instances except in emergencies. If a DevOps Engineer does log in, the Security team must be notified within 15 minutes of the occurrence.

Which solution will meet these requirements?

- A. Install the Amazon Inspector agent on each EC2 instance
- B. Subscribe to Amazon CloudWatch Events notification
- C. Trigger an AWS Lambda function to check if a message is about user login
- D. If it is, send a notification to the Security team using Amazon SNS.
- E. Install the Amazon CloudWatch agent on each EC2 instance
- F. Configure the agent to push all logs to Amazon CloudWatch Logs and set up a CloudWatch metric filter that searches for user login
- G. If a login is found, send a notification to the Security team using Amazon SNS.
- H. Set up AWS CloudTrail with Amazon CloudWatch Log
- I. Subscribe CloudWatch Logs to Amazon Kinesis
- J. Attach AWS Lambda to Kinesis to parse and determine if a log contains a user login
- K. If it does, send a notification to the Security team using Amazon SNS.
- L. Set up a script on each Amazon EC2 instance to push all logs to Amazon S3. Set up an S3 event to trigger an AWS Lambda function, which triggers an Amazon Athena query to run
- M. The Athena query checks for logins and sends the output to the Security team using Amazon SNS.

**Answer:** B

#### NEW QUESTION 166

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
- 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 168

A retail company is currently hosting a Java-based application in its on-premises data center. Management wants the DevOps Engineer to move this application to AWS. Requirements state that while keeping high availability, infrastructure management should be as simple as possible. Also, during deployments of new application versions, while cost is an important metric, the Engineer needs to ensure that at least half of the fleet is available to handle user traffic. What option requires the LEAST amount of management overhead to meet these requirements?

- A. Create an AWS CodeDeploy deployment group and associate it with an Auto Scaling group configured to launch instances across subnets in different Availability Zone
- B. Configure an in-place deployment with a CodeDeploy.HalfAtATime configuration for application deployments.
- C. Create an AWS Elastic Beanstalk Java-based environment using Auto Scaling and load balancing. Configure the network setting for the environment to launch instances across subnets in different Availability Zone
- D. Use "Rolling with additional batch" as a deployment strategy with a batch size of 50%.
- E. Create an AWS CodeDeploy deployment group and associate it with an Auto Scaling group configured to launch instances across subnets in different Availability Zone
- F. Configure an in-place deployment with a custom deployment configuration with the MinimumHealthyHosts option set to type FLEET\_PERCENT and a value of 50.
- G. Create an AWS Elastic Beanstalk Java-based environment using Auto Scaling and load balancing. Configure the network options for the environment to launch instances across subnets in different Availability Zone
- H. Use "Rolling" as a deployment strategy with a batch size of 50%.

**Answer:** D

#### Explanation:

Rolling with batches keep 100% up you need 50%. With rolling deployments, Elastic Beanstalk splits the environment's EC2 instances into batches and deploys the new version of the application to one batch at a time, leaving the rest of the instances in the environment running the old version of the application.  
<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html>

#### NEW QUESTION 171

A company is using AWS CodeDeploy to automate software deployment. The deployment must meet these requirements:

- \*A number of instances must be available to serve traffic during the deployment. Traffic must be balanced across those instances, and the instances must automatically heal in the event of failure.
  - \*A new fleet of instances must be launched for deploying a new revision automatically, with no manual provisioning.
  - \*Traffic must be rerouted to the new environment to half of the new instances at a time. The deployment should succeed if traffic is rerouted to at least half of the instances; otherwise, it should fail.
  - \*Before routing traffic to the new fleet of instances, the temporary files generated during the deployment process must be deleted.
  - \*At the end of a successful deployment, the original instances in the deployment group must be deleted immediately to reduce costs.
- How can a DevOps Engineer meet these requirements?

- A. Use an Application Load Balancer and an in-place deployment
- B. Associate the Auto Scaling group with the deployment group
- C. Use the Automatically copy option, and use CodeDeployDefault.OneAtATime as the deployment configuration
- D. Instruct AWS CodeDeploy to terminate the original Auto Scaling group instances in the deployment group, and use the AllowTraffic hook within appspec.yml to delete the temporary files.
- E. Use an Application Load Balancer and a blue/green deployment
- F. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment group
- G. Use the Automatically copy Auto Scaling group option, create a custom deployment configuration with minimum healthy hosts defined as 50%, and assign the configuration to the deployment group
- H. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BeforeBlockTraffic hook within appspec.yml to delete the temporary files.
- I. Use an Application Load Balancer and a blue/green deployment
- J. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment group
- K. Use the Automatically copy Auto Scaling group option, and use CodeDeployDefault.HalfAtATime as the deployment configuration
- L. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BeforeAllowTraffic hook within appspec.yml to delete the temporary files.
- M. Use an Application Load Balancer and an in-place deployment
- N. Associate the Auto Scaling group and Application Load Balancer target group with the deployment group
- O. Use the Automatically copy AutoScaling group option, and use CodeDeployDefault.AllAtOnce as a deployment configuration
- P. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BlockTraffic hook within appspec.yml to delete the temporary files.

**Answer:** C

#### Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations.html>  
[https://docs.aws.amazon.com/codedeploy/latest/APIReference/API\\_BlueGreenDeploymentConfiguration.html](https://docs.aws.amazon.com/codedeploy/latest/APIReference/API_BlueGreenDeploymentConfiguration.html)

#### NEW QUESTION 175

A company is setting up a centralized logging solution on AWS and has several requirements. The company wants its Amazon CloudWatch Logs and VPC Flow logs to come from different sub accounts and to be delivered to a single auditing account. However, the number of sub accounts keeps changing. The company also needs to index the logs in the auditing account to gather actionable insight. How should a DevOps Engineer implement the solution to meet all of the company's requirements?

- A. Use AWS Lambda to write logs to Amazon ES in the auditing account
- B. Create an Amazon CloudWatch subscription filter and use Amazon Kinesis Data Streams in the sub accounts to stream the logs to the Lambda function deployed in the auditing account.
- C. Use Amazon Kinesis Streams to write logs to Amazon ES in the auditing account

- D. Create a CloudWatch subscription filter and use Kinesis Data Streams in the sub accounts to stream the logs to the Kinesis stream in the auditing account.
- E. Use Amazon Kinesis Firehose with Kinesis Data Streams to write logs to Amazon ES in the auditing account
- F. Create a CloudWatch subscription filter and stream logs from sub accounts to the Kinesis stream in the auditing account.
- G. Use AWS Lambda to write logs to Amazon ES in the auditing account
- H. Create a CloudWatch subscription filter and use Lambda in the sub accounts to stream the logs to the Lambda function deployed in the auditing account.

**Answer:** C

**Explanation:**

<https://aws.amazon.com/pt/blogs/architecture/central-logging-in-multi-account-environments/>

**NEW QUESTION 180**

A company has migrated its container-based applications to Amazon EKS and want to establish automated email notifications. The notifications sent to each email address are for specific activities related to EKS components. The solution will include Amazon SNS topics and an AWS Lambda function to evaluate incoming log events and publish messages to the correct SNS topic.

Which logging solution will support these requirements?

- A. Enable Amazon CloudWatch Logs to log the EKS component
- B. Create a CloudWatch subscription filter for each component with Lambda as the subscription feed destination.
- C. Enable Amazon CloudWatch Logs to log the EKS component
- D. Create CloudWatch Logs Insights queries linked to Amazon CloudWatch Events events that trigger Lambda.
- E. Enable Amazon S3 logging for the EKS component
- F. Configure an Amazon CloudWatch subscription filter for each component with Lambda as the subscription feed destination.
- G. Enable Amazon S3 logging for the EKS component
- H. Configure S3 PUT Object event notifications with AWS Lambda as the destination.

**Answer:** A

**NEW QUESTION 183**

A DevOps Engineer wants to prevent Developers from pushing updates directly to the company's master branch in AWS CodeCommit. These updates should be approved before they are merged.

Which solution will meet these requirements?

- A. Configure an IAM role for the Developers with access to CodeCommit and an explicit deny for write actions when the reference is the master
- B. Allow Developers to use feature branches and create a pull request when a feature is complete
- C. Allow an approver to use CodeCommit to view the changes and approve the pull requests.
- D. Configure an IAM role for the Developers to use feature branches and create a pull request when a feature is complete
- E. Allow CodeCommit to test all code in the feature branches, and dynamically modify the IAM role to allow merging the feature branches into the master
- F. Allow an approver to use CodeCommit to view the changes and approve the pull requests.
- G. Configure an IAM role for the Developers to use feature branches and create a pull request when a feature is complete
- H. Allow CodeCommit to test all code in the feature branches, and issue a new AWS Security Token Service (STS) token allowing a one-time API call to merge the feature branches into the master
- I. Allow an approver to use CodeCommit to view the changes and approve the pull requests.
- J. Configure an IAM role for the Developers with access to CodeCommit and attach an access policy to the CodeCommit repository that denies the Developers role access when the reference is master
- K. Allow Developers to use feature branches and create a pull request when a feature is complete
- L. Allow an approver to use CodeCommit to view the changes and approve the pull requests.

**Answer:** D

**NEW QUESTION 184**

A company uses a complex system that consists of networking, IAM policies, and multiple three-tier applications. Requirements are still being defined for a new system, so the number of AWS components present in the final design is not known. The DevOps Engineer needs to begin defining AWS resources using AWS CloudFormation to automate and version-control the new infrastructure.

What is the best practice for using CloudFormation to create new environments?

- A. Manually construct the networking layer using Amazon VPC and then define all other resources using CloudFormation.
- B. Create a single template to encompass all resources that are required for the system so there is only one template to version-control.
- C. Create multiple separate templates for each logical part of the system, use cross-stack references in CloudFormation, and maintain several templates in version control.
- D. Create many separate templates for each logical part of the system, and provide the outputs from one to the next using an Amazon EC2 instance running SDK for granular control.

**Answer:** C

**NEW QUESTION 189**

A company is running an application on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones.

After a recent application update, users are getting HTTP 502 Bad Gateway errors from the application URL. The DevOps Engineer cannot analyze the problem because Auto Scaling is terminating all EC2 instances shortly after launch for being unhealthy.

What steps will allow the DevOps Engineer access to one of the unhealthy instances to troubleshoot the deployed application?

- A. Create an image from the terminated instance and create a new instance from that image
- B. The Application team can then log into the new instance.
- C. As soon as a new instance is created by AutoScaling, put the instance into a Standby state as this will prevent the instance from being terminated.
- D. Add a lifecycle hook to your Auto Scaling group to move instances in the Terminating state to the Terminating:Wait state.
- E. Edit the Auto Scaling group to enable termination protection as this will protect unhealthy instances from being terminated.

**Answer:** B

**Explanation:**

<https://aws.amazon.com/blogs/aws/auto-scaling-update-lifecycle-standby-detach/>

**NEW QUESTION 193**

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 195**

A company wants to implement a CI/CD pipeline for an application that is deployed on AWS. The company also has a source-code analysis tool hosted on premises that checks for security flaws. The tool has not yet been migrated to AWS and can be accessed only on premises. The company wants to run checks against the source code as part of the pipeline before the code is compiled. The checks take anywhere from minutes to an hour to complete. How can a DevOps Engineer meet these requirements?

- A. Use AWS CodePipeline to create a pipeline
- B. Add an action to the pipeline to invoke an AWS Lambda function after the source stage
- C. Have the Lambda function invoke the source-code analysis tool on premises against the source input from CodePipeline
- D. The function then waits for the execution to complete and places the output in a specified Amazon S3 location.
- E. Use AWS CodePipeline to create a pipeline, then create a custom action type
- F. Create a job worker for the custom action that runs on hardware hosted on premises
- G. The job worker handles running security checks with the on-premises code analysis tool and then returns the job results to CodePipeline
- H. Have the pipeline invoke the custom action after the source stage.
- I. Use AWS CodePipeline to create a pipeline
- J. Add a step after the source stage to make an HTTPS request to the on-premises hosted web service that invokes a test with the source code analysis tool
- K. When the analysis is complete, the web service sends the results back by putting the results in an Amazon S3 output location provided by CodePipeline.
- L. Use AWS CodePipeline to create a pipeline
- M. Create a shell script that copies the input source code to a location on premises
- N. Invoke the source code analysis tool and return the results to CodePipeline
- O. Invoke the shell script by adding a custom script action after the source stage.

**Answer: B**

**NEW QUESTION 196**

A company has a single Developer writing code for an automated deployment pipeline. The Developer is storing source code in an Amazon S3 bucket for each project. The company wants to add more Developers to the team but is concerned about code conflicts and lost work. The company also wants to build a test environment to deploy newer versions of code for testing and allow Developers to automatically deploy to both environments when code is changed in the repository.

What is the MOST efficient way to meet these requirements?

- A. Create an AWS CodeCommit repository for each project, use the master branch for production code, and create a testing branch for code deployed to testing
- B. Use feature branches to develop new features and pull requests to merge code to testing and master branches.
- C. Create another S3 bucket for each project for testing code, and use an AWS Lambda function to promote code changes between testing and production bucket
- D. Enable versioning on all buckets to prevent code conflicts.
- E. Create an AWS CodeCommit repository for each project, and use the master branch for production and test code with different deployment pipelines for each environment
- F. Use feature branches to develop new features.
- G. Enable versioning and branching on each S3 bucket, use the master branch for production code, and create a testing branch for code deployed to testing
- H. Have Developers use each branch for developing in each environment.

**Answer: A**

**NEW QUESTION 197**

A DevOps Engineer is deploying a new web application. The company chooses AWS Elastic Beanstalk for deploying and managing the web application, and Amazon RDS MySQL to handle persistent data. The company requires that new deployments have minimal impact if they fail. The application resources must be at full capacity during deployment, and rolling back a deployment must also be possible.

Which deployment sequence will meet these requirements?

- A. Deploy the application using Elastic Beanstalk and connect to an external RDS MySQL instance using Elastic Beanstalk environment properties
- B. Use Elastic Beanstalk features for a blue/green deployment to deploy the new release to a separate environment, and then swap the CNAME in the two



environments to redirect traffic to the new version.

C. Deploy the application using Elastic Beanstalk, and include RDS MySQL as part of the environment. Use default Elastic Beanstalk behavior to deploy changes to the application, and let rolling updates deploy changes to the application.

D. Deploy the application using Elastic Beanstalk, and include RDS MySQL as part of the environment. Use Elastic Beanstalk immutable updates for application deployments.

E. Deploy the application using Elastic Beanstalk, and connect to an external RDS MySQL instance using Elastic Beanstalk environment properties.

F. Use Elastic Beanstalk immutable updates for application deployments.

**Answer:** A

#### NEW QUESTION 198

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

How should a DevOps Engineer meet these requirements?

A. In both regions, deploy the application on AWS Elastic Beanstalk and use Amazon DynamoDB global tables for session data.

B. Use an Amazon Route 53 weighted routing policy with health checks to distribute the traffic across the regions.

C. In both regions, launch the application in Auto Scaling groups and use DynamoDB for session data.

D. Use a Route 53 failover routing policy with health checks to distribute the traffic across the regions.

E. In both regions, deploy the application in AWS Lambda, exposed by Amazon API Gateway, and use Amazon RDS PostgreSQL with cross-region replication for session data.

F. Deploy the web application with client-side logic to call the API Gateway directly.

G. In both regions, launch the application in Auto Scaling groups and use DynamoDB global tables for session data.

H. Enable an Amazon CloudFront weighted distribution across regions.

I. Point the Amazon Route 53 DNS record at the CloudFront distribution.

**Answer:** B

#### NEW QUESTION 200

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-servicerole.html>

#### NEW QUESTION 204

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 207

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DOP-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DOP-C01 Product From:

<https://www.2passeasy.com/dumps/DOP-C01/>

## Money Back Guarantee

### **DOP-C01 Practice Exam Features:**

- \* DOP-C01 Questions and Answers Updated Frequently
- \* DOP-C01 Practice Questions Verified by Expert Senior Certified Staff
- \* DOP-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* DOP-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year