

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

<https://www.2passeasy.com/dumps/AWS-Certified-DevOps-Engineer-Professional/>



NEW QUESTION 1

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 2

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 3

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 4

A company has containerized all of its in-house quality control applications. The company is running Jenkins on Amazon EC2. which requires patching and upgrading. The compliance officer has requested a DevOps engineer begin encrypting build artifacts since they contain company intellectual property.

What should the DevOps engineer do to accomplish this in the MOST maintainable manner?

- A. Automate patching and upgrading using AWS Systems Manager on EC2 instances and encrypt Amazon EBS volumes by default.
- B. Deploy Jenkins to an Amazon ECS cluster and copy build artifacts to an Amazon S3 bucket with default encryption enabled.
- C. Leverage AWS CodePipeline with a build action and encrypt the artifacts using AWS Secrets Manager.
- D. Use AWS CodeBuild with artifact encryption to replace the Jenkins instance running on Amazon EC2.

Answer: D

NEW QUESTION 5

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 6

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://aws.amazon.com/jp/premiumsupport/knowledge-center/restrict-ec2-iam/>

NEW QUESTION 7

An Application team has three environments for their application: development, pre-production, and production. The team recently adopted AWS CodePipeline. However, the team has had several deployments of misconfigured or nonfunctional development code into the production environment, resulting in user disruption and downtime. The DevOps Engineer must review the pipeline and add steps to identify problems with the application before it is deployed. What should the Engineer do to identify functional issues during the deployment process? (Choose two.)

- A. Use Amazon Inspector to add a test action to the pipeline
- B. Use the Amazon Inspector Runtime Behavior Analysis Inspector rules package to check that the deployed code complies with company security standards before deploying it to production.
- C. Using AWS CodeBuild to add a test action to the pipeline to replicate common user activities and ensure that the results are as expected before progressing to production deployment.
- D. Create an AWS CodeDeploy action in the pipeline with a deployment configuration that automatically deploys the application code to a limited number of instances
- E. The action then pauses the deployment so that the QA team can review the application functionality
- F. When the review is complete, CodeDeploy resumes and deploys the application to the remaining production Amazon EC2 instances.
- G. After the deployment process is complete, run a testing activity on an Amazon EC2 instance in a different region that accesses the application to simulate user behavior
- H. If unexpected results occur, the testing activity sends a warning to an Amazon SNS topic
- I. Subscribe to the topic to get updates.
- J. Add an AWS CodeDeploy action in the pipeline to deploy the latest version of the development code to pre-production
- K. Add a manual approval action in the pipeline so that the QA team can test and confirm the expected functionality
- L. After the manual approval action, add a second CodeDeploy action that deploys the approved code to the production environment.

Answer: BE

Explanation:

<https://docs.aws.amazon.com/codepipeline/latest/userguide/integrations-action-type.html#integrations-test>
<https://docs.aws.amazon.com/codepipeline/latest/userguide/integrations-action-type.html#integrations-deploy>

NEW QUESTION 8

A company has an application that has predictable peak traffic times. The company wants the application instances to scale up only during the peak times. The application stores state in Amazon DynamoDB. The application environment uses a standard Node.js application stack and custom Chef recipes stored in a private Git repository.

Which solution is MOST cost-effective and requires the LEAST amount of management overhead when performing rolling updates of the application environment?

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

Answer: D

NEW QUESTION 9

A DevOps engineer wants to deploy a serverless web application based on AWS Lambda. The deployment must meet the following requirements:

- Provide staging and production environments.
- Restrict the developers from accessing the production environment.
- Avoid hard coding passwords in the Lambda functions
- Store source code in AWS CodeCommit.
- Use AWS CodePipeline to automate the deployment. Which solution will accomplish this?

- A. Create separate staging and production accounts to segregate deployment targets
- B. Use AWS KMS to store environment-specific values Use CodePipeline to automate deployments with AWS CodeDeploy.
- C. Create separate staging and production accounts to segregate deployment targets
- D. Use Lambda environment variables to store environment-specific values
- E. Use CodePipeline to automate deployments with AWS CodeDeploy.
- F. Define tagging conventions for staging and production environments to segregate deployment targets. Use AWS KMS to store environment-specific values Use CodePipeline to automate deployments with AWS CodeDeploy.

- G. Define naming conventions for staging and production environments to segregate deployment targets. Use Lambda environment variables to store environment-specific value
- H. Use CodePipeline to automate deployments with AWS CodeDeploy

Answer: A

NEW QUESTION 10

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

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

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

Answer: BE

NEW QUESTION 10

A Security team is concerned that a Developer can unintentionally attach an Elastic IP address to an Amazon EC2 instance in production. No Developer should be allowed to attach an Elastic IP address to an instance. The Security team must be notified if any production server has an Elastic IP address at any time. How can this task be automated?

- A. Use Amazon Athena to query AWS CloudTrail logs to check for any associate-address attempt
- B. Create an AWS Lambda function to dissociate the Elastic IP address from the instance, and alert the Security team.
- C. Attach an IAM policy to the Developer's IAM group to deny associate-address permission
- D. Create a custom AWS Config rule to check whether an Elastic IP address is associated with any instance tagged as production, and alert the Security team.
- E. Ensure that all IAM groups are associated with Developers do not have associate-address permissions. Create a scheduled AWS Lambda function to check whether an Elastic IP address is associated with any instance tagged as production, and alert the Security team if an instance has an Elastic IP address associated with it.
- F. Create an AWS Config rule to check that all production instances have the EC2 IAM roles that include deny associate-address permission
- G. Verify whether there is an Elastic IP address associated with any instance, and alert the Security team if an instance has an Elastic IP address associated with it.

Answer: B

Explanation:

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-migrate-ipv6.html#vpc-migrate-ipv6-sg-rules>

NEW QUESTION 14

A production account has a requirement that any Amazon EC2 instance that has been logged into manually must be terminated within 24 hours. All applications in the production account are using Auto Scaling groups with Amazon CloudWatch Logs agent configured. How can this process be automated?

- A. Create a CloudWatch Logs subscription to an AWS Step Functions applicatio
- B. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned
- C. Then create a CloudWatch Events rule to trigger a second AWS Lambda function once a day that will terminate all instances with this tag.
- D. Create a CloudWatch alarm that will trigger on the login event
- E. Send the notification to an Amazon SNS topic that the Operations team is subscribed to, and have them terminate the EC2 instance within 24 hours.
- F. Create a CloudWatch alarm that will trigger on the login event
- G. Configure the alarm to send to an Amazon SQS queue
- H. Use a group of worker instances to process messages from the queue, which then schedules the Amazon CloudWatch Events rule to trigger.
- I. Create a CloudWatch Logs subscription in an AWS Lambda function
- J. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned
- K. Create a CloudWatch Events rule to trigger a daily Lambda function that terminates all instances with this tag

Answer: D

Explanation:

<https://boto3.amazonaws.com/v1/documentation/api/latest/guide/cw-example-subscription-filters.html>

NEW QUESTION 15

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

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

Answer: A

NEW QUESTION 16

A DevOps engineer is architecting a continuous development strategy for a company's software as a service (SaaS) web application running on AWS. For application and security reasons, users subscribing to this application are distributed across multiple Application Load Balancers (ALBs), each of which has a dedicated Auto Scaling group and fleet of Amazon EC2 instances. The application does not require a build stage, and when it is committed to AWS CodeCommit, the application must trigger a simultaneous deployment to all ALBs. Auto Scaling groups, and EC2 fleets. Which architecture will meet these requirements with the LEAST amount of configuration?

- A. Create a single AWS CodePipeline pipeline that deploys the application in parallel using unique AWS CodeDeploy applications and deployment groups created for each ALB-Auto Scaling group pair.
- B. Create a single AWS CodePipeline pipeline that deploys the application using a single AWS CodeDeploy application and single deployment group.
- C. Create a single AWS CodePipeline pipeline that deploys the application in parallel using a single AWS CodeDeploy application and unique deployment group for each ALB-Auto Scaling group pair.
- D. Create an AWS CodePipeline pipeline for each ALB-Auto Scaling group pair that deploys the application using an AWS CodeDeploy application and deployment group created for the same ALB-Auto Scaling group pair.

Answer: C

NEW QUESTION 17

A DevOps engineer is tasked with creating a more stable deployment solution for a web application in AWS. Previous deployments have resulted in user-facing bugs, premature user traffic, and inconsistencies between web servers running behind an Application Load Balancer. The current strategy uses AWS CodeCommit to store the code for the application. When developers push to the master branch of the repository. CodeCommit triggers an AWS Lambda deploy function, which invokes an AWS Systems Manager run command to build and deploy the new code to all Amazon EC2 instances. Which combination of actions should be taken to implement a more stable deployment solution? (Select TWO.)

- A. Create a pipeline in AWS CodePipeline with CodeCommit as a source provide
- B. Create parallel pipeline stages to build and test the applicatio
- C. Pass the build artifact to AWS CodeDeploy.
- D. Create a pipeline in AWS CodePipeline with CodeCommit as a source provide
- E. Create separate pipeline stages to build and then test the applicatio
- F. Pass the build artifact to AWS CodeDeploy.
- G. Create and use an AWS CodeDeploy application and deployment group to deploy code updates to the EC2 flee
- H. Select the Application Load Balancer for the deployment group.
- I. Create individual Lambda functions to run all build, test, and deploy actions using AWS CodeDeploy instead of AWS Systems Manager.
- J. Modify the Lambda function to build a single application package to be shared by all instance
- K. Use AWS CodeDeploy instead of AWS Systems Manager to update the code on the EC2 fleet.

Answer: CE

NEW QUESTION 19

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

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

Answer: C

NEW QUESTION 23

A DevOps Engineer is launching a new application that will be deployed using Amazon Route 53, an Application Load Balancer, Auto Scaling, and Amazon DynamoDB. One of the key requirements of this launch is that the application must be able to scale to meet a sudden load increase. During periods of low usage, the infrastructure components must scale down to optimize cost. What steps can the DevOps Engineer take to meet the requirements? (Select TWO.)

- A. Use AWS Trusted Advisor to submit limit increase requests for the Amazon EC2 instances that will be used by the infrastructure.
- B. Determine which Amazon EC2 instance limits need to be raised by leveraging AWS Trusted Advisor, and submit a request to AWS Support to increase those limits.
- C. Enable Auto Scaling for the DynamoDB tables that are used by the application.
- D. Configure the Application Load Balancer to automatically adjust the target group based on the current load.
- E. Create an Amazon CloudWatch Events scheduled rule that runs every 5 minutes to track the current use of the Auto Scaling grou
- F. If usage has changed, trigger a scale-up event to adjust the capacit
- G. Do the same for DynamoDB read and write capacities.

Answer: BC

NEW QUESTION 27

A DevOps engineer is researching the least expensive way to implement an image batch processing cluster on AWS. The application cannot run in Docker containers and must run on Amazon EC2. The batch job stores checkpoint data on an NFS and can tolerate interruptions. Configuring the cluster software from a generic EC2 Linux image takes 30 minutes. What is the MOST cost-effective solution?

- A. Use Amazon EFS for checkpoint dat
- B. To complete the jo
- C. use an EC2 Auto Scaling group and an On-Demand pricing model to provision EC2 instances temporarily.
- D. Use GlusterFS on EC2 instances for checkpoint dat

- E. To run the batch job
- F. configure EC2 instances manually
- G. When the job completes, shut down the instances manually.
- H. Use Amazon EFS for checkpoint data
- I. Use EC2 Fleet to launch EC2 Spot Instances, and utilize user data to configure the EC2 Linux instance on startup.
- J. Use Amazon EFS for checkpoint data
- K. Use EC2 Fleet to launch EC2 Spot Instance
- L. Create a custom AMI for the cluster and use the latest AMI when creating instances.

Answer: A

NEW QUESTION 28

A company is building a solution for storing files containing Personally Identifiable Information (PII) on AWS.

Requirements state:

*All data must be encrypted at rest and in transit.

*All data must be replicated in at least two locations that are at least 500 miles apart. Which solution meets these requirements?

- A. Create primary and secondary Amazon S3 buckets in two separate Availability Zones that are at least 500 miles apart
- B. Use a bucket policy to enforce access to the buckets only through HTTP
- C. Use a bucket policy to enforce Amazon S3 SSE-C on all objects uploaded to the buckets
- D. Configure cross-region replication between the two buckets.
- E. Create primary and secondary Amazon S3 buckets in two separate AWS Regions that are at least 500 miles apart
- F. Use a bucket policy to enforce access to the buckets only through HTTP
- G. Use a bucket policy to enforce S3-Managed Keys (SSE-S3) on all objects uploaded to the buckets
- H. Configure cross-region replication between the two buckets.
- I. Create primary and secondary Amazon S3 buckets in two separate AWS Regions that are at least 500 miles apart
- J. Use an IAM role to enforce access to the buckets only through HTTP
- K. Use a bucket policy to enforce Amazon S3-Managed Keys (SSE-S3) on all objects uploaded to the buckets
- L. Configure cross-region replication between the two buckets.
- M. Create primary and secondary Amazon S3 buckets in two separate Availability Zones that are at least 500 miles apart
- N. Use a bucket policy to enforce access to the buckets only through HTTP
- O. Use a bucket policy to enforce AWS KMS encryption on all objects uploaded to the buckets
- P. Configure cross-region replication between the two buckets
- Q. Create a KMS Customer Master Key (CMK) in the primary region for encrypting objects.

Answer: B

NEW QUESTION 29

A company is using AWS CodeDeploy to manage its application deployments. Recently, the Development team decided to use GitHub for version control, and the team is looking for ways to integrate the GitHub repository with CodeDeploy. The team also needs to develop a way to automate deployment whenever there is a new commit on that repository. The team is currently deploying new application revisions by manually indicating the Amazon S3 location.

How can the integration be achieved in the MOST efficient way?

- A. Create a GitHub webhook to replicate the repository to AWS CodeCommit
- B. Create an AWS CodePipeline pipeline that uses CodeCommit as a source provider and AWS CodeDeploy as a deployment provider
- C. Once configured, commit a change to the GitHub repository to start the first deployment.
- D. Create an AWS CodePipeline pipeline that uses GitHub as a source provider and AWS CodeDeploy as a deployment provider
- E. Connect this new pipeline with the GitHub account and instruct CodePipeline to use webhooks in GitHub to automatically start the pipeline when a change occurs.
- F. Create an AWS Lambda function to check periodically if there has been a new commit within the GitHub repository
- G. If a new commit is found, trigger a CreateDeployment API call to AWS CodeDeploy to start a new deployment based on the last commit ID within the deployment group.
- H. Create an AWS CodeDeploy custom deployment configuration to associate the GitHub repository with the deployment group
- I. During the association process, authenticate the deployment group with GitHub to obtain the GitHub security authentication token
- J. Configure the deployment group options to automatically deploy if a new commit is found
- K. Perform a new commit to the GitHub repository to trigger the first deployment.

Answer: B

NEW QUESTION 31

A company has a website in an AWS Elastic Beanstalk load balancing and automatic scaling environment. This environment has an Amazon RDS MySQL instance configured as its database resource. After a sudden increase in traffic, the website started dropping traffic. An administrator discovered that the application on some instances is not responding as the result of out-of-memory errors. Classic Load Balancer marked those instances as out of service, and the health status of Elastic Beanstalk enhanced health reporting is degraded. However, Elastic Beanstalk did not replace those instances. Because of the diminished capacity behind the Classic Load Balancer, the application response times are slower for the customers.

Which action will permanently fix this issue?

- A. Clone the Elastic Beanstalk environment
- B. When the new environment is up, swap CNAME and terminate the earlier environment.
- C. Temporarily change the maximum number of instances in the Auto Scaling group to allow the group to support more traffic.
- D. Change the setting for the Auto Scaling group health check from Amazon EC2 to Elastic Load Balancing, and increase the capacity of the group.
- E. Write a cron script for restraining the web server process when memory is full, and deploy it with AWS Systems Manager.

Answer: C

NEW QUESTION 36

Company policies require that information about IP traffic going between instances in the production Amazon VPC is captured. The capturing mechanism must always be enabled and the Security team must be notified when any changes in configuration occur.

What should be done to ensure that these requirements are met?

- A. Using the UserData section of an AWS CloudFormation template, install tcpdump on every provisioned Amazon EC2 instance
- B. The output of the tool is sent to Amazon EFS for aggregation and queryin
- C. In addition, scheduling an Amazon CloudWatch Events rule calls an AWS Lambda function to check whether tcpdump is up and running and sends an email to the security organization when there is an exception.
- D. Create a flow log for the production VPC and assign an Amazon S3 bucket as a destination for delivery.Using Amazon S3 Event Notification, set up an AWS Lambda function that is triggered when a new log file gets delivered
- E. This Lambda function updates an entry in Amazon DynamoDB, which is periodically checked by scheduling an Amazon CloudWatch Events rule to notify security when logs have not arrived.
- F. Create a flow log for the production VP
- G. Create a new rule using AWS Config that is triggered by configuration changes of resources of type "'EC2:VPC'. As part of configuring the rule, create an AWS Lambda function that looks up flow logs for a given VP
- H. If the VPC flow logs are not configured, return a "'NON_COMPLIANT' status and notify the security organization.
- I. Configure a new trail using AWS CloudTrail service
- J. Using the UserData section of an AWS CloudFormation template, install tcpdump on every provisioned Amazon EC2 instance
- K. Connect Amazon Athena to the CloudTrail and write an AWS Lambda function that monitors for a flow log disabled even
- L. Once the CloudTrail entry has been spotted, alert the security organization

Answer: C

NEW QUESTION 38

A company is deploying a new mobile game on AWS for its customers around the world. The Development team uses AWS Code services and must meet the following requirements:

- Clients need to send/receive real-time playing data from the backend frequently and with minimal latency
- Game data must meet the data residency requirement

Which strategy can a DevOps Engineer implement to meet their needs?

- A. Deploy the backend application to multiple region
- B. Any update to the code repository triggers a two-stage build and deployment pipeline
- C. A successful deployment in one region invokes an AWSLambda function to copy the build artifacts to an Amazon S3 bucket in another region
- D. After the artifacts are copied, it triggers a deployment pipeline in the new region.
- E. Deploy the backend application to multiple Availability Zones in a single region
- F. Create an Amazon CloudFront distribution to serve the application backend to global customer
- G. Any update to the code repository triggers a two-stage build-and-deployment pipeline
- H. The pipeline deploys the backend application to all Availability Zones.
- I. Deploy the backend application to multiple region
- J. Use AWS Direct Connect to serve the application backend to global customer
- K. Any update to the code repository triggers a two-stage build-and-deployment pipeline in the region
- L. After a successful deployment in the region, the pipeline continues to deploy the artifact to another region.
- M. Deploy the backend application to multiple region
- N. Any update to the code repository triggers a two-stage build-and-deployment pipeline in the region
- O. After a successful deployment in the region, the pipeline invokes the pipeline in another region and passes the build artifact location
- P. The pipeline uses the artifact location and deploys applications in the new region.

Answer: A

NEW QUESTION 39

A company is using an AWS CodeBuild project to build and package an application. The packages are copied to a shared Amazon S3 bucket before being deployed across multiple AWS accounts.

The buildspec.yml file contains the following:

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

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

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

Answer: D

NEW QUESTION 44

A company is using AWS CodeCommit as its source code repository. After an internal audit, the compliance team mandates that any code change that goes into the master branch must be committed by senior developers.

Which solution will meet these requirements?

- A. Create two repositories in CodeCommit: one for working and another for the master
- B. Create separate IAM groups for senior developers and developer
- C. Assign the resource-level permissions on the repositories tied to the IAM group
- D. After the code changes are reviewed, sync the approved files to the master code commit repository.
- E. Create a repository in CodeCommit
- F. Create separate IAM groups for senior developers and developers. Assign code commit permissions for both groups, with code merge permissions for the senior developers group
- G. Create a trigger to notify senior developers with a URL link to approve or deny commit requests delivered through Amazon SNS
- H. Once a senior developer approves the code, the code gets merged to the master branch.

- I. Create a repository in CodeCommit with a working and master branch
- J. Create separate IAM groups for senior developers and developer
- K. Use an IAM policy to assign each IAM group their corresponding branches
- L. Once the code is merged to the working branch, senior developers can pull the changes from the working branch to the master branch.
- M. Create a repository in CodeCommit
- N. Create separate IAM groups for senior developers and developers. Use AWS Lambda triggers on the master branch and get the user name of the developer at the event object of the Lambda function
- O. Validate the user name with the IAM group to approve or deny the commit.

Answer: C

NEW QUESTION 45

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 Auto Scaling is required to solve the issue to ensure the application can handle high traffic loads.

Option D is invalid because there is no Auto Scaling template.

For more information on Auto Scaling, please refer to the below document link: from AWS

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

NEW QUESTION 49

A company is using tagging to allocate AWS costs. The company has Amazon EC2 instances that run in Auto Scaling groups. The Amazon Elastic Block Store (Amazon EBS) volumes that are attached to the EC2 instances are being created without the appropriate cost center tags. A DevOps engineer must ensure that the new EBS volumes are properly tagged.

What is the MOST efficient solution that meets this requirement?

- A. Create a lifecycle hook on the autoscaling:EC2_INSTANCE_TERMINATING instance state that attaches the cost center tags to the EBS volumes.
- B. Update the Auto Scaling group launch template to include the cost center tags for EBS volumes.
- C. Update the Auto Scaling group to include the cost center tag
- D. Set the PropagateAtLaunch property to true.
- E. Use Tag Editor to search for EBS volumes that are missing the tags and to add the cost center tags to the volumes.

Answer: B

NEW QUESTION 50

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 55

A company has several AWS accounts. The accounts are shared and used across multiple teams globally, primarily for Amazon EC2 instances. Each EC2 instance has tags for team, environment, and cost center to ensure accurate cost allocations.

How should a DevOps Engineer help the teams audit their costs and automate infrastructure cost optimization across multiple shared environments and accounts?

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

Answer: C

Explanation:

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

NEW QUESTION 58

A DevOps team needs to query information in application logs that are generated by an application running multiple Amazon EC2 instances deployed with AWS Elastic Beanstalk.

Instance log streaming to Amazon CloudWatch Logs was enabled on Elastic Beanstalk. Which approach would be the MOST cost-efficient?

- A. Use a CloudWatch Logs subscription to trigger an AWS Lambda function to send the log data to an Amazon Kinesis Data Firehouse stream that has an Amazon S3 bucket destinatio
- B. Use Amazon Athena to query the log data from the bucket.
- C. Use a CloudWatch Logs subscription to trigger an AWS Lambda function to send the log data to an Amazon Kinesis Data Firehouse stream that has an Amazon S3 bucket destinatio
- D. Use a new Amazon Redshift cluster and Amazon Redshift Spectrum to query the log data from the bucket.
- E. Use a CloudWatch Logs subscription to send the log data to an Amazon Kinesis Data Firehouse stream that has an Amazon S3 bucket destinatio
- F. Use Amazon Athena to query the log data from the bucket.
- G. Use a CloudWatch Logs subscription to send the log data to an Amazon Kinesis Data Firehouse stream that has an Amazon S3 bucket destinatio
- H. Use a new Amazon Redshift cluster and Amazon Redshift Spectrum to query the log data from the bucket.

Answer: C

Explanation:

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

NEW QUESTION 59

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 63

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

How can this solution be implemented to meet the security requirements using best practices?

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

Answer: C

Explanation:

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

NEW QUESTION 67

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

How can this be accomplished?

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

Answer: D

NEW QUESTION 72

A DevOps Engineer must implement monitoring for a workload running on Amazon EC2 and Amazon RDS MySQL. The monitoring must include: Application logs and operating system metrics for the Amazon EC2 instances Database logs and operating system metrics for the Amazon RDS database Which steps should the Engineer take?

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

Answer: B

NEW QUESTION 76

A company is using AWS CodeBuild, AWS CodeDeploy, and AWS CodePipeline to deploy applications automatically to an Amazon EC2 instance. A DevOps Engineer needs to perform a security assessment scan of the operating system on every application deployment to the environment. How should this be automated?

- A. Use Amazon CloudWatch Events to monitor for Auto Scaling event notifications of new instances and configure CloudWatch Events to trigger an Amazon Inspector scan.
- B. Use Amazon CloudWatch Events to monitor for AWS CodeDeploy notifications of a successful code deployment and configure CloudWatch Events to trigger an Amazon Inspector scan.
- C. Use Amazon CloudWatch Events to monitor for CodePipeline notifications of a successful code deployment and configure CloudWatch Events to trigger an AWS X-Ray scan.
- D. Use Amazon Inspector as a CodePipeline task after the successful use of CodeDeploy to deploy the code to the systems.

Answer: A

NEW QUESTION 77

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

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

Answer: A

NEW QUESTION 79

A DevOps engineer must ensure all IAM entity configurations across multiple AWS accounts in AWS Organizations are compliant with corporate IAM policies. Which combination of steps will accomplish this? (Select TWO.)

- A. Enable AWS Trusted Advisor in Organizations for all accounts to report on noncompliant IAM entities.
- B. Configure an AWS Config aggregator in the Organizations master account for all accounts
- C. Deploy AWS Config rules to the master account in Organizations that match corporate IAM policies.

- D. Apply an SCP in Organizations to ensure compliance of IAM entities.
- E. Deploy AWS Config rules to all accounts in Organizations that match the corporate IAM policies.

Answer: BE

NEW QUESTION 83

A DevOps Engineer is building a continuous deployment pipeline for a serverless application using AWS CodePipeline and AWS CodeBuild. The source, build, and test stages have been created with the deploy stage remaining. The company wants to reduce the risk of an unsuccessful deployment by deploying to a specified subset of customers and monitoring prior to a full release to all customers. How should the deploy stage be configured to meet these requirements?

- A. Use AWS CloudFormation to publish a new version on every stack update
- B. Then set up a CodePipeline approval action for a Developer to test and approve the new version
- C. Finally, use a CodePipeline invoke action to update an AWS Lambda function to use the production alias
- D. Use CodeBuild to use the AWS CLI to update the AWS Lambda function code, then publish a new version of the function and update the production alias to point to the new version of the function.
- E. Use AWS CloudFormation to define the serverless application and AWS CodeDeploy to deploy the AWS Lambda functions using DeploymentPreference: . Canary10Percent15Minutes
- F. Use AWS CloudFormation to publish a new version on every stack update
- G. Use the RoutingConfig property of the AWS::Lambda::Alias resource to update the traffic routing during the stack update.

Answer: C

Explanation:

[https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverle](https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless-application.html)

NEW QUESTION 86

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

Answer: B

NEW QUESTION 87

A company is adopting serverless computing and is migrating some of its existing applications to AWS Lambda. A DevOps engineer must come up with an automated deployment strategy using AWS CodePipeline that should include proper version controls, branching strategies, and rollback methods. Which combination of steps should the DevOps engineer follow when setting up the pipeline? (Select THREE)

- A. Use Amazon S3 as the source code repository
- B. Use AWS CodeCommit as the source code repository
- C. Use AWS CloudFormation to create an AWS Serverless Application Model (AWS SAM) template for deployment.
- D. Use AWS CodeBuild to create an AWS Serverless Application Model (AWS SAM) template for deployment
- E. Use AWS CloudFormation to deploy the application
- F. Use AWS CodeDeploy to deploy the application.

Answer: ABC

NEW QUESTION 92

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 96

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

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

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

Answer: A

NEW QUESTION 99

For auditing, analytics, and troubleshooting purposes, a DevOps Engineer for a data analytics application needs to collect all of the application and Linux system logs from the Amazon EC2 instances before termination. The company, on average, runs 10,000 instances in an Auto Scaling group. The company requires the ability to quickly find logs based on instance IDs and date ranges.

Which is the MOST cost-effective solution?

- A. Create an EC2 Instance-terminate Lifecycle Action on the group, write a termination script for pushing logs into Amazon S3, and trigger an AWS Lambda function based on S3 PUT to create a catalog of log files in an Amazon DynamoDB table with the primary key being Instance ID and sort key being Instance Termination Date.
- B. Create an EC2 Instance-terminate Lifecycle Action on the group, write a termination script for pushing logs into Amazon CloudWatch Logs, create a CloudWatch Events rule to trigger an AWS Lambda function to create a catalog of log files in an Amazon DynamoDB table with the primary key being Instance ID and sort key being Instance Termination Date.
- C. Create an EC2 Instance-terminate Lifecycle Action on the group, create an Amazon CloudWatch Events rule based on it to trigger an AWS Lambda function for storing the logs in Amazon S3, and create a catalog of log files in an Amazon DynamoDB table with the primary key being Instance ID and sort key being Instance Termination Date.
- D. Create an EC2 Instance-terminate Lifecycle Action on the group, push the logs into Amazon Kinesis Data Firehouse, and select Amazon ES as the destination for providing storage and search capability.

Answer: C

Explanation:

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

NEW QUESTION 100

A company updated the AWS CloudFormation template for a critical business application. The stack update process failed due to an error in the updated template, and CloudFormation automatically began the stack rollback process. Later, a DevOps engineer found the application was still unavailable, and that the stack was in the UPDATE_ROLLBACK_FAILED state.

Which combination of actions will allow the stack rollback to complete successfully? (Select TWO)

- A. Attach the AWSCloudFormationFullAccess IAM policy to the CloudFormation role.
- B. Automatically heal the stack resources using CloudFormation drift detection.
- C. Issue a ContinueUpdateRollback command from the CloudFormation console or AWS CLI.
- D. Manually update the resources to match the expectations of the stack.
- E. Update the existing CloudFormation stack using the original template.

Answer: AB

NEW QUESTION 103

A large enterprise is deploying a web application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon RDS Oracle DB instance and Amazon DynamoDB. There are separate environments for development, testing, and production.

What is the MOST secure and flexible way to obtain password credentials during deployment?

- A. Retrieve an access key from an AWS Systems Manager SecureString parameter to access AWS service.
- B. Retrieve the database credentials from a Systems Manager SecureString parameter.
- C. Launch the EC2 instances with an EC2 IAM role to access AWS service.
- D. Retrieve the database credentials from AWS Secrets Manager.
- E. Retrieve an access key from an AWS Systems Manager plaintext parameter to access AWS services. Retrieve the database credentials from a Systems Manager SecureString parameter.
- F. Launch the EC2 instances with an EC2 IAM role to access AWS service.
- G. Store the database passwords in an encrypted config file with the application artifacts.

Answer: B

Explanation:

<https://www.1strategy.com/blog/2019/02/28/aws-parameter-store-vs-aws-secrets-manager/>

NEW QUESTION 104

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

What is likely causing this issue?

- A. The two affected instances failed to fetch the new deployment.

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

Answer: D

NEW QUESTION 107

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 111

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

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

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

Answer: A

NEW QUESTION 112

A company has developed a static website hosted on an Amazon S3 bucket. The website is deployed using AWS CloudFormation. The CloudFormation template defines an S3 bucket and a custom resource that copies content into the bucket from a source location.

The company has decided that it needs to move the website to a new location, so the existing CloudFormation stack must be deleted and re-created. However, CloudFormation reports that the stack could not be deleted cleanly.

What is the MOST likely cause and how can the DevOps Engineer mitigate this problem for this and future versions of the website?

- A. Deletion has failed because the S3 bucket has an active website configuration.
- B. Modify the CloudFormation template to remove the Website Configuration property from the S3 bucket resource.
- C. Deletion has failed because the S3 bucket is not empty.
- D. Modify the custom resource's AWS Lambda function code to recursively empty the bucket when it is deleted.
- E. RequestType
- F. Deletion has failed because the custom resource does not define a deletion policy.
- G. Add a Deletion Policy property to the custom resource definition with a value of RemoveOnDeletion.
- H. Deletion has failed because the S3 bucket is not empty.
- I. Modify the S3 bucket resource in the CloudFormation template to add a Deletion Policy property with a value of Empty.

Answer: D

NEW QUESTION 115

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

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

- A. Use AWS Config to ensure all EC2 instances are managed by Amazon Inspector.
- B. Use AWS Config to ensure all EC2 instances are managed by AWS Systems Manager.
- C. Use AWS Systems Manager to install and manage Amazon Inspector, Systems Manager Patch Manager, and the Amazon CloudWatch agent on all instances.
- D. Use Amazon Inspector to install and manage AWS Systems Manager, Systems Manager Patch Manager, and the Amazon CloudWatch agent on all instances.
- E. Use AWS Systems Manager maintenance windows with Systems Manager Run Command to schedule Systems Manager Patch Manager task.
- F. Use the Amazon CloudWatch agent to schedule Amazon Inspector assessment runs.
- G. Use AWS Systems Manager maintenance windows with Systems Manager Run Command to schedule Systems Manager Patch Manager task.
- H. Use Amazon CloudWatch Events to schedule Amazon Inspector assessment runs.

Answer: BDE

NEW QUESTION 120

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

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

- A. Create an Amazon ECS cluster with the container instances in an Auto Scaling group.

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

Answer: A

Explanation:

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

NEW QUESTION 123

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

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

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

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

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

Answer: C

Explanation:

<https://aws.amazon.com/blogs/devops/using-aws-codepipeline-aws-codebuild-and-aws-lambda-for-serverless-au>

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

NEW QUESTION 128

A retail company has adopted AWS OpsWorks for managing its deployments. In the last three months, the company has discovered that some production instances have been restarting without reason. Upon inspection of the AWS CloudTrail logs, a DevOps Engineer determined that those instances were restarted by OpsWorks. The Engineer now wants automated email notifications whenever OpsWorks restarts an instance when the instance is deemed unhealthy or unable to communicate with the service endpoint.

How can the Engineer meet this requirement?

- A. Create a Chef recipe to place a cron to run a custom script within the Amazon EC2 instances that sends an email to the team by using Amazon SES if the OpsWorks agent detects an instance failure.
- B. Create an Amazon SNS topic and create a subscription for this topic that contains the destination email address
- C. Create an Amazon CloudWatch rule: specify as a source and specify auto-healing in the initiated_by detail
- D. Use the SNS topic as a target
- E. aws.opsworks
- F. Create an Amazon SNS topic and create a subscription for this topic that contains the destination email address
- G. Create an Amazon CloudWatch rule specify as a source and specify instance-replacement in the initiated_by detail
- H. Use the SNS topic as a target
- I. aws.opsworks
- J. Create a subscription for this topic that contains the email address
- K. Enable instance restart notifications within the OpsWorks layer and indicate the destination email address for the notification

Answer: B

NEW QUESTION 132

A company is migrating its public-facing software to AWS. The company plans to use Amazon EC2 to run application code and Amazon RDS to store all application data. The company wants to primarily use one Region with failover capabilities to a secondary Region and Amazon Route 53 to route traffic. The RPO is 2 hours and the RTO is 4 hours.

Which combination of steps should be used to meet these requirements while MINIMIZING cost? (Select THREE.)

- A. Create an AWS CloudFormation template to provision the application server and database instance in a single Region.
- B. Create an AWS CloudFormation template to provision the application tier of the application and a multi-Region database instance.
- C. Configure Amazon CloudWatch Events rules to run every hour
- D. Trigger AWS Lambda functions to create an RDS snapshot and copy it to the secondary Region.
- E. Configure Amazon CloudWatch Events rules to run every 3 hour
- F. Trigger AWS Lambda functions to create an RDS snapshot and copy it to the secondary Region.
- G. In the event of a failure, deploy a new AWS CloudFormation stack in a secondary region to provision the application resources and a new RDS instance using the copied snapshot and a Route 53 failover routing policy.
- H. In the event of a failure, deploy a new AWS CloudFormation stack in a secondary region to provision the application resources and a replica of the RDS database using the copied snapshot and a Route 53 latency-based routing policy.

Answer: CDE

NEW QUESTION 135

An ecommerce company is running an application on AWS. The company wants to create a standby disaster recovery solution in an additional Region that keeps the current application code. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an EC2 Auto

Scaling group across multiple Availability Zones. The database layer is hosted on an Amazon RDS MySQL Multi-AZ DB instance. Amazon Route 53 DNS records point to the ALB.

Which combination of actions will meet these requirements with the LOWEST cost? (Select THREE.)

- A. Configure a failover routing policy for the application DNS entry.
- B. Configure a geolocation routing policy for the application DNS entry.
- C. Create a cross-Region RDS read replica in the new standby Region.
- D. Migrate the database layer to Amazon DynamoDB and enable global replication to the new standby Region.
- E. Provision the ALB and Auto Scaling group in the new standby Region and set the desired capacity to match the active Region.
- F. Provision the ALB and Auto Scaling group in the new standby Region and set the desired capacity to 1.

Answer: AEF

NEW QUESTION 139

A Developer is designing a continuous deployment workflow for a new Development team to facilitate the process for source code promotion in AWS. Developers would like to store and promote code for deployment from development to production while maintaining the ability to roll back that deployment if it fails.

Which design will incur the LEAST amount of downtime?

- A. Create one repository in AWS CodeCommit
- B. Create a development branch to hold merged change
- C. Use AWS CodeBuild to build and test the code stored in the development branch triggered on a new commit
- D. Merge to the master and deploy to production by using AWS CodeDeploy for a blue/green deployment.
- E. Create one repository for each Developer in AWS CodeCommit and another repository to hold the production code
- F. Use AWS CodeBuild to merge development and production repositories, and deploy to production by using AWS CodeDeploy for a blue/green deployment.
- G. Create one repository for development code in AWS CodeCommit and another repository to hold the production code
- H. Use AWS CodeBuild to merge development and production repositories, and deploy to production by using AWS CodeDeploy for a blue/green deployment.
- I. Create a shared Amazon S3 bucket for the Development team to store their code
- J. Set up an Amazon CloudWatch Events rule to trigger an AWS Lambda function that deploys the code to production by using AWS CodeDeploy for a blue/green deployment.

Answer: A

NEW QUESTION 144

You currently have the following setup in AWS

- 1) An Elastic Load Balancer
- 2) Auto Scaling Group which launches EC2 Instances
- 3) AMIs with your code pre-installed

You want to deploy the updates of your app to only a certain number of users. You want to have a cost-effective solution. You should also be able to revert back quickly. Which of the below solutions is the most feasible one?

- A. Create a second ELB, and a new Auto Scaling Group assigned a new Launch Configuration
- B. Create a new AMI with the updated app
- C. Use Route53 Weighted Round Robin records to adjust the proportion of traffic hitting the two ELBs.
- D. Create new AMIs with the new app
- E. Then use the new EC2 instances in half proportion to the older instances.
- F. Redeploy with AWS Elastic Beanstalk and Elastic Beanstalk version
- G. Use Route 53 Weighted Round Robin records to adjust the proportion of traffic hitting the two ELBs
- H. Create a full second stack of instances, cut the DNS over to the new stack of instances, and change the DNS back if a rollback is needed.

Answer: A

Explanation:

The Weighted Routing policy of Route53 can be used to direct a proportion of traffic to your application. The best option is to create a second CLB, attach the new AutoScaling Group and then use Route53 to divert the traffic.

Option B is wrong because just having EC2 instances running with the new code will not help.

Option C is wrong because Elastic beanstalk is good for development environments, and also there is no mention of having 2 environments where environment URLs can be swapped.

Option D is wrong because you still need Route53 to split the traffic.

For more information on Route53 routing policies, please refer to the below link:

➤ <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

NEW QUESTION 148

Two teams are working together on different portions of an architecture and are using AWS CloudFormation to manage their resources. One team administers operating system-level updates and patches, while the other team manages application-level dependencies and updates. The Application team must take the most recent AMI when creating new instances and deploying the application.

What is the MOST scalable method for linking these two teams and processes?

- A. The Operating System team uses CloudFormation to create new versions of their AMIs and lists the Amazon Resource names (ARNs) of the AMIs in an encrypted Amazon S3 object as part of the stack output section
- B. The Application team uses a cross-stack reference to load the encrypted S3 object and obtain the most recent AMI ARNs.
- C. The Operating System team uses CloudFormation stack to create an AWS CodePipeline pipeline that builds new AMIs, then places the latest AMI ARNs in an encrypted Amazon S3 object as part of the pipeline output
- D. The Application team uses a cross-stack reference within their own CloudFormation template to get that S3 object location and obtain the most recent AMI ARNs to use when deploying their application.
- E. The Operating System team uses CloudFormation stack to create an AWS CodePipeline pipeline that builds new AMI
- F. The team then places the AMI ARNs as parameters in AWS Systems Manager Parameter Store as part of the pipeline output
- G. The Application team specifies a parameter of type ssm in their CloudFormation stack to obtain the most recent AMI ARN from the Parameter Store.
- H. The Operating System team maintains a nested stack that includes both the operating system and Application team template

I. The Operating System team uses a stack update to deploy updates to the application stack whenever the Application team changes the application code.

Answer: B

NEW QUESTION 150

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 154

Management has reported an increase in the monthly bill from Amazon Web Services, and they are extremely concerned with this increased cost. Management has asked you to determine the exact cause of this increase. After reviewing the billing report, you notice an increase in the data transfer cost. How can you provide management with a better insight into data transfer use?

- A. Update your Amazon CloudWatch metrics to use five-second granularity, which will give better detailed metrics that can be combined with your billing data to pinpoint anomalies.
- B. Use Amazon CloudWatch Logs to run a map-reduce on your logs to determine high usage and data transfer.
- C. Deliver custom metrics to Amazon CloudWatch per application that breaks down application data transfer into multiple, more specific data points.
- D. Using Amazon CloudWatch metrics, pull your Elastic Load Balancing outbound data transfer metrics monthly, and include them with your billing report to show which application is causing higher bandwidth usage.

Answer: C

Explanation:

You can publish your own metrics to CloudWatch using the AWS CLI or an API. You can view statistical graphs of your published metrics with the AWS Management Console.

CloudWatch stores data about a metric as a series of data points. Each data point has an associated time stamp. You can even publish an aggregated set of data points called a statistic set.

If you have custom metrics specific to your application, you can give a breakdown to the management on the exact issue.

Option A won't be sufficient to provide better insights.

Option B is an overhead when you can make the application publish custom metrics. Option D is invalid because just the ELB metrics will not give the entire picture.

For more information on custom metrics, please refer to the below document link: from AWS

➤ <http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html>

NEW QUESTION 159

A company maintains a stateless web application that is experiencing inconsistent traffic. The company uses AWS CloudFormation to deploy the application. The application runs on Amazon EC2 On-Demand Instances behind an Application Load Balancer (ALB). The instances run across multiple Availability Zones. The company wants to include the use of Spot Instances while continuing to use a small number of On-Demand Instances to ensure that the application remains highly available.

What is the MOST cost-effective solution that meets these requirements?

- A. Add a Spot block resource to the AWS CloudFormation template
- B. Use the diversified allocation strategy with step scaling behind the ALB.
- C. Add a Spot block resource to the AWS CloudFormation template
- D. Use the lowest-price allocation strategy with target tracking scaling behind the ALB.
- E. Add a Spot Fleet resource to the AWS CloudFormation template
- F. Use the capacity-optimized allocation strategy with step scaling behind the ALB.
- G. Add a Spot Fleet resource to the AWS CloudFormation template
- H. Use the diversified allocation strategy with scheduled scaling behind the ALB

Answer: C

NEW QUESTION 161

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 164

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 applicatio
- B. Use CodeDeploy to deploy the application binary to an AWS Lambda function for testin
- C. Use a third-party library on AWS Lambda to simulate the device platfor
- D. Allow a Lambda role to upload to the production Amazon S3 bucke
- E. Make the binary publicly accessibl
- F. Trigger notifications using Amazon SNS.
- G. Use GitHub as the code source and AWS Lambda to compile and package the applicatio
- H. Use another Lambda function to run unit tests and deliver the application binary to a development bucke
- I. Use the binary from the development bucket and install the application on a personal device for testin
- J. Deliver the binary to the production bucket after approva
- K. Trigger notifications using Amazon SNS.
- L. Use an Amazon S3 bucket as the code source and AWS CodeBuild to compile and package the applicatio
- M. Use AWS CodeDeploy to deploy the application binary to a device farm for testin
- N. Deliver the binary to the production S3 bucke
- O. Use an S3 bucket policy to allow public read on the productionS3 bucke
- 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 168

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

- A. Add the instance to an EC2 Auto Scaling group with the minimum, maximum, and desired capacity set to 1.
- B. Add the instance to an EC2 Auto Scaling group with a lifecycle hook to detach the EBS volume when the EC2 instance shuts down or terminates.
- C. Create an Amazon CloudWatch alarm for the StatusCheckFailed_System metric and select the EC2 action to recover the instance
- D. Create an Amazon CloudWatch alarm for the StatusCheckFailedinstance metric and select the EC2 action to reboot the instance

Answer: A

NEW QUESTION 173

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 CloudWatc
- B. Then, set up and share a live CloudWatch dashboar
- 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 bucke
- E. Createan 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 inspecto
- 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-resour>

NEW QUESTION 176

A new zero-day vulnerability was found in OpenSSL requiring the immediate patching of a production web fleet running on Amazon Linux. Currently, OS updates are performed manually on a monthly basis and deployed using updates to the production Auto Scaling Group's launch configuration.
Which method should a DevOps Engineer use to update packages in-place without downtime?

- A. Use AWS CodePipline and AWS CodeBuild to generate new copies of these packages, and update the Auto Scaling group's launch configuration.
- B. Use AWS Inspector to run "yum upgrade" on all running production instances, and manually update the AMI for the next maintenance window.
- C. Use Amazon EC2 Run Command to issue a package update command to all running production instances, and update the AMI for future deployments.
- D. Define a new AWS OpsWorks layer to match the running production instances, and use a recipe to issue a package update command to all running production instances.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/ec2-run-command-is-now-a-cloudwatch-events-target/>

" EC2 Run Command is part of EC2 Systems Manager. It allows you to operate on collections of EC2 instances and on-premises servers reliably and at scale, in a controlled and selective fashion. You can run scripts, install software, collect metrics and log files, manage patches, and much more, on both Windows and Linux."

NEW QUESTION 178

A defect was discovered in production and a new sprint item has been created for deploying a hotfix. However, any code change must go through the following steps before going into production:

*Scan the code for security breaches, such as password and access key leaks. Run the code through extensive, long running unit tests.

Which source control strategy should a DevOps Engineer use in combination with AWS CodePipeline to complete this process?

- A. Create a hotfix tag on the last commit of the master branch
- B. Trigger the development pipeline from the hotfix ta
- C. Use AWS CodeDeploy with Amazon ECS to do a content scan and run unit test
- D. Add a manual approval stage that merges the hotfix tag into the master branch.
- E. Create a hotfix branch from the master branch
- F. Triger the development pipeline from the hotfix branch.Use AWS CodeBuild to do a content scan and run unit test
- G. Add a manual approval stage that merges the hotfix branch into the master branch.
- H. Create a hotfix branch from the master branch
- I. Triger the development pipeline from the hotfix branch.Use AWS Lambda to do a content scan and run unit test
- J. Add a manual approval stage that merges the hotfix branch into the master branch.
- K. Create a hotfix branch from the master branch
- L. Create a separate source stage for the hotfix branch in the production pipelin
- M. Trigger the pipeline from the hotfix branch
- N. Use AWS Lambda to do a content scan and use AWS CodeBuild to run unit test
- O. Add a manual approval stage that merges the hotfix branch into the master branch.

Answer: B

NEW QUESTION 179

An ecommerce company is receiving reports that its order history page is experiencing delays in reflecting the processing status of orders. The order processing system consists of an AWS Lambda function using reserved concurrency. The Lambda function processes order messages from an Amazon SQS queue and inserts processed orders into an Amazon DynamoDB table. The DynamoDB table has Auto Scaling enabled for read and write capacity.

Which actions will diagnose and resolve the delay? (Select TWO.)

- A. Check the ApproximateAgeOfOldestMessage metric for the SQS queue and increase the Lambda function concurrency limit.
- B. Check the ApproximateAgeOfOldestMessage metric for the SQS queue and configure a redrive policy on the SQS queue.
- C. Check the NumberOfMessagesSent metric for the SQS queue and increase the SQS queue visibility timeout.
- D. Check the ThrottledWriteRequests metric for the DynamoDB table and increase the maximum write capacity units for the table's Auto Scaling policy.
- E. Check the Throttles metric for the Lambda function and increase the Lambda function timeout.

Answer: AB

NEW QUESTION 181

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

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

Answer: B

NEW QUESTION 184

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 186

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-DevOps-Engineer-Professional Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-DevOps-Engineer-Professional Product From:

<https://www.2passeasy.com/dumps/AWS-Certified-DevOps-Engineer-Professional/>

Money Back Guarantee

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

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