



Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

A DevOps Engineer must track the health of a stateless RESTful service sitting behind a Classic Load Balancer. The deployment of new application revisions is through a CI/CD pipeline. If the service's latency increases beyond a defined threshold, deployment should be stopped until the service has recovered. Which of the following methods allow for the QUICKEST detection time?

- A. Use Amazon CloudWatch metrics provided by Elastic Load Balancing to calculate average latency. Alarm and stop deployment when latency increases beyond the defined threshold.
- B. Use AWS Lambda and Elastic Load Balancing access logs to detect average latency.
- C. Alarm and stop deployment when latency increases beyond the defined threshold.
- D. Use AWS CodeDeploy's Minimum Healthy Hosts setting to define thresholds for rolling back deployment.
- E. If these thresholds are breached, roll back the deployment.
- F. Use Metric Filters to parse application logs in Amazon CloudWatch Log.
- G. Create a filter for latency. Alarm and stop deployment when latency increases beyond the defined threshold.

Answer: A

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-cloudwatch-metrics.html>
<https://docs.aws.amazon.com/codedeploy/latest/userguide/deployments-stop.html>

NEW QUESTION 2

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

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

Answer: D

NEW QUESTION 3

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

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

Answer: C

Explanation:

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

NEW QUESTION 4

A Security team requires all Amazon EBS volumes that are attached to an Amazon EC2 instance to have AWS Key Management Service (AWS KMS) encryption enabled. If encryption is not enabled, the company's policy requires the EBS volume to be detached and deleted. A DevOps Engineer must automate the detection and deletion of unencrypted EBS volumes. Which method should the Engineer use to accomplish this with the LEAST operational effort?

- A. Create an Amazon CloudWatch Events rule that invokes an AWS Lambda function when an EBS volume is created.
- B. The Lambda function checks the EBS volume for encryption.
- C. If encryption is not enabled and the volume is attached to an instance, the function deletes the volume.
- D. Create an AWS Lambda function to describe all EBS volumes in the region and identify volumes that are attached to an EC2 instance without encryption enabled.
- E. The function then deletes all non-compliant volumes.
- F. The AWS Lambda function is invoked every 5 minutes by an Amazon CloudWatch Events scheduled rule.
- G. Create a rule in AWS Config to check for unencrypted and attached EBS volumes.
- H. Subscribe an AWS Lambda function to the Amazon SNS topic that AWS Config sends change notifications to.
- I. The Lambda function checks the change notification and deletes any EBS volumes that are non-compliant.
- J. Launch an EC2 instance with an IAM role that has permissions to describe and delete volumes.

- K. Run ascript on the EC2 instance every 5 minutes to describe all EBS volumes in all regions and identify volumes that are attached without encryption enable
L. The script then deletes those volumes.

Answer: B

NEW QUESTION 5

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

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

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

Answer: C

Explanation:

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

https://docs.aws.amazon.com/codedeploy/latest/APIReference/API_BlueGreenDeploymentConfiguration.html

NEW QUESTION 6

You have an application consisting of a stateless web server tier running on Amazon EC2 instances behind load balancer, and are using Amazon RDS with read replicas. Which of the following methods should you use to implement a self-healing and cost-effective architecture? Choose 2 answers from the options given below

- A. Set up a third-party monitoring solution on a cluster of Amazon EC2 instances in order to emit custom Cloud Watch metrics to trigger the termination of unhealthy Amazon EC2 instances.
- B. Set up scripts on each Amazon EC2 instance to frequently send ICMP pings to the load balancer in order to determine which instance is unhealthy and replace it.
- C. Set up an Auto Scaling group for the web server tier along with an Auto Scaling policy that uses the Amazon RDS DB CPU utilization Cloud Watch metric to scale the instances.
- D. Set up an Auto Scaling group for the web server tier along with an Auto Scaling policy that uses the Amazon EC2 CPU utilization CloudWatch metric to scale the instances.
- E. Use a larger Amazon EC2 instance type for the web server tier and a larger DB instance type for the data storage layer to ensure that they don't become unhealthy.
- F. Set up an Auto Scaling group for the database tier along with an Auto Scaling policy that uses the Amazon RDS read replica lag CloudWatch metric to scale out the Amazon RDS read replicas.
- G. Use an Amazon RDS Multi-AZ deployment.

Answer: DG

Explanation:

The scaling of EC2 Instances in the Autoscaling group is normally done with the metric of the CPU utilization of the current instances in the Autoscaling group
For more information on scaling in your Autoscaling Group, please refer to the below link:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scaling-simple-step.html>

Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure, Amazon RDS performs an automatic failover to the standby (or to a read replica in the case of Amazon Aurora), so that you can resume database operations as soon as the failover is complete. For more information on RDS

Multi-AZ please refer to the below link:

- <https://aws.amazon.com/rds/details/multi-az/>

Option A is invalid because if you already have in-built metrics from Cloudwatch, why would you want to spend more in using a third-party monitoring solution.

Option B is invalid because health checks are already a feature of AWS CLB

Option C is invalid because the database CPU usage should not be used to scale the web tier.

Option D is invalid because increasing the instance size does not always guarantee that the solution will not become unhealthy.

Option F is invalid because increasing Read-Replica's will not suffice for write operations if the primary DB fails.

NEW QUESTION 7

A company is using an AWS CloudFormation template to deploy web applications. The template requires that manual changes be made for each of the three major environments: production, staging, and development. The current sprint includes the new implementation and configuration of AWS CodePipeline for automated deployments.

What changes should the DevOps Engineer make to ensure that the CloudFormation template is reusable across multiple pipelines?

- A. Use a CloudFormation custom resource to query the status of the CodePipeline to determine which environment is launched
- B. Dynamically alter the launch configuration of the Amazon EC2 instances.
- C. Set up a CodePipeline pipeline for each environment to use input parameter
- D. Use CloudFormation mappings to switch associated UserData for the Amazon EC2 instances to match the environment being launched.
- E. Set up a CodePipeline pipeline that has multiple stages, one for each development environment
- F. Use AWS Lambda functions to trigger CloudFormation deployments to dynamically alter the UserData of the Amazon EC2 instances launched in each environment.
- G. Use CloudFormation input parameters to dynamically alter the LaunchConfiguration and UserData sections of each Amazon EC2 instance every time the CloudFormation stack is updated.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/continuous-delivery-codepipeline-paramet>

NEW QUESTION 8

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

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

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

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

Answer: C

NEW QUESTION 9

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

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

Answer: C

Explanation:

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

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

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

There is no need to create a whole CloudFormation

template for this.

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

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

➤ http://docs.aws.amazon.com/autoscaling/latest/userguide/l_launchConfiguration.html

NEW QUESTION 10

A company is building a web and mobile application that uses a serverless architecture powered by AWS Lambda and Amazon API Gateway. The company wants to fully automate the backend Lambda deployment based on code that is pushed to the appropriate environment branch in an AWS CodeCommit repository.

The deployment must have the following:

*Separate environment pipelines for testing and production.

*Automatic deployment that occurs for test environments only. Which steps should be taken to meet these requirements?

- A. Configure a new AWS CodePipeline service
- B. Create a CodeCommit repository for each environment. Set up CodePipeline to retrieve the source code from the appropriate repository
- C. Set up a deployment step to deploy the Lambda functions with AWS CloudFormation.
- D. Create two AWS CodePipeline configurations for test and production environment
- E. Configure the production pipeline to have a manual approval step
- F. Create a CodeCommit repository for each environment
- G. Set up each CodePipeline to retrieve the source code from the appropriate repository
- H. Set up the deployment step to deploy the Lambda functions with AWS CloudFormation.
- I. Create two AWS CodePipeline configurations for test and production environment
- J. Configure the production pipeline to have a manual approval step
- K. Create one CodeCommit repository with a branch for each environment

- L. Set up each CodePipeline to retrieve the source code from the appropriate branch in the repository
- M. Set up the deployment step to deploy the Lambda functions with AWS CloudFormation.
- N. Create an AWS CodeBuild configuration for test and production environment
- O. Configure the production pipeline to have a manual approval step
- P. Create one CodeCommit repository with a branch for each environment
- Q. Push the Lambda function code to an Amazon S3 bucket
- R. Set up the deployment step to deploy the Lambda functions from the S3 bucket.

Answer: B

NEW QUESTION 10

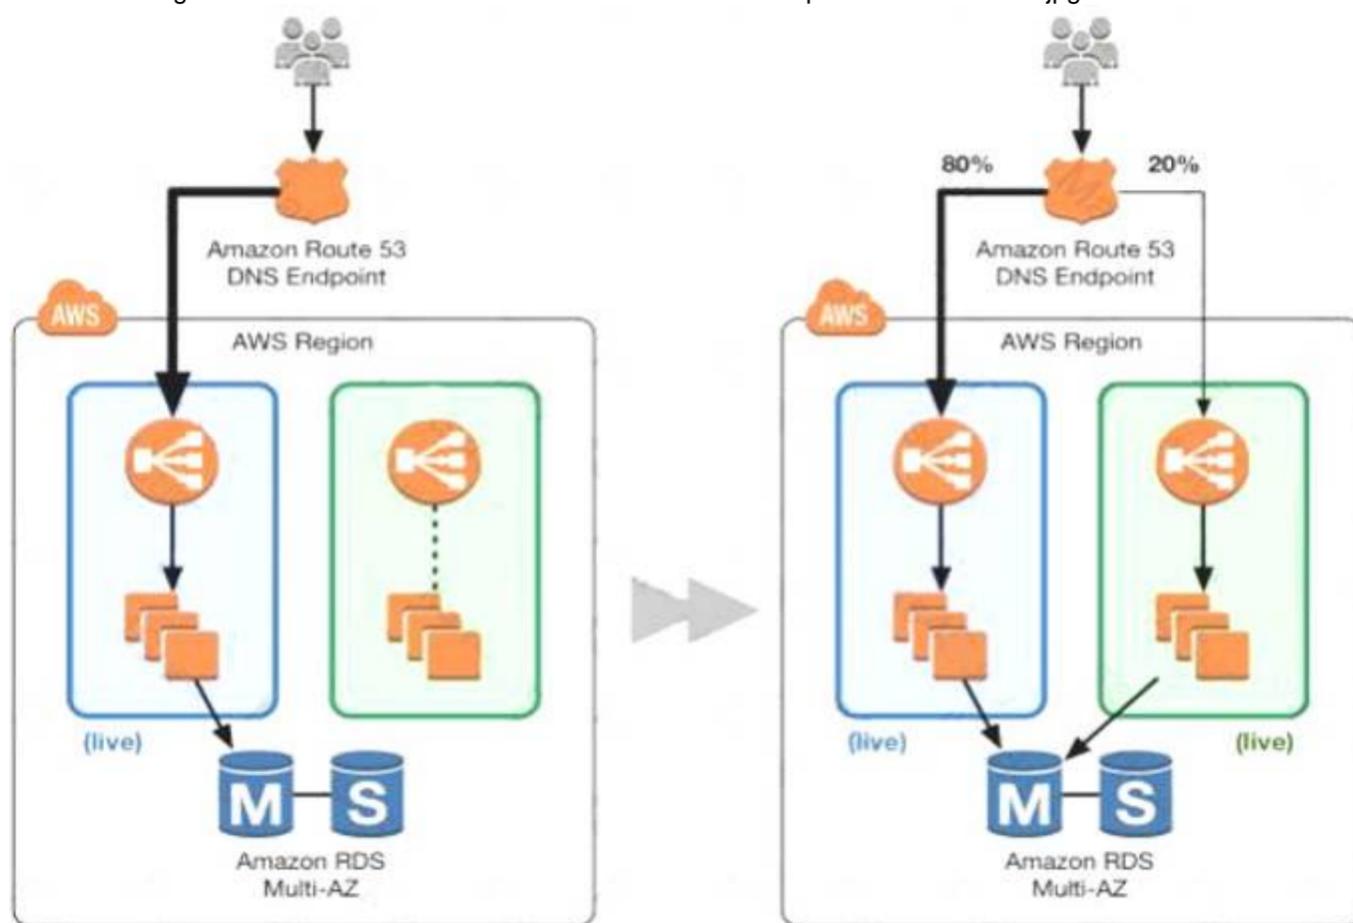
Your application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment?

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

Answer: C

Explanation:

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



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

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

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

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

For more information on blue green deployments, please refer to the below document link: [from AWS >](https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf)

https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 12

A consulting company was hired to assess security vulnerabilities within a client company's application and propose a plan to remediate all identified issues. The architecture is identified as follows: Amazon S3 storage for content, an Auto Scaling group of Amazon EC2 instances behind an Elastic Load Balancer with attached Amazon EBS storage, and an Amazon RDS MySQL database. There are also several AWS Lambda functions that communicate directly with the RDS database using connection string statements in the code.

The consultants identified the top security threat as follows: the application is not meeting its requirement to have encryption at rest.

What solution will address this issue with the LEAST operational overhead and will provide monitoring for potential future violations?

- A. Enable SSE encryption on the S3 buckets and RDS databas
- B. Enable OS-based encryption of data on EBS volume
- C. Configure Amazon Inspector agents on EC2 instances to report on insecure encryption cipher
- D. Set up AWS Config rules to periodically check for non-encrypted S3 objects.
- E. Configure the application to encrypt each file prior to storing on Amazon S3. Enable OS-based encryption of data on EBS volume
- F. Encrypt data on write to RD

- G. Run cron jobs on each instance to check for encrypted data and notify via Amazon SNS
- H. Use S3 Events to call an AWS Lambda function and verify if the file is encrypted.
- I. Enable Secure Sockets Layer (SSL) on the load balancer, ensure that AWS Lambda is using SSL to communicate to the RDS database, and enable S3 encryption
- J. Configure the application to force SSL for incoming connections and configure RDS to only grant access if the session is encrypted
- K. Configure Amazon Inspector agents on EC2 instances to report on insecure encryption ciphers.
- L. Enable SSE encryption on the S3 buckets, EBS volumes, and the RDS databases
- M. Store RDS credentials in EC2 Parameter Store
- N. Enable a policy on the S3 bucket to deny unencrypted put
- O. Set up AWS Config rules to periodically check for non-encrypted S3 objects and EBS volumes, and to ensure that RDS storage is encrypted.

Answer: D

NEW QUESTION 15

A DevOps Engineer must create a Linux AMI in an automated fashion. The newly created AMI identification must be stored in a location where other build pipelines can access the new identification programmatically
What is the MOST cost-effective way to do this?

- A. Build a pipeline in AWS CodePipeline to download and save the latest operating system Open Virtualization Format (OVF) image to an Amazon S3 bucket, then customize the image using the guestfish utility
- B. Use the virtual machine (VM) import command to convert the OVF to an AMI, and store the AMI identification output as an AWS Systems Manager parameter.
- C. Create an AWS Systems Manager automation document with values instructing how the image should be created
- D. Then build a pipeline in AWS CodePipeline to execute the automation document to build the AMI when triggered
- E. Store the AMI identification output as a Systems Manager parameter.
- F. Build a pipeline in AWS CodePipeline to take a snapshot of an Amazon EC2 instance running the latest version of the application
- G. Then start a new EC2 instance from the snapshot and update the running instance using an AWS Lambda function
- H. Take a snapshot of the updated instance, then convert it to an AMI
- I. Store the AMI identification output in an Amazon DynamoDB table.
- J. Launch an Amazon EC2 instance and install Packer
- K. Then configure a Packer build with values defining how the image should be created
- L. Build a Jenkins pipeline to invoke the Packer build when triggered to build an AMI
- M. Store the AMI identification output in an Amazon DynamoDB table.

Answer: D

NEW QUESTION 19

A company has deployed several applications globally. Recently, Security Auditors found that few Amazon EC2 instances were launched without Amazon EBS disk encryption. The Auditors have requested a report detailing all EBS volumes that were not encrypted in multiple AWS accounts and regions. They also want to be notified whenever this occurs in the future.
How can this be automated with the LEAST amount of operational overhead?

- A. Create an AWS Lambda function to set up an AWS Config rule on all the target accounts
- B. Use AWS Config aggregators to collect data from multiple accounts and regions
- C. Export the aggregated report to an Amazon S3 bucket and use Amazon SNS to deliver the notifications.
- D. Set up AWS CloudTrail to deliver all events to an Amazon S3 bucket in a centralized account
- E. Use the S3 event notification feature to invoke an AWS Lambda function to parse AWS CloudTrail logs whenever logs are delivered to the S3 bucket
- F. Publish the output to an Amazon SNS topic using the same Lambda function.
- G. Create an AWS CloudFormation template that adds an AWS Config managed rule for EBS encryption. Use a CloudFormation stack set to deploy the template across all accounts and regions
- H. Store consolidated evaluation results from config rules in Amazon S3. Send a notification using Amazon SNS when non-compliant resources are detected.
- I. Using AWS CLI, run a script periodically that invokes the `aws ec2 describe-volumes` query with a JMESPATH query filter
- J. Then, write the output to an Amazon S3 bucket
- K. Set up an S3 event notification to send events using Amazon SNS when new data is written to the S3 bucket.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/aws-config-update-aggregate-compliance-data-across-accounts-regions/>
<https://docs.aws.amazon.com/config/latest/developerguide/aws-config-managed-rules-cloudformation-templates>

NEW QUESTION 24

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

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

Answer: D

NEW QUESTION 28

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

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

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

Answer: D

NEW QUESTION 30

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 34

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

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

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

Answer: B

NEW QUESTION 39

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

Thread

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 42

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

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

Answer: D

NEW QUESTION 47

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 48

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 50

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

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

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

Answer: BE

NEW QUESTION 54

A company has built a web service that runs on Amazon EC2 instances behind an Application Load Balancer (ALB) the company has deployed the application in us-east-1 Amazon Route 53 provides an external DNS that routes traffic from example.com to the application, created with appropriate health checks.

The company has deployed a second environment for the application in eu-west-1 the company wants traffic to be routed to whichever environment results in the best response time for each user. If there is an outage in one Region, traffic should be directed to the other environment.

Which configuration will achieve this requirements?

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

Answer: B

NEW QUESTION 59

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

Which strategy should be used to meet these requirements?

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

Answer: B

NEW QUESTION 60

A DevOps engineer notices that all Amazon EC2 instances running behind an Application Load Balancer in an Auto Scaling group are failing to respond to user requests. The EC2 instances are also failing target group HTTP health checks.

Upon inspection, the engineer notices the application process was not running in any EC2 instances. There are a significant number of out of memory messages in the system logs. The engineer needs to improve the resilience of the application to cope with a potential application memory leak. Monitoring and notifications should be enabled to alert when there is an issue.

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

- A. Change the Auto Scaling configuration to replace the instances when they fail the load balancer's health checks.
- B. Change the target group health check HealthCheckIntervalSeconds parameter to reduce the interval between health checks.
- C. Change the target group health checks from HTTP to TCP to check if the port where the application is listening is reachable.
- D. Enable the available memory consumption metric within the Amazon CloudWatch dashboard for the entire Auto Scaling group
- E. Create an alarm when the memory utilization is high

- F. Associate an
- G. Amazon SNS topic to the alarm to receive notifications when the alarm goes off.
- H. Use the Amazon CloudWatch agent to collect the memory utilization of the EC2 instances in the Auto Scaling group.
- I. Create an alarm when the memory utilization is high and associate an Amazon SNS topic to receive a notification.

Answer: BE

NEW QUESTION 63

A company plans to stop using Amazon EC2 key pairs for SSH access, and instead plans to use AWS Systems Manager Session Manager. To further enhance security, access to Session Manager must take place over a private network only. Which combinations of actions will accomplish this? (Select TWO.)

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

Answer: BC

NEW QUESTION 65

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

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

Answer: AC

NEW QUESTION 69

A company has multiple child accounts that are part of an organization in AWS Organizations. The security team needs to review every Amazon EC2 security group and their inbound and outbound rules. The security team wants to programmatically retrieve this information from the child accounts using an AWS Lambda function in the master account of the organization. Which combination of access changes will meet these requirements? (Select THREE.)

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

Answer: ADF

NEW QUESTION 73

An application is being deployed with two Amazon EC2 Auto Scaling groups, each configured with an Application Load Balancer. The application is deployed to one of the Auto Scaling groups and an Amazon Route 53 alias record is pointed to the Application Load Balancer of the last deployed Auto Scaling group. Deployments alternate between the two Auto Scaling groups. Home security devices are making requests into the application. The Development team notes that new requests are coming into the old stack days after the deployment. The issue is caused by devices that are not observing the Time to Live (TTL) setting on the Amazon Route 53 alias record. What steps should the DevOps Engineer take to address the issue with requests coming to the old stacks, while creating minimal additional resources?

- A. Create a fleet of Amazon EC2 instances running HAProxy behind an Application Load Balancer
- B. The HAProxy instances will proxy the requests to one of the existing Auto Scaling group
- C. After a deployment the HAProxy instances are updated to send requests to the newly deployed Auto Scaling group.
- D. Reduce the application to one Application Load Balancer
- E. Create two target groups named Blue and Green
- F. Create a rule on the Application Load Balancer pointed to a single target group
- G. Add logic to the deployment to update the Application Load Balancer rule to the target group of the newly deployed Auto Scaling group.
- H. Move the application to an AWS Elastic Beanstalk application with two environments
- I. Perform new deployments on the non-live environment
- J. After a deployment, perform an Elastic Beanstalk CNAME swap to make the newly deployed environment the live environment.
- K. Create an Amazon CloudFront distribution
- L. Set the two existing Application Load Balancers as origins on the distribution
- M. After a deployment, update the CloudFront distribution behavior to send requests to the newly deployed Auto Scaling group.

Answer: B

NEW QUESTION 74

A company is migrating an application to AWS that runs on a single Amazon EC2 instance. Because of licensing limitations, the application does not support horizontal scaling. The application will be using Amazon Aurora for its database. How can the DevOps Engineer architect automated healing to automatically recover from EC2 and Aurora failures, in addition to recovering across Availability Zones (AZs), in the MOST cost-effective manner?

- A. Create an EC2 Auto Scaling group with a minimum and maximum instance count of 1, and have it span across AZ
- B. Use a single-node Aurora instance.
- C. Create an EC2 instance and enable instance recover
- D. Create an Aurora database with a read replica in a second AZ, and promote it to a primary database instance if the primary database instance fails.
- E. Create an Amazon CloudWatch Events rule to trigger an AWS Lambda function to start a new EC2 instance in an available AZ when the instance status reaches a failure state
- F. Create an Aurora database with a read replica in a second AZ, and promote it to a primary database instance when the primary database instance fails.
- G. Assign an Elastic IP address on the instance
- H. Create a second EC2 instance in a second AZ
- I. Create an Amazon CloudWatch Events rule to trigger an AWS Lambda function to move the Elastic IP address to the second instance when the first instance fails
- J. Use a single-node Aurora instance.

Answer: C

NEW QUESTION 77

An application is running on Amazon EC2. It has an attached IAM role that is receiving an AccessDenied error while trying to access a SecureString parameter resource in the AWS Systems Manager Parameter Store. The SecureString parameter is encrypted with a customer-managed Customer Master Key (CMK). What steps should the DevOps Engineer take to grant access to the role while granting least privilege? (Select three.)

- A. Set ssm:GetParameter for the parameter resource in the instance role's IAM policy.
- B. Set kms:Decrypt for the instance role in the customer-managed CMK policy.
- C. Set kms:Decrypt for the customer-managed CMK resource in the role's IAM policy.
- D. Set ssm:DecryptParameter for the parameter resource in the instance role IAM policy.
- E. Set kms:GenerateDataKey for the user on the AWS managed SSM KMS key.
- F. Set kms:Decrypt for the parameter resource in the customer-managed CMK policy.

Answer: ABC

NEW QUESTION 81

A company is implementing an Amazon ECS cluster to run its workload. The company architecture will run multiple ECS services on the cluster, with an Application Load Balancer on the front end, using multiple target groups to route traffic. The Application Development team has been struggling to collect logs that must be collected and sent to an Amazon S3 bucket for near-real time analysis

What must the DevOps Engineer configure in the deployment to meet these requirements? (Select THREE)

- A. Install the Amazon CloudWatch Logs logging agent on the ECS instance
- B. Change the logging driver in the ECS task definition to 'awslogs'.
- C. Download the Amazon CloudWatch Logs container instance from AWS and configure it as a task. Update the application service definitions to include the logging task.
- D. Use Amazon CloudWatch Events to schedule an AWS Lambda function that will run every 60 seconds running the create-export -task CloudWatch Logs command, then point the output to the logging S3 bucket.
- E. Enable access logging on the Application Load Balancer, then point it directly to the S3 logging bucket.
- F. Enable access logging on the target groups that are used by the ECS services, then point it directly to the S3 logging bucket.
- G. Create an Amazon Kinesis Data Firehose with a destination of the S3 logging bucket, then create an Amazon CloudWatch Logs subscription filter for Kinesis

Answer: BDF

NEW QUESTION 85

A company must ensure consistent behavior of an application running on Amazon Linux in its corporate ecosystem before moving into AWS. The company has an existing automated server build system using VMware. The goal is to demonstrate the functionality of the application and its prerequisites on the new target operating system.

The DevOps Engineer needs to use the existing corporate server pipeline and virtualization software to create a server image. The server image will be tested on-premises to resemble the build on Amazon EC2 as closely as possible.

How can this be accomplished?

- A. Download and integrate the latest ISO of CentOS 7 and execute the application deployment on the resulting server.
- B. Launch an Amazon Linux AMI using an AWS OpsWorks deployment agent onto the on-premises infrastructure, then execute the application deployment.
- C. Build an EC2 instance with the latest Amazon Linux operating system, and use the AWS Import/Export service to export the EC2 image to a VMware ISO in Amazon S3. Then import the resulting ISO onto the on-premises system.
- D. Download and integrate the latest ISO of Amazon Linux 2 and execute the application deployment on the resulting server
- E. Confirm that operating system testing results are consistent with EC2 operating system behavior.

Answer: D

NEW QUESTION 86

A company uses federated access for its AWS environment. The available roles are created and managed using AWS CloudFormation from a CI/CD pipeline. All changes should be made to the IAM roles through the pipeline. The security team found that changes are being made to the roles out-of-band and would like to detect when this occurs.

Which action will accomplish this?

- A. Use Amazon Inspector rules to detect and notify when a CloudFormation stack has a configuration change.
- B. Use an AWS Trusted Advisor CloudWatch Events rule to detect and notify when a CloudFormation stack has a configuration change.
- C. Use AWS CloudTrail to detect and notify when a CloudFormation stack has detected a configuration change.
- D. Use an AWS Config rule to detect and notify when a CloudFormation stack has detected a configuration change.

Answer: D

NEW QUESTION 90

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 94

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 AWS Lambda 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 99

A healthcare provider has a hybrid architecture that includes 120 on-premises VMware servers running RedHat and 50 Amazon EC2 instances running Amazon Linux. The company is in the middle of an all-in migration to AWS and wants to implement a solution for collecting information from the on-premises virtual machines and the EC2 instances for data analysis. The information includes:

- Operating system type and version
- Data for installed applications
- Network configuration information, such as MAC and IP addresses
- Amazon EC2 instance AMI ID and IAM profile

How can these requirements be met with the LEAST amount of administration?

- A. Write a shell script to run as a cron job on EC2 instances to collect and push the data to Amazon S3. For on-premises resources, use VMware vSphere to collect the data and write it into a file gateway for storing the data in S3. Finally, use Amazon Athena on the S3 bucket for analytics.
- B. Use a script on the on-premises virtual machines as well as the EC2 instances to gather and push the data into Amazon S3, and then use Amazon Athena for analytics.
- C. Install AWS Systems Manager agents on both the on-premises virtual machines and the EC2 instances. Enable inventory collection and configure resource data sync to an Amazon S3 bucket to analyze the data with Amazon Athena.
- D. Use AWS Application Discovery Service for deploying Agentless Discovery Connector in the VMware environment and Discovery Agents on the EC2 instances for collecting the data
- E. Then use the AWS Migration Hub Dashboard for analytics.

Answer: C

NEW QUESTION 101

A company is using AWS CodePipeline to automate its release pipeline. AWS CodeDeploy is being used in the pipeline to deploy an application to Amazon ECS using the blue/green deployment model. The company wants to implement scripts to shifting traffic. These scripts will complete in 5 minutes or less. If errors are discovered during these tests, the application must be rolled back.

Which strategy will meet these requirements?

- A. Add a stage to the CodePipeline pipeline between the source and deploy stages. Use AWS CodeBuild to create an execution environment and build commands in the buildspec file to invoke test scripts. If errors are found, use the `aws deploy stop-deployment` command to stop the deployment.
- B. Add a stage to the CodePipeline pipeline between the source and deploy stages. Use this stage to execute an AWS Lambda function that will run the test scripts. If errors are found, use the `aws deploy stop-deployment` command to stop the deployment.
- C. Add a hooks section to the CodeDeploy AppSpec file. Use the `AfterAllowTestTraffic` lifecycle event to invoke an AWS Lambda function to run the test script.
- D. If errors are found, exit the Lambda function with an error to trigger rollback.
- E. Add a hooks section to the CodeDeploy AppSpec file. Use the `AfterAllowTraffic` lifecycle event to invoke the test script.
- F. If errors are found, use the `aws deploy stop-deployment` CLI command to stop the deployment.

Answer: A

NEW QUESTION 104

A company is using AWS CodeCommit as its source code repository. After an internal audit, the compliance team mandates that any code change that go 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 105

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 108

A DevOps Engineer is implementing a mechanism for canary testing an application on AWS. The application was recently modified and went through security, unit, and functional testing. The application needs to be deployed on an Auto Scaling group and must use a Classic Load Balancer.

Which design meets the requirement for canary testing?

- A. Create a different Classic Load Balancer and Auto Scaling group for blue/green environment
- B. Use Amazon Route 53 and create weighted A records on Classic Load Balancer.
- C. Create a single Classic Load Balancer and an Auto Scaling group for blue/green environment
- D. Use Amazon Route 53 and create A records for Classic Load Balancer IP
- E. Adjust traffic using A records.
- F. Create a single Classic Load Balancer and an Auto Scaling group for blue/green environment
- G. Create an Amazon CloudFront distribution with the Classic Load Balancer as the origin
- H. Adjust traffic using CloudFront.
- I. Create a different Classic Load Balancer and Auto Scaling group for blue/green environment
- J. Create an Amazon API Gateway with a separate stage for the Classic Load Balancer
- K. Adjust traffic by giving weights to this stage.

Answer: A

NEW QUESTION 113

A company runs an application on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones in us-east-1. The application stores data in an Amazon RDS MySQL Multi-AZ DB instance.

A DevOps Engineer wants to modify the current solution and create a hot standby of the environment in another region to minimize downtime if a problem occurs

in us-east-1.

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

- A. Add a health check to the Amazon Route 53 alias record to evaluate the health of the primary region. Use AWS Lambda, configured with an Amazon CloudWatch Events trigger, to elect the Amazon RDS master in the disaster recovery region.
- B. Create a new Application Load Balancer and Auto Scaling group in the disaster recovery region.
- C. Extend the current Auto Scaling group to the subnets in the disaster recovery region.
- D. Enable multi-region failover for the RDS configuration for the database instance.
- E. Deploy a read replica of the RDS instance in the disaster recovery region.
- F. Create an AWS Lambda function to evaluate the health of the primary region.
- G. If it fails, modify the Amazon Route 53 record to point at the disaster recovery region and elect the RDS master.

Answer: ABE

Explanation:

<https://aws.amazon.com/blogs/database/implementing-a-disaster-recovery-strategy-with-amazon-rds/>

NEW QUESTION 114

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

Which solution is fully automated and will produce the MOST detailed deltas report?

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

Answer: C

NEW QUESTION 119

A security review has identified that an AWS CodeBuild project is downloading a database population script from an Amazon S3 bucket using an unauthenticated request. The security team does not allow unauthenticated requests to S3 buckets for this project.

How can this issue be corrected in the MOST secure manner?

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

Answer: C

NEW QUESTION 122

A DevOps engineer is creating a CI/CD pipeline for an Amazon ECS service. The ECS container instances run behind an Application Load Balancer as the web tier of a three-tier application. An acceptance criterion (or a successful deployment) is the verification that the web tier can communicate with the database and middleware tiers of the application upon deployment.

How can this be accomplished in an automated fashion?

- A. Create a health check endpoint in the web application that tests connectivity to the data and middleware tier
- B. Use this endpoint as the health check URL for the load balancer.
- C. Create an approval step for the quality assurance team to validate connectivity

- D. Reject changes in the pipeline if there is an issue with connecting to the dependent tiers.
- E. Use an Amazon RDS active connection count and an Amazon CloudWatch ELB metric to alarm on a significant change to the number of open connections.
- F. Use Amazon Route 53 health checks to detect issues with the web service and roll back the CI/CD pipeline if there is an error.

Answer: A

NEW QUESTION 123

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

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

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

Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 126

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

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

How can log collection be automated?

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

Answer: D

Explanation:

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

NEW QUESTION 129

A web application with multiple services runs on Amazon EC2 instances behind an Application Load Balancer. The application stores data in an Amazon RDS Multi-AZ DB instance. The instance health check used by the load balancer returns PASS if at least one service is running on the instance.

The company uses AWS CodePipeline with AWS CodeBuild and AWS CodeDeploy steps to deploy code to test and production environments. Recently, a new version was unable to connect to the database server in the test environment. One process was running, so the health checks reported healthy and the application was promoted to production, causing a production outage. The company wants to ensure that test builds are fully functional before a promotion to production.

Which changes should a DevOps Engineer make to the test and deployment process? (Choose two.)

- A. Add an automated functional test to the pipeline that ensures solid test cases are performed.
- B. Add a manual approval action to the CodeDeploy deployment pipeline that requires a Testing Engineer to validate the testing environment.
- C. Refactor the health check endpoint the Elastic Load Balancer is checking to better validate actual application functionality.
- D. Refactor the health check endpoint the Elastic Load Balancer is checking to return a text-based status result and configure the load balancer to check for a valid response.
- E. Add a dependency checking step to the existing testing framework to ensure compatibility.

Answer: DE

NEW QUESTION 132

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

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

- A. Create an Amazon CloudWatch Events rule for the CloudTrail StopLogging event

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

Answer: A

Explanation:

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

NEW QUESTION 136

A DevOps Engineer is building a multi-stage pipeline with AWS CodePipeline to build, verify, stage, test, and deploy an application. There is a manual approval stage required between the test and deploy stages. The development team uses a team chat tool with webhook support. How can the Engineer configure status updates for pipeline activity and approval requests to post to the chat tool?

- A. Create an AWS CloudWatch Logs subscription that filters on "detail-type": "CodePipeline PipelineExecution State Change." Forward that to an Amazon SNS topic
- B. Add the chat webhook URL to the SNS topic as a subscriber and complete the subscription validation.
- C. Create an AWS Lambda function that is triggered by the updating of AWS CloudTrail event
- D. When a "CodePipeline Pipeline Execution State Change" event is detected in the updated events, send the event details to the chat webhook URL.
- E. Create an AWS CloudWatch Events rule that filters on "CodePipeline Pipeline Execution State Change." Forward that to an Amazon SNS topic
- F. Subscribe an AWS Lambda function to the Amazon SNS topic and have it forward the event to the chat webhook URL.
- G. Modify the pipeline code to send event details to the chat webhook URL at the end of each stage. Parametrize the URL so each pipeline can send to a different URL based on the pipeline environment.

Answer: C

NEW QUESTION 140

An ecommerce company is looking for ways to deploy an application on AWS that satisfies the following requirements:

- Has a simple and automated application deployment process.
- Has minimal deployment costs while ensuring that at least half of the instances are available to receive end-user requests.
- If the application fails, an automated healing mechanism will replace the affected instances. Which deployment strategy will meet these requirements?

- A. Create an AWS Elastic Beanstalk environment and configure it to use Auto Scaling and an Elastic Load Balance
- B. Use rolling deployments with a batch size of 50%.
- C. Create an AWS OpsWorks stack
- D. Configure the application layer to use rolling deployments as a deployment strategy
- E. Add an Elastic Load Balancing layer
- F. Enable auto healing on the application layer.
- G. Use AWS CodeDeploy with Auto Scaling and an Elastic Load Balancer Use the CodeDeployDefault.HalfAtATime deployment strategy
- H. Enable an Elastic Load Balancing health check to report the status of the application, and set the Auto Scaling health check to ELB.
- I. Use AWS CodeDeploy with Auto Scaling and an Elastic Load Balance
- J. Use a blue/green deployment strategy
- K. Enable an Elastic Load Balancing health check to report the status of the application, and set the Auto Scaling health check to ELB.

Answer: C

NEW QUESTION 142

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

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

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

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

O. Use custom health checks to run unit tests on each instance with rollback on failure.

Answer: C

NEW QUESTION 143

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 147

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 `update` 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 152

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

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

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

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

Answer: C

NEW QUESTION 153

An application is deployed on Amazon EC2 instances running in an Auto Scaling group. During the bootstrapping process, the instances register their private IP addresses with a monitoring system. The monitoring system performs health checks frequently by sending ping requests to those IP addresses and sending alerts if an instance becomes non-responsive.

The existing deployment strategy replaces the current EC2 instances with new ones. A DevOps engineer has noticed that the monitoring system is sending false alarms during a deployment, and is tasked with stopping these false alarms.

Which solution will meet these requirements without affecting the current deployment method?

- A. Define an Amazon CloudWatch Events target, an AWS Lambda function, and a lifecycle hook attached to the Auto Scaling group
- B. Configure CloudWatch Events to invoke Amazon SNS to send a message to the systems administrator group for remediation.
- C. Define an AWS Lambda function and a lifecycle hook attached to the Auto Scaling group
- D. Configure the lifecycle hook to invoke the Lambda function, which removes the entry of the private IP from the monitoring system upon instance termination.
- E. Define an Amazon CloudWatch Events target, an AWS Lambda function, and a lifecycle hook attached to the Auto Scaling group
- F. Configure CloudWatch Events to invoke the Lambda function, which removes the entry of the private IP from the monitoring system upon instance termination.
- G. Define an AWS Lambda function that will run a script when instance termination occurs in an Auto Scaling group
- H. The script will remove the entry of the private IP from the monitoring system.

Answer: C

NEW QUESTION 158

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

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

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

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

Answer: C

NEW QUESTION 162

A company has multiple development groups working in a single shared AWS account. The Senior Manager of the groups wants to be alerted via a third-party API call when the creation of resources approaches the service limits for the account.

Which solution will accomplish this with the LEAST amount of development effort?

- A. Create an Amazon CloudWatch Event rule that runs periodically and targets an AWS Lambda function. Within the Lambda function, evaluate the current state of the AWS environment and compare deployed resource values to resource limits on the account
- B. Notify the Senior Manager if the account is approaching a service limit.
- C. Deploy an AWS Lambda function that refreshes AWS Trusted Advisor checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically
- D. Create another CloudWatch Events rule with an event pattern matching Trusted Advisor events and a target Lambda function
- E. In the target Lambda function, notify the Senior Manager.
- F. Deploy an AWS Lambda function that refreshes AWS Personal Health Dashboard checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically
- G. Create another CloudWatch Events rule with an event pattern matching Personal Health Dashboard events and a target Lambda function
- H. In the target Lambda function, notify the Senior Manager.
- I. Add an AWS Config custom rule that runs periodically, checks the AWS service limit status, and streams notifications to an Amazon SNS topic
- J. Deploy an AWS Lambda function that notifies the Senior Manager, and subscribe the Lambda function to the SNS topic.

Answer: B

NEW QUESTION 165

A company has thousands of Amazon EC2 instances as well as hundreds of virtual machines on-premises. Developers routinely sign in to the console for on-premises systems to perform troubleshooting. The developers want to sign in to AWS instances to run performance tools, but are unable to do so due to the lack of a central console logging system. A DevOps engineer wants to ensure that console access is logged on all systems.

Which combination of steps will meet these requirements? (Select TWO.)

- A. Attach a role to all AWS instances that contains the appropriate permission
- B. Create an AWS Systems Manager managed-instance activation
- C. Install and configure Systems Manager Agent on on-premises machines.
- D. Enable AWS Systems Manager Session Manager logging to an Amazon S3 bucket
- E. Direct developers to connect to the systems with Session Manager only.
- F. Enable AWS Systems Manager Session Manager logging to AWS CloudTrail
- G. Direct developers to continue normal sign-in procedures for on-premise
- H. Use Session Manager for AWS instances.
- I. Install and configure an Amazon CloudWatch Logs agent on all systems
- J. Create an AWS Systems Manager managed-instance activation.
- K. Set up a Site-to-Site VPN connection between the on-premises and AWS network
- L. Set up a bastion instance to allow developers to sign in to the AWS instances.

Answer: AB

NEW QUESTION 168

A company gives its employees limited rights to AWS. DevOps engineers have the ability to assume an administrator role. For tracking purposes, the security team wants to receive a near-real-time notification when the administrator role is assumed.

How should this be accomplished?

- A. Configure AWS Config to publish logs to an Amazon S3 bucket
- B. Use Amazon Athena to query the logs and send a notification to the security team when the administrator role is assumed.
- C. Configure Amazon GuardDuty to monitor when the administrator role is assumed and send a notification to the security team.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) event rule using an AWS Management Console sign-in events event pattern that publishes a message to an Amazon SNS topic if the administrator role is assumed
- E. [^
- F. Create an Amazon EventBridge (Amazon CloudWatch Events) events rule using an AWS API call that uses an AWS CloudTrail event pattern to trigger an AWS Lambda function that publishes a message to an Amazon SNS topic if the administrator role is assumed.

Answer: D

NEW QUESTION 169

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

Which set of steps should be taken next?

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

Answer: D

NEW QUESTION 171

A media customer has several thousand Amazon EC2 instances in an AWS account. The customer is using a Slack channel for team communications and important updates. A DevOps Engineer was told to send all AWS-scheduled EC2 maintenance notifications to the company Slack channel.

Which method should the Engineer use to implement this process in the LEAST amount of steps?

- A. Integrate AWS Trusted Advisor with AWS Config
- B. Based on the AWS Config rules created, the AWS Config event can invoke an AWS Lambda function to send notifications to the Slack channel.
- C. Integrate AWS Personal Health Dashboard with Amazon CloudWatch Event
- D. Based on the CloudWatch Events created, the event can invoke an AWS Lambda function to send notifications to the Slack channel.
- E. Integrate EC2 events with Amazon CloudWatch monitor
- F. Based on the CloudWatch Alarm created, the alarm can invoke an AWS Lambda function to send EC2 maintenance notifications to the Slack channel.
- G. Integrate AWS Support with AWS CloudTrail
- H. Based on the CloudTrail lookup event created, the event can invoke an AWS Lambda function to pass EC2 maintenance notifications to the Slack channel.

Answer: B

Explanation:

<https://docs.aws.amazon.com/health/latest/ug/cloudwatch-events-health.html>

NEW QUESTION 172

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-serverless.html>

NEW QUESTION 174

An education company has a Docker-based application running on multiple Amazon EC2 instances in an Amazon ECS cluster. When deploying a new version of the application, the Developer, pushes a new image to a private Docker container registry, and then stops and starts all tasks to ensure that they all have the latest version of the application. The Developer discovers that the new tasks are occasionally running with an old image.

How can this issue be prevented?

- A. After pushing the new image, restart ECS Agent, and then start the tasks.
- B. Use "latest" for the Docker image tag in the task definition.
- C. Update the digest on the task definition when pushing the new image.
- D. Use Amazon ECR for a Docker container registry.

Answer: C

NEW QUESTION 175

A company runs a three-tier web application in its production environment, which is built on a single AWS CloudFormation template made up of Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. Data is stored in an Amazon RDS Multi-AZ DB instance with read replicas. Amazon Route 53 manages the application's public DNS record.

A DevOps Engineer must create a workflow to mitigate a failed software deployment by rolling back changes in the production environment when a software cutover occurs for new application software.

What steps should the Engineer perform to meet these requirements with the LEAST amount of downtime?

- A. Use CloudFormation to deploy an additional staging environment and configure the Route 53 DNS with weighted record
- B. During cutover, change the Route 53 A record weights to achieve an even traffic distribution between the two environment
- C. Validate the traffic in the new environment and immediately terminate the old environment if tests are successful.
- D. Use a single AWS Elastic Beanstalk environment to deploy the staging and production environments. Update the environment by uploading the ZIP file with the new application code
- E. Swap the Elastic Beanstalk environment CNAME
- F. Validate the traffic in the new environment and immediately terminate the old environment if tests are successful.
- G. Use a single AWS Elastic Beanstalk environment and an AWS OpsWorks environment to deploy the staging and production environment
- H. Update the environment by uploading the ZIP file with the new application code into the Elastic Beanstalk environment deployed with the OpsWorks stack
- I. Validate the traffic in the new environment and immediately terminate the old environment if tests are successful.
- J. Use AWS CloudFormation to deploy an additional staging environment, and configure the Route 53 DNS with weighted record
- K. During cutover, increase the weight distribution to have more traffic directed to the new staging environment as workloads are successfully validated
- L. Keep the old production environment in place until the new staging environment handles all traffic.

Answer: D

NEW QUESTION 180

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

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

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

Answer: ADF

NEW QUESTION 183

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 188

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 189

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 190

A DevOps Engineer manages an application that has a cross-region failover requirement. The application stores its data in an Amazon Aurora on Amazon RDS database in the primary region with a read replica in the secondary region. The application uses Amazon Route 53 to direct customer traffic to the active region. Which steps should be taken to MINIMIZE downtime if a primary database fails?

- A. Use Amazon CloudWatch to monitor the status of the RDS instance
- B. In the event of a failure, use a CloudWatch Events rule to send a short message service (SMS) to the Systems Operator using Amazon SNS
- C. Have the Systems Operator redirect traffic to an Amazon S3 static website that displays a downtime message
- D. Promote the RDS read replica to the master
- E. Confirm that the application is working normally, then redirect traffic from the Amazon S3 website to the secondary region.
- F. Use RDS Event Notification to publish status updates to an Amazon SNS topic
- G. Use an AWS Lambda function subscribed to the topic to monitor database health
- H. In the event of a failure, the Lambda function promotes the read replica, then updates Route 53 to redirect traffic from the primary region to the secondary region.
- I. Set up an Amazon CloudWatch Events rule to periodically invoke an AWS Lambda function that checks the health of the primary database
- J. If a failure is detected, the Lambda function promotes the read replica
- K. Then, update Route 53 to redirect traffic from the primary to the secondary region.
- L. Set up Route 53 to balance traffic between both regions equally
- M. Enable the Aurora multi-master option, then set up a Route 53 health check to analyze the health of the database
- N. Configure Route 53 to automatically direct all traffic to the secondary region when a primary database fails.

Answer: B

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Events.html

NEW QUESTION 191

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 194

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 Firehose, 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 197

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

What is the MOST secure and flexible approach to manage the application configuration?

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

Answer: D

NEW QUESTION 199

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 201

A company needs to implement a robust CI/CD pipeline to automate the deployment of an application in AWS. The pipeline must support continuous integration, continuous delivery, and automatic rollback upon deployment failure. The entire CI/CD pipeline must be capable of being re-provisioned in alternate AWS accounts or Regions within minutes. A DevOps engineer has already created an AWS CodeCommit repository to store the source code.

Which combination of actions should be taken when building this pipeline to meet these requirements? (Select THREE.)

- A. Configure an AWS CodePipeline pipeline with a build stage using AWS CodeBuild.
- B. Copy the build artifact from CodeCommit to Amazon S3.
- C. Create an Auto Scaling group of Amazon EC2 instances behind an Application Load Balancer (ALB) and set the ALB as the deployment target in AWS CodePipeline.
- D. Create an AWS Elastic Beanstalk environment as the deployment target in AWS CodePipeline.
- E. Implement an Amazon SQS queue to decouple the pipeline components.
- F. Provision all resources using AWS CloudFormation.

Answer: ABD

NEW QUESTION 202

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 services
- B. Retrieve the database credentials from a Systems Manager SecureString parameter.
- C. Launch the EC2 instances with an EC2 IAM role to access AWS services
- 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 services
- 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 206

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

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

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

Answer: BD

NEW QUESTION 209

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 212

.....

Relate Links

100% Pass Your AWS-Certified-DevOps-Engineer-Professional Exam with Exambible Prep Materials

<https://www.exambible.com/AWS-Certified-DevOps-Engineer-Professional-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>