AWS DevOps Engineer Professional.VCEplus.premium.exam.75q

Number: DevOps Engineer Prof
Passing Score: 800
Time Limit: 120 min
File Version: 1.0

AWS DevOps Engineer Professional

Version 1.0

Exam A

QUESTION 1
To run an application, a DevOps Engineer launches an Amazon EC2 instances with public IP addresses in a public subnet. A user data script obtains the application artifacts and installs them on the instances upon launch. A change to the security classification of the application now requires the instances to run with no access to the Internet. While the instances launch successfully and show as healthy, the application does not seem to be installed. Which of the following should successfully install the application while complying with the new rule?

A. Launch the instances in a public subnet with Elastic IP addresses attached. Once the application is installed and running, run a script to disassociate the Elastic IP addresses afterwards.
B. Set up a NAT gateway. Deploy the EC2 instances to a private subnet. Update the private subnet's route table to use the NAT gateway as the default route.
C. Publish the application artifacts to an Amazon S3 bucket and create a VPC endpoint for S3. Assign an IAM instance profile to the EC2 instances so they can read the application artifacts from the S3 bucket.
D. Create a security group for the application instances and whitelist only outbound traffic to the artifact repository. Remove the security group rule once the install is complete.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 2
An IT department manages a portfolio with Windows and Linux (Amazon and Red Hat Enterprise Linux) servers both on-premises and on AWS. An audit reveals that there is no process for updating OS and core application patches, and that the servers have inconsistent patch levels.
Which of the following provides the MOST reliable and consistent mechanism for updating and maintaining all servers at the recent OS and core application patch levels?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 3
A company is setting 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 Create an Amazon CloudWatch subscription filter and use Amazon Kinesis Data Streams in the sub accounts to stream the logs to the Lambda function deploymentin the auditing account.
B. Use Amazon Kinesis Streams to write logs to Amazon ES in the auditing account. Create a CloudWatch subscription filter and use Kinesis Data Streams in the sub accounts to stream the logs to the Kinesis stream in the auditingaccount.
C. Use Amazon Kinesis Firehose with Kinesis Data Streams to write logs to Amazon ES in the auditing account. Create a CloudWatch subscription filter and stream logs from sub accounts to the Kinesis stream in the auditing account.
D. Use AWS Lambda to write logs to Amazon ES in the auditing account. 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.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 4

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

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

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 5
A company has established tagging and configuration standatrds for its infrastructure resources running on AWS. A DevOps Engineer is developing a design that will provide a near-real-time dashboard of the compliance posture with the ability to highlight violations.
Which approach meets the stated requirements?

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

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Reference: https://aws.amazon.com/answers/configuration-management/aws-infrastructure-configuration-management/

QUESTION 6
A production account has a requirement that any Amazon EC2 instance that has been logged into manually must be terminated within 24 hours. All applications in the production account are using Auto Scaling groups with Amazon CloudWatch Logs agent configured.
How can this process be automated?

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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 7

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 AutoScaling 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 environments. Use Amazon Route 53 and create weighted A records on Classic Load Balancer.
B. Create a single Classic Load Balancer and an Auto Scaling group for blue/green environments. Use Amazon Route 53 and create A records for Classic Load Balancer IPs. Adjust traffic using A records.
C. Create a single Classic Load Balancer and an Auto Scaling group for blue/green environments. Create an Amazon CloudFront distribution with the Classic Load Balancer as the origin. Adjust traffic using CloudFront.
D. Create a different Classic Load Balancer and Auto Scaling group for blue/green environments. Create an Amazon API Gateway with a separate stage for the Classic Load Balancer. Adjust traffic by giving weights to this stage.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 8
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.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 9
A company has several AWS accounts. The accounts are shared and used across multiple teams globally, primarily for Amazon EC2 instances. Each EC2 instance has tags for team, environment, and cost center to ensure accurate cost allocations.
How should a DevOps Engineer help the teams audit their costs and automate infrastructure cost optimization across multiple shared environments and accounts?

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

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 10
A company has a hybrid architecture solution in which some legacy systems remain on-premises, while a specific cluster of servers is moved to AWS. The company cannot reconfigure the legacy systems, so the cluster nodes must have a fixed hostname and local IP address for each server that is part of the cluster. The DevOps Engineer must automate the configuration for a six-node cluster with high availability across three Availability Zones (AZs), placing two elastic network interfaces in a specific subnet for each AZ. Each node's hostname and local IP address should remain the same between reboots or instance failures.

Which solution involves the LEAST amount of effort to automate this task?

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

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 11
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.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 12
A financial institution provides security-hardened AMIs of Red Hat Enterprise Linux 7.4 and Windows Server 2016 for its application teams to use in deployments. A DevOps Engineer needs to implement an automated daily check of each AMI to monitor for the latest CVE.
How should the Engineer implement these checks using Amazon Inspector?

A. Install the Amazon Inspector agent in each AMI. Configure AWS Step Functions to launch an Amazon EC2 instance for each operating system from the hardened AMI, and tag the instance with SecurityCheck: True. Once EC2 instanceshave booted up, Step Functions will trigger an Amazon Inspector assessment for all instances with the tag SecurityCheck: True. Implement a scheduled Amazon CloudWatch Events rule that triggers Step Functions once each day.
B. Tag each AMI with SecurityCheck: True. Configure AWS Step Functions to first compose an Amazon Inspector assessment template for all AMIs that have the tag SecurityCheck: True and second to make a call to the AmazonInspector API action StartAssessmentRun. Implement a scheduled Amazon CloudWatch Events rule that triggers Step Functions once each day.
C. Tag each AMI with SecurityCheck: True. Implement a scheduled Amazon Inspector assessment to run once each day for all AMIs with the tag SecurityCheck: True. Amazon Inspector should automatically launch an Amazon EC2instance for each AMI and perform a security assessment.
D. Tag each instance with SecurityCheck: True. Implement a scheduled Amazon Inspector assessment to tun once each day for all instances with the tag SecurityCheck: True. Amazon Inspector should automatically perform an in-placesecurity assessment for each AMI.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 13
A Development team uses AWS CodeCommit for source code control. Developers apply their changes to various feature branches and create pull requests to move those changes to the master branch when they are ready for production. A direct push to the master branch should not be allowed. The team applied the AWS managed policy AWSCodeCommitPowerUser to the Developers' IAM rote, but now members are able to push to the master branch directly on every repository in the AWS account.
What actions should be taken to restrict this?

A. Create an additional policy to include a deny rule for the codecommit: GitPush action, and include a restriction for the specific repositories in the resource statement with a condition for the master reference.
B. Remove the IAM policy and add an AWSCodeCommitReadOnlypolicy. Add an allow rule for the codecommit: GitPush action for the specific repositories in the resource statement with a condition for the master reference.
C. Modify the IAM policy and include a deny rule for the codecommit: GitPush action for the specific repositories in the resource statement with a condition for the master reference.
D. Create an additional policy to include an allow rule for the codecommit: GitPush action and include a restriction for the specific repositories in the resource statement with a condition for the feature branches reference.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 14
A Developer is designing a continuous deployment workflow for a new Development team to facilitate the process for source code promotion in AWS. Developers would like to store and promote code for deployment from development to production while maintaining the ability to roll back that deployment if it fails. Which design will incur the LEAST amount of downtime?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 15
A DevOps Engineer discovered a sudden spike in a website's page load times and found that a recent deployment occurred. A brief diff of the related commit shows that the URL for an external API call was altered and the connecting port changed from 80 to 443. The external API has been verified and works outside the application. The application logs show that the connection is now timing out, resulting in multiple retries and eventual failure of the call. Which debug steps should the Engineer take to determine the root cause of the issue'?

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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 16
An Engineering team manages a Node.js e-commerce application. The current environment consists of the following components:

• Amazon S3 buckets for storing content
• Amazon EC2 for the front-end web servers
• AWS Lambda for executing image processing
• Amazon DynamoDB for storing session-related data

The team expects a significant increase in traffic to the site. The application should handle the additional load without interruption. The team ran initial tests by adding new servers to the EC2 front-end to handle the larger load, but the instances took up to 20 minutes to become fully configured. The team wants to reduce this configuration time.
What changes will the Engineering team need to implement to make the solution the MOST resilient and highly available while meeting the expected increase in demand?

A. Use AWS OpsWorks to automatically configure each new EC2 instance as it is launched. Configure the EC2 instances by using an Auto Scaling group behind an Application Load Balancer across multiple Availability Zones. ImplementAmazon DynamoDB Auto Scaling. Use Amazon Route 53 to point the application DNS record to the Application Load Balancer.
B. Deploy a fleet of EC2 instances, doubling the current capacity, and place them behind an Application Load Balancer. Increase the Amazon DynamoDB read and write capacity units. Add an alias record that contains the Application LoadBalancer endpoint to the existing Amazon Route 53 DNS record that points to the application.
C. Configure Amazon CloudFront and have its origin point to Amazon S3 to host the web application. Implement Amazon DynamoDB Auto Scaling. Use Amazon Route 53 to point the application DNS record to the CloudFront DNS name.
D. Use AWS Elastic Beanstalk with a custom AMI including all web components. Deploy the platform by using an Auto Scaling group behind an Application Load Balancer across multiple Availability Zones. Implement Amazon DynamoDBAuto Scaling. Use Amazon Route 53 to point the application DNS record to the Elastic Beanstalk load balancer.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 17
A DevOps Engineer is working on a project that is hosted on Amazon Linux and has failed a security review. The DevOps Manager has been asked to review the company buildspec.yami file for an AWS CodeBuild project and provide recommendations. The builspec.yami file is configured as follows:

```
env:
variables:
AWS_ACCESS_KEY_ID: AKIAJF7BRFWJBA4GHXNA
AWS_SECRET_ACCESS_KEY: ORjJns3At2mIh4O4tm0+zHxZqz7cNAvMLYRehcI
AWS_DEFAULT_REGION: us-east-1
DB_PASSWORD: cuj5RptFa3va
phases:
build:
commands:
-aws s3 cp s3://db-deploy-bucket/my.cnf.template/tmp/my.cnf
-sed-i '' s/DB_PW/${DB_PASSWORD}/ /tmp/my.cnf
-aws s3 cp s3:// db-deploy-bucket/instance.key/tmp/instance.key
-chmod 600/tmp/instance.key
-scp-i /tmp/instance.key/tmp/my.cnf root@10.25.23:/etc/my.cnf
-ssh- i /tmp/instance.key root@10.25.15.23 /etc/init.d/mysqld restart
```

What changes should be recommended to comply with AWS security best practices? (Select THREE.)

A. Add a post-build command to remove the temporary files from the container before termination to ensure they cannot be seen by other CodeBuild users.
B. Update the CodeBuild project role with the necessary permissions and then remove the AWS credentials from the environment variable.

C. Store the DB_PASSWORD as a SecurityString value in AWS Systems Manager Parameter Store and then remove the DB_PASSWORD from the environment variables.
D. Move the environment variables to the 'db-deploy-bucket' Amazon S3 bucket, add a prebuild stage to download, then export the variables.
E. Use AWS Systems Manager run command versus scp and ssh commands directly to the instance.
F. Scramble the environment variables using XOR followed by Base64, add a section to install, and then run XOR and Base64 to the build phase.

Correct Answer: ADE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 18
A Development team is building more than 40 applications. Each app is a three-tiered web application based on an ELB Application Load Balancer, Amazon EC2, and Amazon RDS. Because the applications will be used internally, the Security team wants to allow access to the 40 applications only from the corporate network and block access from external IP addresses. The corporate network reaches the internet through proxy servers. The proxy servers have 12 proxy IP addresses that are being changed one or two times per month. The Network Infrastructure team manages the proxy servers; they upload the file that contains the latest proxy IP addresses into an Amazon S3 bucket. The DevOps Engineer must build a solution to ensure that the applications are accessible from the corporate network.
Which solution achieves these requirements with MINIMAL impact to application development, MINIMAL operational effort, and the LOWEST infrastructure cost?

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

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 19
A company is implementing AWS CodePipeline to automate its testing process. The company wants to be notified when the execution state fails and used the following custom event pattern in Amazon CloudWatch:

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

Which type of events will match this event pattern?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 20
A company is using several AWS CloudFormation templates for deploying infrastructure as code. In most of the deployments, the company uses Amazon EC2 Auto Scaling groups. A DevOps Engineer needs to update the AMIs for the Auto Scaling group in the template if newer AMIs are available. How can these requirements be met?

A. Manage the AMI mappings in the CloudFormation template. Use Amazon CloudWatch Events for detecting new AMIs and updating the mapping in the template. Reference the map in the launch configuration resource block.
B. Use conditions in the AWS CloudFormation template to check if new AMIs are available and return the AMI ID. Reference the returned AMI ID in the launch configuration resource block.
C. Use an AWS Lambda-backed custom resource in the template to fetch the AMI IDs. Reference the returned AMI ID in the launch configuration resource block.
D. Launch an Amazon EC2 m4 small instance and run a script on it to check for new AMIs. If new AMIs are available, the script should update the launch configuration resource block with the new AMI ID.

Correct Answer: D
Section: (none)
Explanation
Explanation/Reference:

QUESTION 21
A DevOps Engineer administers an application that manages video files for a video production company. The application runs on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. Data is stored in an Amazon RDS PostgreSOL Multi-AZ DB instance, and the video ides are stored in an Amazon S3 bucket. On a typical day 50 GB of new video are added to the S3 bucket. The Engineer must implement a multi-region disaster recovery plan with the least data loss and the lowest recovery times. The current application infrastructure is already described using AWS CloudFormation. Which deployment option should the Engineer choose to meet the uptime and recovery objectives for the system?

A. Launch the application from the CloudFormation template in the second region, which sets the capacity of the Auto Scaling group to 1. Create an Amazon RDS read replica in the second region. In the second region, enable cross-regionreplication between the original S3 bucket and a new S3 bucket. To fail over, promote the read replica as master. Update the CloudFormation stack and increase the capacity of the Auto Scaling group.
B. Launch the application from the CloudFormation template in the second region, witch sets the capacity of the Auto Scaling group to 1. Create a scheduled task to take daily Amazon RDS cross-region snapshots to the second region. Inthe second region, enable cross-region replication between the original S3 bucket and Amazon Glacier. In a disaster, launch a new application stack in the second region and restore the database from the most recent snapshot.
C. Launch the application from the CloudFormation template in the second region which sets the capacity of the Auto Scaling group to 1. Use Amazon CloudWatch Events to schedule a nightly task to take a snapshot of the database, copythe snapshot to the second region, and replace the DB instance in the second region from the snapshot. In the second region, enable cross-region replication between the original S3 bucket and a new S3 bucket. To fail over, increase the capacity of the Auto Scaling group.
D. Use Amazon CloudWatch Events to schedule a nightly task to take a snapshot of the database and copy the snapshot to the second region. Create an AWS Lambda function that copies each object to a new S3 bucket in the secondregion in response to S3 event notifications. In the second region, launch the application from the CloudFormation template and restore the database from the most recent snapshot.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 22
A social networking service runs a web API that allows its partners to search public posts. Post data is stored in Amazon DynamoDB and indexed by AWS Lambda functions, with an Amazon ES domain storing the indexes and providing search functionality to the application.
The service needs to maintain full capacity during deployments and ensure that failed deployments do not cause downtime or reduced capacity, or prevent subsequent deployments. How can these requirements be met? (Select TWO)

A. Run the web application in AWS Elastic Beanstalk with the deployment policy set to All at Once. Deploy the Lambda functions, DynamoDB tables, and Amazon ES domain with an AWS CloudFormation template.
B. Deploy the web application, Lambda functions, DynamoDB tables, and Amazon ES domain in an AWS CloudFormation template. Deploy changes with an AWS CodeDeploy in-place deployment.
C. Run the web application in AWS Elastic Beanstalk with the deployment policy set to Immutable Deploy the Lambda functions, DynamoDB tables, and Amazon ES domain with an AWS CloudFormation template.
D. Deploy the web application, Lambda functions, DynamoDB tables, and Amazon ES domain in an AWS CloudFormation template. Deploy changes with an AWS CodeDeploy blue/green deployment.
E. Run the web.application in AWS Elastic Beanstalk with the deployment policy set to Rolling Deploy the Lambda functions, DynamoDB tables, and Amazon ES domain with an AWS CloudFormation template.

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 23
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 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. Based on the AWS Config rules created, the AWS Config event can invoke an AWS Lambda function to send notifications to the Slack channel.
B. Integrate AWS Personal Health Dashboard with Amazon CloudWatch Events. Based on the CloudWatch Events created, the event can invoke an AWS Lambda function to send notifications to the Slack channel.
C. Integrate EC2 events with Amazon CloudWatch monitoring. Based on the CloudWatch Alarm created, the alarm can invoke an AWS Lambda function to send EC2 maintenance notifications to the Slack channel.
D. Integrate AWS Support with AWS CloudTrail. Based on the CloudTrail lookup event created, the event can invoke an AWS Lambda function to pass EC2 maintenance notifications to the Slack channel.

Correct Answer: C
Section: (none)
Explanation
Explanation/Reference:
Reference: https://yabhinav.github.io/cloud/awslambda-slack-notifications/

QUESTION 24
After conducting a disaster recovery exercise, an Enterprise Architect discovers that a large team of Database and Storage Administrators need more then seven hours of manual effort to make a flagship application's database functional in a different AWS Region. The Architect also discovers that the recovered database is often missing as much as two hours of data transactions. Which solution provides improved RTO and RPO in a cross-region failover scenario?

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

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 25
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 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. Use AWS Config aggregators to collect data from multiple accounts and regions. Export the aggregated report to an Amazon S3 bucket and useAmazon SNS to deliver the notifications.
B. Set up AWS CloudTrail to deliver all events to an Amazon S3 bucket in a centralized account. Use the S3 event notification feature to invoke an AWS Lambda function to parse AWS CloudTrail logs whenever logs are delivered to theS3 bucket. Publish the output to an Amazon SNS topic using the same Lambda function.
C. 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. Store consolidated evaluation results fromconfig rules in Amazon S3. Send a notification using Amazon SNS when non-compliant resources are detected.
D. Using AWS CLI, run a script periodically that invokes the aws ec2 describe-volumes query with a JMESPATH query filter. Then, write the output to an Amazon S3 bucket. Set up an S3 event notification to send events using AmazonSNS when new data is written to the S3 bucket.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 26
A DevOps Engineer has a single Amazon Dynamo DB table that received shipping orders and tracks inventory. The Engineer has three AWS Lambda functions reading from a DymamoDB 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 Dynamo DB 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. Then use this Lambda function to pass the payload to the other three Lambda functions.
D. Have the Lambda functions query the table directly and disable DynamoDB streams. Then have the Lambda functions query from a global secondary index.

Correct Answer: C

Section: (none)
Explanation

Explanation/Reference:

QUESTION 27
A government agency is storing highly confidential files in an encrypted Amazon S3 bucket. The agency has configured federated access and has allowed only a particular on-premises Active Directory user group to access this bucket. The agency wants to maintain audit records and automatically detect and revert any accidental changes administrators make to the IAM policies used for providing this restricted federated access. Which of the following options provide the FASTEST way to meet these requirements?

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

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 28
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 tor 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. Then use the AWS Migration Hub Dashboard for analytics.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 29
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. Confirm that operating system testing results are consistent with EC2 operating system behavior.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 30
A Development team is adding a new country to an e-commerce application. This addition requires that new application features be added to the shipping component of the application. The team has not decided if all new features should be added, as some will take approximately six weeks to build. While the final decision on the shipping component features is being made, other team members are continuing to work on other features of the application. Based on this situation, how should the application feature deployments be managed?

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

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 31
A DevOps Engineer is asked to implement a strategy for deploying updates to a web application with zero downtime. The application infrastructure is defined in AWS CloudFormation and is made up of an Amazon Route 53 record, an Application Load Balancer, Amazon EC2 instances in an EC2 Auto Scaling group, and Amazon DynamoDB tables. To avoid downtime, there must be an active instance serving the application at all times. Which strategies will ensure the deployment happens with zero downtime? (Select TWO.)

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

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 32
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. 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.
B. Create an AWS Systems Manager automation document with values instructing how the image should be created. Then build a pipeline in AWS CodePipeline to execute the automation document to build the AMI when triggered. Storethe AMI identification output as a Systems Manager parameter.
C. Build a pipeline in AWS CodePipeline to take a snapshot of an Amazon EC2 instance running the latest version of the application. Then start a new EC2 instance from the snapshot and update the running instance using an AWSLambda function. Take a snapshot of the updated instance, then convert it to an AMI. Store the AMI identification output in an Amazon DynamoDB table.
D. Launch an Amazon EC2 instance and install Packer. Then configure a Packer build with values defining how the image should be created. Build a Jenkins pipeline to invoke the Packer build when triggered to build an AMI. Store the AMIidentification output in an Amazon DynamoDB table.

Correct Answer: D

Section: (none)
Explanation

Explanation/Reference:

QUESTION 33
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. The HAProxy instances will proxy the requests to one of the existing Auto Scaling groups. After a deployment the HAProxy instances areupdated to send requests to the newly deployed Auto Scaling group.
B. Reduce the application to one Application Load Balancer. Create two target groups named Blue and Green. Create a rule on the Application Load Balancer pointed to a single target group. Add logic to the deployment to update theApplication Load Balancer rule to the target group of the newly deployed Auto Scaling group.
C. Move the application to an AWS Elastic Beanstalk application with two environments. Perform new deployments on the non-live environment. After a deployment, perform an Elastic Beanstalk CNAME swap to make the newly deployedenvironment the live environment.
D. Create an Amazon CloudFront distribution. Set the two existing Application Load Balancers as origins on the distribution. After a deployment, update the CloudFront distribution behavior to send requests to the newly deployed AutoScaling group.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 34
A company has microservices running in AWS Lambda that read data from Amazon DynamoDB. The Lambda code is manually deployed by Developers after successful testing. The company now needs the tests and deployments be automated and run in the cloud. Additionally, traffic to the new versions of each microservice should be incrementally shifted over time after deployment. What solution meets all the requirements, ensuring the MOST developer velocity?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 35
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. Dynamically alter the launch configuration of the Amazon EC2 instances.
B. Set up a CodePipeline pipeline for each environment to use input parameters. Use CloudFormation mappings to switch associated UserData for the Amazon EC2 instances to match the environment being launched.

C. Set up a CodePipeline pipeline that has multiple stages, one for each development environment. Use AWS Lambda functions to trigger CloudFormation deployments to dynamically alter the UserData of the Amazon EC2 instanceslaunched in each environment.
D. Use CloudFormation input parameters to dynamically alter the LaunchConfiguration and UserData sections of each Amazon EC2 instance every time the CloudFormation stack is updated.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 36
An application runs on Amazon EC2 instances behind an Application Load Balancer. Amazon RDS MySOL is used on the backend. The instances run in an Auto Scaling group across multiple Availability Zones. The Application Load Balancer health check ensures the web servers are operating and able to make read/write SQL connections. Amazon Route 53 provides DNS functionality with a record pointing to the Application Load Balancer. A new policy requires a geographically isolated disaster recovery site with an RTO of 4 hours and an RPO of 15 minutes. Which disaster recovery strategy will require the LEAST amount of changes to the application stack?

A. Launch a replica stack of everything except RDS in a different Availability Zone. Create an RDS read-only replica in a new Availability Zone and configure the new stack to point to the local RDS instance. Add the new stack to the Route53 record set with a failover routing policy.
B. Launch a replica stack of everything except RDS in a different region. Create an RDS read-only replica in a new region and configure the new stack to point to the local RDS instance. Add the new stack to the Route 53 record set with alatency routing policy.
C. Launch a replica stack of everything except RDS in a different region. Upon failure, copy the snapshot over from the primary region to the disaster recovery region. Adjust the Amazon Route 53 record set to point to the disaster recoveryregion's Application Load Balancer.
D. Launch a replica stack of everything except RDS in a different region. Create an RDS read-only replica in a new region and configure the new stack to point to the local RDS instance. Add the new stack to the Amazon Route 53 recordset with a failover routing policy.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 37
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”	“rrr”	“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. Have LSI with LastPostDateTime as the sort key and fetch operations for thread.
B. Use ForumName as the primary key and Subject as the sort key. Have LSI with LastPostDateTime as the sort key and the projected attribute thread.
C. Use ForumName as the primary key and Subject as the sort key. Have LSI with Thread as the sort key and the projected attribute LastPostDateTime.
D. Use Subject as the primary key and ForumName as the sort key. Have LSI with Thread as the sort key and fetch operations for LastPostDateTime.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 38
A company used AWS CloudFormation to deploy a three-tier web application that stores data in an Amazon RDS MySOL Multi-AZ DB instance. A DevOps Engineer must upgrade the RDS instance to the latest major version of MySQL while incurring minimal downtime.
How should the Engineer upgrade the instance while minimizing downtime?

A. Update the EngineVersion property of the AWS::RDS:: DBInstance resource type in the CloudFormation template to the latest desired version. Launch a second stack and make the new RDS instance a read replica.

B. Update the DBEngineVersion property of the AWS: : RDS : :DBInstance resource type in the CloudFormation template to the latest desired version. Perform an Update Stack operation. Create a new RDS Read Replicas resource withthe same properties as the instance to be upgraded. Perform a second Update Stack operation.
C. Update the DBEngineVersion property of the AWS: :RDS: :DB:Instance resource type in the CloudFormation template to the latest desired version. Create a new RDS Read Replicas resource with the same properties as the instance tobe upgraded. Perform an Update Stack operation.
D. Update the EngineVersion property of the AWS :: RDS :: DBInstance resource type in the CloudFormation template to the latest version, and perform an Update Stack operation.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 39
A retail company has adopted AWS OpsWorks for managing its deployments. In the last three months: the company has discovered that some production instances have been restarting without reason. Upon inspection of the AWS CloudTrail logs, a DevOps Engineer determined that those instances were restarted by OpsWorks. The Engineer now wants automated email notifications whenever OpsWorks restarts an instance when the instance is deemed unhealthy or unable to communicate with the service endpoint. How can the Engineer meet this requirement?

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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 40
A healthcare services company is concerned about the growing costs of software licensing for an application for monitoring patient wellness. The company wants to create an audit process to ensure that the application is running exclusively on Amazon EC2 Dedicated Hosts. A DevOps Engineer must create a workflow to audit the application to ensure compliance. What steps should the Engineer take to meet this requirement with the LEAST administrative overhead?

A. Use AWS Systems Manager Configuration Compliance. Use calls to the put-compliance- items API action to scan and build a database of noncompliant EC2 instances based on their host placement configuration. Use an AmazonDynamoDB table to store these instance IDs for fast access. Generate a report through Systems Manager by calling the list-compliance- summaries API action.
B. Use custom Java code running on an EC2 instance. Set up EC2 Auto Scaling for the instance depending on the number of instances to be checked. Send the list of noncompliant EC2 instance IDs to an Amazon SQS queue. Set upanother worker instance to process instance IDs from the SQS queue and write them to Amazon DynamoDB. Use an AWS Lambda function to terminate noncompliant instance IDs obtained from the queue, and send them to an Amazon SNS email topic for distribution.
C. Use AWS Config Identify all EC2 instances to be audited by enabling Config Recording on all Amazon EC2 resources for the region. Create a custom AWS Config rule that triggers an AWS Lambda function by using the “config-rule-change-triggered” blueprint. Modify the Lambda evaluate.Compliance () function to verify host placement to return a NON_COMPLIANT result if the instance is not running on an EC2 Dedicated Host. Use the AWS Config report to address noncompliant instances.
D. Use AWS CloudTrail. Identity all EC2 instances to be audited by analyzing all calls to the EC2 RunCommand API action. Invoke an AWS Lambda function that analyzes the host placement of the instance. Store the EC2 instance ID ofnoncompliant resources in an Amazon RDS MySOL DB instance. Generate a report by querying the RDS instance and exporting the query results to a CSV text file.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 41
According to Information Security Policy, changes to the contents of objects inside production Amazon S3 bucket that contain encrypted secrets should only be made by a trusted group of administrators. How should a DevOps Engineer create real-time, automated checks to meet this requirement?
A. Create an AWS Lambda function that is triggered by Amazon S3 data events for object changes and that also checks the IAM user’s membership in an administrator’s IAM role.

B. Create a periodic AWS Config rule to query Amazon S3 Logs for changes and to check the IAM user's membership in an administrator's IAM role.
C. Create a metrics filter for Amazon CloudWatch logs to check for Amazon S3 bucket-level permission changes and to check the IAM user's membership in an administrator's IAM role.
D. Create a periodic AWS Config rule to query AWS CloudTrail logs for changes to the Amazon S3 bucket-level permissions and to check the IAM user's membership in an administrator's IAM role.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 42
A business has an application that consists of five independent AWS Lambda functions.
The DevOps Engineer has built a CI/CD pipeline using AWS CodePipeline and AWS CodeBuild that builds, tests, packages, and deploys each Lambda function in sequence. The pipeline uses an Amazon CloudWatch Events rule to ensure the pipeline execution starts as quickly as possible after a change is made to the application source code.
After working with the pipeline for a few months the DevOps Engineer has noticed the pipeline takes too long to complete. What should the DevOps Engineer implement to BEST improve the speed of the pipeline?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 43
A company uses a complex system that consists of networking, IAM policies, and multiple three-tier applications. Requirements are still being defined for a new system, so the number of AWS components present in the final design is not known. The DevOps Engineer needs to begin defining AWS resources using AWS CloudFormation to automate and version-control the new infrastructure. What is the best practice for using CloudFormation to create new environments?

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

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 44
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. Use Elastic Beanstalk features for a blue/green deployment to deploy the new release to aseparate environment, and then swap the CNAME in the two environments to redirect traffic to the new version.
B. 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.
C. Deploy the application using Elastic Beanstalk, and include RDS MySQL as part of the environment. Use Elastic Beanstalk immutable updates for application deployments.
D. Deploy the application using Elastic Beanstalk, and connect to an external RDS MySQL instance using Elastic Beanstalk environment properties. Use Elastic Beanstalk immutable updates for application deployments.
Correct Answer: A

Section: (none)
Explanation

Explanation/Reference:

QUESTION 45
Am Amazon EC2 instance with no internet access is running in a Virtual Private Cloud (VPC) and needs to download an object from a restricted Amazon S3 bucket. When the DevOps Engineer tries to gain access to the object, an AccessDenied error is received.
What are the possible causes for this error? (Select THREE.)

A. The S3 bucket default encryption is enabled.
B. There is an error in the S3 bucket policy.
C. There is an error in the VPC endpoint policy.
D. The object has been moved to Amazon Glacier.
E. There is an error in the IAM role configuration.
F. S3 versioning is enabled.

Correct Answer: BCE
Section: (none)
Explanation

Explanation/Reference:
Reference: https://aws.amazon.com/premiumsupport/knowledge-center/s3-403-upload-bucket/

QUESTION 46
An application has microservices spread across different AWS accounts and is integrated with an on-premises legacy system for some of its functionality. Because of the segmented architecture and missing logs, every time the application experiences issues, it is taking too long to gather the logs to identify the issues. A DevOps Engineer must fix the log aggregation process and provide a way to centrally analyze the logs. Which is the MOST efficient and cost-effective solution?

A. Collect system logs and application logs by using the Amazon CloudWatch Logs agent. Use the Amazon S3 API to export on-premises logs, and store the logs in an S3 bucket in a central account. Build an Amazon EMR cluster toreduce the logs and derive the root cause.
B. Collect system logs and application logs by using the Amazon CloudWatch Logs agent. Use the Amazon S3 API to import on-premises logs. Store all logs in S3 buckets in individual accounts. Use Amazon Macie to write a query tosearch for the required specific event-related data point.
C. Collect system logs and application logs using the Amazon CloudWatch Logs agent. Install the CloudWatch Logs agent on the on-premises servers. Transfer all logs from AWS to the on-premises data center. Use an AmazonElasticsearch Logstash Kibana stack to analyze logs on premises.
D. Collect system logs and application logs by using the Amazon CloudWatch Logs agent. Install a CloudWatch Logs agent for on-premises resources. Store all logs in an S3 bucket in a central account. Set up an Amazon S3 trigger andan AWS Lambda function to analyze incoming logs and automatically identify anomalies. Use Amazon Athena to run ad hoc queries on the logs in the central account.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 47
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. Then set up a CodePipeline approval action for a Developer to test and approve the new version. Finally, use a CodePipeline invoke action to update an AWSLambda function to use the production alias
B. 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.
C. Use AWS CloudFormation to define the serverless application and AWS CodeDeploy to deploy the AWS Lambda functions using DeploymentPreference: Canary10Percentl5Minutes.
D. Use AWS CloudFormation to publish a new version on every stack update. Use the RoutingConfig property of the AWS : :Lambda: : Alias resource to update the traffic routing during the stack update.
Correct Answer: A

Section: (none)
Explanation

Explanation/Reference:

QUESTION 48
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. Alarm and stop deployment when latency increases beyond the defined threshold.
C. Use AWS CodeDeploy's MinimumHealthyHosts setting to define thresholds for rolling back deployments. If these thresholds are breached, roll back the deployment.
D. Use Metric Filters to parse application logs in Amazon CloudWatch Logs. Create a filter for latency. Alarm and stop deployment when latency increases beyond the defined threshold.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 49
A DevOps Engineer is leading the implementation for automating patching of Windows-based workstations in a hybrid cloud environment by using AWS Systems Manager (SSM). What steps should the Engineer follow to set up Systems Manager to automate patching in this environment? (Select TWO.)

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

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 50
A company needs to introduce automatic DNS failover for a distributed web application to a disaster recovery or standby installation. The DevOps Engineer plans to configure Amazon Route 53 to provide DNS routing to alternate endpoint in the event of an application failure.
What steps should the Engineer take to accomplish this? (Select TWO.)

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

Correct Answer: AC
Section: (none)
Explanation
Explanation/Reference:

QUESTION 51
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 instances. Change the logging driver in the ECS task definition to 'awslogs'.
B. 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.
C. 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.
D. Enable access logging on the Application Load Balancer, then point it directly to the S3 logging bucket.
E. Enable access logging on the target groups that are used by the ECS services, then point it directly to the S3 logging bucket.
F. Create an Amazon Kinesis Data Firehose with a destination of the S3 logging bucket, then create an Amazon CloudWatch Logs subscription filter for Kinesis.

Correct Answer: ACE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 52
A Development team is currently using AWS CodeDeploy to deploy an application revision to an Auto Scaling group. If the deployment process fails, it must be rolled back automatically and a notification must be sent. What is the MOST effective configuration that can satisfy all of the requirements?

A. Create Amazon CloudWatch Events rules for CodeDeploy operations. Configure a CloudWatch Events rule to send out an Amazon SNS message when the deployment fails. Configure CodeDeploy to automatically roll back when thedeployment fails.
B. Use available Amazon CloudWatch metrics for CodeDeploy to create CloudWatch alarms. Configure CloudWatch alarms to send out an Amazon SNS message when the deployment fails. Use AWS CLI to redeploy a previouslydeployed revision.
C. Configure a CodeDeploy agent to create a trigger that will send notification to Amazon SNS topics when the deployment fails. Configure CodeDeploy to automatically roll back when the deployment fails.
D. Use AWS CloudTrail to monitor API calls made by or on behalf of CodeDeploy in the AWS account. Send an Amazon SNS message when deployment fails. Use AWS CLI to redeploy a previously deployed revision.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 53
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. Retrieve the database credentials from a Systems Manager SecureString parameter.
B. Launch the EC2 instances with an EC2 IAM role to access AWS services. Retrieve the database credentials from AWS Secrets Manager.
C. 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.
D. Launch the EC2 instances with an EC2 IAM role to access AWS services. Store the database passwords in an encrypted config file with the application artifacts.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 54
A DevOps Engineer is designing a deployment strategy for a web application. The application will use an Auto Scaling group to launch Amazon EC2 instances using an AMI. The same infrastructure will be deployed in multiple environments (development, test, and quality assurance). The deployment strategy should meet the following requirements:

• Minimize the startup time for the instance
• Allow the same AMI to work in multiple environments
• Store secrets for multiple environments securelyHow should this be accomplished?

A. Preconfigure the AMI using an AWS Lambda function that launches an Amazon EC2 instance, and then runs a script to install the software and create the AMI. Configure an Auto Scaling lifecycle hook to determine which environmentthe instance is launched in, and, based on that finding, run a configuration script. Save the secrets on an .ini file and store them in Amazon S3. Retrieve the secrets using a configuration script in EC2 user data.
B. Preconfigure the AMI by installing all the software using AWS Systems Manager automation and configure Auto Scaling to tag the instances at launch with their specific environment. Then use a bootstrap script in user data to read thetags and configure settings for the environment. Use the AWS Systems Manager Parameter Store to store the secrets using AWS KMS.
C. Use a standard AMI from the AWS Marketplace. Configure Auto Scaling to detect the current environment. Install the software using a script in Amazon EC2 user data. Use AWS Secrets Manager to store the credentials for allenvironments.
D. Preconfigure the AMI by installing all the software and configuration for all environments. Configure Auto Scaling to tag the instances at launch with their environment. Use the Amazon EC2 user data to trigger an AWS Lambda functionthat reads the instance ID and then reconfigures the setting for the proper environment. Use the AWS Systems Manager Parameter Store to store the secrets using AWS KMS.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 55
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. Create an Amazon CloudWatch Alarm for EC2 Instance Terminate Successful and trigger an AWS Lambda function that executes an SSM Run Command scriptto collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
B. Use Auto Scaling lifecycle hooks to put instances in a Terminating: Wait state. Create a Config rule for EC2 Instance-terminate Lifecycle Action 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.
C. Use Auto Scaling lifecycle hooks to put instances in a Terminating: Wait state. Create an Amazon CloudWatch subscription filter for EC2 Instance Terminate Successful and trigger a CloudWatch agent that executes a script to calledlogs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
D. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait state. Create an Amazon CloudWatch Events rule for EC2 'Instance-terminate Lifecycle Action and trigger an AWS Lambda function that executes a SSM RunCommand script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 56
A publishing company used AWS Elastic Beanstalk, Amazon S3, and Amazon DynamoDB to develop a web application. The web application has increased dramatically in popularity, resulting in unpredictable spikes in traffic. A DevOps Engineer has noted that 90% of the requests are duplicate read requests. How can the Engineer improve the performance of the website?

A. Use Amazon ElastiCache for Redis to cache repeated read requests to DynamoDB and AWS Elemental MediaStore to cache images stored in S3.
B. Use Amazon ElastiCache for Memcached to cache repeated read requests to DynamoDB and Varnish to cache images stored in S3.
C. Use DynamoDB Accelerator to cache repeated read requests to DynamoDB and Amazon CloudFront to cache images stored in S3.
D. Use DynamoDB Streams to cache repeated read requests to DynamoDB and API Gateway to cache images stored in S3.

Correct Answer: A
Section: (none)
Explanation
Explanation/Reference:

QUESTION 57

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

• A mechanism to automate the use of the licenses at scale.
• Creation of a dashboard to use in the future to verify which licenses are available at any moment.

What is the MOST effective way to accomplish these requirements'?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 58
A company has developed a static website hosted on an Amazon S3 bucket. The website is deployed using AWS CloudFormation. The CloudFormation template defines an S3 bucket and a custom resource that copies content into the bucket from a source location.
The company has decided that it needs to move the website to a new location, so the existing CloudFormation stack must be deleted and re-created. However, CloudFormation reports that the stack could not be deleted cleanly. What is the MOST likely cause and how can the DevOps Engineer mitigate this problem for this and future versions of the website?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 59
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 regions. Any update to the code repository triggers a two-stage build and deployment pipeline. A successful deployment in one region invokes an AWS Lambda function to copy the buildartifacts to an Amazon S3 bucket in another region. After the artifact is copied, it triggers a deployment pipeline in the new region.
B. Deploy the backend application to multiple Availability Zones in a single region. Create an Amazon CloudFront distribution to serve the application backend to global customers. Any update to the code repository triggers a two-stagebuild-and-deployment pipeline. The pipeline deploys the backend application to all Availability Zones.
C. Deploy the backend application to multiple regions. Use AWS Direct Connect to serve the application backend to global customers. Any update to the code repository triggers a two-stage build-and-deployment pipeline in the region.After a successful deployment in the region, the pipeline continues to deploy the artifact to another region.
D. Deploy the backend application to multiple regions. Any update to the code repository triggers a two-stage build-and-deployment pipeline in the region. After a successful deployment in the region, the pipeline invokes the pipeline inanother region and passes the build artifact location. The pipeline uses the artifact location and deploys applications in the new region.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 60
A Development team is working on a serverless application in AWS. To quickly identify and remediate potential production issues, the team decides to roll out changes to a small number of users as a test before the full release. The DevOps Engineer must develop a solution to minimize downtime and impact.
Which of the following solutions should be used to meet the requirements? (Select TWO.)

A. Create an Application Load Balancer with two target groups. Set up the Application Load Balancer for Amazon API Gateway private integration. Associate one target group to the current version and the other target group to the new version. Configure API Gateway to route 10% of incoming traffic to the new version. As the new version becomes stable, configure API Gateway to send all traffic to the new version and detach the old version from the load balancer. B. Create an alias for an AWS Lambda function pointing to both the current and new versions. Configure the alias to route 10% of incoming traffic to the new version. As the new version is considered stable, update the alias to route all traffic to the new version.

C. Create a rollover record set in AWS Route 53 pointing to the AWS Lambda endpoints for the old and new versions. Configure Route 53 to route 10% of incoming traffic to the new version. As the new version becomes stable, update theDNS record to route all traffic to the new version.

D. Create an ELB Network Load Balancer with two target groups. Set up the Network Load Balancer for Amazon API Gateway private integration Associate one target group with the current version and the other target group with the newversion. Configure the load balancer to route 10% of incoming traffic to the new version. As the new version becomes stable, detach the old version from the load balancer.

E. In Amazon API Gateway, create a canary release deployment by adding canary settings to the stage of a regular deployment. Configure API Gateway to route 10% of the incoming traffic to the canary release. As the canary release isconsidered stable, promote it to a production release

Correct Answer: DE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 61
A company wants to implement a Cl/CD pipeline for an application that is deployed on AWS. The company also has a source-code analysis tool hosted on premises that checks for security flaws. The tool has not yet been migrated to AWS and can be accessed only on premises. The company wants to run checks against the source code as part of the pipeline before the code is compiled. The checks take anywhere from minutes to an hour to complete. How can a DevOps Engineer meet these requirements'?

A. Use AWS CodePipeline to create a pipeline. Add an action to the pipeline to invoke an AWS Lambda function after the source stage. Have the Lambda function invoke the source-code analysis tool on premises against the source inputfrom CodePipeline. The function then waits for the execution to complete and places the output in a specified Amazon S3 location.

B. Use AWS CodePipeline to create a pipeline, then create a custom action type. Create a job worker for the custom action that runs on hardware hosted on premises. The job worker handles running security checks with the on-premisescode analysis tool and then returns the job results to CodePipeline. Have the pipeline invoke the custom action after the source stage.

C. Use AWS CodePipeline to create a pipeline. Add a step after the source stage to make an HTTPS request to the on-premises hosted web service that invokes a test with the source code analysis tool. When the analysis is complete, theweb service sends the results back by putting the results in an Amazon S3 output location provided by CodePipeline.

D. Use AWS CodePipeline to create a pipeline. Create a shell script that copies the input source code to a location on premises. Invoke the source code analysis tool and return the results to CodePipeline. Invoke the shell script by addinga custom script action after the source stage.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 62
A company is adopting AWS CodeDeploy to automate its application deployments for a Java-Apache Tomcat application with an Apache webserver. The Development team started with a proof of concept, created a deployment group for a developer environment, and performed functional tests within the application. After completion, the team will create additional deployment groups for staging and production
The current log level is configured within the Apache settings, but the team wants to change this configuration dynamically when the deployment occurs, so that they can set different log level configurations depending on the deployment group without having a different application revision for each group.
How can these requirements be met with the LEAST management overhead and without requiring different script versions for each deployment group?

A. Tag the Amazon EC2 instances depending on the deployment group. Then place a script into the application revision that calls the metadata service and the EC2 API to identify which deployment group the instance is part of. Use thisinformation to configure the log level settings. Reference the script as part of the Afterinstall lifecycle hook in the appspec yml file.
B. Create a script that uses the CodeDeploy environment variable DEPLOYMENT_GROUP_NAME to identify which deployment group the instances is part of. Use this information to configure the log level settings. Reference this script aspart of the BeforeInstall lifecycle hook in the appspec.yml file
C. Create a CodeDeploy custom environment variable for each environment. Then place a script into the application revision that checks this environment variable to identify which deployment group the instance is part of. Use thisinformation to configure the log level settings. Reference this script as part of the ValidateService lifecycle hook in the appspec yml file.
D. Create a script that uses the CodeDeploy environment variable DEPLOYMENT_GROUP_ID to identify which deployment group the instance is part of to configure the log level settings. Reference this script as part of the Install lifecyclehook in the appspec yml file.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 63
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 recipes. 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 providespermission to access DynamoDB.
B. Create a Docker file that uses the Chef recipes for the application environment based on an official Node.js Docker image. Create an Amazon ECS cluster and a service for the application environment, then create a task based on thisDocker image. Use scheduled scaling to scale the containers at the appropriate times and attach a task-level IAM role that provides permission to access DynamoDB.
C. Configure AWS OpsWorks stacks and use custom Chef cookbooks. Add the Git repository information where the custom recipes are stored, and add a layer in OpsWorks for the Node.js application server. Then configure the customrecipe to deploy the application in the deploy step. Configure time-based instances and attach an Amazon EC2 IAM role that provides permission to access DynamoDB.
D. Configure AWS OpsWorks stacks and push the custom recipes to an Amazon S3 bucket and configure custom recipes to point to the S3 bucket. Then add an application layer type for a standard Node.js application server and configurethe custom recipe to deploy the application in the deploy step from the S3 bucket. Configure time-based instances and attach an Amazon EC2 IAM role that provides permission to access DynamoDB.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 64
The Development team at an online retailer has moved to Business support and want to take advantage of the AWS Health Dashboard and the AWS Health API to automate remediation actions for issues with the health of AWS resources. The first use case is to respond to AWS detecting an IAM access key that is listed on a public code repository site. The automated response will be to delete the IAM access key and send a notification to the Security team.

How should this be achieved?

A. Create an AWS Lambda function to delete the IAM access key. Send AWS CloudTrail logs to AWS CloudWatch logs. Create a CloudWatch Logs metric filter for the AWS_RISK_CREDENTIALS_EXPOSED event with two actions: first,run the Lambda function; second, use Amazon SNS to send a notification to the Security team.
B. Create an AWS Lambda function to delete the IAM access key. Create an AWS Config rule for changes to aws.health and the AWS_RISK_CREDENTIALS_EXPOSED event with two actions: first, run the Lambda function; second, useAmazon SNS to send a notification to the Security team.
C. Use AWS Step Functions to create a function to delete the IAM access key, and then use Amazon SNS to send a notification to the Security team. Create an AWS Personal Health Dashboard rule for theAWS_RISK_CREDENTIALS_EXPOSED event; set the target of the Personal Health Dashboard rule to Step Functions.
D. Use AWS Step Functions to create a function to delete the IAM access key, and then use Amazon SNS to send a notification to the Security team. Create an Amazon CloudWatch Events rule with an aws.health event source and theAWS_RISK_CREDENTIALS_EXPOSED event, set the target of the CloudWatch Events rule to Step Functions.

Correct Answer: C
Section: (none)
Explanation
Explanation/Reference:

QUESTION 65
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. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called. Add theLambda function ARN as a target to the CloudWatch Events rule.
B. 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. Create an AWS Lambda function that usesthe AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called. Add the Lambda function ARN as a target to the CloudWatch Events rule.
C. Create an Amazon CloudWatch Events rule for a scheduled event every 5 minutes. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on an CloudTrail trail in the AWS account. Add the Lambda functionARN as a target to the CloudWatch Events rule.
D. Launch a t2.nano instance with a script running every 5 minutes that uses the AWS SDK to query CloudTrail in the current account. If the CloudTrail trail is disabled, have the script re-enable the trail.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 66
A DevOps Engineer has been asked by the Security team to ensure that AWS CloudTrail files are not tampered with after being created. Currently, there is a process with multiple trails, using AWS IAM to restrict access to specific trails. The Security team wants to ensure they can trace the integrity of each file and make sure there has been no tampering.

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

A. Create an Amazon CloudWatch Events rule that triggers an AWS Lambda function when a new file is delivered. Configure the Lambda function to perform an MD5 hash check on the file, store the name and location of the file, and postthe returned hash to an Amazon DynamoDB table. The Security team can use the values stored in DynamoDB to verify the file authenticity.
B. Enable the CloudTrail file integrity feature on an Amazon S3 bucket. Create an IAM policy that grants the Security team access to the file integrity logs stored in the S3 bucket.
C. Enable the CloudTrail file integrity feature on the trail. Use the digest file created by CloudTrail to verify the integrity of the delivered CloudTrail files.
D. Create an AWS Lambda function that is triggered each time a new file is delivered to the CloudTrail bucket. Configure the Lambda function to execute an MD5 hash check on the file, and store the result on a tag in an Amazon S3object. The Security team can use the information on the tag to verify the integrity of the file.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 67
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. Create a CodeCommit repository for each environment. Set up CodePipeline to retrieve the source code from the appropriate repository. Set up a deployment step to deploy the Lambdafunctions with AWS CloudFormation.
B. Create two AWS CodePipeline configurations for test and production environments. Configure the production pipeline to have a manual approval step. Create a CodeCommit repository for each environment. Set up each CodePipelineto retrieve the source code from the appropriate repository. Set up the deployment step to deploy the Lambda functions with AWS CloudFormation.
C. Create two AWS CodePipeline configurations for test and production environments. Configure the production pipeline to have a manual approval step. Create one CodeCommit repository with a branch for each environment. Set upeach CodePipeline to retrieve the source code from the appropriate branch in the repository. Set up the deployment step to deploy the Lambda functions with AWS CloudFormation.

D. Create an AWS CodeBuild configuration for test and production environments. Configure the production pipeline to have a manual approval step. Create one CodeCommit repository with a branch for each environment. Push theLambda function code to an Amazon S3 bucket. Set up the deployment step to deploy the Lambda functions from the S3 bucket.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 68
A company is using AWS for an application. The Development team must automate its deployments. The team has set up an AWS CodePipeline to deploy the application to Amazon EC2 instances by using AWS CodeDeploy after it has been built using the AWS CodeBuild service.

The team would like to add automated testing to the pipeline to confirm that the application is healthy before deploying it to the next stage of the pipeline using the same code. The team requires a manual approval action before the application is deployed, even if the test is successful. The testing and approval must be accomplished at the lowest costs, using the simplest management solution.

Which solution will meet these requirements?

A. Add a manual approval action after the last deploy action of the pipeline. Use Amazon SNS to inform the team of the stage being triggered. Next, add a test action using CodeBuild to do the required tests. At the end of the pipeline, adda deploy action to deploy the application to the next stage.
B. Add a test action after the last deploy action of the pipeline. Configure the action to use CodeBuild to perform the required tests. If these tests are successful, mark the action as successful. Add a manual approval action that usesAmazon SNS to notify the team, and add a deploy action to deploy the application to the next stage.
C. Create a new pipeline that uses a source action that gets the code from the same repository as the first pipeline. Add a deploy action to deploy the code to a test environment. Use a test action using AWS Lambda to test the deployment.Add a manual approval action by using Amazon SNS to notify the team, and add a deploy action to deploy the application to the next stage.
D. Add a test action after the last deployment action. Use a Jenkins server on Amazon EC2 to do the required tests and mark the action as successful if the tests pass. Create a manual approval action that uses Amazon SQS to notify theteam and add a deploy action to deploy the application to the next stage.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 69
A company is building a solution for storing files containing Personally Identifiable Information (PII) on AWS.

Requirements state:
- All data must be encrypted at rest and in transit.
- Al data must be replicated in at least two locations that are at least 500 miles apart.

Which solution meets these requirements?

A. Create primary and secondary Amazon S3 buckets in two separate Availability Zones that are at least 500 miles apart. Use a bucket policy to enforce access to the buckets only through HTTPS. Use a bucket policy to enforce AmazonS3 SSE-C on all objects uploaded to the bucket. Configure cross-region replication between the two buckets.
B. Create primary and secondary Amazon S3 buckets in two separate AWS Regions that are at least 500 miles apart. Use a bucket policy to enforce access to the buckets only through HTTPS. Use a bucket policy to enforce S3-ManagedKeys (SSE-S3) on all objects uploaded to the bucket. Configure cross-region replication between the two buckets.
C. Create primary and secondary Amazon S3 buckets in two separate AWS Regions that are at least 500 miles apart. Use an IAM role to enforce access to the buckets only through HTTPS. Use a bucket policy to enforce Amazon S3-Managed Keys (SSE-S3) on all objects uploaded to the bucket. Configure cross-region replication between the two buckets.
D. Create primary and secondary Amazon S3 buckets in two separate Availability Zones that are at least 500 miles apart. Use a bucket policy to enforce access to the buckets only through HTTPS. Use a bucket policy to enforce AWS KMSencryption on all objects uploaded to the bucket. Configure cross-region replication between the two buckets. Create a KMS Customer Master Key (CMK) in the primary region for encrypting objects.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
QUESTION 70

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. Associate the Auto Scaling group with the deployment group. Use the `Automatically copy Auto Scaling group` option, and use CodeDeployDefault.OneAtAtime as the deployment configuration. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the AllowTraffic hook within appspec.yml to delete the temporary files.
B. Use an Application Load Balancer and a blue/green deployment. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment group. 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. 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.
C. Use an Application Load Balancer and a blue/green deployment. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment group. Use the `Automatically copy Auto Scaling group` option, and use CodeDeployDefault HalfAtAtime as the deployment configuration. Instruct AWS CodeDeploy to terminate the original isntances in the deployment group, and use the BeforeAllowTraffic hook within appspec.yml to delete the temporary files.
D. Use an Application Load Balancer and an in-place deployment. Associate the Auto Scaling group and Application Load Balancer target group with the deployment group. Use the `Automatically copy Auto Scaling group` option, and use CodeDeployDefault AllatOnce as a deployment configuration. 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.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 71
A DevOps Engineer is working with an application deployed to 12 Amazon EC2 instances across 3 Availability Zones. New instances can be started from an AMI image. On a typical day, each EC2 instance has 30% utilization during business hours and 10% utilization after business hours. The CPU utilization has an immediate spike in the first few minutes of business hours. Other increases in CPU utilization rise gradually.

The Engineer has been asked to reduce costs while retaining the same or higher reliability.

Which solution meets these requirements?

A. Create two Amazon CloudWatch Events rules with schedules before and after business hours begin and end. Create two AWS Lambda functions, one invoked by each rule. The first function should stop nine instances after businesshours end, the second function should restart the nine instances before the business day begins.
B. Create an Amazon EC2 Auto Scaling group using the AMI image, with a scaling action based on the Auto Scaling group's CPU Utilization average with a target of 75%. Create a scheduled action for the group to adjust the minimumnumber of instances to three after business hours end and reset to six before business hours begin.
C. Create two Amazon CloudWatch Events rules with schedules before and after business hours begin and end. Create an AWS CloudFormation stack, which creates an EC2 Auto Scaling group, with a parameter for the number ofinstances. Invoke the stack from each rule, passing a parameter value of three in the morning, and six in the evening.
D. Create an EC2 Auto Scaling group using the AMI image, with a scaling action based on the Auto Scaling group's CPU Utilization average with a target of 75%. Create a scheduled action to terminate nine instances each evening afterthe close of business.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 72
A DevOps Engineer must improve the monitoring of a Finance team payments microservice that handles transactions for an e-commerce platform. The microservice runs on multiple Amazon EC2 instances. The Finance team would like to know the number of payments per minute, and the team would like to be notified when this metric falls below a specified threshold.

How can this be cost-effectively automated?

A. Have the Development team log successful transactions to an application log. Set up Logstash on each instance, which sends logs to an Amazon ES cluster. Create a Kibana dashboard for the Finance team that graphs the metric.

B. Have the Development team post the number of successful transactions to Amazon CloudWatch as a custom metric. Create a CloudWatch alarm when the threshold is breached, and use Amazon SNS to notify the Finance team.
C. Have the Development team log successful transactions to an application log. On each instance, set up the Amazon CloudWatch Logs agent to send application logs to CloudWatch Logs. Use an EC2 instance to monitor a metric filter,and send notifications to the Finance team.
D. Have the Development team log successful transactions to an application log. Set up the Amazon CloudWatch agent on each instance. Create a CloudWatch alarm when the threshold is breached, and use Amazon SNS to notify theFinance team.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 73
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 AZs. Use a single-node Aurora instance.
B. Create an EC2 instance and enable instance recovery. 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.
C. 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. Create an Aurora database with a read replica in a secondAZ, and promote it to a primary database instance when the primary database instance fails.
D. Assign an Elastic IP address on the instance. Create a second EC2 instance in a second AZ. Create an Amazon CloudWatch Events rule to trigger an AWS Lambda function to move the Elastic IP address to the second instance whenthe first instance fails. Use a single-node Aurora instance.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 74
An Application team has three environments for their application: development, pre-production, and production. The team recently adopted AWS CodePipeline. However, the team has had several deployments of misconfigured or nonfunctional development code into the production environment, resulting in user disruption and downtime. The DevOps Engineer must review the pipeline and add steps to identify problems with the application before it is deployed.

What should the Engineer do to identify functional issues during the deployment process? (Choose two.)

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

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 75
A DevOps Engineer is responsible for the deployment of a PHP application. The Engineer is working in a hybrid deployment, with the application running on both on-premises servers and Amazon EC2 instances. The application needs access to a database containing highly confidential information. Application instances need access to database credentials, which must be encrypted at rest and in transit before reaching the instances.

How should the Engineer automate the deployment process while also meeting the security requirements?

A. Use AWS Elastic Beanstalk with a PHP platform configuration to deploy application packages to the instances. Store database credentials on AWS Systems Manager Parameter Store using the Secure String data type. Define an IAMrole for Amazon EC2 allowing access, and decrypt only the database credentials. Associate this role to all the instances.
B. Use AWS CodeDeploy to deploy application packages to the instances. Store database credentials on AWS Systems Manager Parameter Store using the Secure String data type. Define an IAM policy for allowing access, and decryptonly the database credentials. Attach the IAM policy to the role associated to the instance profile for CodeDeploy-managed instances, and to the role used for on-premises instances registration on CodeDeploy.
C. Use AWS CodeDeploy to deploy application packages to the instances. Store database credentials on AWS Systems Manager Parameter Store using the Secure String data type. Define an IAM role with an attached policy that allowsdecryption of the database credentials. Associate this role to all the instances and on-premises servers.
D. Use AWS CodeDeploy to deploy application packages to the instances. Store database credentials in the AppSpec file. Define an IAM policy for allowing access to only the database credentials. Attach the IAM policy to the roleassociated to the instance profile for CodeDeploy-managed instances and the role used for on-premises instances registration on CodeDeploy.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference: