Free Questions for DOP-C01 by dumpshq

Shared by West on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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

Options:

A) Create API Gateway alpha, beta, and production stages. Create a CodeCommit trigger to deploy code to the different stages using an AWS Lambda function.

B) Create API Gateway alpha, beta, and production stages. Create an AWS CodePipeline that pulls code from the CodeCommit repository. Create CodePipeline actions to deploy code to the API Gateway stages.

C) Create Jenkins servers for the alpha, beta, and production stages on Amazon EC2 instances. Create multiple CodeCommit triggers to deploy code to different stages using an AWS Lambda function.

D) Create an AWS CodePipeline pipeline that pulls code from the CodeCommit repository. Create alpha, beta, and production stages with Jenkins servers on CodePipeline.

Answer:

D

Question 2

Question Type: MultipleChoice

An online company uses Amazon EC2 Auto Scaling extensively to provide an excellent customer experience while minimizing the number of running EC2 instances. The company's self-hosted Puppet environment in the application layer manages the configuration of the instances. The IT manager wants the lowest licensing costs and wants to ensure that whenever the EC2 Auto Scaling group scales down, removed EC2 instances are deregistered from the Puppet master as soon as possible. How can the requirement be met?

Options:

A) At instance launch time, use EC2 user data to deploy the AWS CodeDeploy agent. Use CodeDeploy to install the Puppet agent. When the Auto Scaling group scales out, run a script to register the newly deployed instances to the Puppet master. When the Auto Scaling group scales in, use the EC2 Auto Scaling lifecycle hook to trigger de-registration from the Puppet master. EC2_INSTANCE_TERMINATING

B) Bake the AWS CodeDeploy agent into the base AMI. When the Auto Scaling group scales out, use CodeDeploy to install the Puppet agent, and execute a script to register the newly deployed instances to the Puppet master. When the Auto Scaling group scales in, use the CodeDeploy ApplicationStop lifecycle hook to run a script to de-register the instance from the Puppet master.

C) At instance launch time, use EC2 user data to deploy the AWS CodeDeploy agent. When the Auto Scaling group scales out, use CodeDeploy to install the Puppet agent, and run a script to register the newly deployed instances to the Puppet master. When the Auto Scaling group scales in, use the EC2 user data instance stop script to run a script to de-register the instance from the Puppet master.

D) Bake the AWS Systems Manager agent into the base AMI. When the Auto Scaling group scales out, use the AWS Systems Manager to install the Puppet agent, and run a script to register the newly deployed instances to the Puppet master. When the Auto Scaling group scales in, use the Systems Manager instance stop lifecycle hook to run a script to de-register the instance from the Puppet master.

Answer:

C

Question 3

Question Type: MultipleChoice

A DevOps Engineer needs to deploy a scalable three-tier Node.js application in AWS. The application must have zero downtime during deployments and be able to roll back to previous versions. Other applications will also connect to the same MySQL backend database. The CIO has provided the following guidance for logging: *Centrally view all current web access server logs. *Search and filter web and application logs in near-real time. *Retain log data for three months. How should these requirements be met?

Options:

A) Deploy the application using AWS Elastic Beanstalk. Configure the environment type for Elastic Load Balancing and Auto Scaling. Create an Amazon RDS MySQL instance inside the Elastic Beanstalk stack. Configure the Elastic Beanstalk log options to stream logs to Amazon CloudWatch Logs. Set retention to 90 days.

B) Deploy the application on Amazon EC2. Configure Elastic Load Balancing and Auto Scaling. Use an Amazon RDS MySQL instance for the database tier. Configure the application to store log files in Amazon S3. Use Amazon EMR to search and filter the data. Set an Amazon S3 lifecycle rule to expire objects after 90 days.

C) Deploy the application using AWS Elastic Beanstalk. Configure the environment type for Elastic Load Balancing and Auto Scaling. Create the Amazon RDS MySQL instance outside the Elastic Beanstalk stack. Configure the Elastic Beanstalk log options to stream logs to Amazon CloudWatch Logs. Set retention to 90 days.

D) Deploy the application on Amazon EC2. Configure Elastic Load Balancing and Auto Scaling. Use an Amazon RDS MySQL instance for the database tier. Configure the application to load streaming log data using Amazon Kinesis Data Firehouse into Amazon ES. Delete and create a new Amazon ES domain every 90 days.

Answer:

B

Explanation:

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

Question 4

Question Type: MultipleChoice

A healthcare company has a critical application running in AWS. Recently, the company experienced some down time. if it happens again, the company needs to be able to recover its application in another AWS Region. The application uses Elastic Load Balancing and Amazon EC2 instances. The company also maintains a custom AMI that contains its application. This AMI is changed frequently. The workload is required to run in the primary region, unless there is a regional service disruption, in which case traffic should fail over to the new region. Additionally, the cost for the second region needs to be low. The RTO is 2 hours. Which solution allows the company to fail over to another region in the event of a failure, and also meet the above requirements?

Options:

A) Maintain a copy of the AMI from the main region in the backup region. Create an Auto Scaling group with one instance using a launch configuration that contains the copied AMI. Use an Amazon Route 53 record to direct traffic to the load balancer in the backup region in the event of failure, as required. Allow the Auto Scaling group to scale out as needed during a failure.

B) Automate the copying of the AMI in the main region to the backup region. Generate an AWS Lambda function that will create an EC2 instance from the AMI and place it behind a load balancer. Using the same Lambda function, point the Amazon Route 53 record to the load balancer in the backup region. Trigger the Lambda function in the event of a failure.

C) Place the AMI in a replicated Amazon S3 bucket. Generate an AWS Lambda function that can create a launch configuration and assign it to an already created Auto Scaling group. Have one instance in this Auto Scaling group ready to accept traffic. Trigger the Lambda function in the event of a failure. Use an Amazon Route 53 record and modify it with the same Lambda function to point to the load balancer in the backup region.

D) Automate the copying of the AMI to the backup region. Create an AWS Lambda function that can create a launch configuration and assign it to an already created Auto Scaling group. Set the Auto Scaling group maximum size to 0 and only increase it with the Lambda function during a failure. Trigger the Lambda function in the event of a failure. Use an Amazon Route 53 record and modify it with the same Lambda function to point to the load balancer in the backup region.

Answer:

C

Question 5

Question Type: MultipleChoice

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

Options:

A) Enable Amazon Macie in every account. Configure the security account as the Macie Administrator for every member account using invitation/acceptance. Create an Amazon CloudWatch Events rule in the security account to send all findings to Amazon Kinesis Data

Firehouse, which should push the findings to the S3 bucket.

B) Enable Amazon Macie in the security account only. Configure the security account as the Macie Administrator for every member account using invitation/ acceptance. Create an Amazon CloudWatch Events rule in the security account to send all findings to Amazon Kinesis Data Streams. Write and application using KCL to read data from the Kinesis Data Streams and write to the S3 bucket.

C) Enable Amazon GuardDuty in every account. Configure the security account as the GuardDuty Administrator for every member account using invitation/ acceptance. Create an Amazon CloudWatch rule in the security account to send all findings to Amazon Kinesis Data Firehouse, which will push the findings to the S3 bucket.

D) Enable Amazon GuardDuty in the security account only. Configure the security account as the GuardDuty Administrator for every member account using invitation/acceptance. Create an Amazon CloudWatch rule in the security account to send all findings to Amazon Kinesis Data Streams. Write and application using KCL to read data from Kinesis Data Streams and write to the S3 bucket.

Answer:

C

Explanation:

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

Question 6

Question Type: MultipleChoice

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

Options:

A) Use Amazon Athena to query AWS CloudTrail logs to check for any associate-address attempts. Create an AWS Lambda function to dissociate the Elastic IP address from the instance, and alert the Security team.

B) Attach an IAM policy to the Developer's IAM group to deny associate-address permissions. Create a custom AWS Config rule to check whether an Elastic IP address is associated with any instance tagged as production, and alert the Security team.

C) Ensure that all IAM groups are associated with Developers do not have associate-address permissions. Create a scheduled AWS Lambda function to check whether an Elastic IP address is associated with any instance tagged as production, and alert the Security team if an instance has an Elastic IP address associated with it.

D) Create an AWS Config rule to check that all production instances have the EC2 IAM roles that include deny associate-address permissions. Verify whether there is an Elastic IP address associated with any instance, and alert the Security team if an instance has an Elastic IP address associated with it.

Answer:

B

Explanation:

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

Question 7

Question Type: MultipleChoice

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

Options:

A) Use Amazon CloudWatch Events to monitor for Auto Scaling event notifications of new instances and configure CloudWatch Events to trigger an Amazon Inspector scan.

B) Use Amazon CloudWatch Events to monitor for AWS CodeDeploy notifications of a successful code deployment and configure CloudWatch Events to trigger an Amazon Inspector scan.

C) Use Amazon CloudWatch Events to monitor for CodePipeline notifications of a successful code deployment and configure CloudWatch Events to trigger an AWS X-Ray scan.

D) Use Amazon Inspector as a CodePipeline task after the successful use of CodeDeploy to deploy the code to the systems.

Answer:

A

Question 8

Question Type: MultipleChoice

A global company with distributed Development teams built a web application using a microservices architecture running on Amazon ECS. Each application service is independent and runs as a service in the ECS cluster. The container build files and source code reside in a private GitHub source code repository. Separate ECS clusters exist for development, testing, and production environments. Developers are required to push features to branches in the GitHub repository and then merge the changes into an environment-specific branch (development, test, or production). This merge needs to trigger an automated pipeline to run a build and a deployment to the appropriate ECS cluster. What should the DevOps Engineer recommend as an automated solution to these requirements?

Options:

A) Create an AWS CloudFormation stack for the ECS cluster and AWS CodePipeline services. Store the container build files in an Amazon S3 bucket. Use a post-commit hook to trigger a CloudFormation stack update that deploys the ECS cluster. Add a task in the ECS cluster to build and push images to Amazon ECR, based on the container build files in S3.

B) Create a separate pipeline in AWS CodePipeline for each environment. Trigger each pipeline based on commits to the corresponding environment branch in GitHub. Add a build stage to launch AWS CodeBuild to create the container image from the build file and push it to Amazon ECR. Then add another stage to update the Amazon ECS task and service definitions in the appropriate cluster for that environment.

C) Create a pipeline in AWS CodePipeline. Configure it to be triggered by commits to the master branch in GitHub. Add a stage to use the Git commit message to determine which environment the commit should be applied to, then call the create-image Amazon ECR command to build the image, passing it to the container build file. Then add a stage to update the ECS task and service definitions in the appropriate cluster for that environment.

D) Create a new repository in AWS CodeCommit. Configure a scheduled project in AWS CodeBuild to synchronize the GitHub repository to the new CodeCommit repository. Create a separate pipeline for each environment triggered by changes to the CodeCommit repository. Add a stage using AWS Lambda to build the container image and push to Amazon ECR. Then add another stage to update the ECS task and service definitions in the appropriate cluster for that environment.

Answer:

B

Explanation:

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

Question 9

Question Type: MultipleChoice

A DevOps Engineer is reviewing a system that uses Amazon EC2 instances in an Auto Scaling group. This system uses a configuration management tool that runs locally on each EC2 instance. Because of the volatility of the application load, new instances must be fully functional within 3 minutes of entering a running state. Current setup tasks include: Installing the configuration management agent -- 2 minutes Installing the application framework -- 15 minutes Copying configuration data from Amazon S3 -- 2 minutes Running the configuration management agent to configure instances -- 1 minute Deploying the application code from Amazon S3 -- 2 minutes How should the Engineer set up system so it meets the launch time requirement?

Options:

A) Trigger an AWS Lambda function from an Amazon CloudWatch Events rule when a new EC2 instance launches. Have the function install the configuration management agent and the application framework, pull configuration data from Amazon S3, run the agent to configure the instance, and deploy the application from S3.

B) Write a bootstrap script to install the configuration management agent, install and the application framework, pull configuration data from Amazon S3, run the agent to configure the instance, and deploy the application from S3.

C) Build a custom AMI that includes the configuration management agent and application framework. Write a bootstrap script to pull configuration data from Amazon S3, run the agent to configure the instance, and deploy the application from S3.

D) Build a custom AMI that includes the configuration management agent, application framework, and configuration data. Write a bootstrap script to run the agent to configure the instance and deploy the application from Amazon S3.

Answer:

B

Question 10

Question Type: MultipleChoice

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

The buildspec.yml file contains the following:

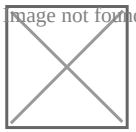

What steps should the DevOps Engineer take to stop this?

Options:

A) Modify the post_build to command to use "-acl public-read and configure a bucket policy that grants read access to the relevant AWS accounts only.

B) Configure a default ACL for the S3 bucket that defines the set of authenticated users as the relevant AWS accounts only and grants

read-only access.

C) Create an S3 bucket policy that grants read access to the relevant AWS accounts and denies read access to the principal '*'

D) Modify the post_build command to remove "-acl authenticated-read and configure a bucket policy that allows read access to the relevant AWS accounts only.

Answer:

C

Question 11

Question Type: MultipleChoice

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

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

Options:

A) Developers should use Docker for local development. When dependencies are changed and a new container is ready, use AWS CodePipeline and AWS CodeBuild to perform functional tests and then upload the new container to Amazon ECR. Use AWS

CloudFormation with the custom container to deploy the new Amazon ECS.

B) Developers should use Docker for local development. Use AWS SMS to import these containers as AMIs for Amazon EC2 whenever dependencies are updated. Use AWS CodePipeline to test new code changes against the Auto Scaling group.

C) Developers should use their native Python environment. When Dependencies are changed and a new container is ready, use AWS CodePipeline and AWS CodeBuild to perform functional tests and then upload the new container to the Amazon ECR. Use AWS CloudFormation with the custom container to deploy the new Amazon ECS.

D) Developers should use their native Python environment. When Dependencies are changed and a new code is ready, use AWS CodePipeline and AWS CodeBuild to perform functional tests and then upload the new container to the Amazon ECR. Use CodePipeline and CodeBuild with the custom container to test new code changes inside AWS Elastic Beanstalk

Answer:

A

Question 12

Question Type: MultipleChoice

An application is running on Amazon EC2. It has an attached IAM role that is receiving an AccessDenied error while trying to access a SecureString parameter resource in the AWS Systems Manager Parameter Store. The SecureString parameter is encrypted with a customer-managed Customer Master Key (CMK),

What steps should the DevOps Engineer take to grant access to the role while granting least privilege? (Select three.)

Options:

A) Set ssm:GetParamter for the parameter resource in the instance role's IAM policy.

B) Set kms:Decrypt for the instance role in the customer-managed CMK policy.

C) Set kms:Decrypt for the customer-managed CMK resource in the role's IAM policy.

D) Set ssm:DecryptParameter for the parameter resource in the instance role IAM policy.

E) Set kms:GenerateDataKey for the user on the AWS managed SSM KMS key.

F) Set kms:Decrypt for the parameter resource in the customer-managed CMK policy.

Answer:

A, B, C