DUMPSARENA

AWS DevOps Engineer Professional (DOP-C01)

Amazon AWS AWS-DevOps-Engineer-Professional-DOP-C01

Version Demo

Total Demo Questions: 20

Total Premium Questions: 557

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

You work for an insurance company and are responsible for the day-to-day operations of your company's online quote system used to provide insurance quotes to members of the public. Your company wants to use the application logs generated by the system to better understand customer behavior. Industry, regulations also require that you retain all application logs for the system indefinitely in order to investigate fraudulent claims in the future. You have been tasked with designing a log management system with the following requirements:

- All log entries must be retained by the system, even during unplanned instance failure.

- The customer insight team requires immediate access to the logs from the past seven days.

- The fraud investigation team requires access to all historic logs, but will wait up to 24 hours before these logs are available.

How would you meet these requirements in a cost-effective manner? (Choose three.)

A. Configure your application to write logs to the instance's ephemeral disk, because this storage is free and has good write performance. Create a script that moves the logs from the instance to Amazon 53 once an hour.

B. Write a script that is configured to be executed when the instance is stopped or terminated and that will upload any remaining logs on the instance to Amazon S3.

C. Create an Amazon S3 lifecycle configuration to move log files from Amazon S3 to Amazon Glacier after seven days.

D. Configure your application to write logs to the instance's default Amazon EBS boot volume, because this storage already exists. Create a script that moves the logs from the instance to Amazon 53 once an hour.

E. Configure your application to write logs to a separate Amazon EBS volume with the "delete on termination" field set to false. Create a script that moves the logs from the instance to Amazon S3 once an hour.

F. Create a housekeeping script that runs on a T2 micro instance managed by an Auto Scaling group for high availability. The script uses the AWS API to identify any unattached Amazon EBS volumes containing log files. Your housekeeping script will mount the Amazon EBS volume, upload all logs to Amazon S3, and then delete the volume.

ANSWER: C E F

QUESTION NO: 2

You are running Amazon CloudTrail on an Amazon S3 bucket and look at your most recent log. You notice that the entries include the ListThings and CreateThings actions and wonder if your devices have been hacked. Based on these entries, what service would you be concerned may have been hacked?

A. Amazon Inspector

B. AWS IoT

C. AWS CodePipeline

D. Amazon Glacier

ANSWER: B

Explanation:

AWS IoT (Internet of Things) is integrated with CloudTrail to capture API calls from the AWS IoT console or from your code to the AWS IoT APIs. AWS IoT provides secure, bi-directional communication between Internet-connected things (such as sensors, actuators, embedded devices, or smart appliances) and the AWS cloud. Using the information collected by CloudTrail, you can determine the request that was made to AWS IoT, the source IP address from which the request was made, who made the request, when it was made, and so on.

Reference: http://docs.aws.amazon.com/iot/latest/developerguide/monitoring_overview.html#iot-usingcloudtrail

QUESTION NO: 3

You have been tasked with deploying a solution for your company that will store images, which the marketing department will use for its campaigns. Employees are able to upload images via a web interface, and once uploaded, each image must be resized and watermarked with the company logo. Image resize and watermark is not time-sensitive and can be completed days after upload if required. How should you design this solution in the most highly available and cost-effective way?

A. Configure your web application to upload images to the Amazon Elastic Transcoder service. Use the Amazon Elastic Transcoder watermark feature to add the company logo as a watermark on your images and then to upload the final images into an Amazon S3 bucket.

B. Configure your web application to upload images to Amazon S3, and send the Amazon S3 bucket URI to an Amazon SQS queue. Create an Auto Scaling group and configure it to use Spot instances, specifying a price you are willing to pay. Configure the instances in this Auto Scaling group to poll the SQS queue for new images and then resize and watermark the image before uploading the final images into Amazon S3.

C. Configure your web application to upload images to Amazon S3, and send the S3 object URI to an Amazon SQS queue. Create an Auto Scaling launch configuration that uses Spot instances, specifying a price you are willing to pay. Configure the instances in this Auto Scaling group to poll the Amazon SQS queue for new images and then resize and watermark the image before uploading the new images into Amazon S3 and deleting the message from the Amazon SQS queue.

D. Configure your web application to upload images to the local storage of the web server. Create a cronjob to execute a script daily that scans this directory for new files and then uses the Amazon EC2 Service API to launch 10 new Amazon EC2 instances, which will resize and watermark the images daily.

ANSWER: C

QUESTION NO: 4

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

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

A. Create an AWS CloudFormation template to provision the application server and database instance in a single Region.

B. Create an AWS CloudFormation template to provision the application tier of the application and a multi-Region database instance.

C. Configure Amazon CloudWatch Events rules to run every hour. Trigger AWS Lambda functions to create an RDS snapshot and copy it to the secondary Region.

D. Configure Amazon CloudWatch Events rules to run every 3 hours. Trigger AWS Lambda functions to create an RDS snapshot and copy it to the secondary Region.

E. In the event of a failure, deploy a new AWS CloudFormation stack in a secondary region to provision the application resources and a new RDS instance using the copied snapshot and a Route 53 failover routing policy.

F. In the event of a failure, deploy a new AWS CloudFormation stack in a secondary region to provision the application resources and a replica of the RDS database using the copied snapshot and a Route 53 latency-based routing policy.

ANSWER: B D E

QUESTION NO: 5

Which of the following Dockerfile commands cannot be overridden at runtime?

A. VOLUME

B. USER

C. ADD

D. CMD

ANSWER: C

Explanation:

When a developer builds an image from a Dockerfile or when she commits it, the developer can set a number of default parameters that take effect when the image starts up as a container. Four of the Dockerfile commands cannot be overridden at runtime: FROM, MAINTAINER, RUN, and ADD. Everything else has a corresponding override in docker run. We'll go through what the developer might have set in each Dockerfile instruction and how the operator can override that setting.

Reference:

https://docs.docker.com/engine/reference/run/#overriding-dockerfile-image-defaults

QUESTION NO: 6

Your application uses Amazon SQS and Auto Scaling to process background jobs. The Auto Scaling policy is based on the number of messages in the queue, with a maximum Instance count of 100. Since the application was launched, the group has never scaled above 50. The Auto Scaling group has now scaled to 100, the queue size is increasing, and very few Jobs are being completed. The number of messages being sent to the queue is at normal levels.

What should you do to identify why the queue size is unusually high, and to reduce it?

A. Temporarily increase the Auto Scaling group's desired value to 200. When the queue size has been reduced, reduce it to 50.

B. Analyze the application logs to identify possible reasons for message processing failure and resolve the cause for failures.

C. Create additional Auto Scaling groups, enabling the processing of the queue to be performed in parallel.

D. Analyze CloudTrail logs for Amazon SQS to ensure that the instances' Amazon EC2 role has permission to receive messages from the queue.

ANSWER: B

QUESTION NO: 7

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? (Choose three.)

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

Explanation:

Reference: https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-paramstore-access.html

QUESTION NO: 8

The development team is creating a social media game which ranks users on a scoreboard. The current implementation uses an Amazon RDS for MySQL database for storing user data; however, the game cannot display scores quickly enough during performance testing.

Which service would provide the fastest retrieval times?

A. Migrate user data to Amazon DynamoDB for managing content.

B. Use AWS Batch to compute and deliver user and score content.

C. Deploy Amazon CloudFront for user and score content delivery.

D. Set up Amazon ElastiCache to deliver user and score content.

ANSWER: D

QUESTION NO: 9

A DevOps team wants to be able to work on the same source code repository. The team has the following requirements for their development workflow and repository access controls:

- Only team members can clone the repository and create new branches.
- A production-ready code state should be isolated from any untested code changes.
- Code changes should be approved by another team member before merging to the production-ready master branch.
- All code change approvals must have an audit record.
- New team members can quickly modify code.

Which combination of actions will these requirements? (Choose three.)

A. Check out the master branch and develop new features locally on a feature branch to keep the production-ready code isolated. Ask team members to review the changes before committing the changes locally.

B. Create an AWS CodeCommit repository and an IAM group with permissions to read/write changes to the repository. Add new team members to this group.

C. Create an AWS CodeCommit repository and an IAM role with permissions to read/write changes to the repository. Attach this IAM role to a single IAM user. Ensure each member of the team uses this IAM user. Provide new team members the credentials to this IAM user.

D. Create a local feature branch from the master branch for new features. Commit the new code and push the changes to the feature branch in the repository.

E. Create a pull request so other team members can review the code changes. Implement any suggestions, pull any additional changes from the master branch, and push to the feature branch again. Merge the master branch with the feature branch.

F. Create a pull request so other team members can review the code changes. Implement any suggestions, pull any additional changes from the master branch, resolve any conflicts, and push to the feature branch again. Merge the feature branch with the master branch.

ANSWER: A B C

QUESTION NO: 10

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

Which combination of actions should be taken to implement a more stable deployment solution? (Choose two.)

A. Create a pipeline in AWS CodePipeline with CodeCommit as a source provider. Create parallel pipeline stages to build and test the application. Pass the build artifact to AWS CodeDeploy.

B. Create a pipeline in AWS CodePipeline with CodeCommit as a source provider. Create separate pipeline stages to build and then test the application. Pass the build artifact to AWS CodeDeploy.

C. Create and use an AWS CodeDeploy application and deployment group to deploy code updates to the EC2 fleet. Select the Application Load Balancer for the deployment group.

D. Create individual Lambda functions to run all build, test, and deploy actions using AWS CodeDeploy instead of AWS Systems Manager.

E. Modify the Lambda function to build a single application package to be shared by all instances. Use AWS CodeDeploy instead of AWS Systems Manager to update the code on the EC2 fleet.

ANSWER: A D

QUESTION NO: 11

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? (Choose 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.

ANSWER: A D F

QUESTION NO: 12

Currently, your deployment process consists of setting your load balancer to point to a maintenance page, turning off ea web application servers, deploying your code, turning the web application servers back on, and removing the maintenance page. Working with your development team, you've agreed that performing rolling deployments of your software would provide a better user experience and a more agile deployment process.

Which techniques could you use to provide a cost-effective rolling deployment process? (Choose two.)

A. Use the Amazon Elastic Cloud Compute (EC2) API to write a service to return a list of servers based on the tags for the application that needs deployment, and use Amazon Simple Queue Service to queue up all servers for a rolling deployment.

B. Re-deploy your application on AWS Elastic Beanstalk, and use Elastic Beanstalk rolling deployments.

C. Re-deploy your application on an AWS OpsWorks stack, and take advantage of OpsWorks rolling deployments.

D. Re-deploy your application using an AWS CloudFormation template, launch a new CloudFormation stack during each deployment, and then tear down the old stack.

E. Re-deploy your application using an AWS CloudFormation template with Auto Scaling group, and use update policies to provide rolling updates.

F. Using Amazon Simple Workflow Service, create a workflow application that talks to the Amazon EC2 API to deploy your new code in a rolling fashion.

ANSWER: B E

QUESTION NO: 13

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 and deploy 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 AWS CodeDeploy 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 AWS CodeDeploy 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 a blue/green deployment.

ANSWER: A

QUESTION NO: 14

A company manages an application that stores logs in Amazon CloudWatch Logs. The company wants to archive the logs in Amazon S3. Logs are rarely accessed after 90 days and must be retained for 10 years.

Which combination of steps should a DevOps engineer take to meet these requirements? (Choose two.)

A. Configure a CloudWatch Logs subscription filter to use AWS Glue to transfer all logs to an S3 bucket.

B. Configure a CloudWatch Logs subscription filter to use Amazon Kinesis Data Firehose to stream all logs to an S3 bucket.

C. Configure a CloudWatch Logs subscription filter to stream all logs to an S3 bucket.

D. Configure the S3 bucket lifecycle policy to transition logs to S3 Glacier after 90 days and to expire logs after 3.650 days.

E. Configure the S3 bucket lifecycle policy to transition logs to Reduced Redundancy after 90 days and to expire logs after 3.650 days.

ANSWER: B C

Explanation:

Reference: https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html

QUESTION NO: 15

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? (Choose 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 UpdatePolicy attribute for the CloudFormation stack and specify the Auto Scaling group that will be updated. Configure MinSuccessfulInstancesPercent and PauseTime to ensure the deployment happens with zero downtime.

ANSWER: A C

QUESTION NO: 16

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 the Finance team.

ANSWER: D

QUESTION NO: 17

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

- Focus first on the development workload.
- The environment must be easy to manage.
- Deployment should be repeatable and reusable for new environments. ▪ Store the code in a GitHub repository.

Which solution will meet these requirements?

A. Set up an Amazon ECS environment. Use AWS CodePipeline to create a pipeline that is triggered on a commit to the GitHub repository. Use AWS CodeBuild to create the container images and AWS CodeDeploy to publish the container image to the ECS environment.

B. Use AWS CodePipeline that triggers on a commit from the GitHub repository, build the container images with AWS CodeBuild, and publish the container images to Amazon ECR. In the final stage, use AWS CloudFormation to create an Amazon ECS environment that gets the container images from the ECR repository.

C. Create a Kubernetes Cluster on Amazon EC2. Use AWS CodePipeline to create a pipeline that is triggered when the code is committed to the repository. Create the container images with a Jenkins server on EC2 and store them in the Docker Hub. Use AWS Lambda from the pipeline to trigger the deployment to the Kubernetes Cluster.

D. Set up an Amazon ECS environment. Use AWS CodePipeline to create a pipeline that is triggered on a commit to the GitHub repository. Use AWS CodeBuild to create the container and store it in the Docker Hub. Use an AWS Lambda function to trigger a deployment and pull the new container image from the Docker Hub.

ANSWER: A

QUESTION NO: 18

When building a Docker image, you are searching through a persistent data volume's logs to provide parameters for the next build. You execute the following command. Which of the operations will cause a failure of the Docker RUNcommand? RUN cat ./data/log/*.error | grep service_status | grep ERROR

A. the first grep command

B. any one of them

C. the second grep command

D. the cat command

ANSWER: C

Explanation:

Some RUN commands depend on the ability to pipe the output of one command into another, using the pipe character (|), as in the following example:

RUN wget -O - https://some.site | wc -l > /number

Docker executes these commands using the /bin/sh -c interpreter, which only evaluates the exit code of the last operation in the pipe to determine success. In the example above this build step succeeds and produces a new image so long as the wc -lcommand succeeds, even if the wget command fails.

Reference:

https://docs.docker.com/engine/userguide/eng-image/dockerfile_best-practices/#run

QUESTION NO: 19

You are building an AWS CloudFormation template for a multi-tier web application. The user data of your Linux web server resource contains a complex script that can take a long time to run. Which techniques could you use to ensure that these servers are fully configured and running before attaching them to the load balancer? (Choose two.)

A. Launch your Linux servers from a nested stack that is called from within the load balancer resource in your AWS CloudFormation template.

B. Add an AWS CloudFormation Wait Condition that depends on the web server resource. When the UserData script finishes on the web servers, use curl to send a signal the Wait Condition at http://169.254.169.254/waithandle/.

C. Add an AWS CloudFormation wait Condition that depends on the web server resource. When the UserData script finishes on the web servers, use curl to signal to the Wait Condition pre-signed URL that they are ready.

D. In your AWS CloudFormation template, position the load balancer resource JSON block directly below your Linux server resource.

E. Add an AWS CloudFormation Wait Condition that depends on the web server resource. When the UserData script finishes on the web servers, use the command "cfn-signal" to signal that they are ready.

ANSWER: C E

QUESTION NO: 20

The Ansible Inventory system allows many attributes to be defined within it. Which item below is not one of these?

A. Group variables

B. Host groups

C. Include vars

D. Children groups

ANSWER: C

Explanation:

Ansible inventory files cannot reference other files for additional data. If this functionality is needed, it must be done in as a script to create a dynamic inventory.

Reference:

http://docs.ansible.com/ansible/intro_inventory.html