DOP-C02$^{Q\&As}$

AWS Certified DevOps Engineer - Professional

Pass Amazon DOP-C02 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.certbus.com/dop-c02.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

A company recently migrated its legacy application from on-premises to AWS. The application is hosted on Amazon EC2 instances behind an Application Load Balancer which is behind Amazon API Gateway. The company wants to ensure users experience minimal disruptions during any deployment of a new version of the application. The company also wants to ensure it can quickly roll back updates if there is an issue.

Which solution will meet these requirements with MINIMAL changes to the application?

A. Introduce changes as a separate environment parallel to the existing one Configure API Gateway to use a canary release deployment to send a small subset of user traffic to the new environment.

B. Introduce changes as a separate environment parallel to the existing one Update the application\\'s DNS alias records to point to the new environment.

C. Introduce changes as a separate target group behind the existing Application Load Balancer Configure API Gateway to route user traffic to the new target group in steps.

D. Introduce changes as a separate target group behind the existing Application Load Balancer Configure API Gateway to route all traffic to the Application Load Balancer which then sends the traffic to the new target group.

Correct Answer: A

API Gateway supports canary deployment on a deployment stage before you direct all traffic to that stage. A parallel environment means we will create a new ALB and a target group that will target a new set of EC2 instances on which the newer version of the app will be deployed. So the canary setting associated to the new version of the API will connect with the new ALB instance which in turn will direct the traffic to the new EC2 instances on which the newer version of the application is deployed.

QUESTION 2

A company uses a single AWS account lo test applications on Amazon EC2 instances. The company has turned on AWS Config in the AWS account and has activated the restricted-ssh AWS Config managed rule.

The company needs an automated monitoring solution that will provide a customized notification in real time if any security group in the account is not compliant with the restricted-ssh rule. The customized notification must contain the name and ID of the noncompliant security group.

A DevOps engineer creates an Amazon Simple Notification Service (Amazon SNS) topic in the account and subscribes the appropriate personnel to the topic.

What should me DevOps engineer do next to meet these requirements?

A. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT tor the restricted-ssh rule. Configure an input transformer for the EventBridge rule Configure the EventBridge rule to publish a notification to the SNS topic.

B. Configure AWS Config to send all evaluation results for the restricted-ssh rule to the SNS topic. Configure a filter policy on the SNS topic to send only notifications that contain the text of NON_COMPLIANT in the notification to subscribers.

C. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule Configure the EventBridge rule to invoke AWS Systems Manager Run Command on the SNS topic to

customize a notification and to publish the notification to the SNS topic

D. Create an Amazon EventBridge rule that matches all AWS Config evaluation results of NON_COMPLIANT Configure an input transformer for the restricted-ssh rule Configure the EventBridge rule to publish a notification to the SNS topic.

Correct Answer: A

Create an Amazon EventBridge (Amazon CloudWatch Events) rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule. Configure an input transformer for the EventBridge (CloudWatch Events) rule. Configure the EventBridge (CloudWatch Events) rule to publish a notification to the SNS topic. This approach uses Amazon EventBridge (previously known as Amazon CloudWatch Events) to filter AWS Config evaluation results based on the restricted-ssh rule and its compliance status (NON_COMPLIANT). An input transformer can be used to customize the information contained in the notification, such as the name and ID of the noncompliant security group. The EventBridge (CloudWatch Events) rule can then be configured to publish a notification to the SNS topic, which will notify the appropriate personnel in real-time.

QUESTION 3

A company releases a new application in a new AWS account. The application includes an AWS Lambda function that processes messages from an Amazon Simple Queue Service (Amazon SQS) standard queue. The Lambda function stores the results in an Amazon S3 bucket for further downstream processing. The Lambda function needs to process the messages within a specific period of time after the messages are published. The Lambda function has a batch size of 10 messages and takes a few seconds to process a batch of messages.

As load increases on the application\\'s first day of service, messages in the queue accumulate at a greater rate than the Lambda function can process the messages. Some messages miss the required processing timelines. The logs show that many messages in the queue have data that is not valid. The company needs to meet the timeline requirements for messages that have valid data.

Which solution will meet these requirements?

A. Increase the Lambda function\\'s batch size. Change the SQS standard queue to an SQS FIFO queue. Request a Lambda concurrency increase in the AWS Region.

B. Reduce the Lambda function\\'s batch size. Increase the SQS message throughput quota. Request a Lambda concurrency increase in the AWS Region.

C. Increase the Lambda function\\'s batch size. Configure S3 Transfer Acceleration on the S3 bucket. Configure an SQS dead-letter queue.

D. Keep the Lambda function\\'s batch size the same. Configure the Lambda function to report failed batch items. Configure an SQS dead-letter queue.

Correct Answer: D

QUESTION 4

A company uses AWS Secrets Manager to store a set of sensitive API keys that an AWS Lambda function uses. When the Lambda function is invoked, the Lambda function retrieves the API keys and makes an API call to an external service. The Secrets Manager secret is encrypted with the default AWS Key Management Service (AWS KMS) key.

A DevOps engineer needs to update the infrastructure to ensure that only the Lambda function\\'s execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege.

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

A. Update the default KMS key for Secrets Manager to allow only the Lambda function\\'s execution role to decrypt.

B. Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function\\'s execution role to decrypt. Update Secrets Manager to use the new customer managed key.

C. Create a KMS customer managed key that trusts Secrets Manager and allows the account\\'s :root principal to decrypt. Update Secrets Manager to use the new customer managed key.

D. Ensure that the Lambda function\\'s execution role has the KMS permissions scoped on the resource level. Configure the permissions so that the KMS key can encrypt the Secrets Manager secret.

E. Remove all KMS permissions from the Lambda function\\'s execution role.

Correct Answer: BD

The requirement is to update the infrastructure to ensure that only the Lambda function\\'s execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege, which means granting the minimum permissions necessary to perform a task. To do this, the DevOps engineer needs to use the following steps: Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function\\'s execution role to decrypt. A customer managed key is a symmetric encryption key that is fully managed by the customer. The customer can define the key policy, which specifies who can use and manage the key. By creating a customer managed key, the DevOps engineer can restrict the decryption permission to only the Lambda function\\'s execution role, and prevent other principals from accessing the secret values. The customer managed key also needs to trust Secrets Manager, which means allowing Secrets Manager to use the key to encrypt and decrypt secrets on behalf of the customer. Update Secrets Manager to use the new customer managed key. Secrets Manager allows customers to choose which KMS key to use for encrypting each secret. By default, Secrets Manager uses the default KMS key for Secrets Manager, which is a service-managed key that is shared by all customers in the same AWS Region. By updating Secrets Manager to use the new customer managed key, the DevOps engineer can ensure that only the Lambda function\\'s execution role can decrypt the secret values using that key. Ensure that the Lambda function\\'s execution role has the KMS permissions scoped on the resource level. The Lambda function\\'s execution role is an IAM role that grants permissions to the Lambda function to access AWS services and resources. The role needs to have KMS permissions to use the customer managed key for decryption. However, to apply the principle of least privilege, the role should have the permissions scoped on the resource level, which means specifying the ARN of the customer managed key as a condition in the IAM policy statement. This way, the role can only use that specific key and not any other KMS keys in the account.

QUESTION 5

You have an application running a specific process that is critical to the application\\'s functionality, and have added the health check process to your Auto Scaling Group. The instances are showing healthy but the application itself is not working as it should. What could be the issue with the health check, since it is still showing the instances as healthy.

A. You do not have the time range in the health check properly configured

B. It is not possible for a health check to monitor a process that involves the application

C. The health check is not configured properly

D. The health check is not checking the application process

Correct Answer: D

If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use this information. For example, if you determine that an instance is not functioning as expected, you can

set the health status of the instance to Unhealthy. The next time that Auto Scaling performs a health check on the instance, it will determine that the instance is unhealthy and then launch a replacement instance.

QUESTION 6

You have just recently deployed an application on EC2 instances behind an ELB. After a couple of weeks, customers are complaining on receiving errors from the application. You want to diagnose the errors and are trying to get errors from the ELB access logs. But the ELB access logs are empty. What is the reason for this.

A. You do not have the appropriate permissions to access the logs

B. You do not have your CloudWatch metrics correctly configured

C. ELB Access logs are only available for a maximum of one week

D. Access logging is an optional feature of Elastic Load Balancing that is disabled by default

Correct Answer: D

Clastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Cach log contains information such as the time the request was received, the client\\'s IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot issues. Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer. Clastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time.

QUESTION 7

A company runs an application on one Amazon EC2 instance. Application metadata is stored in Amazon S3 and must be retrieved if the instance is restarted. The instance must restart or relaunch automatically if the instance becomes unresponsive.

Which solution will meet these requirements?

A. Create an Amazon CloudWatch alarm for the StatusCheckFailed metric. Use the recover action to stop and start the instance. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.

B. Configure AWS OpsWorks, and use the auto healing feature to stop and start the instance. Use a lifecycle event in OpsWorks to pull the metadata from Amazon S3 and update it on the instance.

C. Use EC2 Auto Recovery to automatically stop and start the instance in case of a failure. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.

D. Use AWS CloudFormation to create an EC2 instance that includes the UserData property for the EC2 resource. Add a command in UserData to retrieve the application metadata from Amazon S3.

Correct Answer: B

QUESTION 8

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 WebsiteConfiguration property 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: B

QUESTION 9

If designing a single playbook to run across multiple Linux distributions that have distribution specific commands, what would be the best method to allow a successful run?

A. Enable fact gathering and use the `when\\' conditional to match the distribution to the task.

B. This is not possible, a separate playbook for each target Linux distribution is required.

C. Use `ignore_errors: true\\' in the tasks.

D. Use the `shell\\' module to write your own checks for each command that is ran.

Correct Answer: A

Ansible provides a method to only run a task when a condition is met using the `when\\' declarative. With gather facts enabled, the play has access to the distribution name of the Linux system, thus, tasks can be tailored to a specific distribution and ran only when the condition is met, e.g.: ` - when: ansible_os_family == "Debian"\\'.

Reference: http://docs.ansible.com/ansible/playbooks_conditionals.html

QUESTION 10

You have an application which consists of EC2 instances in an Auto Scaling group. Between a particular time frame every day, there is an increase in traffic to your website. Hence users are complaining of a poor response time on the

application. You have configured your Auto Scaling group to deploy one new EC2 instance when CPU utilization is greater than 60% for 2 consecutive periods of 5 minutes.

What is the least cost-effective way to resolve this problem?

A. Decrease the consecutive number of collection periods

B. Increase the minimum number of instances in the Auto Scaling group

C. Decrease the collection period to ten minutes

D. Decrease the threshold CPU utilization percentage at which to deploy a new instance

Correct Answer: B

If you increase the minimum number of instances, then they will be running even though the load is not high on the website. Hence you are incurring cost even though there is no need. All of the remaining options are possible options which can be used to increase the number of instances on a high load. For more information on On-demand scaling, please refer to the below link.

Reference: http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scale-based-on-demand.html

QUESTION 11

A company has developed an AWS Lambda function that handles orders received through an API. The company is using AWS CodeDeploy to deploy the Lambda function as the final stage of a CI/CD pipeline.

A DevOps engineer has noticed there are intermittent failures of the ordering API for a few seconds after deployment. After some investigation the DevOps engineer believes the failures are due to database changes not having fully propagated before the Lambda function is invoked

How should the DevOps engineer overcome this?

A. Add a BeforeAllowTraffic hook to the AppSpec file that tests and waits for any necessary database changes before traffic can flow to the new version of the Lambda function.

B. Add an AfterAllowTraffic hook to the AppSpec file that forces traffic to wait for any pending database changes before allowing the new version of the Lambda function to respond.

C. Add a BeforeAllowTraffic hook to the AppSpec file that tests and waits for any necessary database changes before deploying the new version of the Lambda function.

D. Add a validateService hook to the AppSpec file that inspects incoming traffic and rejects the payload if dependent services such as the database are not yet ready.

Correct Answer: A

https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html#appspec-hooks-lambda

QUESTION 12

A software team is using AWS CodePipeline to automate its Java application release pipeline. The pipeline consists of a source stage, then a build stage, and then a deploy stage. Each stage contains a single action that has a runOrder value of 1.

The team wants to integrate unit tests into the existing release pipeline. The team needs a solution that deploys only the code changes that pass all unit tests.

Which solution will meet these requirements?

A. Modify the build stage. Add a test action that has a runOrder value of 1. Use AWS CodeDeploy as the action provider to run unit tests.

B. Modify the build stage. Add a test action that has a runOrder value of 2. Use AWS CodeBuild as the action provider to run unit tests.

C. Modify the deploy stage. Add a test action that has a runOrder value of 1. Use AWS CodeDeploy as the action provider to run unit tests.

D. Modify the deploy stage. Add a test action that has a runOrder value of 2. Use AWS CodeBuild as the action provider to run unit tests.

Correct Answer: B

B is correct: Runorder value of 2 ensure that we do unit tests after we build artifacts.

A: The unit tests would run in parellel with the build step, which is incorrect. We can only test after we have done building C and D: The unit tests would not run before the deploy step.

QUESTION 13

Which answer is the proper syntax for specifying two target hosts on the command line when running an Ansible Playbook?

A. ansible-playbook -h host1.example.com -i all playbook.yml

B. ansible-playbook -i host1.example.com playbook.yml

C. ansible-playbook -h host1.example.com,host2.example.com playbook.yml

D. ansible-playbook -i host1.example.com,host2.example.com playbook.yml

Correct Answer: D

Ansible uses the `-i\\' flag for accepting an inventory file or host. To allow Ansible to determine if you are passing a host list versus an inventory file the list must be comma separated. If a single host is specified, a trailing comma must be present.

Reference: http://docs.ansible.com/ansible/intro_inventory.html#inventory

---

QUESTION 14

An AWS CodePipeline pipeline has implemented a code release process. The pipeline is integrated with AWS CodeDeploy to deploy versions of an application to multiple Amazon EC2 instances for each CodePipeline stage.

During a recent deployment, the pipeline failed due to a CodeDeploy issue. The DevOps team wants to improve monitoring and notifications during deployment to decrease resolution times.

What should the DevOps engineer do to create notifications. When issues are discovered?

A. Implement Amazon CloudWatch Logs for CodePipeline and CodeDeploy, create an AWS Config rule to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.

B. Implement Amazon EventBridge for CodePipeline and CodeDeploy, create an AWS Lambda function to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.

C. Implement AWS CloudTrail to record CodePipeline and CodeDeploy API call information, create an AWS Lambda function to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.

D. Implement Amazon EventBridge for CodePipeline and CodeDeploy, create an Amazon Inspector assessment target to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.

Correct Answer: B

AWS CloudWatch Events can be used to monitor events across different AWS resources, and a CloudWatch Event Rule can be created to trigger an AWS Lambda function when a deployment issue is detected in the pipeline. The Lambda function can then evaluate the issue and send a notification to the appropriate stakeholders through an Amazon SNS topic. This approach allows for real-time notifications and faster resolution times.

QUESTION 15

A company is running a custom-built application that processes records. All the components run on Amazon EC2 instances that run in an Auto Scaling group. Each record\\'s processing is a multistep sequential action that is compute-intensive. Each step is always completed in 5 minutes or less.

A limitation of the current system is that if any steps fail, the application has to reprocess the record from the beginning. The company wants to update the architecture so that the application must reprocess only the failed steps.

What is the MOST operationally efficient solution that meets these requirements?

A. Create a web application to write records to Amazon S3. Use S3 Event Notifications to publish to an Amazon Simple Notification Service (Amazon SNS) topic. Use an EC2 instance to poll Amazon SNS and start processing. Save intermediate results to Amazon S3 to pass on to the next step.

B. Perform the processing steps by using logic in the application. Convert the application code to run in a container. Use AWS Fargate to manage the container instances. Configure the container to invoke itself to pass the state from one step to the next.

C. Create a web application to pass records to an Amazon Kinesis data stream. Decouple the processing by using the Kinesis data stream and AWS Lambda functions.

D. Create a web application to pass records to AWS Step Functions. Decouple the processing into Step Functions tasks and AWS Lambda functions.

Correct Answer: D