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.leads4pass.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 uses an organization in AWS Organizations that has all features enabled. The company uses AWS Backup in a primary account and uses an AWS Key Management Service (AWS KMS) key to encrypt the backups.

The company needs to automate a cross-account backup of the resources that AWS Backup backs up in the primary account. The company configures cross-account backup in the Organizations management account. The company creates a new AWS account in the organization and configures an AWS Backup backup vault in the new account. The company creates a KMS key in the new account to encrypt the backups. Finally, the company configures a new backup plan in the primary account. The destination for the new backup plan is the backup vault in the new account.

When the AWS Backup job in the primary account is invoked, the job creates backups in the primary account. However, the backups are not copied to the new account\\'s backup vault.

Which combination of steps must the company take so that backups can be copied to the new account\\'s backup vault? (Select TWO.)

A. Edit the backup vault access policy in the new account to allow access to the primary account.

B. Edit the backup vault access policy in the primary account to allow access to the new account.

C. Edit the backup vault access policy in the primary account to allow access to the KMS key in the new account.

D. Edit the key policy of the KMS key in the primary account to share the key with the new account.

E. Edit the key policy of the KMS key in the new account to share the key with the primary account.

Correct Answer: AE

To enable cross-account backup, the company needs to grant permissions to both the backup vault and the KMS key in the destination account. The backup vault access policy in the destination account must allow the primary account to

copy backups into the vault. The key policy of the KMS key in the destination account must allow the primary account to use the key to encrypt and decrypt the backups. These steps are described in the AWS documentation. Therefore, the

correct answer is A and E.

References:

1: Creating backup copies across AWS accounts -AWS Backup

2: Using AWS Backup with AWS Organizations -AWS Backup

QUESTION 2

A DevOps engineer has developed an AWS Lambda function The Lambda function starts an AWS CloudFormation drift detection operation on all supported resources for a specific CloudFormation stack The Lambda function then exits Its invocation The DevOps engineer has created an Amazon EventBrdge scheduled rule that Invokes the Lambda function every hour. An Amazon Simple Notification Service (Amazon SNS) topic already exists In the AWS account. The DevOps engineer has subscribed to the SNS topic to receive notifications

The DevOps engineer needs to receive a notification as soon as possible when drift is detected in this specific stack configuration.

Which solution Will meet these requirements?

A. Configure the existing EventBridge rule to also target the SNS topic Configure an SNS subscription filter policy to match the Cloud Formation stack. Attach the subscription filter policy to the SNS tomc

B. Create a second Lambda function to query the CloudFormation API for the drift detection results for the stack Configure the second Lambda function to publish a message to the SNS topic If drift ts detected Adjust the existing EventBridge rule to also target the second Lambda function

C. Configure Amazon GuardDuty in the account with drift detection for all CloudFormation stacks. Create a second EventBndge rule that reacts to the GuardDuty drift detection event finding for the specific CloudFormation stack. Configure the SNS topic as a target of the second EventBridge rule.

D. Configure AWS Config in the account. Use the cloudformation-stack-drift-detection-check managed rule. Create a second EventBndge rule that reacts to a compliance change event for the CloudFormaUon stack. Configure the SNS topc as a target of the second EventBridge rule.

Correct Answer: D

A comprehensive and detailed explanation is: Option A is incorrect because EventBridge rules cannot filter events based on the message body or attributes of the target service. Therefore, configuring an SNS subscription filter policy to match the CloudFormation stack will not work. The SNS topic will receive all events from the EventBridge rule, regardless of the stack name or drift status. Option B is incorrect because it introduces unnecessary complexity and cost. Creating a second Lambda function to query the CloudFormation API for the drift detection results is redundant, since CloudFormation already publishes drift detection events to EventBridge. Moreover, invoking two Lambda functions every hour will incur more charges than invoking one. Option C is incorrect because GuardDuty does not provide drift detection for CloudFormation stacks. GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior in AWS accounts and workloads. It does not monitor or report on configuration changes or drifts in CloudFormation stacks. Option D is correct because it leverages AWS Config and its managed rule for drift detection. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. It can detect configuration changes and drifts in CloudFormation stacks using the cloudformation-stack-drift-detection-check managed rule. This rule triggers an AWS Config event when a stack drifts from its expected template configuration. By creating a second EventBridge rule that reacts to this event for the specific stack, the DevOps engineer can configure the SNS topic as a target and receive a notification as soon as possible when drift is detected. References: AWS Config Amazon SNS subscription filter policies Amazon EventBridge rules

QUESTION 3

Which of the following is an invalid variable name in Ansible?

A. host1st_ref

B. host-first-ref

C. Host1stRef

D. host_first_ref

Correct Answer: B

Variable names can contain letters, numbers and underscores and should always start with a letter. Invalid variable examples, `host first ref\\', `1st_host_ref\\'\\'.

Reference: http://docs.ansible.com/ansible/playbooks_variables.html#what-makes-a-valid-variable-name

QUESTION 4

Which is the proper syntax for referencing a variable\\'s value in an Ansible task?

A. ${variable_name}

B. { variable_name }

C. "{{ variable_name }}"

D. @variable_name

Correct Answer: C

We use the variable\\'s name to reference the variable which we encapsulate in curly brackets `{{ }}\\'; however, the YAML syntax dictates that a string beginning with a curly bracket denotes a dictionary value. To get around this, it is proper to wrap the variable declaration in quotes.

Reference: http://docs.ansible.com/ansible/playbooks_variables.html#hey-wait-a-yaml-gotcha

QUESTION 5

A company needs a strategy for failover and disaster recovery of its data and application. The application uses a MySQL database and Amazon EC2 instances. The company requires a maximum RPO of 2 hours and a maximum RTO of 10 minutes for its data and application at all times.

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

A. Create an Amazon Aurora Single-AZ cluster in multiple AWS Regions as the data store. Use Aurora\\'s automatic recovery capabilities in the event of a disaster.

B. Create an Amazon Aurora global database in two AWS Regions as the data store. In the event of a failure, promote the secondary Region to the primary for the application. Update the application to use the Aurora cluster endpoint in the secondary Region.

C. Create an Amazon Aurora cluster in multiple AWS Regions as the data store. Use a Network Load Balancer to balance the database traffic in different Regions.

D. Set up the application in two AWS Regions. Use Amazon Route 53 failover routing that points to Application Load Balancers in both Regions. Use health checks and Auto Scaling groups in each Region.

E. Set up the application in two AWS Regions. Configure AWS Global Accelerator to point to Application Load Balancers (ALBs) in both Regions. Add both ALBs to a single endpoint group. Use health checks and Auto Scaling groups in each Region.

Correct Answer: BE

To meet the requirements of failover and disaster recovery, the company should use the following deployment strategies: Create an Amazon Aurora global database in two AWS Regions as the data store. In the event of a failure, promote the secondary Region to the primary for the application. Update the application to use the Aurora cluster endpoint in the secondary Region. This strategy can provide a low RPO and RTO for the data, as Aurora global database replicates data with minimal latency across Regions and allows fast and easy failover12. The company can use the Amazon Aurora cluster endpoint to connect to the current primary DB cluster without needing to change any

application code1. Set up the application in two AWS Regions. Configure AWS Global Accelerator to point to Application Load Balancers (ALBs) in both Regions. Add both ALBs to a single endpoint group. Use health checks and Auto Scaling groups in each Region. This strategy can provide high availability and performance for the application, as AWS Global Accelerator uses the AWS global network to route traffic to the closest healthy endpoint3. The company can also use static IP addresses that are assigned by Global Accelerator as a fixed entry point for their application1. By using health checks and Auto Scaling groups, the company can ensure that their application can scale up or down based on demand and handle any instance failures4. The other options are incorrect because: Creating an Amazon Aurora Single-AZ cluster in multiple AWS Regions as the data store would not provide a fast failover or disaster recovery solution, as the company would need to manually restore data from backups or snapshots in another Region in case of a failure. Creating an Amazon Aurora cluster in multiple AWS Regions as the data store and using a Network Load Balancer to balance the database traffic in different Regions would not work, as Network Load Balancers do not support cross-Region routing. Moreover, this strategy would not provide a consistent view of the data across Regions, as Aurora clusters do not replicate data automatically between Regions unless they are part of a global database. Setting up the application in two AWS Regions and using Amazon Route 53 failover routing that points to Application Load Balancers in both Regions would not provide a low RTO, as Route 53 failover routing relies on DNS resolution, which can take time to propagate changes across different DNS servers and clients. Moreover, this strategy would not provide deterministic routing, as Route 53 failover routing depends on DNS caching behavior, which can vary depending on different factors.