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

AWS Certified DevOps Engineer - Professional

Pass Amazon DOP-C02 Exam with 100% Guarantee

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

QUESTION 1

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

https://aws.amazon.com/blogs/mt/how-to-set-up-aws-opsworks-stacks-auto-healing-notifications-in-amazon-cloudwatch-events/

QUESTION 2

A development team is building an ecommerce application and is using Amazon Simple Notification Service (Amazon SNS) to send order messages to multiple endpoints. One of the endpoints is an external HTTP endpoint that is not always available. The development team needs to receive a notification if an order message is not delivered to the HTTP endpoint.

What should a DevOps engineer do to meet these requirements?

A. Create an Amazon Simple Queue Service (Amazon SQS) queue. On the SNS topic, configure a redrive policy that sends undelivered messages to the SQS queue. Create an Amazon CloudWatch alarm for the new SQS queue to notify the development team when messages are delivered to the queue.

B. Create an Amazon Simple Queue Service (Amazon SQS) queue. On the HTTP endpoint subscription of the SNS topic, configure a redrive policy that sends undelivered messages to the SQS queue. Create an Amazon CloudWatch alarm for the new SQS queue to notify the development team when messages are delivered to the queue.

C. On the SNS topic, configure an HTTPS delivery policy that will retry delivery until the order message is delivered successfully. Configure the backoffFunction parameter in the policy to notify the development team when a message cannot be delivered within the set constraints.

D. On the HTTP endpoint subscription of the SNS topic, configure an HTTPS delivery policy that will retry delivery until the order message is delivered successfully. Configure the backoffFunction parameter in the policy to notify the development team when a message cannot be delivered within the set constraints.

Correct Answer: C

QUESTION 3

A DevOps engineer is planning to deploy a Ruby-based application to production. The application needs to interact with an Amazon RDS for MySQL database and should have automatic scaling and high availability. The stored data in the database is critical and should persist regardless of the state of the application stack.

The DevOps engineer needs to set up an automated deployment strategy for the application with automatic rollbacks. The solution also must alert the application team when a deployment fails.

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

A. Deploy the application on AWS Elastic Beanstalk. Deploy an Amazon RDS for MySQL DB instance as part of the Elastic Beanstalk configuration.

B. Deploy the application on AWS Elastic Beanstalk. Deploy a separate Amazon RDS for MySQL DB instance outside of Elastic Beanstalk.

C. Configure a notification email address that alerts the application team in the AWS Beanstalk configuration.

D. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule to monitor AWS Health events. Use an Amazon Simple Notification Service (Amazon SNS) topic as a target to alert the application team.

E. Use the immutable deployment method to deploy new application versions.

F. Use the rolling deployment method to deploy new application versions.

Correct Answer: AEF

QUESTION 4

A growing company manages more than 50 accounts in an organization in AWS Organizations. The company has configured its applications to send logs to Amazon CloudWatch Logs.

A DevOps engineer needs to aggregate logs so that the company can quickly search the logs to respond to future security incidents. The DevOps engineer has created a new AWS account for centralized monitoring.

Which combination of steps should the DevOps engineer take to make the application logs searchable from the monitoring account? (Select THREE.)

A. In the monitoring account, download an AWS CloudFormation template from CloudWatch to use in Organizations. Use CloudFormation StackSets in the organization\\'s management account to deploy the CloudFormation template to the entire organization.

B. Create an AWS CloudFormation template that defines an IAM role. Configure the role to allow logs-amazonaws.com to perform the logs:Link action if the aws:ResourceAccount property is equal to the monitoring account ID. Use CloudFormation StackSets in the organization\\'s management account to deploy the CloudFormation template to the entire organization.

C. Create an IAM role in the monitoring account. Attach a trust policy that allows logs.amazonaws.com to perform the iam:CreateSink action if the aws:PrincipalOrgld property is equal to the organization ID.

D. In the organization\\'s management account, enable the logging policies for the organization.

E. use CloudWatch Observability Access Manager in the monitoring account to create a sink. Allow logs to be shared with the monitoring account. Configure the monitoring account data selection to view the Observability data from the organization ID.

F. In the monitoring account, attach the CloudWatchLogsReadOnlyAccess AWS managed policy to an IAM role that can be assumed to search the logs.

Correct Answer: BCF

To aggregate logs from multiple accounts in an organization, the DevOps engineer needs to create a cross-account subscription that allows the monitoring account to receive log events from the sharing accounts.

To enable cross-account subscription, the DevOps engineer needs to create an IAM role in each sharing account that grants permission to CloudWatch Logs to link the log groups to the destination in the monitoring account. This can be done

using a CloudFormation template and StackSets to deploy the role to all accounts in the organization.

The DevOps engineer also needs to create an IAM role in the monitoring account that allows CloudWatch Logs to create a sink for receiving log events from other accounts. The role must have a trust policy that specifies the organization ID

as a condition.

Finally, the DevOps engineer needs to attach the CloudWatchLogsReadOnlyAccess policy to an IAM role in the monitoring account that can be used to search the logs from the cross-account subscription.

References:

1: Cross-account log data sharing with subscriptions

2: Create an IAM role for CloudWatch Logs in each sharing account

3: AWS CloudFormation StackSets

4: Create an IAM role for CloudWatch Logs in your monitoring account

5: CloudWatchLogsReadOnlyAccess policy

QUESTION 5

Amazon Inspector agent collects telemetry data during assessment run and sends this data to Amazon Inspector dedicated S3 bucket for analysis. How can you access telemetry data out of Amazon Inspector and how can you benefit from this data in securing your resources?

A. Telemetry data is kept in S3 and encrypted with a pre-assessment test key configured in KMS, as long as you have access to that key you can download and decrypt telemetry data.

B. Telemetry data is stored in Amazon Inspector dedicated S3 bucket that does NOT belong to your account, Amazon Inspector currently does NOT provide an API or an S3 bucket access mechanism to collected telemetry. Data is retained temporarily only to allow for assistance with support requests.

C. Telemetry data is saved on S3 bucket in your account, therefore telemetry data is accessible with proper permissions on that bucket.

D. Telemetry data is deleted immediately after assessment run, therefore data can NOT be accessed or analyzed by any other tools.

Correct Answer: B

The telemetry data stored in S3 is retained only to allow for assistance with support requests and is not used or aggregated by Amazon for any other purpose. After 30 days, telemetry data is permanently deleted per a standard Amazon Inspector-dedicated S3 bucket lifecycle policy. At present, Amazon Inspector does not provide an API or an S3 bucket access mechanism to collected telemetry.

Reference: https://docs.aws.amazon.com/inspector/latest/userguide/inspector_agents.html