Free Questions for DOP-C01 by ebraindumps

Shared by Livingston on 24-05-2024

Question 1

Question Type: MultipleChoice

A devops team uses AWS CloudFormation to build their infrastructure. The security team is concerned about sensitive parameters, such as passwords, being exposed.

Which combination of steps will enhance the security of AWS CloudFormation? (Select THREE.)

Options:

A- Create a secure string with AWS KMS and choose a KMS encryption key. Reference the ARN of the secure string, and give AWS CloudFormation permission to the KMS key for decryption.

B- Create secrets using the AWS Secrets Manager AWS::SecretsManager::Secret resource type. Reference the secret resource return attributes in resources that need a password, such as an Amazon RDS database.

C- Store sensitive static data as secure strings in the AWS Systems Manager Parameter Store. Use dynamic references in the resources that need access to the data.

D- Store sensitive static data in the AWS Systems Manager Parameter Store as strings. Reference the stored value using types of Systems Manager parameters.

E- Use AWS KMS to encrypt the CloudFormation template.

F- Use the CloudFormation NoEcho parameter property to mask the parameter value.

Answer:

A, B, D

Question 2

Question Type: MultipleChoice

A DevOps engineer is troubleshooting deployments to a new application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. Instances sometimes come online before they are ready, which is leading to increased error rates among users. The current health check configuration gives instances a 60-second grace period and considers instances healthy after two 200 response codes from /index.php, a page that may respond intermittently during the deployment process. The development team wants instances to come online as soon as possible.

Which strategy would address this issue?

Options:

A- Increase the instance grace period from 60 seconds to 180 seconds, and the consecutive health check requirement from 2 to 3.

B- Increase the instance grace period from 60 seconds to 120 seconds, and change the response code requirement from 200 to 204.

C- Modify the deployment script to create a /health-check.php file when the deployment begins, then modify the health check path to point to that file.

D- Modify the deployment script to create a /health-check.php file when all tasks are complete, then modify the health check path to point to that file.

Answer:

D

Question 3

Question Type: MultipleChoice

A company's security team discovers that IAM access keys were exposed in a public code repository. Moving forward, the DevOps team wants to implement a solution that will automatically disable any keys that are suspected of being compromised, and notify the security team.

Which solution will accomplish this?

Options:

A- Create an Amazon CloudWatch Events event for Amazon Made. Create an Amazon SNS topic with two subscriptions: one to notify the security team and another to trigger an AWS Lambda function that disables the access keys.

B- Enable Amazon GuardDuty and set up an Amazon CloudWatch Events rule event for GuardDuty. Trigger an AWS Lambda function to check if the event relates to compromised keys. If so, send a notification to the security team and disable the access keys.

C- Run an AWS CloudWatch Events rule every 5 minutes to invoke an AWS Lambda function that checks to see if the compromised tag for any access key is set to true. If so. notify the security team and disable the access keys.

D- Set up AWS Config and create an AWS CloudTrail event for AWS Config. Create an Amazon SNS topic with two subscriptions: one to notify the security team and another to trigger an AWS Lambda function that disables the access keys.

Answer:

B

Question 4

Question Type: MultipleChoice

A company uses AWS CodePipeline to manage and deploy infrastructure as code. The infrastructure is defined in AWS CloudFormation templates and is primarily comprised of multiple Amazon EC2 instances and Amazon RDS databases. The Security team has observed many operators creating inbound security group rules with a source CIDR of 0 0 0 0/0 and would like to proactively stop the deployment of rules with open CIDRs

The DevOps Engineer will implement a predeptoyment step that runs some security checks over the CloudFormation template before the pipeline processes it. This check should allow only inbound security group rules with a source CIDR of 0.0.0.0/0 if the rule has the description "Security Approval Ref XXXXX (where XXXXX is a preallocated reference). The pipeline step should fail if this condition is not met and the deployment should be blocked

How should this be accomplished?

Options:

A- Enable a SCP in AWS Organizations. The policy should deny access to the API call Create Security GroupRule if the rule specifies 0.0.0.0/0 without a description referencing a security approval

B- Add an initial stage to CodePipeline called Security Check. This stage should call an AWS Lambda function that scans the CloudFormation template and fails the pipeline if it finds 0.0.0.0/0 in a security group without a description referencing a security approval

C- Create an AWS Config rule that is triggered on creation or edit of resource type EC2 SecurityGroup. This rule should call an AWS Lambda function to send a failure notification if the security group has any rules with a source CIDR of 0.0.0.0/0 without a description referencing a security approval.

D- Modify the IAM role used by CodePipeline. The IAM policy should deny access.

Answer:

B

Question 5

Question Type: **MultipleChoice**

A DevOps Engineer has been asked to recommend a tool to deploy the components of a three-tier web application. This application will use Amazon DynamoDB as a database

Which deployment requires the LEAST amount of operational management?

Options:

A- Use AWS CloudFormation to create a Classic Load Balancer and an Auto Scaling group. Use AWS OpsWorks to create the application and database resources Deploy application updates with OpsWorks using lifecycle events

B- Use AWS OpsWorks to create a Classic Load Balancer, an Auto Scaling group application, and database resources Deploy application updates using OpsWorks lifecycle events

C- Use AWS OpsWorks to create a Classic Load Balancer Auto Scaling and application resources Use AWS CloudFormation to create the database resources Deploy application updates using CloudFormation rolling updates

D- Use AWS CloudFormation to create a Classic Load Balancer an Auto Scaling group and database resources Deploy application updates using CloudFormation rolling updates

Answer:

B

Question 6

Question Type: MultipleChoice

A company indexes all of its Amazon CloudWatch Logs on Amazon ES and uses Kibana to view a dashboard for actionable insight. The company wants to restrict user access to Kibana by user

Which actions can a DevOps Engineer take to meet this requirement? (Select TWO.)

Options:

A- Create a proxy server with user authentication in an Auto Scaling group and restrict access of the Amazon ES endpoint to an Auto Scaling group tag

B- Create a proxy server with user authentication and an Elastic IP address and restrict access of the Amazon ES endpoint to the IP address

C- Create a proxy server with AWS IAM user and restrict access of the Amazon ES endpoint to the IAM user

D- Use AWS SSO to offer user name and password protection for Kibana

E- Use Amazon Cognito to offer user name and password protection for Kibana

Answer:

B, E

Question 7

Question Type: MultipleChoice

A DevOps Engineer has several legacy applications that all generate different log formats. The Engineer must standardize the formats before writing them to Amazon S3 for querying and analysis.

How can this requirement be met at the LOWEST cost?

Options:

A- Have the application send its logs to an Amazon EMR cluster and normalize the logs before sending them to Amazon S3

B- Have the application send its logs to Amazon QuickSight then use the Amazon QuickSight SPICE engine to normalize the logs Do the analysis directly from Amazon QuickSight.

C- Keep the logs in Amazon S3 and use Amazon Redshift Spectrum to normalize the logs in place

D- Use Amazon Kinesis Agent on each server to upload the logs and have Amazon Kinesis Data Firehose use an AWS Lambda function to normalize the logs before writing them to Amazon S3

Answer:

D

Question 8

Question Type: **MultipleChoice**

You have an application consisting of a stateless web server tier running on Amazon EC2 instances behind load balancer, and are using Amazon RDS with read replicas. Which of the following methods should you use to implement a self-healing and cost-effective architecture? Choose 2 answers from the optionsgiven below

Options:

A- Set up a third-party monitoring solution on a cluster of Amazon EC2 instances in order to emit custom Cloud Watch metrics to trigger the termination of unhealthy Amazon EC2 instances.

B- Set up scripts on each Amazon EC2 instance to frequently send ICMP pings to the load balancer in order to determine which instance is unhealthy and replace it.

C- Set up an Auto Scalinggroup for the web server tier along with an Auto Scaling policy that uses the Amazon RDS DB CPU utilization Cloud Watch metric to scale the instances.

D- Set up an Auto Scalinggroup for the web server tier along with an Auto Scaling policy that uses the Amazon EC2 CPU utilization CloudWatch metric to scale the instances.

E- Use a larger Amazon EC2 instance type for the web server tier and a larger DB instance type for the data storage layer to ensure that they don't become unhealthy.

F- Set up an Auto Scalinggroup for the database tier along with an Auto Scaling policy that uses the Amazon RDS read replica lag CloudWatch metric to scale out the Amazon RDS read replicas.

G- Use an Amazon RDS Multi-AZ deployment.

Answer:

D, G

Explanation:

The scaling of CC2 Instances in the Autoscaling group is normally done with the metric of the CPU

utilization of the current instances in the Autoscaling group

For more information on scaling in your Autoscaling Group, please refer to the below link:

* http://docs.aws.a mazon.com/autoscaling/latest/userguide/as-scaling-si mple-step.html

Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Cach AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure, Amazon RDS performs an automatic failover to the standby for to a read replica in the case of Amazon Aurora), so that you can resume database operations as soon as the failover is complete. For more information on RDS Multi-AZ please refer to the below link:

https://aws.amazon.com/rds/details/multi-az/

Option A is invalid because if you already have in-built metrics from Cloudwatch, why would you want to spend more in using a a third-party monitoring solution.

Option B is invalid because health checks are already a feature of AWS CLB

Option C is invalid because the database CPU usage should not be used to scale the web tier.

Option C is invalid because increasing the instance size does not always guarantee that the solution will not become unhealthy.

Option F is invalid because increasing Read-Replica's will not suffice for write operations if the primary DB fails.

Question 9

Question Type: MultipleChoice

Management has reported an increase in the monthly bill from Amazon Web Services, and they are extremely concerned with this increased cost. Management has asked you to determine the exact cause of this increase. After reviewing the billing report, you notice an increase in the data transfer cost. How can you provide management with a better insight into data transfer use?

Options:

A- Update your Amazon CloudWatch metrics to use five-second granularity, which will give better detailed metrics that can be combined with your billing data to pinpoint anomalies.

B- Use Amazon CloudWatch Logs to run a map-reduce on your logs to determine high usage and data transfer.

C- Deliver custom metrics to Amazon CloudWatch per application that breaks down application data transfer into multiple, more specific data points.
D- Using Amazon CloudWatch metrics, pull your Elastic Load Balancing outbound data transfer metrics monthly, and include them with your billing report to show which application is causing higher bandwidth usage.

Answer:

C

Explanation:

You can publish your own metrics to CloudWatch using the AWS CLI or an API. You can view statistical graphs of your published metrics with the AWS Management Console.

CloudWatch stores data about a metric as a series of data points. Each data point has an associated time stamp. You can even publish an aggregated set of data points called a statistic set.

If you have custom metrics specific to your application, you can give a breakdown to the management on the exact issue.

Option A won't be sufficient to provide better insights.

Option B is an overhead when you can make the application publish custom metrics

Option D is invalid because just the ELB metrics will not give the entire picture

For more information on custom metrics, please refer to the below document link: from AWS

http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publ ishingMetrics.htmI