

## [2017-Feb-NEWExam Pass 100% !Amazon AWS Certified DevOps Engineer - Professional VCE & PDF 190q Instant Download[Q11-Q20

2017 Feb NEW Amazon AWS Certified DevOps Engineer - Professional Exam Dumps (PDF & VCE) Updated Today!1.|NEW Amazon AWS Certified DevOps Engineer - Professional PDF and VCE Dumps 190Q&As Download:

<http://www.braindump2go.com/aws-devops-engineer-professional.html2>.|NEW Amazon AWS Certified DevOps Engineer - Professional Exam Questions and Answers Download:<https://1drv.ms/f/s!Avi7wzKf6QBjgh4jQ6bIe9aoaNIC> QUESTION 11Your application uses CloudFormation to orchestrate your application's resources. During your testing phase before the application went live, your Amazon RDS instance type was changed and caused the instance to be re-created, resulting In the loss of test data.How should you prevent this from occurring in the future? A. Within the AWS CloudFormation parameter with which users can select the Amazon RDS instance type, set AllowedValues to only contain the current instance type.B. Use an AWS CloudFormation stack policy to deny updates to the instance. Only allow UpdateStack permission to IAM principals that are denied SetStackPolicy. C. In the AWS CloudFormation template, set the AWS::RDS::DBInstance's DBInstanceClass property to be read-only.D. Subscribe to the AWS CloudFormation notification "BeforeResourceUpdate," and call CancelStackUpdate if the resource identified is the Amazon RDS instance.E. In the AWS CloudFormation template, set the DeletionPolicy of the AWS::RDS::DBInstance's DeletionPolicy property to "Retain." Answer: E QUESTION 12Your company develops a variety of web applications using many platforms and programming languages with different application dependencies. Each application must be developed and deployed quickly and be highly evadable to satisfy your business requirements.Which of the following methods should you use to deploy these applications rapidly? A. Develop the applications in Docker containers, and then deploy them to Elastic Beanstalk environments with Auto Scaling and Elastic Load Balancing.B. Use the AWS CloudFormation Docker import service to build and deploy the applications with high availability in multiple Availability Zones.C. Develop each application's code in DynamoDB, and then use hooks to deploy it to Elastic Beanstalk environments with Auto Scaling and Elastic Load Balancing.D. Store each application's code in a Git repository, develop custom package repository managers for each application's dependencies, and deploy to AWS OpsWorks in multiple Availability Zones. Answer: A QUESTION 13You have a large number of web servers in an Auto Scaling group behind a load balancer. On an hourly basis, you want to filter and process the logs to collect data on unique visitors, and then put that data in a durable data store in order to run reports. Web servers in the Auto Scaling group are constantly launching and terminating based on your scaling policies, but you do not want to lose any of the log data from these servers during a stop/termination initiated by a user or by Auto Scaling. What two approaches will meet these requirements?Choose 2 answers A. Install an Amazon Cloudwatch Logs Agent on every web server during the bootstrap process. Create a CloudWatch log group and define Metric Filters to create custom metrics that track unique visitors from the streaming web server logs. Create a scheduled task on an Amazon EC2 instance that runs every hour to generate a new report based on the Cloudwatch custom metrics.B. On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to Amazon Glacier. Ensure that the operating system shutdown procedure triggers a logs transmission when the Amazon EC2 instance is stopped/terminated. Use Amazon Data Pipeline to process the data in Amazon Glacier and run reports every hour.C. On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to an Amazon S3 bucket. Ensure that the operating system shutdown procedure triggers a logs transmission when the Amazon EC2 instance is stopped/terminated. Use AWS Data Pipeline to move log data from the Amazon S3 bucket to Amazon Redshift In order to process and run reports every hour.D. Install an AWS Data Pipeline Logs Agent on every web server during the bootstrap process. Create a log group object in AWS Data Pipeline, and define Metric Filters to move processed log data directly from the web servers to Amazon Redshift and run reports every hour. Answer: AC QUESTION 14You have been tasked with deploying a scalable distributed system using AWS OpsWorks. Your distributed system is required to scale on demand. As it is distributed, each node must hold a configuration file that includes the hostnames of the other instances within the layer. How should you configure AWS OpsWorks to manage scaling this application dynamically? A. Create a Chef Recipe to update this configuration file, configure your AWS OpsWorks stack to use custom cookbooks, and assign this recipe to the Configure LifeCycle Event of the specific layer.B. Update this configuration file by writing a script to poll the AWS OpsWorks service API for new instances.Configure your base AMI to execute this script on Operating System startup.C. Create a Chef Recipe to update this configuration file, configure your AWS OpsWorks stack to use custom cookbooks, and assign this recipe to execute when instances are launched.D. Configure your AWS OpsWorks layer to use the AWS-provided recipe for distributed host configuration, and configure the instance hostname and file path parameters in your recipes settings. Answer: A QUESTION 15 You have an application running on an Amazon EC2 instance and you are using IAM roles to securely access AWS Service APIs. How can you configure your application running on that instance to retrieve the API keys for use with the AWS SDKs? A. When

assigning an EC2 IAM role to your instance in the console, in the "Chosen SDK" drop-down list, select the SDK that you are using, and the instance will configure the correct SDK on launch with the API keys.B. Within your application code, make a GET request to the IAM Service API to retrieve credentials for your user.C. When using AWS SDKs and Amazon EC2 roles, you do not have to explicitly retrieve API keys, because the SDK handles retrieving them from the Amazon EC2 MetaData service.D. Within your application code, configure the AWS SDK to get the API keys from environment variables, because assigning an Amazon EC2 role stores keys in environment variables on launch. Answer: C

QUESTION 16When an Auto Scaling group is running in Amazon Elastic Compute Cloud (EC2), your application rapidly scales up and down in response to load within a 10-minute window; however, after the load peaks, you begin to see problems in your configuration management system where previously terminated Amazon EC2 resources are still showing as active.What would be a reliable and efficient way to handle the cleanup of Amazon EC2 resources within your configuration management system?Choose 2 answers A. Write a script that is run by a daily cron job on an Amazon EC2 instance and that executes API Describe calls of the EC2 Auto Scaling group and removes terminated instances from the configuration management system.B. Configure an Amazon Simple Queue Service (SQS) queue for Auto Scaling actions that has a script that listens for new messages and removes terminated instances from the configuration management system.C. Use your existing configuration management system to control the launching and bootstrapping of instances to reduce the number of moving parts in the automation.D. Write a small script that is run during Amazon EC2 instance shutdown to de-register the resource from the configuration management system.E. Use Amazon Simple Workflow Service (SWF) to maintain an Amazon DynamoDB database that contains a whitelist of instances that have been previously launched, and allow the Amazon SWF worker to remove information from the configuration management system. Answer: AD

QUESTION 17You have enabled Elastic Load Balancing HTTP health checking. After looking at the AWS Management Console, you see that all instances are passing health checks, but your customers are reporting that your site is not responding.What is the cause? A. The HTTP health checking system is misreporting due to latency in inter-instance metadata synchronization.B. The health check in place is not sufficiently evaluating the application function.C. The application is returning a positive health check too quickly for the AWS Management Console to respond.D. Latency in DNS resolution is interfering with Amazon EC2 metadata retrieval. Answer: B

QUESTION 18You use Amazon CloudWatch as your primary monitoring system for your web application. After a recent software deployment, your users are getting Intermittent 500 Internal Server Errors when using the web application. You want to create a CloudWatch alarm, and notify an on-call engineer when these occur.How can you accomplish this using AWS services? Choose 3 answers A. Deploy your web application as an AWS Elastic Beanstalk application. Use the default Elastic Beanstalk Cloudwatch metrics to capture 500 Internal Server Errors. Set a CloudWatch alarm on that metric.B. Install a CloudWatch Logs Agent on your servers to stream web application logs to CloudWatch.C. Use Amazon Simple Email Service to notify an on-call engineer when a CloudWatch alarm is triggered.D. Create a CloudWatch Logs group and define metric filters that capture 500 Internal Server Errors. Set a CloudWatch alarm on that metric.E. Use Amazon Simple Notification Service to notify an on-call engineer when a CloudWatch alarm is triggered.F. Use AWS Data Pipeline to stream web application logs from your servers to CloudWatch. Answer: BDE

QUESTION 19After a daily scrum with your development teams, you've agreed that using Blue/Green style deployments would benefit the team. Which technique should you use to deliver this new requirement? A. Re-deploy your application on AWS Elastic Beanstalk, and take advantage of Elastic Beanstalk deployment types.B. Using an AWS CloudFormation template, re-deploy your application behind a load balancer, launch a new AWS CloudFormation stack during each deployment, update your load balancer to send half your traffic to the new stack while you test, after verification update the load balancer to send 100% of traffic to the new stack, and then terminate the old stack.C. Re-deploy your application behind a load balancer that uses Auto Scaling groups, create a new identical Auto Scaling group, and associate it to the load balancer. During deployment, set the desired number of instances on the old Auto Scaling group to zero, and when all instances have terminated, delete the old Auto Scaling group.D. Using an AWS OpsWorks stack, re-deploy your application behind an Elastic Load Balancing load balancer and take advantage of OpsWorks stack versioning, during deployment create a new version of your application, tell OpsWorks to launch the new version behind your load balancer, and when the new version is launched, terminate the old OpsWorks stack. Answer: C

QUESTION 20You have a complex system that involves networking, IAM policies, and multiple, three-tier applications. You are still receiving requirements for the new system, so you don't yet know how many AWS components will be present in the final design.You want to start using AWS CloudFormation to define these AWS resources so that you can automate and version-control your infrastructure.How would you use AWS CloudFormation to provide agile new environments for your customers in a cost-effective, reliable manner? A. Manually create one template to encompass all the resources that you need for the system, so you only have a single template to version-control.B. Create multiple separate templates for each logical part of the system, create nested stacks in AWS CloudFormation, and maintain several templates to version-control.C. Create multiple separate templates for each logical part of

the system, and provide the outputs from one to the next using an Amazon Elastic Compute Cloud (EC2) instance running the SDK for finer granularity of control.D. Manually construct the networking layer using Amazon Virtual Private Cloud (VPC) because this does not change often, and then use AWS CloudFormation to define all other ephemeral resources. Answer: B

!!!RECOMMEND!!! 1.|NEW Amazon AWS Certified DevOps Engineer - Professional PDF and VCE Dumps 190Q&As Download: <http://www.braindump2go.com/aws-devops-engineer-professional.html> 2.|NEW Amazon AWS Certified DevOps Engineer - Professional Study Guide Video: YouTube Video: [YouTube.com/watch?v=nRhDUfpWQgE](https://www.youtube.com/watch?v=nRhDUfpWQgE)