[2017 PDF&VCE Pass AWS-DevOps-Engineer-Professional Exam By Exercising Lead2pass Latest AWS-DevOps-Engineer-Professional VCE And PDF Dumps (61-80)

Lead2pass 2017 September New Amazon AWS-DevOps-Engineer-Professional Exam Dumps! 100% Free Download! 100% Pass Guaranteed! Good news, Lead2pass has updated the AWS-DevOps-Engineer-Professional exam dumps. With all the questions and answers in your hands, you will pass the Amazon AWS-DevOps-Engineer-Professional exam easily. Following questions and answers are all new published by **Amazon** Official Exam Center:

https://www.lead2pass.com/aws-devops-engineer-professional.html QUESTION 61You have been asked to use your departments existing continuous Integration (CI) tool to test a three-tier web architecture defined In an AWS CloudFormation template. The tool already supports AWS APIs and can launch new AWS CloudFormation stacks after polling version control. The CI tool reports on the success of the AWS CloudFormation stack creation by using the Describe Stacks API to look for the CREATE COMPLETE status.The architecture tiers defined in the template consist of: - One load balancer- Five Amazon EC2 instances running the web application- One multi-AZ Amazon ROS instance How would you implement this? Choose 2 answers A.   Define a WaitCondition and a WaitConditionHandle for the output of a UserData command that does sanity checking of the application's post-install state.B.   Define a CustomResource and write a script that runs architecture-level Integration tests through the load balancer to the application and database for the state of multiple tiers.C.   Define a WaitCondition and use a WaitConditionHandle that leverages the AWS SDK to run the DescribeStacks API call until the CREATE COMPLETE status is returned.D.   Define a CustomResource that leverages the AWS SDK to run the DescribeStacks API call until the 'CREATE COMPLETE status is returned.E.   Define a UserDataHandle for the output of a UserData command that does sanity checking of the application's post-install state and runs integration tests on the state of multiple tiers through the load balancer to the application.F.   Define a UserDataHandle for the output of a CustomResource that does sanity checking of the application's post-install state.Answer: CE QUESTION 62You are building a large, multi-tenant SaaS (software-as-a-service) application with a component that fetches data to process from a customer-specific Amazon S3 bucket in their account. How should you ensure that your application follows security best practices and limits risk when fetching data from customer-owned Amazon S3 buckets? A.   Have users create an IAM user with a policy that grants read-only access to the Amazon S3 bucket required by your application, and store the corresponding access keys in an encrypted database that holds their account data.B.   Have users create a cross-account lAM role with a policy that grants read-only access to the Amazon S3 bucket required by your application to the AWS account ID running your production Sass application.C.   Have users create an Amazon S3 bucket policy that grants read-only access to the Amazon S3 bucket required by your application, and securely store the corresponding access keys in the database holding their account data.D.   Have users create an Amazon S3 bucket policy that grants read-only access to the Amazon S3 bucket required by your application and limits access to the public IP address of the SaaS application. Answer: B QUESTION 63You have a fleet of Elastic Compute Cloud (EC2) instances in an Auto Scaling group. All of these instances are running Microsoft Windows Server 2012 backed by Amazon Elastic Block Store (EBS). These instances were launched through AWS CloudFormation. You have determined that your instances are underutilized, and in order to save some money, have decided to modify the instance type of the fleet. In which of the following ways can you achieve the desired result during a scheduled maintenance window? Choose 2 answers A.   Create a new Auto Scaling launch configuration specifying the new instance type, associate it to the existing Auto Scaling group, and terminate the running instances.B.   Identify the new instance type in the user data and restart the running instances one at a time.C.   Use the AWS Command Line Interface (CLI) to modify the instance type of each running instance.D.   Change the instance type in the AWS CloudFormation template that was used to create the Amazon EC2 instances, and then update the stack.E.   Take snapshots of the running instances, and launch new instances based on those snapshots. Answer: AD QUESTION 64You run a large number of applications on Amazon EC2 instances. Each application has associated metadata, such as cost center, support contact, and application ID. Many applications usually co-exist on each Amazon EC2 instance, so the amount of metadata per instance can range from 10 to 200 items. The customer wants to be able to quickly access this metadata using an API without logging into the instances. Which of the following options will satisfy their requirements? Choose 2 answers A.   Create individual Amazon EC2 tags for each metadata item, and associate them with the Amazon EC2 instances. Access the metadata by using the ec2-describe-instance API call.B.   Create compound Amazon EC2 tags for the metadata items, where multiple items are joined together in individual tags, and associate them with the Amazon EC2 instances. Access the metadata by using the ec2-describe-tags API call.C.   Create a DynamoDB table to hold the metadata, and associate it with the Amazon EC2 instance IDs running the applications. Access the metadata by querying the database via the DynamoDB API.D.   As part of the Amazon EC2 Instance bootstrapping process, add the metadata to the Amazon

EC2 user data. Access the metadata by using the ec2-describe-instance API call.E.   As part of the Amazon EC2 instance bootstrapping process, add the metadata to the Amazon EC2 user data. Access the metadata by accessing its loopback address from a management instance in the same VPC. Answer: BC QUESTION 65You have an application running on multiple Amazon EC2 instances within an Auto Scaling group. You notice that instances are being re-spawned as their health checks are failing in Amazon EC2. However, before you have a chance to diagnose the issue, the affected instances are being terminated by the Auto Scaling service. You receive notifications of health checks failing and investigate within 20 minutes. However, this is not enough time to troubleshoot the issue.What should you change that will enable you to troubleshoot the instance before it is terminated by the Auto Scaling service, while keeping costs minimal? A.   Install the Amazon CloudWatch Logs Agent on the instance and configure application and system logs to be sent to the CloudWatch Logs service.B.   Configure an Amazon SNS topic and associate it with your Auto Scaling group's CloudWatch alarms.Configure an Amazon SQS queue as a subscriber of this topic, and then create a fleet of Amazon EC2 workers that poll this queue and instruct the Amazon EC2 Auto Scaling API to remove the instance from the Auto Scaling group when an alarm is triggered.C.   Create an Auto Scaling Group lifecycle hook to hold the instance in a terminating:wait state until you have completed any troubleshooting. When you have completed troubleshooting, wait for the terminating state to expire, or notify to Scaling to complete the lifecycle hook and terminate the Instance.D.   Change the "DeleteOnTermination" flag to false in the Auto Scaling group configuration to ensure that instances are not deleted in the future. Answer: C QUESTION 66You set up a scalable continuous integration platform on AWS. The platform consists of a master node that can delegate project build jobs to multiple slave nodes, all running on Amazon EC2. The build output will be stored in Amazon S3. You always have five slave nodes deployed. Each slave node can handle 10 build jobs simultaneously. Your master node publishes a custom Amazon CloudWatch metric with the name "RunningBuildiobs" that Slows you to programmatically track how many build jobs are running across your platform.Which two configuration options will allow you to flexibly scale your platform to support more than 50 simultaneous build jobs while minimizing costs? Choose 2 answers A.   Place your fleet of slave nodes in an Auto Scaling group. Configure a CloudWatch alarm that triggers an Auto Scaling policy to launch Amazon EC2 Instances when "RunningBuildJobs" is greater than 45 for more than five minutes.B.   Configure a CloudWatch alarm that sends an alert when "RunningBuildJobs" is greater than 45 for more than five minutes. Use Amazon Simple Queue Service to process additional build jobs when the CloudWatch alarm is triggered.C.   Configure your fleet of slave nodes to fully utilize all of your purchased Amazon EC2 Heavy Utilization Reserved Instances. Configure a CloudWatch alarm that launches new Amazon EC2 instances when "RunningBuildJobs" is less than 40 for more than five minutes.D.   Run your fleet of slave nodes in an Auto Scaling group. Configure a Cloudwatch alarm that launches new Amazon EC2 Dedicated Instances when "RunningBuildJobs" is less than 40 for more than five minutes.E.   Place your fleet of slave nodes in an Auto Scaling group. Configure a CloudWatch alarm that triggers an Auto Scaling policy to terminate Amazon EC2 instances when "RunningBuildJobs" is less than 40 for more than five minutes. Answer: AE QUESTION 67You have just come from your Chief Information Security Officer's (CISO) office with the instructions to provide an audit report of all AWS network rules used by the organization's Amazon EC2 instances. You have discovered that a single Describe-Security-Groups API call will return all of an account's security groups and rules within a region. You create the following pseudo-code to create the required report: - Parse "aws ec2 describe-security-groups" output- For each security group- Create report of ingress and egress rules Which two additional pieces of logic should you include to meet the CISO's requirements? Choose 2 answers A.   Parse security groups in each region.B.   Parse security groups in each Availability Zone and region.C.   Evaluate VPC network access control lists.D.   Evaluate AWS CloudTrail logs.E.   Evaluate Elastic Load Balancing access control lists.F.   Parse CloudFront access control lists. Answer: AC QUESTION 68You are responsible for a large-scale video transcoding system that operates with an Auto Scaling group of video transcoding workers. The Auto Scaling group is configured with a minimum of 750 Amazon EC2 instances and a maximum of 1000 Amazon EC2 instances. You are using Amazon SQS to pass a message containing the URI for a video stored in Amazon S3 to the transcoding workers. An Amazon CloudWatch alarm has notified you that the queue depth is becoming very large. How can you resolve the alarm without the risk of increasing the time to transcode videos? Choose 2 answers. A.   Create a second queue in Amazon SQS.B.   Adjust the Amazon CloudWatch alarms for a higher queue depth.C.   Create a new Auto Scaling group with a launch configuration that has a larger Amazon EC2 instance typeD.   Add an additional Availability Zone to the Auto Scaling group configuration.E.   Change the Amazon CloudWatch alarm so that it monitors the CPU utilization of the Amazon EC2 instances rather than the Amazon SQS queue depth.F.   Adjust the Auto Scaling group configuration to increase the maximum number of Amazon EC2 instances. Answer: CF QUESTION 69You have been tasked with deploying a solution for your company that will store images, which the marketing department will use for its campaigns. Employees are able to upload images via a web interface, and once uploaded, each image must be resized and watermarked with the company logo. Image resize and watermark is not time-sensitive and can be completed days after upload if required. How should

you design this solution in the most highly available and cost-effective way? A.   Configure your web application to upload images to the Amazon Elastic Transcoder service. Use the Amazon Elastic Transcoder watermark feature to add the company logo as a watermark on your images and then to upload the final images into an Amazon S3 bucket.B.   Configure your web application to upload images to Amazon S3, and send the Amazon S3 bucket URI to an Amazon SQS queue. Create an Auto Scaling group and configure it to use Spot instances, specifying a price you are willing to pay. Configure the instances in this Auto Scaling group to poll the SQS queue for new images and then resize and watermark the image before uploading the final images into Amazon S3.C. Configure your web application to upload images to Amazon S3, and send the S3 object URI to an Amazon SQS queue. Create an Auto Scaling launch configuration that uses Spot instances, specifying a price you are willing to pay. Configure the instances in this Auto Scaling group to poll the Amazon SQS queue for new images and then resize and watermark the image before uploading the new images into Amazon S3 and deleting the message from the Amazon SQS queue.D.   Configure your web application to upload images to the local storage of the web server. Create a cronjob to execute a script daily that scans this directory for new files and then uses the Amazon EC2 Service API to launch 10 new Amazon EC2 instances, which will resize and watermark the images daily. Answer: C QUESTION 70You run a small online consignment marketplace. Interested sellers complete an online application in order to allow them to sell their products on your website. Once approved, they can post their product using a custom interface. From that pant, you manage the shopping cart process so that when a buyer decides to buy a product, you handle the billing and coordinate the shipping. Part of this process requires sending emails to the buyer and the seller at different stages.Your system has been running on AWS for a few months. Occasionally, products are shipped before payment cleared and emails are sent out of order. Furthermore, sometimes credit cards are being charged twice. How can you resolve these problems? A.   Use the Amazon Simple Queue Service (SQS), and use a different set of workers for each task.B.   Use the Amazon Simple Workflow Service (SWF), and use a different set of workers for each task.C.   Use the Simple Email Service (SES) to control the correct order of email delivery.D.   Use the AWS Data Pipeline service to control the process flow of the various tasks.E.   Use the Amazon Simple Queue Service (SQS), and use a single set of workers for each task. Answer: B QUESTION 71Your application has an Auto Scaling group of m3.large instances running an application that receives messages born an Amazon SQS queue. After a while, the number of instances reaches the maximum set for the group and the number of messages on the queue continues to increase. You have discovered that a third- party library used by the application has a bug that causes a memory leak. What cost-effective steps can you take to continue message processing while the library developer fixes the bug? A.   Enable Elastic Load Balancing health checks for the Auto Scaling group. When Elastic Load Balancing has detected a failure, Auto Scaling will terminate the failing application's instance and launch a new one.B.   Use Amazon EC2 instance memory usage CloudWatch metrics to raise alerts when they reach a defined level and send a message to Auto Scaling to fail the instance health check.C.   Use application monitoring on the instance to restart the application when memory usage reaches a defined level.D.   Create a new Auto Scaling launch configuration to use the r3.large instance type. Update the Auto Scaling group with the new launch configuration. Answer: D QUESTION 72You are in charge of a large-scale highly available multi-tier web application infrastructure. This architecture consists of Amazon Route53 with a load balancer and multiple Amazon EC2 instances. You have been tasked to come up with a process to provide Blue/Green style deployments. Which technique should you use to deliver this new requirement? A.   Using Elastic Beanstalk re-deploy your application and configure Elastic Beanstalk Deployment types, and then use Amazon Route53's alias resource record set to swap between Elastic Beanstalk deployment types.B.   Re-deploy your application behind a load balancer using an AWS CloudFormation template, launch a new AWS CloudFormation stack during each deployment, update your Amazon Route53 alias resource record set to point to the new load balancer, and finally, terminate your old AWS CloudFormation stack.C.   Re-deploy your application behind a load balancer using Auto Scaling groups, create a new identical Auto Scaling group, and associate it to the load balancer. During deployment, create a new Amazon Route53 hosted zone, add this new load balancer to the zone in an alias resource record set, and then remove your old Auto Scaling group.D.   Re-deploy your application behind a load balancer using an OpsWorks stack, and use AWS 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 launches, update your Amazon Route53 alias resource retort to point to the new load balancer. Answer: B QUESTION 73Your application uses Amazon SQS and Auto Scaling to process background jobs. The Auto Scaling policy is based on the number of messages in the queue, with a maximum Instance count of 100. Since the application was launched, the group has never scaled above 50. The Auto Scaling group has now scaled to 100, the queue size is increasing, and very few Jobs are being completed. The number of messages being sent to the queue is at normal levels.What should you do to identify why the queue size is unusually high, and to reduce it? A.   Temporarily increase the Auto Scaling group's desired value to 200. When the queue size has been reduced, reduce it to 50.B.   Analyze the application logs to identify possible reasons for message processing failure and resolve the cause for failures.C.   Create additional Auto Scaling groups, enabling the

processing of the queue to be performed in parallel.D.   Analyze CloudTrail logs for Amazon SQS to ensure that the instances' Amazon EC2 role has permission to receive messages from the queue. Answer: B QUESTION 74You have a web application that is currently running on a collection of micro instance types in a single AZ behind a single load balancer. You have an Auto Scaling group configured to scale from 2 to 64 instances. When reviewing your CloudWatch metrics, you see that sometimes your Auto Scaling group is running 64 micro instances. The web application is reading and writing to a DynamoDB-configured backend and configured with 800 Write Capacity Units and 800 Read Capacity Units. Your customers are complaining that they are experiencing long load times when viewing your website. You have investigated the DynamoDB CloudWatch metrics; you are under the provisioned Read and write Capacity Units and there is no throttling. How do you scale your service to improve the load times and ensure the principles of high availability? A.   Change your Auto Scaling group configuration to include multiple AZs.B.   Change your Auto Scaling group configuration to include multiple AZs, and increase the number of Read Capacity Units in your DynamoDB table by a factor of three, because you will need to be calling DynarnoDB from three AZs.C.   Add a second load balancer to your Auto Scaling group so that you can support more inbound connections per second.D.   Change your Auto Scaling group configuration to use larger instances and include multiple AZ's instead of one. Answer: D QUESTION 75Your social media marketing application has a component written in Ruby running on AWS Elastic Beanstalk. This application component posts messages to social media sites in support of various marketing campaigns. Your management now requires you to record replies to these social media messages to analyze the effectiveness of the marketing campaign in comparison to past and future efforts. You' ve already developed a new application component to interface with the social media site APIs in order to read the replies. Which process should you use to record the social media replies in a durable data store that can be accessed at any time for analysis of historical data? A.   Deploy the new application component in an Auto Scaling group of Amazon Elastic Compute Cloud (EC2) Instances, read the data from the social media sites, store it with Amazon Elastic Block Store, and use AWS Data Pipeline to publish it to Amazon Kinesis for analytics.B.   Deploy the new application component as an Elastic Beanstalk application, read the data from the social media sites, store it in Amazon DynamoDB, and use Apache Hive with Amazon Elastic MapReduce for analytics.C.   Deploy the new application component in an Auto Scaling group of Amazon EC2 instances, read the data from the social media sites, store it in Amazon Glacier, and use AWS Data Pipeline to publish it to Amazon Redshift for analytics.D.   Deploy the new application component as an Amazon Elastic Beanstalk application, read the data from the social media site, store it with Amazon Elastic Block Store, and use Amazon Kinesis to stream the data to Amazon CloudWatch for analytics. Answer: B QUESTION 76A web application is being actively developed by multiple development teams within your organization. You have created a self-service portal-driven by AWS CloudFormation and the AWS APIs-that allows testers to select a code branch containing a new feature that they want to test. The portal will then provision an environment and deploy the right branch of code to it.Recently you have noticed that a large number of environments contain broken builds. You want to introduce a set of automated browser tests that are executed on a new environment before the environment is available to the tester. This way a tester does not waste time trying to test new features in a broken environment. Select a suitable way to implement such a feature into the existing self-service portal: A.   Specify your automated tests in the "tests" section of the AWS CloudFormation template. AWS CloudFormation will then execute the tests on your behalf as part of the environment build.B.   Configure a centralized test server that hosts an automated browser testing framework. Use an AWS CloudFormation custom resource to notify the centralized test server, via an Amazon SNS topic, that a new environment has been initialized. The centralized test server can then execute the tests before sending the results back to the AWS CloudFormation service.C.   Pass the test scripts to the cfn-init service via the "tests" section of the AWS::CloudFormation::Init metadata. Cfn-init will then execute these tests and return the result to the AWS CloudFormation service.D.   Configure a centralized test server that hosts an automated browser testing framework. Include an Amazon SES email resource under the outputs section of your AWS CloudFormation template. This we send an email to your centralized test server, informing it that the environment is ready for tests. Answer: B QUESTION 77You have written a server-side Node.Js application and a web application with an HTML/JavaScript front end that uses the Angular.js framework. The server-side application connects to an Amazon Redshift cluster, issues queries, and then returns the results to the front end for display. Your user base is very large and distributed, but it is important to keep the cost of running this application low. Which deployment strategy is both technically valid and the most cost-effective? A.   Deploy an AWS Elastic Beanstalk application with two environments: one for the Node.js application and another for the web front end. Launch an Amazon Redshift cluster, and point your application to its Java Database Connectivity (JDBC) endpoint.B.   Deploy an AWS OpsWorks stack with three layers: a static web server layer for your front end, a Node.js app server layer for your server-side application, and a Redshift DB layer for your Amazon Redshift duster.C.   Upload the HTML, CSS, images, and JavaScript for the front end to an Amazon Simple Storage Service (S3) bucket. Create an Amazon CloudFront distribution with this bucket as its origin. Use AWS Elastic Beanstalk to deploy the Node.js application. Launch an

Amazon Redshift cluster, and point your application to its JDBC endpoint.D.   Upload the HTML, CSS, images, and JavaScript for the front end, plus the Node.js code for the server-side application, to an Amazon S3 bucket. Create a CloudFront distribution with this bucket as its origin. Launch an Amazon Redshift cluster, and point your application to its JDBC endpoint.E.   Upload the HTML, CSS, images, and JavaScript for the front end to an Amazon S3 bucket. Use AWS Elastic Beanstalk to deploy the Node.js application. Launch an Amazon Redshift cluster, and point your application to its JDBC endpoint. Answer: C QUESTION 78You are building an AWS CloudFormation template for a multi-tier web application. The user data of your Linux web server resource contains a complex script that can take a long time to run. Which techniques could you use to ensure that these servers are fully configured and running before attaching them to the load balancer? Choose 2 answers A.   Launch your Linux servers from a nested stack that is called from within the load balancer resource in your AWS CloudFormation template.B.   Add an AWS CloudFormation Wait Condition that depends on the web server resource. When the UserData script finishes on the web servers, use curl to send a signal the Wait Condition at http://169.254.169.254/waithandle/.C.   Add an AWS CloudFormation wait Condition that depends on the web server resource. When the UserData script finishes on the web servers, use curl to signal to the Wait Condition pre-signed URL that they are ready.D.   In your AWS CloudFormation template, position the load balancer resource JSON block directly below your Linux server resource.E.   Add an AWS CloudFormation Wait Condition that depends on the web server resource. When the UserData script finishes on the web servers, use the command "cfn-signal" to signal that they are ready. Answer: CE QUESTION 79Customers have recently been complaining that your web application has randomly stopped responding. During a deep dive of your logs, the team has discovered a major bug in your new Java web application. This bug is causing a memory leak that eventually causes the application to crash. Your web application runs on Amazon EC2 and was built with AWS CloudFormation. Which techniques should you use to help detect these problems faster, as well as help eliminate the server's unresponsiveness? Choose 2 answers A.   Update your AWS CloudFormation configuration and enable a CustomResource that uses cfn-signal to detect memory leaks.B.   Update your CloudWatch metric granularity config for all Amazon EC2 memory metrics to support five- second granularity. Create a CloudWatch alarm that triggers an Amazon SNS notification to page your team when the application memory becomes too large.C.   Update your AWS CloudFormation configuration to take advantage of Auto Scaling groups. Configure an Auto Scaling group policy to trigger off your custom CloudWatch metrics.D.   Create a custom CloudWatch metric that you push your JVM memory usage to. Create a Cloudwatch alarm that triggers an Amazon SNS notification to page your team when the application memory usage becomes too large.E.   Update your AWS CloudFormation configuration to take advantage of CloudWatch metrics Agent.Configure the CloudWatch Metrics Agent to monitor memory usage and trigger an Amazon SNS alarm. Answer: CD QUESTION 80You have an ASP.NET web application running in Amazon Elastic Beanstalk. Your next version of the application requires a third-party Windows Installer package to be installed on the instance on first boot and before the application launches.Which options are possible? Choose 2 answers A.   In the application's Global.asax file, run msiexec.exe to install the package using Process.Start() in the Application Start event handler.B.   In the source bundle's .ebextensions folder, create a file with a .config extension. In the file, under the "packages" section and "msi" package manager, include the package's URL.C.   Launch a new Amazon EC2 instance from the AMI used by the environment. Log into the instance, install the package and run sysprep. Create a new AMI. Configure the environment to use the new AMI.D.   In the environment's configuration, edit the instances configuration and add the package's URL to the "Packages" section.E.   In the source bundle's .ebextensions folder, create a "Packages" folder. Place the package in the folder. Answer: BD More free Lead2pass **AWS-DevOps-Engineer-Professional** exam new questions on Google Drive: https://drive.google.com/open?id=0B3Syig5i8gpDbVZ1cTB3QnNPQlk Once there are some changes on AWS-DevOps-Engineer-Professional exam questions, we will update the study materials timely to make sure that our customer can download the latest edition. 2017 **Amazon AWS-DevOps-Engineer-Professional** (All 190 Q&As) exam dumps (PDF&VCE) from Lead2pass: https://www.lead2pass.com/aws-devops-engineer-professional.html [100% Exam Pass Guaranteed]