

[2017 PDF&VCE Lead2pass Latest Amazon AWS-DevOps-Engineer-Professional Exam Questions Free Download (41-60)]

Lead2pass 2017 September New Amazon [AWS-DevOps-Engineer-Professional Exam Dumps! 100% Free Download! 100% Pass Guaranteed!](#) Pass AWS-DevOps-Engineer-Professional exam with the latest Lead2pass AWS-DevOps-Engineer-Professional dumps: Lead2pass AWS-DevOps-Engineer-Professional exam questions and answers in PDF are prepared by our experts. Moreover, they are based on the recommended syllabus that covering all the AWS-DevOps-Engineer-Professional exam objectives. Following questions and answers are all new published by **Amazon** Official Exam Center:

<https://www.lead2pass.com/aws-devops-engineer-professional.html> QUESTION 41 Your company currently runs a large multi-tier web application. One component is an API service that all other components of your application rely on to perform read/write operations. This service must have high availability and zero downtime during deployments. Which technique should you use to provide cost-effective, zero-downtime deployments for this component? A. Use an AWS CloudFormation template to re-deploy your application behind a load balancer, and launch a new AWS CloudFormation stack during each deployment. Update your load balancer to send traffic to the new stack, and then deploy your software. Leave your old stacks running, and tag their resources with the version for rollback. B. Re-deploy your application on Elastic Beanstalk. During deployment, create a new version of your application, and create a new environment running that version in Elastic Beanstalk. Finally, take advantage of the Elastic Beanstalk Swap CNAME operation to switch to the new environment. 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 your Amazon Route53 zone. Configure Amazon Route53 to auto-weight traffic over to the new Auto Scaling group when all instances are marked as healthy. D. Re-deploy your application behind a load balancer using an AWS OpsWorks stack and use AWS OpsWorks stack versioning, during deployment create a new version of your application, tell AWS OpsWorks to launch the new version behind your load balancer, and when the new version is launched, terminate the old AWS OpsWorks stack. Answer: B QUESTION 42 You want to build a new search tool feature for your monitoring system that will allow your information security team to quickly audit all API calls in your AWS accounts. What combination of AWS services can you use to develop and automate the backend processes supporting this tool? Choose 3 answers. A. Create an Amazon CloudSearch domain for API call logs. Configure the search domain so that it can be used to index API call logs for the search tool. B. Use AWS CloudTrail to store API call logs in an Amazon S3 bucket. Configure an Amazon Simple Notification Service topic called "log-notification" that notifies subscribers when new logs are available. Subscribe an Amazon SQS queue to the topic. C. Use Amazon Cloudwatch to ship AWS CloudTrail logs to your monitoring system. D. Create an AWS Elastic Beanstalk application in worker role mode that uses an Amazon Simple Email Service (SES) domain to facilitate batch processing new API call log files retrieved from an Amazon S3 bucket into a search index. E. Use AWS CloudTrail to store API call logs in an Amazon S3 bucket. Configure Amazon Simple Email Service (SES) to notify subscribers when new logs are available. Subscribe an Amazon SQS queue to the email domain. F. Create Amazon Cloudwatch custom metrics for the API call logs. Configure a Cloudwatch search domain so that it can be used to index API call logs for the search tool. G. Create an AWS Elastic Beanstalk application in worker role mode that uses an Amazon SQS queue to facilitate batch processing new API call log files retrieved from an Amazon S3 bucket into a search index. Answer: ABG

QUESTION 43 You are using AWS Elastic Beanstalk to deploy your application and must make data stored on an Amazon Elastic Block Store (EBS) volume snapshot available to the Amazon Elastic Compute Cloud (EC2) instances. How can you modify your Elastic Beanstalk environment so that the data is added to the Amazon EC2 instances every time you deploy your application? A. Add commands to a configuration file in the .ebextensions folder of your deployable archive that mount an additional Amazon EBS volume on launch. Also add a "BlockDeviceMappings" option, and specify the snapshot to use for the block device in the Auto Scaling launch configuration. B. Add commands to a configuration file in the .ebextensions folder of your deployable archive that uses the create-volume Amazon EC2 API or CLI to create a new ephemeral volume based on the specified snapshot and then mounts the volume on launch. C. Add commands to the Amazon EC2 user data that will be executed by eb-init, which uses the create-volume Amazon EC2 API or CLI to create a new Amazon EBS volume based on the specified snapshot, and then mounts the volume on launch. D. Add commands to the Chef recipe associated with your environment, use the create-volume Amazon EC2 API or CLI to create a new Amazon EBS volume based on the specified snapshot, and then mount the volume on launch. Answer: A

QUESTION 44 You would like to run automated, continuous application level integration tests on all members of an Auto Scaling group. Which two options should you use? Choose 2 answers. A. Use the AWS SDK to run the DescribeInstances API call on the Auto Scaling group, and then iterate over the members and remotely connect to each Amazon EC2 instance and run the integration tests. B. Use the AWS SDK to run the DescribeAutoScalingInstances API call on the Auto Scaling Group, iterate over the members

using the Describe Instances API call, remotely connect to each Amazon EC2 instance, and then run the integration tests.C. Set up a custom CloudWatch metric with the output of your integration tests that are run by a scheduled process on each instance, and then set up a CloudWatch alert for any failures.D. Using an Auto Cycle Group lifecycle policy, define a scheduled task to run integration tests when a new Amazon EC2 instance enters the InService state.E. Set up a custom CloudWatch metric that uses the output of the DescribeAutoScalingInstances API call to determine the HealthCheck status of the Amazon EC2 instances.F. Using the Auto Cycle Group lifecycle policy, define a scheduled task to run integration tests on individual instances using the Amazon EC2 user data to export test data to CloudWatch Logs. Answer: BC

QUESTION 45Your application Amazon Elastic Compute Cloud (EC2) instances bootstrap by using a master configuration file that is kept in a version-enabled Amazon Simple Storage Service (S3) bucket. Which one of the following methods should you use to securely install the current configuration version onto the instances in a cost-effective way? A. Create an Amazon DynamoDB table to store the different versions of the configuration file. Associate AWS Identity and Access Management (IAM) EC2 roles to the Amazon EC2 instances, and reference the DynamoDB table to get the latest file from Amazon Simple Storage Service (S3).B. Associate an IAM S3 role to the bucket, list the object versions using the Amazon S3 API, and then get the latest object.C. Associate an IAM EC2 role to the instances, list the object versions using the Amazon S3 API, and then get the latest object.D. Associate an IAM EC2 role to the instances, and then simply get the object from Amazon S3, because the default is the current version.E. Store the IAM credentials in the Amazon EC2 user data for each instance, and then simply get the object from S3, because the default is the current version. Answer: D

QUESTION 46Your company operates a website for promoters to sell tickets for entertainment events. You are using a load balancer in front of an Auto Scaling group of web servers. Promotion of popular events can cause surges of website visitors. During scaling-out at these times, newly launched instances are unable to complete configuration quickly enough, leading to user disappointment. What options should you choose to improve scaling yet minimize costs? Choose 2 answers. A. Create an AMI with the application pre-configured. Create a new Auto Scaling launch configuration using this new AMI, and configure the Auto Scaling group to launch with this AMI.B. Use Auto Scaling pre-warming to launch instances before they are required. Configure pre-warming to use the CPU trend CloudWatch metric for the group.C. Publish a custom CloudWatch memo from your application on the number of tickets sold, and create an Auto Scaling policy based on this.D. Use the history of past scaling events for similar event sales to predict future scaling requirements. Use the Auto Scaling scheduled scaling feature to vary the size of the fleet.E. Configure an Amazon S3 bucket for website hosting. Upload into the bucket an HTML holding page with its 'x-amz-website-redirect-location' metadata property set to the load balancer endpoint. Configure Elastic Load Balancing to redirect to the holding page when the load on web servers is above a certain level. Answer: AD

QUESTION 47You're responsible for a popular file sharing application that uses Elastic Load Balancing to distribute traffic to an Amazon EC2 application tier deployed in an Auto Scaling group that runs across multiple Availability Zones. You currently record the number of user file transfers to a log file on the application server, and then write data points from the logs to an Amazon RDS MySQL instance.You aren't happy with how your application scales, and want to implement a new scaling policy based on the average number of user file transfers in a 10-minute period instead of average CPU utilization in the last five minutes. What steps should you take to ensure that your application tier scales based on this new policy? Choose 2 answers A. Create a new CloudWatch alarm based on the Elastic Load Balancing "RequestCount" metric that triggers an Auto Scaling action to scale the application tier.B. Create a new CloudWatch alarm based on a custom metric streaming from the Amazon RDS MySQL instance that triggers an Auto Scaling action to scale the application tier.C. Create a new CloudWatch alarm based on a custom metric published from file transfer logs streaming to CloudWatch that triggers an Auto Scaling action to scale the application tier.D. Create a new Auto Scaling launch configuration that includes an Amazon EC2 user data script that installs a CloudWatch Logs Agent on newly launched instances in the application tier. The agent will be configured to stream the file transfers log tile to CloudWatch.E. Create a new Auto Scaling launch configuration for the application tier that scales based on an Auto Scaling policy that reads the file transfer log data from the Amazon RIDS MySQL instance.F. Create a new Auto Scaling launch configuration that includes an Amazon EC2 user data script that installs an Amazon RDS Logs Agent on newly launched instances in the application tier. The agent will be configured to stream the file transfer data points to the Auto Scaling group. Answer: CD

QUESTION 48Your DevOps team is responsible for a multi-tier, Windows-based web application consisting of web servers, Amazon RDS database instances, and a load balancer behind Amazon Route53. You've been asked by your manager to build a cost-effective rolling deployment solution for this web application.What method should you use? A. Re-deploy your application on an AWS OpsWorks stack. Use the AWS OpsWorks done stack feature to allow updates between duplicate stacks.B. Re-deploy your application on Elastic Beanstalk and take advantage of Elastic BeanStalk rolling updates.C. Re-deploy your application using an AWS CloudFormation template, launch a new AWS CloudFormation stack during each deployment, and then tear down the old stack.D. Re-deploy your application using an AWS CloudFormation template. Use

AWS CloudFormation rolling deployment policies, create a new policy for your AWS CloudFormation stack, and initiate an update stack operation to deploy new code. Answer: D

QUESTION 49 You recently encountered a major bug in your Windows-based web application during a deployment cycle. During this failed deployment, it took the team four hours to roll back to a previously working state, which left customers with a poor user experience. During the post-mortem, your team discussed the need to provide a quicker way to roll back failed deployments. You currently run your web application on Amazon EC2 using Windows 2012R2 and use Elastic Load Balancing for your load balancing needs. Which technique should you use to solve this problem?

A. Create deployable versioned bundles of your application. Store the bundles on Amazon S3. Re-deploy your web application on Elastic Beanstalk, and enable the Elastic Beanstalk auto-rollback feature tied to CloudWatch metrics that define failure.

B. Re-deploy your web application using an AWS OpsWorks stack, and use the AWS OpsWorks auto-rollback feature to initiate a rollback during failures.

C. Create deployable versioned bundles of your application. Store the bundle on Amazon S3. Re-deploy your web application using an AWS OpsWorks stack, and use AWS OpsWorks application versioning to initiate a rollback during failures.

D. Re-deploy your web application using Elastic Beanstalk, and use the Elastic Beanstalk application versions when deploying. During failures, re-deploy the previous version to the Elastic Beanstalk environment.

E. Re-deploy your web application using Elastic Beanstalk, and use the Elastic Beanstalk API to trigger a FailedDeployment API call to initiate a rollback to the previous version.

Answer: D

QUESTION 50 You have a high-traffic application running behind a load balancer with clients that are very sensitive to latency. How should you determine which back-end Amazon Elastic Compute Cloud application instances are causing increased latency so that they can be replaced?

A. By using the Elastic Load Balancing Latency CloudWatch metric.

B. By using the HTTP X-Forwarded-For header for requests from the load balancer.

C. By running a distributed load test to the load balancer.

D. By using the load balancer access logs.

Answer: D

QUESTION 51 Your company operates an application consisting of an AWS CloudFormation stack that contains a load balancer, an Auto Scaling group of web servers, and an Amazon RDS instance. To save time and costs, you update the current test stack when testing minor changes, and create a new stack for major changes. As part of the testing procedure of your application, each version needs to be registered once and only once with a Configuration Management Database (CMDB). What cost-effective solution should you choose to perform this registration?

A. Use Auto Scaling Leader Node functionality to notify the registration application from the UserData script of a single Instance. Use the AWS CloudFormation cfn-hup helper application to receive template updates on the leader node, which then notifies the CMDB.

B. Define an AWS::CloudFormation::CustomResource in the AWS CloudFormation template, with the application version as one of its properties. Modify the CMDB to subscribe to the resource's creation and update notifications.

C. Define an AWS::CloudFormation::HttpRequest in the AWS CloudFormation template, and configure it to notify the CMDB on stack creation and update.

D. Define an AWS::EC2::Instance resource in the AWS CloudFormation template that is configured to run a UserData script to notify the CMDB and then terminate itself on completion.

Answer: B

QUESTION 52 You manage a three-tier web application consisting of an autoscaled web proxy tier, an autoscaled application tier, and an Amazon RDS database tier. You use a load balancer to distribute requests from end users to the web proxy tier and another, internal load balancer to distribute requests between the web tier and the application tier. After deploying a small database schema update, you notice that all of your web and application instances have been terminated. What may have caused this?

A. Your load balancers use an HTTP health check, and the page relies on retrieving data from your database.

B. Your load balancer use TCP health checks to provide application-level health checks.

C. The cooldown period of the Auto Scaling group is too short, so the instances don't have enough time to recover from an issue.

D. Your Auto Scaling group health check type is set to "EC2" to check that the instances themselves are healthy.

Answer: A

QUESTION 53 Your organization has decided to implement a third-party configuration management tool that uses a master server from which nodes pull configuration. You have built a custom base Amazon Machine Image that already has the third-party configuration management agent installed. You want to use the same base AMI in Development, Test and Production environments, each of which has its own master server. How should you configure your Amazon EC2 instances to register with the correct master server on launch?

A. Create a tag for all instances that specifies their environment. When launching instances, provide an Amazon EC2 UserData script that gets this tag by querying the MetaData Service and registers the agent with the master.

B. Use Amazon CloudFormation to describe your environment. Configure an input parameter for the master server hostname/address, and use this parameter within an Amazon EC2 UserData script that registers the agent with the master.

C. Create a script on your third-party configuration management master server that queries the Amazon EC2 API for new instances and registers them with it.

D. Use Amazon Simple Workflow Service to automate the process of registering new instances with your master server. Use an Environment tag in Amazon EC2 to register instances with the correct master server.

Answer: B

QUESTION 54 You have been asked to handle a large data migration from multiple Amazon RDS MySQL instances to a DynamoDB table. You have been given a short amount of time to complete the data migration. What will allow you to complete this complex data

processing workflow? A. Create an Amazon Kinesis data stream, pipe in all of the Amazon RDS data, and direct the data toward a DynamoDB table. B. Write a script in your language of choice, install the script on an Amazon EC2 instance, and then use Auto Scaling groups to ensure that the latency of the migration pipelines never exceeds four seconds in any 15- minute period. C. Write a bash script to run on your Amazon RDS instance that will export data into DynamoDB. D. Create a data pipeline to export Amazon RDS data and import the data into DynamoDB. Answer: D

QUESTION 55 Your application requires a fault-tolerant, low-latency and repeatable method to load configurations files via Auto Scaling when Amazon Elastic Compute Cloud (EC2) instances launch. Which approach should you use to satisfy these requirements? A. Securely copy the content from a running Amazon EC2 instance. B. Use an Amazon EC2 UserData script to copy the configurations from an Amazon Storage Services (S3) bucket. C. Use a script via cfn-init to pull content hosted in an Amazon ElastiCache cluster. D. Use a script via cfn-init to pull content hosted on your on-premises server. E. Use an Amazon EC2 UserData script to pull content hosted on your on-premises server. Answer: B

QUESTION 56 Currently, your deployment process consists of setting your load balancer to point to a maintenance page, turning off ea web application servers, deploying your code, turning the web application servers back on, and removing the maintenance page. Working with your development team, you've agreed that performing rolling deployments of your software would provide a better user experience and a more agile deployment process. Which techniques could you use to provide a cost-effective rolling deployment process? Choose 2 answers. A. Use the Amazon Elastic Cloud Compute (EC2) API to write a service to return a list of servers based on the tags for the application that needs deployment, and use Amazon Simple Queue Service to queue up all servers for a rolling deployment. B. Re-deploy your application on AWS Elastic Beanstalk, and use Elastic Beanstalk rolling deployments. C. Re-deploy your application on an AWS OpsWorks stack, and take advantage of OpsWorks rolling deployments. D. Re-deploy your application using an AWS CloudFormation template, launch a new CloudFormation stack during each deployment, and then tear down the old stack. E. Re-deploy your application using an AWS CloudFormation template with Auto Scaling group, and use update policies to provide rolling updates. F. Using Amazon Simple Workflow Service, create a workflow application that talks to the Amazon EC2 API to deploy your new code in a rolling fashion. Answer: BE

QUESTION 57 You manage a web advertising platform on a single AWS account. This platform produces real-time ad-click data that you store as objects in an Amazon S3 bucket called "click-data." Your advertising partners want to use Amazon Elastic MapReduce in their own AWS accounts to do analytics on the ad-click data. They've asked for immediate access to the ad-click data so that they can run analytics. Which two choices are required to facilitate secure access to this data? Choose 2 answers. A. Create a cross-account TAM role with a trust policy that contains partner AWS account IDs and a unique external ID. B. Create a new IAM group for AWS Data Pipeline users with a trust policy that contains partner AWS account IDs. C. Configure an Amazon S3 bucket policy for the "click-data" bucket that allows Read-Only access to the objects, and associate this policy with an IAM role. D. Configure the Amazon S3 bucket access control list to allow access to the partners Amazon Elastic MapReduce cluster. E. Configure AWS Data Pipeline in the partner AWS accounts to use the web Identity Federation API to access data in the "click-data" bucket. F. Configure AWS Data Pipeline to transfer the data from the "click-data" bucket to the partner's Amazon Elastic MapReduce cluster. Answer: AC

QUESTION 58 You run a SIP-based telephony application that uses Amazon EC2 for its web tier and uses MySQL on Amazon RDS as its database. The application stores only the authentication profile data for its existing users in the database and therefore is read-intensive. Your monitoring system shows that your web instances and the database have high CPU utilization. Which of the following steps should you take in order to ensure the continual availability of your application? Choose 2 answers A. Use a CloudFront RTMP download distribution with the application tier as the origin for the distribution. B. Set up an Auto Scaling group for the application tier and a policy that scales based on the Amazon EC2 CloudWatch CPU utilization metric. C. Vertically scale up the Amazon EC2 instances manually. D. Set up an Auto Scaling group for the application tier and a policy that scales based on the Amazon RDS CloudWatch CPU utilization metric. E. Switch to General Purpose (SSD) Storage from Provisioned IOPS Storage (PIOPS) for the Amazon RDS database. F. Use multiple Amazon RDS read replicas. Answer: BF

QUESTION 59 Your team is responsible for an AWS Elastic Beanstalk application. The business requires that you move to a continuous deployment model, thus releasing updates to the application multiple times per day with zero downtime. What should you do to enable this and still be able to roll back to the previous version almost immediately in an emergency? A. Enable roiling updates in the Elastic Beanstalk environment and set an appropriate pause time for application startup. B. Create a second Elastic Beanstalk environment that runs the new application version, and swap the environment CNAMEs. C. Configure the application to poll for a new application version in your code repository; download and install the new version to each running Elastic Beanstalk instance. D. Create a second Elastic Beanstalk environment with the new application version, and configure the old environment to use the HTTP 301 response code to redirect clients to the new environment. Answer: B

QUESTION 60 Your Company wants to perform A/B testing on a new website feature for 20 percent of its users. The website uses CloudFront for whole site delivery, with some

content cached for up to 24 hours. How do you enable this testing for the required proportion of users while minimizing performance impact?

A. Configure the web servers to handle two domain names. The feature is switched on or off depending on which domain name is used for a request. Configure a CloudFront origin for each domain name, and configure the CloudFront distribution to use one origin for 20 percent of users and the other origin for the other 80 percent.

B. Configure the CloudFront distribution to forward a cookie specific to this feature. For requests where the cookie is not set, the web servers set its value to "on" for 20 percent of responses and "off" for 80 percent. For requests where the cookie is set, the web servers use its value to determine whether the feature should be on or off for the response.

C. Create a second stack of web servers that host the website with the feature on. Using Amazon Route53, create two resource record sets with the same name: one with a weighting of "1" and a value of this new stack; the other a weighting of "4" and a value of the existing stack. Use the resource record set's name as the CloudFront distribution's origin.

D. Invalidate all of the CloudFront distribution's cache items that the feature affects. On future requests, the web servers create responses with the feature on for 20 percent of users, and off for 80 percent. The web servers set "Cache-Control: no-cache" on all of these responses.

Answer: B

More free Lead2pass **AWS-DevOps-Engineer-Professional** exam new questions on Google Drive: <https://drive.google.com/open?id=0B3Syig5i8gpDbVZ1cTB3QnNPQlk> Comparing with others', you will find our AWS-DevOps-Engineer-Professional exam questions are more helpful and precise since all the AWS-DevOps-Engineer-Professional exam content is regularly updated and has been checked for accuracy by our team of Amazon expert professionals. **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]