

[2017 PDF&VCE Lead2pass Latest Amazon AWS-DevOps-Engineer-Professional Exam Questions Free Download (21-40)

Lead2pass 2017 August New Amazon [AWS-DevOps-Engineer-Professional Exam Dumps! 100% Free Download! 100% Pass Guaranteed!](#) Lead2pass provides 100% pass AWS-DevOps-Engineer-Professional exam questions and answers for your Amazon AWS-DevOps-Engineer-Professional exam. We provide Amazon AWS-DevOps-Engineer-Professional exam questions from Lead2pass dumps and answers for the training of AWS-DevOps-Engineer-Professional practice test. Following questions and answers are all new published by **Amazon** Official Exam Center:

<https://www.lead2pass.com/aws-devops-engineer-professional.html> QUESTION 21 Your development team wants account-level access to production instances in order to do live debugging of a highly secure environment. Which of the following should you do?

A. Place the credentials provided by Amazon Elastic Compute Cloud (EC2) into a secure Amazon Sample Storage Service (S3) bucket with encryption enabled. Assign AWS Identity and Access Management (IAM) users to each developer so they can download the credentials file. B. Place an internally created private key into a secure S3 bucket with server-side encryption using customer keys and configuration management, create a service account on all the instances using this private key, and assign IAM users to each developer so they can download the file. C. Place each developer's own public key into a private S3 bucket, use instance profiles and configuration management to create a user account for each developer on all instances, and place the user's public keys into the appropriate account. D. Place the credentials provided by Amazon EC2 onto an MFA encrypted USB drive, and physically share it with each developer so that the private key never leaves the office. Answer: C

QUESTION 22 As part of your continuous deployment process, your application undergoes an I/O load performance test before it is deployed to production using new AMIs. The application uses one Amazon Elastic Block Store (EBS) PIOPS volume per instance and requires consistent I/O performance. Which of the following must be carried out to ensure that I/O load performance tests yield the correct results in a repeatable manner?

A. Ensure that the I/O block sizes for the test are randomly selected. B. Ensure that the Amazon EBS volumes have been pre-warmed by reading all the blocks before the test. C. Ensure that snapshots of the Amazon EBS volumes are created as a backup. D. Ensure that the Amazon EBS volume is encrypted. E. Ensure that the Amazon EBS volume has been pre-warmed by creating a snapshot of the volume before the test. Answer: B

QUESTION 23 After reviewing the last quarter's monthly bills, management has noticed an increase in the overall bill from Amazon. After researching this increase in cost, you discovered that one of your new services is doing a lot of GET Bucket API calls to Amazon S3 to build a metadata cache of all objects in the applications bucket. Your boss has asked you to come up with a new cost-effective way to help reduce the amount of these new GET Bucket API calls. What process should you use to help mitigate the cost?

A. Update your Amazon S3 buckets' lifecycle policies to automatically push a list of objects to a new bucket, and use this list to view objects associated with the application's bucket. B. Create a new DynamoDB table. Use the new DynamoDB table to store all metadata about all objects uploaded to Amazon S3. Any time a new object is uploaded, update the application's internal Amazon S3 object metadata cache from DynamoDB. C. Using Amazon SNS, create a notification on any new Amazon S3 objects that automatically updates a new DynamoDB table to store all metadata about the new object. Subscribe the application to the Amazon SNS topic to update its internal Amazon S3 object metadata cache from the DynamoDB table. D. Upload all images to Amazon SQS, set up SQS lifecycles to move all images to Amazon S3, and initiate an Amazon SNS notification to your application to update the application's Internal Amazon S3 object metadata cache. E. Upload all images to an ElastiCache filecache server. Update your application to now read all file metadata from the ElastiCache filecache server, and configure the ElastiCache policies to push all files to Amazon S3 for long-term storage. Answer: C

QUESTION 24 Your current log analysis application takes more than four hours to generate a report of the top 10 users of your web application. You have been asked to implement a system that can report this information in real time, ensure that the report is always up to date, and handle increases in the number of requests to your web application. Choose the option that is cost-effective and can fulfill the requirements.

A. Publish your data to CloudWatch Logs, and configure your application to autoscale to handle the load on demand. B. Publish your log data to an Amazon S3 bucket. Use AWS CloudFormation to create an Auto Scaling group to scale your post-processing application which is configured to pull down your log files stored an Amazon S3. C. Post your log data to an Amazon Kinesis data stream, and subscribe your log-processing application so that is configured to process your logging data. D. Configure an Auto Scaling group to increase the size of your Amazon EMR duster. E. Create a multi-AZ Amazon RDS MySQL cluster, post the logging data to MySQL, and run a map reduce job to retrieve the required information on user counts. Answer: C

QUESTION 25 You are using Elastic Beanstalk to manage your e-commerce store. The store is based on an open source e-commerce platform and is deployed across multiple instances in an Auto Scaling group. Your development team often creates new "extensions" for the e-commerce store. These extensions include PHP source code

as well as an SQL upgrade script used to make any necessary updates to the database schema. You have noticed that some extension deployments fail due to an error when running the SQL upgrade script. After further investigation, you realize that this is because the SQL script is being executed on all of your Amazon EC2 instances. How would you ensure that the SQL script is only executed once per deployment regardless of how many Amazon EC2 instances are running at the time? A. Use a "Container command" within an Elastic Beanstalk configuration file to execute the script, ensuring that the "leader only" flag is set to true. B. Make use of the Amazon EC2 metadata service to query whether the instance is marked as the leader in the Auto Scaling group. Only execute the script if "true" is returned. C. Use a "Solo Command" within an Elastic Beanstalk configuration file to execute the script. The Elastic Beanstalk service will ensure that the command is only executed once. D. Update the Amazon RDS security group to only allow write access from a single instance in the Auto Scaling group; that way, only one instance will successfully execute the script on the database. Answer: A

QUESTION 26 You are administering a continuous integration application that polls version control for changes and then launches new Amazon EC2 instances for a full suite of build tests. What should you do to ensure the lowest overall cost while being able to run as many tests in parallel as possible? A. Perform syntax checking on the continuous integration system before launching a new Amazon EC2 instance for build test, unit and integration tests. B. Perform syntax and build tests on the continuous integration system before launching the new Amazon EC2 instance unit and integration tests. C. Perform all tests on the continuous integration system, using AWS OpsWorks for unit, integration, and build tests. D. Perform syntax checking on the continuous integration system before launching a new AWS Data Pipeline for coordinating the output of unit, integration, and build tests. Answer: B

QUESTION 27 You are doing a load testing exercise on your application hosted on AWS. While testing your Amazon RDS MySQL DB instance, you notice that when you hit 100% CPU utilization on it, your application becomes non-responsive. Your application is read-heavy. What are methods to scale your data tier to meet the application's needs? Choose 3 answers A. Add Amazon RDS DB read replicas, and have your application direct read queries to them. B. Add your Amazon RDS DB instance to an Auto Scaling group and configure your CloudWatch metric based on CPU utilization. C. Use an Amazon SQS queue to throttle data going to the Amazon RDS DB instance. D. Use ElastiCache in front of your Amazon RDS DB to cache common queries. E. Shard your data set among multiple Amazon RDS DB instances. F. Enable Multi-AZ for your Amazon RDS DB instance. Answer: ADE

QUESTION 28 Your mobile application includes a photo-sharing service that is expecting tens of thousands of users at launch. You will leverage Amazon Simple Storage Service (S3) for storage of the user images, and you must decide how to authenticate and authorize your users for access to these images. You also need to manage the storage of these images. Which two of the following approaches should you use? Choose 2 answers A. Create an Amazon S3 bucket per user, and use your application to generate the S3 URI for the appropriate content. B. Use AWS Identity and Access Management (IAM) user accounts as your application-level user database, and offload the burden of authentication from your application code. C. Authenticate your users at the application level, and use AWS Security Token Service (STS) to grant token-based authorization to S3 objects. D. Authenticate your users at the application level, and send an SMS token message to the user. Create an Amazon S3 bucket with the same name as the SMS message token, and move the user's objects to that bucket. E. Use a key-based naming scheme comprised from the user IDs for all user objects in a single Amazon S3 bucket. Answer: CE

QUESTION 29 You have an Auto Scaling group of Instances that processes messages from an Amazon Simple Queue Service (SQS) queue. The group scales on the size of the queue. Processing involves calling a third-party web service. The web service is complaining about the number of failed and repeated calls it is receiving from you. You have noticed that when the group scales in, instances are being terminated while they are processing. What cost-effective solution can you use to reduce the number of incomplete process attempts? A. Create a new Auto Scaling group with minimum and maximum of 2 and instances running web proxy software. Configure the VPC route table to route HTTP traffic to these web proxies. B. Modify the application running on the instances to enable termination protection while it processes a task and disable it when the processing is complete. C. Increase the minimum and maximum size for the Auto Scaling group, and change the scaling policies so they scale less dynamically. D. Modify the application running on the instances to put itself into an Auto Scaling Standby state while it processes a task and return itself to InService when the processing is complete. Answer: D

QUESTION 30 The operations team and the development team want a single place to view both operating system and application logs. How should you implement this using AWS services? Choose 2 answers A. Using AWS CloudFormation, create a CloudWatch Logs LogGroup and send the operating system and application logs of interest using the CloudWatch Logs Agent. B. Using AWS CloudFormation and configuration management, set up remote logging to send events via UDP packets to CloudTrail. C. Using configuration management, set up remote logging to send events to Amazon Kinesis and insert these into Amazon CloudSearch or Amazon Redshift, depending on available analytic tools. D. Using AWS CloudFormation, create a CloudWatch Logs LogGroup. Because the Cloudwatch Log agent automatically sends all operating system logs, you only have to configure the application logs for sending off-machine. E. Using AWS CloudFormation, merge the

application logs with the operating system logs, and use IAM Roles to allow both teams to have access to view console output from Amazon EC2. Answer: AC QUESTION 31 The project you are working on currently uses a single AWS CloudFormation template to deploy its AWS infrastructure, which supports a multi-tier web application. You have been tasked with organizing the AWS CloudFormation resources so that they can be maintained in the future, and so that different departments such as Networking and Security can review the architecture before it goes to Production. How should you do this in a way that accommodates each department, using their existing workflows? A. Organize the AWS CloudFormation template so that related resources are next to each other in the template, such as VPC subnets and routing rules for Networking and security groups and IAM information for Security. B. Separate the AWS CloudFormation template into a nested structure that has individual templates for the resources that are to be governed by different departments, and use the outputs from the networking and security stacks for the application template that you control. C. Organize the AWS CloudFormation template so that related resources are next to each other in the template for each department's use, leverage your existing continuous integration tool to constantly deploy changes from all parties to the Production environment, and then run tests for validation. D. Use a custom application and the AWS SDK to replicate the resources defined in the current AWS CloudFormation template, and use the existing code review system to allow other departments to approve changes before altering the application for future deployments. Answer: B QUESTION 32 You currently run your infrastructure on Amazon EC2 instances behind an Auto Scaling group. All logs for your application are currently written to ephemeral storage. Recently your company experienced a major bug in code that made it through testing and was ultimately deployed to your fleet. This bug triggered your Auto Scaling group to scale up and back down before you could successfully retrieve the logs off your server to better assist you in troubleshooting the bug. Which technique should you use to make sure you are able to review your logs after your instances have shut down? A. Configure the ephemeral policies on your Auto Scaling group to back up on terminate. B. Configure your Auto Scaling policies to create a snapshot of all ephemeral storage on terminate. C. Install the CloudWatch Logs Agent on your AMI, and configure CloudWatch Logs Agent to stream your logs. D. Install the CloudWatch monitoring agent on your AMI, and set up new SNS alert for CloudWatch metrics that triggers the CloudWatch monitoring agent to backup all logs on the ephemeral drive. E. Install the CloudWatch monitoring agent on your AMI, Update your Auto Scaling policy to enable automated CloudWatch Log copy. Answer: C QUESTION 33 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? 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 QUESTION 34 During metric analysis, your team has determined that the company's website is experiencing response times during peak hours that are higher than anticipated. You currently rely on Auto Scaling to make sure that you are scaling your environment during peak windows. How can you improve your Auto Scaling policy to reduce this high response time? Choose 2 answers. A. Push custom metrics to CloudWatch to monitor your CPU and network bandwidth from your servers, which will allow your Auto Scaling policy to have better fine-grain insight. B. Increase your Auto Scaling group's number of max servers. C. Create a script that runs and monitors your servers; when it detects an anomaly in load, it posts to an Amazon SNS topic that triggers Elastic Load Balancing to add more servers to the load balancer. D. Push custom metrics to CloudWatch for your application that include more detailed information about your web application, such as how many requests it is handling and how many are waiting to be processed. E. Update the CloudWatch metric used for your Auto Scaling policy, and enable sub-minute granularity to allow auto scaling to trigger faster. Answer: BD QUESTION 35 You are responsible for your company's large multi-tiered Windows-based web application running on Amazon EC2 instances situated behind a load balancer. While reviewing metrics, you've started noticing an upwards trend for slow customer page load time. Your manager has asked you to come up with a solution to ensure that customer load time is not affected by too many requests per second. Which technique would you use to solve this issue? A. Re-deploy your infrastructure using an AWS CloudFormation template. Configure Elastic Load Balancing health checks to initiate a new AWS CloudFormation stack when health checks return failed. B. Re-deploy your infrastructure using an AWS CloudFormation template. Spin up a second AWS CloudFormation stack. Configure Elastic Load Balancing SpillOver functionality to spill over any slow connections to the second AWS CloudFormation stack. C. Re-deploy your infrastructure using AWS CloudFormation, Elastic Beanstalk, and Auto Scaling. Set up your Auto Scaling group policies to scale based on the number of requests per second as well as the current customer

load time.D. Re-deploy your application using an Auto Scaling template. Configure the Auto Scaling template to spin up a new Elastic Beanstalk application when the customer load time surpasses your threshold. Answer: C

QUESTION 36 Your company has multiple applications running on AWS. Your company wants to develop a tool that notifies on-call teams immediately via email when an alarm is triggered in your environment. You have multiple on-call teams that work different shifts, and the tool should handle notifying the correct teams at the correct times. How should you implement this solution?

A. Create an Amazon SNS topic and an Amazon SQS queue. Configure the Amazon SQS queue as a subscriber to the Amazon SNS topic. Configure CloudWatch alarms to notify this topic when an alarm is triggered. Create an Amazon EC2 Auto Scaling group with both minimum and desired Instances configured to 0. Worker nodes in this group spawn when messages are added to the queue. Workers then use Amazon Simple Email Service to send messages to your on-call teams.

B. Create an Amazon SNS topic and configure your on-call team email addresses as subscribers. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to this new topic. Notifications will be sent to on-call users when a CloudWatch alarm is triggered.

C. Create an Amazon SNS topic and configure your on-call team email addresses as subscribers. Create a secondary Amazon SNS topic for alarms and configure your CloudWatch alarms to notify this topic when triggered. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggered. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to the first topic so that on-call engineers receive alerts.

D. Create an Amazon SNS topic for each on-call group, and configure each of these with the team member emails as subscribers. Create another Amazon SNS topic and configure your CloudWatch alarms to notify this topic when triggered. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggered. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to the correct team topic when on shift. Answer: D

QUESTION 37 Your company releases new features with high frequency while demanding high application availability. As part of the application's A/B testing, logs from each updated Amazon EC2 instance of the application need to be analyzed in near real-time, to ensure that the application is working flawlessly after each deployment. If the logs show any anomalous behavior, then the application version of the instance is changed to a more stable one. Which of the following methods should you use for shipping and analyzing the logs in a highly available manner?

A. Ship the logs to Amazon S3 for durability and use Amazon EMR to analyze the logs in a batch manner each hour.

B. Ship the logs to Amazon CloudWatch Logs and use Amazon EMR to analyze the logs in a batch manner each hour.

C. Ship the logs to an Amazon Kinesis stream and have the consumers analyze the logs in a live manner.

D. Ship the logs to a large Amazon EC2 instance and analyze the logs in a live manner.

E. Store the logs locally on each instance and then have an Amazon Kinesis stream pull the logs for live analysis. Answer: C

QUESTION 38 You have a code repository that uses Amazon S3 as a data store. During a recent audit of your security controls, some concerns were raised about maintaining the integrity of the data in the Amazon S3 bucket. Another concern was raised around securely deploying code from Amazon S3 to applications running on Amazon EC2 in a virtual private cloud. What are some measures that you can implement to mitigate these concerns? Choose 2 answers.

A. Add an Amazon S3 bucket policy with a condition statement to allow access only from Amazon EC2 instances with RFC 1918 IP addresses and enable bucket versioning.

B. Add an Amazon S3 bucket policy with a condition statement that requires multi-factor authentication in order to delete objects and enable bucket versioning.

C. Use a configuration management service to deploy AWS Identity and Access Management user credentials to the Amazon EC2 instances. Use these credentials to securely access the Amazon S3 bucket when deploying code.

D. Create an Amazon Identity and Access Management role with authorization to access the Amazon S3 bucket, and launch all of your application's Amazon EC2 instances with this role.

E. Use AWS Data Pipeline to lifecycle the data in your Amazon S3 bucket to Amazon Glacier on a weekly basis.

F. Use AWS Data Pipeline with multi-factor authentication to securely deploy code from the Amazon S3 bucket to your Amazon EC2 instances. Answer: BD

QUESTION 39 You have an application consisting of a stateless web server tier running on Amazon EC2 instances behind a 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.

A. Set up a third-party monitoring solution on a cluster of Amazon EC2 instances in order to emit custom CloudWatch 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 Scaling group for the web server tier along with an Auto Scaling policy that uses the Amazon RDS DB CPU utilization CloudWatch metric to scale the instances.

D. Set up an Auto Scaling group 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 Scaling group 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: AD

QUESTION 40 Your

application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment? A. Set up Amazon Route 53 health checks to fail over from any Amazon EC2 instance that is currently being deployed to. B. Using AWS CloudFormation, create a test stack for validating the code, and then deploy the code to each production Amazon EC2 instance. C. Create a new load balancer with new Amazon EC2 instances, carry out the deployment, and then switch DNS over to the new load balancer using Amazon Route 53 after testing. D. Launch more Amazon EC2 instances to ensure high availability, de-register each Amazon EC2 instance from the load balancer, upgrade it, and test it, and then register it again with the load balancer. Answer: C More free Lead2pass

AWS-DevOps-Engineer-Professional exam new questions on Google Drive:

<https://drive.google.com/open?id=0B3Syig5i8gpDbVZ1cTB3QnNPQlk> Lead2pass is the leader in

AWS-DevOps-Engineer-Professional certification test questions with training materials for Amazon

AWS-DevOps-Engineer-Professional exam dumps. Lead2pass Amazon training tools are constantly being revised and updated. We

100% guarantee Amazon AWS-DevOps-Engineer-Professional exam questions with quality and reliability which will help you pass

Amazon AWS-DevOps-Engineer-Professional exam. **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]