

[2017 PDF&VCE New Lead2pass Amazon AWS-DevOps-Engineer-Professional New Questions Free Download (141-160)]

Lead2pass 2017 September New Amazon [AWS-DevOps-Engineer-Professional Exam Dumps!](#) [100% Free Download!](#) [100% Pass Guaranteed!](#) Test your preparation for Amazon AWS-DevOps-Engineer-Professional with these actual AWS-DevOps-Engineer-Professional new questions below. Exam questions are a sure method to validate one's preparation for actual certification exam. Following questions and answers are all new published by **Amazon** Official Exam Center:

<https://www.lead2pass.com/aws-devops-engineer-professional.html> QUESTION 141 Which is not a restriction on AWS EBS Snapshots? A. Snapshots which are shared cannot be used as a basis for other snapshots. B. You cannot share a snapshot containing an AWS Access Key ID or AWS Secret Access Key. C. You cannot share unencrypted snapshots. D. Snapshot restorations are restricted to the region in which the snapshots are created. Answer: A Explanation: Snapshots shared with other users are usable in full by the recipient, including but limited to the ability to base modified volumes and snapshots.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modifying-snapshot-permissions.html> QUESTION 142 You need to deploy a new application version to production. Because the deployment is high-risk, you need to roll the new version out to users over a number of hours, to make sure everything is working correctly. You need to be able to control the proportion of users seeing the new version of the application down to the percentage point. You use ELB and EC2 with Auto Scaling Groups and custom AMIs with your code pre-installed assigned to Launch Configurations. There are no database-level changes during your deployment. You have been told you cannot spend too much money, so you must not increase the number of EC2 instances much at all during the deployment, but you also need to be able to switch back to the original version of code quickly if something goes wrong. What is the best way to meet these requirements? A. Create a second ELB, Auto Scaling Launch Configuration, and Auto Scaling Group using the Launch Configuration. Create AMIs with all code pre-installed. Assign the new AMI to the second Auto Scaling Launch Configuration. Use Route53 Weighted Round Robin Records to adjust the proportion of traffic hitting the two ELBs. B. Use the Blue-Green deployment method to enable the fastest possible rollback if needed. Create a full second stack of instances and cut the DNS over to the new stack of instances, and change the DNS back if a rollback is needed. C. Create AMIs with all code pre-installed. Assign the new AMI to the Auto Scaling Launch Configuration, to replace the old one. Gradually terminate instances running the old code (launched with the old Launch Configuration) and allow the new AMIs to boot to adjust the traffic balance to the new code. On rollback, reverse the process by doing the same thing, but changing the AMI on the Launch Config back to the original code. D. Migrate to use AWS Elastic Beanstalk. Use the established and well-tested Rolling Deployment setting AWS provides on the new Application Environment, publishing a zip bundle of the new code and adjusting the wait period to spread the deployment over time. Re-deploy the old code bundle to rollback if needed. Answer: A Explanation: Only Weighted Round Robin DNS Records and reverse proxies allow such fine-grained tuning of traffic splits. The Blue-Green option does not meet the requirement that we mitigate costs and keep overall EC2 fleet size consistent, so we must select the 2 ELB and ASG option with WRR DNS tuning. This method is called A/B deployment and/or Canary deployment.

<https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf> QUESTION 143 What is required to achieve gigabit network throughput on EC2? You already selected cluster-compute, 10GB instances with enhanced networking, and your workload is already network-bound, but you are not seeing 10 gigabit speeds. A. Enable bplex networking on your servers, so packets are non-blocking in both directions and there's no switching overhead. B. Ensure the instances are in different VPCs so you don't saturate the Internet Gateway on any one VPC. C. Select PIOPS for your drives and mount several, so you can provision sufficient disk throughput. D. Use a placement group for your instances so the instances are physically near each other in the same Availability Zone. Answer: D Explanation: You are not guaranteed 10 gigabit performance, except within a placement group. A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

QUESTION 144 If I want CloudFormation stack status updates to show up in a continuous delivery system in as close to real time as possible, how should I achieve this? A. Use a long-poll on the Resources object in your CloudFormation stack and display those state changes in the UI for the system. B. Use a long-poll on the `ListStacks` API call for your CloudFormation stack and display those state changes in the UI for the system. C. Subscribe your continuous delivery system to an SNS topic that you also tell your CloudFormation stack to publish events into. D. Subscribe your continuous delivery system to an SQS queue that you also tell your CloudFormation stack to publish events into. Answer: C Explanation: Use NotificationARNs.member.N when making a CreateStack call to push stack events into SNS in nearly real-time.

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-monitor-stack.html> QUESTION 145

What does it mean if you have zero IOPS and a non-empty I/O queue for all EBS volumes attached to a running EC2 instance? A. The I/O queue is buffer flushing. B. Your EBS disk head(s) is/are seeking magnetic stripes. C. The EBS volume is unavailable. D. You need to re-mount the EBS volume in the OS. Answer: C Explanation: This is the definition of Unavailable from the EC2 and EBS SLA. "Unavailable" and "Unavailability" mean...

For Amazon EBS, when all of your attached volumes perform zero read write IO, with pending IO in the queue. <https://aws.amazon.com/ec2/sla/> QUESTION 146

From a compliance and security perspective, which of these statements is true? A. You do not ever need to rotate access keys for AWS IAM Users. B. You do not ever need to rotate access keys for AWS IAM Roles, nor AWS IAM Users. C. None of the other statements are true. D. You do not ever need to rotate access keys for AWS IAM Roles. Answer: D Explanation: IAM Role Access Keys are auto-rotated by AWS on your behalf; you do not need to rotate them. The application is granted the permissions for the actions and resources that you've defined for the role through the security credentials associated with the role. These security credentials are temporary and we rotate them automatically. We make new credentials available at least five minutes prior to the expiration of the old credentials.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html> QUESTION 147

Which of these configuration or deployment practices is a security risk for RDS? A. Storing SQL function code in plaintext. B. Non-Multi-AZ RDS instance. C. Having RDS and EC2 instances exist in the same subnet. D. RDS in a public subnet. Answer: D Explanation: Making RDS accessible to the public internet in a public subnet poses a security risk, by making your database directly addressable and spammable. DB instances deployed within a VPC can be configured to be accessible from the Internet or from EC2 instances outside the VPC. If a VPC security group specifies a port access such as TCP port 22, you would not be able to access the DB instance because the firewall for the DB instance provides access only via the IP addresses specified by the DB security groups the instance is a member of and the port defined when the DB instance was created.

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.RDSSecurityGroups.html> QUESTION 148

Which of these is not a reason a Multi-AZ RDS instance will failover? A. An Availability Zone outage. B. A manual failover of the DB instance was initiated using Reboot with failover. C. To autoscale to a higher instance class. D. The primary DB instance fails. Answer: C Explanation: The primary DB instance switches over automatically to the standby replica if any of the > following conditions occur: An Availability Zone outage, the primary DB instance fails, the DB instance's server type is changed, the operating system of the DB instance is, undergoing software patching, a manual failover of the DB instance was initiated using Reboot with failover

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html> QUESTION 149

You need to create an audit log of all changes to customer banking data. You use DynamoDB to store this customer banking data. It's important not to lose any information due to server failures. What is an elegant way to accomplish this? A. Use a DynamoDB StreamSpecification and stream all changes to AWS Lambda. Log the changes to AWS CloudWatch Logs, removing sensitive information before logging. B. Before writing to DynamoDB, do a pre-write acknowledgment to disk on the application server, removing sensitive information before logging. Periodically rotate these log files into S3. C. Use a DynamoDB StreamSpecification and periodically flush to an EC2 instance store, removing sensitive information before putting the objects. Periodically flush these batches to S3. D. Before writing to DynamoDB, do a pre-write acknowledgment to disk on the application server, removing sensitive information before logging. Periodically pipe these files into CloudWatch Logs. Answer: A Explanation: All suggested periodic options are sensitive to server failure during or between periodic flushes. Streaming to Lambda and then logging to CloudWatch Logs will make the system resilient to instance and Availability Zone failures. <http://docs.aws.amazon.com/lambda/latest/dg/with-ddb.html> QUESTION 150

You need your API backed by DynamoDB to stay online during a total regional AWS failure. You can tolerate a couple minutes of lag or slowness during a large failure event, but the system should recover with normal operation after those few minutes. What is a good approach? A. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standby in another region. Create an Auto Scaling Group behind an ELB in each of the two regions DynamoDB is running in. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records. B. Set up a DynamoDB Multi-Region table. Create an Auto Scaling Group behind an ELB in each of the two regions DynamoDB is running in. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records. C. Set up a DynamoDB Multi-Region table. Create a cross-region ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross-region ELB. D. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standby in another region. Create a cross-region ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross-region ELB. Answer: A Explanation: There is no such thing as a cross-region ELB, nor such thing as a cross-region Auto Scaling Group, nor such thing as a DynamoDB Multi-Region Table. The only option that makes sense is the cross-regional replication version with two ELBs and ASGs with Route53 Failover and Latency DNS.

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.CrossRegionRepl.html> QUESTION 151 You have an asynchronous processing application using an Auto Scaling Group and an SQS Queue. The Auto Scaling Group scales according to the depth of the job queue. The completion velocity of the jobs has gone down, the Auto Scaling Group size has maxed out, but the inbound job velocity did not increase. What is a possible issue? A. Some of the new jobs coming in are malformed and unprocessable. B. The routing tables changed and none of the workers can process events anymore. C. Someone changed the IAM Role Policy on the instances in the worker group and broke permissions to access the queue. D. The scaling metric is not functioning correctly. Answer: A Explanation: The IAM Role must be fine, as if it were broken, NO jobs would be processed since the system would never be able to get any queue messages. The same reasoning applies to the routing table change. The scaling metric is fine, as instance count increased when the queue depth increased due to more messages entering than exiting. Thus, the only reasonable option is that some of the recent messages must be malformed and unprocessable.

[https://github.com/andrew-templeton/cloudacademy/blob/fca920b45234bbe99cc0e8efb9c65134884dd48](https://github.com/andrew-templeton/cloudacademy/blob/fca920b45234bbe99cc0e8efb9c65134884dd489/questions/null) 9/questions/null QUESTION 152 Your company wants to understand where cost is coming from in the company's production AWS account. There are a number of applications and services running at any given time. Without expending too much initial development time, how best can you give the business a good understanding of which applications cost the most per month to operate? A. Create an automation script which periodically creates AWS Support tickets requesting detailed intra-month information about your bill. B. Use custom CloudWatch Metrics in your system, and put a metric data point whenever cost is incurred. C. Use AWS Cost Allocation Tagging for all resources which support it. Use the Cost Explorer to analyze costs throughout the month. D. Use the AWS Price API and constantly running resource inventory scripts to calculate total price based on multiplication of consumed resources over time. Answer: C Explanation: Cost Allocation Tagging is a built-in feature of AWS, and when coupled with the Cost Explorer, provides a simple and robust way to track expenses. You can also use tags to filter views in Cost Explorer. Note that before you can filter views by tags in Cost Explorer, you must have applied tags to your resources and activate them, as described in the following sections. For more information about Cost Explorer, see Analyzing Your Costs with Cost Explorer.

<http://docs.aws.amazon.com/awssaccountbilling/latest/aboutv2/cost-alloc-tags.html> QUESTION 153 There is a very serious outage at AWS. EC2 is not affected, but your EC2 instance deployment scripts stopped working in the region with the outage. What might be the issue? A. The AWS Console is down, so your CLI commands do not work. B. S3 is unavailable, so you can't create EBS volumes from a snapshot you use to deploy new volumes. C. AWS turns off the `DeployCode` API call when there are major outages, to protect from system floods. D. None of the other answers make sense. If EC2 is not affected, it must be some other issue. Answer: B Explanation: S3 stores all snapshots. If S3 is unavailable, snapshots are unavailable. Amazon EC2 also uses Amazon S3 to store snapshots (backup copies) of the data volumes. You can use snapshots for recovering data quickly and reliably in case of application or system failures. You can also use snapshots as a baseline to create multiple new data volumes, expand the size of an existing data volume, or move data volumes across multiple Availability Zones, thereby making your data usage highly scalable. For more information about using data volumes and snapshots, see Amazon Elastic Block Store.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonS3.html> QUESTION 154 Which of the following tools does not directly support AWS OpsWorks, for monitoring your stacks? A. AWS Config B. Amazon CloudWatch Metrics C. AWS CloudTrail D. Amazon CloudWatch Logs Answer: A Explanation: You can monitor your stacks in the following ways: AWS OpsWorks uses Amazon CloudWatch to provide thirteen custom metrics with detailed monitoring for each instance in the stack; AWS OpsWorks integrates with AWS CloudTrail to log every AWS OpsWorks API call and store the data in an Amazon S3 bucket; You can use Amazon CloudWatch Logs to monitor your stack's system, application, and custom logs.

<http://docs.aws.amazon.com/opsworks/latest/userguide/monitoring.html> QUESTION 155 What is a circular dependency in AWS CloudFormation? A. When a Template references an earlier version of itself. B. When Nested Stacks depend on each other. C. When Resources form a DependOn loop. D. When a Template references a region, which references the original Template. Answer: C Explanation: To resolve a dependency error, add a DependsOn attribute to resources that depend on other resources in your template. In some cases, you must explicitly declare dependencies so that AWS CloudFormation can create or delete resources in the correct order. For example, if you create an Elastic IP and a VPC with an Internet gateway in the same stack, the Elastic IP must depend on the Internet gateway attachment. For additional information, see DependsOn Attribute.

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/troubleshooting.html#troubleshootin> g-errors-dependence-error QUESTION 156 You need to run a very large batch data processing job one time per day. The source data exists entirely in S3, and the output of the processing job should also be written to S3 when finished. If you need to version control this processing job and all setup and teardown logic for the system, what approach should you use? A. Model an AWS EMR job in AWS Elastic Beanstalk. B. Model an AWS EMR job in AWS CloudFormation. C. Model an AWS EMR job in AWS OpsWorks. D. Model an AWS EMR

job in AWS CLI Composer. Answer: BExplanation:To declaratively model build and destroy of a cluster, you need to use AWS CloudFormation. OpsWorks and Elastic Beanstalk cannot directly model EMR Clusters. The CLI is not declarative, and CLI Composer does not exist.<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-emr-cluster.html>

QUESTION 157What is true of the way that encryption works with EBS? A. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot creates an encrypted volume when specified / requested.B. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot creates an encrypted volume when specified / requested.C. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot always creates an encrypted volume.D. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot always creates an encrypted volume. Answer: CExplanation:Snapshots that are taken from encrypted volumes are automatically encrypted. Volumes that are created from encrypted snapshots are also automatically encrypted. Your encrypted volumes and any associated snapshots always remain protected. For more information, see Amazon EBS Encryption. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html>

QUESTION 158When thinking of AWS OpsWorks, which of the following is true? A. Stacks have many layers, layers have many instances.B. Instances have many stacks, stacks have many layers.C. Layers have many stacks, stacks have many instances.D. Layers have many instances, instances have many stacks. Answer: AExplanation:The stack is the core AWS OpsWorks component. It is basically a container for AWS resources--Amazon EC2 instances, Amazon RDS database instances, and so on--that have a common purpose and should be logically managed together. You define the stack's constituents by adding one or more layers. A layer represents a set of Amazon EC2 instances that serve a particular purpose, such as serving applications or hosting a database server. An instance represents a single computing resource, such as an Amazon EC2 instance.<http://docs.aws.amazon.com/opsworks/latest/userguide/welcome.html>

QUESTION 159When thinking of AWS Elastic Beanstalk, which statement is true? A. Worker tiers pull jobs from SNS.B. Worker tiers pull jobs from HTTP.C. Worker tiers pull jobs from JSON.D. Worker tiers pull jobs from SQS. Answer: D Explanation:Elastic Beanstalk installs a daemon on each Amazon EC2 instance in the Auto Scaling group to process Amazon SQS messages in the worker environment. The daemon pulls data off the Amazon SQS queue, inserts it into the message body of an HTTP POST request, and sends it to a user-configurable URL path on the local host. The content type for the message body within an HTTP POST request is application/json by default.
<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.html>

QUESTION 160Your company needs to automate 3 layers of a large cloud deployment. You want to be able to track this deployment's evolution as it changes over time, and carefully control any alterations. What is a good way to automate a stack to meet these requirements? A. Use OpsWorks Stacks with three layers to model the layering in your stack.B. Use CloudFormation Nested Stack Templates, with three child stacks to represent the three logical layers of your cloud.C. Use AWS Config to declare a configuration set that AWS should roll out to your cloud.D. Use Elastic Beanstalk Linked Applications, passing the important DNS entries between layers using the metadata interface. Answer: BExplanation:Only CloudFormation allows source controlled, declarative templates as the basis for stack automation. Nested Stacks help achieve clean separation of layers while simultaneously providing a method to control all layers at once when needed.
<https://blogs.aws.amazon.com/application-management/post/Tx1T9JYQOS8AB9I/Use-Nested-Stacks-to-Create-Reusable-Templates-and-Support-Role-Specialization>

More free Lead2pass **AWS-DevOps-Engineer-Professional** exam new questions on Google Drive: <https://drive.google.com/open?id=0B3Syig5i8gpDbVZ1cTB3QnNPQIk> These Amazon AWS-DevOps-Engineer-Professional exam questions are all a small selection of questions. If you want to practice more questions for actual AWS-DevOps-Engineer-Professional exam, use the links at the end of this document. Also you can find links for AWS-DevOps-Engineer-Professional VCE software that is great for preparation and self-assessment for 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]