

[2017 PDF&VCE 100% Pass Lead2pass AWS-DevOps-Engineer-Professional New Questions Free Version (121-140)]

Lead2pass 2017 September New Amazon [AWS-DevOps-Engineer-Professional Exam Dumps! 100% Free Download! 100% Pass Guaranteed!](#) Lead2pass is constantly updating AWS-DevOps-Engineer-Professional exam dumps. We will provide our customers with the latest and the most accurate exam questions and answers that cover a comprehensive knowledge point, which will help you easily prepare for AWS-DevOps-Engineer-Professional exam and successfully pass your exam. You just need to spend 20-30 hours on studying the exam dumps. Following questions and answers are all new published by **Amazon** Official Exam Center:

<https://www.lead2pass.com/aws-devops-engineer-professional.html> QUESTION 121 Your CTO thinks your AWS account was hacked. What is the only way to know for certain if there was unauthorized access and what they did, assuming your hackers are very sophisticated AWS engineers and doing everything they can to cover their tracks? A. Use CloudTrail Log File Integrity Validation. AB. Use AWS Config SNS Subscriptions and process events in real time. C. Use CloudTrail backed up to AWS S3 and Glacier. D. Use AWS Config Timeline forensics. Answer: A Explanation: You must use CloudTrail Log File Validation (default or custom implementation), as any other tracking method is subject to forgery in the event of a full account compromise by sophisticated enough hackers. Validated log files are invaluable in security and forensic investigations. For example, a validated log file enables you to assert positively that the log file itself has not changed, or that particular user credentials performed specific API activity. The CloudTrail log file integrity validation process also lets you know if a log file has been deleted or changed, or assert positively that no log files were delivered to your account during a given period of time.

<http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html> QUESTION 122 Which of these is not a Pseudo Parameter in AWS CloudFormation? A. AWS::StackName B. AWS::AccountId C. AWS::StackArn D. AWS::NotificationARNs Answer: C Explanation: This is the complete list of Pseudo Parameters: AWS::AccountId, AWS::NotificationARNs, AWS::NoValue, AWS::Region, AWS::StackId, AWS::StackName

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/pseudo-parameter-reference.html> QUESTION 123 What is the scope of an EBS volume? A. VPC B. Region C. Placement Group D. Availability Zone Answer: D Explanation: An Amazon EBS volume is tied to its Availability Zone and can be attached only to instances in the same Availability Zone.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/resources.html> QUESTION 124 You are experiencing performance issues writing to a DynamoDB table. Your system tracks high scores for video games on a marketplace. Your most popular game experiences all of the performance issues. What is the most likely problem? A. DynamoDB's vector clock is out of sync, because of the rapid growth in request for the most popular game. B. You selected the Game ID or equivalent identifier as the primary partition key for the table. C. Users of the most popular video game each perform more read and write requests than average. D. You did not provision enough read or write throughput to the table. Answer: B Explanation: The primary key selection dramatically affects performance consistency when reading or writing to DynamoDB. By selecting a key that is tied to the identity of the game, you forced DynamoDB to create a hotspot in the table partitions, and over-request against the primary key partition for the popular game. When it stores data, DynamoDB divides a table's items into multiple partitions, and distributes the data primarily based upon the partition key value. The provisioned throughput associated with a table is also divided evenly among the partitions, with no sharing of provisioned throughput across partitions.

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuidelinesForTables.html#GuidelinesForTables.UniformWorkload> QUESTION 125

You meet once per month with your operations team to review the past month's data. During the meeting, you realize that 3 weeks ago, your monitoring system which pings over HTTP from outside AWS recorded a large spike in latency on your 3-tier web service API. You use DynamoDB for the database layer, ELB, EBS, and EC2 for the business logic tier, and SQS, ELB, and EC2 for the presentation layer. Which of the following techniques will NOT help you figure out what happened? A. Check your CloudTrail log history around the spike's time for any API calls that caused slowness. B. Review CloudWatch Metrics graphs to determine which component(s) slowed the system down. C. Review your ELB access logs in S3 to see if any ELBs in your system saw the latency. D. Analyze your logs to detect bursts in traffic at that time. Answer: B Explanation: Metrics data are available for 2 weeks. If you want to store metrics data beyond that duration, you can retrieve it using our GetMetricStatistics API as well as a number of applications and tools offered by AWS partners.

<https://aws.amazon.com/cloudwatch/faqs/> QUESTION 126 Which of these is not an intrinsic function in AWS CloudFormation? A. Fn::Split B. Fn::FindInMap C. Fn::Select D. Fn::GetAZs Answer: A Explanation: This is the complete list of Intrinsic Functions...: Fn::Base64, Fn::And, Fn::Equals, Fn::If, Fn::Not, Fn::Or, Fn::FindInMap, Fn::GetAtt, Fn::GetAZs, Fn::Join, Fn::Select, Ref <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference.html> QUESTION 127 For

AWS CloudFormation, which is true? A. Custom resources using SNS have a default timeout of 3 minutes. B. Custom resources using SNS do not need a `ServiceToken` property. C. Custom resources using Lambda and `Code.ZipFile` allow inline nodejs resource composition. D. Custom resources using Lambda do not need a `ServiceToken` property. Answer: C. Explanation: Code is a property of the AWS::Lambda::Function resource that enables you to specify the source code of an AWS Lambda (Lambda) function. You can point to a file in an Amazon Simple Storage Service (Amazon S3) bucket or specify your source code as inline text (for nodejs runtime environments only).

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-custom-resources.html> QUESTION 128 Your API requires the ability to stay online during AWS regional failures. Your API does not store any state, it only aggregates data from other sources - you do not have a database. What is a simple but effective way to achieve this uptime goal? A. Use a CloudFront distribution to serve up your API. Even if the region your API is in goes down, the edge locations CloudFront uses will be fine. B. Use an ELB and a cross-zone ELB deployment to create redundancy across datacenters. Even if a region fails, the other AZ will stay online. C. Create a Route53 Weighted Round Robin record, and if one region goes down, have that region redirect to the other region. D. Create a Route53 Latency Based Routing Record with Failover and point it to two identical deployments of your stateless API in two different regions. Make sure both regions use Auto Scaling Groups behind ELBs. Answer: D. Explanation: Latency Based Records allow request distribution when all is well with both regions, and the Failover component enables fallbacks between regions. By adding in the ELB and ASG, your system in the surviving region can expand to meet 100% of demand instead of the original fraction, whenever failover occurs. <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html>

QUESTION 129 You are designing an enterprise data storage system. Your data management software system requires mountable disks and a real filesystem, so you cannot use S3 for storage. You need persistence, so you will be using AWS EBS Volumes for your system. The system needs as low-cost storage as possible, and access is not frequent or high throughput, and is mostly sequential reads. Which is the most appropriate EBS Volume Type for this scenario? A. gp1 B. io1 C. standard D. gp2. Answer: C. Explanation: standard volumes, or Magnetic volumes, are best for: Cold workloads where data is infrequently accessed, or scenarios where the lowest storage cost is important.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html> QUESTION 130 You need to deploy an AWS stack in a repeatable manner across multiple environments. You have selected CloudFormation as the right tool to accomplish this, but have found that there is a resource type you need to create and model, but is unsupported by CloudFormation. How should you overcome this challenge? A. Use a CloudFormation Custom Resource Template by selecting an API call to proxy for create, update, and delete actions. CloudFormation will use the AWS SDK, CLI, or API method of your choosing as the state transition function for the resource type you are modeling. B. Submit a ticket to the AWS Forums. AWS extends CloudFormation Resource Types by releasing tooling to the AWS Labs organization on GitHub. Their response time is usually 1 day, and they complete requests within a week or two. C. Instead of depending on CloudFormation, use Chef, Puppet, or Ansible to author Heat templates, which are declarative stack resource definitions that operate over the OpenStack hypervisor and cloud environment. D. Create a CloudFormation Custom Resource Type by implementing create, update, and delete functionality, either by subscribing a Custom Resource Provider to an SNS topic, or by implementing the logic in AWS Lambda. Answer: D. Explanation: Custom resources provide a way for you to write custom provisioning logic in AWS CloudFormation template and have AWS CloudFormation run it during a stack operation, such as when you create, update or delete a stack. For more information, see Custom Resources.

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-custom-resources.html> QUESTION 131 You run a 2000-engineer organization. You are about to begin using AWS at a large scale for the first time. You want to integrate with your existing identity management system running on Microsoft Active Directory, because your organization is a power-user of Active Directory. How should you manage your AWS identities in the most simple manner? A. Use a large AWS Directory Service Simple AD. B. Use a large AWS Directory Service AD Connector. C. Use an Sync Domain running on AWS Directory Service. D. Use an AWS Directory Sync Domain running on AWS Lambda. Answer: B. Explanation: You must use AD Connector as a power-user of Microsoft Active Directory. Simple AD only works with a subset of AD functionality. Sync Domains do not exist; they are made up answers. AD Connector is a directory gateway that allows you to proxy directory requests to your on-premises Microsoft Active Directory, without caching any information in the cloud. AD Connector comes in 2 sizes; small and large. A small AD Connector is designed for smaller organizations of up to 500 users. A large AD Connector is designed for larger organizations of up to 5,000 users. <https://aws.amazon.com/directoryservice/details/>

QUESTION 132 When thinking of AWS OpsWorks, which of the following is not an instance type you can allocate in a stack layer? A. 24/7 instances B. Spot instances C. Time-based instances D. Load-based instances. Answer: B. Explanation: AWS OpsWorks supports the following instance types, which are characterized by how they are started and stopped. 24/7 instances are started manually and run until you stop them. Time-based

instances are run by AWS OpsWorks on a specified daily and weekly schedule. They allow your stack to automatically adjust the number of instances to accommodate predictable usage patterns. Load-based instances are automatically started and stopped by AWS OpsWorks, based on specified load metrics, such as CPU utilization. They allow your stack to automatically adjust the number of instances to accommodate variations in incoming traffic. Load-based instances are available only for Linux-based stacks.

<http://docs.aws.amazon.com/opsworks/latest/userguide/welcome.html> QUESTION 133 Which of these is not a CloudFormation Helper Script? A. cfn-signal B. cfn-hup C. cfn-request D. cfn-get-metadata Answer: C Explanation: This is the complete list of CloudFormation Helper Scripts: cfn-init, cfn-signal, cfn-get-metadata, cfn-hup

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-helper-scripts-reference.html> QUESTION 134 Your team wants to begin practicing continuous delivery using CloudFormation, to enable automated builds and deploys of whole, versioned stacks or stack layers. You have a 3-tier, mission-critical system. Which of the following is NOT a best practice for using CloudFormation in a continuous delivery environment? A. Use the AWS CloudFormation `ValidateTemplate` call before publishing changes to AWS. B. Model your stack in one template, so you can leverage CloudFormation's state management and dependency resolution to propagate all changes. C. Use CloudFormation to create brand new infrastructure for all stateless resources on each push, and run integration tests on that set of infrastructure. D. Parametrize the template and use

`Mappings` to ensure your template works in multiple Regions. Answer: B Explanation: Putting all resources in one stack is a bad idea, since different tiers have different life cycles and frequencies of change. For additional guidance about organizing your stacks, you can use two common frameworks: a multi-layered architecture and service-oriented architecture (SOA).

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html#organizingstacks> QUESTION 135 You need to replicate API calls across two systems in real time. What tool should you use as a buffer and transport mechanism for API call events? A. AWS SQS B. AWS Lambda C. AWS Kinesis D. AWS SNS Answer: C Explanation: AWS Kinesis is an event stream service. Streams can act as buffers and transport across systems for in-order programmatic events, making it ideal for replicating API calls across systems. A typical Amazon Kinesis Streams application reads data from an Amazon Kinesis stream as data records. These applications can use the Amazon Kinesis Client Library, and they can run on Amazon EC2 instances. The processed records can be sent to dashboards, used to generate alerts, dynamically change pricing and advertising strategies, or send data to a variety of other AWS services. For information about Streams features and pricing, see Amazon Kinesis Streams.

<http://docs.aws.amazon.com/kinesis/latest/dev/introduction.html> QUESTION 136 You are building a Ruby on Rails application for internal, non-production use which uses MySQL as a database. You want developers without very much AWS experience to be able to deploy new code with a single command line push. You also want to set this up as simply as possible. Which tool is ideal for this setup? A. AWS CloudFormation B. AWS OpsWorks C. AWS ELB + EC2 with CLI Push D. AWS Elastic Beanstalk Answer: D Explanation: Elastic Beanstalk's primary mode of operation exactly supports this use case out of the box. It is simpler than all the other options for this question. With Elastic Beanstalk, you can quickly deploy and manage applications in the AWS cloud without worrying about the infrastructure that runs those applications. AWS Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.

http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_Ruby_rails.html QUESTION 137 What is the scope of AWS IAM? A. Global B. Availability Zone C. Region D. Placement Group Answer: A Explanation: IAM resources are all global; there is not regional constraint. <https://aws.amazon.com/iam/faqs/> QUESTION 138 You are building a mobile app for consumers to post cat pictures online. You will be storing the images in AWS S3. You want to run the system very cheaply and simply. Which one of these options allows you to build a photo sharing application without needing to worry about scaling expensive uploads

processes, authentication/authorization and so forth? A. Build the application out using AWS Cognito and web identity federation to allow users to log in using Facebook or Google Accounts. Once they are logged in, the secret token passed to that user is used to directly access resources on AWS, like AWS S3. B. Use JWT or SAML compliant systems to build authorization policies. Users log in with a username and password, and are given a token they can use indefinitely to make calls against the photo infrastructure. C. Use AWS API Gateway with a constantly rotating API Key to allow access from the client-side. Construct a custom build of the SDK and include S3 access in it. D. Create an AWS OAuth Service Domain and grant public sign up and access to the domain. During setup, add at least one major social media site as a trusted Identity Provider for users. Answer: A Explanation: The short answer is that Amazon Cognito is a superset of the functionality provided by web identity federation. It supports the same providers, and you configure your app and authenticate with those providers in the same way. But Amazon Cognito includes a variety of additional features. For example, it enables your users to start using the app as a guest user and later sign in using one of the supported identity providers.

processes, authentication/authorization and so forth? A. Build the application out using AWS Cognito and web identity federation to allow users to log in using Facebook or Google Accounts. Once they are logged in, the secret token passed to that user is used to directly access resources on AWS, like AWS S3. B. Use JWT or SAML compliant systems to build authorization policies. Users log in with a username and password, and are given a token they can use indefinitely to make calls against the photo infrastructure. C. Use AWS API Gateway with a constantly rotating API Key to allow access from the client-side. Construct a custom build of the SDK and include S3 access in it. D. Create an AWS OAuth Service Domain and grant public sign up and access to the domain. During setup, add at least one major social media site as a trusted Identity Provider for users. Answer: A Explanation: The short answer is that Amazon Cognito is a superset of the functionality provided by web identity federation. It supports the same providers, and you configure your app and authenticate with those providers in the same way. But Amazon Cognito includes a variety of additional features. For example, it enables your users to start using the app as a guest user and later sign in using one of the supported identity providers.

<https://blogs.aws.amazon.com/security/post/Tx3SYCORF5EKRC0/How-Does-Amazon-Cognito-Relate-to-Existing-Web-Identity-Federatio> QUESTION 139 Your CTO has asked you to make sure that you know what all users of your AWS account are doing to change resources at all times. She wants a report of who is doing what over time, reported to her once per week, for as broad a resource type group as possible. How should you do this? A. Create a global AWS CloudTrail Trail. Configure a script to aggregate the log data delivered to S3 once per week and deliver this to the CTO. B. Use CloudWatch Events Rules with an SNS topic subscribed to all AWS API calls. Subscribe the CTO to an email type delivery on this SNS Topic. C. Use AWS IAM credential reports to deliver a CSV of all uses of IAM User Tokens over time to the CTO. D. Use AWS Config with an SNS subscription on a Lambda, and insert these changes over time into a DynamoDB table. Generate reports based on the contents of this table. Answer: A Explanation: This is the ideal use case for AWS CloudTrail. CloudTrail provides visibility into user activity by recording API calls made on your account. CloudTrail records important information about each API call, including the name of the API, the identity of the caller, the time of the API call, the request parameters, and the response elements returned by the AWS service. This information helps you to track changes made to your AWS resources and to troubleshoot operational issues. CloudTrail makes it easier to ensure compliance with internal policies and regulatory standards. <https://aws.amazon.com/cloudtrail/faqs/> QUESTION 140 What is the order of most-to-least rapidly-scaling (fastest to scale first)? a) EC2 + ELB + Auto Scaling b) Lambda c) RDS A. B, A, CB. C, B, AC. C, A, BD. A, C, B Answer: A Explanation: Lambda is designed to scale instantly. EC2 + ELB + Auto Scaling require single-digit minutes to scale out. RDS will take at least 15 minutes, and will apply OS patches or any other updates when applied. <https://aws.amazon.com/lambda/faqs/> More free Lead2pass **AWS-DevOps-Engineer-Professional** exam new questions on Google Drive: <https://drive.google.com/open?id=0B3Syig5i8gpDbVZ1cTB3QnNPQlk> Lead2pass is no doubt your best choice. Using the Amazon AWS-DevOps-Engineer-Professional exam dumps can let you improve the efficiency of your studying so that it can help you save much more time. 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]