

[2017 PDF&VCE New Lead2pass Amazon AWS-DevOps-Engineer-Professional New Questions Free Download (101-120)]

Lead2pass 2017 September New Amazon [AWS-DevOps-Engineer-Professional Exam Dumps!](#) 100% Free Download! 100% Pass [Guaranteed!](#) In recent years, many people choose to take Amazon AWS-DevOps-Engineer-Professional certification exam which can make you get the Amazon certificate and that is the passport to get a better job and get promotions. How to prepare for Amazon AWS-DevOps-Engineer-Professional exam and get the certificate? Please refer to Amazon AWS-DevOps-Engineer-Professional exam questions and answers on Lead2pass. Following questions and answers are all new published by **Amazon** Official Exam Center: <https://www.lead2pass.com/aws-devops-engineer-professional.html> QUESTION 101What is the maximum supported single-volume throughput on EBS? A. 320MiB/sB. 160MiB/sC. 40MiB/sD. 640MiB/sAnswer: AExplanation:The ceiling throughput for PIOPS on EBS is 320MiB/s.<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html> QUESTION 102For AWS Auto Scaling, what is the first transition state a new instance enters after leaving steady state when scaling out due to increased load? A. EnteringStandbyB. PendingC. Terminating:WaitD. Detaching Answer: BExplanation: When a scale out event occurs, the Auto Scaling group launches the required number of EC2 instances, using its assigned launch configuration. These instances start in the Pending state. If you add a lifecycle hook to your Auto Scaling group, you can perform a custom action here. For more information, see Lifecycle Hooks.

<http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AutoScalingGroupLifecycle.html> QUESTION 103When a user is detaching an EBS volume from a running instance and attaching it to a new instance, which of the below mentioned options should be followed to avoid file system damage? A. Unmount the volume firstB. Stop all the I/O of the volume before processingC. Take a snapshot of the volume before detachingD. Force Detach the volume to ensure that all the data stays intact Answer: AExplanation:When a user is trying to detach an EBS volume, the user can either terminate the instance or explicitly remove the volume. It is a recommended practice to unmount the volume first to avoid any file system damage. QUESTION 104A user is creating a new EBS volume from an existing snapshot. The snapshot size shows 10 GB. Can the user create a volume of 30 GB from that snapshot? A. Provided the original volume has set the change size attribute to trueB. YesC. Provided the snapshot has the modify size attribute set as trueD. No Answer: BExplanation:A user can always create a new EBS volume of a higher size than the original snapshot size. The user cannot create a volume of a lower size. When the new volume is created the size in the instance will be shown as the original size. The user needs to change the size of the device with resize2fs or other OS specific commands. QUESTION 105How long are the messages kept on an SQS queue by default? A. If a message is not read, it is never deletedB. 2 weeksC. 1 dayD. 4 days Answer: DExplanation:The SQS message retention period is configurable and can be set anywhere from 1 minute to 2 weeks. The default is 4 days and once the message retention limit is reached your messages will be automatically deleted. The option for longer message retention provides greater flexibility to allow for longer intervals between message production and consumption. QUESTION 106A user has attached an EBS volume to a running Linux instance as a "/dev/sdf" device. The user is unable to see the attached device when he runs the command "df -h". What is the possible reason for this? A. The volume is not in the same AZ of the instanceB. The volume is not formattedC. The volume is not attached as a root deviceD.

The volume is not mounted Answer: DExplanation:When a user creates an EBS volume and attaches it as a device, it is required to mount the device. If the device/volume is not mounted it will not be available in the listing. QUESTION 107When using Amazon SQS how much data can you store in a message? A. 8 KBB. 2 KBC. 16 KBD. 4 KB Answer: AExplanation:With Amazon SQS version 2008-01-01, the maximum message size for both SOAP and Query requests is 8KB. If you need to send messages to the queue that are larger than 8 KB, AWS recommends that you split the information into separate messages. Alternatively, you could use Amazon S3 or Amazon SimpleDB to hold the information and include the pointer to that information in the Amazon SQS message. If you send a message that is larger than 8KB to the queue, you will receive a MessageTooLong error with HTTP code 400. QUESTION 18What is the maximum time messages can be stored in SQS? A. 14 daysB. one monthC. 4 daysD. 7 days Answer: AExplanation:A message can be stored in the Simple Queue Service (SQS) from 1 minute up to a maximum of 14 days. QUESTION 109In DynamoDB, a secondary index is a data structure that contains a subset of attributes from a table, along with an alternate key to support _____ operations. A. None of the aboveB. BothC. QueryD. Scan Answer: CExplanation:In DynamoDB, a secondary index is a data structure that contains a subset of attributes from a table, along with an alternate key to support Query operations. QUESTION 110A user has created a new EBS volume from an existing snapshot. The user mounts the volume on the instance to which it is attached. Which of the below mentioned options is a required step before the user can mount the volume? A. Run a cyclic check on the device for data consistencyB. Create the file system of the volumeC. Resize the volume as per the original snapshot sizeD. No step is required. The user can directly mount the device Answer: DExplanation:

When a user is trying to mount a blank EBS volume, it is required that the user first creates a file system within the volume. If the volume is created from an existing snapshot then the user needs not to create a file system on the volume as it will wipe out the existing data. **QUESTION 111** You need your CI to build AMIs with code pre-installed on the images on every new code push. You need to do this as cheaply as possible. How do you do this? A. Bid on spot instances just above the asking price as soon as new commits come in, perform all instance configuration and setup, then create an AMI based on the spot instance. B. Have the CI launch a new on-demand EC2 instance when new commits come in, perform all instance configuration and setup, then create an AMI based on the on-demand instance. C. Purchase a Light Utilization Reserved Instance to save money on the continuous integration machine. Use these credits whenever you create AMIs on instances. D. When the CI instance receives commits, attach a new EBS volume to the CI machine. Perform all setup on this EBS volume so you don't need a new EC2 instance to create the AMI. **Answer: A** Explanation: Spot instances are the cheapest option, and you can use minimum run duration if your AMI takes more than a few minutes to create. Spot instances are also available to run for a predefined duration - in hourly increments up to six hours in length - at a significant discount (30-45%) compared to On-Demand pricing plus an additional 5% during off-peak times for a total of up to 50% savings. <https://aws.amazon.com/ec2/spot/pricing/> **QUESTION 112** When thinking of DynamoDB, what are true of Global Secondary Key properties? A. The partition key and sort key can be different from the table. B. Only the partition key can be different from the table. C. Either the partition key or the sort key can be different from the table, but not both. D. Only the sort key can be different from the table. **Answer: A** Explanation: Global secondary index -- an index with a partition key and a sort key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions.

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html> **QUESTION 113** You need to process long-running jobs once and only once. How might you do this? A. Use an SNS queue and set the visibility timeout to long enough for jobs to process. B. Use an SQS queue and set the reprocessing timeout to long enough for jobs to process. C. Use an SQS queue and set the visibility timeout to long enough for jobs to process. D. Use an SNS queue and set the reprocessing timeout to long enough for jobs to process. **Answer: C** Explanation: The message timeout defines how long after a successful receive request SQS waits before allowing jobs to be seen by other components, and proper configuration prevents duplicate processing.

<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/MessageLifecycle.html> **QUESTION 114** You are getting a lot of empty receive requests when using Amazon SQS. This is making a lot of unnecessary network load on your instances. What can you do to reduce this load? A. Subscribe your queue to an SNS topic instead. B. Use as long of a poll as possible, instead of short polls. C. Alter your visibility timeout to be shorter. D. Use `sqsd` on your EC2 instances. **Answer: B** Explanation: One benefit of long polling with Amazon SQS is the reduction of the number of empty responses, when there are no messages available to return, in reply to a `ReceiveMessage` request sent to an Amazon SQS queue. Long polling allows the Amazon SQS service to wait until a message is available in the queue before sending a response.

<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-long-polling.html> **QUESTION 115** You need to know when you spend \$1000 or more on AWS. What's the easy way for you to see that notification? A. AWS CloudWatch Events tied to API calls, when certain thresholds are exceeded, publish to SNS. B. Scrape the billing page periodically and pump into Kinesis. C. AWS CloudWatch Metrics + Billing Alarm + Lambda event subscription. When a threshold is exceeded, email the manager. D. Scrape the billing page periodically and publish to SNS. **Answer: C** Explanation: Even if you're careful to stay within the free tier, it's a good idea to create a billing alarm to notify you if you exceed the limits of the free tier. Billing alarms can help to protect you against unknowingly accruing charges if you inadvertently use a service outside of the free tier or if traffic exceeds your expectations. <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/free-tier-alarms.html> **QUESTION 116** You need to grant a vendor access to your AWS account. They need to be able to read protected messages in a private S3 bucket at their leisure. They also use AWS. What is the best way to accomplish this? A. Create an IAM User with API Access Keys. Grant the User permissions to access the bucket. Give the vendor the AWS Access Key ID and AWS Secret Access Key for the User. B. Create an EC2 Instance Profile on your account. Grant the associated IAM role full access to the bucket. Start an EC2 instance with this Profile and give SSH access to the instance to the vendor. C. Create a cross-account IAM Role with permission to access the bucket, and grant permission to use the Role to the vendor AWS account. D. Generate a signed S3 PUT URL and a signed S3 PUT URL, both with wildcard values and 2 year durations. Pass the URLs to the vendor. **Answer: C** Explanation: When third parties require access to your organization's AWS resources, you can use roles to delegate access to them. For example, a third party might provide a service for managing your AWS resources. With IAM roles, you can grant these third parties access to your AWS resources without sharing your AWS security credentials. Instead, the third party can access your AWS resources by assuming a role that you create in your AWS account.

http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_common-scenarios_third-party.html QUESTION 117 Your serverless architecture using AWS API Gateway, AWS Lambda, and AWS DynamoDB experienced a large increase in traffic to a sustained 400 requests per second, and dramatically increased in failure rates. Your requests, during normal operation, last 500 milliseconds on average. Your DynamoDB table did not exceed 50% of provisioned throughput, and Table primary keys are designed correctly. What is the most likely issue? A. Your API Gateway deployment is throttling your requests. B. Your AWS API Gateway Deployment is bottlenecking on request (de)serialization. C. You did not request a limit increase on concurrent Lambda function executions. D. You used Consistent Read requests on DynamoDB and are experiencing semaphore lock. Answer: C Explanation: AWS API Gateway by default throttles at 500 requests per second steady-state, and 1000 requests per second at spike. Lambda, by default, throttles at 100 concurrent requests for safety. At 500 milliseconds (half of a second) per request, you can expect to support 200 requests per second at 100 concurrency. This is less than the 400 requests per second your system now requires. Make a limit increase request via the AWS Support Console. AWS Lambda: Concurrent requests safety throttle per account -> 100

http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_lambda QUESTION 118 Why are more frequent snapshots or EBS Volumes faster? A. Blocks in EBS Volumes are allocated lazily, since while logically separated from other EBS Volumes, Volumes often share the same physical hardware. Snapshotting the first time forces full block range allocation, so the second snapshot doesn't need to perform the allocation phase and is faster. B. The snapshots are incremental so that only the blocks on the device that have changed after your last snapshot are saved in the new snapshot. C. AWS provisions more disk throughput for burst capacity during snapshots if the drive has been pre-warmed by snapshotting and reading all blocks. D. The drive is pre-warmed, so block access is more rapid for volumes when every block on the device has already been read at least one time. Answer: B Explanation: After writing data to an EBS volume, you can periodically create a snapshot of the volume to use as a baseline for new volumes or for data backup. If you make periodic snapshots of a volume, the snapshots are incremental so that only the blocks on the device that have changed after your last snapshot are saved in the new snapshot. Even though snapshots are saved incrementally, the snapshot deletion process is designed so that you need to retain only the most recent snapshot in order to restore the volume. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html> QUESTION 119 For AWS CloudFormation, which stack state refuses UpdateStack calls? A. <code>UPDATE_ROLLBACK_FAILED</code> B. <code>UPDATE_ROLLBACK_COMPLETE</code> C. <code>UPDATE_COMPLETE</code> D. <code>CREATE_COMPLETE</code> Answer: A Explanation: When a stack is in the UPDATE_ROLLBACK_FAILED state, you can continue rolling it back to return it to a working state (to UPDATE_ROLLBACK_COMPLETE). You cannot update a stack that is in the UPDATE_ROLLBACK_FAILED state. However, if you can continue to roll it back, you can return the stack to its original settings and try to update it again.

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-continueu> pdater rollback.html QUESTION 120 You need to migrate 10 million records in one hour into DynamoDB. All records are 1.5KB in size. The data is evenly distributed across the partition key. How many write capacity units should you provision during this batch load? A. 6667 B. 4166 C. 5556 D. 2778 Answer: C Explanation: You need 2 units to make a 1.5KB write, since you round up. You need 20 million total units to perform this load. You have 3600 seconds to do so. Divide and round up for 5556.

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ProvisionedThroughput.html> More free Lead2pass AWS-DevOps-Engineer-Professional exam new questions on Google Drive:

<https://drive.google.com/open?id=0B3Syig5i8gpDbVZ1cTB3QnNPQlk> Lead2pass is a good website that provides all candidates with the latest IT certification exam materials. Lead2pass will provide you with the exam questions and verified answers that reflect the actual exam. The Amazon AWS-DevOps-Engineer-Professional exam dumps are developed by experienced IT professionals. 99.9% of hit rate. Guarantee you success in your AWS-DevOps-Engineer-Professional exam with our exam materials. 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]