

[2017 New Lead2pass Latest AWS Certified Solutions Architect - Associate PDF Guarantee 100% Pass AWS Certified Solutions Architect–Associate (476-500)]

2017 August Amazon Official New Released AWS Certified Solutions Architect ? Associate Dumps in Lead2pass.com! 100% Free Download! 100% Pass Guaranteed! Are you worrying about the AWS Certified Solutions Architect - Associate exam? With the complete collection of AWS Certified Solutions Architect - Associate exam questions and answers, Lead2pass has assembled to take you through your AWS Certified Solutions Architect - Associate exam preparation. Each Q & A set will test your existing knowledge of AWS Certified Solutions Architect - Associate fundamentals, and offer you the latest training products that guarantee you passing AWS Certified Solutions Architect - Associate exam easily. Following questions and answers are all new published by Amazon Official Exam Center: <https://www.lead2pass.com/aws-certified-solutions-architect-associate.html> QUESTION 476 In Amazon EC2, while sharing an Amazon EBS snapshot, can the snapshots with AWS Marketplace product codes be public? A. Yes, but only for US-based providers. B. Yes, they can be public. C. No, they cannot be made public. D. Yes, they are automatically made public by the system. Answer: C Explanation: Snapshots with AWS Marketplace product codes can't be made public. Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/ebs-modifying-snapshot-permissions.html> QUESTION 477 An organization has created an application which is hosted on the AWS EC2 instance. The application stores images to S3 when the end user uploads to it. The organization does not want to store the AWS secure credentials required to access the S3 inside the instance. Which of the below mentioned options is a possible solution to avoid any security threat? A. Use the IAM based single sign between the AWS resources and the organization application. B. Use the IAM role and assign it to the instance. C. Since the application is hosted on EC2, it does not need credentials to access S3. D. Use the X.509 certificates instead of the access and the secret access keys. Answer: B Explanation: The AWS IAM role uses temporary security credentials to access AWS services. Once the role is assigned to an instance, it will not need any security credentials to be stored on the instance. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html> QUESTION 478 Can resource record sets in a hosted zone have a different domain suffix (for example, www.blog.acme.com and www.acme.ca)? A. Yes, it can have for a maximum of three different TLDs. B. Yes. C. Yes, it can have depending on the TLD. D. No Answer: D Explanation: The resource record sets contained in a hosted zone must share the same suffix. For example, the example.com hosted zone can contain resource record sets for www.example.com and www.aws.example.com subdomains, but it cannot contain resource record sets for a www.example.ca subdomain. Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/AboutHostedZones.html> QUESTION 479 You are running PostgreSQL on Amazon RDS and it seems to be all running smoothly deployed in one availability zone. A database administrator asks you if DB instances running PostgreSQL support Multi-AZ deployments. What would be a correct response to this question? A. Yes. B. Yes but only for small db instances. C. No. D. Yes but you need to request the service from AWS. Answer: A Explanation: Amazon RDS supports DB instances running several versions of PostgreSQL. Currently we support PostgreSQL versions 9.3.1, 9.3.2, and 9.3.3. You can create DB instances and DB snapshots, point-in-time restores and backups. DB instances running PostgreSQL support Multi-AZ deployments, Provisioned IOPS, and can be created inside a VPC. You can also use SSL to connect to a DB instance running PostgreSQL. You can use any standard SQL client application to run commands for the instance from your client computer. Such applications include pgAdmin, a popular Open Source administration and development tool for PostgreSQL, or psql, a command line utility that is part of a PostgreSQL installation. In order to deliver a managed service experience, Amazon RDS does not provide host access to DB instances, and it restricts access to certain system procedures and tables that require advanced privileges. Amazon RDS supports access to databases on a DB instance using any standard SQL client application. Amazon RDS does not allow direct host access to a DB instance via Telnet or Secure Shell (SSH). Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_PostgreSQL.html QUESTION 480 A user has launched 10 EC2 instances inside a placement group. Which of the below mentioned statements is true with respect to the placement group? A. All instances must be in the same AZ. B. All instances can be across multiple regions. C. The placement group cannot have more than 5 instances. D. All instances must be in the same region. Answer: A Explanation: A placement group is a logical grouping of EC2 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. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html> QUESTION 481 Which of the following AWS CLI commands is syntactically incorrect? 1. \$ aws ec2 describe-instances 2. \$ aws ec2 start-instances --instance-ids i-1348636c3. \$ aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError -message "Script Failure" 4. \$ aws

sqs receive-message --queue-url <https://queue.amazonaws.com/546419318123/Test> A. 3B. 4C. 2D. 1 Answer: A Explanation: The following CLI command is missing a hyphen before "-message". aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError -message "Script Failure" It has been added below in redaws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError ---message "Script Failure" Reference: <http://aws.amazon.com/cli/>

QUESTION 482 An organization has developed a mobile application which allows end users to capture a photo on their mobile device, and store it inside an application. The application internally uploads the data to AWS S3. The organization wants each user to be able to directly upload data to S3 using their Google ID. How will the mobile app allow this? A. Use the AWS Web identity federation for mobile applications, and use it to generate temporary security credentials for each user. B. It is not possible to connect to AWS S3 with a Google ID. C. Create an IAM user every time a user registers with their Google ID and use IAM to upload files to S3. D. Create a bucket policy with a condition which allows everyone to upload if the login ID has a Google part to it. Answer: A Explanation: For Amazon Web Services, the Web identity federation allows you to create cloud-backed mobile apps that use public identity providers, such as login with Facebook, Google, or Amazon. It will create temporary security credentials for each user, which will be authenticated by the AWS services, such as S3. Reference: <http://docs.aws.amazon.com/STS/latest/UsingSTS/CreatingWIF.html>

QUESTION 483 You are architecting an auto-scalable batch processing system using video processing pipelines and Amazon Simple Queue Service (Amazon SQS) for a customer. You are unsure of the limitations of SQS and need to find out. What do you think is a correct statement about the limitations of Amazon SQS? A. It supports an unlimited number of queues but a limited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 weeks. B. It supports an unlimited number of queues and unlimited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 days. C. It supports an unlimited number of queues but a limited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 days. D. It supports an unlimited number of queues and unlimited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 weeks. Answer: B Explanation: Amazon Simple Queue Service (Amazon SQS) is a messaging queue service that handles message or workflows between other components in a system. Amazon SQS supports an unlimited number of queues and unlimited number of messages per queue for each user. Please be aware that Amazon SQS automatically deletes messages that have been in the queue for more than 4 days. Reference: <http://aws.amazon.com/documentation/sqs/>

QUESTION 484 An online gaming site asked you if you can deploy a database that is a fast, highly scalable NoSQL database service in AWS for a new site that he wants to build. Which database should you recommend? A. Amazon DynamoDB. B. Amazon RDS. C. Amazon Redshift. D. Amazon SimpleDB Answer: A Explanation: Amazon DynamoDB is ideal for database applications that require very low latency and predictable performance at any scale but don't need complex querying capabilities like joins or transactions. Amazon DynamoDB is a fully-managed NoSQL database service that offers high performance, predictable throughput and low cost. It is easy to set up, operate, and scale. With Amazon DynamoDB, you can start small, specify the throughput and storage you need, and easily scale your capacity requirements on the fly. Amazon DynamoDB automatically partitions data over a number of servers to meet your request capacity. In addition, DynamoDB automatically replicates your data synchronously across multiple Availability Zones within an AWS Region to ensure high-availability and data durability. Reference: https://aws.amazon.com/running_databases/#dynamodb_anchor

QUESTION 485 You have been doing a lot of testing of your VPC Network by deliberately failing EC2 instances to test whether instances are failing over properly. Your customer who will be paying the AWS bill for all this asks you if he being charged for all these instances. You try to explain to him how the billing works on EC2 instances to the best of your knowledge. What would be an appropriate response to give to the customer in regards to this? A. Billing commences when Amazon EC2 AMI instance is completely up and billing ends as soon as the instance starts to shutdown. B. Billing only commences only after 1 hour of uptime and billing ends when the instance terminates. C. Billing commences when Amazon EC2 initiates the boot sequence of an AMI instance and billing ends when the instance shuts down. D. Billing commences when Amazon EC2 initiates the boot sequence of an AMI instance and billing ends as soon as the instance starts to shutdown. Answer: C Explanation: Billing commences when Amazon EC2 initiates the boot sequence of an AMI instance. Billing ends when the instance shuts down, which could occur through a web services command, by running "shutdown -h", or through instance failure. Reference: <http://aws.amazon.com/ec2/faqs/#Billing>

QUESTION 486 You log in to IAM on your AWS console and notice the following message. "Delete your root access keys." Why do you think IAM is requesting this? A. Because the root access keys will expire as soon as you log out. B. Because the root access keys expire after 1 week. C. Because the root access keys are the same for all users. D. Because they provide unrestricted access to your AWS resources. Answer: D Explanation: In AWS an access key is required in order to sign requests that you make using the command-line interface (CLI), using the AWS SDKs, or using direct API calls. Anyone who has the access key for your root account has

unrestricted access to all the resources in your account, including billing information. One of the best ways to protect your account is to not have an access key for your root account. We recommend that unless you must have a root access key (this is very rare), that you do not generate one. Instead, AWS best practice is to create one or more AWS Identity and Access Management (IAM) users, give them the necessary permissions, and use IAM users for everyday interaction with AWS. Reference:

<http://docs.aws.amazon.com/general/latest/gr/aws-access-keys-best-practices.html#root-password> QUESTION 487 Once again your customers are concerned about the security of their sensitive data and with their latest enquiry ask about what happens to old storage devices on AWS. What would be the best answer to this question? A. AWS reformats the disks and uses them again. B. AWS uses the techniques detailed in DoD 5220.22-M to destroy data as part of the decommissioning process. C. AWS uses their own proprietary software to destroy data as part of the decommissioning process. D. AWS uses a 3rd party security organization to destroy data as part of the decommissioning process. Answer: B Explanation: When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual ") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.

Reference: <http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf> QUESTION 488 Your company has been storing a lot of data in Amazon Glacier and has asked for an inventory of what is in there exactly. So you have decided that you need to download a vault inventory. Which of the following statements is incorrect in relation to Vault Operations in Amazon Glacier? A. You can use Amazon Simple Notification Service (Amazon SNS) notifications to notify you when the job completes. B. A vault inventory refers to the list of archives in a vault. C. You can use Amazon Simple Queue Service (Amazon SQS) notifications to notify you when the job completes. D. Downloading a vault inventory is an asynchronous operation. Answer: C Explanation: Amazon Glacier supports various vault operations. A vault inventory refers to the list of archives in a vault. For each archive in the list, the inventory provides archive information such as archive ID, creation date, and size. Amazon Glacier updates the vault inventory approximately once a day, starting on the day the first archive is uploaded to the vault. A vault inventory must exist for you to be able to download it. Downloading a vault inventory is an asynchronous operation. You must first initiate a job to download the inventory. After receiving the job request, Amazon Glacier prepares your inventory for download. After the job completes, you can download the inventory data. Given the asynchronous nature of the job, you can use Amazon Simple Notification Service (Amazon SNS) notifications to notify you when the job completes. You can specify an Amazon SNS topic for each individual job request or configure your vault to send a notification when specific vault events occur. Amazon Glacier prepares an inventory for each vault periodically, every 24 hours. If there have been no archive additions or deletions to the vault since the last inventory, the inventory date is not updated. When you initiate a job for a vault inventory, Amazon Glacier returns the last inventory it generated, which is a point-in-time snapshot and not real-time data. You might not find it useful to retrieve vault inventory for each archive upload. However, suppose you maintain a database on the client-side associating metadata about the archives you upload to Amazon Glacier. Then, you might find the vault inventory useful to reconcile information in your database with the actual vault inventory. Reference: <http://docs.aws.amazon.com/amazonglacier/latest/dev/working-with-vaults.html> QUESTION 489 A

customer enquires about whether all his data is secure on AWS and is especially concerned about Elastic Map Reduce (EMR) so you need to inform him of some of the security features in place for AWS. Which of the below statements would be an incorrect response to your customers enquiry? A. Amazon EMR customers can choose to send data to Amazon S3 using the HTTPS protocol for secure transmission. B. Amazon S3 provides authentication mechanisms to ensure that stored data is secured against unauthorized access. C. Every packet sent in the AWS network uses Internet Protocol Security (IPsec). D. Customers may encrypt the input data before they upload it to Amazon S3. Answer: C Explanation: Amazon S3 provides authentication mechanisms to ensure that stored data is secured against unauthorized access. Unless the customer who is uploading the data specifies otherwise, only that customer can access the data. Amazon EMR customers can also choose to send data to Amazon S3 using the HTTPS protocol for secure transmission. In addition, Amazon EMR always uses HTTPS to send data between Amazon S3 and Amazon EC2. For added security, customers may encrypt the input data before they upload it to Amazon S3 (using any common data compression tool); they then need to add a decryption step to the beginning of their cluster when Amazon EMR fetches the data from Amazon S3. Reference: <https://aws.amazon.com/elasticmapreduce/faqs/> QUESTION 490 You are in the process of building an online gaming site for a client and one of the requirements is that it must be able to process vast amounts of data easily. Which AWS Service would be very helpful in processing all this data? A. Amazon S3 B. AWS Data Pipeline C. AWS Direct Connect D. Amazon EMR Answer: D

Explanation: Managing and analyzing high data volumes produced by online games platforms can be difficult. The back-end infrastructures of online games can be challenging to maintain and operate. Peak usage periods, multiple players, and high volumes

of write operations are some of the most common problems that operations teams face. Amazon Elastic MapReduce (Amazon EMR) is a service that processes vast amounts of data easily. Input data can be retrieved from web server logs stored on Amazon S3 or from player data stored in Amazon DynamoDB tables to run analytics on player behavior, usage patterns, etc. Those results can be stored again on Amazon S3, or inserted in a relational database for further analysis with classic business intelligence tools.

Reference: http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_games_10.pdf QUESTION 491 You need to change some settings on Amazon Relational Database Service but you do not want the database to reboot immediately which you know might happen depending on the setting that you change. Which of the following will cause an immediate DB instance reboot to occur? A. You change storage type from standard to PIOPS, and Apply Immediately is set to true. B. You change the DB instance class, and Apply Immediately is set to false. C. You change a static parameter in a DB parameter group. D. You change the backup retention period for a DB instance from 0 to a nonzero value or from a nonzero value to 0, and Apply Immediately is set to false. Answer: A Explanation: A DB instance outage can occur when a DB instance is rebooted, when the DB instance is put into a state that prevents access to it, and when the database is restarted. A reboot can occur when you manually reboot your DB instance or when you change a DB instance setting that requires a reboot before it can take effect. A DB instance reboot occurs immediately when one of the following occurs: You change the backup retention period for a DB instance from 0 to a nonzero value or from a nonzero value to 0 and set Apply Immediately to true. You change the DB instance class, and Apply Immediately is set to true. You change storage type from standard to PIOPS, and Apply Immediately is set to true. A DB instance reboot occurs during the maintenance window when one of the following occurs: You change the backup retention period for a DB instance from 0 to a nonzero value or from a nonzero value to 0, and Apply Immediately is set to false. You change the DB instance class, and Apply Immediately is set to false. Reference:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Troubleshooting.html#CHAP_Troubleshooting.Security

QUESTION 492 What does the following policy for Amazon EC2 do? {"Statement":[{"Effect":"Allow","Action":"ec2:Describe*","Resource":"*"}]} A. Allow users to use actions that start with "Describe" over all the EC2 resources. B. Share an AMI with a partner. C. Share an AMI within the account. D. Allow a group to only be able to describe, run, stop, start, and terminate instances. Answer: A Explanation: You can use IAM policies to control the actions that your users can perform against your EC2 resources. For instance, a policy with the following statement will allow users to perform actions whose name start with "Describe" against all your EC2 resources. {"Statement":[{"Effect":"Allow","Action":"ec2:Describe*","Resource":"*"}]} Reference:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/UsingIAM.html> QUESTION 493 You are setting up a very complex financial services grid and so far it has 5 Elastic IP (EIP) addresses. You go to assign another EIP address, but all accounts are limited to 5 Elastic IP addresses per region by default, so you aren't able to. What is the reason for this? A. For security reasons. B. Hardware restrictions. C. Public (IPv4) internet addresses are a scarce resource. D. There are only 5 network interfaces per instance. Answer: C Explanation: Public (IPv4) internet addresses are a scarce resource. There is only a limited amount of public IP space available, and Amazon EC2 is committed to helping use that space efficiently. By default, all accounts are limited to 5 Elastic IP addresses per region. If you need more than 5 Elastic IP addresses, AWS asks that you apply for your limit to be raised. They will ask you to think through your use case and help them understand your need for additional addresses. Reference:

http://aws.amazon.com/ec2/faqs/#How_many_instances_can_I_run_in_Amazon_EC2 QUESTION 494 Amazon RDS provides high availability and failover support for DB instances using _____. A. customized deployments B. Appstream customizations C. log events D. Multi-AZ deployments Answer: D Explanation: Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. Multi-AZ deployments for Oracle, PostgreSQL, MySQL, and MariaDB DB instances use Amazon technology, while SQL Server DB instances use SQL Server Mirroring. Reference:

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html> QUESTION 495 A major customer has asked you to set up his AWS infrastructure so that it will be easy to recover in the case of a disaster of some sort. Which of the following is important when thinking about being able to quickly launch resources in AWS to ensure business continuity in case of a disaster? A. Create and maintain AMIs of key servers where fast recovery is required. B. Regularly run your servers, test them, and apply any software updates and configuration changes. C. All items listed here are important when thinking about disaster recovery. D. Ensure that you have all supporting custom software packages available in AWS. Answer: C Explanation: In the event of a disaster to your AWS infrastructure you should be able to quickly launch resources in Amazon Web Services (AWS) to ensure business continuity. The following are some key steps you should have in place for preparation: 1. Set up Amazon EC2 instances to replicate or mirror data. 2. Ensure that you have all supporting custom software packages available in AWS. 3. Create and maintain AMIs of key servers where fast recovery is required. 4. Regularly run these servers, test them, and apply any software updates and configuration changes. 5. Consider automating the provisioning of AWS resources. Reference:

http://d36cz9buwru1tt.cloudfront.net/AWS_Disaster_Recovery.pdf QUESTION 496 What does Amazon DynamoDB provide? A. A predictable and scalable MySQL database B. A fast and reliable PL/SQL database cluster C. A standalone Cassandra database, managed by Amazon Web Services D. A fast, highly scalable managed NoSQL database service Answer: D Explanation: Amazon DynamoDB is a managed NoSQL database service offered by Amazon. It automatically manages tasks like scalability for you while it provides high availability and durability for your data, allowing you to concentrate in other aspects of your application. Reference: https://aws.amazon.com/running_databases/ QUESTION 497 You want to use AWS Import/Export to send data from your S3 bucket to several of your branch offices. What should you do if you want to send 10 storage units to AWS? A. Make sure your disks are encrypted prior to shipping. B. Make sure you format your disks prior to shipping. C. Make sure your disks are 1TB or more. D. Make sure you submit a separate job request for each device. Answer: D Explanation: When using Amazon Import/Export, a separate job request needs to be submitted for each physical device even if they belong to the same import or export job. Reference: <http://docs.aws.amazon.com/AWSImportExport/latest/DG/Concepts.html> QUESTION 498 What would be the best way to retrieve the public IP address of your EC2 instance using the CLI? A. Using tags B. Using traceroute C. Using ipconfig D. Using instance metadata Answer: D Explanation: To determine your instance's public IP address from within the instance, you can use instance metadata. Use the following command to access the public IP address: For Linux use, \$ curl <http://169.254.169.254/latest/meta-data/public-ipv4>, and for Windows use, \$ wget <http://169.254.169.254/latest/meta-data/public-ipv4>. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-instance-addressing.html> QUESTION 499 You need to measure the performance of your EBS volumes as they seem to be under performing. You have come up with a measurement of 1,024 KB I/O but your colleague tells you that EBS volume performance is measured in IOPS. How many IOPS is equal to 1,024 KB I/O? A. 16 B. 256 C. 8 D. 4 Answer: D Explanation: Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration. IOPS are input/output operations per second. Amazon EBS measures each I/O operation per second (that is 256 KB or smaller) as one IOPS. I/O operations that are larger than 256 KB are counted in 256 KB capacity units. For example, a 1,024 KB I/O operation would count as 4 IOPS. When you provision a 4,000 IOPS volume and attach it to an EBS-optimized instance that can provide the necessary bandwidth, you can transfer up to 4,000 chunks of data per second (provided that the I/O does not exceed the 128 MB/s per volume throughput limit of General Purpose (SSD) and Provisioned IOPS (SSD) volumes). Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSPerformance.html> QUESTION 500 Having set up a website to automatically be redirected to a backup website if it fails, you realize that there are different types of failovers that are possible. You need all your resources to be available the majority of the time. Using Amazon Route 53 which configuration would best suit this requirement? A. Active-active failover B. None. Route 53 can't failover C. Active-passive failover D. Active-active-passive and other mixed configurations. Answer: A Explanation: You can set up a variety of failover configurations using Amazon Route 53 alias: weighted, latency, geolocation routing, and failover resource record sets. Active-active failover: Use this failover configuration when you want all of your resources to be available the majority of the time. When a resource becomes unavailable, Amazon Route 53 can detect that it's unhealthy and stop including it when responding to queries. Active-passive failover: Use this failover configuration when you want a primary group of resources to be available the majority of the time and you want a secondary group of resources to be on standby in case all of the primary resources become unavailable. When responding to queries, Amazon Route 53 includes only the healthy primary resources. If all of the primary resources are unhealthy, Amazon Route 53 begins to include only the healthy secondary resources in response to DNS queries. Active-active-passive and other mixed configurations: You can combine alias and non-alias resource record sets to produce a variety of Amazon Route 53 behaviors. Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html> At Lead2pass, we are positive that our Amazon AWS Certified Solutions Architect - Associate dumps with questions and answers PDF provide most in-depth solutions for individuals that are preparing for the Amazon AWS Certified Solutions Architect - Associate exam. Our updated AWS Certified Solutions Architect - Associate braindumps will allow you the opportunity to know exactly what to expect on the exam day and ensure that you can pass the exam beyond any doubt. AWS Certified Solutions Architect ? Associate new questions on Google Drive: <https://drive.google.com/open?id=0B3Syig5i8gpDR1h2VU4tOHhDcW8> **2017 Amazon AWS Certified Solutions Architect ? Associate exam dumps (All 680 Q&As) from Lead2pass:** <https://www.lead2pass.com/aws-certified-solutions-architect-associate.html> [100% Exam Pass Guaranteed]