

[Lead2pass New AWS Certified Solutions Architect - Associate Exam Dumps Free Download In Lead2pass 100% AWS Certified Solutions Architect - Associate Exam Questions (501-525)]

[2017 October Amazon Official New Released AWS Certified Solutions Architect ? Associate Dumps in Lead2pass.com!](#) 100% Free Download! 100% Pass Guaranteed! The Amazon AWS Certified Solutions Architect - Associate PDF, AWS Certified Solutions Architect - Associate VCE and AWS Certified Solutions Architect - Associate exam questions and answers at Lead2pass are written and prepared by Amazon affiliated trainers and lecturers with decades of experience in the IT field. This ensures that you are equipped with the latest and most current information to give you a better chance of passing the Amazon AWS Certified Solutions Architect - Associate exam. Following questions and answers are all new published by Amazon Official Exam Center:

<https://www.lead2pass.com/aws-certified-solutions-architect-associate.html> QUESTION 501 AWS CloudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CloudFormation takes care of provisioning and configuring those resources for you. What formatting is required for this template? A. JSON-formatted document B. CSS-formatted document C. XML-formatted document D. HTML-formatted document Answer: A Explanation: You can write an AWS CloudFormation template (a JSON-formatted document) in a text editor or pick an existing template. The template describes the resources you want and their settings. For example, suppose you want to create an Amazon EC2. Your template can declare an instance Amazon EC2 and describe its properties, as shown in the following example: {"AWSTemplateFormatVersion": "2010-09-09", "Description": "A simple Amazon EC2 instance", "Resources": {"MyEC2Instance": {"Type": "AWS::EC2::Instance", "Properties": {"ImageId": "ami-2f726546", "InstanceType": "t1.micro"}}}} Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-what-is-howdoesitwork.html> QUESTION 502 True or False: In Amazon Route 53, you can create a hosted zone for a top-level domain (TLD). A. FALSE B. False, Amazon Route 53 automatically creates it for you. C. True, only if you send an XML document with a CreateHostedZoneRequest element for TLD. D. TRUE Answer: A Explanation: In Amazon Route 53, you cannot create a hosted zone for a top-level domain (TLD). Reference: http://docs.aws.amazon.com/Route53/latest/APIReference/API_CreateHostedZone.html QUESTION 503 You decide that you need to create a number of Auto Scaling groups to try and save some money as you have noticed that at certain times most of your EC2 instances are not being used. By default, what is the maximum number of Auto Scaling groups that AWS will allow you to create? A. 12 B. Unlimited C. 20 D. 2 Answer: C Explanation: Auto Scaling is an AWS service that allows you to increase or decrease the number of EC2 instances within your application's architecture. With Auto Scaling, you create collections of EC2 instances, called Auto Scaling groups. You can create these groups from scratch, or from existing EC2 instances that are already in production. Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_autoscaling QUESTION 504 A user needs to run a batch process which runs for 10 minutes. This will only be run once, or at maximum twice, in the next month, so the processes will be temporary only. The process needs 15 X-Large instances. The process downloads the code from S3 on each instance when it is launched, and then generates a temporary log file. Once the instance is terminated, all the data will be lost. Which of the below mentioned pricing models should the user choose in this case? A. Spot instance B. Reserved instance C. On-demand instance D. EBS optimized instance Answer: A Explanation: In Amazon Web Services, the spot instance is useful when the user wants to run a process temporarily. The spot instance can terminate the instance if the other user outbids the existing bid. In this case all storage is temporary and the data is not required to be persistent. Thus, the spot instance is a good option to save money. Reference: <http://aws.amazon.com/ec2/purchasing-options/spot-instances/> QUESTION 505 Which of the following is NOT a characteristic of Amazon Elastic Compute Cloud (Amazon EC2)? A. It can be used to launch as many or as few virtual servers as you need. B. It increases the need to forecast traffic by providing dynamic IP addresses for static cloud computing. C. It eliminates your need to invest in hardware up front, so you can develop and deploy applications faster. D. It offers scalable computing capacity in the Amazon Web Services (AWS) cloud. Answer: B Explanation: Amazon Elastic Compute Cloud (Amazon EC2) provides scalable computing capacity in the Amazon Web Services (AWS) cloud. Using Amazon EC2 eliminates your need to invest in hardware up front, so you can develop and deploy applications faster. You can use Amazon EC2 to launch as many or as few virtual servers as you need, configure security and networking, and manage storage. Amazon EC2 enables you to scale up or down to handle changes in requirements or spikes in popularity, reducing your need to forecast traffic. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html> QUESTION 506 You have been storing massive amounts of data on Amazon Glacier for the past 2 years and now start to wonder if there are any limitations on this. What is the correct answer

to your question? A. The total volume of data is limited but the number of archives you can store are unlimited. B. The total volume of data is unlimited but the number of archives you can store are limited. C. The total volume of data and number of archives you can store are unlimited. D. The total volume of data is limited and the number of archives you can store are limited.

Answer: C
Explanation: An archive is a durably stored block of information. You store your data in Amazon Glacier as archives. You may upload a single file as an archive, but your costs will be lower if you aggregate your data. TAR and ZIP are common formats that customers use to aggregate multiple files into a single file before uploading to Amazon Glacier. The total volume of data and number of archives you can store are unlimited. Individual Amazon Glacier archives can range in size from 1 byte to 40 terabytes. The largest archive that can be uploaded in a single upload request is 4 gigabytes. For items larger than 100 megabytes, customers should consider using the Multipart upload capability. Archives stored in Amazon Glacier are immutable, i.e. archives can be uploaded and deleted but cannot be edited or overwritten. Reference: <https://aws.amazon.com/glacier/faqs/> QUESTION 507

You are setting up your first Amazon Virtual Private Cloud (Amazon VPC) so you decide to use the VPC wizard in the AWS console to help make it easier for you. Which of the following statements is correct regarding instances that you launch into a default subnet via the VPC wizard? A. Instances that you launch into a default subnet receive a public IP address and 10 private IP addresses. B. Instances that you launch into a default subnet receive both a public IP address and a private IP address. C. Instances that you launch into a default subnet don't receive any IP addresses and you need to define them manually. D. Instances that you launch into a default subnet receive a public IP address and 5 private IP addresses. Answer: B

Explanation: Instances that you launch into a default subnet receive both a public IP address and a private IP address. Instances in a default subnet also receive both public and private DNS hostnames. Instances that you launch into a nondefault subnet in a default VPC don't receive a public IP address or a DNS hostname. You can change your subnet's default public IP addressing behavior. Reference: <http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/default-vpc.html> QUESTION 508

A user has configured ELB with two EBS backed EC2 instances. The user is trying to understand the DNS access and IP support for ELB. Which of the below mentioned statements may not help the user understand the IP mechanism supported by ELB? A. The client can connect over IPV4 or IPV6 using Dualstack. B. Communication between the load balancer and back-end instances is always through IPV4. C. ELB DNS supports both IPV4 and IPV6. D. The ELB supports either IPV4 or IPV6 but not both. Answer: D

Explanation: Elastic Load Balancing supports both Internet Protocol version 6 (IPv6) and Internet Protocol version 4 (IPv4). Clients can connect to the user's load balancer using either IPv4 or IPv6 (in EC2-Classic) DNS. However, communication between the load balancer and its back-end instances uses only IPv4. The user can use the Dualstack-prefixed DNS name to enable IPv6 support for communications between the client and the load balancers. Thus, the clients are able to access the load balancer using either IPv4 or IPv6 as their individual connectivity needs dictate. Reference: <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/UserScenariosForEC2.html> QUESTION 509

Does AWS CloudFormation support Amazon EC2 tagging? A. Yes, AWS CloudFormation supports Amazon EC2 tagging. B. No, CloudFormation doesn't support any tagging. C. No, it doesn't support Amazon EC2 tagging. D. It depends if the Amazon EC2 tagging has been defined in the template. Answer: A

Explanation: In AWS CloudFormation, Amazon EC2 resources that support the tagging feature can also be tagged in an AWS template. The tag values can refer to template parameters, other resource names, resource attribute values (e.g. addresses), or values computed by simple functions (e.g., a concatenated list of strings). Reference: <http://aws.amazon.com/cloudformation/faqs/> QUESTION 510

An existing client comes to you and says that he has heard that launching instances into a VPC (virtual private cloud) is a better strategy than launching instances into a EC2-classic which he knows is what you currently do. You suspect that he is correct and he has asked you to do some research about this and get back to him. Which of the following statements is true in regards to what ability launching your instances into a VPC instead of EC2-Classic gives you? A. All of the things listed here. B. Change security group membership for your instances while they're running. C. Assign static private IP addresses to your instances that persist across starts and stops. D. Define network interfaces, and attach one or more network interfaces to your instances. Answer: A

Explanation: By launching your instances into a VPC instead of EC2-Classic, you gain the ability to: Assign static private IP addresses to your instances that persist across starts and stops. Assign multiple IP addresses to your instances. Define network interfaces, and attach one or more network interfaces to your instances. Change security group membership for your instances while they're running. Control the outbound traffic from your instances (egress filtering) in addition to controlling the inbound traffic to them (ingress filtering). Add an additional layer of access control to your instances in the form of network access control lists (ACL). Run your instances on single-tenant hardware. Reference: http://media.amazonwebservices.com/AWS_Cloud_Best_Practices.pdf QUESTION 511

Amazon S3 allows you to set per-file permissions to grant read and/or write access. However you have decided that you want an entire bucket with 100 files already in it to be accessible to the public. You don't want to go through 100 files individually and set permissions. What would be the best way

to do this? A. Move the bucket to a new region.B. Add a bucket policy to the bucket.C. Move the files to a new bucket.D. Use Amazon EBS instead of S3 Answer: BExplanation:Amazon S3 supports several mechanisms that give you flexibility to control who can access your data as well as how, when, and where they can access it. Amazon S3 provides four different access control mechanisms: AWS Identity and Access Management (IAM) policies, Access Control Lists (ACLs), bucket policies, and query string authentication. IAM enables organizations to create and manage multiple users under a single AWS account. With IAM policies, you can grant IAM users fine-grained control to your Amazon S3 bucket or objects. You can use ACLs to selectively add (grant) certain permissions on individual objects.Amazon S3 bucket policies can be used to add or deny permissions across some or all of the objects within a single bucket.With Query string authentication, you have the ability to share Amazon S3 objects through URLs that are valid for a specified period of time.Reference: <http://aws.amazon.com/s3/details/#security> QUESTION 512A user is accessing an EC2 instance on the SSH port for IP 10.20.30.40. Which one is a secure way to configure that the instance can be accessed only from this IP? A. In the security group, open port 22 for IP 10.20.30.40B. In the security group, open port 22 for IP 10.20.30.40/32C. In the security group, open port 22 for IP 10.20.30.40/24D. In the security group, open port 22 for IP 10.20.30.40/0 Answer: BExplanation:In AWS EC2, while configuring a security group, the user needs to specify the IP address in CIDR notation. The CIDR IP range 10.20.30.40/32 says it is for a single IP 10.20.30.40. If the user specifies the IP as 10.20.30.40 only, the security group will not accept and ask it in a CIRD format. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html> QUESTION 513Which of the following statements is true of creating a launch configuration using an EC2 instance? A. The launch configuration can be created only using the Query APIs.B. Auto Scaling automatically creates a launch configuration directly from an EC2 instance.C. A user should manually create a launch configuration before creating an Auto Scaling group.D. The launch configuration should be created manually from the AWS CLI. Answer: BExplanation:You can create an Auto Scaling group directly from an EC2 instance. When you use this feature, Auto Scaling automatically creates a launch configuration for you as well.Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/create-lc-with-instanceID.html> QUESTION 514You need to set up a high level of security for an Amazon Relational Database Service (RDS) you have just built in order to protect the confidential information stored in it. What are all the possible security groups that RDS uses? A. DB security groups, VPC security groups, and EC2 security groups.B. DB security groups only.C. EC2 security groups only.D. VPC security groups, and EC2 security groups. Answer: AExplanation:A security group controls the access to a DB instance. It does so by allowing access to IP address ranges or Amazon EC2 instances that you specify.Amazon RDS uses DB security groups, VPC security groups, and EC2 security groups. In simple terms, a DB security group controls access to a DB instance that is not in a VPC, a VPC security group controls access to a DB instance inside a VPC, and an Amazon EC2 security group controls access to an EC2 instance and can be used with a DB instance.Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html> QUESTION 515You have been using T2 instances as your CPU requirements have not been that intensive. However you now start to think about larger instance types and start looking at M1 and M3 instances. You are a little confused as to the differences between them as they both seem to have the same ratio of CPU and memory. Which statement below is incorrect as to why you would use one over the other? A. M3 instances are less expensive than M1 instances.B. M3 instances are configured with more swap memory than M1 instances.C. M3 instances provide better, more consistent performance than M1 instances for most use-cases.D. M3 instances also offer SSD-based instance storage that delivers higher I/O performance. Answer: BExplanation:Amazon EC2 allows you to set up and configure everything about your instances from your operating system up to your applications. An Amazon Machine Image (AMI) is simply a packaged-up environment that includes all the necessary bits to set up and boot your instance. M1 and M3 Standard instances have the same ratio of CPU and memory, some reasons below as to why you would use one over the other.M3 instances provide better, more consistent performance than M1 instances for most use-cases. M3 instances also offer SSD-based instance storage that delivers higher I/O performance. M3 instances are also less expensive than M1 instances. Due to these reasons, we recommend M3 for applications that require general purpose instances with a balance of compute, memory, and network resources.However, if you need more disk storage than what is provided in M3 instances, you may still find M1 instances useful for running your applications.Reference: <https://aws.amazon.com/ec2/faqs/> QUESTION 516You have set up an Elastic Load Balancer (ELB) with the usual default settings, which route each request independently to the application instance with the smallest load. However, someone has asked you to bind a user's session to a specific application instance so as to ensure that all requests coming from the user during the session will be sent to the same application instance. AWS has a feature to do this. What is it called? A. Connection drainingB. Proxy protocolC. TaggingD. Sticky session Answer: DExplanation:An Elastic Load Balancer(ELB) by default, routes each request independently to the application instance with the smallest load. However, you can use the sticky session feature (also known as session affinity), which enables the load balancer to bind a user's session to a specific application instance.

This ensures that all requests coming from the user during the session will be sent to the same application instance. The key to managing the sticky session is determining how long your load balancer should consistently route the user's request to the same application instance. If your application has its own session cookie, then you can set Elastic Load Balancing to create the session cookie to follow the duration specified by the application's session cookie. If your application does not have its own session cookie, then you can set Elastic Load Balancing to create a session cookie by specifying your own stickiness duration. You can associate stickiness duration for only HTTP/HTTPS load balancer listeners. An application instance must always receive and send two cookies: A cookie that defines the stickiness duration and a special Elastic Load Balancing cookie named AWSELB, that has the mapping to the application instance. Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/TerminologyandKeyConcepts.html#session-stickiness>

QUESTION 517A user wants to achieve High Availability with PostgreSQL DB. Which of the below mentioned functionalities helps achieve HA? A. Multi AZB. Read ReplicaC. Multi regionD. PostgreSQL does not support HA Answer:

AExplanation:The Multi AZ feature allows the user to achieve High Availability. For Multi AZ, Amazon RDS automatically provisions and maintains a synchronous "standby" replica in a different Availability Zone. Reference:

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html>

QUESTION 518A user has created an application which will be hosted on EC2. The application makes calls to DynamoDB to fetch certain data. The application is using the DynamoDB SDK to connect with from the EC2 instance. Which of the below mentioned statements is true with respect to the best practice for security in this scenario? A. The user should create an IAM user with DynamoDB access and use its credentials within the application to connect with DynamoDBB. The user should attach an IAM role with DynamoDB access to the EC2 instanceC.

The user should create an IAM role, which has EC2 access so that it will allow deploying the applicationD. The user should create an IAM user with DynamoDB and EC2 access. Attach the user with the application so that it does not use the root account credentials Answer: BExplanation:With AWS IAM a user is creating an application which runs on an EC2 instance and makes requests to AWS, such as DynamoDB or S3 calls. Here it is recommended that the user should not create an IAM user and pass the user's credentials to the application or embed those credentials inside the application. Instead, the user should use roles for EC2 and give that role access to DynamoDB /S3. When the roles are attached to EC2, it will give temporary security credentials to the application hosted on that EC2, to connect with DynamoDB / S3.Reference:

http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_WorkingWithGroupsAndUsers.html

QUESTION 519After setting up several database instances in Amazon Relational Database Service (Amazon RDS) you decide that you need to track the performance and health of your databases. How can you do this? A. Subscribe to Amazon RDS events to be notified when changes occur with a DB instance, DB snapshot, DB parameter group, or DB security group.B. Use the free Amazon CloudWatch service to monitor the performance and health of a DB instance.C. All of the items listed will track the performance and health of a database.D. View, download, or watch database log files using the Amazon RDS console or Amazon RDS APIs.You can also query some database log files that are loaded into database tables. Answer: CExplanation:Amazon Relational Database Service (Amazon RDS) is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. It provides cost-efficient, resizeable capacity for an industry-standard relational database and manages common database administration tasks. There are several ways you can track the performance and health of a database or a DB instance. You can:Use the free Amazon CloudWatch service to monitor the performance and health of a DB instance. Subscribe to Amazon RDS events to be notified when changes occur with a DB instance, DB snapshot, DB parameter group, or DB security group.View, download, or watch database log files using the Amazon RDS console or Amazon RDS APIs. You can also query some database log files that are loaded into database tables. Use the AWS CloudTrail service to record AWS calls made by your AWS account. The calls are recorded in log files and stored in an Amazon S3 bucket.Reference:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Monitoring.html

QUESTION 520You are building a system to distribute confidential documents to employees. Using CloudFront, what method could be used to serve content that is stored in S3, but not publically accessible from S3 directly? A. Add the CloudFront account security group "amazon-cf/amazon-cf-sg" to the appropriate S3 bucket policy.B. Create a S3 bucket policy that lists the CloudFront distribution ID as the Principal and the target bucket as the Amazon Resource Name (ARN).C. Create an Identity and Access Management (IAM) User for CloudFront and grant access to the objects in your S3 bucket to that IAM User.D. Create an Origin Access Identity (OAI) for CloudFront and grant access to the objects in your S3 bucket to that OAI. Answer: DExplanation:You restrict access to Amazon S3 content by creating an origin access identity, which is a special CloudFront user. You change Amazon S3 permissions to give the origin access identity permission to access your objects, and to remove permissions from everyone else. When your users access your Amazon S3 objects using CloudFront URLs, the CloudFront origin access identity gets the objects on your users' behalf. If your users try to access

objects using Amazon S3 URLs, they're denied access. The origin access identity has permission to access objects in your Amazon S3 bucket, but users don't. Reference:

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html> QUESTION 521A user has attached 1 EBS volume to a VPC instance. The user wants to achieve the best fault tolerance of data possible. Which of the below mentioned options can help achieve fault tolerance? A. Attach one more volume with RAID 1 configuration.B. Attach one more volume with RAID 0 configuration.C. Connect multiple volumes and stripe them with RAID 6 configuration.D. Use the EBS volume as a root device. Answer: AExplanation:The user can join multiple provisioned IOPS volumes together in a RAID 1 configuration to achieve better fault tolerance. RAID 1 does not provide a write performance improvement; it requires more bandwidth than non-RAID configurations since the data is written simultaneously to multiple volumes. Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html> QUESTION 522A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. The user has 3 elastic IPs and is trying to assign one of the Elastic IPs to the VPC instance from the console. The console does not show any instance in the IP assignment screen. What is a possible reason that the instance is unavailable in the assigned IP console? A. The IP address may be attached to one of the instancesB. The IP address belongs to a different zone than the subnet zoneC. The user has not created an internet gatewayD. The IP addresses belong to EC2 Classic; so they cannot be assigned to VPC Answer: DExplanation:A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. When the user is launching an instance he needs to select an option which attaches a public IP to the instance. If the user has not selected the option to attach the public IP then it will only have a private IP when launched. If the user wants to connect to an instance from the internet he should create an elastic IP with VPC. If the elastic IP is a part of EC2 Classic it cannot be assigned to a VPC instance. Reference:

<http://docs.aws.amazon.com/AmazonVPC/latest/GettingStartedGuide/LaunchInstance.html> QUESTION 523A user is aware that a huge download is occurring on his instance. He has already set the Auto Scaling policy to increase the instance count when the network I/O increases beyond a certain limit. How can the user ensure that this temporary event does not result in scaling? A. The network I/O are not affected during data downloadB. The policy cannot be set on the network I/O.C. There is no way the user can stop scaling as it is already configuredD. Suspend scaling Answer: DExplanation:The user may want to stop the automated scaling processes on the Auto Scaling groups either to perform manual operations or during emergency situations. To perform this, the user can suspend one or more scaling processes at any time. Once it is completed, the user can resume all the suspended processes. Reference:

http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html QUESTION 524Select a true statement about Amazon EC2 Security Groups (EC2-Classic). A. After you launch an instance in EC2-Classic, you can't change its security groups.B. After you launch an instance in EC2-Classic, you can change its security groups only once.C. After you launch an instance in EC2-Classic, you can only add rules to a security group.D. After you launch an instance in EC2-Classic, you cannot add or remove rules from a security group. Answer: AExplanation:After you launch an instance in EC2-Classic, you can't change its security groups. However, you can add rules to or remove rules from a security group, and those changes are automatically applied to all instances that are associated with the security group. Reference:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-network-security.html> QUESTION 525A user has created photo editing software and hosted it on EC2. The software accepts requests from the user about the photo format and resolution and sends a message to S3 to enhance the picture accordingly. Which of the below mentioned AWS services will help make a scalable software with the AWS infrastructure in this scenario? A. AWS Simple Notification ServiceB. AWS Simple Queue ServiceC. AWS Elastic TranscoderD. AWS Glacier Answer: BExplanation:Amazon Simple Queue Service (SQS) is a fast, reliable, scalable, and fully managed message queuing service. SQS provides a simple and cost-effective way to decouple the components of an application. The user can configure SQS, which will decouple the call between the EC2 application and S3. Thus, the application does not keep waiting for S3 to provide the data. Reference: <http://aws.amazon.com/sqs/faqs/> Lead2pass guarantees your AWS

Certified Solutions Architect - Associate exam success with our exam resources. Our AWS Certified Solutions Architect - Associate braindumps are the latest and developed by experienced IT certification professionals working in today's prospering companies and data centers. All our AWS Certified Solutions Architect - Associate braindumps include AWS Certified Solutions Architect - Associate real exam questions which guarantee your 100% success of AWS Certified Solutions Architect - Associate exam in your first try. More 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 796 Q&As) from Lead2pass:

<https://www.lead2pass.com/aws-certified-solutions-architect-associate.html> [100% Exam Pass Guaranteed]