

## [Lead2pass New Free Lead2pass Amazon AWS Certified Solutions Architect - Associate PDF Exam Questions And Answers Download (576-600)]

2017 October Amazon Official New Released AWS Certified Solutions Architect ? Associate Dumps in Lead2pass.com! 100% Free Download! 100% Pass Guaranteed! Amazing,100% candidates have passed the AWS Certified Solutions Architect - Associate exam by practising the preparation material of Lead2pass, because the braindumps are the latest and cover every aspect of AWS Certified Solutions Architect - Associate exam. Download the braindumps for an undeniable success in 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 576A user has created a CloudFormation stack. The stack creates AWS services, such as EC2 instances, ELB, AutoScaling, and RDS. While creating the stack it created EC2, ELB and AutoScaling but failed to create RDS. What will CloudFormation do in this scenario? A. Rollback all the changes and terminate all the created servicesB. It will wait for the user's input about the error and correct the mistake after the inputC.

CloudFormation can never throw an error after launching a few services since it verifies all the steps before launchingD. It will warn the user about the error and ask the user to manually create RDS Answer: AExplanation:AWS CloudFormation is an application management tool which provides application modeling, deployment, configuration, management and related activities.

The AWS CloudFormation stack is a collection of AWS resources which are created and managed as a single unit when AWS CloudFormation instantiates a template. If any of the services fails to launch, CloudFormation will rollback all the changes and terminate or delete all the created services.Reference: <http://aws.amazon.com/cloudformation/faqs/> QUESTION 577A major client who has been spending a lot of money on his internet service provider asks you to set up an AWS Direct Connection to try and save him some money. You know he needs high-speed connectivity. Which connection port speeds are available on AWS Direct

Connect? A. 500Mbps and 1GbpsB. 1Gbps and 10GbpsC. 100Mbps and 1GbpsD. 1Gbps Answer: BExplanation:AWS Direct Connect is a network service that provides an alternative to using the internet to utilize AWS cloud services.Using AWS Direct Connect, data that would have previously been transported over the Internet can now be delivered through a private network connection between AWS and your datacenter or corporate network.1Gbps and 10Gbps ports are available. Speeds of 50Mbps, 100Mbps, 200Mbps, 300Mbps, 400Mbps, and 500Mbps can be ordered from any APN partners supporting AWS Direct Connect.

Reference: <https://aws.amazon.com/directconnect/faqs/> QUESTION 578In Amazon EC2, what is the limit of Reserved Instances per Availability Zone each month? A. 5B. 20C. 50D. 10 Answer: BExplanation:There are 20 Reserved Instances per Availability Zone in each month. Reference: [http://docs.aws.amazon.com/general/latest/gr/aws\\_service\\_limits.html](http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html) QUESTION 579You have just finished setting up an advertisement server in which one of the obvious choices for a service was Amazon Elastic Map Reduce(EMR) and are now troubleshooting some weird cluster states that you are seeing. Which of the below is not an Amazon EMR cluster state? A. STARTINGB. STOPPEDC. RUNNINGD. WAITING Answer: BExplanation:Amazon Elastic Map Reduce (EMR) is a web service that enables businesses, researchers, data analysts, and developers to easily and cost-effectively process vast amounts of data. Amazon EMR historically referred to an Amazon EMR cluster (and all processing steps assigned to it) as a "cluster". Every cluster has a unique identifier that starts with "j-". The different cluster states of an Amazon EMR cluster are listed below.

STARTING ?The cluster provisions, starts, and configures EC2 instances. BOOTSTRAPPING ?Bootstrap actions are being executed on the cluster.RUNNING ?A step for the cluster is currently being run.WAITING ?The cluster is currently active, but has no steps to run. TERMINATING - The cluster is in the process of shutting down.TERMINATED - The cluster was shut down without error.TERMINATED\_WITH\_ERRORS - The cluster was shut down with errors.Reference:

<https://aws.amazon.com/elasticmapreduce/faqs/> QUESTION 580The AWS CloudHSM service defines a resource known as a high-availability (HA) \_\_\_\_\_, which is a virtual partition that represents a group of partitions, typically distributed between several physical HSMs for high-availability. A. proxy groupB. partition groupC. functional groupD. relational group Answer: BExplanation:The AWS CloudHSM service defines a resource known as a high-availability (HA) partition group, which is a virtual partition that represents a group of partitions, typically distributed between several physical HSMs for high-availability.

Reference: <http://docs.aws.amazon.com/cloudhsm/latest/userguide/configuring-ha.html> QUESTION 581Is it possible to get a history of all EC2 API calls made on your account for security analysis and operational troubleshooting purposes? A. Yes, by default, the history of your API calls is logged.B. Yes, you should turn on the CloudTrail in the AWS console.C. No, you can only get a history of VPC API calls.D. No, you cannot store history of EC2 API calls on Amazon. Answer: BExplanation:To get a history of all EC2 API calls (including VPC and EBS) made on your account, you simply turn on CloudTrail in the AWS Management Console.Reference: <https://aws.amazon.com/ec2/faqs/> QUESTION 582You have just set up your first Elastic Load Balancer (ELB) but it does not seem to be configured properly. You discover that before you start using ELB, you have to configure the listeners for

your load balancer. Which protocols does ELB use to support the load balancing of applications? A. HTTP and HTTPS B. HTTP, HTTPS, TCP, SSL and SSH C. HTTP, HTTPS, TCP, and SSL D. HTTP, HTTPS, TCP, SSL and SFTP Answer: CExplanation: Before you start using Elastic Load Balancing(ELB), you have to configure the listeners for your load balancer. A listener is a process that listens for connection requests. It is configured with a protocol and a port number for front-end (client to load balancer) and back-end (load balancer to back-end instance) connections. Elastic Load Balancing supports the load balancing of applications using HTTP, HTTPS (secure HTTP), TCP, and SSL (secure TCP) protocols. The HTTPS uses the SSL protocol to establish secure connections over the HTTP layer. You can also use SSL protocol to establish secure connections over the TCP layer. The acceptable ports for both HTTPS/SSL and HTTP/TCP connections are 25, 80, 443, 465, 587, and 1024-65535. Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-listener-config.html> QUESTION 583 After setting up some EC2 instances you now need to set up a monitoring solution to keep track of these instances and to send you an email when the CPU hits a certain threshold. Which statement below best describes what thresholds you can set to trigger a CloudWatch Alarm?

A. Set a target value and choose whether the alarm will trigger when the value is greater than (>), greater than or equal to (>=), less than (<), or less than or equal to (<=) that value. B. Thresholds need to be set in IAM not CloudWatch C. Only default thresholds can be set you can't choose your own thresholds. D. Set a target value and choose whether the alarm will trigger when the value hits this threshold

Answer: AExplanation: Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, and set alarms. When you create an alarm, you first choose the Amazon CloudWatch metric you want it to monitor. Next, you choose the evaluation period (e.g., five minutes or one hour) and a statistical value to measure (e.g., Average or Maximum). To set a threshold, set a target value and choose whether the alarm will trigger when the value is greater than (>), greater than or equal to (>=), less than (<), or less than or equal to (<=) that value. Reference: <http://aws.amazon.com/cloudwatch/faqs/>

QUESTION 584 After moving an E-Commerce website for a client from a dedicated server to AWS you have also set up auto scaling to perform health checks on the instances in your group and replace instances that fail these checks. Your client has come to you with his own health check system that he wants you to use as it has proved to be very useful prior to his site running on AWS. What do you think would be an appropriate response to this given all that you know about auto scaling?

A. It is not possible to implement your own health check system. You need to use AWSs health check system. B. It is not possible to implement your own health check system due to compatibility issues. C. It is possible to implement your own health check system and then send the instance's health information directly from your system to Cloud Watch. D. It is possible to implement your own health check system and then send the instance's health information directly from your system to Cloud Watch but only in the US East (N. Virginia) region.

Answer: CExplanation: Auto Scaling periodically performs health checks on the instances in your group and replaces instances that fail these checks. By default, these health checks use the results of EC2 instance status checks to determine the health of an instance. If you use a load balancer with your Auto Scaling group, you can optionally choose to include the results of Elastic Load Balancing health checks.

Auto Scaling marks an instance unhealthy if the calls to the Amazon EC2 action `DescribeInstanceStatus` returns any other state other than running, the system status shows impaired, or the calls to Elastic Load Balancing action `DescribeInstanceHealth` returns `OutOfService` in the instance state field.

After an instance is marked unhealthy because of an Amazon EC2 or Elastic Load Balancing health check, it is scheduled for replacement. You can customize the health check conducted by your Auto Scaling group by specifying additional checks or by having your own health check system and then sending the instance's health information directly from your system to Auto Scaling. Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/healthcheck.html>

QUESTION 585 When does the billing of an Amazon EC2 system begin?

A. It starts when the Status column for your distribution changes from Creating to Deployed. B. It starts as soon as you click the create instance option on the main EC2 console. C. It starts when your instance reaches 720 instance hours. D. It starts when Amazon EC2 initiates the boot sequence of an AMI instance.

Answer: DExplanation: Billing commences when Amazon EC2 initiates the boot sequence of an AMI instance. Billing ends when the instance terminates, which could occur through a web services command, by running "shutdown -h", or through instance failure. When you stop an instance, Amazon shuts it down but doesn't charge hourly usage for a stopped instance, or data transfer fees, but charges for the storage for any Amazon EBS volumes.

Reference: <http://aws.amazon.com/ec2/faqs/> QUESTION 586 You have just discovered that you can upload your objects to Amazon S3 using Multipart Upload API. You start to test it out but are unsure of the benefits that it would provide. Which of the following is not a benefit of using multipart uploads?

A. You can begin an upload before you know the final object size. B. Quick recovery from any network issues. C. Pause and resume object uploads. D. It's more secure than normal upload.

Answer: DExplanation: Multipart upload in Amazon S3 allows you to upload a single object as a set of parts. Each part is a contiguous portion of the object's data. You can upload these object parts independently and in any order. If transmission of any part fails, you can re-transmit that part without affecting other parts. After all parts of your object are uploaded, Amazon S3

assembles these parts and creates the object. In general, when your object size reaches 100 MB, you should consider using multipart uploads instead of uploading the object in a single operation. Using multipart upload provides the following advantages: Improved throughput--You can upload parts in parallel to improve throughput. Quick recovery from any network issues--Smaller part size minimizes the impact of restarting a failed upload due to a network error. Pause and resume object uploads--You can upload object parts over time. Once you initiate a multipart upload there is no expiry; you must explicitly complete or abort the multipart upload. Begin an upload before you know the final object size--You can upload an object as you are creating it. Reference:

<http://docs.aws.amazon.com/AmazonS3/latest/dev/uploadobjusingmpu.html> QUESTION 587 What is the data model of DynamoDB? A. Since DynamoDB is schema-less, there is no data model. B. "Items", with Keys and one or more Attribute; and "Attribute", with Name and Value. C. "Table", a collection of Items; "Items", with Keys and one or more Attribute; and "Attribute", with Name and Value. D. "Database", which is a set of "Tables", which is a set of "Items", which is a set of "Attributes". Answer: C Explanation: The data model of DynamoDB is: "Table", a collection of Items; "Items", with Keys and one or more Attribute; "Attribute", with Name and Value. Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DataModel.html> QUESTION 588 What happens to Amazon EBS root device volumes, by default, when an instance terminates? A. Amazon EBS root device volumes are moved to IAM. B. Amazon EBS root device volumes are copied into Amazon RDS. C. Amazon EBS root device volumes are automatically deleted. D. Amazon EBS root device volumes remain in the database until you delete them. Answer: C Explanation: By default, Amazon EBS root device volumes are automatically deleted when the instance terminates. Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/terminating-instances.html> QUESTION 589 Which of the following would you use to list your AWS Import/Export jobs? A. Amazon RDSB. AWS Import/Export Web Service ToolC. Amazon S3 REST API D. AWS Elastic Beanstalk Answer: C Explanation: You can list AWS Import/Export jobs with the ListJobs command using the command line client or REST API. Reference: <http://docs.aws.amazon.com/AWSImportExport/latest/DG/ListingYourJobs.html>

QUESTION 590 A gaming company comes to you and asks you to build them infrastructure for their site. They are not sure how big they will be as with all start ups they have limited money and big ideas. What they do tell you is that if the game becomes successful, like one of their previous games, it may rapidly grow to millions of users and generate tens (or even hundreds) of thousands of writes and reads per second. After considering all of this, you decide that they need a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Which of the following databases do you think would best fit their needs? A. Amazon DynamoDBB. Amazon RedshiftC. Any non-relational database.D. Amazon SimpleDB Answer: A Explanation: Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Amazon DynamoDB enables customers to offload the administrative burdens of operating and scaling distributed databases to AWS, so they don't have to worry about hardware provisioning, setup and configuration, replication, software patching, or cluster scaling. Today's web-based applications generate and consume massive amounts of data. For example, an online game might start out with only a few thousand users and a light database workload consisting of 10 writes per second and 50 reads per second. However, if the game becomes successful, it may rapidly grow to millions of users and generate tens (or even hundreds) of thousands of writes and reads per second. It may also create terabytes or more of data per day. Developing your applications against Amazon DynamoDB enables you to start small and simply dial-up your request capacity for a table as your requirements scale, without incurring downtime. You pay highly cost-efficient rates for the request capacity you provision, and let Amazon DynamoDB do the work over partitioning your data and traffic over sufficient server capacity to meet your needs. Amazon DynamoDB does the database management and administration, and you simply store and request your data. Automatic replication and failover provides built-in fault tolerance, high availability, and data durability. Amazon DynamoDB gives you the peace of mind that your database is fully managed and can grow with your application requirements.

Reference: <http://aws.amazon.com/dynamodb/faqs/> QUESTION 591 Mike is appointed as Cloud Consultant in Netcrak Inc. Netcrak has the following VPCs set-up in the US East Region: A VPC with CIDR block 10.10.0.0/16, a subnet in that VPC with CIDR block 10.10.1.0/24 A VPC with CIDR block 10.40.0.0/16, a subnet in that VPC with CIDR block 10.40.1.0/24 Netcrak Inc is trying to establish network connection between two subnets, a subnet with CIDR block 10.10.1.0/24 and another subnet with CIDR block 10.40.1.0/24. Which one of the following solutions should Mike recommend to Netcrak Inc? A. Create 2 Virtual Private Gateways and configure one with each VPC.B. Create one EC2 instance in each subnet, assign Elastic IPs to both instances, and configure a set up Site-to-Site VPN connection between both EC2 instances.C. Create a VPC Peering connection between both VPCs.D. Create 2 Internet Gateways, and attach one to each VPC. Answer: C Explanation: A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IP addresses. EC2 instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between

your own VPCs, or with a VPC in another AWS account within a single region. AWS uses the existing infrastructure of a VPC to create a VPC peering connection; it is neither a gateway nor a VPN connection, and does not rely on a separate piece of physical hardware. Reference: <http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-peering.html> QUESTION 592A favored client needs you to quickly deploy a database that is a relational database service with minimal administration as he wants to spend the least amount of time administering it. Which database would be the best option? A. Amazon SimpleDBB. Your choice of relational AMIs on Amazon EC2 and EBS.C. Amazon RDSD. Amazon Redshift Answer: CExplanation:Amazon Relational Database Service (Amazon RDS) is a web service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while managing time-consuming database administration tasks, freeing you up to focus on your applications and business. Amazon RDS gives you access to the capabilities of a familiar MySQL, Oracle, SQL Server, or PostgreSQL database engine. This means that the code, applications, and tools you already use today with your existing databases can be used with Amazon RDS. Amazon RDS automatically patches the database software and backs up your database, storing the backups for a user-defined retention period and enabling point-in-time recovery. Reference:

[https://aws.amazon.com/running\\_databases/#rds\\_anchor](https://aws.amazon.com/running_databases/#rds_anchor) QUESTION 593You're trying to delete an SSL certificate from the IAM certificate store, and you're getting the message "Certificate: <certificate-id> is being used by CloudFront." Which of the following statements is probably the reason why you are getting this error? A. Before you can delete an SSL certificate, you need to either rotate SSL certificates or revert from using a custom SSL certificate to using the default CloudFront certificate.B. You can't delete SSL certificates . You need to request it from AWS.C. Before you can delete an SSL certificate, you need to set up the appropriate access level in IAMD. Before you can delete an SSL certificate you need to set up https on your server. Answer: AExplanation: CloudFront is a web service that speeds up distribution of your static and dynamic web content, for example, .html, .css, .php, and image files, to end users. Every CloudFront web distribution must be associated either with the default CloudFront certificate or with a custom SSL certificate. Before you can delete an SSL certificate, you need to either rotate SSL certificates (replace the current custom SSL certificate with another custom SSL certificate) or revert from using a custom SSL certificate to using the default CloudFront certificate. Reference: <http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Troubleshooting.html> QUESTION 594You have been asked to design the storage layer for an application. The application requires disk performance of at least 100,000 IOPS in addition, the storage layer must be able to survive the loss of an individual disk. EC2 instance, or Availability Zone without any data loss. The volume you provide must have a capacity of at least 3 TB. Which of the following designs will meet these objectives'? A. Instantiate a c3.8xlarge instance in us-east-1. Provision 4x1TB EBS volumes, attach them to the instance, and configure them as a single RAID 5 volume. Ensure that EBS snapshots are performed every 15 minutes.B. Instantiate a c3.8xlarge instance in us-east-1. Provision 3xiTB EBS volumes, attach them to the Instance, and configure them as a single RAID 0 volume. Ensure that EBS snapshots are performed every 15 minutes.C. Instantiate an i2.8xlarge instance in us-east-1a. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instance. Provision 3x1TB EBS volumes, attach them to the instance, and configure them as a second RAID 0 volume. Configure synchronous, block-level replication from the ephemeral-backed volume to the EBS-backed volume.D. Instantiate a c3.8xlarge instance in us-east-1. Provision an AWS Storage Gateway and configure it for 3 TB of storage and 100,000 IOPS. Attach the volume to the instance. E. Instantiate an i2.8xlarge instance in us-east-1a. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instance. Configure synchronous, block- level replication to an identically configured instance in us-east-1b. Answer: CExplanation:QUESTION 595You need to set up security for your VPC and you know that Amazon VPC provides two features that you can use to increase security for your VPC: Security groups and network access control lists (ACLs). You start to look into security groups first. Which statement below is incorrect in relation to security groups? A. Are stateful: Return traffic is automatically allowed, regardless of any rules.B. Evaluate all rules before deciding whether to allow traffic.C. Support allow rules and deny rules.D. Operate at the instance level (first layer of defense). Answer: CExplanation:Amazon VPC provides two features that you can use to increase security for your VPC: Security groups--Act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level and supports allow rules only. Network access control lists (ACLs)--Act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level and supports allow rules and deny rules. Reference:

[http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_Security.html](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Security.html) QUESTION 596You are setting up some IAM user policies and have also become aware that some services support resource-based permissions, which let you attach policies to the service's resources instead of to IAM users or groups. Which of the below statements is true in regards to resource-level permissions? A. All services support resource-level permissions for all actions.B. Resource-level permissions are supported by Amazon CloudFrontC. All services support resource-level permissions only for some actions.D. Some services support resource-level permissions only for some actions. Answer: DExplanation:AWS Identity and Access Management is a web service

that enables Amazon Web Services (AWS) customers to manage users and user permissions in AWS. The service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon RDS, and the AWS Management Console. With IAM, you can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access. In addition to supporting IAM user policies, some services support resource-based permissions, which let you attach policies to the service's resources instead of to IAM users or groups. Resource-based permissions are supported by Amazon S3, Amazon SNS, and Amazon SQS. The resource-level permissions service supports IAM policies in which you can specify individual resources using Amazon Resource Names (ARNs) in the policy's Resource element. Some services support resource-level permissions only for some actions. Reference:

[http://docs.aws.amazon.com/IAM/latest/UserGuide/Using\\_SpecificProducts.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SpecificProducts.html) QUESTION 597A user wants to increase the durability and availability of the EBS volume. Which of the below mentioned actions should he perform? A. Take regular snapshots.B. Create an AMI.C. Create EBS with higher capacity.D. Access EBS regularly. Answer: AExplanation:In Amazon Web Services, Amazon EBS volumes that operate with 20 GB or less of modified data since their most recent snapshot can expect an annual failure rate (AFR) between 0.1% and 0.5%. For this reason, to maximize both durability and availability of their Amazon EBS data, the user should frequently create snapshots of the Amazon EBS volumes.Reference:

[http://media.amazonwebservices.com/AWS\\_Storage\\_Options.pdf](http://media.amazonwebservices.com/AWS_Storage_Options.pdf) QUESTION 598In relation to AWS CloudHSM, High-availability (HA) recovery is hands-off resumption by failed HA group members.Prior to the introduction of this function, the HA feature provided redundancy and performance, but required that a failed/lost group member be \_\_\_\_\_ reinstated. A. automatically B. periodicallyC. manuallyD. continuosly Answer: CExplanation:In relation to AWS CloudHSM, High-availability (HA) recovery is hands-off resumption by failed HA group members.Prior to the introduction of this function, the HA feature provided redundancy and performance, but required that a failed/lost group member be manually reinstated. Reference:

<http://docs.aws.amazon.com/cloudhsm/latest/userguide/ha-best-practices.html> QUESTION 599You have created a Route 53 latency record set from your domain to a machine in Northern Virginia and a similar record to a machine in Sydney.When a user located in US visits your domain he will be routed to: A. Northern VirginiaB. SydneyC. Both, Northern Virginia and SydneyD. Depends on the Weighted Resource Record Sets Answer: AExplanation:If your application is running on Amazon EC2 instances in two or more Amazon EC2 regions, and if you have more than one Amazon EC2 instance in one or more regions, you can use latency-based routing to route traffic to the correct region and then use weighted resource record sets to route traffic to instances within the region based on weights that you specify.For example, suppose you have three Amazon EC2 instances with Elastic IP addresses in the US East (Virginia) region and you want to distribute requests across all three IPs evenly for users for whom US East (Virginia) is the appropriate region. Just one Amazon EC2 instance is sufficient in the other regions, although you can apply the same technique to many regions at once. Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/Tutorials.html>

QUESTION 600Any person or application that interacts with AWS requires security credentials. AWS uses these credentials to identify who is making the call and whether to allow the requested access. You have just set up a VPC network for a client and you are now thinking about the best way to secure this network. You set up a security group called `vpccsecuritygroup`. Which following statement is true in respect to the initial settings that will be applied to this security group if you choose to use the default settings for this group? A. Allow all inbound traffic and allow no outbound traffic.B. Allow no inbound traffic and allow all outbound traffic.C. Allow inbound traffic on port 80 only and allow all outbound traffic.D. Allow all inbound traffic and allow all outbound traffic. Answer: BExplanation:Amazon VPC provides advanced security features such as security groups and network access control lists to enable inbound and outbound filtering at the instance level and subnet level. AWS assigns each security group a unique ID in the form `sg-xxxxxxxx`. The following are the initial settings for a security group that you create:Allow no inbound trafficAllow all outbound trafficReference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html>

You can pass Amazon AWS Certified Solutions Architect - Associate exam if you get a complete hold of AWS Certified Solutions Architect - Associate braindumps in Lead2pass. What's more, all the AWS Certified Solutions Architect - Associate Certification exam Q and As provided by Lead2pass are the latest. More AWS Certified Solutions Architect ? Associate new questions on Google Drive: <https://drive.google.com/open?id=0B3Syig5i8gpDVm1nMUwwQ1pkRE0> **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]