

[Lead2pass New Lead2pass Latest AWS Certified Solutions Architect - Associate Free Dumps Guarantee AWS Certified Solutions Architect - Associate Certification Exam 100% Success (526-550)]

[2017 October Amazon Official New Released AWS Certified Solutions Architect ? Associate Dumps in Lead2pass.com!](#) 100% Free Download! 100% Pass Guaranteed! Lead2pass is one of the leading exam preparation material providers. Its updated AWS Certified Solutions Architect - Associate braindumps in PDF can ensure most candidates pass the exam without too much effort. If you are struggling for the AWS Certified Solutions Architect - Associate exam, it will be a wise choice that get help from Lead2pass. Following questions and answers are all new published by Amazon Official Exam Center:

<https://www.lead2pass.com/aws-certified-solutions-architect-associate.html> QUESTION 526 Which one of the following answers is not a possible state of Amazon CloudWatch Alarm? A. INSUFFICIENT_DATA B. ALARM C. OK D. STATUS_CHECK_FAILED Answer: D Explanation: Amazon CloudWatch Alarms have three possible states: OK: The metric is within the defined threshold ALARM: The metric is outside of the defined threshold INSUFFICIENT_DATA: The alarm has just started, the metric is not available, or not enough data is available for the metric to determine the alarm state Reference:

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/AlarmThatSendsEmail.html> QUESTION 527 An accountant asks you to design a small VPC network for him and, due to the nature of his business, just needs something where the workload on the network will be low, and dynamic data will be accessed infrequently. Being an accountant, low cost is also a major factor. Which EBS volume type would best suit his requirements? A. Magnetic B. Any, as they all perform the same and cost the same. C. General Purpose (SSD) D. Magnetic or Provisioned IOPS (SSD) Answer: A Explanation: You can choose between three EBS volume types to best meet the needs of their workloads: General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic. General Purpose (SSD) is the new, SSD-backed, general purpose EBS volume type that we recommend as the default choice for customers. General Purpose (SSD) volumes are suitable for a broad range of workloads, including small to medium sized databases, development and test environments, and boot volumes. Provisioned IOPS (SSD) volumes offer storage with consistent and low-latency performance, and are designed for I/O intensive applications such as large relational or NoSQL databases. Magnetic volumes provide the lowest cost per gigabyte of all EBS volume types. Magnetic volumes are ideal for workloads where data is accessed infrequently, and applications where the lowest storage cost is important. Reference: <https://aws.amazon.com/ec2/faqs/>

QUESTION 528 A user is planning to launch a scalable web application. Which of the below mentioned options will not affect the latency of the application? A. Region. B. Provisioned IOPS. C. Availability Zone. D. Instance size. Answer: C Explanation: In AWS, the instance size decides the I/O characteristics. The provisioned IOPS ensures higher throughput, and lower latency. The region does affect the latency; latency will always be less when the instance is near to the end user. Within a region the user uses any AZ and this does not affect the latency. The AZ is mainly for fault toleration or HA. Reference: <https://aws.amazon.com/ec2/faqs/>

QUESTION 529 Which of the following strategies can be used to control access to your Amazon EC2 instances? A. DB security groups B. IAM policies C. None of these D. EC2 security groups Answer: D Explanation: IAM policies allow you to specify what actions your IAM users are allowed to perform against your EC2 Instances. However, when it comes to access control, security groups are what you need in order to define and control the way you want your instances to be accessed, and whether or not certain kind of communications are allowed or not. Reference: http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf

QUESTION 530 A user has launched one EC2 instance in the US East region and one in the US West region. The user has launched an RDS instance in the US East region. How can the user configure access from both the EC2 instances to RDS? A. It is not possible to access RDS of the US East region from the US West region B. Configure the US West region's security group to allow a request from the US East region's instance and configure the RDS security group's ingress rule for the US East EC2 group C. Configure the security group of the US East region to allow traffic from the US West region's instance and configure the RDS security group's ingress rule for the US East EC2 group D. Configure the security group of both instances in the ingress rule of the RDS security group Answer: C Explanation: The user cannot authorize an Amazon EC2 security group if it is in a different AWS Region than the RDS DB instance. The user can authorize an IP range or specify an Amazon EC2 security group in the same region that refers to an IP address in another region. In this case allow IP of US West inside US East's security group and open the RDS security group for US East region. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/UsingIAM.html>

QUESTION 531 In Amazon EC2, if your EBS volume stays in the detaching state, you can force the detachment by clicking _____. A. Force Detach B. Detach Instance C. AttachVolume D. AttachInstance Answer: A Explanation: If your volume stays in the detaching state, you can force the detachment by clicking Force Detach. Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithSecurityGroups.html

QUESTION 531 In Amazon EC2, if your EBS volume stays in the detaching state, you can force the detachment by clicking _____. A. Force Detach B. Detach Instance C. AttachVolume D. AttachInstance Answer: A Explanation: If your volume stays in the detaching state, you can force the detachment by clicking Force Detach. Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithSecurityGroups.html

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/ebs-detaching-volume.html> QUESTION 532 Do you need to shutdown your EC2 instance when you create a snapshot of EBS volumes that serve as root devices? A. No, you only need to shutdown an instance before deleting it. B. Yes C. No, the snapshot would turn off your instance automatically. D. No Answer: B Explanation: Yes, to create a snapshot for Amazon EBS volumes that serve as root devices, you should stop the instance before taking the snapshot. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html> QUESTION 533 An organization has a statutory requirement to protect the data at rest for data stored in EBS volumes. Which of the below mentioned options can the organization use to achieve data protection? A. Data replication. B. Data encryption. C. Data snapshot. D. All the options listed here. Answer: D Explanation: For protecting the Amazon EBS data at REST, the user can use options, such as Data Encryption (Windows / Linux / third party based), Data Replication (AWS internally replicates data for redundancy), and Data Snapshot (for point in time backup). Reference: http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf QUESTION 534 A client of yours has a huge amount of data stored on Amazon S3, but is concerned about someone stealing it while it is in transit. You know that all data is encrypted in transit on AWS, but which of the following is wrong when describing server-side encryption on AWS? A. Amazon S3 server-side encryption employs strong multi-factor encryption. B. Amazon S3 server-side encryption uses one of the strongest block ciphers available, 256-bit Advanced Encryption Standard (AES-256), to encrypt your data. C. In server-side encryption, you manage encryption/decryption of your data, the encryption keys, and related tools. D. Server-side encryption is about data encryption at rest--that is, Amazon S3 encrypts your data as it writes it to disks. Answer: C Explanation: Amazon S3 encrypts your object before saving it on disks in its data centers and decrypts it when you download the objects. You have two options depending on how you choose to manage the encryption keys: Server-side encryption and client-side encryption. Server-side encryption is about data encryption at rest--that is, Amazon S3 encrypts your data as it writes it to disks in its data centers and decrypts it for you when you access it. As long as you authenticate your request and you have access permissions, there is no difference in the way you access encrypted or unencrypted objects. Amazon S3 manages encryption and decryption for you. For example, if you share your objects using a pre-signed URL, that URL works the same way for both encrypted and unencrypted objects. In client-side encryption, you manage encryption/decryption of your data, the encryption keys, and related tools. Server-side encryption is an alternative to client-side encryption in which Amazon S3 manages the encryption of your data, freeing you from the tasks of managing encryption and encryption keys. Amazon S3 server-side encryption employs strong multi-factor encryption. Amazon S3 encrypts each object with a unique key. As an additional safeguard, it encrypts the key itself with a master key that it regularly rotates. Amazon S3 server-side encryption uses one of the strongest block ciphers available, 256-bit Advanced Encryption Standard (AES-256), to encrypt your data. Reference: <http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingServerSideEncryption.html> QUESTION 535 A user is running a batch process which runs for 1 hour every day. Which of the below mentioned options is the right instance type and costing model in this case if the user performs the same task for the whole year? A. EBS backed instance with on-demand instance pricing. B. EBS backed instance with heavy utilized reserved instance pricing. C. EBS backed instance with low utilized reserved instance pricing. D. Instance store backed instance with spot instance pricing. Answer: A Explanation: For Amazon Web Services, the reserved instance helps the user save money if the user is going to run the same instance for a longer period. Generally if the user uses the instances around 30-40% annually it is recommended to use RI. Here as the instance runs only for 1 hour daily it is not recommended to have RI as it will be costlier. The user should use on-demand with EBS in this case. Reference: <http://aws.amazon.com/ec2/purchasing-options/reserved-instances/> QUESTION 536 You have just set up a large site for a client which involved a huge database which you set up with Amazon RDS to run as a Multi-AZ deployment. You now start to worry about what will happen if the database instance fails. Which statement best describes how this database will function if there is a database failure? A. Updates to your DB Instance are synchronously replicated across Availability Zones to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure. B. Your database will not resume operation without manual administrative intervention. C. Updates to your DB Instance are asynchronously replicated across Availability Zones to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure. D. Updates to your DB Instance are synchronously replicated across S3 to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure. Answer: A Explanation: Amazon Relational Database Service (Amazon RDS) is a managed 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. When you create or modify your DB Instance to run as a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous "standby" replica in a different Availability Zone. Updates to your DB Instance are synchronously replicated across Availability Zones to the standby in order to keep both in sync and protect your latest database updates against DB

Instance failure. During certain types of planned maintenance, or in the unlikely event of DB Instance failure or Availability Zone failure, Amazon RDS will automatically failover to the standby so that you can resume database writes and reads as soon as the standby is promoted. Since the name record for your DB Instance remains the same, your application can resume database operation without the need for manual administrative intervention. With Multi-AZ deployments, replication is transparent: you do not interact directly with the standby, and it cannot be used to serve read traffic. If you are using Amazon RDS for MySQL and are looking to scale read traffic beyond the capacity constraints of a single DB Instance, you can deploy one or more Read Replicas. Reference: <http://aws.amazon.com/rds/faqs/> QUESTION 537 Which IAM role do you use to grant AWS Lambda permission to access a DynamoDB Stream? A. Dynamic role B. Invocation role C. Execution role D. Event Source role Answer: C Explanation: You grant AWS Lambda permission to access a DynamoDB Stream using an IAM role known as the "execution role". Reference: <http://docs.aws.amazon.com/lambda/latest/dg/intro-permission-model.html> QUESTION 538 Name the disk storage supported by Amazon Elastic Compute Cloud (EC2). A. None of these B. Amazon AppStream store C. Amazon SNS store D. Amazon Instance Store Answer: D Explanation: Amazon EC2 supports the following storage options: Amazon Elastic Block Store (Amazon EBS) Amazon EC2 Instance Store Amazon Simple Storage Service (Amazon S3) Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html> QUESTION 539 You are signed in as root user on your account but there is an Amazon S3 bucket under your account that you cannot access. What is a possible reason for this? A. An IAM user assigned a bucket policy to an Amazon S3 bucket and didn't specify the root user as a principal B. The S3 bucket is full. C. The S3 bucket has reached the maximum number of objects allowed. D. You are in the wrong availability zone Answer: A Explanation: With IAM, you can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access. In some cases, you might have an IAM user with full access to IAM and Amazon S3. If the IAM user assigns a bucket policy to an Amazon S3 bucket and doesn't specify the root user as a principal, the root user is denied access to that bucket. However, as the root user, you can still access the bucket by modifying the bucket policy to allow root user access. Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/iam-troubleshooting.html#testing2> QUESTION 540 You have a number of image files to encode. In an Amazon SQS worker queue, you create an Amazon SQS message for each file specifying the command (jpeg-encode) and the location of the file in Amazon S3. Which of the following statements best describes the functionality of Amazon SQS? A. Amazon SQS is a distributed queuing system that is optimized for horizontal scalability, not for single-threaded sending or receiving speeds. B. Amazon SQS is for single-threaded sending or receiving speeds. C. Amazon SQS is a non-distributed queuing system. D. Amazon SQS is a distributed queuing system that is optimized for vertical scalability and for single-threaded sending or receiving speeds. Answer: A Explanation: Amazon SQS is a distributed queuing system that is optimized for horizontal scalability, not for single-threaded sending or receiving speeds. A single client can send or receive Amazon SQS messages at a rate of about 5 to 50 messages per second. Higher receive performance can be achieved by requesting multiple messages (up to 10) in a single call. It may take several seconds before a message that has been to a queue is available to be received. Reference: http://media.amazonwebservices.com/AWS_Storage_Options.pdf QUESTION 541 A user is observing the EC2 CPU utilization metric on CloudWatch. The user has observed some interesting patterns while filtering over the 1 week period for a particular hour. The user wants to zoom that data point to a more granular period. How can the user do that easily with CloudWatch? A. The user can zoom a particular period by selecting that period with the mouse and then releasing the mouse B. The user can zoom a particular period by specifying the aggregation data for that period C. The user can zoom a particular period by double clicking on that period with the mouse D. The user can zoom a particular period by specifying the period in the Time Range Answer: A Explanation: Amazon CloudWatch provides the functionality to graph the metric data generated either by the AWS services or the custom metric to make it easier for the user to analyse. The AWS CloudWatch console provides the option to change the granularity of a graph and zoom in to see data over a shorter time period. To zoom, the user has to click in the graph details pane, drag on the graph area for selection, and then release the mouse button. Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/zoom_in_on_graph.html QUESTION 542 A scope has been handed to you to set up a super fast gaming server and you decide that you will use Amazon DynamoDB as your database. For efficient access to data in a table, Amazon DynamoDB creates and maintains indexes for the primary key attributes. 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. How many types of secondary indexes does DynamoDB support? A. 2 B. 16 C. 4 D. As many as you need. Answer: A Explanation: DynamoDB supports two types of secondary indexes: Local secondary index -- an index that has the same hash key as the table, but a different range key. A local secondary index is "local" in the sense that every partition of a local secondary index is scoped to a table partition that has the same hash key. Global secondary index -- an index with a hash and range 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. Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html> QUESTION 543 Select the correct statement: Within Amazon EC2, when using Linux instances, the device name /dev/sda1 is _____. A. reserved for EBS volumes B. recommended for EBS volumes C. recommended for instance store volumes D. reserved for the root device Answer: D Explanation: Within Amazon EC2, when using a Linux instance, the device name /dev/sda1 is reserved for the root device.

Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/device_naming.html QUESTION 544 The common use cases for DynamoDB Fine-Grained Access Control (FGAC) are cases in which the end user wants _____. A. to change the hash keys of the table directly B. to check if an IAM policy requires the hash keys of the tables directly C. to read or modify any codecommit key of the table directly, without a middle-tier service D. to read or modify the table directly, without a middle-tier service Answer: D Explanation: FGAC can benefit any application that tracks information in a DynamoDB table, where the end user (or application client acting on behalf of an end user) wants to read or modify the table directly, without a middle-tier service. For instance, a developer of a mobile app named Acme can use FGAC to track the top score of every Acme user in a DynamoDB table. FGAC allows the application client to modify only the top score for the user that is currently running the application. Reference:

http://aws.amazon.com/dynamodb/faqs/#security_anchor QUESTION 545 A user has set up the CloudWatch alarm on the CPU utilization metric at 50%, with a time interval of 5 minutes and 10 periods to monitor. What will be the state of the alarm at the end of 90 minutes, if the CPU utilization is constant at 80%? A. ALERT B. ALARM C. OK D. INSUFFICIENT_DATA Answer: B Explanation: In this case the alarm watches a metric every 5 minutes for 10 intervals. Thus, it needs at least 50 minutes to come to the "OK" state. Till then it will be in the INSUFFICIENT_DATA state. Since 90 minutes have passed and CPU utilization is at 80% constant, the state of alarm will be "ALARM". Reference:

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/AlarmThatSendsEmail.html> QUESTION 546 You 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 have already looked into security groups and you are now trying to understand ACLs. Which statement below is incorrect in relation to ACLs? A. Supports allow rules and deny rules. B. Is stateful: Return traffic is automatically allowed, regardless of any rules. C. Processes rules in number order when deciding whether to allow traffic. D. Operates at the subnet level (second layer of defense). Answer: B Explanation: 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 Network access control lists (ACLs)--Act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level Security groups are stateful: (Return traffic is automatically allowed, regardless of any rules) Network ACLs are stateless: (Return traffic must be explicitly allowed by rules)

Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Security.html QUESTION 547 A user comes to you and wants access to Amazon CloudWatch but only wants to monitor a specific LoadBalancer. Is it possible to give him access to a specific set of instances or a specific LoadBalancer? A. No because you can't use IAM to control access to CloudWatch data for specific resources. B. Yes. You can use IAM to control access to CloudWatch data for specific resources. C. No because you need to be Sysadmin to access CloudWatch data. D. Yes. Any user can see all CloudWatch data and needs no access rights. Answer: A Explanation: Amazon CloudWatch integrates with AWS Identity and Access Management (IAM) so that you can specify which CloudWatch actions a user in your AWS Account can perform. For example, you could create an IAM policy that gives only certain users in your organization permission to use GetMetricStatistics. They could then use the action to retrieve data about your cloud resources. You can't use IAM to control access to CloudWatch data for specific resources. For example, you can't give a user access to CloudWatch data for only a specific set of instances or a specific LoadBalancer. Permissions granted using IAM cover all the cloud resources you use with CloudWatch. In addition, you can't use IAM roles with the Amazon CloudWatch command line tools. Using Amazon CloudWatch with IAM doesn't change how you use CloudWatch. There are no changes to CloudWatch actions, and no new CloudWatch actions related to users and access control. Reference:

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/UsingIAM.html> QUESTION 548 A user is planning to make a mobile game which can be played online or offline and will be hosted on EC2. The user wants to ensure that if someone breaks the highest score or they achieve some milestone they can inform all their colleagues through email. Which of the below mentioned AWS services helps achieve this goal? A. AWS Simple Workflow Service. B. AWS Simple Email Service. C. Amazon Cognito D. AWS Simple Queue Service. Answer: B Explanation: Amazon Simple Email Service (Amazon SES) is a highly scalable and cost-effective email-sending service for businesses and developers. It integrates with other AWS services, making it easy to send emails from applications that are hosted on AWS. Reference: <http://aws.amazon.com/ses/faqs/> QUESTION 549 You have multiple VPN connections and want to provide secure communication between sites using the AWS VPN CloudHub.

Which statement is the most accurate in describing what you must do to set this up correctly? A. Create a virtual private gateway with multiple customer gateways, each with unique Border Gateway Protocol (BGP) Autonomous System Numbers (ASNs) B. Create a virtual private gateway with multiple customer gateways, each with a unique set of keys C. Create a virtual public gateway with multiple customer gateways, each with a unique Private subnet D. Create a virtual private gateway with multiple customer gateways, each with unique subnet id
Answer: A
Explanation: If you have multiple VPN connections, you can provide secure communication between sites using the AWS VPN CloudHub. The VPN CloudHub operates on a simple hub-and-spoke model that you can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing Internet connections who'd like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectivity between these remote offices. To use the AWS VPN CloudHub, you must create a virtual private gateway with multiple customer gateways, each with unique Border Gateway Protocol (BGP) Autonomous System Numbers (ASNs). Customer gateways advertise the appropriate routes (BGP prefixes) over their VPN connections. These routing advertisements are received and re-advertised to each BGP peer, enabling each site to send data to and receive data from the other sites. The routes for each spoke must have unique ASNs and the sites must not have overlapping IP ranges. Each site can also send and receive data from the VPC as if they were using a standard VPN connection. Reference:

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPN_CloudHub.html
QUESTION 550 You need to create an Amazon Machine Image (AMI) for a customer for an application which does not appear to be part of the standard AWS AMI template that you can see in the AWS console. What are the alternative possibilities for creating an AMI on AWS? A. You can purchase an AMIs from a third party but cannot create your own AMI. B. You can purchase an AMIs from a third party or can create your own AMI. C. Only AWS can create AMIs and you need to wait till it becomes available. D. Only AWS can create AMIs and you need to request them to create one for you.
Answer: B
Explanation: You can purchase an AMIs from a third party, including AMIs that come with service contracts from organizations such as Red Hat. You can also create an AMI and sell it to other Amazon EC2 users. After you create an AMI, you can keep it private so that only you can use it, or you can share it with a specified list of AWS accounts. You can also make your custom AMI public so that the community can use it. Building a safe, secure, usable AMI for public consumption is a fairly straightforward process, if you follow a few simple guidelines. Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html>
There is no doubt that Lead2pass is the top IT certificate exam material provider. All the braindumps are the latest and tested by senior Amazon lecturers and experts. Get the AWS Certified Solutions Architect - Associate exam braindumps in Lead2pass, and there would be no suspense to pass the exam. 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]