

[Lead2pass New Free Lead2pass Amazon AWS Certified Solutions Architect - Associate PDF Exam Questions And Answers Download (626-650)]

2017 October Amazon Official New Released AWS Certified Solutions Architect ? Associate Dumps in Lead2pass.com! 100% Free Download! 100% Pass Guaranteed! Whether you are a student attempting to pass AWS Certified Solutions Architect - Associate exam to be eligible for a post-graduate job, or a working professional hoping to improve your work credentials and earn that dream promotion Lead2pass is here to help. We have AWS Certified Solutions Architect - Associate exam dumps and brain dumps, so passing AWS Certified Solutions Architect - Associate exam is not an easy feat. Following questions and answers are all new published by Amazon Official Exam Center: <https://www.lead2pass.com/aws-certified-solutions-architect-associate.html>

QUESTION 626 Your company has multiple IT departments, each with their own VPC. Some VPCs are located within the same AWS account, and others in a different AWS account. You want to peer together all VPCs to enable the IT departments to have full access to each others' resources. There are certain limitations placed on VPC peering. Which of the following statements is incorrect in relation to VPC peering?

- A. Private DNS values cannot be resolved between instances in peered VPCs.
- B. You can have up to 3 VPC peering connections between the same two VPCs at the same time.
- C. You cannot create a VPC peering connection between VPCs in different regions.
- D. You have a limit on the number active and pending VPC peering connections that you can have per VPC.

Answer: B

Explanation: To create a VPC peering connection with another VPC, you need to be aware of the following limitations and rules:

- You cannot create a VPC peering connection between VPCs that have matching or overlapping CIDR blocks.
- You cannot create a VPC peering connection between VPCs in different regions.
- You have a limit on the number active and pending VPC peering connections that you can have per VPC.
- VPC peering does not support transitive peering relationships; in a VPC peering connection, your VPC will not have access to any other VPCs that the peer VPC may be peered with. This includes VPC peering connections that are established entirely within your own AWS account.
- You cannot have more than one VPC peering connection between the same two VPCs at the same time.
- The Maximum Transmission Unit (MTU) across a VPC peering connection is 1500 bytes.
- A placement group can span peered VPCs; however, you will not get full-bisection bandwidth between instances in peered VPCs.
- Unicast reverse path forwarding in VPC peering connections is not supported.
- You cannot reference a security group from the peer VPC as a source or destination for ingress or egress rules in your security group. Instead, reference CIDR blocks of the peer VPC as the source or destination of your security group's ingress or egress rules.
- Private DNS values cannot be resolved between instances in peered VPCs.

Reference: <http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/vpc-peering-overview.html#vpc-peering-limits>

QUESTION 627 A company wants to review the security requirements of Glacier. Which of the below mentioned statements is true with respect to the AWS Glacier data security?

- A. All data stored on Glacier is protected with AES-256 serverside encryption.
- B. All data stored on Glacier is protected with AES-128 serverside encryption.
- C. The user can set the serverside encryption flag to encrypt the data stored on Glacier.
- D. The data stored on Glacier is not encrypted by default.

Answer: A

Explanation: For Amazon Web Services, all the data stored on Amazon Glacier is protected using serverside encryption. AWS generates separate unique encryption keys for each Amazon Glacier archive, and encrypts it using AES-256. The encryption key then encrypts itself using AES-256 with a master key that is stored in a secure location.

Reference: http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf

QUESTION 628 You are architecting a highly-scalable and reliable web application which will have a huge amount of content. You have decided to use Cloudfront as you know it will speed up distribution of your static and dynamic web content and know that Amazon CloudFront integrates with Amazon CloudWatch metrics so that you can monitor your web application. Because you live in Sydney you have chosen the the Asia Pacific (Sydney) region in the AWS console. However you have set up this up but no CloudFront metrics seem to be appearing in the CloudWatch console. What is the most likely reason from the possible choices below for this?

- A. Metrics for CloudWatch are available only when you choose the same region as the application you are monitoring.
- B. You need to pay for CloudWatch for it to become active.
- C. Metrics for CloudWatch are available only when you choose the US East (N. Virginia).
- D. Metrics for CloudWatch are not available for the Asia Pacific region as yet.

Answer: C

Explanation: CloudFront is a global service, and metrics are available only when you choose the US East (N. Virginia) region in the AWS console. If you choose another region, no CloudFront metrics will appear in the CloudWatch console.

Reference: <http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/monitoring-using-cloudwatch.html>

QUESTION 629 A friend wants you to set up a small BitTorrent storage area for him on Amazon S3. You tell him it is highly unlikely that AWS would allow such a thing in their infrastructure. However you decide to investigate. Which of the following statements best describes using BitTorrent with Amazon S3?

- A. Amazon S3 does not support the BitTorrent protocol because it is used for pirated software.
- B. You can use the BitTorrent protocol but only for objects that are less than 100 GB in size.
- C. You can use the BitTorrent protocol

but you need to ask AWS for specific permissions first.D. You can use the BitTorrent protocol but only for objects that are less than 5 GB in size. Answer: DExplanation: BitTorrent is an open, peer-to-peer protocol for distributing files. You can use the BitTorrent protocol to retrieve any publicly-accessible object in Amazon S3. Amazon S3 supports the BitTorrent protocol so that developers can save costs when distributing content at high scale. Amazon S3 is useful for simple, reliable storage of any data. The default distribution mechanism for Amazon S3 data is via client/server download. In client/server distribution, the entire object is transferred point-to-point from Amazon S3 to every authorized user who requests that object. While client/server delivery is appropriate for a wide variety of use cases, it is not optimal for everybody. Specifically, the costs of client/server distribution increase linearly as the number of users downloading objects increases. This can make it expensive to distribute popular objects. BitTorrent addresses this problem by recruiting the very clients that are downloading the object as distributors themselves: Each client downloads some pieces of the object from Amazon S3 and some from other clients, while simultaneously uploading pieces of the same object to other interested "peers." The benefit for publishers is that for large, popular files the amount of data actually supplied by Amazon S3 can be substantially lower than what it would have been serving the same clients via client/server download. Less data transferred means lower costs for the publisher of the object. Reference:

<http://docs.aws.amazon.com/AmazonS3/latest/dev/S3Torrent.html> QUESTION 630 After a major security breach your manager has requested a report of all users and their credentials in AWS. You discover that in IAM you can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, MFA devices, and signing certificates. Which following statement is incorrect in regards to the use of credential reports? A. Credential reports are downloaded XML files.B. You can get a credential report using the AWS Management Console, the AWS CLI, or the IAM API.C. You can use the report to audit the effects of credential lifecycle requirements, such as password rotation.D. You can generate a credential report as often as once every four hours. Answer: AExplanation: To access your AWS account resources, users must have credentials. You can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, MFA devices, and signing certificates. You can get a credential report using the AWS Management Console, the AWS CLI, or the IAM API. You can use credential reports to assist in your auditing and compliance efforts. You can use the report to audit the effects of credential lifecycle requirements, such as password rotation. You can provide the report to an external auditor, or grant permissions to an auditor so that he or she can download the report directly. You can generate a credential report as often as once every four hours. When you request a report, IAM first checks whether a report for the account has been generated within the past four hours. If so, the most recent report is downloaded. If the most recent report for the account is more than four hours old, or if there are no previous reports for the account, IAM generates and downloads a new report. Credential reports are downloaded as comma-separated values (CSV) files. You can open CSV files with common spreadsheet software to perform analysis, or you can build an application that consumes the CSV files programmatically and performs custom analysis. Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/credential-reports.html> QUESTION 631 In the most recent company meeting, your CEO focused on the fact that everyone in the organization needs to make sure that all of the infrastructure that is built is truly scalable. Which of the following statements is incorrect in reference to scalable architecture? A. A scalable service is capable of handling heterogeneity.B. A scalable service is resilient.C. A scalable architecture won't be cost effective as it grows.D. Increasing resources results in a proportional increase in performance. Answer: CExplanation: In AWS it is critical to build a scalable architecture in order to take advantage of a scalable infrastructure. The cloud is designed to provide conceptually infinite scalability. However, you cannot leverage all that scalability in infrastructure if your architecture is not scalable. Both have to work together. You will have to identify the monolithic components and bottlenecks in your architecture, identify the areas where you cannot leverage the on-demand provisioning capabilities in your architecture, and work to refactor your application, in order to leverage the scalable infrastructure and take advantage of the cloud. Characteristics of a truly scalable application: Increasing resources results in a proportional increase in performance A scalable service is capable of handling heterogeneityA scalable service is operationally efficientA scalable service is resilientA scalable service should become more cost effective when it grows (Cost per unit reduces as the number of units increases)Reference:

http://media.amazonaws.com/AWS_Cloud_Best_Practices.pdf QUESTION 632 A user has defined an AutoScaling termination policy to first delete the instance with the nearest billing hour. AutoScaling has launched 3 instances in the US-East-1A region and 2 instances in the US-East-1B region. One of the instances in the US-East-1B region is running nearest to the billing hour. Which instance will AutoScaling terminate first while executing the termination action? A. Random Instance from US-East-1AB. Instance with the nearest billing hour in US-East-1BC. Instance with the nearest billing hour in US-East-1AD. Random instance from US-East-1B Answer: CExplanation: Even though the user has configured the termination policy, before AutoScaling selects an instance to terminate, it first identifies the Availability Zone that has more instances than the other

Availability Zones used by the group. Within the selected Availability Zone, it identifies the instance that matches the specified termination policy. Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/us-termination-policy.html>

QUESTION 633 A user has configured a website and launched it using the Apache web server on port 80. The user is using ELB with the EC2 instances for Load Balancing. What should the user do to ensure that the EC2 instances accept requests only from ELB? A. Configure the security group of EC2, which allows access to the ELB source security group B. Configure the EC2 instance so that it only listens on the ELB port C. Open the port for an ELB static IP in the EC2 security group D. Configure the security group of EC2, which allows access only to the ELB listener Answer: A Explanation: When a user is configuring ELB and registering the EC2 instances with it, ELB will create a source security group. If the user wants to allow traffic only from ELB, he should remove all the rules set for the other requests and open the port only for the ELB source security group. Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/using-elb-security-groups.html> QUESTION 634 A user is planning a highly available application deployment with EC2. Which of the below mentioned options will not help to achieve HA?

A. Elastic IP address B. PIOPSC. AMID. Availability Zones Answer: B Explanation: In Amazon Web Service, the user can achieve HA by deploying instances in multiple zones. The elastic IP helps the user achieve HA when one of the instances is down but still keeps the same URL. The AMI helps launching the new instance. The PIOPS is for the performance of EBS and does not help for HA. Reference: http://media.amazonwebservices.com/AWS_Web_Hosting_Best_Practices.pdf

QUESTION 635 You are playing around with setting up stacks using JSON templates in CloudFormation to try and understand them a little better. You have set up about 5 or 6 but now start to wonder if you are being charged for these stacks. What is AWS's billing policy regarding stack resources? A. You are not charged for the stack resources if they are not taking any traffic. B. You are charged for the stack resources for the time they were operating (even if you deleted the stack right away) C. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 60 minutes) D. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 30 minutes) Answer: B Explanation: A stack is a collection of AWS resources that you can manage as a single unit. In other words, you can create, update, or delete a collection of resources by creating, updating, or deleting stacks. All the resources in a stack are defined by the stack's AWS CloudFormation template. A stack, for instance, can include all the resources required to run a web application, such as a web server, a database, and networking rules. If you no longer require that web application, you can simply delete the stack, and all of its related resources are deleted. You are charged for the stack resources for the time they were operating (even if you deleted the stack right away). Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacks.html> QUESTION 636 You have been given a scope to set up an AWS Media Sharing Framework for a new start up photo sharing company similar to flickr. The first thing that comes to mind about this is that it will obviously need a huge amount of persistent data storage for this framework. Which of the following storage options would be appropriate for persistent storage?

A. Amazon Glacier or Amazon S3 B. Amazon Glacier or AWS Import/Export C. AWS Import/Export or Amazon CloudFront D. Amazon EBS volumes or Amazon S3 Answer: D Explanation: Persistent storage--If you need persistent virtual disk storage similar to a physical disk drive for files or other data that must persist longer than the lifetime of a single Amazon EC2 instance, Amazon EBS volumes or Amazon S3 are more appropriate. Reference:

http://media.amazonwebservices.com/AWS_Storage_Options.pdf QUESTION 637 After deploying a new website for a client on AWS, he asks if you can set it up so that if it fails it can be automatically redirected to a backup website that he has stored on a dedicated server elsewhere. You are wondering whether Amazon Route 53 can do this. Which statement below is correct in regards to Amazon Route 53? A. Amazon Route 53 can't help detect an outage. You need to use another service. B. Amazon Route 53 can help detect an outage of your website and redirect your end users to alternate locations. C. Amazon Route 53 can help detect an outage of your website but can't redirect your end users to alternate locations. D. Amazon Route 53 can't help detect an outage of your website, but can redirect your end users to alternate locations. Answer: B Explanation: With DNS Failover, Amazon Route 53 can help detect an outage of your website and redirect your end users to alternate locations where your application is operating properly. Reference: <http://aws.amazon.com/about-aws/whats-new/2013/02/11/announcing-dns-failover-for-route-53/>

QUESTION 638 In Route 53, what does a Hosted Zone refer to? A. A hosted zone is a collection of geographical load balancing rules for Route 53. B. A hosted zone is a collection of resource record sets hosted by Route 53. C. A hosted zone is a selection of specific resource record sets hosted by CloudFront for distribution to Route 53. D. A hosted zone is the Edge Location that hosts the Route 53 records for a user. Answer: B Explanation: A Hosted Zone refers to a selection of resource record sets hosted by Route 53.

Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/AboutHostedZones.html> QUESTION 639 Which of the following statements is true of Amazon EC2 security groups? A. You can change the outbound rules for EC2-Classic. Also, you can add and remove rules to a group at any time. B. You can modify an existing rule in a group. However, you can't add and remove rules to a group. C. None of the statements are correct. D. You can't change the outbound rules for EC2-Classic. However,

you can add and remove rules to a group at any time. Answer: DExplanation:When dealing with security groups, bear in mind that you can freely add and remove rules from a group, but you can't change the outbound rules for EC2-Classic. If you're using the Amazon EC2 console, you can modify existing rules, and you can copy the rules from an existing security group to a new security group.Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-network-security.html> QUESTION 640A web design company currently runs several FTP servers that their 250 customers use to upload and download large graphic files They wish to move this system to AWS to make it more scalable, but they wish to maintain customer privacy and Keep costs to a minimum.What AWS architecture would you recommend? A. ASK their customers to use an 53 client instead of an FTP client. Create a single 53 bucket Create an IAM user for each customer Put the IAM Users in a Group that has an IAM policy that permits access to sub-directories within the bucket via use of the 'username' Policy variable.B. Create a single 53 bucket with Reduced Redundancy Storage turned on and ask their customers to use an 53 client instead of an FTP client Create a bucket for each customer with a Bucket Policy that permits access only to that one customer.C. Create an auto-scaling group of FTP servers with a scaling policy to automatically scale-in when minimum network traffic on the auto-scaling group is below a given threshold. Load a central list of ftp users from 53 as part of the user Data startup script on each Instance.D. Create a single 53 bucket with Requester Pays turned on and ask their customers to use an 53 client instead of an FTP client Create a bucket for each customer with a Bucket Policy that permits access only to that one customer. Answer: AExplanation:QUESTION 641While creating a network in the VPC, which of the following is true of a NAT device? A. You have to administer the NAT Gateway Service provided by AWS.B. You can choose to use any of the three kinds of NAT devices offered by AWS for special purposes.C. You can use a NAT device to enable instances in a private subnet to connect to the Internet.D. You are recommended to use AWS NAT instances over NAT gateways, as the instances provide better availability and bandwidth. Answer: CExplanation:You can use a NAT device to enable instances in a private subnet to connect to the Internet (for example, for software updates) or other AWS services, but prevent the Internet from initiating connections with the instances. AWS offers two kinds of NAT devices ?a NAT gateway or a NAT instance. We recommend NAT gateways, as they provide better availability and bandwidth over NAT instances. The NAT Gateway service is also a managed service that does not require your administration efforts. A NAT instance is launched from a NAT AMI. You can choose to use a NAT instance for special purposes. Reference:

<http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-nat.html> QUESTION 642You need to create a management network using network interfaces for a virtual private cloud (VPC) network. Which of the following statements is incorrect pertaining to Best Practices for Configuring Network Interfaces. A. You can detach secondary (ethN) network interfaces when the instance is running or stopped.However, you can't detach the primary (eth0) interface.B. Launching an instance with multiple network interfaces automatically configures interfaces, private IP addresses, and route tables on the operating system of the instance.C. You can attach a network interface in one subnet to an instance in another subnet in the same VPC, however, both the network interface and the instance must reside in the same Availability Zone.D. Attaching another network interface to an instance is a valid method to increase or double the network bandwidth to or from the dual-homed instance Answer: DExplanation:Best Practices for Configuring Network InterfacesYou can attach a network interface to an instance when it's running (hot attach), when it's stopped (warm attach), or when the instance is being launched (cold attach). You can detach secondary (ethN) network interfaces when the instance is running or stopped. However, you can't detach the primary (eth0) interface.You can attach a network interface in one subnet to an instance in another subnet in the same VPC, however, both the network interface and the instance must reside in the same Availability Zone. When launching an instance from the CLI or API, you can specify the network interfaces to attach to the instance for both the primary (eth0) and additional network interfaces. Launching an instance with multiple network interfaces automatically configures interfaces, private IP addresses, and route tables on the operating system of the instance. A warm or hot attach of an additional network interface may require you to manually bring up the second interface, configure the private IP address, and modify the route table accordingly. (Instances running Amazon Linux automatically recognize the warm or hot attach and configure themselves.) Attaching another network interface to an instance is not a method to increase or double the network bandwidth to or from the dual-homed instance.Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#use-network-and-security-appliances-in-your-vpc> QUESTION 643All Amazon EC2 instances are assigned two IP addresses at launch. Which are those? A. 2 Elastic IP addressesB. A private IP address and an Elastic IP addressC. A public IP address and an Elastic IP addressD. A private IP address and a public IP address Answer: DExplanation:In Amazon EC2-Classic every instance is given two IP Addresses: a private IP address and a public IP addressReference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-instance-addressing.html#differences> QUESTION 644Your manager has asked you to set up a public subnet with instances that can send and receive internet traffic, and a private subnet that can't receive traffic directly from the internet, but can initiate traffic to the internet (and receive responses)

through a NAT instance in the public subnet. Hence, the following 3 rules need to be allowed:Inbound SSH traffic.Web servers in the public subnet to read and write to MS SQL servers in the private subnet Inbound RDP traffic from the Microsoft Terminal Services gateway in the public private subnet What are the respective ports that need to be opened for this? A. Ports 22,1433,3389 B. Ports 21,1433,3389C. Ports 25,1433,3389D. Ports 22,1343,3999 Answer: AExplanation: A network access control list (ACL) is an optional layer of security that acts as a firewall for controlling traffic in and out of a subnet. You might set up network ACLs with rules similar to your security groups in order to add an additional layer of security to your VPC.The following ports are recommended by AWS for a single subnet with instances that can receive and send Internet traffic and a private subnet that can't receive traffic directly from the Internet. However, it can initiate traffic to the Internet (and receive responses) through a NAT instance in the public subnet.Inbound SSH traffic. Port 22Web servers in the public subnet to read and write to MS SQL servers in the private subnet. Port 1433 Inbound RDP traffic from the Microsoft Terminal Services gateway in the public private subnet. Port 3389 Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Appendix_NACLs.html#VPC_Appendix_NACLs_Scenario_2 QUESTION 645You want to establish a dedicated network connection from your premises to AWS in order to save money by transferring data directly to AWS rather than through your internet service provider. You are sure there must be some other benefits beyond cost savings. Which of the following would not be considered a benefit if you were to establish such a connection? A. ElasticityB. Compatibility with all AWS services.C. Private connectivity to your Amazon VPC.D. Everything listed is a benefit. Answer: DExplanation: AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS.Using AWS Direct Connect, you can establish private connectivity between AWS and your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections. You could expect the following benefits if you use AWS Direct Connect.Reduced bandwidth costsConsistent network performanceCompatibility with all AWS servicesPrivate connectivity to your Amazon VPCElasticitySimplicityReference: <http://aws.amazon.com/directconnect/> QUESTION 646A user has launched an EC2 instance. The instance got terminated as soon as it was launched. Which of the below mentioned options is not a possible reason for this? A. The user account has reached the maximum volume limitB. The AMI is missing. It is the required partC. The snapshot is corruptD. The user account has reached the maximum EC2 instance limit Answer: DExplanation: When the user account has reached the maximum number of EC2 instances, it will not be allowed to launch an instance. AWS will throw an 'Instance Limit Exceeded' error. For all other reasons, such as "AMI is missing part", "Corrupt Snapshot" or "Volume limit has reached" it will launch an EC2 instance and then terminate it.Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_InstanceStraightToTerminated.html QUESTION 647Can I change the EC2 security groups after an instance is launched in EC2-Classic? A. Yes, you can change security groups after you launch an instance in EC2-Classic.B. No, you cannot change security groups after you launch an instance in EC2-Classic.C. Yes, you can only when you remove rules from a security group.D. Yes, you can only when you add rules to a security group. Answer: BExplanation: 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 648You can seamlessly join an EC2 instance to your directory domain. What connectivity do you need to be able to connect remotely to this instance? A. You must have IP connectivity to the instance from the network you are connecting from.B. You must have the correct encryption keys to connect to the instance remotely.C. You must have enough bandwidth to connect to the instance.D. You must use MFA authentication to be able to connect to the instance remotely. Answer: AExplanation: You can seamlessly join an EC2 instance to your directory domain when the instance is launched using the Amazon EC2 Simple Systems Manager. If you need to manually join an EC2 instance to your domain, you must launch the instance in the proper region and security group or subnet, then join the instance to the domain. To be able to connect remotely to these instances, you must have IP connectivity to the instances from the network you are connecting from. In most cases, this requires that an Internet gateway be attached to your VPC and that the instance has a public IP address. Reference: http://docs.aws.amazon.com/directoryservice/latest/admin-guide/join_a_directory.html QUESTION 649George has launched three EC2 instances inside the US-East-1a zone with his AWS account. Ray has launched two EC2 instances in the US-East-1a zone with his AWS account. Which of the below mentioned statements will help George and Ray understand the availability zone (AZ) concept better? A. All the instances of George and Ray can communicate over a private IP with a minimal costB. The US-East-1a region of George and Ray can be different availability zonesC. All the instances of George and Ray can communicate over a private IP without any costD. The instances of George and Ray will be running in the same data centre Answer: BExplanation: Each AWS region has multiple, isolated locations known as Availability Zones. To ensure that the AWS resources are

distributed across the Availability Zones for a region, AWS independently maps the Availability Zones to identifiers for each account. In this case the Availability Zone US-East-1a where George's EC2 instances are running might not be the same location as the US-East-1a zone of Ray's EC2 instances. There is no way for the user to coordinate the Availability Zones between accounts.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html> QUESTION 650 You are in the process of moving your friend's WordPress site onto AWS to try and save him some money, and you have told him that he should probably also move his domain name. He asks why he can't leave his domain name where it is and just have his infrastructure on AWS. What would be an incorrect response to his question? A. Route 53 offers low query latency for your end users. B. Route 53 is designed to automatically answer queries from the optimal location depending on network conditions. C. The globally distributed nature of AWS's DNS servers helps ensure a consistent ability to route your end users to your application. D. Route 53 supports Domain Name System Security Extensions (DNSSEC). Answer: D Explanation: Amazon Route 53 provides highly available and scalable Domain Name System (DNS), domain name registration, and health-checking web services. Route 53 is built using AWS's highly available and reliable infrastructure. The globally distributed nature of our DNS servers helps ensure a consistent ability to route your end users to your application by circumventing any internet or network related issues. Route 53 is designed to provide the level of dependability required by important applications. Using a global anycast network of DNS servers around the world, Route 53 is designed to automatically answer queries from the optimal location depending on network conditions. As a result, the service offers low query latency for your end users. Amazon Route 53 does not support Domain Name System Security Extensions (DNSSEC) at this time.

Reference: <https://aws.amazon.com/route53/faqs/> Your focus should be getting the best dumps to prepare for AWS Certified Solutions Architect - Associate exam. That is where Lead2pass comes in. We have collected an extensive library of exam dumps from Amazon certification. 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]