[Lead2pass New Free Lead2pass Amazon AWS Certified Solutions Architect - Associate PDF Exam Questions And Answers Download (551-575)

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 551An ERP application is deployed across multiple AZs in a single region. In the event of failure, the Recovery Time Objective (RTO) must be less than 3 hours, and the Recovery Point Objective (RPO) must be 15 minutes the customer realizes that data corruption occurred roughly 1.5 hours ago. What DR strategy could be used to achieve this RTO and RPO in the event of this kind of failure? A. Take hourly DB backups to 53, with transaction logs stored in 53 every 5 minutes.B. Use synchronous database master-slave replication between two availability zones.C. Take hourly DB backups to EC2 Instance store volumes with transaction logs stored In 53 every 5 minutes.D. Take 15 minute DB backups stored In Glacier with transaction logs stored in 53 every 5 minutes. Answer: AExplanation: QUESTION 552Which one of the below is not an AWS Storage Service? A. Amazon S3B. Amazon GlacierC. Amazon CloudFrontD. Amazon EBS Answer: CExplanation: AWS Storage Services are: Amazon S3Amazon GlacierAmazon EBSAWS Storage GatewayReference: https://console.aws.amazon.com/console QUESTION 553You are very concerned about security on your network because you have multiple programmers testing APIs and SDKs and you have no idea what is happening. You think CloudTrail may help but are not sure what it does. Which of the following statements best describes the AWS service CloudTrail? A. With AWS CloudTrail you can get a history of AWS API calls and related events for your account.B. With AWS CloudTrail you can get a history of IAM users for your account.C. With AWS CloudTrail you can get a history of S3 logfiles for your With AWS CloudTrail you can get a history of CloudFormation JSON scripts used for your account. Answer: A Explanation: With AWS CloudTrail, you can get a history of AWS API calls for your account, including API calls made via the AWS Management Console, the AWS SDKs, the command line tools, and higher-level AWS services. You can also identify which users and accounts called AWS APIs for services that support CloudTrail, the source IP address the calls were made from, and when the calls occurred. You can identify which users and accounts called AWS for services that support CloudTrail, the source IP address the calls were made from, and when the calls occurred. You can integrate CloudTrail into applications using the API, automate trail creation for your organization, check the status of your trails, and control how administrators turn CloudTrail logging on and off.Reference: http://docs.aws.amazon.com/awscloudtrail/latest/userguide/what is cloud trail top level.html QUESTION 554A user has deployed an application on his private cloud. The user is using his own monitoring tool. He wants to configure it so that whenever there is an error, the monitoring tool will notify him via SMS. Which of the below mentioned AWS services will help in this scenario? A. AWS SESB. AWS SNSC. None because the user infrastructure is in the private cloud.D. AWS SMS Answer: BExplanation: Amazon Simple Notification Service (Amazon SNS) is a fast, flexible, and fully managed push messaging service. Amazon SNS can be used to make push notifications to mobile devices. Amazon SNS can deliver notifications by SMS text message or email to the Amazon Simple Queue Service (SQS) queues or to any HTTP endpoint. In this case user can use the SNS apis to send SMS.Reference: http://aws.amazon.com/sns/ QUESTION 555After setting up an EC2 security group with a cluster of 20 EC2 instances, you find an error in the security group settings. You quickly make changes to the security group settings. When will the changes to the settings be effective? A. The settings will be effective immediately for all the instances in the security group.B. The settings will be effective only when all the instances are restarted.C. The settings will be effective for all the instances only after 30 minutes.D. The settings will be effective only for the new instances added to the security group. Answer: A Explanation: Amazon Redshift applies changes to a cluster security group immediately. So if you have associated the cluster security group with a cluster, inbound cluster access rules in the updated cluster security group apply immediately. Reference: http://docs.aws.amazon.com/redshift/latest/mgmt/working-with-security-groups.html QUESTION 556Regarding Amazon Route 53, 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 to route traffic to the correct region and then use route traffic to instances within the region, based on probabilities that you specify. A. weighted-based routing; alias resource record setsB. latency-based routing; weighted resource record setsC. weighted-based routing; weighted resource record setsD. latency-based routing; alias resource record sets Answer: BExplanation:Regarding Amazon Route 53, 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. Reference:

http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/Tutorials.html QUESTION 557You have a lot of data stored in the AWS Storage Gateway and your manager has come to you asking about how the billing is calculated, specifically the Virtual Tape Shelf usage. What would be a correct response to this? A. You are billed for the virtual tape data you store in Amazon Glacier and are billed for the size of the virtual tape. B. You are billed for the virtual tape data you store in Amazon Glacier and billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.C. You are billed for the virtual tape data you store in Amazon S3 and billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.D. You are billed for the virtual tape data you store in Amazon S3 and are billed for the size of the virtual tape. Answer: BExplanation: The AWS Storage Gateway is a service connecting an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and AWS's storage infrastructure. AWS Storage Gateway billing is as follows. Volume storage usage (per GB per month): You are billed for the Cached volume data you store in Amazon S3. You are only billed for volume capacity you use, not for the size of the volume you create. Snapshot Storage usage (per GB per month): You are billed for the snapshots your gateway stores in Amazon S3. These snapshots are stored and billed as Amazon EBS snapshots. Snapshots are incremental backups, reducing your storage charges. When taking a new snapshot, only the data that has changed since your last snapshot is stored. Virtual Tape Library usage (per GB per month): You are billed for the virtual tape data you store in Amazon S3. You are only billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape. Virtual Tape Shelf usage (per GB per month): You are billed for the virtual tape data you store in Amazon Glacier. You are only billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape. Reference:

https://aws.amazon.com/storagegateway/faqs/ QUESTION 558You are configuring a new VPC for one of your clients for a cloud migration project, and only a public VPN will be in place. After you created your VPC, you created a new subnet, a new internet gateway, and attached your internet gateway to your VPC. When you launched your first instance into your VPC, you realized that you aren't able to connect to the instance, even if it is configured with an elastic IP. What should be done to access the instance? A. A route should be created as 0.0.0.0/0 and your internet gateway as target.B. Attach another ENI to the instance and connect via new ENI.C. A NAT instance should be created and all traffic should be forwarded to NAT instance.D. A NACL should be created that allows all outbound traffic. Answer: AExplanation:All traffic should be routed via Internet Gateway. So, a route should be created with 0.0.0.0/0 as a source, and your Internet Gateway as your target.Reference:

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario1.html QUESTION 559A user is currently building a website which will require a large number of instances in six months, when a demonstration of the new site will be given upon launch. Which of the below mentioned options allows the user to procure the resources beforehand so that they need not worry about infrastructure availability during the demonstration? A. Procure all the instances as reserved instances beforehand.B. Launch all the instances as part of the cluster group to ensure resource availability.C. Pre-warm all the instances one month prior to ensure resource availability.D. Ask AWS now to procure the dedicated instances in 6 months. Answer: AExplanation: Amazon Web Services has massive hardware resources at its data centers, but they are finite. The best way for users to maximize their access to these resources is by reserving a portion of the computing capacity that they require. This can be done through reserved instances. With reserved instances, the user literally reserves the computing capacity in the Amazon Web Services cloud. Reference: http://media.amazonwebservices.com/AWS Building Fault Tolerant Applications.pdf QUESTION 560You receive a bill from AWS but are confused because you see you are incurring different costs for the exact same storage size in different regions on Amazon S3. You ask AWS why this is so. What response would you expect to receive from AWS? A. We charge less in different time zones.B. We charge less where our costs are less.C. This will balance out next bill.D. It must be a mistake. Answer: B Explanation: Amazon S3 is storage for the internet. It's a simple storage service that offers software developers a highly-scalable, reliable, and low-latency data storage infrastructure at very low costs.AWS charges less where their costs are less.For example, their costs are lower in the US Standard Region than in the US West (Northern California) Region.Reference:

https://aws.amazon.com/s3/faqs/ QUESTION 561You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups. Which RAID setup is not recommended for Amazon EBS? A. RAID 5 onlyB. RAID 5 and RAID 6C. RAID 1 onlyD. RAID 1 and RAID 6 Answer: BExplanation:With Amazon EBS, you can use any of the standard RAID configurations that you can use with a traditional bare metal server, as long as that particular RAID configuration is supported by the operating system for your instance. This is because all RAID is accomplished at the software level. For greater I/O performance than you can achieve with a single volume, RAID 0 can stripe multiple volumes together; for on-instance redundancy, RAID 1 can mirror two volumes together. RAID 5 and

RAID 6 are not recommended for Amazon EBS because the parity write operations of these RAID modes consume some of the IOPS available to your volumes. Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html QUESTION 562You receive the following request from a client to quickly deploy a static website for them, specifically on AWS. The requirements are low-cost, reliable, online storage, and a reliable and cost-effective way to route customers to the website, as well as a way to deliver content with low latency and high data transfer speeds so that visitors to his website don't experience unnecessary delays. What do you think would be the minimum AWS services that could fulfill the client's request? A. Amazon Route 53, Amazon CloudFront and Amazon VPC.B. Amazon S3, Amazon Route 53 and Amazon RDSC. Amazon S3, Amazon Route 53 and Amazon CloudFrontD. Amazon S3 and Amazon Route 53. Answer: CExplanation: You can easily and inexpensively use AWS to host a website that uses client-side technologies (such as HTML, CSS, and JavaScript) and does not require server-side technologies (such as PHP and ASP.NET). This type of site is called a static website, and is used to display content that does not change frequently. Before you create and deploy a static website, you must plan your architecture to ensure that it meets your requirements. Amazon S3, Amazon Route 53, and Amazon CloudFront would be required in this instance. Reference: http://docs.aws.amazon.com/gettingstarted/latest/swh/website-hosting-intro.html QUESTION 563What is the default maximum number of Access Keys per user? A. 10B. 15C. 2D. 20 Answer: CExplanation: The default maximum number of Access Keys per user is 2.Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html QUESTION 564What is the network performance offered by the c4.8xlarge instance in Amazon EC2? A. 20 GigabitB. 10 GigabitC. Very High but variableD. 5 Gigabit Answer: BExplanation: Networking performance offered by the c4.8xlarge instance is 10 Gigabit. Reference: http://aws.amazon.com/ec2/instance-types/ QUESTION 565Doug has created a VPC with CIDR 10.201.0.0/16 in his AWS account. In this VPC he has created a public subnet with CIDR block 10.201.31.0/24. While launching a new EC2 from the console, he is not able to assign the private IP address 10.201.31.6 to this instance. Which is the most likely reason for this issue? A. Private IP address 10.201.31.6 is blocked via ACLs in Amazon infrastructure as a part of platform security.B. Private address IP 10.201.31.6 is currently assigned to another interface.C. Private IP address 10.201.31.6 is not part of the associated subnet's IP address range. D. Private IP address 10.201.31.6 is reserved by Amazon for IP networking purposes. Answer: BExplanation:In Amazon VPC, you can assign any Private IP address to your instance as long as it is:Part of the associated subnet's IP address rangeNot reserved by Amazon for IP networking purposesNot currently assigned to another interfaceReference: http://aws.amazon.com/vpc/faqs/ OUESTION 566You need to create a JSON-formatted text file for AWS CloudFormation. This is your first template and the only thing you know is that the templates include several major sections but there is only one that is required for it to work. What is the only section required? A. MappingsB. OutputsC. ResourcesD. Conditions Answer: CExplanation: AWS CloudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CloudFormation takes care of provisioning and configuring those resources for you. A template is a JSON-formatted text file that describes your AWS infrastructure. Templates include several major sections. The Resources section is the only section that is required. The first character in the template must be an open brace ({), and the last character must be a closed brace (}). The following template fragment shows the template structure and sections. Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-anatomy.html QUESTION 567You are planning and configuring some EBS volumes for an application. In order to get the most performance out of your EBS volumes, you should attach them to an instance with enough ______ to support your volumes. A. RedundancyB. StorageC. BandwidthD. Memory Answer: CExplanation: When you plan and configure EBS volumes for your application, it is important to consider the configuration of the instances that you will attach the volumes to. In order to get the most performance out of your EBS volumes, you should attach them to an instance with enough bandwidth to support your volumes, such as an EBS-optimized instance or an instance with 10 Gigabit network connectivity. This is especially important when you use General Purpose (SSD) or Provisioned IOPS (SSD) volumes, or when you stripe multiple volumes together in a RAID configuration. Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-ec2-config.html QUESTION 568Can a single EBS volume be attached to multiple EC2 instances at the same time? A. YesB. NoC. Only for high-performance EBS volumes.D. Only when the instances are located in the US regions. Answer: BExplanation: You can't attach an EBS volume to multiple EC2 instances. This is because it is equivalent to using a single hard drive with many computers at the same time. Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html QUESTION 569How long does an AWS free usage tier EC2 last for? A. ForeverB. 12 Months upon signupC. 1 Month upon signupD. 6 Months upon signup Answer: BExplanation: The AWS free usage tier will expire 12 months from the date you sign up. When your free usage expires or if your application use exceeds the free usage tiers, you simply pay the standard, pay-as-you-go service rates. Reference: http://aws.amazon.com/free/faqs/

QUESTION 570A user is hosting a website in the US West-1 region. The website has the highest client base from the Asia-Pacific (Singapore / Japan) region. The application is accessing data from S3 before serving it to client. Which of the below mentioned regions gives a better performance for S3 objects? A. JapanB. SingaporeC. US EastD. US West-1 Answer: DExplanation: Access to Amazon S3 from within Amazon EC2 in the same region is fast. In this aspect, though the client base is Singapore, the application is being hosted in the US West-1 region. Thus, it is recommended that S3 objects be stored in the US-West-1 region. Reference: http://media.amazonwebservices.com/AWS Storage Options.pdf QUESTION 571Which of the following statements is true of tagging an Amazon EC2 resource? A. You don't need to specify the resource identifier while terminating a resource.B. You can terminate, stop, or delete a resource based solely on its tags.C. You can't terminate, stop, or delete a resource based solely on its tags.D. You don't need to specify the resource identifier while stopping a resource. Answer: CExplanation: You can assign tags only to resources that already exist. You can't terminate, stop, or delete a resource based solely on its tags; you must specify the resource identifier. Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Using Tags.html QUESTION 572 You have been setting up an Amazon Virtual Private Cloud (Amazon VPC) for your company, including setting up subnets. Security is a concern, and you are not sure which is the best security practice for securing subnets in your VPC. Which statement below is correct in describing the protection of AWS resources in each subnet? A. You can use multiple layers of security, including security groups and network access control lists (ACL).B. You can only use access control lists (ACL).C. You don't need any security in subnets.D. You can use multiple layers of security, including security groups, network access control lists (ACL) and CloudHSM. Answer: AExplanation: A subnet is a range of IP addresses in your VPC. You can launch AWS resources into a subnet that you select. Use a public subnet for resources that must be connected to the Internet, and a private subnet for resources that won't be connected to the Internet. To protect the AWS resources in each subnet, you can use multiple layers of security, including security groups and network access control lists (ACL). Reference:

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC Introduction.html QUESTION 573You have been asked to tighten up the password policies in your organization after a serious security breach, so you need to consider every possible security measure. Which of the following is not an account password policy for IAM Users that can be set? A. Force IAM users to contact an account administrator when the user has allowed his or her password to expire.B. A minimum password length.C. Force IAM users to contact an account administrator when the user has entered his password incorrectly.D. Prevent IAM users from reusing previous passwords. Answer: CExplanation:IAM users need passwords in order to access the AWS Management Console. (They do not need passwords if they will access AWS resources programmatically by using the CLI, AWS SDKs, or the APIs.)You can use a password policy to do these things:Set a minimum password length.Require specific character types, including uppercase letters, lowercase letters, numbers, and non-alphanumeric characters. Be sure to remind your users that passwords are case sensitive.Allow all IAM users to change their own passwords.Require IAM users to change their password after a specified period of time (enable password expiration).Prevent IAM users from reusing previous passwords.Force IAM users to contact an account administrator when the user has allowed his or her password to expire.Reference:

http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_ManagingPasswordPolicies.html QUESTION 574Your organization is in the business of architecting complex transactional databases. For a variety of reasons, this has been done on EBS. What is AWS's recommendation for customers who have architected databases using EBS for backups? A. Backups to Amazon S3 be performed through the database management system.B. Backups to AWS Storage Gateway be performed through the database management system.C. If you take regular snapshots no further backups are required.D. Backups to Amazon Glacier be performed through the database management system. Answer: AExplanation: Data stored in Amazon EBS volumes is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge. However, Amazon EBS replication is stored within the same availability zone, not across multiple zones; therefore, it is highly recommended that you conduct regular snapshots to Amazon S3 for long-term data durability. For customers who have architected complex transactional databases using EBS, it is recommended that backups to Amazon S3 be performed through the database management system so that distributed transactions and logs can be checkpointed. AWS does not perform backups of data that are maintained on virtual disks attached to running instances on Amazon EC2.Reference: http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf QUESTION 575You 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. How many EC2 instances would be sufficient to distribute requests in other regions? A. 3B. 9C. 2D. 1 Answer: DExplanation: 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 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=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]