

[Lead2pass New Free Lead2pass Amazon AWS Certified Solutions Architect - Associate PDF Exam Questions And Answers Download (676-700)]

2017 October Amazon Official New Released AWS Certified Solutions Architect ? Associate Dumps in Lead2pass.com! 100% Free Download! 100% Pass Guaranteed! I was very confused about selecting the best practice test website when preparing for my AWS Certified Solutions Architect - Associate certification exam. Luckily, a friend told me about Lead2pass.com. I passed the AWS Certified Solutions Architect - Associate exam from the first try. Excellent website for free exam dumps! Following questions and answers are all new published by Amazon Official Exam Center:

<https://www.lead2pass.com/aws-certified-solutions-architect-associate.html> QUESTION 676 An application hosted at the EC2 instance receives an HTTP request from ELB. The same request has an X-Forwarded-For header, which has three IP addresses. Which system's IP will be a part of this header? A. Previous Request IP address. B. Client IP address. C. All of the answers listed here. D. Load Balancer IP address. Answer: C Explanation: When a user sends a request to ELB over HTTP/HTTPS, the request header log at the instance will only receive the IP of ELB. This is because ELB is the interceptor between the EC2 instance and the client request. To get the client IP, use the header X-Forwarded-For in header. The client IP address in the X-Forwarded-For request header is followed by the IP addresses of each successive proxy that passes along the request. The last IP address is the IP address that connects to the back-end application instance. e.g. if the HTTP request already has a header when it reaches the Load Balancer, the IP address from which the request came is appended at the end of the header followed by the IP address of the Load Balancer. In such cases, the X-Forwarded-For request header takes the following form: X-Forwarded-For: clientIPAddress, previousRequestIPAddress, LoadBalancerIPAddress. Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/TerminologyandKeyConcepts.html> QUESTION 677 You need to develop and run some new applications on AWS and you know that Elastic Beanstalk and CloudFormation can both help as a deployment mechanism for a broad range of AWS resources. Which of the following statements best describes the differences between Elastic Beanstalk and CloudFormation? A. Elastic Beanstalk uses Elastic load balancing and CloudFormation doesn't. B. CloudFormation is faster in deploying applications than Elastic Beanstalk. C. Elastic Beanstalk is faster in deploying applications than CloudFormation. D. CloudFormation is much more powerful than Elastic Beanstalk, because you can actually design and script custom resources Answer: D Explanation: These services are designed to complement each other. AWS Elastic Beanstalk provides an environment to easily develop and run applications in the cloud. It is integrated with developer tools and provides a one-stop experience for you to manage the lifecycle of your applications. AWS CloudFormation is a convenient deployment mechanism for a broad range of AWS resources. It supports the infrastructure needs of many different types of applications such as existing enterprise applications, legacy applications, applications built using a variety of AWS resources and container-based solutions (including those built using AWS Elastic Beanstalk). AWS CloudFormation introduces two new concepts: The template, a JSON-format, text-based file that describes all the AWS resources you need to deploy to run your application and the stack, the set of AWS resources that are created and managed as a single unit when AWS CloudFormation instantiates a template. Reference:

<http://aws.amazon.com/cloudformation/faqs/> QUESTION 678 You need to set up a security certificate for a client's e-commerce website as it will use the HTTPS protocol. Which of the below AWS services do you need to access to manage your SSL server certificate? A. AWS Directory Service B. AWS Identity & Access Management C. AWS CloudFormation D. Amazon Route 53 Answer: B Explanation: AWS Identity and Access Management (IAM) is a web service that enables Amazon Web Services (AWS) customers to manage users and user permissions in AWS. All your SSL server certificates are managed by AWS Identity and Access management (IAM). Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/ManagingServerCerts.html> QUESTION

679 When controlling access to Amazon EC2 resources, each Amazon EBS Snapshot has a _____ attribute that controls which AWS accounts can use the snapshot. A. createVolumePermission B. LaunchPermission C. SharePermission D. RequestPermission Answer: A Explanation: Each Amazon EBS Snapshot has a createVolumePermission attribute that you can set to one or more AWS Account IDs to share the AMI with those AWS Accounts. To allow several AWS Accounts to use a particular EBS snapshot, you can use the snapshots's createVolumePermission attribute to include a list of the accounts that can use it. Reference: <http://docs.aws.amazon.com/AmazonEBS/latest/dev/using-iam.html> QUESTION

680 With which AWS orchestration service can you implement Chef recipes? A. CloudFormation B. Elastic Beanstalk C. Opsworks D. Lambda Answer: C QUESTION 681 You work for a construction company that has their production environment in AWS. The production environment consists of 3 identical web servers that are launched from a standard Amazon linux AMI using Auto Scaling. The web servers are launched in to the same public subnet and belong to the same security group. They also sit behind the same ELB. You decide to do some test and dev and you launch a 4th EC2 instance in to the same subnet and same security group. Annoyingly your

4th instance does not appear to have internet connectivity. What could be the cause of this? A. You need to update your routing table so as to provide a route out for this instance. B. Assign an elastic IP address to the fourth instance. C. You have not configured a NAT in the public subnet. D. You have not configured a routable IP address in the host OS of the fourth instance. Answer: C QUESTION 682 You need to add a route to your routing table in order to allow connections to the internet from your subnet. What route should you add? A. Destination: 192.168.1.258/0 --> Target: your Internet gateway B. Destination: 0.0.0.0/33 --> Target: your virtual private gateway C. Destination: 0.0.0.0/0 --> Target: 0.0.0.0/24 D. Destination: 10.0.0.0/32 --> Target: your virtual private gateway E. Destination: 0.0.0.0/0 --> Target: your Internet gateway Answer: E QUESTION 683 You have developed a new web application in us-west-2 that requires six Amazon Elastic Compute Cloud (EC2) instances running at all times. You have three availability zones available in that region (us-west-2a, us-west-2b, and us-west-2c). You need 100 percent fault tolerance if any single Availability Zone in us-west-2 becomes unavailable. How would you do this, each answer has 2 answers, select the answer with BOTH correct answers. A. Answer 1 - Us-west-2a with two EC2 instances, us-west-2b with two EC2 instances, and us-west-2c with two EC2 instances. Answer 2 - Us-west-2a with six EC2 instances, us-west-2b with six EC2 instances, and us-west-2c with no EC2 instances B. Answer 1 - Us-west-2a with six EC2 instances, us-west-2b with six EC2 instances, and us-west-2c with no EC2 instances. Answer 2 - Us-west-2a with three EC2 instances, us-west-2b with three EC2 instances, and us-west-2c with three EC2 instances C. Answer 1 - Us-west-2a with three EC2 instances, us-west-2b with three EC2 instances, and us-west-2c with no EC2 instances. Answer 2 - Us-west-2a with three EC2 instances, us-west-2b with three EC2 instances, and us-west-2c with three EC2 instances D. Answer 1 - Us-west-2a with three EC2 instances, us-west-2b with three EC2 instances, and us-west-2c with three EC2 instances. Answer 2 - Us-west-2a with four EC2 instances, us-west-2b with two EC2 instances, and us-west-2c with two EC2 instances. Answer: B QUESTION 684 You work for a major news network in Europe. They have just released a new app which allows users to report on events as and when they happen using their mobile phone. Users are able to upload pictures from the app and then other users will be able to view these pics. Your organization expects this app to grow very quickly, essentially doubling its user base every month. The app uses S3 to store the media and you are expecting sudden and large increases in traffic to S3 when a major news event takes place (as people will be uploading content in huge numbers). You need to keep your storage costs to a minimum however and it does not matter if some objects are lost. Which storage media should you use to keep costs as low as possible? A. S3 - Infrequently Accessed Storage B. S3 - Reduced Redundancy Storage (RRS) C. Glacier D. S3 - Provisioned IOPS. Answer: B QUESTION 685 You work for a famous bakery who are deploying a hybrid cloud approach. Their legacy IBM AS400 servers will remain on premise within their own datacenter however they will need to be able to communicate to the AWS environment over a site to site VPN connection. What do you need to do to establish the VPN connection? A. Connect to the environment using AWS Direct Connect B. Assign a public IP address to your Amazon VPC Gateway C. Create a dedicated NAT and deploy this to the public subnet D. Update your route table to add a route for the NAT to 0.0.0.0/0. Answer: B QUESTION 686 Your company has decided to set up a new AWS account for test and dev purposes. They already use AWS for production, but would like a new account dedicated for test and dev so as to not accidentally break the production environment. You launch an exact replica of your production environment using a CloudFormation template that your company uses in production. However CloudFormation fails. You use the exact same CloudFormation template in production, so the failure is something to do with your new AWS account. The CloudFormation template is trying to launch 60 new EC2 instances in a single AZ. After some research you discover that the problem is; A. For all new AWS accounts there is a soft limit of 20 EC2 instances per region. You should submit the limit increase form and retry the template after your limit has been increased B. For all new AWS accounts there is a soft limit of 20 EC2 instances per availability zone. You should submit the limit increase form and retry the template after your limit has been increased C. You cannot launch more than 20 instances in your default VPC, instead reconfigure the CloudFormation template to provision the instances in a custom VPC D. Your CloudFormation template is configured to use the parent account and not the new account. Change the account number in the CloudFormation template and relaunch the template. Answer: A QUESTION 687 You are a solutions architect who has been asked to do some consulting for a US company that produces re-useable rocket parts. They have a new web application that needs to be built and this application must be stateless. Which three services could you use to achieve this? A. AWS Storage Gateway, ElastiCache & ELB B. ELB, ElastiCache & RDS C. CloudWatch, RDS & DynamoDB D. RDS, DynamoDB & ElastiCache. Answer: D QUESTION 688 You run an automobile reselling company that has a popular online store on AWS. The application sits behind an Auto Scaling group and requires new instances of the Auto Scaling group to identify their public and private IP addresses. How can you achieve this? A. By using Ipconfig for windows or Ifconfig for Linux B. By using a cloud watch metric C. Using a Curl or Get Command to get the latest meta-data from <http://169.254.169.254/latest/meta-data/> D. Using a Curl or Get Command to get the latest user-data from <http://169.254.169.254/latest/user-data/> Answer: C QUESTION 689 You are a solutions architect working for a biotech company

who is pioneering research in immunotherapy. They have developed a new cancer treatment that may be able to cure up to 94% of cancers. They store their research data on S3, however recently an intern accidentally deleted some critical files. You've been asked to prevent this from happening in the future. What options below can prevent this? A. Make sure the interns can only access data on S3 using signed URLs. B. Enable S3 versioning on the bucket & enable Enable Multifactor Authentication (MFA) on the bucket. C. Use S3 Infrequently Accessed storage to store the data on. D. Create an IAM bucket policy that disables deletes.

Answer: B QUESTION 690 You are a security architect working for a large antivirus company. The production environment has recently been moved to AWS and is in a public subnet. You are able to view the production environment over HTTP however when your customers try to update their virus definition files over a custom port, that port is blocked. You log in to the console and you allow traffic in over the custom port. How long will this take to take effect? A. Straight away but to the new instances only. B. Immediately. C. After a few minutes this should take effect. D. Straight away to the new instances, but old instances must be stopped and restarted before the new rules apply. Answer: B QUESTION 691

You have been asked to identify a service on AWS that is a durable key value store. Which of the services below meets this definition? A. Mobile Hub B. Kinesis C. Simple Storage Service (S3) D. Elastic File Service (EFS) Answer: C QUESTION 692

By definition a public subnet within a VPC is one that; A. In it's routing table it has at least one route that uses an Internet Gateway (IGW). B. Has at least one route in it's routing table that routes via a Network Address Translation (NAT) instance. C. Where the the Network Access Control List (NACL) permitting outbound traffic to 0.0.0.0/0. D. Has had the public subnet check box ticked when setting up this subnet in the VPC console. Answer: A QUESTION 693

You work in the genomics industry and you process large amounts of genomic data using a nightly Elastic Map Reduce (EMR) job. This job processes a single 3 Tb file which is stored on S3. The EMR job runs on 3 on-demand core nodes and four on-demand task nodes. The EMR job is now taking longer than anticipated and you have been asked to advise how to reduced the completion time? A. Use four Spot Instances for the task nodes rather than four On-Demand instances. B. You should reduce the input split size in the MapReduce job configuration and then adjust the number of simultaneous mapper tasks so that more tasks can be processed at once. C. Store the file on Elastic File Service instead of S3 and then mount EFS as an independent volume for your core nodes. D. Configure an independent VPC in which to run the EMR jobs and then mount EFS as an independent volume for your core nodes. E. Enable termination protection for the job flow. Answer: B QUESTION 694

You work for a toy company that has a busy online store. As you are approaching christmas you find that your store is getting more and more traffic. You ensure that the web tier of your store is behind an Auto Scaling group, however you notice that the web tier is frequently scaling, sometimes multiple times in an hour, only to scale back after peak usage. You need to prevent this so that Auto Scaling does not scale as rapidly, just to scale back again. What option would help you to achieve this? A. Configure Auto Scaling to terminate your oldest instances first, then adjust your CloudWatch alarm. B. Configure Auto Scaling to terminate your newest instances first, then adjust your CloudWatch alarm. C. Change your Auto Scaling so that it only scales at scheduled times. D. Modify the Auto Scaling group cool-down timers & modify the Amazon CloudWatch alarm period that triggers your Auto Scaling scale down policy. Answer: D QUESTION 695

You are a student currently learning about the different AWS services. Your employer asks you to tell him a bit about Amazon's glacier service. Which of the following best describes the use cases for Glacier? A. Infrequently accessed data & data archives B. Hosting active databases C. Replicating Files across multiple availability zones and regions D. Frequently Accessed Data Answer: A QUESTION 696

You are a systems administrator and you need to monitor the health of your production environment. You decide to do this using Cloud Watch, however you notice that you cannot see the health of every important metric in the default dash board. Which of the following metrics do you need to design a custom cloud watch metric for, when monitoring the health of your EC2 instances? A. CPU Usage B. Memory usage C. Disk read operations D. Network in E. Estimated charges Answer: B QUESTION 697

You working in the media industry and you have created a web application where users will be able to upload photos they create to your website. This web application must be able to call the S3 API in order to be able to function. Where should you store your API credentials whilst maintaining the maximum level of security. A. Save the API credentials to your php files. B. Don't save your API credentials. Instead create a role in IAM and assign this role to an EC2 instance when you first create it. C. Save your API credentials in a public Github repository. D. Pass API credentials to the instance using instance userdata. Answer: B QUESTION 698

You work for a cosmetic company which has their production website on AWS. The site itself is in a two- tier configuration with web servers in the front end and database servers at the back end. The site uses using Elastic Load Balancing and Auto Scaling. The databases maintain consistency by replicating changes to each other as and when they occur. This requires the databases to have extremely low latency. Your website needs to be highly redundant and must be designed so that if one availability zone goes offline and Auto Scaling cannot launch new instances in the remaining Availability Zones the site will not go offline. How can the current architecture be enhanced to ensure this? A. Deploy your site in three different AZ's within the same region. Configure the Auto Scaling minimum to handle 50 percent of the peak load per zone. B.

Deploy your website in 2 different regions. Configure Route53 with a failover routing policy and set up health checks on the primary site.C. Deploy your site in three different AZ's within the same region. Configure the Auto Scaling minimum to handle 33 percent of the peak load per zone.D. Deploy your website in 2 different regions. Configure Route53 with Weighted Routing. Assign a weight of 25% to region 1 and a weight of 75% to region 2. Answer: A QUESTION 699You have been asked to create VPC for your company. The VPC must support both Internet-facing web applications (ie they need to be publicly accessible) and internal private applications (i.e. they are not publicly accessible and can be accessed only over VPN). The internal private applications must be inside a private subnet. Both the internet-facing and private applications must be able to leverage at least three Availability Zones for high availability.At a minimum, how many subnets must you create within your VPC to achieve this? A. 5 B. 3C. 4D. 6 Answer: D QUESTION 700You are hosting a MySQL database on the root volume of an EC2 instance. The database is using a large amount of IOPs and you need to increase the IOPs available to it.What should you do? A. Migrate the database to an S3 bucket.B. Migrate the database to Glacier.C. Add 4 additional EBS SSD volumes and create a RAID 10 using these volumes.D. Use Cloud Front to cache the database. Answer: C Thanks for the high quality AWS Certified Solutions Architect - Associate study guide. Will be back soon for more dumps. 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]