

[2017 New Lead2pass AWS Certified Solutions Architect - Associate New Questions Free Download (351-375)]

2017 August Amazon Official New Released AWS Certified Solutions Architect ? Associate Dumps in Lead2pass.com! 100% Free Download! 100% Pass Guaranteed! How to 100% pass AWS Certified Solutions Architect - Associate exam? Lead2pass provides the guaranteed AWS Certified Solutions Architect - Associate exam dumps to boost up your confidence in AWS Certified Solutions Architect - Associate exam. Successful candidates have provided their reviews about our AWS Certified Solutions Architect - Associate dumps. Now Lead2pass supplying the new version of AWS Certified Solutions Architect - Associate VCE and PDF dumps. We ensure our AWS Certified Solutions Architect - Associate exam questions are the most complete and authoritative compared with others', which will ensure your AWS Certified Solutions Architect - Associate exam pass. Following questions and answers are all new published by Amazon Official Exam Center:

<https://www.lead2pass.com/aws-certified-solutions-architect-associate.html> QUESTION 351A large real-estate brokerage is exploring the option of adding a cost-effective location based alert to their existing mobile application. The application backend infrastructure currently runs on AWS. Users who opt in to this service will receive alerts on their mobile device regarding real-estate offers in proximity to their location. For the alerts to be relevant, delivery time needs to be in the low minute count. The existing mobile app has 5 million users across the US. Which one of the following architectural suggestions would you make to the customer?

A. The mobile application will submit its location to a web service endpoint utilizing Elastic Load Balancing and EC2 instances; DynamoDB will be used to store and retrieve relevant offers. EC2 instances will communicate with mobile carriers/device providers to push alerts back to mobile application.

B. Use AWS DirectConnect or VPN to establish connectivity with mobile carriers. EC2 instances will receive the mobile applications' location through carrier connection; ROS will be used to store and retrieve relevant offers. EC2 instances will communicate with mobile carriers to push alerts back to the mobile application.

C. The mobile application will send device location using SQS. EC2 instances will retrieve the relevant offers from DynamoDB. AWS Mobile Push will be used to send offers to the mobile application.

D. The mobile application will send device location using AWS Mobile Push. EC2 instances will retrieve the relevant offers from DynamoDB. EC2 instances will communicate with mobile carriers/device providers to push alerts back to the mobile application.

Answer: C
Explanation: AWS using SQS to store the message from mobile apps, and using AWS Mobile Push to send offers to mobile apps.

QUESTION 352A company is building a voting system for a popular TV show. Viewers win watch the performances then visit the show's website to vote for their favorite performer. It is expected that in a short period of time after the show has finished, the site will receive millions of visitors. The visitors will first login to the site using their Amazon.com credentials and then submit their vote. After the voting is completed, the page will display the vote totals. The company needs to build the site such that it can handle the rapid influx of traffic while maintaining good performance but also wants to keep costs to a minimum. Which of the design patterns below should they use?

A. Use CloudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers; the web servers will first call the Login With Amazon service to authenticate the user, then process the user's vote and store the result into a multi-AZ Relational Database Service instance.

B. Use CloudFront and the static website hosting feature of S3 with the Javascript SDK to call the Login With Amazon service to authenticate the user; use IAM Roles to gain permissions to a DynamoDB table to store the user's vote.

C. Use CloudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers; the web servers will first call the Login with Amazon service to authenticate the user; the web servers will process the user's vote and store the result into a DynamoDB table using IAM Roles for EC2 instances to gain permissions to the DynamoDB table.

D. Use CloudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers; the web servers will first call the Login With Amazon service to authenticate the user; the web servers will process the user's vote and store the result into an SQS queue using IAM Roles for EC2 instances to gain permissions to the SQS queue. A set of application servers will then retrieve the items from the queue and store the result into a DynamoDB table.

Answer: D

QUESTION 353 You are developing a new mobile application and are considering storing user preferences in AWS. This would provide a more uniform cross-device experience to users using multiple mobile devices to access the application. The preference data for each user is estimated to be 50KB in size. Additionally, 5 million customers are expected to use the application on a regular basis. The solution needs to be cost-effective, highly available, scalable, and secure. How would you design a solution to meet the above requirements?

A. Setup an RDS MySQL instance in 2 availability zones to store the user preference data. Deploy a public-facing application on a server in front of the database to manage security and access credentials.

B. Setup a DynamoDB table with an item for each user having the necessary attributes to hold the user preferences. The mobile application will query the user preferences directly from the DynamoDB table. Utilize STS, Web Identity Federation, and DynamoDB Fine Grained Access Control to authenticate and authorize access.

C. Setup an RDS MySQL instance with multiple read replicas in 2 availability zones to store the

user preference data .The mobile application will query the user preferences from the read replicas. Leverage the MySQL user management and access privilege system to manage security and access credentials.D. Store the user preference data in S3 Setup a DynamoDB table with an item for each user and an item attribute pointing to the user' S3 object. The mobile application will retrieve the S3 URL from DynamoDB and then access the S3 object directly utilize STS, Web identity Federation, and S3 ACLs to authenticate and authorize access. Answer: BExplanation:

<https://aws.amazon.com/blogs/aws/fine-grained-access-control-for-amazon-dynamodb/>Here are some of the things that you can build using fine-grained access control:A mobile app that displays information for nearby airports, based on the user's location. The app can access and display attributes such as airline names, arrival times, and flight numbers. However, it cannot access or display pilot names or passenger counts.A mobile game which stores high scores for all users in a single table. Each user can update their own scores, but has no access to the other ones. QUESTION 354Your team has a tomcat-based Java application you need to deploy into development, test and production environments. After some research, you opt to use Elastic Beanstalk due to its tight integration with your developer tools and RDS due to its ease of management. Your QA team lead points out that you need to roll a sanitized set of production data into your environment on a nightly basis. Similarly, other software teams in your org want access to that same restored data via their EC2 instances in your VPC .The optimal setup for persistence and security that meets the above requirements would be the following. A. Create your RDS instance as part of your Elastic Beanstalk definition and alter its security group to allow access to it from hosts in your application subnets.B. Create your RDS instance separately and add its IP address to your application's DB connection strings in your code Alter its security group to allow access to it from hosts within your VPC's IP address block.C. Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variable. Create a security group for client machines and add it as a valid source for DB traffic to the security group of the RDS instance itself.D. Create your RDS instance separately and pass its DNS name to your's DB connection string as an environment variable Alter its security group to allow access to It from hosts In your application subnets. Answer: AExplanation: Elastic Beanstalk provides support for running Amazon RDS instances in your Elastic Beanstalk environment. This works great for development and testing environments, but is not ideal for a production environment because it ties the lifecycle of the database instance to the lifecycle of your application's environment.

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo.RDS.html> QUESTION 355You are looking to migrate your Development (Dev) and Test environments to AWS. You have decided to use separate AWS accounts to host each environment. You plan to link each accounts bill to a Master AWS account using Consolidated Billing. To make sure you Keep within budget you would like to implement a way for administrators in the Master account to have access to stop, delete and/or terminate resources in both the Dev and Test accounts. Identify which option will allow you to achieve this goal. A. Create IAM users in the Master account with full Admin permissions. Create cross-account roles in the Dev and Test accounts that grant the Master account access to the resources in the account by inheriting permissions from the Master account.B. Create IAM users and a cross-account role in the Master account that grants full Admin permissions to the Dev and Test accounts.C. Create IAM users in the Master account Create cross-account roles in the Dev and Test accounts that have full Admin permissions and grant the Master account access.D. Link the accounts using Consolidated Billing. This will give IAM users in the Master account access to resources in the Dev and Test accounts Answer: CExplanation:http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-user.html QUESTION 356Your customer is willing to consolidate their log streams (access logs application logs security logs etc.) in one single system. Once consolidated, the customer wants to analyze these logs in real time based on heuristics. From time to time, the customer needs to validate heuristics, which requires going back to data samples extracted from the last 12 hours?What is the best approach to meet your customer's requirements? A. Send all the log events to Amazon SQS. Setup an Auto Scaling group of EC2 servers to consume the logs and apply the heuristics.B. Send all the log events to Amazon Kinesis develop a client process to apply heuristics on the logsC. Configure Amazon Cloud Trail to receive custom logs, use EMR to apply heuristics the logsD. Setup an Auto Scaling group of EC2 syslogd servers, store the logs on S3 use EMR to apply heuristics on the logs Answer: BExplanation:Amazon Kinesis Streams allows for real-time data processing. With Amazon Kinesis Streams, you can continuously collect data as it is generated and promptly react to critical information about your business and operations.<https://aws.amazon.com/kinesis/streams/>

QUESTION 357You deployed your company website using Elastic Beanstalk and you enabled log file rotation to S3. An Elastic Map Reduce job is periodically analyzing the logs on S3 to build a usage dashboard that you share with your CIO. You recently improved overall performance of the website using Cloud Front for dynamic content delivery and your website as the originAfter this architectural change, the usage dashboard shows that the traffic on your website dropped by an order of magnitude. How do you fix your usage dashboard'? A. Enable Cloud Front to deliver access logs to S3 and use them as input of the Elastic Map Reduce job.B. Turn on Cloud Trail and use trail log tiles on S3 as input of the Elastic Map Reduce jobC. Change your log collection

process to use Cloud Watch ELB metrics as input of the Elastic Map Reduce jobD. Use Elastic Beanstalk "Rebuild Environment" option to update log delivery to the Elastic Map Reduce job.E. Use Elastic Beanstalk "Restart App server(s)" option to update log delivery to the Elastic Map Reduce job. Answer: AExplanation:

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/AccessLogs.html>

QUESTION 358You are running a successful multitier web application on AWS and your marketing department has asked you to add a reporting tier to the application. The reporting tier will aggregate and publish status reports every 30 minutes from user-generated information that is being stored in your web application's database. You are currently running a Multi-AZ RDS MySQL instance for the database tier. You also have implemented ElastiCache as a database caching layer between the application tier and database tier. Please select the answer that will allow you to successfully implement the reporting tier with as little impact as possible to your database. A. Continually send transaction logs from your master database to an S3 bucket and generate the reports off the S3 bucket using S3 byte range requests. B. Generate the reports by querying the synchronously replicated standby RDS MySQL instance maintained through Multi-AZ.C. Launch a RDS Read Replica connected to your Multi AZ master database and generate reports by querying the Read Replica.D. Generate the reports by querying the ElastiCache database caching tier. Answer: CExplanation:Amazon RDS allows you to use read replicas with Multi-AZ deployments. In Multi-AZ deployments for MySQL, Oracle, SQL Server, and PostgreSQL, the data in your primary DB Instance is synchronously replicated to a standby instance in a different Availability Zone (AZ). Because of their synchronous replication, Multi-AZ deployments for these engines offer greater data durability benefits than do read replicas. (In all Amazon RDS for Aurora deployments, your data is automatically replicated across 3 Availability Zones.)You can use Multi-AZ deployments and read replicas in conjunction to enjoy the complementary benefits of each. You can simply specify that a given Multi-AZ deployment is the source DB Instance for your Read replicas. That way you gain both the data durability and availability benefits of Multi-AZ deployments and the read scaling benefits of read replicas.Note that for Multi-AZ deployments, you have the option to create your read replica in an AZ other than that of the primary and the standby for even more redundancy. You can identify the AZ corresponding to your standby by looking at the "Secondary Zone" field of your DB Instance in the AWS Management Console. QUESTION 359A web company is looking to implement an intrusion detection and prevention system into their deployed VPC. This platform should have the ability to scale to thousands of instances running inside of the VPC.How should they architect their solution to achieve these goals? A. Configure an instance with monitoring software and the elastic network interface (ENI) set to promiscuous mode packet sniffing to see an traffic across the VPC.B. Create a second VPC and route all traffic from the primary application VPC through the second VPC where the scalable virtualized IDS/IPS platform resides.C. Configure servers running in the VPC using the host-based 'route' commands to send all traffic through the platform to a scalable virtualized IDS/IPS.D. Configure each host with an agent that collects all network traffic and sends that traffic to the IDS/IPS platform for inspection. Answer: BExplanation:A - promiscuous mode is not allowedC - there is no 'route' commandD - The company need IPS so agent will not work QUESTION 360A web-startup runs its very successful social news application on Amazon EC2 with an Elastic Load Balancer, an Auto-Scaling group of Java/Tomcat application-servers, and DynamoDB as data store. The main web-application best runs on m2 x large instances since it is highly memory-bound Each new deployment requires semi-automated creation and testing of a new AMI for the application servers which takes quite a while and is therefore only done once per week. Recently, a new chat feature has been implemented in nodejs and wails to be integrated in the architecture. First tests show that the new component is CPU bound Because the company has some experience with using Chef, they decided to streamline the deployment process and use AWS Ops Works as an application life cycle tool to simplify management of the application and reduce the deployment cycles.What configuration in AWS Ops Works is necessary to integrate the new chat module in the most cost-efficient and flexible way? A. Create one AWS Ops Works stack, create one AWS Ops Works layer, create one custom recipe B. Create one AWS Ops Works stack create two AWS Ops Works layers create one custom recipeC. Create two AWS Ops Works stacks create two AWS Ops Works layers create one custom recipeD. Create two AWS Ops Works stacks create two AWS Ops Works layers create two custom recipe Answer: B QUESTION 361Your firm has uploaded a large amount of aerial image data to S3 In the past, in your on-premises environment, you used a dedicated group of servers to oaten process this data and used Rabbit MQ -An open source messaging system to get job information to the servers. Once processed the data would go to tape and be shipped offsite. Your manager told you to stay with the current design, and leverage AWS archival storage and messaging services to minimize cost. Which is correct? A. Use SQS for passing job messages use Cloud Watch alarms to terminate EC2 worker instances when they become idle. Once data is processed, change the storage class of the S3 objects to Reduced Redundancy Storage.B. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SOS Once data is processed,C. Change the storage class of the S3 objects to Reduced Redundancy Storage. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SQS Once data is processed, change the storage class of the S3 objects to

Glacier.D. Use SNS to pass job messages use Cloud Watch alarms to terminate spot worker instances when they become idle. Once data is processed, change the storage class of the S3 object to Glacier. Answer: C Explanation: The question key part to focus on is ?and leverage AWS archival storage and messaging services to minimize cost.? For that the storage that is the lowest cost in the answers is Glacier, in addition, the messaging cost is less for SQS than for SNS if they both exceed 1 million transactions which is free. The only answer that satisfies the above two criteria is answer C. Also, there does not seem to be an urgency in speed of messaging therefore SQS satisfies that need. SNS being more real time delivery mechanism. QUESTION 362 What does Amazon S3 stand for? A. Simple Storage Solution. B. Storage Storage Storage (triple redundancy Storage). C. Storage Server Solution. D. Simple Storage Service. Answer: D QUESTION 363 You must assign each server to at least ____ security group A. 3 B. 2 C. 4 D. 1 Answer: D Explanation: Your AWS account automatically has a default security group per region for EC2-Classical. When you create a VPC, we automatically create a default security group for the VPC. If you don't specify a different security group when you launch an instance, the instance is automatically associated with the appropriate default security group.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html> QUESTION 364 Before I delete an EBS volume, what can I do if I want to recreate the volume later? A. Create a copy of the EBS volume (not a snapshot) B. Store a snapshot of the volume C. Download the content to an EC2 instance D. Back up the data in to a physical disk Answer: B Explanation: After you no longer need an Amazon EBS volume, you can delete it. After deletion, its data is gone and the volume can't be attached to any instance. However, before deletion, you can store a snapshot of the volume, which you can use to re-create the volume later. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-deleting-volume.html> QUESTION 365 Select the most correct answer: The device name /dev/sda1 (within Amazon EC2) is ____ A. Possible for EBS volumes B. Reserved for the root device C. Recommended for EBS volumes D. Recommended for instance store volumes Answer: B Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/RootDeviceStorage.html> The root device is typically /dev/sda1 (Linux) or xvda (Windows). QUESTION 366 If I want an instance to have a public IP address, which IP address should I use? A. Elastic IP Address B. Class B IP Address C. Class A IP Address D. Dynamic IP Address Answer: A QUESTION 367 What does RRS stand for when talking about S3? A. Redundancy Removal System B. Relational Rights Storage C. Regional Rights Standard D. Reduced Redundancy Storage Answer: D QUESTION 368 All Amazon EC2 instances are assigned two IP addresses at launch, out of which one can only be reached from within the Amazon EC2 network? A. Multiple IP address B. Public IP address C. Private IP address D. Elastic IP Address Answer: C QUESTION 369 What does Amazon SWF stand for? A. Simple Web Flow B. Simple Work Flow C. Simple Wireless Forms D. Simple Web Form Answer: B QUESTION 370 What is the Reduced Redundancy option in Amazon S3? A. Less redundancy for a lower cost. B. It doesn't exist in Amazon S3, but in Amazon EBS. C. It allows you to destroy any copy of your files outside a specific jurisdiction. D. It doesn't exist at all Answer: A QUESTION 371 Fill in the blanks: Resources that are created in AWS are identified by a unique identifier called an _____. A. Amazon Resource Number B. Amazon Resource Nametag C. Amazon Resource Name D. Amazon Resource Namespace Answer: C

QUESTION 372 If I write the below command, what does it do? `ec2-run ami-e3a5408a -n 20 -g appserver` A. Start twenty instances as members of appserver group. B. Creates 20 rules in the security group named appserver C. Terminate twenty instances as members of appserver group. D. Start 20 security groups Answer: A QUESTION 373 While creating an Amazon RDS DB, your first task is to set up a DB _____ that controls what IP addresses or EC2 instances have access to your DB Instance. A. Security Pool B. Secure Zone C. Security Token Pool D. Security Group Answer: D QUESTION 374 When you run a DB Instance as a Multi-AZ deployment, the "_____" serves database writes and reads A. secondary B. backup C. stand by D. primary Answer: D QUESTION 375 Every user you create in the IAM system starts with _____. A. Partial permissions B. Full permissions C. No permissions Answer: C

All the AWS Certified Solutions Architect - Associate braindumps are updated. Get a complete hold of AWS Certified Solutions Architect - Associate PDF dumps and AWS Certified Solutions Architect - Associate practice test with free VCE player through Lead2pass and boost up your skills. 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 680 Q&As) from Lead2pass:**

<https://www.lead2pass.com/aws-certified-solutions-architect-associate.html> [100% Exam Pass Guaranteed]