

پروژه سوم

رشته: پردازش ابری

ماژول: ارزیابی دانش و مهارت‌های AWS Solution Architect

تعداد سوالات : ۳۴ سوال

مدت زمان: ۴۵ دقیقه

نام و شماره رقابت کننده:

QUESTION 1

A large financial institution is gradually moving their infrastructure and applications to AWS. The company has data needs that will utilize all of RDS, DynamoDB, Redshift, and ElastiCache. Which description best describes Amazon Redshift?

- ☐ Can be used to significantly improve latency and throughput for many read-heavy application workloads.
- ☐ Near real-time complex querying on massive data sets.
- ☐ Cloud-based relational database.
- ☐ Key-value and document database that delivers single-digit millisecond performance at any scale.

QUESTION 2

You suspect that one of the AWS services your company is using has gone down. Which service can provide you proactive and transparent notifications about the status of your specific AWS environment?

- ☐ AWS Trusted Advisor
- ☐ Amazon Inspector
- ☐ AWS Personal Health Dashboard
- ☐ AWS Organizations

QUESTION 3

Your development team has created a gaming application that uses DynamoDB to store user statistics and provide fast game updates back to users. The team has begun testing the application but needs a consistent data set to perform tests with. The testing process alters the dataset, so the baseline data needs to be retrieved upon each new test. Which AWS service can meet this need by exporting data from DynamoDB and importing data into DynamoDB?

☐ Elastic Map Reduce

☐ DAX

☐ Redshift

☐ AWS Import/Export

QUESTION 4

You have two EC2 instances running in the same VPC, but in different subnets. You are removing the secondary ENI from an EC2 instance and attaching it to another EC2 instance. You want this to be fast and with limited disruption. So you want to attach the ENI to the EC2 instance when it's running. What is this called?

☐ synchronous attach

☐ cold attach

☐ warm attach

☐ hot attach

QUESTION 5

You are working as a Solutions Architect in a large healthcare organization. You have many Auto Scaling groups that you need to create. One requirement is that you need to reuse some software licenses and therefore need to use dedicated hosts on EC2 instances in your Auto Scaling groups. What step must you take to meet this requirement?

- ☐ Create the Dedicated Host EC2 instances, and then add them to an existing Auto Scaling group.
- ☐ Make sure your launch configurations are using Dedicated Hosts.
- ☐ Use a launch template with your Auto Scaling group and select the Dedicated Host option.
- ☐ Create your launch configuration, but manually change the instances to Dedicated

QUESTION 6

You have multiple EC2 instances housing applications in a VPC in a single Availability Zone. Your EC2 workloads need low-latency network performance, high network throughput, and a tightly-coupled node-to-node communication. What's the best measure you can do to ensure this throughput?

- ☐ Use Elastic Network Interfaces
- ☐ Use Auto Scaling Groups
- ☐ Increase the size of the instances
- ☐ Launch your instances in a cluster placement group

QUESTION 7

A software company is looking for compute capacity in the cloud for a fault-tolerant and flexible application. The application is not mission-critical, so occasional downtime is acceptable. What type of EC2 servers can be used to meet these requirements at the lowest cost?

☐ On-Demand

☐ Spot

☐ Dedicated Hosts

☐ Reserved

QUESTION 8

The company you work for has reshuffled teams a bit and you've been moved from the AWS IAM team to the AWS network team. One of your first assignments is to review the subnets in the main VPCs. You have recommended that the company add some private subnets and segregate databases from public traffic. What differentiates a public subnet from a private subnet?

☐ If a subnet's traffic is routed to an internet gateway, the subnet is known as a public subnet.

☐ Public subnets are associated with public Availability zones.

☐ Public subnets are meant to house EC2 instances with public IP addresses.

☐ A public subnet has a public IP address.

QUESTION 9

Your boss has tasked you with decoupling your existing web frontend from the backend. Both applications run on EC2 instances. After you investigate the existing architecture, you find that (on average) the backend resources are processing about 50,000 requests per second and will need something that supports their extreme level of message processing. It's also important that each request is processed only 1 time. What can you do to decouple these resources?

- ☐ Upsize your EC2 instances to reduce the message load on the backend servers.
- ☐ Use SQS FIFO to decouple the applications.
- ☐ Use S3 to store the messages being sent between the EC2 instances.
- ☐ Use SQS Standard. Include a unique ordering ID in each message, and have the backend application use this to deduplicate messages.

QUESTION 10

After an IT Steering Committee meeting, you have been put in charge of configuring a hybrid environment for the company's compute resources. You weigh the pros and cons of various technologies based on the requirements you are given. The main requirements to drive this selection are overall cost considerations and the ability to reuse existing internet connections. Which technology best meets these requirements?

- ☐ AWS Direct Gateway
- ☐ VPC Peering
- ☐ AWS Direct Connect
- ☐ AWS Managed VPN

QUESTION 11

A gaming company is creating an application which simply provides a leaderboard for specific games. The leaderboard will use DynamoDB for data, and simply needs to be updated in near real-time. An EC2 instance will be configured to house the application which will be accessed by subscribers from the Internet. Which step is **NOT** necessary for internet traffic to flow to and from the Internet?

- ☐ Ensure that your network access control lists and security group rules allow the relevant traffic to flow to and from your instance.
- ☐ Add a route to your subnet's route table that directs internet-bound traffic to the internet gateway.
- ☐ Attach an internet gateway to your VPC.
- ☐ A route in the route table to the DynamoDB table.

QUESTION 12

You work for an organization that has multiple AWS accounts in multiple regions and multiple applications. You have been tasked with making sure that all your firewall rules across these multiple accounts and regions are consistent. You need to do this as quickly and efficiently as possible. Which AWS service would help you achieve this?

- ☐ AWS Web Application Firewall (AWS WAF)
- ☐ AWS Network Firewall
- ☐ Amazon Detective
- ☐ AWS Firewall Manager

QUESTION 13

The company you work for has reshuffled teams a bit and you've been moved from the AWS IAM team to the AWS Network team. One of your first assignments is to review the subnets in the main VPCs. What are two key concepts regarding subnets?

(Choose 2)

☐ Every subnet you create is associated with the main route table for the VPC.

☐ Private subnets can only hold databases.

☐ Each subnet is associated with one security group.

☐ Each subnet maps to a single Availability Zone.

☐ A subnet spans all the Availability Zones in a Region.

QUESTION 14

Your application is housed on an Auto Scaling Group of EC2 instances. The application is backed by the Multi-AZ MySQL RDS database and an additional read replica. You need to simulate some failures for disaster recovery drills. Which event will NOT cause an RDS to perform a failover to the standby replica?

☐ Storage failure on primary

☐ Compute unit failure on primary

☐ Loss of network connectivity to primary

☐ Read replica failure

QUESTION 15

A consultant is hired by a small company to configure an AWS environment. The consultant begins working with the VPC and launching EC2 instances within the VPC. The initial instances will be placed in a public subnet. The consultant begins to create security groups. How many security groups can be attached to an EC2 instance?

- ☐ Instances in private subnets cannot have multiple security groups.
- ☐ You can assign two security groups to an instance.
- ☐ You can assign up to five security groups to the instance.
- ☐ You can only assign one security group to an instance.

QUESTION 16

You have been evaluating the NACLs in your company. Currently, you are looking at the default network ACL. Which statement is true regarding subnets and NACLs?

- ☐ You have to delete the default NACL before creating a custom NACL to associate with a subnet.
- ☐ Only public subnets can use the default NACL.
- ☐ The default NACL will always be associated with each subnet.
- ☐ Each subnet in your VPC must be associated with a network ACL. If you don't explicitly associate a subnet with a network ACL, the subnet is automatically associated with the default network ACL.

QUESTION 17

You work for an online cloud education provider that provides hands-on labs for training students. Recently, you noticed a spike in CPU activity for one of your EC2 instances and you suspect it is being used to mine bitcoin rather than for educational purposes. Somehow, your production environment has been compromised and you need to quickly identify the root cause of this compromise. Which AWS service would be best suited to identify the root cause?

☐ AWS Trusted Advisor

☐ Amazon Detective

☐ Amazon CloudWatch

☐ AWS Artifact

QUESTION 18

You have configured an Auto Scaling Group of EC2 instances fronted by an Application Load Balancer and backed by an RDS database. You want to begin monitoring the EC2 instances using CloudWatch metrics. Which metric is not readily available out of the box?

☐ DiskReadOps

☐ Memory utilization

☐ NetworkIn

☐ CPU utilization

QUESTION 19

A large, big-box hardware chain is setting up a new inventory management system. They have developed a system using IoT sensors which captures the removal of items from the store shelves in near real-time and want to use this information to update their inventory system. The company wants to analyze this data in the hopes of being ahead of demand and properly managing logistics and delivery of in-demand items.

Which AWS service can be used to capture this data as close to real-time as possible, while being able to both transform and load the streaming data into Amazon S3 or Elasticsearch?

☐ Amazon Kinesis Data Firehose

☐ Amazon Redshift

☐ Amazon Aurora

☐ Amazon Kinesis Data Streams

QUESTION 20

Bill is a cloud solutions architect for a small technology startup company. The company started out completely on-premises, but Bill has finally convinced them to explore shifting their application to AWS. The application is fairly complex and leverages message brokers that communicate using AMQP 1.0 protocols to exchange data between nodes and complete workloads.

Which service should Bill use to design the new AWS cloud-based architecture?

☐ Amazon SQS

☐ AWS Batch

☐ Amazon SNS

☐ Amazon MQ

QUESTION 21

You work in healthcare for an IVF clinic. You host an application on AWS, which allows patients to track their medication during IVF cycles. The application also allows them to view test results, which contain sensitive medical data. You have a regulatory requirement that the application is secure and you must use a firewall managed by AWS that enables control and visibility over VPC-to-VPC traffic and prevents the VPCs hosting your sensitive application resources from accessing domains using unauthorized protocols. What AWS service would support this?

☐ AWS Firewall Manager

☐ AWS Network Firewall

☐ AWS PrivateLink

☐ AWS WAF

QUESTION 22

Your company is storing highly sensitive data in S3 Buckets. The data includes personal and financial information. An audit has determined that this data must be stored in a secured manner and any data stored in the buckets already or data coming into the buckets must be analyzed and alerts sent out flagging improperly stored data. Which AWS service can be used to meet this requirement?

☐ AWS Inspector

☐ Amazon Macie

☐ AWS Trusted Advisor

☐ AWS GuardDuty

QUESTION 23

A previous cloud engineer deployed several Amazon EC2 instances within your AWS account. You have recently taken over control of the account and have noticed there is a significant amount of idle and underutilized instances in place. Hundreds of instances are within the account, and you do not have the time to go through each instance and manually check all of them.

Which AWS service allows you to kick off the collection of metrics and generate recommendations for incorrectly sized EC2 instances?

- ☐ AWS Compute Optimizer
- ☐ Amazon CloudWatch dashboards
- ☐ AWS Budgets
- ☐ AWS Cost and Usage Reports

QUESTION 24

A travel company has deployed a web application that serves travel updates to users all over the world. This application uses an Amazon RDS database, which is very read-heavy and can have performance issues at certain times of the year. What can you do to enhance performance and reduce the load on your source DB instance?

- ☐ Configure multi-Region RDS
- ☐ Add read replicas
- ☐ Place CloudFront in front of the Database.
- ☐ Configure RDS Multi-AZ

QUESTION 25

You are working as a Solutions Architect in a large healthcare organization. You have many Auto Scaling Groups that utilize launch configurations. Many of these launch configurations are similar yet have subtle differences. You'd like to use multiple versions of these launch configurations. An ideal approach would be to have a default launch configuration and then have additional versions that add additional features. Which option best meets these requirements?

- ☐ Create the launch configurations in CloudFormation and version the templates accordingly.
- ☐ Store the launch configurations in S3 and turn on versioning.
- ☐ Use launch templates instead.
- ☐ Simply create the needed versions. Launch configurations already have versioning.

QUESTION 26

An international company has many clients around the world. These clients need to transfer gigabytes to terabytes of data quickly and on a regular basis to an S3 bucket. Which S3 feature will enable these long distance data transfers in a secure and fast manner?

- ☐ AWS Snowmobile
- ☐ Multipart upload
- ☐ Cross-account replication
- ☐ Transfer Acceleration

QUESTION 27

A professional baseball league has chosen to use a key-value and document database for storage, processing, and data delivery. Many of the data requirements involve high-speed processing of data such as a Doppler radar system which samples the position of the baseball 2000 times per second. Which AWS data storage can meet these requirements?

- ☐ Redshift
- ☐ DynamoDB
- ☐ S3
- ☐ RDS

QUESTION 28

You have configured a VPC with both a public and a private subnet. You need to deploy a web server and a database. You want the web server to be accessed from the Internet by customers. Which is the proper configuration for this architecture?

- ☐ Database outside the VPC for decoupling from web server, and web server in public subnet for internet access.
- ☐ Web server in public subnet, database in private subnet.
- ☐ Web server outside of VPC for internet access, database in private subnet.
- ☐ Both web server and database in public subnets to facilitate internet access.

QUESTION 29

You are working for a large financial institution and have been tasked with creating a relational database solution to deal with a read-heavy workload. The database needs to be highly available within the Oregon region and quickly recover if an Availability Zone goes offline. Which of the following would you select to meet these requirements?

(Choose 2)

- ☐ Enable Multi-AZ support for the RDS database.
- ☐ Create a read replica and point your read workloads to the new endpoint RDS provides.
- ☐ Split your database into multiple RDS instances across different regions. In the event of a failure, point your application to the new region.
- ☐ Use an Amazon Aurora global database to ensure a region failure won't break the application.
- ☐ Using RDS, create a read replica. If an AZ fails, RDS will automatically cut over to the read replica.

QUESTION 30

You work for the national weather service and are in the process of doing some research on climate change. You have thousands of weather stations around the country that all have temperature sensors that report temperatures every 15 seconds. This time-series data is then fed to AWS where you need to run machine learning on it in real time. Which AWS service would best suit your requirements?

- ☐ Amazon Rekognition
- ☐ Amazon Forecast
- ☐ Amazon Transcribe
- ☐ Amazon Lex

QUESTION 31

An accounting company has big data applications for analyzing actuary data. The company is migrating some of its services to the cloud, and for the foreseeable future, will be operating in a hybrid environment. They need a storage service that provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. Which AWS service can meet these requirements?

☐ Glacier

☐ S3

☐ EBS

☐ EFS

QUESTION 32

A small development team with very limited AWS knowledge has begun the process of creating and deploying a new frontend application based on React within AWS. The application is simple and does not need any backend processing via traditional databases. The application does, however, require GraphQL interactions to complete the required processing of data. Which AWS service can the team use to complete this?

☐ Host the application in AWS Lambda instead and perform the processing using DynamoDB.

☐ Stand up a full stack application easily via AWS Amplify.

☐ Leverage API Gateway for any GraphQL calls. It supports GraphQL and REST API.

☐ Deploy a GraphQL interface via AWS AppSync.

QUESTION 33

Your company is storing stack traces for application errors in an S3 bucket. The engineers using these stack traces review them when addressing application issues. It has been decided that the files only need to be kept for four weeks; then, they must be purged. How can you meet this requirement in S3?

- ☐ Add an S3 Lifecycle rule to archive these files to Glacier after one month.
- ☐ Create a bucket policy to purge the rules after one month.
- ☐ Configure the S3 Lifecycle rules to purge the files after a month.
- ☐ Write a cron job to purge the files after one month.

QUESTION 34

You have configured an Auto Scaling Group of EC2 instances. You have begun testing the scaling of the Auto Scaling Group using a stress tool to force the CPU utilization metric being used to force scale out actions. The stress tool is also being manipulated by removing stress to force a scale in. But you notice that these actions are only taking place in five-minute intervals. What is happening?

- ☐ The Auto Scaling Group is following the default cooldown procedure.
- ☐ A load balancer is managing the load and limiting the effectiveness of stressing the servers.
- ☐ The stress tool is configured to run for five minutes.
- ☐ Auto Scaling Groups can only scale in intervals of five minutes or greater.