- **Vendor:** Amazon
- **Exam Code:** SAA-C02
- **Exam Name:** AWS Certified Solutions Architect - Associate (SAA-C02) Exam
- **New Updated Questions from Braindump2go (Updated in April/2021)**

QUESTION 637
A restaurant reservation application needs to access a waiting list.
When a customer tries to reserve a table, and none are available, the customer application will put the user on the waiting list, and the application will notify the customer when a table becomes free.
The waiting list must preserve the order in which customers were added to the waiting list.
Which service should the solutions architect recommend to store this waiting list?

A. Amazon Simple Notification Service (Amazon SNS)
B. AWS Step Functions invoking AWS Lambda functions
C. A FIFO queue in Amazon Simple Queue Service (Amazon SQS)
D. A standard queue in Amazon Simple Queue Service (Amazon SQS)

Answer: C

QUESTION 638
A solutions architect is designing an architecture that includes web application and database tiers. The web tier must be capable of auto scaling. The solutions architect has decided to separate each tier into its own subnets. The design includes two public subnets and four private subnets. The security team requires that tiers be able to communicate with each other only when there is a business need and that all other network traffic be blocked.
What should the solutions architect do to meet these requirements?

A. Create an Amazon GuardDuty source'destmation rule set to control communication
B. Create one security group for all tiers to limit traffic to only the required source and destinations
C. Create specific security groups for each tier to limit traffic to only the required source and destinations
D. Create network ACLs in all six subnets to limit traffic to the sources and destinations required for the application to function

Answer: D

QUESTION 639
A company has hired a solutions architect to design a reliable architecture for its application.
The application consists of one Amazon RDS DB instance and two manually provisioned Amazon EC2 instances that run web servers.
The EC2 instances are located in a single Availability Zone.
An employee recently deleted the DB instance and the application was unavailable for 24 hours as a result.
The company is concerned with the overall reliability of its environment.
What should the solutions architect do to maximize reliability of the application's infrastructure?

A. Delete one EC2 instance and enable termination protection on the other EC2 instance.

Update the DB instance to be Muto-AZ and enable deletion protection

B. Update the DB instance to be Multiple-AZ and enable deletion protection.
Place the EC2 instances behind an Application Load Balancer and run them in an EC2 Auto Seating group across multiple Availability Zones

C. Create an additional DB instance along with an Amazon API Gateway and an AWS Lambda function.
Configure the application to invoke the Lambda function through API Gateway.
Have the Lambda function write the data to the two DB instances

D. Place the EC2 instances in an EC2 Auto Scaling group that has multiple subnets located in multiple Availability Zones.
Use Spot Instances instead of On-Demand instances.
Set up Amazon CloudWatch alarms to monitor the health of the instances.
Update the DB instance to be Multi-AZ and enable deletion protection

Answer: B

QUESTION 640
company's human resources (HR) department saves its sensitive documents in an Amazon S3 bucket named confidential_bucket.
An IAM policy grants permission for ail S3 actions to a group of which each HR employee is a member.
A solutions architect needs to make the objects secure and accessible outside the company's AWS account and on-premises IP CIDR range.
The solutions architect adds the following S3 bucket policy:
("Version": "2008-10-17", "Statement": [
{ "Effect": "Deny", "Principal": { "AWS": --
"Action": "s3:"", "Resource": "arn:aws:s3:::confidential_bucket/*", "Condition": {
"StringNotLike": {
"aws:sourceVpce": "vpce-C12345789" }, "NotlpAddress": { "aws:SourceIp": [
"10.100.0.0/24", "172.31.0.0/24"
J } }
} J }
What is the effect of the added bucket policy?

A. The added bucket policy will have no effect because all the users from the HR department have full permissions in Amazon S3.

B. The bucket can be accessible only if the source IP is from CIDR blocks 10.100.0.0/24 or 172.31.0.0/24, or if it comes from services within the VPC that has endpoint vpce-0123456789 attached.

C. The bucket can be accessible only if the request is coming from the CIDR blocks 10.100.0.0/24 and 172.31.0.0/24, or if it comes from services within the VPC that has endpoint vpce-0123456789 attached.

D. The bucket can be accessible only if the request is coming from the CIDR blocks 10.100.0.0/24 or 172.31.0.0/24, and if it comes from services within the VPC that has endpoint vpce-0123456789 attached.

Answer: D

QUESTION 641
A company is planning to make a series of schema changes to tables on its Amazon Aurora DB cluster.
A solutions architect needs to test the changes in the most cost-effective manner possible.
What should the solutions architect do to meet these requirements?

A. Create a clone of the current Aurora DB cluster.
Perform the schema changes on the clone.
Once the changes are tested and performance is acceptable, apply the same changes on the original cluster.
Delete the clone.

B. Create an Amazon RDS for MySQL replica.
Perform the schema changes on the replica.
Once the changes are tested and performance is acceptable, apply the same changes on the replica.
Once the changes are tested and performance is acceptable, apply the same changes on the primary DB instance.
Delete the replica.

C. Create an additional Aurora Replica Perform the schema changes on the Aurora Replica.
Once the changes are tested and performance is acceptable, apply the same changes on the primary DB instance.
Delete the Aurora Replica.

D. Take a snapshot of the current Aurora DB cluster.
Restore the snapshot of the cluster to a new cluster.
Perform the schema changes on the restored cluster.
Once the changes are tested and performance is acceptable, apply the same changes on the origin cluster.
Delete the restored cluster.

Answer: A

QUESTION 642
An ecommerce application places orders in an Amazon Simple Queue Service (Amazon SQS) queue.
When a message is received, the Amazon EC2 worker instances process the request The EC2 instancesare in an Auto Scaling group 236.
How should the architecture be designed to scale the auto scaling group with the LEAST amount of operational overhead?

A. Use an Amazon CloudWatch alarm on the EC2 CPU to scale the Auto Scaling group up and down
B. Use an Amazon EC2 Auto Scaling health check for messages processed on the EC2 instances to scale up or down.
C. Use an Amazon CloudWatch alarm based on the number of messages in the queue to scale the Auto Scaling group up or down
D. Use an Amazon CloudWatch alarm based on the CPU to scale the Auto Scaling group up or down

Answer: C

QUESTION 643
A solutions architect is using an AWS Cloud Formation template to deploy a three-tier web application. The web application consists of a web tier and an application tier that stores and retrieves user data in Amazon DynamoDB tables. The web and application tiers are hosted on Amazon EC2 instances, and the database tier is not publicly accessible. The application EC2 instances need to access the DynamoDB tables without exposing API credentials in the template. What should the solutions architect do to meet these requirements?

A. Create an IAM role to read the DynamoOB tables.
Associate the role with the application instances by reference an instance profile

B. Create an IAM role that has the required permissions to read and write from the DynamoOB tables.
Add the role to the EC2 instance profile and associate the instance profile with the apphcanon instances

C. Use the parameter section in the AWS CkHidFormaton template to have the user input access and secret keys from an already-created IAM user mat has the required permissions to read and write from the DynamoOB tables

D. Create an IAM user m the AWS CioudFormation template that has the required permissions to read and write from the DynamoOB tables.
Use the GetAti function to retrieve the access and secret keys and pass them to the application instances through the user data

Answer: B

QUESTION 644
A company has an application that scans millions of connected devices for security threats and pushes the scan logs to an Amazon S3 bucket.
A total of 70 GB of data is generated each week, and the company needs to store 3 years of data for historical reporting.
The company must process, aggregate, and enrich the data from Amazon S3 by performing complex analytical queries and joins in the least amount of time.
The aggregated dataset is visualized on an Amazon QuickSight dashboard.
What should a solutions architect recommend to meet these requirements?

A. Create and run an ETL job in AWS Glue to process the data from Amazon S3 and load it into Amazon Redshift.
Perform the aggregation queries on Amazon Redshift.
B. Use AWS Lambda functions based on S3 PutObject event triggers to copy the incremental changes to Amazon DynamoDB.
Perform the aggregation queries on DynamoDB.
C. Use AWS Lambda functions based on S3 PutObject event triggers to copy the incremental changes to Amazon Aurora MySQL.
Perform the aggregation queries on Aurora MySQL
D. Use AWS Glue to catalog the data in Amazon S3.
Perform the aggregation queries on the cataloged tables by using Amazon Athena.
Query the data directly from Amazon S3.

Answer: A

QUESTION 645
A company is using various types of Amazon EC2 On-Demand Instances.
The company suspects that these instances have greater CPU and memory capacity than its workloads require.
Which actions should the company take to obtain recommendations to optimize cost? (Select TWO.)

A. Use AWS Trusted Advisor for instance type recommendations
B. Use AWS Compute Optimizer for instance type recommendations
C. Use AWS Budgets for instance type recommendations
D. Use Cost Explorer right sizing recommendations
E. Use Amazon Inspector to identify under utilized EC2 instances.

Answer: AD

QUESTION 646
A company receives data from different sources and implements multiple applications to consume this data.
There are many short-running jobs that run only on the weekend. The data arrives in batches rather than throughout the entire weekend.
The company needs an environment on AWS to ingest and process this data while maintaining the order of the transactions.
Which combination of AWS services meets these requirements in the MOST cost-effective manner?

A. Amazon Kinesis Data Streams with AWS Lambda
B. Amazon Kinesis Data Streams with Amazon EC2 Auto Scaling
C. Amazon Simple Queue Service (Amazon SQS) with AWS Lambda

D. Amazon Simple Queue Service (Amazon SQS) with Amazon EC2 Auto Scaling

Answer: C

QUESTION 647
A global company plans to track and store information about local allergens in an Amazon DynamoOB table and query this data from its website.
The company anticipates that website traffic will fluctuate.
The company estimates that the combined read and write capacity units will range from 10 10.000 per second, depending on the severity of the conditions for the given day.
A solutions architect must design a solution that avoids throttling issues and manages capacity efficiently.
What should the solutions architect do to meet these requirements MOST cost-effectively?

A. Use provisioned capacity mode.
Set the table's read capacity units to 10,000.
B. Use provisioned capacity mode and a scaling policy in DynamoDB auto scaling
C. Use on demand capacity made for a couple of months.
Then switch to provisioned capacity mode
D. Use on-demand capacity mode only.
Configure DynamoDB Accelerator (DAX) to be in front of the tab

Answer: B

QUESTION 648
A company is using various types of amazon EC-2 On-Demand instances.
The company suspects that these instances have greater cpu and memory capacity than its workloads require.
Which actions should the company take to obtain recommendation to optimize cost?

A. Use AWS Trusted Advisor for instance type recommendations.
B. Use AWS Compute Optimizer for instance type recommendations.
C. Use AWS Budgets for instance type recommendations.
D. Use Cost Explorer rightsizing recommendations.
E. Use Amazon Inspector to identify underutilized EC2 instances.

Answer: AD

QUESTION 649
A company has an AWS Direct Connect connection from its corporate data center to its VPC in the useast-1 Region.
The company recently acquired a corporation that has several VPCs and a Direct Connect connection between its on-premises data center and the eu-west-2 Region.
The CIDR blocks for the VPCs of the company and the corporation do not overlap. The company requires connectivity between two Regions and the data centers.
The company needs a solution that is scalable while reducing operational overhead.
What should a solutions architect do to meet these requirements?

A. Set up inter-Region VPC peering between the VPC m us-east-1 and the VPCs in eu-west-2
B. Create private virtual interfaces from the Direct Connect connection in us-east-1 to the VPCs in eu-west-2
C. Establish VPN appliances in a fully meshed VPN network hosted by Amazon EC2
Use AWS VPN CloudHub to send and receive data between the data centers and each VPC
D. Connect the existing Direct Connect connection to a Direct Connect gateway
Route traffic from the virtual private gateways of the VPCs in each Region to the Direct Connect gateway

Answer: D

QUESTION 650
A user owns a MySQL database that is accessed by various clients who expect, at most, 100 ms latency on requests.
Once a record is stored in the database, it is rarely changed. Clients only access one record al a time.
Database access has been increasing exponentially due to increased client demand.
The resultant load will soon exceed the capacity of the most expensive hardware available for purchase.
The user wants to migrate to AWS, and is willing to change database systems.
Which service would alleviate the database load issue and offer virtually unlimited scalability for the future?

A. Amazon RDS
B. Amazon DynamoDB
C. Amazon Redshift
D. AWS Data Pipeline

Answer: B

QUESTION 651
A company has several web servers that need to frequently access a common Amazon RDS MySQL Multi-AZ instance.
The company wants a secure method for the web servers to connect to the database while meeting a security requirement to rotate user credentials frequently.
A company has several web servers that need to frequently access a common Amazon ROS MySQL Muto-AZ DB instance.
The company wants a secure method for the web servers to connect to the database while meeting a security requirement to rotate user credentials frequently.
Which solution meets these requirements?

A. Store the database user credentials in AWS Secrets Manager.
Grant the necessary IAM permissions to allow the web servers to access AWS Secrets Manager
B. Store the database user credentials m AWS Systems Manager OpsCenter.
Grant the necessary IAM permissions to allow the web servers to access OpsCenter
C. Store the database user credentials in a secure Amazon S3 bucket.
Grant the necessary IAM permissions to allow the web servers to retrieve credentials and access the database
D. Store the database user credentials in fries encrypted with AWS Key Management Service (AWS KMS) on the web server file system.
The web server should be able to decrypt the files and access the database

Answer: A

QUESTION 652
A company has an asynchronous web application where Amazon API Gateway triggers AWS Lambda functions to perform write and update operations on an Amazon RDS DB instance.
During periods of extreme use API Gateway and Lambda scale in response to the incoming workload but service outages occur due to congestion with Amazon RDS.
The company is seeking a cost-effective design to alleviate this congestion.
What should a solutions architect recommend'?

A. implement RDS storage autoscaling with a larger instance type
B. Create read replicas to alleviate me read requests on the database
C. Use Amazon Kinesis to poll the incoming requests from API Gateway to the Lambda functions
D. Use Amazon Simple Queue Service (Amazon SQS) to buffer the incoming requests before delivering them to the Lambda functions

Answer: D

QUESTION 653
A solutions architect is designing a highly available website that is served by multiple web servers hosted outside of

AWS.
If an instance becomes unresponsive, the architect needs to remove it from the rotation.
What is the MOST efficient way to fulfill this requirement?

A. Use Amazon CloudWatch to monitor utilization.
B. Use Amazon API Gateway to monitor availability.
C. Use an Amazon Elastic Load Balancer.
D. Use Amazon Route 53 health checks.

Answer: C

QUESTION 654
A company has thousands of edge devices that collectively generate 1 TB of status alerts each day Each alert is approximately 2 KB in size.
A solutions architect needs to implement a solution to ingest and store the alerts for future analysis.
The company wants a highly available solution However the company needs to minimize costs and does not want to manage additional infrastructure.
Additionally, the company wants to keep 14 days of data available for immediate analysis and archive any data older than 14 days.
What is the MOST operationally efficient solution that meets these requirements?

A. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts.
Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon S3 bucket.
Set up an S3 LifecycJe configuration to transition data to Amazon S3 Glacier after 14 days
B. Launch Amazon EC2 instances across two Availability Zones and place them behind an Elastic Load Balancer to ingest the alerts.
Create a script on the EC2 instances that will store the alerts in an Amazon S3 bucket.
Set up an S3 Lifecycle configuration to transition data to Amazon S3 Glacier after 14 days
C. Create an Amazon Kinesis Data Firehose delivery stream to ingest the alerts.
Configure the Kinesis Data Firehose stream to deliver the alerts to an Amazon Elasticsearch Service (Amazon ES) cluster.
Set up the Amazon ES cluster to take manual snapshots every day and delete data from the cluster that is older than 14 days
D. Create an Amazon Simple Queue Service (Amazon SQS) standard queue to ingest the alerts and set the message retention period to 14 days.
Configure consumers to poll the SQS queue check the age of the message and analyze the message data as needed if the message is 14 days old, the consumer should copy the message to an Amazon S3 bucket and delete the message from the SQS queue

Answer: A

QUESTION 655
A company is using Amazon DynamoDB to stage its product catalog which is 1 GB.
Since a product entry on average consists of100 KB of data, and the average traffic is about 250 requests per second, the database administrator has provisioned 3.000 RCUs of read capacity throughput.
However, some products are very popular and users are experiencing delays or timeouts due to throttling.
What improvement offers a long-term solution to this problem?

A. Increase the throughput provisioning to 6.000 read capacity units (RCUs)
B. Use Amazon DynamoDB Accelerator to maintain the frequently read items
C. Augment Amazon DynamoDB by storing only the key product attributes, with the details stored on Amazon S3
D. Change the partition key to consist of a hash of product key and product type instead of just the product key

Answer: B

QUESTION 656
A company needs guaranteed Amazon EC2 capacity in three specific Availability Zones in a specific AWS Region for an upcoming event that will last 1 week.
What should the company do to guarantee the EC2 capacity?

A. Purchase Reserved Instances that specify the Region needed.
B. Create an On-Demand Capacity Reservation that specifies the Region needed.
C. Purchase Reserved Instances that specify the Region and three Availability Zones needed.
D. Create an On-Demand Capacity Reservation that specifies the Region and three Availability Zones needed.

Answer: A

QUESTION 657
A solutions architect is moving the static content from a public website hosted on Amazon EC2 instances to an Amazon S3 bucket.
An Amazon CloudFront distribution will be used to deliver the static assets.
The security group used by the EC2 instances restricts access to a limited set of IP ranges.
Access to the static content should be similarly restricted.
Which combination of steps will meet these requirements? (Select TWO.)

A. Create an origin access identity (OAI) and associate it with the distribution.
Change the permissions in the bucket policy so that only the OAI can read the objects.
B. Create an AWS WAF web ACL that includes the same IP restrictions that exist in the EC2 security group.
Associate this new web ACL with the CloudFront distribution.
C. Create a new security group that includes the same IP restrictions that exist in the current EC2 security group.
Associate this new security group with the CloudFront distribution.
D. Create a new security group that includes the same IP restrictions that exist in the current EC2 security group.
Associate this new security group with the S3 bucket hosting the static content.
E. Create a new IAM role and associate the role with the distribution.
Change the permissions either on the S3 bucket or on the files within the S3 bucket so that only the newly created IAM role has read and download permissions.

Answer: AB