



Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.



SAA-C02 Dumps
SAA-C02 Braindumps
SAA-C02 Real Questions
SAA-C02 Practice Test
SAA-C02 Actual Questions



Amazon

SAA-C02

AWS Certified Solutions Architect - Associate - 2024



<https://killexams.com/pass4sure/exam-detail/SAA-C02>

Question: 980

A company's web application is using multiple Linux Amazon EC2 instances and storing data on Amazon EBS volumes. The company is looking for a solution to increase the resiliency of the application in case of a failure and to provide storage that complies with atomicity, consistency, isolation, and durability (ACID).

What should a solutions architect do to meet these requirements?

- A. Launch the application on EC2 instances in each Availability Zone. Attach EBS volumes to each EC2 instance.
- B. Create an Application Load Balancer with Auto Scaling groups across multiple Availability Zones. Mount an instance store on each EC2 instance.
- C. Create an Application Load Balancer with Auto Scaling groups across multiple Availability Zones. Store data on Amazon EFS and mount a target on each instance.
- D. Create an Application Load Balancer with Auto Scaling groups across multiple Availability Zones. Store data using Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA).

Answer: C

Question: 981

A solutions architect is implementing a document review application using an Amazon S3 bucket for storage. The solution must prevent an accidental deletion of the documents and ensure that all versions of the documents are available. Users must be able to download, modify, and upload documents.

Which combination of actions should be taken to meet these requirements? (Choose two.)

- A. Enable a read-only bucket AC
- C. Enable versioning on the bucket.
- D. Attach an IAM policy to the bucket.
- E. Enable MFA Delete on the bucket.
- F. Encrypt the bucket using AWS KM

Answer: BD

Question: 982

A company hosts a static website on-premises and wants to migrate the website to AWS. The website should load as quickly as possible for users around the world. The company also wants the most cost-effective solution.

What should a solutions architect do to accomplish this?

- A. Copy the website content to an Amazon S3 bucket. Configure the bucket to serve static webpage content. Replicate the S3 bucket to multiple AWS Regions.
- B. Copy the website content to an Amazon S3 bucket. Configure the bucket to serve static webpage content. Configure Amazon CloudFront with the S3 bucket as the origin.
- C. Copy the website content to an Amazon EBS-backed Amazon EC2 instance running Apache HTTP Server. Configure Amazon Route 53 geolocation routing policies to select the closest origin.
- D. Copy the website content to multiple Amazon EBS-backed Amazon EC2 instances running Apache HTTP Server in multiple AWS Regions. Configure Amazon CloudFront geolocation routing policies to select the closest origin.

Answer: B

Question: 983

An application runs on Amazon EC2 instances across multiple Availability Zones. The instances run in an Amazon EC2 Auto Scaling group behind an Application Load Balancer. The application performs best when the CPU utilization of the EC2 instances is at or near 40%.

What should a solutions architect do to maintain the desired performance across all instances in the group?

- A. Use a simple scaling policy to dynamically scale the Auto Scaling group.
- B. Use a target tracking policy to dynamically scale the Auto Scaling group.
- C. Use an AWS Lambda function to update the desired Auto Scaling group capacity.
- D. Use scheduled scaling actions to scale up and scale down the Auto Scaling group.

Answer: B

Question: 984

A company's production application runs online transaction processing (OLTP) transactions on an Amazon RDS MySQL DB instance. The company is launching a new reporting tool that will access the same data. The reporting tool must be highly available and not impact the performance of the production application.

How can this be achieved?

- A. Create hourly snapshots of the production RDS DB instance.
- B. Create a Multi-AZ RDS Read Replica of the production RDS DB instance.
- C. Create multiple RDS Read Replicas of the production RDS DB instance. Place the Read Replicas in an Auto Scaling group.
- D. Create a Single-AZ RDS Read Replica of the production RDS DB instance. Create a second Single-AZ RDS Read Replica from the replica.

Answer: B

Question: 985

A company's legacy application is currently relying on a single-instance Amazon RDS MySQL database without encryption. Due to new compliance requirements, all existing and new data in this database must be encrypted.

How should this be accomplished?

- A. Create an Amazon S3 bucket with server-side encryption enabled. Move all the data to Amazon S3. Delete the RDS instance.
- B. Enable RDS Multi-AZ mode with encryption at rest enabled. Perform a failover to the standby instance.
- C. Take a Snapshot of the RDS instance. Create an encrypted copy of the snapshot. Restore the RDS instance from the encrypted snapshot.
- D. Create an RDS read replica with encryption at rest enabled. Promote the read replica to master and switch the over to the new master. Delete the old RDS instance.

Answer: C

Question: 986

Organizers for a global event want to put daily reports online as static HTML pages. The pages are expected to generate millions of views from users around the world. The files are stored in an Amazon S3 bucket. A solutions architect has been asked to design an efficient and effective solution.

Which action should the solutions architect take to accomplish this?

- A. Generate presigned URLs for the files.
- B. Use cross-Region replication to all Regions.
- C. Use the geoproximity feature of Amazon Route 53.
- D. Use Amazon CloudFront with the S3 bucket as its origin.

Answer: D

Question: 987

An application running on AWS uses an Amazon Aurora Multi-AZ deployment for its database. When evaluating performance metrics, a solutions architect discovered that the database reads are causing high I/O and adding latency to the write requests against the database.

What should the solutions architect do to separate the read requests from the write requests?

- A. Enable read-through caching on the Amazon Aurora database.
- B. Update the application to read from the Multi-AZ standby instance.
- C. Create a read replica and modify the application to use the appropriate endpoint.
- D. Create a second Amazon Aurora database and link it to the primary database as a read replica.

Answer: C

Question: 988

A solutions architect is designing a web application that will run on Amazon EC2 instances behind an Application Load Balancer (ALB). The company strictly requires that the application be resilient against malicious internet activity and attacks, and protect against new common vulnerabilities and exposures.

What should the solutions architect recommend?

- A. Leverage Amazon CloudFront with the ALB endpoint as the origin .
- B. Deploy an appropriate managed rule for AWS WAF and associate it with the AL
- D. Subscribe to AWS Shield Advanced and ensure common vulnerabilities and exposures are blocked.
- E. Configure network ACLs and security groups to allow only ports 80 and 443 to access the EC2 instances.

Answer: B

Explanation:

My answer here was C not A as I can see above .

the right answer can be either shield advanced or WAF. Basically shield advanced have WAF included . But it costs more cause it has also automatic remediation . so right answer according to cost is B

Question: 989

A solutions architect is designing the cloud architecture for a new application being deployed on AWS. The process should run in parallel while adding and removing application nodes as needed based on the number of jobs to be processed. The processor application is stateless. The solutions architect must ensure that the application is loosely coupled and the job items are durably stored.

Which design should the solutions architect use?

- A. Create an Amazon SNS topic to send the jobs that need to be processed. Create an Amazon Machine Image (AMI) that consists of the processor application. Create a launch configuration that uses the AMI.
- B. Create an Auto Scaling group using the launch configuration. Set the scaling policy for the Auto Scaling group to add and remove nodes based on CPU usage.
- C. Create an Amazon SQS queue to hold the jobs that need to be processed. Create an Amazon Machine Image (AMI) that consists of the processor application. Create a launch configuration that uses the AMI.
- D. Create an Auto Scaling group using the launch configuration. Set the scaling policy for the Auto Scaling group to add and remove nodes based on network usage.
- E. Create an Amazon SQS queue to hold the jobs that need to be processed. Create an Amazon Machine Image (AMI) that consists of the processor application. Create a launch template that uses the AMI.
- F. Create an Auto Scaling group using the launch template. Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue.
- G. Create an Amazon SNS topic to send the jobs that need to be processed. Create an Amazon Machine Image (AMI) that consists of the processor application. Create a launch template that uses the AMI.
- H. Create an Auto Scaling group using the launch template. Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of messages published to the SNS topic.

Answer: C

Question: 990

A solutions architect at an ecommerce company wants to back up application log data to Amazon S3. The solutions architect is unsure how frequently the logs will be accessed or which logs will be accessed the most. The company wants to keep costs as low as possible by using the appropriate S3 storage class.

Which S3 storage class should be implemented to meet these requirements?

- A. S3 Glacier
- B. S3 Intelligent-Tiering
- C. S3 Standard-Infrequent Access (S3 Standard-IA)
- D. S3 One Zone-Infrequent Access (S3 One Zone-IA)

Answer: B

Question: 991

A company has deployed an API in a VPC behind an internet-facing Application Load Balancer (ALB). An application that consumes the API as a client is deployed in a second account in private subnets behind a NAT gateway. When requests to the client application increase, the NAT gateway costs are higher than expected. A solutions architect has configured the ALB to be internal.

Which combination of architectural changes will reduce the NAT gateway costs? (Choose two.)

- A. Configure a VPC peering connection between the two VPCs. Access the API using the private address.
- B. Configure an AWS Direct Connect connection between the two VPCs. Access the API using the private address.
- C. Configure a ClassicLink connection for the API into the client VPC.
- D. Access the API using the ClassicLink address.
- E. Configure a PrivateLink connection for the API into the client VPC.
- F. Access the API using the PrivateLink address.
- G. Configure an AWS Resource Access Manager connection between the two accounts. Access the API using the private address.

Answer: AD

Question: 992

A solutions architect is designing a new service behind Amazon API Gateway. The request patterns for the service will be unpredictable and can change suddenly from 0 requests to over 500 per second. The total size of the data that needs to be persisted in a database is currently less than 1 GB with unpredictable future growth. Data can be queried using simple key-value requests.

Which combination of AWS services would meet these requirements? (Choose two.)

- A. AWS Fargate
- B. AWS Lambda
- C. Amazon DynamoDB
- D. Amazon EC2 Auto Scaling
- E. MySQL-compatible Amazon Aurora

Answer: BC

Question: 993

A solutions architect needs to ensure that API calls to Amazon DynamoDB from Amazon EC2 instances in a VPC do not traverse the internet.

What should the solutions architect do to accomplish this? (Choose two.)

- A. Create a route table entry for the endpoint.
- B. Create a gateway endpoint for DynamoDB
- D. Create a new DynamoDB table that uses the endpoint.
- E. Create an ENI for the endpoint in each of the subnets of the VPC
- G. Create a security group entry in the default security group to provide access.

Answer: AB

Question: 994

A company has been storing analytics data in an Amazon RDS instance for the past few years. The company asked a solutions architect to find a solution that allows users to access this data using an API. The expectation is that the application will experience periods of inactivity but could receive bursts of traffic within seconds.

Which solution should the solutions architect suggest?

- A. Set up an Amazon API Gateway and use Amazon EC2
- C. Set up an Amazon API Gateway and use AWS Elastic Beanstalk.
- D. Set up an Amazon API Gateway and use AWS Lambda functions.
- E. Set up an Amazon API Gateway and use Amazon EC2 with Auto Scaling.

Answer: C

Question: 995

A company allows its developers to attach existing IAM policies to existing IAM roles to enable faster experimentation and agility.

However, the security operations team is concerned that the developers could attach the existing administrator policy, when would allow the developers to circumvent any other security policies.

How should a solutions architect address this issue?

- A. Create an Amazon SNS topic to send an alert every time a developer creates a new policy.
- B. Use service control policies to disable IAM activity across all accounts in the organizational unit.
- C. Prevent the developers from attaching any policies and assign all IAM duties to the security operations team.
- D. Set an IAM permissions boundary on the developer IAM role that explicitly denies attaching the administrator policy.

Answer: D

Question: 996

A company is planning to use Amazon S3 to store images uploaded by its users. The images must be encrypted at rest in Amazon S3. The company does not want to spend time managing and rotating the keys, but it does want to control who can access those keys.

What should a solutions architect use to accomplish this?

- A. Server-Side Encryption with keys stored in an S3 bucket
- B. Server-Side Encryption with Customer-Provided Keys (SSE-C)
- C. Server-Side Encryption with Amazon S3-Managed Keys (SSE-S3)
- D. Server-Side Encryption with AWS KMS-Managed Keys (SSE-KMS)

Answer: D

Question: 997

A company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. On the first day of every month at midnight, the application becomes much slower when the month-end financial calculation batch executes. This causes the CPU utilization of the EC2 instances to immediately peak to 100%, which disrupts the application.

What should a solutions architect recommend to ensure the application is able to handle the workload and avoid downtime?

- A. Configure an Amazon CloudFront distribution in front of the AL
- C. Configure an EC2 Auto Scaling simple scaling policy based on CPU utilization.
- D. Configure an EC2 Auto Scaling scheduled scaling policy based on the monthly schedule.
- E. Configure Amazon ElastiCache to remove some of the workload from the EC2 instances.

Answer: C

Question: 998

A company is migrating from an on-premises infrastructure to the AWS Cloud. One of the company's applications stores files on a Windows file server farm that uses Distributed File System Replication (DFS) to keep data in sync. A solutions architect needs to replace the file server farm.

Which service should the solutions architect use?

- A. Amazon EFS
- B. Amazon FSx
- C. Amazon S3
- D. AWS Storage Gateway

Answer: B

Question: 999

An Amazon EC2 administrator created the following policy associated with an IAM group containing several users:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:TerminateInstances",
      "Resources": "*",
      "Condition": {
        "IpAddress": {
          "aws:SourceIp": "10.100.100.0/24"
        }
      }
    },
    {
      "Effect": "Deny",
      "Action": "ec2:*",
      "Resources": "*",
      "Condition": {
        "StringNotEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    }
  ]
}
```

What is the effect of this policy?

- A. Users can terminate an EC2 instance in any AWS Region except us-east-1.
- B. Users can terminate an EC2 instance with the IP address 10.100.100.1 in the us-east-1 Region.
- C. Users can terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100.100.254.
- D. Users cannot terminate an EC2 instance in the us-east-1 Region when the user's source IP is 10.100.100.254.

Answer: A

Question: 1000

A data science team requires storage for nightly log processing. The size and number of logs is unknown and will persist for 24 hours only.

What is the MOST cost-effective solution?

A. Amazon S3 Glacier

B. Amazon S3 Standard

C. Amazon S3 Intelligent-Tiering

D. Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA)

Answer: B



SAMPLE QUESTIONS

*These questions are for demo purpose only. **Full version** is up to date and contains actual questions and answers.*

Killexams.com is an online platform that offers a wide range of services related to certification exam preparation. The platform provides actual questions, exam dumps, and practice tests to help individuals prepare for various certification exams with confidence. Here are some key features and services offered by Killexams.com:

Actual Exam Questions: *Killexams.com provides actual exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these actual questions, candidates can familiarize themselves with the content and format of the real exam.*

Exam Dumps: *Killexams.com offers exam dumps in PDF format. These dumps contain a comprehensive collection of questions and answers that cover the exam topics. By using these dumps, candidates can enhance their knowledge and improve their chances of success in the certification exam.*

Practice Tests: *Killexams.com provides practice tests through their desktop VCE exam simulator and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice tests cover a wide range of questions and enable candidates to identify their strengths and weaknesses.*

Guaranteed Success: *Killexams.com offers a success guarantee with their exam dumps. They claim that by using their materials, candidates will pass their exams on the first attempt or they will refund the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exams.*

Updated Content: *Killexams.com regularly updates its question bank and exam dumps to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.*

Technical Support: *Killexams.com provides free 24x7 technical support to assist candidates with any queries or issues they may encounter while using their services. Their certified experts are available to provide guidance and help candidates throughout their exam preparation journey.*

For More exams visit <https://killexams.com/vendors-exam-list>
Kill your exam at First Attempt....Guaranteed!