

Amazon

SAA-C01 Exam

Amazon AWS Certified Solutions Architect Associate Exam

**Questions & Answers
Demo**

Version: 19.0

Question: 1

A Solutions Architect is designing an application that will encrypt all data in an Amazon Redshift cluster. Which action will encrypt the data at rest?

- A. Place the Redshift cluster in a private subnet.
- B. Use the AWS KMS Default Customer master key.
- C. Encrypt the Amazon EBS volumes.
- D. Encrypt the data using SSL/TLS.

Answer: B

Explanation:

https://docs.aws.amazon.com/en_pv/redshift/latest/mgmt/security-server-side-encryption.html

Question: 2

A website experiences unpredictable traffic. During peak traffic times, the database is unable to keep up with the write request. Which AWS service will help decouple the web application from the database?

- A. Amazon SQS
- B. Amazon EFS
- C. Amazon S3
- D. AWS Lambda

Answer: A

Question: 3

A legacy application needs to interact with local storage using iSCSI. A team needs to design a reliable storage solution to provision all new storage on AWS. Which storage solution meets the legacy application requirements?

- A. AWS Snowball storage for the legacy application until the application can be rearchitected.
- B. AWS Storage Gateway in cached mode for the legacy application storage to write data to Amazon S3.
- C. AWS Storage Gateway in stored mode for the legacy application storage to write data to Amazon S3.
- D. An Amazon S3 volume mounted on the legacy application server locally using the File Gateway service.

Answer: B

Explanation:

<https://docs.aws.amazon.com/storagegateway/latest/userguide/StorageGatewayConcepts.html>

Cached volumes – You store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally. Cached volumes offer a substantial cost savings on primary storage and minimize the need to scale your storage on-premises. You also retain low-latency access to your frequently accessed data. Stored volumes – If you need low-latency access to your entire dataset, first configure your on-premises gateway to store all your data locally. Then asynchronously back up point-in-time snapshots of this data to Amazon S3. This configuration provides durable and inexpensive offsite backups that you can recover to your local data center or Amazon EC2. For example, if you need replacement capacity for disaster recovery, you can recover the backups to Amazon EC2
<https://docs.aws.amazon.com/storagegateway/latest/userguide/WhatIsStorageGateway.html>

Question: 4

A Solutions Architect is designing an architecture for a mobile gaming application. The application is expected to be very popular. The Architect needs to prevent the Amazon RDS MySQL database from becoming a bottleneck due to frequently accessed queries. Which service or feature should the Architect add to prevent a bottleneck?

- A. Multi-AZ feature on the RDS MySQL Database
- B. ELB Classic Load Balancer in front of the web application tier
- C. Amazon SQS in front of RDS MySQL Database
- D. Amazon ElastiCache in front of the RDS MySQL Database

Answer: D

Explanation:

Elasticache (Redis and Memcached) is an in-memory cache for RDS DB instances and it helps improve performance by diverting frequently accessed read queries to the elasticache.

Question: 5

A company is launching an application that it expects to be very popular. The company needs a database that can scale with the rest of the application. The schema will change frequently. The application cannot afford any downtime for database changes. Which AWS service allows the company to achieve these objectives?

- A. Amazon Redshift
- B. Amazon DynamoDB
- C. Amazon RDS MySQL
- D. Amazon Aurora

Answer: B

Explanation:

Amazon DynamoDB transactions simplify the developer experience of making coordinated, all-or-nothing changes to multiple items both within and across tables. Transactions provide atomicity, consistency, isolation, and durability (ACID) in DynamoDB, enabling you to maintain data correctness in your applications easily. Refer pdf -Page 629 <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/dynamodb-dg.pdf>

Question: 6

A Solution Architect is designing a disaster recovery solution for a 5 TB Amazon Redshift cluster. The recovery site must be at least 500 miles (805 kilometers) from the live site. How should the Architect meet these requirements?

- A. Use AWS CloudFormation to deploy the cluster in a second region.
- B. Take a snapshot of the cluster and copy it to another Availability Zone.
- C. Modify the Redshift cluster to span two regions.
- D. Enable cross-region snapshots to a different region.

Answer: D

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2019/10/amazon-redshift-improves-performance-of-inter-region-snapshot-transfers/>

Question: 7

A customer has written an application that uses Amazon S3 exclusively as a data store. The application works well until the customer increases the rate at which the application is updating information. The customer now reports that outdated data occasionally appears when the application accesses objects in Amazon S3. What could be the problem, given that the application logic is otherwise correct?

- A. The application is reading parts of objects from Amazon S3 using a range header.
- B. The application is reading objects from Amazon S3 using parallel object requests.
- C. The application is updating records by writing new objects with unique keys.
- D. The application is updating records by overwriting existing objects with the same keys.

Answer: D

Explanation:

Amazon S3 offers eventual consistency for overwrite PUTS and DELETES in all Regions. Updates to a single key are atomic. For example, if you PUT to an existing key, a subsequent read might return the old data or the updated data, but it never returns corrupted or partial data. Amazon S3 achieves high availability by replicating data across multiple servers within AWS data centers. If a PUT request is successful, your data is safely stored. However, information about the changes must replicate across

Amazon S3, which can take some time, and so you might observe the following behaviors: A process writes a new object to Amazon S3 and immediately lists keys within its bucket. Until the change is fully propagated, the object might not appear in the list.

<https://docs.aws.amazon.com/AmazonS3/latest/dev/optimizing-performance.html>

<https://docs.aws.amazon.com/AmazonS3/latest/dev/Introduction.html>

Question: 8

A Solutions Architect is designing a new social media application. The application must provide a secure method for uploading profile photos. Each user should be able to upload a profile photo into a shared storage location for one week after their profile is created.

Which approach will meet all of these requirements?

- A. Use Amazon Kinesis with AWS CloudTrail for auditing the specific times when profile photos are uploaded.
- B. Use Amazon EBS volumes with IAM policies restricting user access to specific time periods.
- C. Use Amazon S3 with the default private access policy and generate presigned URLs each time a new site profile is created.
- D. Use Amazon CloudFront with AWS CloudTrail for auditing the specific times when profile photos are uploaded.

Answer: C

Question: 9

An application requires block storage for file updates. The data is 500 GB and must continuously sustain 100 MiB/s of aggregate read/write operations.

Which storage option is appropriate for this application?

- A. Amazon S3
- B. Amazon EFS
- C. Amazon EBS
- D. Amazon Glacier

Answer: C

Explanation:

<https://aws.amazon.com/efs/when-to-choose-efs/>

<https://docs.aws.amazon.com/efs/latest/ug/performance.html>

Question: 10

A mobile application serves scientific articles from individual files in an Amazon S3 bucket. Articles older than 30 days are rarely read. Articles older than 60 days no longer need to be available through the

application, but the application owner would like to keep them for historical purposes. Which cost-effective solution BEST meets these requirements?

- A. Create a Lambda function to move files older than 30 days to Amazon EBS and move files older than 60 days to Amazon Glacier.
- B. Create a Lambda function to move files older than 30 days to Amazon Glacier and move files older than 60 days to Amazon EBS.
- C. Create lifecycle rules to move files older than 30 days to Amazon S3 Standard Infrequent Access and move files older than 60 days to Amazon Glacier.
- D. Create lifecycle rules to move files older than 30 days to Amazon Glacier and move files older than 60 days to Amazon S3 Standard Infrequent Access.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/archive-s3-to-glacier/>