

➤ **Vendor: Amazon**

➤ **Exam Code: SAA-C02**

➤ **Exam Name: AWS Certified Solutions Architect - Associate (SAA-C02) Exam**

➤ **New Updated Questions from [Braindump2go](#) (Updated in [Auguts/2020](#))**

Visit Braindump2go and Download Full Version SAA-C02 Exam Dumps

QUESTION 274

A company is building a media-sharing application and decides to use Amazon S3 for storage.

When a media file is uploaded the company starts a multi-step process to create thumbnails, identify objects in the images, transcode videos into standard formats and resolutions and extract and store the metadata to an Amazon DynamoDB table.

The metadata is used for searching and navigation. The amount of traffic is variable. The solution must be able to scale to handle spikes in load without unnecessary expenses.

What should a solutions architect recommend to support this workload?

- A. Build the processing into the website or mobile app used to upload the content to Amazon S3. Save the required data to the DynamoDB table when the objects are uploaded
- B. Trigger AWS Step Functions when an object is stored in the S3 bucket. Have the Step Functions perform the steps needed to process the object and then write the metadata to the DynamoDB table
- C. Trigger an AWS Lambda function when an object is stored in the S3 bucket. Have the Lambda function start AWS Batch to perform the steps to process the object. Place the object data in the DynamoDB table when complete
- D. Trigger an AWS Lambda function to store an initial entry in the DynamoDB table when an object is uploaded to Amazon S3. Use a program running on an Amazon EC2 instance in an Auto Scaling group to poll the index for unprocessed use the program to perform the processing

Answer: D

QUESTION 275

A company has recently updated its internal security standards.

The company must now ensure all Amazon S3 buckets and Amazon Elastic Block Store (Amazon EBS) volumes are encrypted with keys created and periodically rotated by internal security specialists.

The company is looking for a native, software-based AWS service to accomplish this goal.

What should a solutions architect recommend as a solution?

- A. Use AWS Secrets Manager with customer master keys (CMKs) to store master key material and apply a routine to create a new CMK periodically and replace it in AWS Secrets Manager.
- B. Use AWS Key Management Service (AWS KMS) with customer master keys (CMKs) to store master key material and apply a routine to re-create a new key periodically and replace it in AWS KMS.
- C. Use an AWS CloudHSM cluster with customer master keys (CMKs) to store master key material and apply a routine to re-create a new key periodically and replace it in the CloudHSM cluster nodes.
- D. Use AWS Systems Manager Parameter Store with customer master keys (CMKs) keys to store

[SAA-C02 Exam Dumps](#) [SAA-C02 Exam Questions](#) [SAA-C02 PDF Dumps](#) [SAA-C02 VCE Dumps](#)

<https://www.braindump2go.com/saa-c02.html>

master key material and apply a routine to re-create a new periodically and replace it in the Parameter Store.

Answer: B

QUESTION 276

A solution architect must design a solution that uses Amazon CloudFront with an Amazon S3 to store a static website. The company security policy requires that all websites traffic be inspected by AWS WAF. How should the solution architect company with these requirements?

- A. Configure an S3 bucket policy to accept requests coming from the AWS WAF Amazon Resource Name (ARN) only
- B. Configure Amazon CloudFront to forward all incoming requests to AWS WAF before requesting content from the S3 origin,
- C. Configure a security group that allows Amazon CloudFront IP addresses to access Amazon S3 only Associate AWS WAF to CloudFront.
- D. Configure Amazon CloudFront and Amazon S3 to use an origin access identity (OAI) to restrict access to the S3 bucket. Enable AWS WAF on the distribution.

Answer: B

QUESTION 277

A company has copied 1 PB of data from a colocation facility to an Amazon S3 bucket in the us-east-1 Region using an AWS Direct Connect link.

The company now wants to copy the data to another S3 bucket in the us-west-2 Region.

The colocation facility does not allow the use AWS Snowball.

What should a solutions architect recommend to accomplish this?

- A. Order a Snowball Edge device to copy the data from one Region to another Region.
- B. Transfer contents from the source S3 bucket to a target S3 bucket using the S3 console.
- C. Use the aws S3 sync command to copy data from the source bucket to the destination bucket.
- D. Add a cross-Region replication configuration to copy objects across S3 buckets in different Reg.

Answer: A

QUESTION 278

A company has hired a new cloud engineer who should not have access to an Amazon S3 bucket named Company Confidential.

The cloud engineer must be able to read from and write to an S3 bucket called AdminTools.

Which IAM policy will meet these requirements?

A.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3:::AdminTools"
    },
    {
      "Effect": "Allow",
      "Action": [ "s3:GetObject", "s3:PutObject" ],
      "Resource": "arn:aws:s3:::AdminTools/*"
    },
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3:::CompanyConfidential/*",
        "arn:aws:s3:::CompanyConfidential"
      ]
    }
  ]
}
```

B.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:ListBucket",
      "Resource": [
        "arn:aws:s3:::AdminTools",
        "arn:aws:s3:::CompanyConfidential/*"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [ "s3:GetObject", "s3:PutObject", "s3:DeleteObject" ],
      "Resource": "arn:aws:s3:::AdminTools/*"
    },
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": "arn:aws:s3:::CompanyConfidential"
    }
  ]
}
```

C.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [ "s3:GetObject", "s3:PutObject" ],
      "Resource": "arn:aws:s3:::AdminTools/*"
    },
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3:::CompanyConfidential/*",
        "arn:aws:s3:::CompanyConfidential"
      ]
    }
  ]
}
```

D.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3::AdminTools/*"
    },
    {
      "Effect": "Allow",
      "Action": [ "s3:GetObject", "s3:PutObject", "s3:DeleteObject" ],
      "Resource": "arn:aws:s3::AdminTools/*"
    },
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3::CompanyConfidential",
        "arn:aws:s3::CompanyConfidential/*",
        "arn:aws:s3::AdminTools/*"
      ]
    }
  ]
}
```

Answer: A**QUESTION 279**

An engineering team is developing and deploying AWS Lambda functions.

The team needs to create roles and manage policies in AWS IAM to configure the permissions of the Lambda functions.

How should the permissions for the team be configured so they also adhere to the concept of least privilege?

- A. Create an IAM role with a managed policy attached.
Allow the engineering team and the Lambda functions to assume this role
- B. Create an IAM group for the engineering team with an IAMFullAccess policy attached.
Add all the users from the team to this IAM group
- C. Create an execution role for the Lambda functions.
Attach a managed policy that has permission boundaries specific to these Lambda functions
- D. Create an IAM role with a managed policy attached that has permission boundaries specific to the Lambda functions.
Allow the engineering team to assume this role.

Answer: A**QUESTION 280**

A company needs a secure connection between its on-premises environment and AWS.

This connection does not need high bandwidth and will handle a small amount of traffic.

The connection should be set up quickly.

What is the MOST cost-effective method to establish this type of connection?

- A. Implement a client VPN
- B. Implement AWS Direct Connect
- C. Implement a bastion host on Amazon EC2 53D.
- D. Implement an AWS Site-to-Site VPN connection.

Answer: D**QUESTION 281**

A company is building a payment application that must be highly available even during regional service disruptions.

A solutions architect must design a data storage solution that can be easily replicated and used in other AWS Regions.

The application also requires low-latency atomicity, consistency, isolation, and durability (ACID) transactions that need

[SAA-C02 Exam Dumps](#) [SAA-C02 Exam Questions](#) [SAA-C02 PDF Dumps](#) [SAA-C02 VCE Dumps](#)

<https://www.braindump2go.com/saa-c02.html>

to be immediately available to generate reports.
The development team also needs to use SQL.
Which data storage solution meets these requirements?

- A. Amazon Aurora Global Database
- B. Amazon DynamoDB global tables
- C. Amazon S3 with cross-Region replication and Amazon Athena
- D. MySQL on Amazon EC2 instances with Amazon Elastic Block Store (Amazon EBS) snapshot replication

Answer: C

QUESTION 282

A solutions architect is using Amazon S3 to design the storage architecture of a new digital media application. The media files must be resilient to the loss of an Availability Zone. Some files are accessed frequently while other files are rarely accessed in an unpredictable pattern. The solutions architect must minimize the costs of storing and retrieving the media files. Which storage option meets these requirements?

- A. S3 Standard
- 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 283

A company uses a legacy on-premises analytics application that operates on gigabytes of csv files and represents months of data. The legacy application cannot handle the growing size of csv files. New csv files are added daily from various data sources to a central on-premises storage location. The company wants to continue to support the legacy application while users learn AWS analytics services. To achieve this, a solutions architect wants to maintain two synchronized copies of all the csv files on-premises and in Amazon S3. Which solution should the solutions architect recommend?

- A. Deploy AWS DataSync on-premises.
Configure DataSync to continuously replicate the csv files between the company's on-premises storage and the company's S3 bucket
- B. Deploy an on-premises file gateway.
Configure data sources to write the csv files to the file gateway.
Point the legacy analytics application to the file gateway.
The file gateway should replicate the csv files to Amazon S3
- C. Deploy an on-premises volume gateway.
Configure data sources to write the csv files to the volume gateway.
Point the legacy analytics application to the volume gateway.
The volume gateway should replicate data to Amazon S3.
- D. Deploy AWS DataSync on-premises.
Configure DataSync to continuously replicate the csv files between on-premises and Amazon Elastic File System (Amazon EFS).
Enable replication from Amazon EFS to the company's S3 bucket.

Answer: B

QUESTION 284

An application allows users at a company's headquarters to access product data.

[SAA-C02 Exam Dumps](#) [SAA-C02 Exam Questions](#) [SAA-C02 PDF Dumps](#) [SAA-C02 VCE Dumps](#)

<https://www.braindump2go.com/saa-c02.html>

The product data is stored in an Amazon RDS MySQL DB instance.

The operations team has isolated an application performance slowdown and wants to separate read traffic from write traffic.

A solutions architect needs to optimize the application's performance quickly.

What should the solutions architect recommend?

- A. Change the existing database to a Multi-AZ deployment.
Serve the read requests from the primary Availability Zone.
- B. Change the existing database to a Multi-AZ deployment.
Serve the read requests from the secondary Availability Zone.
- C. Create read replicas for the database.
Configure the read replicas with half of the compute and storage resources as the source database.
- D. Create read replicas for the database.
Configure the read replicas with the same compute and storage resources as the source database.

Answer: D