

- **Vendor: Amazon**
- **Exam Code: SAA-C02**
- **Exam Name: AWS Certified Solutions Architect - Associate (SAA-C02) Exam**
- **New Updated Questions from [Braindump2go](#)**
- **(Updated in [February/2022](#))**

### **Visit Braindump2go and Download Full Version SAA-C02 Exam Dumps**

#### **QUESTION 911**

A company is implementing a new business application. The application runs on two Amazon EC2 instances and uses an Amazon S3 bucket for document storage. A solutions architect needs to ensure that the EC2 instances can access the S3 bucket.

What should the solutions architect do to meet this requirement?

- A. Create an IAM role that grants access to the S3 bucket.  
Attach the role to the EC2 Instances.
- B. Create an IAM policy that grants access to the S3 bucket.  
Attach the policy to the EC2 Instances
- C. Create an IAM group that grants access to the S3 bucket.  
Attach the group to the EC2 instances
- D. Create an IAM user that grants access to the S3 bucket.  
Attach the user account to the EC2 Instances

**Answer: C**

#### **QUESTION 912**

A company runs an on-premises application that is powered by a MySQL database. The company is migrating the application to AWS to increase the application's elasticity and availability. The current architecture shows heavy read activity on the database during times of normal operation. Every 4 hours the company's development team pulls a full export of the production database to populate a database in the staging environment. During this period, users experience unacceptable application latency. The development team is unable to use the staging environment until the procedure completes.

A solutions architect must recommend replacement architecture that alleviates the application latency issue. The replacement architecture also must give the development team the ability to continue using the staging environment without delay.

Which solution meets these requirements?

- A. Use Amazon Aurora MySQL with Multi-AZ Aurora Replicas for production.  
Populate the staging database by implementing a backup and restore process that uses the mysqldump utility.
- B. Use Amazon Aurora MySQL with Multi-AZ Aurora Replicas for production.  
Use database cloning to create the staging database on-demand
- C. Use Amazon RDS for MySQL with a Multi-AZ deployment and read replicas for production.  
Use the standby instance for the staging database.
- D. Use Amazon RDS for MySQL with a Multi-AZ deployment and read replicas for production.

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

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

Populate the staging database by implementing a backup and restore process that uses the mysqldump utility.

**Answer: C**

**QUESTION 913**

The DNS provider that hosts a company's domain name records is experiencing outages that cause service disruption for a website running on AWS. The company needs to migrate to a more resilient managed DNS service and wants the service to run on AWS.

What should a solutions architect do to rapidly migrate the DNS hosting service?

- A. Create an Amazon Route 53 public hosted zone for the domain name.  
Import the zone file containing the domain records hosted by the previous provider.
- B. Create an Amazon Route 53 private hosted zone for the domain name.  
Import the zone file containing the domain records hosted by the previous provider
- C. Create a Simple AD directory in AWS. Enable zone transfer between the DNS provider and AWS Directory Service for Microsoft Active Directory for the domain records.
- D. Create an Amazon Route 53 Resolver inbound endpoint in the VPC.  
Specify the IP addresses that the provider's DNS will forward DNS queries to Configure the provider's DNS to forward DNS queries for the domain to the IP addresses that are specified in the inbound endpoint.

**Answer: B**

**QUESTION 914**

A company uses Amazon API Gateway to manage its REST APIs that third-party service providers access. The company must protect the REST APIs from SQL injection and cross site scripting attacks.

What is the MOST operationally efficient solution that meets these requirements?

- A. Configure AWS Shield
- B. Configure AWS WAF
- C. Set up API Gateway with an Amazon CloudFront distribution.  
Configure AWS Shield in CloudFront
- D. Set up API Gateway with an Amazon CloudFront distribution.  
Configure AWS WAF in CloudFront.

**Answer: B**

**QUESTION 915**

A company hosts its web application on AWS using seven Amazon EC2 instances. The company requires that the IP addresses of all healthy EC2 instances be returned in response to DNS queries. Which policy should be used to meet this requirement?

- A. Simple routing policy
- B. Latency routing policy
- C. Multivalue routing policy
- D. Geolocation routing policy

**Answer: C**

**Explanation:**

<https://aws.amazon.com/premiumsupport/knowledge-center/multivalue-versus-simple-policies/>

"Use a multivalue answer routing policy to help distribute DNS responses across multiple resources.

For example, use multivalue answer routing when you want to associate your routing records with a Route 53 health check."

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policy-multivalue>

**QUESTION 916**

A company is hosting a web application on AWS using a single Amazon EC2 instance that stores user-uploaded documents in an Amazon EBS volume. For better scalability and availability, the company duplicated the architecture and created a second EC2 instance and EBS volume in another Availability Zone placing both behind an Application Load Balancer. After completing this change, users reported that, each time they refreshed the website, they could see one subset of their documents or the other, but never all of the documents at the same time. What should a solutions architect propose to ensure users see all of their documents at once?

- A. Copy the data so both EBS volumes contain all the documents.
- B. Configure the Application Load Balancer to direct a user to the server with the documents
- C. Copy the data from both EBS volumes to Amazon EFS.  
Modify the application to save new documents to Amazon EFS
- D. Configure the Application Load Balancer to send the request to both servers.  
Return each document from the correct server.

**Answer: C**

**Explanation:**

Amazon EFS provides file storage in the AWS Cloud. With Amazon EFS, you can create a file system, mount the file system on an Amazon EC2 instance, and then read and write data to and from your file system. You can mount an Amazon EFS file system in your VPC, through the Network File System versions 4.0 and 4.1 (NFSv4) protocol. We recommend using a current generation Linux NFSv4.1 client, such as those found in the latest Amazon Linux, Redhat, and Ubuntu AMIs, in conjunction with the Amazon EFS Mount Helper. For instructions, see Using the amazon-efs-utils Tools.

For a list of Amazon EC2 Linux Amazon Machine Images (AMIs) that support this protocol, see NFS Support. For some AMIs, you'll need to install an NFS client to mount your file system on your Amazon EC2 instance. For instructions, see Installing the NFS Client. You can access your Amazon EFS file system concurrently from multiple NFS clients, so applications that scale beyond a single connection can access a file system. Amazon EC2 instances running in multiple Availability Zones within the same AWS Region can access the file system, so that many users can access and share a common data source.

**QUESTION 917**

A company is designing an application to run in a VPC on AWS. The application consists of Amazon EC2 instances that run in private subnets as part of an Auto Scaling group. The application also includes a Network Load Balancer that extends across public subnets. The application stores data in an Amazon RDS DB instance.

The company has attached a security group that is named "web-servers" to the EC2 instances. The company has attached a security group that is named "database" to the DB Instance.

How should a solutions architect configure the communication between the EC2 instances and the DB instance?

- A. Configure the "web-servers" security group to allow access to the DB instance's current IP addresses.  
Configure the "database" security group to allow access from the current set of IP addresses in use by the EC2 instances
- B. Configure the "web-servers" security group to allow access to the "database" security group.  
Configure the "database" security group to allow access from the "web-servers" security group
- C. Configure the "web-servers" security group to allow access to the DB instance's current IP addresses.  
Configure the "database" security group to allow access from the Auto Scaling group
- D. Configure the "web servers" security group to allow access to the "database" security group.  
Configure the "database" security group to allow access from the Auto Scaling group

**Answer: C**

**QUESTION 918**

A company wants an AWS Lambda function to call a third-party API and save the response to a private Amazon RDS DB instance in the same private subnet.

What should a solutions architect do to meet these requirements?

- A. Create a NAT gateway. In the route table for the private subnet, add a route to the NAT gateway. Attach the Lambda function to the private subnet. Create an IAM role that includes the AWSLambdaBasicExecutionRole permissions policy. Attach the role to the Lambda function B Create an internet gateway In the route table for the private subnet, add a route to the internet gateway. Attach the Lambda function to the private subnet. Create an IAM role that includes me AWSLambdaBasicExecutionRole permissions policy. Attach the role to the Lambda function
- B. Create a NAT gateway in the route table for the private subnet add a route to the NAT gateway. Attach the Lambda function to the private subnet. Create an IAM role that includes the AWS LambdaVPCAccessExecutionRole permissions policy. Attach the role to the Lambda function
- C. Create an internet gateway in the route table for the private subnet, add a route to the internet gateway. Attach the Lambda function to the private subnet. Create an IAM role that includes the AWSLambdaVPCAccessExecutionRole permissions policy. Attach the role to the Lambda function

**Answer: B**

#### **QUESTION 919**

A company has a stateless asynchronous application that runs in an Apache Hadoop cluster.

The application is invoked on demand to run extract, transform and load (ETL) jobs several times a day.

A solutions architect needs to migrate this application to the AWS Cloud by designing an Amazon EMR cluster for the workload.

The cluster must be available immediately to process jobs.

Which implementation meets these requirements MOST cost-effectively?

- A. Use zonal Reserved Instances for the master nodes and the ewe nodes.  
Use a Spot Fleet for tire task nodes
- B. Use zonal Reserved Instances for the master nodes.  
Use Spot instances for the core nodes and the task nodes
- C. Use regional Reserved Instances for the master nodes.  
Use a Spot Fleer for the core nodes and the task nodes
- D. Use regional Reserved Instances for the master nodes.  
Use On-Demand Capacity Reservations for the core nodes and the task nodes.

**Answer: A**

#### **QUESTION 920**

A media company is using video conversion tools that run on Amazon EC2 instances. The video conversion tools run on a combination of Windows EC2 instances and Linux EC2 instances Each video file is tens of gigabytes in size. The video conversion tools must process the video files in the shortest possible amount of time. The company needs a single, centralized file storage solution that can be mounted on all the EC2 Instances that host the video conversion tools.

Which solution will meet these requirements?

- A. Deploy Amazon FSx for Windows File Server with hard disk drive (HDD) storage
- B. Deploy Amazon FSx for Windows File Server wild solid stale drive (SSD) storage
- C. Deploy Amazon Elastic File System (Amazon EFS) with Max I/O performance mode
- D. Deploy Amazon Elastic File System (Amazon EFS) with General Purpose performance mode

**Answer: D**

#### **QUESTION 921**

A company's order system sends requests from clients to Amazon EC2 instances. The EC2 instances process the orders and then store the orders in a database on Amazon RDS. Users report that they must reprocess orders when the system tails. The company wants a resilient solution that can process orders automatically if a system outage occurs.

What should a solutions architect do to meet these requirements?

- A. Move the EC2 instances Into an Auto Scaling group.  
Create an Amazon EventBridge (Amazon CloudWatch Events) rule to target an Amazon Elastic Container Service (Amazon ECS) task
- B. Move the EC2 instances into an Auto Seating group behind an Application Load Balancer (ALB).  
Update the order system to send message to the ALB endpoint
- C. Move the EC2 instances into an Auto Scaling group.  
Configure the order system to send messages to an Amazon Simple Queue Service (Amazon SGS) queue.  
Configure the EC2 instances to consume messages from the queue.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic.  
Create an AWS Lambda function, and subscribe the function to the SNS topic.  
Configure (he order system to send messages to the SNS topic.  
Send a command to the EC2 instances to process the messages by using AWS Systems Manager Run Command

**Answer: C**

#### **QUESTION 922**

A company has an application that processes customer of tiers. The company hosts the application on an Amazon EC2 instance that saves the orders to an Amazon Aurora database. Occasionally when traffic is high, the workload does not process orders fast enough.

What should a solutions architect do to write the orders reliably to the database as quickly as possible?

- A. Increase the instance size of the EC2 instance when baffle ls high.  
Write orders to Amazon Simple Notification Service (Amazon SNS).  
Subscribe the database endpoint to the SNS topic
- B. Write orders to an Amazon Simple Queue Service (Amazon SOS) queue.  
Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read born the SQS queue and process orders into the database
- C. Write orders to Amazon Simple Notification Service (Amazon SNS).  
Subscribe the database endpoint to the SNS topic.  
Use EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SNS topic.
- D. Write orders to an Amazon Simple Queue Service (Amazon SQS) queue when the EC2 instance reaches CPU threshold limits.  
Use scheduled scaling of EC2 instances in an Auto Scaling group behind an Application Load Balancer to read from the SQS queue and process orders into the database

**Answer: B**

#### **QUESTION 923**

A company runs a critical customer-lacing application on Amazon Elastic Kubernetes Service (Amazon EKS). The application has a microservices architecture. The company needs to implement a solution that collects, aggregates and summarizes metrics and logs from the application in a centralized location.

Which solution meets these requirements?

- A. Run the Amazon CloudWatch agent in the existing EKS cluster.  
View the metrics and logs in the CloudWatch console.
- B. Run AWS App Mesh in the existing EKS duster.  
View the metrics and logs m the App Mesh console
- C. Configure AWS CloudTrail to capture data events.  
Query CloudTrail by using Amazon OpenSearch Service (Amazon Elasticsearch Service)
- D. Configure Amazon CloudWatch Container Insights in the existing EKS cluster.  
View the metrics and logs in the CloudWatch console.

**Answer: C**



**QUESTION 924**

A company is deploying a new application to Amazon Elastic Kubernetes Service (Amazon EKS) with an AWS Fargate cluster. The application needs a storage solution for data persistence. The solution must be highly available and fault tolerant. The solution also must be shared between multiple application containers. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create Amazon Elastic Block Store (Amazon EBS) volumes in the same Availability Zones where EKS worker nodes are placed. Register the volumes in a StorageClass object on an EKS cluster. Use EBS Multi-Attach to share the data between containers
- B. Create an Amazon Elastic File System (Amazon EFS) file system. Register the file system in a StorageClass object on an EKS cluster. Use the same file system for all containers
- C. Create an Amazon Elastic Block Store (Amazon EBS) volume. Register the volume in a StorageClass object on an EKS cluster. Use the same volume for all containers.
- D. Create Amazon Elastic File System (Amazon EFS) file systems in the same Availability Zones where EKS worker nodes are placed. Register the file systems in a StorageClass object on an EKS cluster. Create an AWS Lambda function to synchronize the data between file systems

**Answer: B**

**QUESTION 925**

A company is storing sensitive user information in an Amazon S3 bucket. The company wants to provide secure access to this bucket from the application tier running on Amazon EC2 instances inside a VPC.

Which combination of steps should a solutions architect take to accomplish this? (Select TWO.)

- A. Configure a VPC gateway endpoint (or Amazon S3 within the VPC)
- B. Create a bucket policy to make the objects in the S3 bucket public
- C. Create a bucket policy that limits access to only the application tier running in the VPC
- D. Create an IAM user with an S3 access policy and copy the IAM credentials to the EC2 instance
- E. Create a NAT instance and have the EC2 instances use the NAT instance to access the S3 bucket

**Answer: AD**

**QUESTION 926**

A company is preparing to store confidential data in Amazon S3. For compliance reasons the data must be encrypted at rest. Encryption key usage must be logged for auditing purposes. Keys must be rotated every year. Which solution meets these requirements and the MOST operationally efficient?

- A. Server-side encryption with customer-provided keys (SSE-C)
- B. Server-side encryption with Amazon S3 managed keys (SSE-S3)
- C. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with manual rotation
- D. Server-side encryption with AWS KMS (SSE-KMS) customer master keys (CMKs) with automatic rotation

**Answer: D**

**Explanation:**

<https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html>

When you enable automatic key rotation for a customer managed key, AWS KMS generates new cryptographic material for the KMS key every year. AWS KMS also saves the KMS key's older cryptographic material in perpetuity so it can be used to decrypt data that the KMS key encrypted.

Key rotation in AWS KMS is a cryptographic best practice that is designed to be transparent and easy to use. AWS KMS supports optional automatic key rotation only for customer managed CMKs. Enable and disable key rotation. Automatic key rotation is disabled by default on customer managed CMKs. When you enable (or re-enable) key rotation, AWS KMS automatically rotates the CMK 365 days after the enable date and every 365 days thereafter.

**QUESTION 927**

A survey company has gathered data for several years from areas in the United States. The company hosts the data in an Amazon S3 bucket that is 3 TB in size and growing. The company has started to share the data with a European marketing firm that has S3 buckets. The company wants to ensure that its data transfer costs remain as low as possible.

Which solution will meet these requirements?

- A. Configure the Requester Pays feature on the company's S3 bucket
- B. Configure S3 Cross-Region Replication from the company's S3 bucket to one of the marketing firm's S3 buckets.
- C. Configure cross-account access for the marketing firm so that the marketing firm has access to the company's S3 bucket.
- D. Configure the company's S3 bucket to use S3 Intelligent-Tiering Sync the S3 bucket to one of the marketing firm's S3 buckets

**Answer: A**

**QUESTION 928**

A company hosts an application on AWS Lambda functions that are invoked by an Amazon API Gateway API. The Lambda functions save customer data to an Amazon Aurora MySQL database. Whenever the company upgrades the database, the Lambda functions fail to establish database connections until the upgrade is complete. The result is that customer data is not recorded for some of the event.

A solutions architect needs to design a solution that stores customer data that is created during database upgrades. Which solution will meet these requirements?

- A. Provision an Amazon RDS proxy to sit between the Lambda functions and the database.  
Configure the Lambda functions to connect to the RDS proxy
- B. Increase the run time of the Lambda functions to the maximum.  
Create a retry mechanism in the code that stores the customer data in the database
- C. Persist the customer data to Lambda local storage.  
Configure new Lambda functions to scan the local storage to save the customer data to the database.
- D. Store the customer data in an Amazon Simple Queue Service (Amazon SQS) FIFO queue.  
Create a new Lambda function that polls the queue and stores the customer data in the database

**Answer: C**

**QUESTION 929**

A company has an application that loads documents into an Amazon S3 bucket and converts the documents into another format. The application stores the converted documents in another S3 bucket and saves the document name and URLs in an Amazon DynamoDB table. The DynamoDB entries are used during subsequent days to access the documents. The company uses a DynamoDB Accelerator (DAX) cluster in front of the table.

Recently, traffic to the application has increased. Document processing tasks are timing out during the scheduled DAX maintenance window. A solutions architect must ensure that the documents continue to load during the maintenance window.

What should the solutions architect do to accomplish this goal?

- A. Modify the application to write to the DAX cluster.  
Configure the DAX cluster to write to the DynamoDB table when the maintenance window is complete
- B. Enable Amazon DynamoDB Streams for the DynamoDB table.  
Modify the application to write to the stream.

Configure the stream to load the data when the maintenance window is complete.

- C. Convert the application to an AWS Lambda function.  
Configure the Lambda function runtime to be longer than the maintenance window.  
Create an Amazon CloudWatch alarm to monitor Lambda timeouts
- D. Modify the application to write the document name and URLs to an Amazon Simple Queue Service (Amazon SQS) queue.  
Create an AWS Lambda function to read the SQS queue and write to DynamoDB.

**Answer: C**

#### **QUESTION 930**

A company is creating a new application that will store a large amount of data. The data will be analyzed hourly and will be modified by several Amazon EC2 Linux instances that are deployed across multiple Availability Zones. The needed amount of storage space will continue to grow for the next 6 Months.

Which storage solution should a solutions architect recommend to meet these requirements?

- A. Store the data in Amazon S3 Glacier.  
Update the S3 Glacier vault policy to allow access to the application instances
- B. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume.  
Mount the EBS volume on the application instances.
- C. Store the data in an Amazon Elastic File System (Amazon EFS) file system.  
Mount the file system on the application instances.
- D. Store the data in an Amazon Elastic Block Store (Amazon EBS).  
Provisioned IOPS volume shared between the application instances.

**Answer: C**

#### **QUESTION 931**

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. What should the solutions architect do to maximize reliability of the application infrastructure?

- A. Delete one EC2 instance and enable termination protection on the other EC2 instance.  
Update the DB instance to be Multi-AZ, and enable deletion protection.
- B. Update the DB instance to be Multi-AZ, and enable deletion protection.  
Place the EC2 instances behind an Application Load Balancer, and run them in an EC2 Auto Scaling 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**

**Explanation:**

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-spot-instances.html>

#### **QUESTION 932**

A company has developed a new content-sharing application that runs on Amazon Elastic Container Service (Amazon ECS). The application runs on Amazon Linux Docker tasks that use the Amazon EC2 launch type. The application requires a storage solution that has the following characteristics:

- Accessibility (or multiple ECS tasks through bind mounts)

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

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



- Resiliency across Availability Zones
- Burstable throughput of up to 3 Gbps
- Ability to be scaled up over time

Which storage solution meets these requirements?

- A. Launch an Amazon FSx for Windows File Server Multi-AZ instance.  
Configure the ECS task definitions to mount the Amazon FSx instance volume at launch.
- B. Launch an Amazon Elastic File System (Amazon EFS) instance.  
Configure the ECS task definitions to mount the EFS Instance volume at launch.
- C. Create a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach set to enabled.  
Attach the EBS volume to the ECS EC2 instance.  
Configure ECS task definitions to mount the EBS instance volume at launch.
- D. Launch an EC2 instance with several Provisioned IOPS SSD (PIS) Amazon Elastic Block Store (Amazon EBS) volumes attached in a RAID 0 configuration.  
Configure the EC2 instance as an NFS storage server.  
Configure ECS task definitions to mount the volumes at launch.

**Answer: B**

### **QUESTION 933**

A company has 150 TB of archived image data stored on-premises that needs to be moved to the AWS Cloud within the next month.

The company's current network connection allows up to 100 Mbps uploads for this purpose during the night only.

What is the MOST cost-effective mechanism to move this data and meet the migration deadline?

- A. Use AWS Snowmobile to ship the data to AWS.
- B. Order multiple AWS Snowball devices to ship the data to AWS.
- C. Enable Amazon S3 Transfer Acceleration and securely upload the data.
- D. Create an Amazon S3 VPC endpoint and establish a VPN to upload the data

**Answer: B**

#### **Explanation:**

eg. 6 hrs night

$6 \text{ hrs} \times 60 \text{ min/hr} = 360 \text{ min}$

$360 \text{ min} \times 60 \text{ sec/min} = 21600 \text{ sec}$

$100 \text{ Mbps} \times 21600 \text{ s} = 2160000 \text{ Mb}$

or 2160 Gb or 2.1 TB can only be done

So, for 150 TB, we can use 2 X Snowball Edge Storage Optimised devices.

Size of Snowball Edge Storage Optimised device = 80 TB  
Size of Snowball Edge Compute Optimised device = 40 TB

Size of Snowcone = 8 TB

Size of Snowmobile = 100 PB (1 PB = 1000 TB)

Q: How should I choose between Snowmobile and Snowball?

To migrate large datasets of 10PB or more in a single location, you should use Snowmobile. For datasets less than 10PB or distributed in multiple locations, you should use Snowball. In addition, you should evaluate the amount of available bandwidth in your network backbone. If you have a high speed backbone with hundreds of Gb/s of spare throughput, then you can use Snowmobile to migrate the large datasets all at once. If you have limited bandwidth on your backbone, you should consider using multiple Snowballs to migrate the data incrementally.