

- **Vendor: Amazon**
- **Exam Code: SAA-C03**
- **Exam Name: AWS Certified Solutions Architect - Associate (SAA-C03)**
- **New Updated Questions from [Braindump2go](https://www.braindump2go.com)**
- **(Updated in [September/2022](https://www.braindump2go.com))**

### **Visit Braindump2go and Download Full Version SAA-C03 Exam Dumps**

#### **QUESTION 66**

A company has a data ingestion workflow that includes the following components:

- An Amazon Simple Notification Service (Amazon SNS) topic that receives notifications about new data deliveries.
- An AWS Lambda function that processes and stores the data

The ingestion workflow occasionally fails because of network connectivity issues.

When tenure occurs the corresponding data is not ingested unless the company manually reruns the job.

What should a solutions architect do to ensure that all notifications are eventually processed?

- A. Configure the Lambda function for deployment across multiple Availability Zones
- B. Modify the Lambda function's configuration to increase the CPU and memory allocations for the function
- C. Configure the SNS topic's retry strategy to increase both the number of retries and the wait time between retries
- D. Configure an Amazon Simple Queue Service (Amazon SQS) queue as the on failure destination.  
Modify the Lambda function to process messages in the queue.

**Answer: A**

#### **QUESTION 67**

A company has a service that produces event data. The company wants to use AWS to process the event data as it is received.

The data is written in a specific order that must be maintained throughout processing.

The company wants to implement a solution that minimizes operational overhead.

How should a solutions architect accomplish this?

- A. Create an Amazon Simple Queue Service (Amazon SQS) FIFO queue to hold messages.  
Set up an AWS Lambda function to process messages from the queue.
- B. Create an Amazon Simple Notification Service (Amazon SNS) topic to deliver notifications containing payloads to process.  
Configure an AWS Lambda function as a subscriber.
- C. Create an Amazon Simple Queue Service (Amazon SQS) standard queue to hold messages.  
Set up an AWS Lambda function to process messages from the queue independently.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic to deliver notifications containing payloads to process.  
Configure an Amazon Simple Queue Service (Amazon SQS) queue as a subscriber.

**[SAA-C03 Exam Dumps](#) [SAA-C03 Exam Questions](#) [SAA-C03 PDF Dumps](#) [SAA-C03 VCE Dumps](#)**

**<https://www.braindump2go.com/saa-c03.html>**

**Answer: A**

**Explanation:**

The details are revealed in below url:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFO-queues.html>

FIFO (First-In-First-Out) queues are designed to enhance messaging between applications when the order of operations and events is critical, or where duplicates can't be tolerated. Examples of situations where you might use FIFO queues include the following: To make sure that user-entered commands are run in the right order. To display the correct product price by sending price modifications in the right order. To prevent a student from enrolling in a course before registering for an account.

**QUESTION 68**

A company is migrating an application from on-premises servers to Amazon EC2 instances. As part of the migration design requirements, a solutions architect must implement infrastructure metric alarms. The company does not need to take action if CPU utilization increases to more than 50% for a short burst of time. However, if the CPU utilization increases to more than 50% and read IOPS on the disk are high at the same time, the company needs to act as soon as possible. The solutions architect also must reduce false alarms. What should the solutions architect do to meet these requirements?

- A. Create Amazon CloudWatch composite alarms where possible.
- B. Create Amazon CloudWatch dashboards to visualize the metrics and react to issues quickly.
- C. Create Amazon CloudWatch Synthetics canaries to monitor the application and raise an alarm.
- D. Create single Amazon CloudWatch metric alarms with multiple metric thresholds where possible.

**Answer: A**

**QUESTION 69**

A company wants to migrate its on-premises data center to AWS. According to the company's compliance requirements, the company can use only the ap-northeast-3 Region. Company administrators are not permitted to connect VPCs to the internet.

Which solutions will meet these requirements? (Choose two.)

- A. Use AWS Control Tower to implement data residency guardrails to deny internet access and deny access to all AWS Regions except ap-northeast-3.
- B. Use rules in AWS WAF to prevent internet access.  
Deny access to all AWS Regions except ap-northeast-3 in the AWS account settings.
- C. Use AWS Organizations to configure service control policies (SCPS) that prevent VPCs from gaining internet access.  
Deny access to all AWS Regions except ap-northeast-3.
- D. Create an outbound rule for the network ACL in each VPC to deny all traffic from 0.0.0.0/0.  
Create an IAM policy for each user to prevent the use of any AWS Region other than ap-northeast-3.
- E. Use AWS Config to activate managed rules to detect and alert for internet gateways and to detect and alert for new resources deployed outside of ap-northeast-3.

**Answer: AC**

**QUESTION 70**

A company uses a three-tier web application to provide training to new employees. The application is accessed for only 12 hours every day. The company is using an Amazon RDS for MySQL DB instance to store information and wants to minimize costs.

What should a solutions architect do to meet these requirements?

- A. Configure an IAM policy for AWS Systems Manager Session Manager.  
Create an IAM role for the policy.  
Update the trust relationship of the role.

Set up automatic start and stop for the DB instance.

- B. Create an Amazon ElastiCache for Redis cache cluster that gives users the ability to access the data from the cache when the DB instance is stopped.  
Invalidate the cache after the DB instance is started.
- C. Launch an Amazon EC2 instance.  
Create an IAM role that grants access to Amazon RDS.  
Attach the role to the EC2 instance.  
Configure a cron job to start and stop the EC2 instance on the desired schedule.
- D. Create AWS Lambda functions to start and stop the DB instance.  
Create Amazon EventBridge (Amazon CloudWatch Events) scheduled rules to invoke the Lambda functions.  
Configure the Lambda functions as event targets for the rules

**Answer: D**

#### **QUESTION 71**

A company sells ringtones created from clips of popular songs. The files containing the ringtones are stored in Amazon S3 Standard and are at least 128 KB in size. The company has millions of files, but downloads are infrequent for ringtones older than 90 days. The company needs to save money on storage while keeping the most accessed files readily available for its users.

Which action should the company take to meet these requirements MOST cost-effectively?

- A. Configure S3 Standard-Infrequent Access (S3 Standard-IA) storage for the initial storage tier of the objects.
- B. Move the files to S3 Intelligent-Tiering and configure it to move objects to a less expensive storage tier after 90 days.
- C. Configure S3 inventory to manage objects and move them to S3 Standard-Infrequent Access (S3 Standard-1A) after 90 days.
- D. Implement an S3 Lifecycle policy that moves the objects from S3 Standard to S3 Standard-Infrequent Access (S3 Standard-1A) after 90 days.

**Answer: D**

#### **QUESTION 72**

A company needs to save the results from a medical trial to an Amazon S3 repository. The repository must allow a few scientists to add new files and must restrict all other users to read-only access. No users can have the ability to modify or delete any files in the repository. The company must keep every file in the repository for a minimum of 1 year after its creation date.

Which solution will meet these requirements?

- A. Use S3 Object Lock In governance mode with a legal hold of 1 year
- B. Use S3 Object Lock in compliance mode with a retention period of 365 days.
- C. Use an IAM role to restrict all users from deleting or changing objects in the S3 bucket Use an S3 bucket policy to only allow the IAM role
- D. Configure the S3 bucket to invoke an AWS Lambda function every time an object is added  
Configure the function to track the hash of the saved object to that modified objects can be marked accordingly

**Answer: B**

#### **QUESTION 73**

A large media company hosts a web application on AWS. The company wants to start caching confidential media files so that users around the world will have reliable access to the files. The content is stored in Amazon S3 buckets. The company must deliver the content quickly, regardless of where the requests originate geographically.

Which solution will meet these requirements?

- A. Use AWS DataSync to connect the S3 buckets to the web application.

**[SAA-C03 Exam Dumps](#) [SAA-C03 Exam Questions](#) [SAA-C03 PDF Dumps](#) [SAA-C03 VCE Dumps](#)**

**<https://www.braindump2go.com/saa-c03.html>**

- B. Deploy AWS Global Accelerator to connect the S3 buckets to the web application.
- C. Deploy Amazon CloudFront to connect the S3 buckets to CloudFront edge servers.
- D. Use Amazon Simple Queue Service (Amazon SQS) to connect the S3 buckets to the web application.

**Answer: C**

**Explanation:**

CloudFront uses a local cache to provide the response, AWS Global accelerator proxies requests and connects to the application all the time for the response.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html#private-content-granting-permissions-to-oai>

**QUESTION 74**

A company produces batch data that comes from different databases. The company also produces live stream data from network sensors and application APIs. The company needs to consolidate all the data into one place for business analytics. The company needs to process the incoming data and then stage the data in different Amazon S3 buckets. Teams will later run one-time queries and import the data into a business intelligence tool to show key performance indicators (KPIs).

Which combination of steps will meet these requirements with the LEAST operational overhead? (Choose two.)

- A. Use Amazon Athena for one-time queries.  
Use Amazon QuickSight to create dashboards for KPIs.
- B. Use Amazon Kinesis Data Analytics for one-time queries.  
Use Amazon QuickSight to create dashboards for KPIs.
- C. Create custom AWS Lambda functions to move the individual records from the databases to an Amazon Redshift cluster.
- D. Use an AWS Glue extract, transform, and load (ETL) job to convert the data into JSON format.  
Load the data into multiple Amazon OpenSearch Service (Amazon Elasticsearch Service) clusters.
- E. Use blueprints in AWS Lake Formation to identify the data that can be ingested into a data lake.  
Use AWS Glue to crawl the source, extract the data, and load the data into Amazon S3 in Apache Parquet format.

**Answer: AB**

**QUESTION 75**

A gaming company has a web application that displays scores. The application runs on Amazon EC2 instances behind an Application Load Balancer. The application stores data in an Amazon RDS for MySQL database. Users are starting to experience long delays and interruptions that are caused by database read performance. The company wants to improve the user experience while minimizing changes to the application's architecture.

What should a solutions architect do to meet these requirements?

- A. Use Amazon ElastiCache in front of the database.
- B. Use RDS Proxy between the application and the database.
- C. Migrate the application from EC2 instances to AWS Lambda.
- D. Migrate the database from Amazon RDS for MySQL to Amazon DynamoDB.

**Answer: C**

**QUESTION 76**

A business's backup data totals 700 terabytes (TB) and is kept in network attached storage (NAS) at its data center. This backup data must be available in the event of occasional regulatory inquiries and preserved for a period of seven years. The organization has chosen to relocate its backup data from its on-premises data center to Amazon Web Services (AWS). Within one month, the migration must be completed. The company's public internet connection provides 500 Mbps of dedicated capacity for data transport.

What should a solutions architect do to ensure that data is migrated and stored at the LOWEST possible cost?

- A. Order AWS Snowball devices to transfer the data.  
Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
- B. Deploy a VPN connection between the data center and Amazon VPC.  
Use the AWS CLI to copy the data from on premises to Amazon S3 Glacier.
- C. Provision a 500 Mbps AWS Direct Connect connection and transfer the data to Amazon S3.  
Use a lifecycle policy to transition the files to Amazon S3 Glacier Deep Archive.
- D. Use AWS DataSync to transfer the data and deploy a DataSync agent on premises.  
Use the DataSync task to copy files from the on-premises NAS storage to Amazon S3 Glacier.

**Answer: A**

#### **QUESTION 77**

A company wants to direct its users to a backup static error page if the company's primary website is unavailable. The primary website's DNS records are hosted in Amazon Route 53. The domain is pointing to an Application Load Balancer (ALB). The company needs a solution that minimizes changes and infrastructure overhead. Which solution will meet these requirements?

- A. Update the Route 53 records to use a latency routing policy.  
Add a static error page that is hosted in an Amazon S3 bucket to the records so that the traffic is sent to the most responsive endpoints.
- B. Set up a Route 53 active-passive failover configuration.  
Direct traffic to a static error page that is hosted in an Amazon S3 bucket when Route 53 health checks determine that the ALB endpoint is unhealthy.
- C. Set up a Route 53 active-active configuration with the ALB and an Amazon EC2 instance that hosts a static error page as endpoints.  
Configure Route 53 to send requests to the instance only if the health checks fail for the ALB.
- D. Update the Route 53 records to use a multivalue answer routing policy.  
Create a health check.  
Direct traffic to the website if the health check passes.  
Direct traffic to a static error page that is hosted in Amazon S3 if the health check does not pass.

**Answer: B**

#### **QUESTION 78**

A corporation has recruited a new cloud engineer who should not have access to the CompanyConfidential Amazon S3 bucket. The cloud engineer must have read and write permissions on an S3 bucket named AdminTools. Which IAM policy will satisfy these criteria?

A. {  
 "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"  
 }  
 ]  
}

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**

**Explanation:**

[https://docs.amazonaws.cn/en\\_us/IAM/latest/UserGuide/reference\\_policies\\_examples\\_s3\\_rw-bucket.html](https://docs.amazonaws.cn/en_us/IAM/latest/UserGuide/reference_policies_examples_s3_rw-bucket.html)

#### QUESTION 79

A new employee has joined a company as a deployment engineer. The deployment engineer will be using AWS CloudFormation templates to create multiple AWS resources.

A solutions architect wants the deployment engineer to perform job activities while following the principle of least privilege.

Which steps should the solutions architect do in conjunction to reach this goal? (Select two.)

- A. Have the deployment engineer use AWS account root user credentials for performing AWS CloudFormation stack operations.

**[SAA-C03 Exam Dumps](#) [SAA-C03 Exam Questions](#) [SAA-C03 PDF Dumps](#) [SAA-C03 VCE Dumps](#)**

**<https://www.braindump2go.com/saa-c03.html>**

- B. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the PowerUsers IAM policy attached.
- C. Create a new IAM user for the deployment engineer and add the IAM user to a group that has the Administrate/Access IAM policy attached.
- D. Create a new IAM User for the deployment engineer and add the IAM user to a group that has an IAM policy that allows AWS CloudFormation actions only.
- E. Create an IAM role for the deployment engineer to explicitly define the permissions specific to the AWS CloudFormation stack and launch stacks using Dial IAM role.

**Answer:** DE

**Explanation:**

[https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html)

[https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_users.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_users.html)

#### **QUESTION 80**

A company runs a high performance computing (HPC) workload on AWS. The workload required low-latency network performance and high network throughput with tightly coupled node-to-node communication. The Amazon EC2 instances are properly sized for compute and storage capacity, and are launched using default options. What should a solutions architect propose to improve the performance of the workload?

- A. Choose a cluster placement group while launching Amazon EC2 instances.
- B. Choose dedicated instance tenancy while launching Amazon EC2 instances.
- C. Choose an Elastic Inference accelerator while launching Amazon EC2 instances.
- D. Choose the required capacity reservation while launching Amazon EC2 instances.

**Answer:** A

**Explanation:**

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-placementgroup.html>

A cluster placement group is a logical grouping of instances within a single Availability Zone that benefit from low network latency, high network throughput.

#### **QUESTION 81**

A company wants to use the AWS Cloud to make an existing application highly available and resilient. The current version of the application resides in the company's data center. The application recently experienced data loss after a database server crashed because of an unexpected power outage. The company needs a solution that avoids any single points of failure. The solution must give the application the ability to scale to meet user demand. Which solution will meet these requirements?

- A. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones.  
Use an Amazon RDS DB instance in a Multi-AZ configuration.
- B. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group in a single Availability Zone.  
Deploy the database on an EC2 instance.  
Enable EC2 Auto Recovery.
- C. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones.  
Use an Amazon RDS DB instance with a read replica in a single Availability Zone.  
Promote the read replica to replace the primary DB instance if the primary DB instance fails.
- D. Deploy the application servers by using Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones.  
Deploy the primary and secondary database servers on EC2 instances across multiple Availability Zones.  
Use Amazon Elastic Block Store (Amazon EBS) Multi-Attach to create shared storage between the instances.

**Answer:** A

**[SAA-C03 Exam Dumps](#) [SAA-C03 Exam Questions](#) [SAA-C03 PDF Dumps](#) [SAA-C03 VCE Dumps](#)**

**<https://www.braindump2go.com/saa-c03.html>**



**QUESTION 82**

A company wants to run a gaming application on Amazon EC2 instances that are part of an Auto Scaling group in the AWS Cloud. The application will transmit data by using UDP packets. The company wants to ensure that the application can scale out and in as traffic increases and decreases. What should a solutions architect do to meet these requirements?

- A. Attach a Network Load Balancer to the Auto Scaling group
- B. Attach an Application Load Balancer to the Auto Scaling group.
- C. Deploy an Amazon Route 53 record set with a weighted policy to route traffic appropriately
- D. Deploy a NAT instance that is configured with port forwarding to the EC2 instances in the Auto Scaling group.

**Answer: B**

**QUESTION 83**

A solutions architect is designing a customer-facing application for a company. The application's database will have a clearly defined access pattern throughout the year and will have a variable number of reads and writes that depend on the time of year. The company must retain audit records for the database for 7 days. The recovery point objective (RPO) must be less than 5 hours.

Which solution meets these requirements?

- A. Use Amazon DynamoDB with auto scaling.  
Use on-demand backups and Amazon DynamoDB Streams.
- B. Use Amazon Redshift. Configure concurrency scaling.  
Activate audit logging.  
Perform database snapshots every 4 hours.
- C. Use Amazon RDS with Provisioned IOPS.  
Activate the database auditing parameter.  
Perform database snapshots every 5 hours.
- D. Use Amazon Aurora MySQL with auto scaling.  
Activate the database auditing parameter

**Answer: B**

**QUESTION 84**

A company hosts a two-tier application on Amazon EC2 instances and Amazon RDS. The application's demand varies based on the time of day. The load is minimal after work hours and on weekends. The EC2 instances run in an EC2 Auto Scaling group that is configured with a minimum of two instances and a maximum of five instances. The application must be available at all times, but the company is concerned about overall cost.

Which solution meets the availability requirement MOST cost-effectively?

- A. Use all EC2 Spot Instances.  
Stop the RDS database when it is not in use.
- B. Purchase EC2 Instance Savings Plans to cover five EC2 instances.  
Purchase an RDS Reserved DB Instance
- C. Purchase two EC2 Reserved Instances.  
Use up to three additional EC2 Spot Instances as needed.  
Stop the RDS database when it is not in use.
- D. Purchase EC2 Instance Savings Plans to cover two EC2 instances.  
Use up to three additional EC2 On-Demand Instances as needed.  
Purchase an RDS Reserved DB Instance.

**Answer: D**

**QUESTION 85**

A company has an ecommerce checkout workflow that writes an order to a database and calls a service to process the

**[SAA-C03 Exam Dumps](#)** **[SAA-C03 Exam Questions](#)** **[SAA-C03 PDF Dumps](#)** **[SAA-C03 VCE Dumps](#)**

**<https://www.braindump2go.com/saa-c03.html>**

payment. Users are experiencing timeouts during the checkout process.

When users resubmit the checkout form, multiple unique orders are created for the same desired transaction.

How should a solutions architect refactor this workflow to prevent the creation of multiple orders?

- A. Configure the web application to send an order message to Amazon Kinesis Data Firehose. Set the payment service to retrieve the message from Kinesis Data Firehose and process the order.
- B. Create a rule in AWS CloudTrail to invoke an AWS Lambda function based on the logged application path request. Use Lambda to query the database, call the payment service, and pass in the order information.
- C. Store the order in the database. Send a message that includes the order number to Amazon Simple Notification Service (Amazon SNS). Set the payment service to poll Amazon SNS, retrieve the message, and process the order.
- D. Store the order in the database. Send a message that includes the order number to an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Set the payment service to retrieve the message and process the order. Delete the message from the queue.

**Answer: D**

#### **QUESTION 86**

A company is planning to build a high performance computing (HPC) workload as a service solution that is hosted on AWS.

A group of 16 Amazon EC2 Linux instances requires the lowest possible latency for node-to-node communication.

The instances also need a shared block device volume for high-performing storage.

Which solution will meet these requirements?

- A. Use a duster placement group. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon EBS) volume to all the instances by using Amazon EBS Multi-Attach.
- B. Use a cluster placement group. Create shared file systems across the instances by using Amazon Elastic File System (Amazon EFS).
- C. Use a partition placement group. Create shared tile systems across the instances by using Amazon Elastic File System (Amazon EFS).
- D. Use a spread placement group. Attach a single Provisioned IOPS SSD Amazon Elastic Block Store (Amazon EBS) volume to all the instances by using Amazon EBS Multi-Attach.

**Answer: A**

#### **QUESTION 87**

A company has an event-driven application that invokes AWS Lambda functions up to 800 times each minute with varying runtimes.

The Lambda functions access data that is stored in an Amazon Aurora MySQL DB cluster.

The company is noticing connection timeouts as user activity increases. The database shows no signs of being overloaded. CPU, memory, and disk access metrics are all low.

Which solution will resolve this issue with the LEAST operational overhead?

- A. Adjust the size of the Aurora MySQL nodes to handle more connections. Configure retry logic in the Lambda functions for attempts to connect to the database.
- B. Set up Amazon ElastiCache for Redis to cache commonly read items from the database. Configure the Lambda functions to connect to ElastiCache for reads.
- C. Add an Aurora Replica as a reader node.

Configure the Lambda functions to connect to the reader endpoint of the OB cluster rather than to the writer endpoint.

- D. Use Amazon ROS Proxy to create a proxy.  
Set the DB cluster as the target database.  
Configure the Lambda functions to connect to the proxy rather than to the DB cluster.

**Answer: D**

#### **QUESTION 88**

A company is building a containerized application on premises and decides to move the application to AWS. The application will have thousands of users soon after it is deployed.

The company is unsure how to manage the deployment of containers at scale. The company needs to deploy the containerized application in a highly available architecture that minimizes operational overhead.

Which solution will meet these requirements?

- A. Store container images in an Amazon Elastic Container Registry (Amazon ECR) repository.  
Use an Amazon Elastic Container Service (Amazon ECS) cluster with the AWS Fargate launch type to run the containers.  
Use target tracking to scale automatically based on demand.
- B. Store container images in an Amazon Elastic Container Registry (Amazon ECR) repository.  
Use an Amazon Elastic Container Service (Amazon ECS) cluster with the Amazon EC2 launch type to run the containers.  
Use target tracking to scale automatically based on demand.
- C. Store container images in a repository that runs on an Amazon EC2 instance.  
Run the containers on EC2 instances that are spread across multiple Availability Zones.  
Monitor the average CPU utilization in Amazon CloudWatch.  
Launch new EC2 instances as needed.
- D. Create an Amazon EC2 Amazon Machine Image (AMI) that contains the container image.  
Launch EC2 instances in an Auto Scaling group across multiple Availability Zones.  
Use an Amazon CloudWatch alarm to scale out EC2 instances when the average CPU utilization threshold is breached.

**Answer: A**

#### **QUESTION 89**

A company's application is having performance issues. The application is stateful and needs to complete in-memory tasks on Amazon EC2 instances. The company used AWS CloudFormation to deploy infrastructure and used the M5 EC2 Instance family. As traffic increased, the application performance degraded. Users are reporting delays when they attempt to access the application.

Which solution will resolve these issues in the MOST operationally efficient way?

- A. Replace the EC2 instances with T3 EC2 instances that run in an Auto Scaling group.  
Make the changes by using the AWS Management Console.
- B. Modify the CloudFormation templates to run the EC2 instances in an Auto Scaling group.  
Increase the desired capacity and the maximum capacity of the Auto Scaling group manually when an increase is necessary.
- C. Modify the CloudFormation templates.  
Replace the EC2 instances with R5 EC2 instances.  
Use Amazon CloudWatch built-in EC2 memory metrics to track the application performance for future capacity planning.
- D. Modify the CloudFormation templates.  
Replace the EC2 instances with R5 EC2 instances.  
Deploy the Amazon CloudWatch agent on the EC2 instances to generate custom application latency metrics for future capacity planning.

**Answer: D**

**QUESTION 90**

An ecommerce company has an order-processing application that uses Amazon API Gateway and an AWS Lambda function.

The application stores data in an Amazon Aurora PostgreSQL database.

During a recent sales event, a sudden surge in customer orders occurred.

Some customers experienced timeouts and the application did not process the orders of those customers.

A solutions architect determined that the CPU utilization and memory utilization were high on the database because of a large number of open connections.

The solutions architect needs to prevent the timeout errors while making the least possible changes to the application.

Which solution will meet these requirements?

- A. Configure provisioned concurrency for the Lambda function.  
Modify the database to be a global database in multiple AWS Regions.
- B. Use Amazon RDS Proxy to create a proxy for the database.  
Modify the Lambda function to use the RDS Proxy endpoint instead of the database endpoint.
- C. Create a read replica for the database in a different AWS Region.  
Use query string parameters in API Gateway to route traffic to the read replica.
- D. Migrate the data from Aurora PostgreSQL to Amazon DynamoDB by using AWS Database Migration Service (AWS DMS) Modify the Lambda function to use the OynamoDB table.

**Answer: D**