

➤ **Vendor: Amazon**

➤ **Exam Code: AWS-Developer-Associate**

➤ **Exam Name: AWS Certified Developer - Associate (DVA-C01)**

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

**[Visit Braindump2go and Download Full Version AWS-Developer-Associate Exam Dumps](#)**

**QUESTION 580**

A Developer has discovered that an application responsible for processing messages in an Amazon SQS queue is routinely falling behind. The application is capable of processing multiple messages in one execution, but is only receiving one message at a time.

What should the Developer do to increase the number of messages the application receives?

- A. Call the ChangeMessageVisibility API for the queue and set MaxNumberOfMessages to a value greater than the default of 1.
- B. Call the AddPermission API to set MaxNumberOfMessages for the ReceiveMessage action to a value greater than the default of 1.
- C. Call the ReceiveMessage API to set MaxNumberOfMessages to a value greater than the default of 1.
- D. Call the SetQueueAttributes API for the queue and set MaxNumberOfMessages to a value greater than the default of 1.

**Answer: C**

**Explanation:**

[https://docs.aws.amazon.com/AWSSimpleQueueService/latest/APIReference/API\\_ReceiveMessage.html](https://docs.aws.amazon.com/AWSSimpleQueueService/latest/APIReference/API_ReceiveMessage.html)

**QUESTION 581**

A Developer is investigating an application's performance issues. The application consists of hundreds of microservices, and a single API call can potentially have a deep call stack. The Developer must isolate the component that is causing the issue.

Which AWS service or feature should the Developer use to gather information about what is happening and isolate the fault?

- A. AWS X-Ray
- B. VPC Flow Logs
- C. Amazon GuardDuty
- D. Amazon Macie

**Answer: A**

**Explanation:**

<https://docs.aws.amazon.com/xray/latest/devguide/xray-concepts.html>

**QUESTION 582**

A Company runs continuous integration/continuous delivery (CI/CD) pipelines for its application on AWS CodePipeline. A Developer must write unit tests and run them as part of the pipelines before staging the artifacts for testing.

**[AWS-Developer-Associate Exam Dumps](#) [AWS-Developer-Associate Exam Questions](#)**

**[AWS-Developer-Associate PDF Dumps](#) [AWS-Developer-Associate VCE Dumps](#)**

**<https://www.braindump2go.com/aws-developer-associate.html>**

How should the Developer incorporate unit tests as part of CI/CD pipelines?

- A. Create a separate CodePipeline pipeline to run unit tests
- B. Update the AWS CodeBuild specification to include a phase for running unit tests
- C. Install the AWS CodeDeploy agent on an Amazon EC2 instance to run unit tests
- D. Create a testing branch in AWS CodeCommit to run unit tests

**Answer:** D

**QUESTION 583**

An application has the following requirements:

- Performance efficiency of seconds with up to a minute of latency.
- The data storage size may grow up to thousands of terabytes.
- Per-message sizes may vary between 100 KB and 100 MB.
- Data can be stored as key/value stores supporting eventual consistency.

What is the MOST cost-effective AWS service to meet these requirements?

- A. Amazon DynamoDB
- B. Amazon S3
- C. Amazon RDS (with a MySQL engine)
- D. Amazon ElastiCache

**Answer:** A

**Explanation:**

<https://aws.amazon.com/nosql/key-value/>

**QUESTION 584**

An application is experiencing performance issues based on increased demand. This increased demand is on read-only historical records pulled from an Amazon RDS-hosted database with custom views and queries. A Developer must improve performance without changing the database structure.

Which approach will improve performance and MINIMIZE management overhead?

- A. Deploy Amazon DynamoDB, move all the data, and point to DynamoDB.
- B. Deploy Amazon ElastiCache for Redis and cache the data for the application.
- C. Deploy Memcached on Amazon EC2 and cache the data for the application.
- D. Deploy Amazon DynamoDB Accelerator (DAX) on Amazon RDS to improve cache performance.

**Answer:** B

**QUESTION 585**

A Developer has an Amazon DynamoDB table that must be in provisioned mode to comply with user requirements. The application needs to support the following:

- Average item size: 10 KB
- Item reads each second: 10 strongly consistent
- Item writes each second: 2 transactional

Which read and write capacity cost-effectively meets these requirements?

- A. Read 10; write 2
- B. Read 30; write 40
- C. Use on-demand scaling
- D. Read 300; write 400

**Answer:** A

**[AWS-Developer-Associate Exam Dumps](#) **[AWS-Developer-Associate Exam Questions](#)****

**[AWS-Developer-Associate PDF Dumps](#) **[AWS-Developer-Associate VCE Dumps](#)****

**<https://www.braindump2go.com/aws-developer-associate.html>**

**QUESTION 586**

A company wants to containerize an existing three-tier web application and deploy it to Amazon ECS Fargate. The application is using session data to keep track of user activities. Which approach would provide the BEST user experience?

- A. Provision a Redis cluster in Amazon ElastiCache and save the session data in the cluster.
- B. Create a session table in Amazon Redshift and save the session data in the database table.
- C. Enable session stickiness in the existing Network Load Balancer and manage the session data in the container.
- D. Use an Amazon S3 bucket as data store and save the session data in the bucket.

**Answer: C**

**QUESTION 587**

An application is using a single-node Amazon ElastiCache for Redis instance to improve read performance. Over time, demand for the application has increased exponentially, which has increased the load on the ElastiCache instance. It is critical that this cache layer handles the load and is resilient in case of node failures. What can the Developer do to address the load and resiliency requirements?

- A. Add a read replica instance.
- B. Migrate to a Memcached cluster.
- C. Migrate to an Amazon Elasticsearch Service cluster.
- D. Vertically scale the ElastiCache instance.

**Answer: A**

**Explanation:**

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Replication.Redis.Groups.html>

**QUESTION 588**

A Developer is designing an AWS Lambda function that create temporary files that are less than 10 MB during execution. The temporary files will be accessed and modified multiple times during execution. The Developer has no need to save or retrieve these files in the future. Where should the temporary file be stored?

- A. the /tmp directory
- B. Amazon EFS
- C. Amazon EBS
- D. Amazon S3

**Answer: A**

**Explanation:**

<https://docs.aws.amazon.com/lambda/latest/dg/lambda-dg.pdf> (23)

**QUESTION 589**

A Developer is writing an application that runs on Amazon EC2 instances in an Auto Scaling group. The application data is stored in an Amazon DynamoDB table and records are constantly updated by all instances. An instance sometimes retrieves old data. The Developer wants to correct this by making sure the reads are strongly consistent. How can the Developer accomplish this?

- A. Set ConsistentRead to true when calling GetItem.
- B. Create a new DynamoDB Accelerator (DAX) table.
- C. Set Consistency to strong when calling UpdateTable.
- D. Use the GetShardIterator command.

**[AWS-Developer-Associate Exam Dumps](#) **[AWS-Developer-Associate Exam Questions](#)****

**[AWS-Developer-Associate PDF Dumps](#) **[AWS-Developer-Associate VCE Dumps](#)****

**<https://www.braindump2go.com/aws-developer-associate.html>**

**Answer:** A

**Explanation:**

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadConsistency.html>

**QUESTION 590**

A Developer has an application that must accept a large amount of incoming data streams and process the data before sending it to many downstream users.

Which serverless solution should the Developer use to meet these requirements?

- A. Amazon RDS MySQL stored procedure with AWS Lambda
- B. AWS Direct Connect with AWS Lambda
- C. Amazon Kinesis Data Streams with AWS Lambda
- D. Amazon EC2 bash script with AWS Lambda

**Answer:** C

**Explanation:**

<https://aws.amazon.com/kinesis/data-analytics/faqs/>

**QUESTION 591**

A company is using Amazon API Gateway to manage its public-facing API. The CISO requires that the APIs be used by test account users only.

What is the MOST secure way to restrict API access to users of this particular AWS account?

- A. Client-side SSL certificates for authentication
- B. API Gateway resource policies
- C. Cross-origin resource sharing (CORS)
- D. Usage plans

**Answer:** B

**Explanation:**

<https://aws.amazon.com/blogs/compute/control-access-to-your-apis-using-amazon-api-gateway-resource-policies/>

**QUESTION 592**

A Developer is migrating existing applications to AWS. These applications use MongoDB as their primary data store, and they will be deployed to Amazon EC2 instances. Management requires that the Developer minimize changes to applications while using AWS services.

Which solution should the Developer use to host MongoDB in AWS?

- A. Install MongoDB on the same instance where the application is running.
- B. Deploy Amazon DocumentDB in MongoDB compatibility mode.
- C. Use Amazon API Gateway to translate API calls from MongoDB to Amazon DynamoDB.
- D. Replicate the existing MongoDB workload to Amazon DynamoDB.

**Answer:** D

**[AWS-Developer-Associate Exam Dumps](#) **[AWS-Developer-Associate Exam Questions](#)****

**[AWS-Developer-Associate PDF Dumps](#) **[AWS-Developer-Associate VCE Dumps](#)****

**<https://www.braindump2go.com/aws-developer-associate.html>**