



# DVA-C01<sup>Q&As</sup>

AWS Certified Developer - Associate (DVA-C01)





**Pass Amazon DVA-C01 Exam with 100% Guarantee**

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/aws-certified-developer-associate.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon  
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





### QUESTION 1

A company is developing an application that will be accessed through the Amazon API Gateway REST API. Registered users should be the only ones who can access certain resources of this API. The token being used should expire automatically and needs to be refreshed periodically.

How can a developer meet these requirements?

- A. Create an Amazon Cognito identity pool, configure the Amazon Cognito Authorizer in API Gateway, and use the temporary credentials generated by the identity pool
- B. Create and maintain a database record for each user with a corresponding token and use an AWS Lambda authorizer in API Gateway
- C. Create an Amazon Cognito user pool, configure the Cognito Authorizer in API Gateway, and use the identity or access token
- D. Create an IAM user for each API user, attach an invoke permissions policy to the API, and use an IAM authorizer in API Gateway.

Correct Answer: C

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/cognito-custom-scopes-api-gateway/>

---

### QUESTION 2

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.

Correct Answer: C

---

### QUESTION 3

A company experienced partial downtime during the last deployment of a new application. AWS Elastic Beanstalk split the environment's Amazon EC2 instances into batches and deployed a new version one batch at a time after taking them

out of service. Therefore, full capacity was not maintained during deployment.

The developer plans to release a new version of the application, and is looking for a policy that will maintain full capacity and minimize the impact of the failed deployment.



Which deployment policy should the developer use?

- A. Immutable
- B. All at Once
- C. Rolling
- D. Rolling with an Additional Batch

Correct Answer: D

Reference: <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html>

---

#### QUESTION 4

A gaming company has deployed a web portal on AWS Elastic Beanstalk. The company sometimes needs to deploy new versions three or four times in a day. The company needs to deploy new features for all users as quickly as possible. The solution must minimize performance impact and must maximize availability.

What solution will meet these requirements?

- A. Use a rolling deployment policy to deploy to Amazon EC2 instances.
- B. Use an immutable deployment policy to deploy to Amazon EC2 instances.
- C. Use an all-at-once deployment policy to deploy to Amazon EC2 instances.
- D. Use a canary deployment strategy to deploy changes to Amazon EC2 instances.

Correct Answer: A

Reference: <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html>

---

#### QUESTION 5

An application uses Amazon Kinesis Data Streams to ingest and process large streams of data records in real time. Amazon EC2 instances consume and process the data from the shards of the Kinesis data stream by using Amazon Kinesis Client Library (KCL). The application handles the failure scenarios and does not require standby workers. The application reports that a specific shard is receiving more data than expected. To adapt to the changes in the rate of data flow, the "hot" shard is resharded.

Assuming that the initial number of shards in the Kinesis data stream is 4, and after resharding the number of shards increased to 6, what is the maximum number of EC2 instances that can be deployed to process data from all the shards?

- A. 12
- B. 6
- C. 4



D. 1

Correct Answer: B

Typically, when you use the KCL, you should ensure that the number of instances does not exceed the number of shards (except for failure standby purposes). Each shard is processed by exactly one KCL worker and has exactly one corresponding record processor, so you never need multiple instances to process one shard. However, one worker can process any number of shards, so it's fine if the number of shards exceeds the number of instances. <https://docs.aws.amazon.com/streams/latest/dev/kinesis-record-processor-scaling.html>

[Latest DVA-C01 Dumps](#)

[DVA-C01 PDF Dumps](#)

[DVA-C01 Practice Test](#)