

QUESTIONS & ANSWERS

Kill your exam at first Attempt



DVA-C01 Dumps
DVA-C01 Braindumps
DVA-C01 Real Questions
DVA-C01 Practice Test
DVA-C01 dumps free



Amazon

DVA-C01

AWS Certified Developer -Associate (DVA-C01)

<http://killexams.com/pass4sure/exam-detail/DVA-C01>



Question #224 Section 2

A startup s photo-sharing site is deployed in a VPC. An ELB distributes web traffic across two subnets. ELB session stickiness is configured to use the AWS- generated session cookie, with a session TTL of 5 minutes. The webserver Auto Scaling Group is configured as: min-size=4, max-size=4. The startups preparing for a public launch, by running load-testing software installed on a single EC2 instance running in us-west-2 After 60 minutes of load-testing, the webserver logs show:

WEBSERVER LOGS	# of HTTP requests from load-tester	# of HTTP requests from private beta users
webserver #1 (subnet in us-west-2a):	19,210	434
webserver #2 (subnet in us-west-2a):	21,790	490
webserver #3 (subnet in us-west-2b):	0	410
webserver #4 (subnet in us-west-2b):	0	428

Which recommendations can help ensure load-testing HTTP requests are evenly distributed across the four webservers? (Choose two.)

- A. Launch and run the load-tester EC2 instance from us-east-1 instead.
- B. Re-configure the load-testing software to re-resolve DNS for each web request.
- C. Use a 3rd-party load-testing service which offers globally-distributed test clients.
- D. Configure ELB and Auto Scaling to distribute across us-west-2a and us-west-2c.
- E. Configure ELB session stickiness to use the app-specific session cookie.

Answer: BE

Question #225 Section 2

A development team uses AWS Elastic Beanstalk to deploy a Java-based web application. The team wants to ensure that the changes to the source code and the configuration are always deployed on new instances. The team configures the Elastic Beanstalk environment to use immutable updates. However, an error occurs the first time a change is deployed with the new update policy. What is the MOST likely cause of this issue?

- A. Immutable updates are not supported for Java-based applications.
- B. The account has reached its on-demand instance limit.
- C. Immutable updates are only supported for m4.large and larger instance types.
- D. The developer must also modify the .ebextensions/immutable-updates.config file to enable immutable updates.

Answer: D

Reference:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environmentmgmt-updates-immutable.html>

Question #226 Section 2

A developer tested an application locally and then deployed it to AWS Lambda. While testing the application remotely, the Lambda function fails with an access denied message. How can this issue be addressed?

- A. Update the Lambda function's execution role to include the missing permissions.
- B. Update the Lambda function's resource policy to include the missing permissions.
- C. Include an IAM policy document at the root of the deployment package and redeploy the Lambda function.
- D. Redeploy the Lambda function using an account with access to the AdministratorAccess policy.

Answer: A

Reference:

<https://aws.amazon.com/premiumsupport/knowledge-center/access-denied-lambda-s3-bucket/>

Question #227 Section 2

An application contains two components: one component to handle HTTP requests, and another component to handle background processing tasks. Each component must scale independently. The developer wants to deploy this application using AWS Elastic Beanstalk.

How should this application be deployed, based on these requirements?

- A. Deploy the application in a single Elastic Beanstalk environment.
- B. Deploy each component in a separate Elastic Beanstalk environment.
- C. Use multiple Elastic Beanstalk environments for the HTTP component, but one environment for the background task component.
- D. Use multiple Elastic Beanstalk environments for the background task component, but one environment for the HTTP component.

Answer: D

Question #228 Section 2

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

Answer: D

Reference:

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

Question #229 Section 2

An application running on multiple Amazon EC2 instances pulls messages from a standard Amazon SQS queue. A requirement for the application is that all messages must be encrypted at rest.

Developers are instructed to use methods that allow for centralized key management and minimize possible support requirements whenever possible.

Which of the following solutions supports these requirements?

- A. Encrypt individual messages by using client-side encryption with customer managed keys, then write to the SQS queue.
- B. Encrypt individual messages by using SQS Extended Client and the Amazon S3 encryption client.
- C. Create an SQS queue, and encrypt the queue by using server-side encryption with AWS KMS.
- D. Create an SQS queue, and encrypt the queue by using client-side encryption.

Answer: B

Question #230 Section 2

A company is developing a serverless ecommerce web application. The application needs to make coordinated, all-or-nothing changes to multiple items in the company's inventory table in Amazon DynamoDB.

Which solution will meet these requirements?

- A. Enable transactions for the DynamoDB table. Use the BatchWriteItem operation to update the items.
- B. Use the TransactWriteItems operation to group the changes. Update the items in the table.
- C. Set up a FIFO queue using Amazon SQS. Group the changes in the queue. Update the table based on the grouped changes.
- D. Create a transaction table in an Amazon Aurora DB cluster to manage the transactions. Write a backend process to sync the Aurora DB table and the DynamoDB table.

Answer: B

Reference:

<https://aws.amazon.com/blogs/mobile/appsync-caching-transactions/>

Question #231 Section 2

How can a developer use a debugger for AWS Lambda code that is deployed with AWS Serverless Application Model (AWS SAM)?

- A. Download the Lambda code locally and use the AWS CLI to execute it
- B. Use the Lambda console to connect the debugger
- C. Use AWS SAM to invoke a function locally in debug mode
- D. Connect a third-party-compatible integrated development environment (IDE) to the Lambda debugger endpoint

Answer: C

Reference:

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverless-sam-cli-using-debugging.html>

Question #232 Section 2

An application takes longer than expected to process an Amazon SQS message.

What should the developer do to the application so that other instances do not pick up the same message?

- A. Make a ReceiveMessage call to get the same message again from the queue
- B. Issue a DeleteMessage call to delete the message from the queue
- C. Use SendMessage to pass the message to the dead letter queue
- D. Send a ChangeMessageVisibility call to extend VisibilityTimeout

Answer: A

Question #233 Section 2

A developer is building a WebSocket API using Amazon API Gateway. The payload sent to this API is JSON that includes an action key. This key can have three different values: create, update, and remove. The developer must integrate with different routes based on the value of the action key of the incoming JSON payload.

How can the developer accomplish this task with the LEAST amount of configuration?

- A. Deploy the WebSocket API to three stages for the respective routes: create, update, and remove
- B. Create a new route key and set the name as action
- C. Set the value of the route selection expression to action
- D. Set the value of the route selection expression to \$request.body.action

Answer: B

Question #234 Section 2

A development team is creating a new application designed to run on AWS. While the test and production environments will run on Amazon EC2 instances, developers will each run their own environment on their laptops.

Which of the following is the simplest and MOST secure way to access AWS services from the local development machines?

- A. Use an IAM role to assume a role and execute API calls using the role.
- B. Create an IAM user to be shared with the entire development team; provide the development team with the access key.
- C. Create an IAM user for each developer on the team; provide each developer with a unique access key.
- D. Set up a federation through an Amazon Cognito user pool.

Answer: D

Question #235 Section 2

A developer wants to ensure the Amazon EC2 instances in AWS Elastic Beanstalk execute a certain set of commands before the application is ready to

use.

Which Elastic Beanstalk feature will allow the developer to accomplish this?

- A. Rolling update
- B. Immutable update
- C. User data
- D. .ebextensions

Answer: D

Reference:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/customize-containers-ec2.html>

Question #236 Section 2

A developer is planning to use an Amazon API Gateway and AWS Lambda to provide a REST API. The developer will have three distinct environments to manage: development, test, and production.

How should the application be deployed while minimizing the number of resources to manage?

- A. Create a separate API Gateway and separate Lambda function for each environment in the same Region.
- B. Assign a Region for each environment and deploy API Gateway and Lambda to each Region.
- C. Create one API Gateway with multiple stages with one Lambda function with multiple aliases.
- D. Create one API Gateway and one Lambda function, and use a REST parameter to identify the environment.

Answer: C

Reference:

<https://aws.amazon.com/blogs/compute/using-api-gateway-stage-variables-to-manage-lambda-functions/>

Question #237 Section 2

A developer is creating an application to process a large number of requests. Requests must be processed in order, and each request should be processed only once.

How should Amazon SQS be deployed to achieve this?

- A. Configure First in First out (FIFO) delivery in a standard Amazon SQS queue to process requests.
- B. Use an SQS FIFO queue to process requests.
- C. Use the SetOrder attribute to ensure sequential request processing.
- D. Convert the standard queue to a FIFO queue by renaming the queue to use the .fifo suffix.

Answer: B

Reference:

<https://medium.com/awesome-cloud/aws-difference-between-sqs-standard-and-fifo-first-in-first-out-queues-28d1ea5e153>

For More exams visit <https://killexams.com/vendors-exam-list>



Kill your exam at First Attempt....Guaranteed!