DUMPSARENA

AWS Certified Developer Associate (DVA-C01)

Amazon AWS AWS-Certified-Developer-Associate-DVA-C01

Version Demo

Total Demo Questions: 20

Total Premium Questions: 712

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

A Developer has created a software package to be deployed on multiple EC2 instances using IAM roles.

What actions could be performed to verify IAM access to get records from Amazon Kinesis Streams? (Choose two.)

A. Use the AWS CLI to retrieve the IAM group.

B. Query Amazon EC2 metadata for in-line IAM policies.

C. Request a token from AWS STS, and perform a describe action.

D. Perform a get action using the –-dry-run argument.

E. Validate the IAM role policy with the IAM policy simulator.

ANSWER: B E

QUESTION NO: 2

A company's security policies require all database passwords to be rotated every 30 days The company uses different database platforms, including Amazon Aurora databases and proprietary NoSQL document databases, for different applications A developer needs to implement a solution for password rotation

Which solution will meet these requirements?

A. Create an AWS Lambda rotation function that has appropriate IAM permissions Store the password in AWS Secrets Manager Configure Secrets Manager to rotate the password by using the Lambda function

B. Encrypt the existing password with AWS Key Management Service (AWS KMS) Export the existing password Generate a random password with AWS KMS Use the AWS KMS password renewal feature to replace the existing password with the new password.

C. Create an AWS Lambda rotation function that has appropriate IAM permissions Store the password in AWS Systems Manager Parameter Store Configure Parameter Store to rotate the password by using the Lambda function

D. Integrate AWS Systems Manager Parameter Store with a Key Management Interoperability Protocol (KMIP)-compliant third-party secret manager to enable third-party database password rotation on AWS

ANSWER: C

QUESTION NO: 3

A Developer wants to encrypt new objects that are being uploaded to an Amazon S3 bucket by an application. There must be an audit trail of who has used the key during this process. There should be no change to the performance of the application.

Which type of encryption meets these requirements?

A. Server-side encryption using S3-managed keys

B. Server-side encryption with AWS KMS-managed keys

C. Client-side encryption with a client-side symmetric master key

D. Client-side encryption with AWS KMS-managed keys

ANSWER: B

QUESTION NO: 4

An application is real-time processing millions of events that are received through an API.

What service could be used to allow multiple consumers to process the data concurrently and MOST cost-effectively?

A. Amazon SNS with fanout to an SQS queue for each application

B. Amazon SNS with fanout to an SQS FIFO (first-in, first-out) queue for each application

C. Amazon Kinesis Firehose

D. Amazon Kinesis Streams

ANSWER: D

Explanation:

Reference: https://aws.amazon.com/kinesis/data-streams/getting-started/

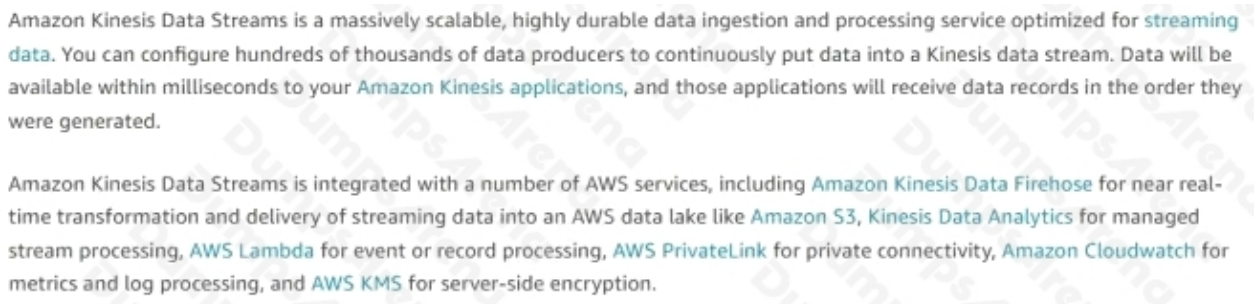

QUESTION NO: 5

A Developer wants to use AWS X-Ray to trace a user request end-to-end throughput the software stack. The Developer made the necessary changes in the application, tested it, and found that the application is able to send the traces to AWS X-Ray. However, when the application is deployed to an EC2 instance, the traces are not available.

Which of the following could create this situation? (Choose two.)

A. The traces are reaching X-Ray, but the Developer does not have access to view the records.

B. The X-Ray daemon is not installed on the EC2 instance.

C. The X-Ray endpoint specified in the application configuration is incorrect.

D. The instance role does not have “xray:BatchGetTraces” and “xray:GetTraceGraph” permissions.

E. The instance role does not have “xray:PutTraceSegments” and “xray:PutTelemetryRecords” permissions.

ANSWER: B E

QUESTION NO: 6

A company is using an AWS Lambda function to process records from an Amazon Kinesis data stream The company recently observed slow processing of the records. A developer notices that the iterator age metric for the function is Increasing and that the Lambda run duration is constantly above normal.

Which actions should the developer take to increase the processing speed? (Select TWO.)

A. Increase the number of shards of the Kinesis data stream

B. Decrease the timeout of the Lambda function

C. Increase the memory that is allocated to the Lambda function.

D. Decrease the number of shards of the Kinesis data stream

E. Increase the timeout of the Lambda function

ANSWER: D E

QUESTION NO: 7

A company is creating a REST service using an Amazon API Gateway with AWS Lambda integration. The service must run different versions for testing purposes.

What would be the BEST way to accomplish this?

A. Use an X-Version header to denote which version is being called and pass that header to the Lambda function(s)

B. Create an API Gateway Lambda authorizer to route API clients to the correct API version

C. Create an API Gateway resource policy to isolate versions and provide context to the Lambda function(s)

D. Deploy the API versions as unique stages with unique endpoints and use stage variables to provide further context

ANSWER: C

QUESTION NO: 8

A company has a two-tier application running on an Amazon EC2 server that handles all of its AWS based e-commerce activity During peak times, the backend servers that process orders are overloaded with requests. This results in some orders failing to process. A developer needs to create a solution that will re-factor the application.

Which steps will allow for more flexibility during peak times, while still remaining cost-effective? (Select TWO.)

A. Increase the backend T2 EC2 instance sees to xl to handle the largest possible load throughout the year

B. implement an Amazon SQS queue to decouple the front-end and backend servers

C. Use an Amazon SNS queue to decouple the front-end and backend servers.

D. Migrate the backend servers to on-premises and pull from an Amazon SNS queue

E. Modify the backend servers to pull from an Amazon SQS queue.

ANSWER: C D

QUESTION NO: 9

An application is processing clickslream data using Amazon Kinesis. The clickstream data feed into Kinesis experiences periodic spikes. The PutRecords API call occasionally fails and the logs show that the failed call returns the response shown below.

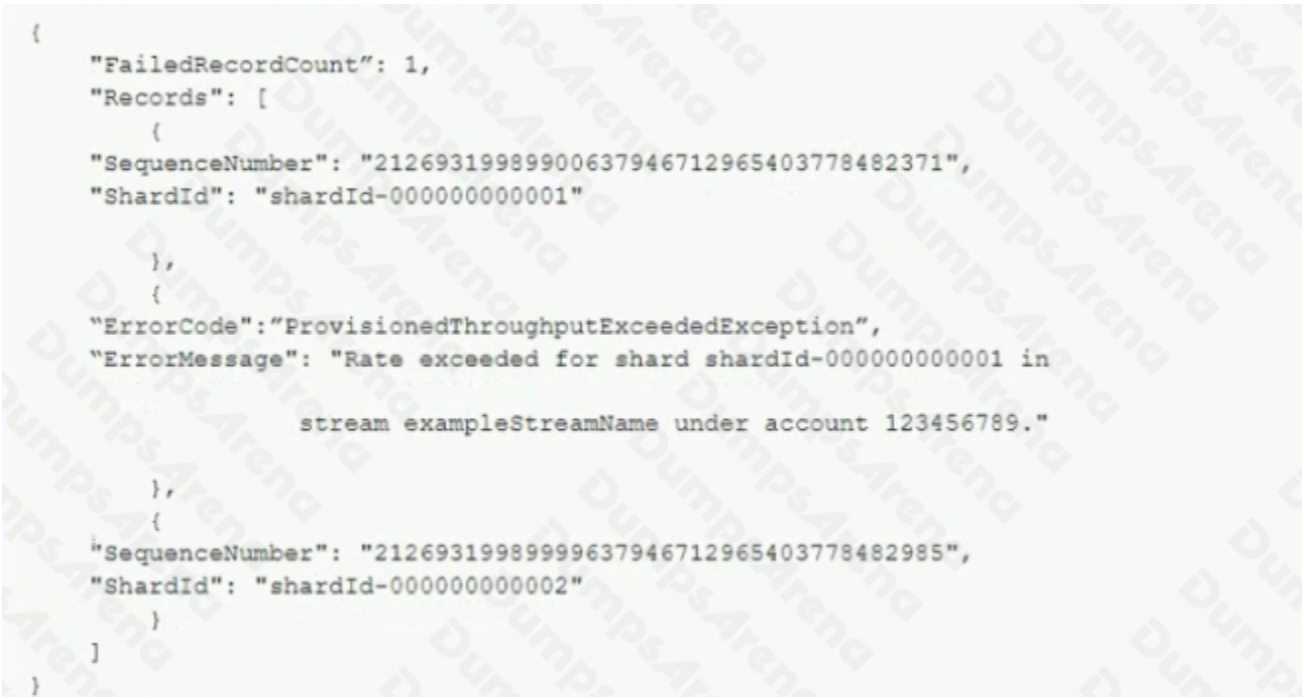

Which techniques will help mitigate this exception? (Select TWO.)

A. Implement retries with exponential backoff

B. Use a PutRecord API instead of PutRecords

C. Reduce the frequency and/or size of the requests

D. Use Amazon SNS instead of Kinesis.

E. Reduce the number of KCL consumers.

ANSWER: A C

QUESTION NO: 10

A developer is troubleshooting a three-tier application, which is deployed on Amazon EC2 instances. There is a connectivity problem between the application servers and the database servers.

Which AWS services or tools should be used to identify the faulty component? (Choose two.)

A. AWS CloudTrail

B. AWS Trusted Advisor

C. Amazon VPC Flow Logs

D. Network access control lists

E. AWS Config rules

ANSWER: C D

QUESTION NO: 11

A developer is building a backend system for the long-term storage of information from an inventory management system. The information needs to be stored so that other teams can build tools to report and analyze the data

How should the developer implement this solution to achieve the FASTEST running time?

A. Create an AWS Lambda function that writes to Amazon S3 synchronously Increase the function's concurrency to match the highest expected value of concurrent scans and requests.

B. Create an AWS Lambda function that writes to Amazon S3 asynchronously Configure a dead-letter queue to collect unsuccessful invocations

C. Create an AWS Lambda function that writes to Amazon S3 synchronously Set the inventory system to retry failed requests.

D. Create an AWS Lambda function that writes to an Amazon ElastiCache for Redis cluster asynchronously Configure a dead-letter queue to collect unsuccessful invocations.

ANSWER: A

QUESTION NO: 12

The Lambda function below is being called through an API using Amazon API Gateway. The average execution time for the Lambda function is about 1 second. The pseudocode for the Lambda function is as shown in the exhibit.

```
include "3rd party encryption module"
include "math module"
lambda_ handler(event, context)
     rds_host = "rds-instance-endpoint"
     name = db_username
     password = db_password
     db_name = db_name
# Connect to the RDS Database
conn = RDSConnection(rds_host, user=name, passwd=password,
db=db_name, connect_timeout=5)
#Perform some Processing reading data from the RDS database
#Code Block
#Code Block
#Code Block
```

What two actions can be taken to improve the performance of this Lambda function without increasing the cost of the solution? (Choose two.)

A. Package only the modules the Lambda function requires

B. Use Amazon DynamoDB instead of Amazon RDS

C. Move the initialization of the variable Amazon RDS connection outside of the handler function

D. Implement custom database connection pooling with the Lambda function

E. Implement local caching of Amazon RDS data so Lambda can re-use the cache

ANSWER: A C

QUESTION NO: 13

A Developer has written code for an application and wants to share it with other Developers on the team to receive feedback. The shared application code needs to be stored long-term with multiple versions and batch change tracking.

Which AWS service should the Developer use?

A. AWS CodeBuild

B. Amazon S3

C. AWS CodeCommit

D. AWS Cloud9

ANSWER: C

Explanation:

Reference: https://docs.aws.amazon.com/codecommit/latest/userguide/codecommit-user.pdf

QUESTION NO: 14

A Developer is making changes to a custom application that is currently using AWS Elastic Beanstalk.

After the Developer completes the changes, what solutions will update the Elastic Beanstalk environment with the new application version? (Choose two.)

A. Package the application code into a .zip file, and upload, then deploy the packaged application from the AWS Management Console

B. Package the application code into a .tar file, create a new application version from the AWS Management Console, then update the environment by using AWS CLI

C. Package the application code into a .tar file, and upload and deploy the packaged application from the AWS Management Console

D. Package the application code into a .zip file, create a new application version from the packaged application by using AWS CLI, then update the environment by using AWS CLI

E. Package the application code into a .zip file, create a new application version from the AWS Management Console, then rebuild the environment by using AWS CLI

ANSWER: C D

QUESTION NO: 15

A developer at a company writes an AWS CloudFormation template. The template refers to subnets that were created by a separate AWS CloudFormation template that the company's network team wrote. When the developer attempts to launch the stack for the first time, the launch fails.

Which template coding mistakes could have caused this failure? (Choose two.)

A. The developer's template does not use the Ref intrinsic function to refer to the subnets.

B. The developer's template does not use the ImportValue intrinsic function to refer to the subnets.

C. The Mappings section of the developer's template does not refer to the subnets.

D. The network team's template does not export the subnets in the Outputs section.

E. The network team's template does not export the subnets in the Mappings section.

ANSWER: B D

Explanation:

Reference: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference.html

QUESTION NO: 16

A Developer is building a serverless application using AWS Lambda and must create a REST API using an HTTP GET method.

What needs to be defined to meet this requirement? (Choose two.)

A. A Lambda@Edge function

B. An Amazon API Gateway with a Lambda function

C. An exposed GET method in an Amazon API Gateway

D. An exposed GET method in the Lambda function

E. An exposed GET method in Amazon Route 53

ANSWER: B C

Explanation:

Reference: https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-getting-started-with-rest-apis.html

QUESTION NO: 17

A company hosts a monolithic application on Amazon EC2 instances. The company starts converting some features of the application to a serverless architecture by using Amazon API Gateway and AWS Lambda After the migration, some users report problems with payment processing

Upon inspection, a developer discovers that the Lambda function that calls the external payment API is taking longer than expected Therefore, the API Gateway requests are timing out

What should the developer do to resolve this issue in the serverless architecture?

A. Use the EC2 instances to make the API calls to the payment API

B. Use Amazon Simple Queue Service (Amazon SQS) with API Gateway and the Lambda function to asynchronously call the payment API

C. Increase the API Gateway timeout duration to match the payment API time

D. Increase the Lambda function's memory to increase the network bandwidth and increase the speed of the payment API calls

ANSWER: B

QUESTION NO: 18

A company has developed a new serverless application using AWS Lambda functions that will be deployed using the AWS Serverless Application Model (AWS SAM) CLI.

Which step should the developer complete prior to deploying the application?

A. Compress the application to a .zip file and upload it into AWS Lambda

B. Test the new AWS Lambda function by first tracing it in AWS X-Ray

C. Bundle the serverless application using a SAM package

D. Create the application environment using the eb create my-env command

ANSWER: A

Explanation:

Reference: https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverless-application-model.pdf

QUESTION NO: 19

An application is using a custom library to make HTTP calls directly to AWS service endpoints. The application is experiencing transient errors that are causing processes to stop when each error is first encountered. A request has been made to make the application more resilient by adding error retries and exponential backoff.

How should a developer implement the changes with MINIMAL custom code?

A. Add a Retry-After HTTP header to API requests

B. Use the AWS CLI to configure the retry settings in a named profile

C. Change the custom library to retry on 5xx errors only

D. Use an AWS SDK and set retry-specific configurations

ANSWER: B

Explanation:

Reference: https://docs.aws.amazon.com/cli/latest/topic/config-vars.html

QUESTION NO: 20

A company is hosting a workshop for external users and wants to share the reference documents with the external users for 7 days. The company stores the reference documents in an Amazon S3 bucket that the company owns.

What is the MOST secure way to share the documents with the external users?

A. Use S3 presigned URLs to share the documents with the external users. Set an expiration time of 7 days.

B. Move the documents to an Amazon WorkDocs folder. Share the links of the WorkDocs folder with the external users.

C. Create temporary IAM users that have read-only access to the S3 bucket. Share the access keys with the external users. Expire the credentials after 7 days.

D. Create a role that has read-only access to the S3 bucket. Share the Amazon Resource Name (ARN) of this role with the external users.

ANSWER: A

Explanation:

Reference: https://docs.aws.amazon.com/AmazonS3/latest/userguide/ShareObjectPreSignedURL.html