

DUMPS ARENA

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.com>

sales@dumpsarena.com

dumpsarena.com

QUESTION NO: 1

A company runs an e-commerce website that uses Amazon DynamoDB where pricing for items is dynamically updated in real time. At any given time, multiple updates may occur simultaneously for pricing information on a particular product. This is causing the original editor's changes to be overwritten without a proper review process.

Which DynamoDB write option should be selected to prevent this overwriting?

- A. Concurrent writes
- B. Conditional writes
- C. Atomic writes
- D. Batch writes

ANSWER: B**QUESTION NO: 2**

A company maintains an application that uses an Amazon RDS DB instance for its database. A developer needs to implement encryption at rest for the database. Which combination of steps should the developer take to meet this requirement? (Choose two.)

- A. Enable encryption on the DB instance in the AWS Management Console.
- B. Stop the DB instance.
- C. Restore the DB instance from the encrypted snapshot.
- D. Take a snapshot of the DB instance, and create an encrypted copy of the snapshot.
- E. Create a customer managed key in AWS Key Management Service (AWS KMS).

ANSWER: C D**Explanation:**

Reference: <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html>

QUESTION NO: 3

A developer is running an application on an Amazon EC2 instance. When the application tries to read an Amazon S3 bucket the application fails. The developer notices that the associated IAM role is missing the S3 read permission. The developer needs to give the application the ability to read the S3 bucket.

Which solution will meet this requirement with the LEAST application disruption?

- A. Add the permission to the role. Terminate the existing EC2 instance. Launch a new EC2 instance
- B. Add the permission to the role so that the change will take effect automatically
- C. Add the permission to the role. Hibernate and restart the existing EC2 instance.
- D. Add the permission to the S3 bucket. Restart the EC2 instance.

ANSWER: D

Explanation:

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/ec2-instance-access-s3-bucket/>

QUESTION NO: 4

An organization is using Amazon CloudFront to ensure that its users experience low-latency access to its web application. The organization has identified a need to encrypt all traffic between users and CloudFront, and all traffic between CloudFront and the web application.

How can these requirements be met? (Choose two.)

- A. Use AWS KMS to encrypt traffic between CloudFront and the web application.
- B. Set the Origin Protocol Policy to "HTTPS Only".
- C. Set the Origin's HTTP Port to 443.
- D. Set the Viewer Protocol Policy to "HTTPS Only" or "Redirect HTTP to HTTPS".
- E. Enable the CloudFront option Restrict Viewer Access.

ANSWER: C D

QUESTION NO: 5

A developer is deploying an application on Amazon EC2 instances that run in Account

A. The application needs to read data from an existing Amazon Kinesis data stream in Account B.
Which actions should the developer take to provide the application with access to the stream? (Choose two.)
Update the instance profile role in Account A with stream read permissions

- B. Create an IAM role with stream read permissions in Account B
- C. Add a trust policy to the instance profile role and IAM role in Account B to allow the instance profile role to assume the IAM role
- D. Add a trust policy to the instance profile role and IAM role in Account B to allow reads from the stream

E. Add a resource-based policy in Account B to allow read access from the instance profile role

ANSWER: A B

Explanation:

Reference: <https://docs.databricks.com/administration-guide/cloud-configurations/aws/iam-kinesis.html>

QUESTION NO: 6

A company is using AWS CodePipeline to deliver one of its applications. The delivery pipeline is triggered by changes to the master branch of an AWS CodeCommit repository and uses AWS CodeBuild to implement the test and build stages of the process and AWS CodeDeploy to deploy the application.

The pipeline has been operating successfully for several months and there have been no modifications. Following a recent change to the application's source code, AWS CodeDeploy has not deployed the updates application as expected.

What are the possible causes? (Choose two.)

- A. The change was not made in the master branch of the AWS CodeCommit repository.
- B. One of the earlier stages in the pipeline failed and the pipeline has terminated.
- C. One of the Amazon EC2 instances in the company's AWS CodePipeline cluster is inactive.
- D. The AWS CodePipeline is incorrectly configured and is not executing AWS CodeDeploy.
- E. AWS CodePipeline does not have permissions to access AWS CodeCommit.

ANSWER: B C

QUESTION NO: 7

A developer must extend an existing application that is based on the AWS Services Application Model (AWS SAM). The developer has used the AWS SAM CLI to create the project. The project contains different AWS Lambda functions.

Which combination of commands must the developer use to redeploy the AWS SAM application (Select TWO.)

- A. Sam init
- B. Sam validate
- C. Sam build
- D. Sam deploy
- E. Sam publish

ANSWER: A D

QUESTION NO: 8

A developer is creating a script to automate the deployment process for a serverless application. The developer wants to use an existing AWS Serverless Application Model (AWS SAM) template for the application

What should the developer use for the project? (Select TWO)

- A. Callaws cloudformation package to create the deployment package Call aws cloudformation deploy to deploy the package afterward.
- B. Call sam package to create the deployment package Call sam deploy to deploy the package afterward
- C. Callaws s3 cp to upload the AWS SAM template to Amazon S3 Call aws lambda update-function-code to create the application.
- D. Create a ZIP package locally and call aws serverlessrepo create-applicarion to create the application.
Create a ZIP package and upload it to Amazon S3 Callaws cloudfonnation create-stack to create the application

ANSWER: B C**QUESTION NO: 9**

A development team is designing a mobile app that requires multi-factor authentication.

Which steps should be taken to achieve this? (Choose two.)

- A. Use Amazon Cognito to create a user pool and create users in the user pool.
- B. Send multi-factor authentication text codes to users with the Amazon SNS Publish API call in the app code.
- C. Enable multi-factor authentication for the Amazon Cognito user pool.
- D. Use AWS IAM to create IAM users.
- E. Enable multi-factor authentication for the users created in AWS IAM.

ANSWER: C E**Explanation:**

Reference: <https://docs.aws.amazon.com/cognito/latest/developerguide/user-pool-settings-mfa.html#:~:text=To%20configure%20MFA%20in%20the,the%20risk%2Dbased%20adaptive%20authentication>
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_mfa_enable_virtual.html

QUESTION NO: 10

A developer is troubleshooting connectivity issues between an AWS Lambda function and an Amazon EC2 instance that runs Amazon Linux 2. The Lambda function and the EC2 instance cannot communicate with each other even though the Lambda function is configured to access resources in the EC2 instance's subnet.

How can the developer inspect the network traffic between the Lambda function and the EC2 instance?

- A. Inspect the VPC flow logs for network activity.
- B. Use the traceroute command on the EC2 instance to check connectivity.
- C. Analyze the Amazon CloudWatch metrics for network traffic.
- D. Use the telnet command on the EC2 instance to check connectivity.

ANSWER: A

Explanation:

Reference: <https://aws.amazon.com/blogs/aws/vpc-flow-logs-log-and-view-network-traffic-flows/>

QUESTION NO: 11

A developer has code stored in an Amazon S3 bucket. The code must be deployed as an AWS Lambda function across multiple accounts in the same Region as the S3 bucket. The Lambda function will be deployed using an AWS CloudFormation template that is run for each account.

What is the MOST secure approach to allow access to the Lambda code in the S3 bucket?

- A. Grant the CloudFormation execution role S3 list and get permissions. Add a bucket policy to Amazon S3 with the Principal of "AWS": [account numbers].
- B. Grant the CloudFormation execution role S3 get permissions. Add a bucket policy to Amazon S3 with the Principal of "".
- C. Use a service-based link to grant the Lambda function S3 list and get permissions by explicitly adding the S3 bucket's account number in the resource.
- D. Use a service-based link to grant the Lambda function S3 get permissions and add a Resource of "*" to allow access to the S3 bucket.

ANSWER: A

QUESTION NO: 12

According to best practice, how should access keys be managed in AWS? (Choose two.)

- A. Use the same access key in all applications for consistency.
- B. Delete all access keys for the account root user.
- C. Leave unused access keys in the account for tracking purposes.
- D. Embed and encrypt access keys in code for continuous deployment.
- E. Use Amazon IAM roles instead of access keys where possible.

ANSWER: B E**Explanation:**

Reference: <https://docs.aws.amazon.com/general/latest/gr/aws-access-keys-best-practices.html#iam-user-access-keys>

QUESTION NO: 13

A developer is receiving an intermittent `ProvisionedThroughputExceededException` error from an application that is based on Amazon DynamoDB. According to the Amazon CloudWatch metrics for the table, the application is not exceeding the provisioned throughput

What could be the cause of the issue?

- A.** The DynamoDB table storage size is larger than the provisioned size.
- B.** The application is exceeding capacity on a particular hash key.
- C.** The DynamoDB table is exceeding the provisioned scaling operations.
- D.** The application is exceeding capacity on a particular sort key.

ANSWER: B**QUESTION NO: 14**

A company has a large number of documents that are stored securely in Amazon S3. The company is creating an application that occasionally will read these documents. The application will be deployed on Amazon EC2 instances.

The company's security requirements mandate that no long-term credentials can be stored on the EC2 instances and that only the needed documents can be accessed. Only authorized users and applications can access the documents. Access must be logged by Amazon S3, and each document must follow S3 Lifecycle policies for archival and destruction.

What should a developer do to meet these requirements?

- A.** Create an event to invoke an AWS Lambda function when a document is uploaded. Configure the function to write the documents to an Amazon Elastic File System (Amazon EFS) file system. Configure the EC2 instances to mount the EFS file system. Configure the application to access the documents that are stored in the file system as needed.
- B.** Create a user that has programmatic credentials, and attach a policy that allows read access to the S3 bucket. Use the AWS CLI to configure those credentials for the EC2 instances to use. Create an Amazon Machine Image (AMI), and add the access key and secret access key to the user data section to create environment variables. Use the AMI to launch each EC2 instance that runs the application. Add application code to use the keys that are stored in the environment variables to access the S3 bucket objects as needed.
- C.** Modify the S3 bucket, make the bucket public, and make each object public. Add application code to make REST calls to access the objects in the S3 bucket as needed.
- D.** Create an IAM role with permissions to read objects from Amazon S3. Attach the role to the EC2 instances as an instance profile. Add application code to access the objects in the S3 bucket as needed.

ANSWER: D**QUESTION NO: 15**

A company wants to implement a continuous integration for its workloads on AWS. The company wants to trigger unit test in its pipeline for commits-on its code repository, and wants to be notified of failure events in the pipeline.

How can these requirements be met?

- A.** Store the source code in AWS CodeCommit. Create a CodePipeline to automate unit testing. Use Amazon SNS to trigger notifications of failure events.
- B.** Store the source code in GitHub. Create a CodePipeline to automate unit testing. Use Amazon SES to trigger notifications of failure events.
- C.** Store the source code on GitHub. Create a CodePipeline to automate unit testing. Use Amazon CloudWatch to trigger notifications of failure events.
- D.** Store the source code in AWS CodeCommit. Create a CodePipeline to automate unit testing. Use Amazon CloudWatch to trigger notification of failure events.

ANSWER: A**QUESTION NO: 16**

A company has an application where reading objects from Amazon S3 is based on the type of user. The user types are registered user and guest user. The company has 25,000 users and is growing. Information is pulled from an S3 bucket depending on the user type.

Which approaches are recommended to provide access to both user types? (Choose two.)

- A.** Provide a different access key and secret access key in the application code for registered users and guest users to provide read access to the objects.
- B.** Use S3 bucket policies to restrict read access to specific IAM users.
- C.** Use Amazon Cognito to provide access using authenticated and unauthenticated roles.
- D.** Create a new IAM user for each user and grant read access.
- E.** Use the AWS IAM service and let the application assume the different roles using the AWS Security Token Service (AWS STS) AssumeRole action depending on the type of user and provide read access to Amazon S3 using the assumed role.

ANSWER: A B**QUESTION NO: 17**

A developer at a company needs to create a small application that makes the same API call once each day at a designated time. The company does not have infrastructure in the AWS Cloud yet, but the company wants to implement this functionality on AWS.

Which solution meets these requirements in the MOST operationally efficient manner?

- A. Use a Kubernetes cron job that runs on Amazon Elastic Kubernetes Service (Amazon EKS).
- B. Use an Amazon Linux crontab scheduled job that runs on Amazon EC2.
- C. Use an AWS Lambda function that is invoked by an Amazon EventBridge (Amazon CloudWatch Events) scheduled event.
- D. Use an AWS Batch job that is submitted to an AWS Batch job queue.

ANSWER: D

QUESTION NO: 18

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: 19

An application adds a processing date to each transaction that it receives. The application writes each transaction to an Amazon DynamoDB table by using the PutItem operation. Each transaction has a unique ID (transactionID). Sometimes the application receives transactions more than once.

A developer notices that duplicate transactions in DynamoDB have the latest processing date instead of the date when the transaction was first received. Duplicate records happen infrequently, and most of the transactions are unique.

What is the MOST cost-effective solution that the developer can implement to ensure that PutItem does not update an existing record?

- A. Call the GetItem operation first to confirm that the record does not exist. Then call PutItem.
- B. Enable the TTL attribute on the DynamoDB table.

- C. Implement a conditional put by using the `attribute_exists(transactionID)` condition expression.
- D. Implement a conditional put by using the `attribute_not_exists(transactionID)` condition expression.

ANSWER: C

QUESTION NO: 20

An organization must store thousands of sensitive audio and video files in an Amazon S3 bucket. Organizational security policies require that all data written to this bucket be encrypted.

How can compliance with this policy be ensured?

- A. Use AWS Lambda to send notifications to the security team if unencrypted objects are put in the bucket.
- B. Configure an Amazon S3 bucket policy to prevent the upload of objects that do not contain the `x-amz-server-side-encryption` header.
- C. Create an Amazon CloudWatch event rule to verify that all objects stored in the Amazon S3 bucket are encrypted.
- D. Configure an Amazon S3 bucket policy to prevent the upload of objects that contain the `x-amz-server-side-encryption` header.

ANSWER: B